

Haptic Augmentation of the Cursor: Transforming Virtual Actions into Physical Actions

Ian Oakley

PhD Thesis submitted to University of Glasgow

**Research conducted in the Department of Computing
Science**

13/05/03

Supervisors

Stephen Brewster

Philip Gray

Abstract

This thesis demonstrates, through the exploration of two very different examples, the general claim that haptic feedback relating to a user's representation in a computer system (typically a cursor) can lead to increases in objective performance and subjective experience. Design guidelines covering each of these two topics are also presented, to ensure that the research described here can be readily adopted by other researchers, designers and system developers.

The first topic to be investigated was desktop user interfaces. This thesis describes the design of a variety of different forms of haptic feedback for use with number of different Graphical User Interface (GUI) widgets, or widget groups. Two empirical evaluations of these designs are also described in some depth. The results of these studies indicate that although haptic feedback can provide improvements in objective performance, it can also reduce performance and increase subjective workload if inappropriately applied. From these results, and from the previous literature, detailed guidelines were drawn up covering the addition of haptic feedback to GUIs. The goal of these guidelines is to support the creation of performance-enhancing haptic feedback.

The second topic examined was communication in interactive collaborative systems. The design of a suite of haptic communication is presented in detail, before two studies investigating different aspects of its use. The first study focuses on the subjective influence of the haptic communication as a whole, while the second is a more thorough look at one particular form of the feedback and includes objective measurements. The combined results of these studies indicate that haptic feedback has a valuable potential for increasing the quality of a user's subjective experience. Observations from these studies also reveal insights into the role of haptic feedback in communication. A set of guidelines summing up this research and the previous literature relevant to this topic are then presented. As research on this domain is in its infancy, the goal of these guidelines is to concisely present the main issues and potential benefits that respectively restrict and drive this work.

Thesis Statement

Haptic augmentation of cursor interactions can improve user task performance and subjective experience.

Table of Contents

1. Introduction	1
1.1. Demonstration of general claim	6
1.1.1. Haptics in graphical user interfaces	6
1.1.2. Haptics in collaborative systems	7
1.2. Contents of this thesis	7
2. Haptics	12
2.1 Introduction	12
2.2 Psychophysics of Touch	12
2.3. Haptic Technologies and Devices	17
2.3.1. Tactile hardware	17
2.3.2. Kinaesthetic hardware	20
2.3.3. The PHANToM	30
2.3.4. Point Force	31
2.4. Application Domains	33
2.5. Terminology	36
2.6. Conclusion	36
3. Haptic Augmentation of the Desktop	38
3.1. Introduction	38
3.2. Motivations	39
3.3. Single Target Haptic Augmentation	41

3.4. Multi Target Haptic Augmentation	47
3.5. Summary	49
3.6. Conclusions	51
4. Design, Implementation and Evaluation of Desktop Haptic Feedback	52
4.1. Introduction	52
4.2. Comparison of different forms of haptic feedback	53
4.2.1. Motivations	53
4.2.2. Experimental Task	53
4.2.3. Haptic Effects	55
4.2.4. System	61
4.2.5. Participants	62
4.2.6. Experimental Design	62
4.2.7. Measures	63
4.2.8. Procedure	66
4.2.9. Hypotheses	67
4.2.10. Results	67
4.2.11. Discussion	76
4.2.12. Conclusions	83
4.3. Multi-Target Augmentation of a Menu System	83
4.3.1. Introduction and Motivations	83
4.3.2. Experimental Task	84
4.3.3. Haptic Feedback	86
4.3.4. System	88
4.3.5. Participants	89
4.3.6. Experimental Design	89
4.3.7. Measures	89
4.3.8. Procedure	90

4.3.9. Hypotheses	91
4.3.10. Results	91
4.3.11. Discussion	94
4.3.12. Conclusions	97
4.4. Further Work on this Topic	97
4.4.1. Haptic Feedback	98
4.4.2. Evaluation of this feedback	99
4.5. Design Guidelines for haptic widgets	103
4.5.1. Introduction	103
4.5.2. Literature Review	104
4.5.3. Methodology for Creation of Guidelines	105
4.5.4. Guidelines	106
4.5.5. Conclusions from Guidelines	109
4.6. Conclusions	110
5. Haptics in Communication	111
5.1. Introduction	111
5.2. Motivations	112
5.3. Literature Review	113
5.3.1. CSCW and Groupware	113
5.3.2. Shared Editors	115
5.3.3. Problems with Shared Editors	117
5.3.4. Solutions to these Problems	120
5.3.5. Evaluation in CSCW	124
5.3.6. Techniques for Experimental Evaluation	126
5.3.7. Haptics in Communication	128
5.3.8. Summary of Haptics in Communication	135
5.4. Conclusions	136

6. Design, Implementation and Evaluation of Haptic Communication	138
6.1. Introduction	138
6.2. Design of Haptic Communication	139
6.2.1. Push	139
6.2.2. Gesture	140
6.2.3. Proximity	142
6.2.4. Locate and Grab	143
6.2.5. Discussion of Haptic Cursor Communication	144
6.3. Generic Implementation Details	145
6.3.1. Haptic Simulation	145
6.3.2. Implementation of Networked Haptics	145
6.4. Initial Exploratory Study	146
6.4.1. Introduction and Motivations	146
6.4.2. Experimental Task	147
6.4.3. Haptic Feedback	147
6.4.4. Collaborative Editor	148
6.4.5. Participants	151
6.4.6. Experimental Design	151
6.4.7. Procedure	151
6.4.8. Measures	152
6.4.9. Hypotheses	154
6.4.10. Results	155
6.4.11. Discussion	159
6.4.12. Conclusions	162
6.5. Haptic Gesturing Study	162
6.5.1. Introduction and Motivations	162
6.5.2. Literature Review	163
6.5.3. Haptic Feedback	165

6.5.4. Materials	165
6.5.5. Task	166
6.5.6. Participants	167
6.5.7. Experimental Design	168
6.5.8. Measures	168
6.5.9. Procedure	169
6.5.10. Hypotheses	169
6.5.11. Results	169
6.5.12. Discussion	174
6.5.13. Conclusions	176
6.6. Guidelines for Haptic Communication and Collaboration	177
6.6.1. Design Considerations for Haptic Communication	178
6.6.2. Benefits Conveyed by Haptic Communication	181
6.6.3. Conclusions from Guidelines	183
6.7. Conclusions From Haptic Communication	183
7. Conclusions	185
7.1. Introduction	185
7.2. Summary of research	185
7.2.1. Haptics in desktop user interfaces	186
7.2.2. Haptics in collaborative systems	187
7.3. Limitations of this work	189
7.3.1. Haptics in desktop user interfaces	190
7.3.2. Haptics in collaborative systems	192
7.4. Future work	193
7.4.1. Guidelines for the addition of haptic cues to user interfaces	193
7.5. Summary of contribution	194
7.5.1. Haptics in desktop user interfaces	195
7.5.2. Haptics in collaborative systems	196

7.6. Conclusions	197
8. References	198
Appendix A. Materials for Gathering TLX Data	221
Appendix B. Questionnaire Administered in Buttons Study	224
Appendix C. Raw Data from Buttons Study	226
Appendix D. Questionnaire Administered in Menu Study	240
Appendix E. Raw Data from Menu Study	242
Appendix F. Problem used in Shared Editor Study	246
Appendix G. Instructions and User Manual from Shared Editor Study	248
Appendix H. Raw Data from CHASE Study	252
Appendix I. Images Used in Gesture Study	258
Appendix J. Instructions from Gesture Study	261
Appendix K. Questionnaire Administered in Gesture Study	265
Appendix L. Raw Data from Gesture Study	267
Appendix M. Sample of Images Used in Rating Process in Gesture Study	277

Acknowledgements

Firstly, thanks must go to Stephen Brewster and Philip Gray, my supervisors throughout this long PhD process. Their comments, perspectives and thoughts have helped shape this thesis in innumerable ways. I certainly wouldn't have reached the end without their support and advice. I must also extend thanks to all my various office mates during my time at Glasgow, but especially Marilyn McGee and Andrew Crossan; we three dungeon dwellers managed to remain friends (well, at least on speaking terms) despite our cramped and dingy quarters. I must also thank them for uncomplainingly reading all the various drafts I foisted on them over the years.

Sile O'Modhrain was kind enough to give me a job before I had finished this document, which very much kept the wolf from the door. She didn't even complain too often that I hadn't completed it yet, and amazingly, when she did, she seemed satisfied with my unconvincing replies suggesting that it would be imminent, that I was just in the process of scribbling the final words. To Sile, I am very grateful.

I must also thank Ann Tanner for supporting me throughout my PhD. Its been tough time, especially this final year where various aspects of the real world have intruded on my abstract academic bliss. She kept me working when I really couldn't face it anymore. On a lighter note, I must also thank her for all the proof reading she did on my endless drafts. Thanks also to my parents, who kept me focused, with relatively incessant questions regarding my progress. Sometimes its just what you need. Other times, however.....

The work described here was funded by ESPRC project GR/L79212 and EPSRC studentship 98700418. Thanks must also go the SHEFC REVELATION Project, SensAble Technologies and Virtual Presence Ltd.

Declaration

All publications relating to this thesis have been co-authored with Stephen Brewster and Philip Gray. The two studies described in Chapter 4 have been published at ACM CHI'2000 and ACM CHI'2001 [150, 151]. The first of these papers included an additional author: Marilyn McGee. The design guidelines relating to the integration of haptic feedback into GUIs were published at the 16th British HCI Group Annual Conference (BHCI'2002) [147], and also included an extra author: Alison Adams. The design of the haptic communication presented at the start of Chapter 6 appears in *Haptic Human-Computer Interaction*, Springer LNCS Vol 2058 [149]. The first collaborative study conducted, which examines all aspects of the haptic communication, appears in the proceedings of EuroHaptics 2001 [148]. This thesis only exploits those parts of collaborative papers that are directly attributable to the author.

List of Figures

- Figure 2.1.** Lederman's Exploratory Procedures (reprinted without permission from Lederman *et al.* [114]). **16**
- Figure 2.2.** The Tractile vibrotactile device (reprinted without permission from Campbell *et al.* [38]). **19**
- Figure 2.3.** The Brailnote, a tactile display for visually impaired users (reprinted without permission from www.brailnote.com). **19**
- Figure 2.4.** Akamatsu's haptic feedback mouse (reprinted without permission from Akamatsu [4]). **22**
- Figure 2.5.** The Maglev haptic device (reprinted without permission from www-2.cs.cmu.edu/afs/cs/project/msl/www/haptic/haptic_device.html). **23**
- Figure 2.6.** Rutgers Master II-ND force feedback glove, a typical body based device (reprinted without permission from Bouzit *et al.* [23]). **24**
- Figure 2.7.** A typical ground based haptic device, the HapticMaster (reprinted without permission from www.fcs-cs.com/robotics/). **25**
- Figure 2.8.** The Wingman Extreme vibration device (reprinted without permission from www.logitech.com). **26**
- Figure 2.9.** The Wingman force feedback joystick (reprinted without permission from www.logitech.com). **26**
- Figure 2.10.** The Wingman force feedback mouse (reprinted without permission from www.logitech.com). **27**
- Figure 2.11.** The Cybergrasp glove based force feedback device (reprinted without permission from www.immersion.com). **28**

Figure 2.12. The Laparoscopic Surgical Workstation, a haptic device for the simulation of laproscopic surgery (reprinted without permission from www.immersion.com).	28
Figure 2.13. The SPIDAR haptic device (reprinted without permission from Bouguila <i>et al.</i> [22]).	29
Figure 2.14. The Delta haptic device (reprinted without permission from www.forcedimension.com)	29
Figure 2.15. The CyberForce haptic device (reprinted without permission from www.immersion.com).	30
Figure 2.16. The PHANToM 1.0 haptic device (reprinted without permission from www.sensable.com).	30
Figure 4.1. Screen shot of buttons experiment (annotated with lines indicating the path of the corner buttons).	54
Figure 4.2. 3D visualisation of the texture effect.	56
Figure 4.3. Cross-section of the texture effect with force orientation marked.	57
Figure 4.4. 3D visualisation of the recess effect.	59
Figure 4.5. Cross section of the recess effect.	59
Figure 4.6. Plan view of the gravity well effect with force direction and magnitude illustrated.	60
Figure 4.7. Timing measures used in buttons study.	64
Figure 4.8. Mean leaving time in different Gravity Well magnitude conditions in buttons study.	68
Figure 4.9. Mean total trial per trial in buttons study.	68
Figure 4.10. Mean approach time per trial in buttons study.	69
Figure 4.11. Mean stopping time per trial in buttons study.	69

Figure 4.12. Mean clicking time per trial in buttons study.	70
Figure 4.13. Mean leaving time per trial in buttons study.	70
Figure 4.14. Mean total slide-over errors in different Gravity Well magnitude conditions in buttons study.	71
Figure 4.15. Mean total slide-over errors in different Texture magnitude conditions in buttons study.	71
Figure 4.16. Mean total slide-over errors in buttons study.	72
Figure 4.17. Mean total slip-off errors in buttons study.	73
Figure 4.18. Mean total off-target errors in buttons study.	73
Figure 4.19. Mean total wrong-target errors in buttons study.	74
Figure 4.20. TLX workload results from buttons study.	75
Figure 4.21. Mean ratings from post-experiment questionnaire in buttons study.	76
Figure 4.22. Mean selection time per trial in buttons study.	78
Figure 4.23. Mean total time on target in buttons study.	78
Figure 4.24. Screenshot of menu experiment.	85
Figure 4.25. Force profile of two haptic barriers from menu study.	86
Figure 4.26. Modification of force according to user direction in menu study.	87
Figure 4.27. Mean time per trial in menu study.	92
Figure 4.28. Mean total slide-over errors in menu study.	92
Figure 4.29. Mean total wrong-target errors in menu study.	93
Figure 4.30. Mean total skip-ahead actions in menu study.	93
Figure 4.31. Mean total modified-slide-over errors in menu study.	94

Figure 4.32. TLX workload results from menu study.	95
Figure 4.33. Screenshot of toolbar study.	99
Figure 4.34. Screenshot of icons study.	100
Figure 5.1. Illustration of different awareness widgets (reprinted without permission from Gutwin <i>et al.</i> [84]).	123
Figure 5.2. The inTouch haptic communication device (reprinted without permission from www.media.mit.edu/tangible).	131
Figure 5.3. Illustration of handJive haptic communication device (reprinted without permission from Fogg <i>et al.</i> [67]).	132
Figure 5.4. The ComTouch haptic communication device (reprinted without permission from Chang <i>et al.</i> [40]).	132
Figure 5.5. Illustration of interface for haptic communication (reprinted without permission from Basdogan <i>et al.</i> [13]).	133
Figure 6.1. Visualisation of haptic gesturing algorithm.	141
Figure 6.2. Screenshot of CHASE.	149
Figure 6.3. Architecture of the CHASE system.	150
Figure 6.4. TLX workload results from CHASE study.	157
Figure 6.5. ITC Presence questionnaire results from CHASE study.	158
Figure 6.6. CS usability questionnaire results from CHASE study.	159
Figure 6.7. Custom collaborative questionnaire results from CHASE study.	159
Figure 6.8. Encoder's view in gesture study.	166
Figure 6.9. Decoder's view in gesture study.	167
Figure 6.10. Mean time per trial in gesture study.	170

Figure 6.11. TLX workload results from gesture study.	171
Figure 6.12. Mean image rating per condition in gesture study.	172
Figure I.1. Images used in practice session in gesture study.	259
Figure I.2. Images used in experimental session in gesture study.	260
Figure J.1. Annotated encoder view in gesture study.	262
Figure J.2. Annotated decoder view in gesture study.	263
Figure J.3. Scale used in image rating process.	264
Figure M.1. Original image (K42) used in rating process.	278
Figure M.2. First nine copies used in rating process for image K42.	279
Figure M.3. Second nine copies used in rating process for image K42.	280
Figure M.4. Original image (K11) used in rating process.	281
Figure M.5. First nine copies used in rating process for image K11.	282
Figure M.6. Second nine copies used in rating process for image K11.	283

List of Tables

Table 2.1. Definitions of terms used in this thesis.	36
Table 4.1. The eight different orders of condition presentation in buttons experiment.	63
Table 4.2. Results from pairwise comparisons of slip-offs in buttons experiment.	74
Table 4.3. Results from pairwise comparisons of slide-overs in buttons experiment.	74
Table 4.4. Results from pairwise comparisons of off-targets in buttons experiment.	74
Table 4.5. Results from ANOVA analysis of TLX results in buttons experiment.	76
Table 4.6. Results from ANOVA analysis of TLX results in menu experiment.	95
Table 4.7. Data from toolbar study reported by Adams [2].	101
Table 4.8. Statistics from toolbar study reported by Adams [2].	101
Table 4.9. Data from icons study reported by Adams [2].	102
Table 4.10. Statistics from icons study reported by Adams [2].	102
Table 5.1. Johansen's [106] classification of groupware applications.	114
Table 6.1. Questions from custom collaborative questionnaire.	154
Table 6.2. Results from ANOVA analysis of TLX results in gesture experiment.	171
Table 6.3. Rater Coefficients of Correlation in gesture experiment.	172

Table A.1. Workload descriptions given to subjects when filling in the workload charts. The descriptions are taken from [111].	222
Table A.2. Form given out to subjects to capture TLX workload data.	223
Table C.1. TLX data from gravity well condition.	227
Table C.2. TLX data from recess condition.	227
Table C.3. TLX data from friction condition.	228
Table C.4. TLX data from texture condition.	228
Table C.5. TLX data from visual condition.	229
Table C.6. Participant Preference for each condition.	229
Table C.7. Approach and stopping time in gravity well condition.	230
Table C.8. Approach and stopping time in recess condition.	230
Table C.9. Approach and stopping time in friction condition.	231
Table C.10. Approach and stopping time in texture condition.	231
Table C.11. Approach and stopping time in visual condition.	232
Table C.12. Clicking and leaving time in gravity well condition.	232
Table C.13. Clicking and leaving time in recess condition.	233
Table C.14. Clicking and leaving time in friction condition.	233
Table C.15. Clicking and leaving time in texture condition.	234
Table C.16. Clicking and leaving time in visual condition.	234
Table C.17. Slip-off and slide-over errors in gravity well condition.	235
Table C.18. Slip-off and slide-over errors in recess condition.	235
Table C.19. Slip-off and slide-over errors in friction condition.	236

Table C.20. Slip-off and slide-over errors in texture condition.	236
Table C.21. Slip-off and slide-over errors in visual condition.	237
Table C.22. Off-target and wrong-target errors in gravity well condition.	237
Table C.23. Off-target and wrong-target errors in recess condition.	238
Table C.24. Off-target and wrong-target errors in friction condition.	238
Table C.25. Off-target and wrong-target errors in texture condition.	239
Table C.26. Off-target and wrong-target errors in visual condition.	239
Table E.1. Results from TLX in visual condition.	243
Table E.2. Results from TLX in adjusted condition.	243
Table E.3. Results from TLX in haptic condition.	244
Table E.4. Objective measures from visual condition.	244
Table E.5. Objective measures from adjusted condition.	245
Table E.6. Objective measures from haptic condition.	245
Table H.1. Results from custom collaborative questionnaire in haptic condition.	253
Table H.2. Results from custom collaborative questionnaire in visual condition.	253
Table H.3. Results from TLX in haptic condition.	254
Table H.4. Results from TLX in visual condition.	254
Table H.5. Results from ITC presence questionnaire in haptic condition.	255
Table H.6. Results from ITC presence questionnaire in visual condition.	256
Table H.7. Results from usability questionnaire in haptic condition.	257

Table H.8. Results from usability questionnaire in visual condition.	257
Table L.1. TLX results from decoder visual condition.	268
Table L.2. TLX results from decoder haptic condition.	268
Table L.3. TLX results from decoder combined condition.	269
Table L.4. TLX results from encoder condition.	269
Table L.5. Trial times (in seconds) in each condition.	270
Table L.6. Similarity ratings (from three raters) for first 90 encoder images.	271
Table L.7. Similarity ratings (from three raters) for second 90 encoder images.	272
Table L.8. Similarity ratings (from three raters) for third 90 encoder images.	273
Table L.9. Similarity ratings (from three raters) for fourth 90 encoder images.	274
Table L.10. Similarity ratings (from three raters) for fifth 90 encoder images.	275
Table L.11. Similarity ratings (from three raters) for sixth 90 encoder images.	276

1. Introduction

Haptic, according to Webster's dictionary [198], is a term that "*refers to something of, or relating to the sense of touch*". Haptic, or touch, cues are therefore a ubiquitous part of everyday life. They are a continual and essential source of information during the performance of virtually any physical activity ranging from reading a book, where we almost subconsciously hold and turn the pages, to participating in a sport, where proficiency in haptic interaction is highly prized and honed to near perfection. Haptic feedback is of critical importance whenever we come into contact with objects, perhaps picking them up or simply manipulating them in some way, and we instinctively rely on the inherent haptic cues we receive from these actions to inform us about the world and regulate our performance.

Given this fact, it is unsurprising that haptic feedback has been explicitly incorporated into the physical interfaces to devices that could operate equally well without it; into electronic devices. Although there is little requirement for them to do so, as alternatives such as pressure or capacitive [153] sensors allow users to input data without experiencing any haptic stimuli (and indeed, lacking moving parts, these sensor systems are often more reliable and long-lasting), the interfaces to electronic devices almost universally exhibit haptic feedback in the basic form of their construction. Controls protrude from the surface to which they are fixed, affording haptic notification when they are encountered, and haptic feedback is explicitly engineered and apparent in the dynamics of such objects. Dials feature a reassuring resistance to motion when turned, and pressing a button results in an affirming click. Millar and Zeleznik [136] classify these types of feedback as

“guidance” and “follow-through”. The former supports the haptic discovery of objects while the later indicates that an action has been achieved. Standard computer keyboards possess these properties and, in work forming a powerful argument for the adoption of this kind of feedback, it has been demonstrated that the follow-through cues provided by depressing and releasing keys is vital to the performance of both novice and expert typists [12].

As with the keyboard, the controls on a mouse are also augmented with haptic cues: pressing a button results in a clicking sensation, and the scroll-wheels [207] that are becoming commonplace feature a series of physical detents (or bumps) when they are turned. However, in contrast, no explicit and designed haptic feedback is associated with mouse motion. Kinaesthetic feedback, in the form of our awareness of the position and orientation of our limbs is present; it is a property of our haptic sensory system. Although this self-referential feedback is no doubt valuable to a user, it is obviously incapable of representing the complexities of interface state; as it is not actively generated through the mouse it is essentially the same at any given moment, for any given task, and cannot compare with the richness that can be achieved through the construction of specific, carefully designed haptic cues. This absence of intentional, mechanical, haptic feedback separates mouse input from input with the vast majority of other controls, and this difference becomes more dramatic when the importance of cursor motion is considered. Arguably the concepts behind GUIs are fundamentally built on interaction through cursor motion. Although the depression or release of a button usually signifies intent, this is mediated by the position of the cursor. In GUIs, objects are inherently spatially segregated and cursor position is the chief determinant as to what action is taken in response to a command.

The lack of active, purposefully generated, haptic feedback in mouse motion is both explained and its significance exacerbated by the fact that it forms a rich input channel: it involves two independent spatial axes, which are of sufficiently high resolution to support varied and expressive physical interaction. This richness enables the performance of complex sets of actions and tasks, a potentially demanding scenario where it would seem likely that users could strongly benefit from haptic cues designed to aid them; the inherent richness of the input channel highlights the absence of meaningful haptic feedback accompanying it. Examples of

the range of tasks that can be performed are targeting key objects, for instance when interacting with buttons; using absolute displacement to indicate change, as when interacting with scrollbars; and specifying isometric parameters, as when interacting with AutoScroll widgets. The lack of haptic feedback relating to mouse motion is also explained by the complexity of the input channel: mouse motion is an interaction technique intended to be general in scope and work in a variety of situations, and these desirable features separate it from the vast majority of other input techniques and devices which typically behave in a single way and serve a single purpose. This non-specificity of function means that no single haptic cue can be engineered into the hardware of the mouse, as any given cue would be inappropriate in almost every situation. In short, there is no simple physical mapping between mouse state and meaningful haptic cue. This is true because mouse-movement is an input technique that gains its context, and therefore the information about the haptic cues that might be appropriate, not from its physical environment, but instead from the virtual environment, the desktop, that its virtual representation, in the form of the cursor, inhabits.

This fact presents substantial challenges to the creation of haptic feedback to support mouse motion. Historically, the most significant of these has been the demanding constraints placed on the hardware which renders the haptic cues. As the feedback produced must be controlled from within the virtual environment, the generation and construction of the cues must be based in software. The hardware, then, must be able to dynamically and precisely produce feedback that satisfactorily matches the input capabilities of mouse motion: two active Degrees Of Freedom (DOF). Until relatively recently, hardware capable of producing this kind of feedback has not been available, and this fact has doubtlessly contributed significantly to the largely unexamined nature of this topic.

Now, however, suitable haptic devices are available, but further problems are emerging. The most significant of these is simply that there is little knowledge or expertise on what haptic feedback is appropriate for the kind of virtual tasks that the cursor is used for. Haptic cues in the design of physical controls have a long and established pedigree [146], and there are a lot of *de facto* standards and common sense knowledge regarding what cues are appropriate in different situations. The

ubiquitous click upon press in a switch or button is an excellent example of this fact, and the kind of evolved understanding that this represents is a strong contributory factor to the general high quality and significant value of these kinds of cues. Unfortunately, this knowledge about physical design does not seem to be directly transferable to cursor interactions in virtual environments. An example of this comes from a recent study of button selection [160]. In this work, the authors contrasted user performance operating a haptically rendered button that needed to be physically pressed, much like a button in the real world, with one that was activated through the depression of a physical button mounted on the haptic device. Despite its weaker metaphorical connection to real world interaction, they found this second technique to be superior.

One reason why the design heuristics and principles that are successful and appropriate in physical design do not generalise to the design of virtual haptic stimuli in a computer mediated scenario is simply that the tasks that are undertaken in these different situations do not match. Most physical interfaces in the real world are relatively simple in construction; they feature a number of buttons, and perhaps one or more dials or sliders, and no more complex constructs. Arguably, even very complex physical interfaces, such as those in the cockpits of planes, are composed of little more than a large number of relatively straightforward basic components. Even simple virtual interfaces, on the other hand, often possess many different and more complex controls. These include menus, movable objects and a range of buttons possessing subtly different standards for interaction (e.g. double-click or single-click to activate). More complicated virtual tasks feature yet more complex controls such as grids and rulers, which support a variety of different forms of behaviour and interaction, and multi-dimensional selection tools such as colour pickers. These kinds of advanced virtual interaction devices do not even possess a physical equivalent that haptic feedback could be based upon and it is hard to envisage how knowledge relating to the construction of physical controls can be generalised to these sophisticated interfaces. This lack of knowledge about what haptic feedback to apply to virtual interactions is a barrier to the adoption of this technology that is arguably as strong as the historical lack of suitable display devices.

In work very relevant to this argument, Maclean and colleagues [124, 176] have investigated the role of software driven haptic feedback in general-purpose control devices. Typically, this work has focused on the scenario of media control, and considers the design and architecture of a system that would allow users to operate a single control in order to access all functions of, for instance, their stereo. Examples of the kind of operations this control might have to perform are to adjust the volume, skip forward to a specific track, or tune a radio. The work suggests that a single physical interface device (typically a dial or rotational knob) could meaningfully control all these functions if it possessed dynamically generated haptic cues. In supporting their work, they state [124]:

“It [software controlled haptic feedback] further enables powerful metaphoric operations: interaction with a target via a mediating virtual physical model can be designed to permit seamless, intuitive control”

While these ideas have not yet undergone formal evaluation, the authors do present a compelling case for their validity. A similar, and equally persuasive, example of haptic feedback in a general purpose control has recently appeared in a commercially available vehicle [8]. In this car a single haptically augmented dial is used to interact the vast majority of the electrical systems. While neither of these cases provides formal evidence for their claims, the academic and design interest that they represent does suggest that further investigations into these ideas may prove valuable. This thesis intends to do this, to investigate the addition of haptic cues to a virtual user interface from the perspective of creating controls which feature a beneficial physical representation.

In short, the main point of this discussion is that cursor interactions are extremely complex and flexible and, in contrast to the majority of other controls and interfaces, are not supported by any haptic feedback. Haptic feedback has attracted interest as a mechanism for creating a physical aspect to general purpose controls, and this thesis extends this idea. It takes the position that the addition of haptic cues to cursor interactions will lead to increases in performance and subjective experience. A further point from this discussion, and correspondingly a further aspect of this thesis, is concerned with how the addition of haptic feedback to these interactions might be achieved. It seems likely that existing knowledge about the design of physical

interfaces will not generalise effectively to this new situation. Consequently a secondary goal of this thesis is to further the current understanding of design in this topic. This will be achieved through the compilation of design guidelines.

1.1. Demonstration of general claim

This central claim of this thesis is a general one: that haptic cues relating a user's representation in a computer system (typically termed a cursor, or avatar) can provide measurable benefits in terms of user performance and subjective experience. In order to demonstrate this claim, this thesis proposes to evaluate it empirically in two very different scenarios. This mechanism, essentially proof by example, is a historically established one, and while it cannot provide an absolute proof, it can form a convincing and well-justified argument. This is what is sought here.

From the discussion in the introduction, a distinction can be made between relatively simple controls, such as GUI buttons which are based on analogous real world objects, and more complex interaction devices which are not based on physical objects. The two domains studied to construct the proof by example also follow this distinction; one is based around relatively simple interactions with a direct real-world metaphor, while the other looks at a form of interaction which is inherently and uniquely virtual. The significant advantage of this choice is that it ensures that the two domains go some way towards being representative of overall design space available for cursor interactions. They will also be radically different from one another, a beneficial property for two examples attempting to demonstrate a general claim; they have the potential to show the wide applicability of the idea. The two areas studied are introduced and discussed briefly below.

1.1.1. Haptics in graphical user interfaces

GUIs are presently the dominant form of computer interface for the vast majority of work and leisure activities [173]. The most common cursor interactions involve simple GUI widgets designed to allow users to issue commands: to execute programs, handle files, select options and adjust views, among many others. The kind of controls users interact with to achieve these operations are buttons, icons, menus, scrollbars and so forth. As an everyday aspect of computer use, and one that involves interaction with virtual objects often based around existing physical

controls, this domain is an obvious choice for study in this thesis. An investigation into this topic has the potential to reveal benefits and insights applicable to a large number of users, and the importance of haptic cues in analogous real world objects (such as buttons) strongly suggests that user performance or subjective experience could be improved through the addition of feedback in this novel modality.

1.1.2. Haptics in collaborative systems

In order to investigate a domain that has a vastly different focus from the examination of GUIs, and one that supports interactions not based on a direct real world metaphor, this thesis turned to communication in interactive collaborative systems [50]. Unlike GUIs, interactive collaborative systems are not an everyday part of computer use, and consequently this research may not be generally applicable to most users. However, with the advent of networking technologies such as the Internet, these kinds of tools are developing rapidly, and have substantial potential. Furthermore, research on this topic suggests that users find communication in collaborative environments challenging [57, 84, 103], in part due to the reduced richness of the computer mediated medium; typically they can talk, and may be able to see one another, but this is often no substitute for occupying the same physical space. As novel channel for communication, haptic feedback may prove useful by providing new ways for users to express themselves to one another.

1.2. Contents of this thesis

The goal of this thesis is to demonstrate the general claim that haptic cues relating to a user's avatar can provide improvements in task performance and subjective experience through the exploration of two specific examples. The structure of this thesis reflects this fact. After an initial literature review, in Chapter 2, that covers the general research relating to haptics, there are two distinct threads within the thesis. Chapters 3 and 4 deal with the first of the examples chosen: the addition of haptic cues to desktop interactions. Similarly Chapters 5 and 6 cover the second example: haptics cues in communication and collaboration scenarios. Chapter 7 summarises the whole of this work, and draws conclusions from this thesis. A more detailed description of the contents of each chapter is presented below.

Chapter 2 reviews the general literature relating to haptic feedback in computer interfaces. It sets the context for this thesis by exploring the background literature relating to touch, and is important as the research discussed within it played a crucial role informing the design of the various novel forms of haptic feedback presented in the subsequent chapters. The chapter itself is split into a review of three distinct topics, each important for any design effort in this field. The topics are: psychophysics, hardware and applications. The section examining psychophysics considers the psychological research on the human sense of touch. This substantial body of work describes the physical capabilities and perceptual processes that underlie the sense of touch, and consequently is a vital resource for the design of any haptic stimuli: it describes what we perceive and how we do it. The section covering hardware looks at the various display technologies for the production of haptic cues. An understanding of this research is important, as currently available devices are not capable of producing haptic sensations that match the richness of those present in real world interactions. Just as the psychophysics of touch imposes constraints on the creation of haptic stimuli, so do the capabilities of the output devices available. The designer of a haptic cue must consider not only what would make an effective percept, but also whether or not the cue can be adequately rendered. Indeed, using currently available technologies, the choice of output device is inextricably linked to the specific qualities of the desired stimuli. The section on applications briefly examines where haptic feedback has been previously applied in computer interfaces. This review serves to provide context to the designs and evaluations described in the following chapters. It is presented from the perspective that the previous successes and failures of haptic feedback in computer interfaces can provide valuable insight into where haptic feedback might prove effective in novel situations. This chapter concludes with the statement that the application of haptic feedback in user interfaces is challenging, at least in part due to the complexities of haptic psychophysics and hardware. It suggests that design guidelines relating to different aspects of this topic would be beneficial, and facilitate the adoption of this technology.

Chapter 3 presents an in-depth literature review of haptic feedback in desktop user interfaces, the first of the two specific domains to be examined in this thesis. As such it serves as a detailed exploration of haptic feedback in this scenario, setting the

stage for the empirical work presented in the next chapter. The literature review is split into a section covering the motivations for the addition of haptics to desktop user interfaces, and two sections describing the previous literature on this topic. Respectively, these sections discuss work termed in this thesis as *single-target* and *multi-target*, referring to work that considers the simple case of individual haptically active targets and work that examines the more complex and realistic situation when many haptically active targets are presented simultaneously. This chapter closes by summarising this body of work, and describing its main weaknesses and omissions. Critically, these are a lack of work directly comparing the different forms of haptic feedback that can be used in GUIs, an absence of an effective technique for applying haptic feedback in multi-target situations and a lack of design guidelines, or any other tool, that distils the disparate research on this topic into a palatable form that interface designers can easily adopt.

Chapter 4 is composed of empirical and design work on the topic of haptic feedback in desktop user interfaces. The goal of this chapter is to demonstrate that haptic feedback relating to a user's cursor can yield performance increases and improvements in subjective experience in this domain. The work presented in this chapter addresses the limitations identified in the previous chapter and includes the detailed description of two studies, some additional design work and a set of design guidelines. The two studies examine two of the topics identified as important in the previous chapter. The first compares four different forms of haptic feedback in a simple single-target situation in order to characterise the influence of the different types of feedback. The second study introduces and examines a novel approach to the mediation of haptic feedback in a multi-target scenario. The evaluation yielded promising results, and this idea is extended with the design of another haptic widget set. This is followed by a brief description of work strongly related to this thesis, where this design was evaluated in two different situations. This chapter closes by presenting a set of design guidelines drawn not only from the studies described here, but also from the wider body of literature. These guidelines are the first of their kind on this topic and are intended to support the integration of this novel form of feedback into the creations and products of designers and system developers working in real world scenarios. Overall, this chapter achieves its objectives: the work presented here furthers knowledge on this topic, provides a compelling argument for

the integration of haptic cues into desktop scenarios and, in the form of the design guidelines, describes how this might be achieved by system developers working on real problems and products.

Chapter 5 contains a literature review covering the second of the examples to be examined in detail in this thesis: haptics in communication and collaboration scenarios. The role of this chapter in this thesis is to provide a context in which haptic feedback can be applied to a collaborative environment. To this end, it begins with the motivations for adding haptic feedback to these environments, and then focuses in on the specific class of applications that are considered here: shared editors. Shared editors are then described in some detail, with a specific focus on the communication problems apparent in these systems and the previous work that has attempted to enhance the relatively basic communication available through visual, audio or semantic augmentations of the presented environment. With the domain for the empirical work described, the review then turns to the topic of evaluation methodologies and techniques for assessing the effectiveness of collaborative systems. This section is important, as evaluation in collaborative systems is a difficult task, due to the fact that the assessment is typically focused on the quality of communication supported and not on some more directly measurable feature of the system. Finally, the chapter includes a detailed examination of the role of touch in communication. It covers a diverse range of work from psychological perspectives on interpersonal touch, to descriptions of the collaborative and communication systems that have been built to explore the potential of sending and receiving haptic messages.

Chapter 6 builds on the literature review presented in the previous chapter and describes the design, implementation and evaluation of a novel form of haptic feedback for shared editors. The goal of this chapter is to show that haptic cues presented in this collaborative scenario can improve user performance and experience. The design of the feedback involved the augmentation of telepointers – essentially distributed cursors – with five different haptic interactions. Two studies exploring the influence of this feedback on users are then described. The first examines users engaged in a complex high-level design task, and observationally and subjectively contrasts user experience and behaviour with and without the haptic

feedback. This study was intended to provide a broad overview of the influence of the haptic feedback, and the qualitative measures employed reflect this fact. The second study follows on directly from observations in the first study, and is a focused investigation of one specific aspect of one of the haptic interactions. The more constrained nature of this second study is reflected in the objective measures used; these include task performance time and quality of solution as well as a subjective assessment. Although an interesting observational study, the results of this experiment do not indicate that the haptic feedback provides an objective improvement in task performance. Drawing on the insights gained from these studies, and from the previous literature, guidelines for the application of haptic feedback in collaborative systems are then presented. Given the embryonic nature of research on this topic, the goal of these guidelines is to present a concise description of the main design issues that influence the creation of haptic communication, and the advantages that it can offer users. On the whole, this chapter achieves its objectives. The first study shows a broad spread of subjective improvements due to the haptic communication, while the observations of both experiments provides insight into the inner workings of communication through touch. Finally, the guidelines summarise the research on this topic, presenting its main tenets in a palatable form for the consumption by future researchers or designers.

Chapter 7 summarises the contribution of this thesis. It also discusses the limitations of the work presented here, and comments on areas that deserve further study.

2. Haptics

2.1 Introduction

In order to design haptic interfaces the general background behind such systems must be considered. This chapter achieves this by reviewing the existing haptic literature from three stances. The first stance considers psychophysical research, describing and classifying the perception of touch. The second relates how psychological research has been instantiated in the design of haptic devices, and briefly describes and classifies these devices, focusing on their capabilities and limitations. The final stance looks at previous applications in which haptic feedback has been applied, in order to gain a perspective on what it can offer a user. This chapter then moves on to specify and define the terminology used in this thesis. This is an important task as the field of haptics is currently at a stage in which there is no universally agreed lexicon of terms: different researchers are likely to assign terms with widely varying meanings. This chapter closes with a brief discussion of the absence of explicit design knowledge in this domain, and how such understanding would facilitate the integration of this kind of feedback into real world systems.

2.2 Psychophysics of Touch

The psychophysics of touch, the study of perception through touch, plays a vital role in the design and construction of any haptic interface. The knowledge drawn from this field delineates the requirements for the creation of haptic interfaces at all levels from the design of devices, to the presentation of stimuli using these devices. Failing to take heed of the psychophysical literature in the process of designing of a haptic

device is likely to lead to a final product whose capabilities do not match those possessed by the human perceptual system, and which is consequently unable to produce realistic stimuli. Similarly the costs associated with ignoring the psychophysical literature when creating haptic stimuli are that those stimuli may be indistinguishable from one another, beneath perceptual thresholds, or may fundamentally interfere with a user's intended actions.

The psychophysical literature reveals several crucial facts concerning the sense of touch. Firstly, a distinction is made between two separate perceptual systems, typically gathered together under the single banner of touch [178], and often both (or either) referred to as haptic. The first of these systems, the cutaneous or tactile system, refers to information sensed through the medium of skin, and encompasses such disparate sensations as texture, temperature and pain [31]. Perhaps the most important aspect of the tactile sense is its ability to perceive vibrations or perturbations of the skin [194]. This enables it to distinguish a wide variety of textural information. The second system, known as the kinaesthetic system, relates to stimuli originating from an intimate knowledge of the internal state of the body [31]. Awareness of the positions of the limbs, the forces exerted by the muscles, and the resistances opposing these forces are all kinaesthetic stimuli. The information gained from these two senses is in the most part different, with one crucial overlap: contact. Physical contact with an object can be felt through the tactile sense via the pressure it exerts on the skin, and also through the kinaesthetic sense due to the displacement (or lack thereof) of the limbs. Other than this similarity, the tactile and kinaesthetic senses yield distinct but complementary information. The tactile sense provides information about the fine grained details of an object, such as its texture, while the kinaesthetic sense informs us about the larger scale details of an object, such as its shape, weight, and how compliant or stiff it is.

There is a large body of literature cataloguing the abilities of the tactile and kinaesthetic systems, and this is briefly summarised here. This review focuses on the abilities of the hand, and to some extent the upper limbs, as these areas of the body possess the greatest tactile and kinaesthetic acuity, and have traditionally attracted the most study. These areas are also the most relevant to this thesis. The hand itself is an elaborate structure consisting of nineteen bones, connected by low friction joints

and attached by tendons to approximately forty muscles [79]. This framework provides twenty-two Degrees Of Freedom (DOF). The tactile and kinaesthetic senses are equally sophisticated.

Tactile sensory capabilities peak on the finger pad. On this most sensitive area of the body, the location of a point can be determined to within 0.15 mm [119] and the resolution at which two points can be discriminated is approximately 1 mm [107]. A single dot 2 microns high can be detected on a smooth surface, as can a grating of 0.06 microns in height [112]. Vibrations up to 10 kHz can be sensed and differences of 320 Hz can be discriminated [140]. Pressure of 10^{-5} Newtons (N) to the skin can be detected. Emphasising the importance of tactile cues in the performance of physical tasks, studies in which participants have had to perform physical manipulations while wearing thick gloves such as work gloves [167] or space-suit gloves [44] have found that task completion times were prolonged by between 10–80%, depending on the thickness of the gloves.


The kinaesthetic sense is also highly accurate. The resolution of joints, calculated by Just Noticeable Difference (JND), in the hand is approximately 2.5 degrees, in the wrist and elbow 2 degrees and in the shoulder 0.8 degrees [188]. When sensing forces the JND, or Weber fraction, is approximately 7% regardless of test conditions. Similarly, the JND for sensing viscosity is approximately 12%, for mass 20%, velocity 11% and acceleration 17% [15]. Fingertip positional resolution when grasping objects of 1 to 80 mm in size is between 0.5 and 2.5 mm [62]. In the hands, force control resolution, the accuracy at which a person can produce forces, is 0.04 N, or 1%, whichever value is higher [180]. The maximum controllable force that can be exerted at fingertip is between 50 and 100 N, depending on whether the shoulder muscles are in use [188]. However, typical forces used in exploration and manipulation are between 5 and 15 N [178]. The force required to present a stiff wall – a wall that feels solid – has been estimated at approximately 25 N, though it has been suggested that as little as 5 N, a fairly compliant surface, is sufficient for the suspension of disbelief [188]. One unique and important aspect of the kinaesthetic sense is that it allows duplex communication; it can act as a closed loop system. Through our kinaesthetic sense we can both perceive and respond to stimuli simultaneously and dynamically; indeed it is often beneficial to actively explore in

this way [74]. This observation characterises how individuals use their sense of touch, and exerts a considerable influence on the design of haptic devices and stimuli.

In landmark work extending this idea of touch as an active exploratory sense, Lederman *et al.* [113, 114] recorded and classified the mechanisms by which we employ both our tactile and kinaesthetic senses to investigate and perceive objects. They observed eight distinct Exploratory Procedures (EP): lateral motion, pressure, static contact, unsupported holding, enclosure, contour following, part motion and function. These are illustrated in Figure 2.1. Lederman *et al.* state that each of these techniques is distinct, and that each facilitates the extraction of different information. For instance, lateral motion, the act of rubbing the fingertips rapidly back and forth on the surface of an object provides information regarding the surface texture of the object. Similarly, exerting pressure on an object yields information concerning its hardness. Static contact, holding the skin still against a object is an effective method of revealing its temperature. Unsupported holding, or picking an object up, often also hefting it up and down, provides information as to its mass. Enclosure, the act of covering an object as completely as possible with the hands, serves to reveal information relevant to that object's overall shape and volume. Contour following, tracing the edges of an object with the fingertip, provides information on its fine details; its precise shape. The final two EPs apply only to specific subsets of objects. Part motion refers to the EP of exploring the relative movement of subparts of an object, while function refers to the act of testing an object for its suitability for various purposes; for instance trying an object on as a glove, or squeezing part of an object thought to make noise.

Although this work contributes little to the basic psychophysical literature (it does not, for instance, provide a numerical account of some perceptual acuity), it is of vital importance as it exposes novel aspects of haptic perception, offering an explanation for some of the differences between the poor performance observed in previous abstract studies of haptic object recognition [115] and clearly observable superior real world performance [111]. While not addressing every aspect of haptic perception, the treatment and explanation, of these processes as active has substantial implications for the creation of any haptic stimuli; it suggests this property should be

Figure 2.1. Lederman's Exploratory Procedures (reprinted without permission from Lederman *et al.* [114]).



exploited. In short, this dynamic view of haptic perception seems both informative and useful, and should be considered in tandem with the more traditional psychophysical data during the construction of any system incorporating haptic cues.

A final important point regarding the characteristics of the human sense of touch is that haptic experiences are extremely susceptible to influences from cues presented in other modalities; specifically visual and auditory cues [199]. Studies have shown,

for instance, that a haptic surface presented with no visual cues will be judged softer than the presentation of that same haptic surface in conjunction with a graphical representation that appears stiffer; one in which the graphical cursor remains on top of a surface, irrespective of inward motion of the haptic cursor [179]. Similar results have been obtained with auditory stimuli [54]. A knocking sound on contact with a wall will increase subjective judgements of the stiffness of that wall. The presence of these cross-modal interactions is significant because they have the potential to be leveraged to create stimuli more compelling than solely haptic equivalents.

2.3. Haptic Technologies and Devices

The ultimate goal in the creation of a haptic device is to generate realistic sensations of touch; to match the stimulus production capabilities of the device against those of the human perceptual system, as described by the body of psychophysical research [188]. By achieving this aim a virtual stimulus will be indistinguishable from a real stimulus. This goal has been achieved, to some extent, in the domains of both visual and auditory display. These technologies have been developed to such an extent that they convincingly match human sensory abilities, creating engrossing illusions of reality. The minimum refresh rate for a visual display is 30 Hz because presentation of still images at this speed fools the human visual system into believing that it perceives a moving picture. Similarly, audio speakers are capable of producing sound in a range of frequencies and amplitudes that match the abilities of the human ear. Haptic technologies are not yet at this level; the available hardware does not match the performance of the human sense of touch. Consequently, the design of haptic stimuli currently depends as heavily on the restrictions imposed by the hardware, as it does on basic capabilities of our perceptual system. A brief review of the current hardware is presented below, drawing the distinction between devices appealing to the tactile sense and devices appealing to the kinaesthetic sense. The limitations of these devices are also outlined.

2.3.1. Tactile hardware

Existing tactile devices present stimuli to the surface of the skin; they perturb the skin, often with very small forces, and make no attempt to move any part of the body. They are typically designed for application to the hand or fingertip, the areas of greatest tactile acuity, but devices have been made which apply stimuli to other


areas of the body. Examples of such systems are those that integrate tactile output devices into belts or jackets in order to present spatially distinct cues to the torso to support navigation tasks [193]. Irrespective of the desired target area of the body, there are several different mechanisms by which tactile stimuli can be produced.

Shimoga [172] reviews tactile display technologies with an eye to telemanipulation, the remote control of robotic devices. He is concerned with any mechanism for the display of tactile information, and divides the available displays into five camps: visual display, neurological stimulation, pneumatic stimulation, vibro-tactile stimulation, and electrotactile stimulation. In his classification, visual displays are devices which present a graphical representation of tactile information [154] and neurological stimulation encompasses a set of devices which can apply electrical stimulation directly to the brain in order to create the illusion of tactile stimuli [175]. Neither of these two techniques fall into the definition of tactile hardware used in this thesis, as they do attempt to produce tactile stimuli, instead aiming to provide tactile information through some other means.

Electrotactile, pneumatic and vibrotactile displays all present information directly to the skin, and yield stimuli that are perceivable by the tactile sense. Electrotactile and pneumatic devices are usually attached to the skin, typically mounted on the hand or fingertip, and often encapsulated in gloves. Examples of devices that use these technologies are those developed by Zhu [208] and Sato [168]. Electrotactile displays work by applying small currents to the skin, creating tingling sensations of varying intensity. Pneumatic displays work using either air jets aimed towards the skin, or air pockets resting against the skin. Activating the air jets or pressuring the air pockets results in differing tactile sensations of pressure.

Vibro-tactile devices are the most common sort of tactile device and create feedback by vibrating an object against a user's skin. These displays are typically embedded within an object and the stimuli are presented when a user places his or her skin against the display; they are often constructed in such a way that a user is not required to don special equipment. This flexibility makes them suitable for a wider variety of situations than other tactile display technologies. The simplest form of vibrotactile display involves a single oscillating element, which a user comes in contact with to perceive the vibrations. These displays have been constructed by

Figure 2.2. The Tractile vibrotactile device (reprinted without permission from Campbell *et al.* [38]).



adapting small speakers, and also by more complex technologies such as specialised vibrating pins. Single element vibrotactile displays have been integrated into a variety of devices, including both mice [3] and isometric trackpoints [38] (pictured in Figure 2.2). More sophisticated devices using arrays of oscillating pins have also been produced, and these have the potential to display more complex information. Indeed, such systems have been used to display Braille [71]. The BrailleNote [101], a typical device, is illustrated in Figure 2.3. One significant limitation to these multi-pin devices is that they typically have very low refresh rates (when the whole array is considered), often in the order of one or more seconds, or greater.

The range of tactile hardware currently available fails to provide accurate tactile feedback – no device has been produced that can even begin to match the capacity of the human sensory system. The stimuli that existing tactile devices are able to

Figure 2.3. The Braillenote, a tactile display for visually impaired users (reprinted without permission from www.brailnote.com).



produce does not attain the resolution and sensitivity required to simulate realistic tactile cues. Arguably then, the aim of existing tactile display technologies is not to provide realistic feedback but instead to facilitate feedback that is sufficiently realistic that the mapping a user must make between stimulus and meaning is not wholly symbolic. For instance, while current tactile devices are unable to present realistic feedback to support the action of gripping an object between thumb and forefinger, they can instead apply their unrealistic feedback to the appropriate sections of the hand. This basic feedback provides confirmation that an object has been gripped, and conveys information as to what parts of the hand are in contact with the object, despite bearing little resemblance to the sensation of gripping an object in the real world.

2.3.2. Kinaesthetic hardware

Devices designed to apply kinaesthetic stimuli are known as force feedback devices. To provide realistic kinaesthetic stimuli they need only present the net forces and torques involved in interacting with virtual objects – they are not concerned with feedback that palpates the skin. Typically this means that a force feedback device consists of some armature designed to apply precisely controlled forces to a specific part of its mechanism. To experience stimuli presented by the device, a user holds this part, known as a surrogate, and through it can feel the forces the device exerts. The experience for a user can be likened to exploring the world through the intermediary of a tool; the feel of the tool against the skin remains the same, while the forces felt through the tool vary according to the objects it encounters. Fundamental to the design of force feedback devices (and the stimuli they display) is the concept of *distal attribution*. This term, coined by Loomis [120], refers to the ability people possess to extend their sense of touch through an object they are manipulating so that subjectively they do not sense their contact with the object, but instead the object's contact with the environment. Examples of this phenomenon are in driving, where users sense the vehicle's contact with the road and not their contact with the seat or steering wheel, and in tool-mediated manipulation (such as using a screwdriver), where users perceive the tool's interaction with the environment rather than the force the tool exerts on their body. This principle underpins the ability of force feedback devices to create compelling simulations of reality, and is important in the design of stimuli for display on such devices.

The generation of kinaesthetic feedback has been found to be highly sensitive to refresh rate issues. Early work by Minsky *et al.* [137] indicated that the refresh rate required to smoothly render a surface was a function of the desired stiffness of that surface; stiffer objects required higher refresh rates, typically of 500 Hz or 1000 Hz. This observation is a consequence of the sensitivity of the kinaesthetic and tactile senses, combined with the properties of the basic spring algorithms used to generate surfaces. In a haptically presented surface rendered by a linear spring the force applied (f) is a function of the stiffness of the surface (k) multiplied by the distance a user is displaced within that surface (d): $f = k * d$. A compliant surface, with a low value of k , yields forces which change gradually over distance and can be rendered at lower refresh rates than a stiff surface, with a high value of k , in which forces change very rapidly over small distances. Rapidly changing forces require high refresh rates to be smoothly rendered. Minsky *et al.* [137] observed that instabilities, in the form of vibrations, were the consequence of attempting to produce a stiff surface with a low refresh rate. With a stiff surface, the change in force between one refresh of the system and the next can become too great to provide the illusion of a smooth transition. As the tactile system is extremely sensitive to vibrations [194], these unwanted by-products rapidly destroy the illusion of a haptically realistic surface. Needless to say, these high refresh rate requirements have strongly influenced the development of force-feedback hardware and software, all of which must function in a control loop with a duration of 1-2 milliseconds. This demanding performance requirement is one significant reason why force-feedback technologies only became popular in the relatively recent past.

A second crucial consequence of the spring model is that to render surfaces, forces must be uniformly applied to the surrogate as it changes position. The ability to do this enables a device to exert constant feedback, for instance, as a user pushes against, and sinks into, a surface. This consideration has heavily influenced the physical design of force feedback technologies.


In order to provide a detailed and structured review of force feedback devices, here they are classified into the following categories: enabling technology, grounding mechanism, and number of degrees of freedom possessed. In the next few

paragraphs, each of these distinctions is explained and discussed, and their advantages and limitations are highlighted.

Three main technologies exist to enable force feedback. Electromagnets can be used to influence the positioning of an iron surrogate, brakes can be used to create resistance, and motors can be used to apply forces. Early electromagnetic systems took the form of augmented mice incorporating solenoids that were used in conjunction with an iron mouse mat [3]. Akamatsu & Sate describe a typical device, which is illustrated in Figure 2.4. Activating the solenoid generated a crude velocity damping effect. The latest devices to employ this technique are sophisticated 6 degree of freedom free-floating devices such as the Maglev [19] (shown in Figure 2.5). To operate these devices a user holds a surrogate floating in a electromagnetic field. The device applies forces by adjusting the electromagnets generating this field, influencing the position of the surrogate. The absence of moving parts in these systems means that they can provide a high quality of feedback. However, they are typically limited to a small workspace, a consequence of the short effective range of magnetic fields.

Braking systems provide force feedback by restricting a user's motion [163]. While this technique allows the presentation of virtual objects, it does not support the exertion of active forces - forces actually pushing or pulling a user about. This is a substantial limitation in the majority of situations, but can be useful in some cases, for instance in systems designed to prevent surgical error and reduce unnecessary trauma [171]. In such a system the movements of a surgeon could be restricted to a

Figure 2.4. Akamatsu's haptic feedback mouse (reprinted without permission from Akamatsu [4]).



pre-defined critical area, or small motions due to tremor could be damped, reducing the potential for unnecessary damage to tissue. As such a system would not have the ability to actively move the surgeon, it could not itself be a source of surgical error. Motor-based systems are the most common and successful force feedback technology [157, 165]. There are two basic mechanisms for enabling motor-based force feedback. Forces can either be transmitted from the motors to the surrogate via a system of gears, or via cabling. Both of these approaches are used, although cable systems typically provide a better quality of feedback, as they have low friction and are back-drivable.

The majority of force feedback devices suffer the disadvantage that any force they apply entails the production of an equal and opposite force, which must be grounded, or absorbed, to prevent an imbalance. In other words, as a device generates a force in a certain direction, it experiences an equal force in the opposite direction, and must possess the (typically mechanical) ability to negate, or ground, this or be subject to motion in accordance with this force. Grounding can be achieved by applying the extraneous, opposite, force to either the user, or to the earth. This provides a classification of devices into body-based or ground-based devices. Body-based devices are typically designed to be held or worn, while ground-based devices have a relatively stationary aspect and often incorporate a heavy base to lend stability against the application of grounding forces.

Figure 2.5. The Maglev haptic device (reprinted without permission from www-2.cs.cmu.edu/afs/cs/project/msl/www/haptic/haptic_device.html).





Figure 2.6. Rutgers Master II-ND force feedback glove, a typical body based device (reprinted without permission from Bouzit *et al.* [23]).



There are advantages and disadvantages to both approaches. Body-based devices provide great freedom of movement and potentially a workspace that coincides with that of the user. However they suffer the disadvantage that users will always be subjected to the grounding forces generated by their interactions with the virtual world, and this may reduce the effectiveness of the presented stimuli. This issue becomes especially significant with regard to the large forces required to simulate rigid objects [188]. As users push hard against an object, a similarly large but opposite force will have to be grounded on some part of their body. Furthermore with complex and consequently often heavy devices, user fatigue can become a problem, preventing prolonged use. Finally, the engineering constraints for sophisticated body-based devices such as exoskeleton arms that are designed to fit around parts of the human body are very complex. A consequence of this is high prices. The Rutgers Master II-ND force feedback glove [23], a recent and typical high-end body-based device, is illustrated in Figure 2.6.

Ground-based devices, on the other hand, suffer the disadvantage that they are not portable during the course of a user's interaction with them; the user is confined within the workspace dictated by the physical range of the device. This is offset, however, by the fact that users will feel no discrepant grounding forces, nor suffer fatigue from the mass of the device, and finally by the far simpler engineering constraints imposed in the creation of a ground-based device. Ground-based devices are not required to fit around the human body, thus simplifying the design. This is reflected in lower costs. Figure 2.7 shows the HapticMaster, a ground based device.

Figure 2.7. A typical ground based haptic device, the HapticMaster (reprinted without permission from www.fcs-cs.com/robotics/).



Beyond these distinctions, force feedback devices are characterised chiefly by the number of independent axes down, or around, which they can exert force; in short, the number of degrees of freedom that they possess. Available devices range from those capable of producing non-directional forces, such as large-scale vibrations, to six DOF devices that can independently activate high fidelity forces along and around all three spatial axes. A range of these devices is described below, separated into consumer devices, mainly aimed towards the entertainment market, and the much more sophisticated specialist and research devices.

The games industry has recently adopted force feedback with enthusiasm. A variety of devices are now available, and are becoming ubiquitous. The simplest force feedback devices, available for game consoles, are body-based, motor-activated devices that allow a variable sensation of vibration to be presented [98] (illustrated in Figure 2.8) A user holds the game controller in his or her hand, and motors spin weights to vibrate the entire device. Typically the frequency and amplitude of this spin can be adjusted, and the feedback is primarily used to indicate an impact of some sort. A number of mice using similar technology have recently been introduced

Figure 2.8. The Wingman Extreme vibration device (reprinted without permission from www.logitech.com).



[99]. These devices are again intended for gaming use, but also support a haptically enabled desktop, in which moving over targets is mapped to various sensations of vibration. A similar technology is also used in the vast majority of mobile phones. This allows users to configure their device to signify the presence of an incoming call or message with a haptic buzz, rather than an audible ring tone.

More sophisticated motor activated, ground based gaming devices that can apply directional feedback are also popular. These range from one degree of freedom force feedback devices, such as force feedback steering wheels [98], to more general

Figure 2.9. The Wingman force feedback joystick (reprinted without permission from www.logitech.com).



Figure 2.10. The Wingman force feedback mouse (reprinted without permission from www.logitech.com).



purpose force feedback joysticks, which activate force along two axes [98] (shown in Figure 2.9). Although nowhere near as popular, two DOF force feedback mice are also available. A typical example of this type of device is the Wingman Force Feedback Mouse [118] (illustrated in Figure 2.10). This device is mounted on a fixed platform and forces are applied through the mounting. Forces are enabled in games, but the haptic augmentation of other tasks is also supported. The force feedback has been integrated into the desktop, with the goal of providing targeting benefits, and into applications such as drawing packages, to aid complex manipulations. One potential reason why mice featuring force-feedback have not attained the same level of commercial success as similarly marketed and produced gaming technologies is the fact that adding haptic cues to a GUI is a challenging task, and one not yet fully addressed by current research. Indeed, some studies have [4] have indicated that the addition of different haptic cues can lead to either increases or decreases in task performance levels in typical GUI tasks.

The scope of haptic devices within specialist markets and a research context is much greater. Body-based glove systems such as CyberGrasp [97] (shown in Figure 2.11) present one DOF force feedback separately to each finger, allowing the creation of a compelling sensation of grasping an object. Such systems are designed to support the manipulation of objects in virtual environments. Force feedback devices for the simulation of specific surgical procedures have also been created. Examples include devices that have been designed explicitly to support simulations of laparoscopic surgery, such as the laparoscopic surgical workstation pictured in Figure 2.12 [100].

Figure 2.11. The Cybergrasp glove based force feedback device (reprinted without permission from www.immersion.com).



Figure 2.12. The Laparoscopic Surgical Workstation, a haptic device for the simulation of laproscopic surgery (reprinted without permission from www.immersion.com).



Many more general-purpose force feedback devices also exist. Such systems are usually ground-based and motor activated. Examples include SPIDAR [22] (shown in Figure 2.13), a 3 DOF system consisting of a ring surrogate suspended in space by cables. Forces are exerted down the cables to move the ring. This system has a relatively novel design, and the advantage that it can scale up to large workspaces,

Figure 2.13. The SPIDAR haptic device (reprinted without permission from Bouguila *et al.* [22]).

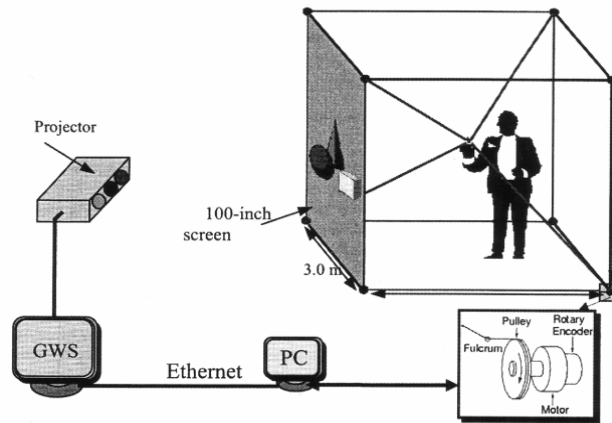


Figure 2.14. The Delta haptic device (reprinted without permission from www.forcedimension.com/).



encompassing body sized spaces. Other examples typically exert force down a robotic exoskeleton. Examples are Delta [69] (illustrated in Figure 2.14), a six DOF device with a relatively small workspace, designed to be mounted on a desk, and the Haptic Master [192] (shown earlier in Figure 2.7), a floor mounted device with a workspace designed to be shoulder-sized, and the capacity to exert extremely strong forces (of 100 N or greater). Devices merging several kinaesthetic technologies have also been created. The CyberForce [96] (shown in Figure 2.15) is a glove system capable of providing a one DOF grasping sensation to each finger, mounted on the tip of a three DOF exoskeleton device. The most popular and widespread force feedback controller in use in a research context is the PHANToM from SensAble Technologies Inc. [128]. Two very similar models of this device (the PHANToM 1.0 and 1.5) are used extensively in the studies reported in this thesis, and are described in detail below. Differences between the two models are highlighted.

2.3.3. The PHANToM

All PHANToM devices take the form of a triple jointed robot arm. A PHANToM 1.0 is pictured in Figure 2.16. Both the 1.0 and 1.5 models can exert force along all three spatial axes and are operated by manipulating a surrogate attached to the tip of the arm. This surrogate is interchangeable; with modification any appropriately sized object can be used. Typically the surrogates are chosen to match the application domain – in a surgical simulator, the handle from a relevant medical instrument can be used. Each surrogate attaches to a PHANToM via a frictionless gimbal with 3 rotational degrees of freedom, all coinciding on a single point. This point is also the

Figure 2.15. The CyberForce haptic device (reprinted without permission from www.immersion.com).



Figure 2.16. The PHANToM 1.0 haptic device (reprinted without permission from www.sensable.com).



focus of the force feedback. This spatial conjunction of the passive rotational degrees of freedom and the active axial forces leads to the beneficial property that the net force applied can always be sensed without bias, irrespective of the orientation of the gimbal. A user is able to rotate the surrogate freely and this action will not affect the displayed force.

The PHANToM 1.0 and 1.5 are of extremely high fidelity and each model can sense the position at the tip of its arm at a resolution of 0.03 mm. They can exert a peak force of 8.5 N in strength independently on each axis. The maximum continuous force is 1.4 N. The backdrive friction, the impediment to free motion produced by the presence of the motors, is small at 0.04 N. When operating they have an inertia (an apparent mass at the tip) of less than 75 grams. The PHANToM 1.0 has a footprint of 18 by 25 centimetres and a workspace of 13 by 18 by 25 centimetres. The PHANToM 1.5 has a footprint 25 by 33 centimetres and a workspace of 19.5 by 27 by 37.5 centimetres. Drivers guarantee a refresh rate of 1000 Hz, sufficient to provide a smooth haptic sensation in demanding conditions [137]. Another larger version of the PHANToM exists with similar, but spatially expanded statistics. Recently, prototype versions of the PHANToM have become available which are able to produce 6 DOF force feedback. These devices are not discussed nor used in this thesis, as most of the work described here predates their appearance.

Standard versions of the PHANToM cannot exert rotational forces, a deficiency discussed in the next section, but with this substantial exception, the feedback it produces can match the majority of the requirements of the human kinaesthetic system. The PHANToM lacks somewhat in the strength of the forces that it can both sustain and maximally produce. It has been observed, however, that the forces a user typically exerts when interacting with a PHANToM are small, typically around 1 N [128]. This is substantially lower than the exploratory forces observed in real world interactions, and is probably a consequence of user perceptions of the fragility of the device itself.

2.3.4. Point Force

A crucial consequence of not incorporating rotational force feedback into a kinaesthetic device is in terms of the representation of a user's presence in the simulated world, the properties of the virtual avatar through which they experience

forces. A device that applies only axial forces is capable of representing stimuli that would be experienced by a single point of contact, or a sphere, but not of any more complex object. Massie [128] coins the term ‘point force’ to describe such interfaces. Adding rotational DOFs enables the simulation of more complex avatars, more reminiscent of real world tools. An illustrative example comes from the surgical domain, where exploration is often conducted using thin cylindrical manipulators. The rotational forces experienced by such a tool provide a surgeon with vital information. However, a point force device is not able to simulate these key forces; it could not represent the rotational forces that result from multiple points of contact. Instead it would only be able to display the forces accessible from a single point at the tip of the tool. This fact limits the usefulness of point-force devices in a number of domains.

Addressing this weakness, Massie [128] discusses the implications of point force at length. He justifies devices that produce point force with the assertion that: ‘*many meaningful haptic interactions involve little torque*’. He states that if users can be provided with three passive rotational degrees of freedom then they can comfortably interact with the remaining three axial ones, and that a large number of haptic interactions do not require the presence of rotational forces. He states that although a point force interface may not be suitable for the simulation of all interactions, there are many that it can convey compellingly.

Accepting this assertion, there are also several advantages to point force interfaces. Foremost is that of simplicity. Hardware design is made easier if only axial forces are desired. The creation of simulations which treat the user as a point also avoid the majority of the issues involved in complex collision detection [206], which, at the high refresh rates required to produce kinaesthetic stimuli, can be very challenging.

There are several other features of point force interactions. If a user is represented as a sphere in the virtual world, varying the size of this sphere can effectively change the resolution at which the haptic scene is rendered. As the sphere becomes larger it will cease to fit in small gaps, providing an illusion of smoothness. Using a sphere as a cursor can also smooth unintentional defects such as those arising from inaccuracies in floating point mathematics. Another consequence is that a friction effect cannot be accurately simulated using a point force interface – a totally realistic

friction would involve torque about the centre of the sphere representing the user. Fortunately this is a minor component of friction and its absence has not prevented successful and compelling friction implementations in point force systems [128, 135]. A final, and potentially cognitively disturbing, side effect of point force systems is that as users are represented by a single sphere (or point) in the world they can easily move their hands inside virtual objects while the simulation constrains the tip of the surrogate, the sphere, to remain on the surface. This physically impossible situation can lead to confusion on the part of some users while others immediately capitalise on their new found freedom to explore an object in novel ways [128].

2.4. Application Domains

Haptic feedback has attracted a large amount of interest and proven successful in a number of application domains. Viewed as a whole these previous successes exert a substantial influence on the creation of novel systems as they can be viewed as a practical, but very loosely defined, set of *ad hoc* guidelines as to the use of haptic feedback. Studying these topics can provide an indication of where, and where not, haptic feedback should be applied. While any insights gained from such an examination are unlikely to have the benefits of a grounded or rigorous theoretical understanding, nor invariably be accurate, a broad knowledge of the field seems likely to positively inform the development of systems in novel application domains. It is to this end that this section briefly reviews some of the areas in which the role of haptic feedback has previously been investigated.

One of the earliest application areas in which haptic feedback flourished was telemanipulation [64]. Telemanipulation is the remote control of some distant robotic device, called a slave, by manipulating a local robotic device, a master. For instance, if both the slave and master devices were robot arms then movements a user would make with the master device would be reflected in the slave, potentially providing a remote control arm with a transparent interface. Interest in telemanipulation has traditionally stemmed from the desire to perform complex, human controlled, physical manipulations in inhospitable environments, such as deep under water [196] or in space [49] or for dangerous activities such as bomb disposal [68]. In this kind of manipulation, haptic feedback has typically been used to provide a representation

of the objects the slave physically encounters aimed at our sense of touch, theorising that this direct representation would be beneficial.

Another early use of force feedback was scientific visualisation, specifically molecular docking [29]. Molecular docking systems allow chemists to virtually assemble molecules according to the rules governing chemical bonds. They can then manipulate atoms, typically by virtually picking them up and repositioning them with the cursor. As they do this there is a continual evaluation of the constraints on the atom being repositioned, preventing the creation of impossible structures. Haptic feedback has been shown to be a useful addition to the visualisation of these constraints, as the resistance to motion possessed by an atom can be simply translated into a force vector. Haptic molecular docking systems such as GROPE III [29] and more recently Sculpt [197] have proved highly successful, providing chemists with what some authors have described as “intelligence amplification” [140] whereby haptic feedback conveys a greatly increased understanding of the constraints and tensions in place between the atoms.

Haptic feedback has also been applied to other visualisation tasks. It has been integrated into scientific data visualisation packages [60], to allow researchers to haptically explore their data. There is also a growing body of literature detailing the potential benefits of haptic interfaces for the visually impaired, for instance in the presentation of mathematical data [205], or musical notation [39]. Haptics has also been integrated into 3D design systems. Researchers have considered haptics as a mechanism for supporting object assembly [41] and as an integral mechanism for creating the objects themselves [42].

However, perhaps the largest research effort in the field of haptics today is that of surgical training. The majority of surgical procedures are essentially skilled physical tasks, in which practitioners rely heavily of their sense of touch [31]. Traditional methods for training surgeons involve a lengthy apprenticeship period in which a substantial number of operations are firstly observed and then performed. This process often follows the routine: see one, do one, teach one. This method is not without danger to patients [76] and can also prolong operation times and incur substantial costs. Furthermore suitable patients for training are not available on demand, and this can lead to the omission, or reduction, of training in particular

procedures. Haptic virtual reality simulations have the potential to augment traditional training methods, by providing simulations of procedures in which students can practice. These simulations are by necessity highly complex – to provide valuable training they must achieve realism. This will often entail simulating the feel of a variety of medical implements as they grasp or cut and also simulating the behaviour of deformable organic surfaces. These factors mean that surgical simulations require high precision hardware and are often extremely algorithmically complex. Dedicated hardware to support surgical simulations is continually under development [100], as are new algorithms for modelling objects with complex behaviours, such as deformable surfaces, or for increasing the efficiency of collision detection to facilitate the use of arbitrary cursor objects [33]. Recent surgical training simulators are undergoing evaluations that seem promising [76].

Systems that support the performance of actual surgical procedures are also gradually gaining acceptance [171]. These systems typically either restrict a surgeon's movements to a specific critical area in order to prevent the occurrence of trauma through surgical error, facilitate laparoscopic or microsurgical operations that have previously only been achievable through more substantial surgeries, or are designed to support the remote execution of medical procedures [36]. This kind of tele-surgery promises significant benefits for isolated communities, and in dangerous or inhospitable situations such as battlefield scenarios.

A final application area where force feedback has proven popular is simply interacting with virtual worlds, typically for entertainment. The most obvious market for this sort of interaction is the gaming market and modern game consoles are beginning to be supplied as standard with rudimentary force feedback controllers, while personal computers can be equipped with fairly advanced 2 degree of freedom force feedback joysticks at little cost. One of motivations for force feedback in entertainment is to heighten the sense of realism, to increase a user's involvement with the simulated environment. Indeed, research has suggested that haptic feedback is linked to increased feelings of virtual presence [61], of immersion in a virtual world, but otherwise this aspect of haptic feedback has been relatively unexplored.

2.5. Terminology

Due to the infancy of the field, the literature on haptics is characterised by a wealth of terminology, often applied to different ends. The specific terminology adopted in this thesis is presented in Table 1.1. Originally presented by McGee [129] (and subsequently in Oakley *et al.* [151]), these definitions draw on concepts introduced earlier in this chapter, and hold throughout this thesis.

2.6. Conclusion

This chapter has broadly reviewed research in the field of haptics with the goal of providing an informed background knowledge of the topic in order to support the design of novel haptic cues for use in novel situations. To achieve this goal, three distinct areas were reviewed: haptic psychophysics, haptic hardware and haptic applications. A thorough understanding of the first of these is critical for the creation of any haptic cue as it describes and explains the very sophisticated capabilities of the human sense of touch. It provides background regarding what and how we feel. A review of haptic hardware, the second of these areas, is also important as the strengths and limitations possessed by current display devices fundamentally affects the nature of the stimuli that can be rendered. Understanding these capabilities is vital for the creation of functional and useful haptic cues. Finally, an understanding of the role of haptics in a variety of application areas is important, as it can serve as an informal, rule of thumb, guide as to where and how haptic information is beneficial, and where and how it is not.

Table 2.1. Definitions of terms used in this thesis.

Term	Definition
Haptic	Relating to the sense of touch.
Proprioceptive	Relating to sensory information about the state of the body (including cutaneous, kinaesthetic, and vestibular sensations).
Vestibular	Pertaining to the perception of head orientation, acceleration, and deceleration.
Kinaesthetic	Meaning the feeling of motion. Relating to sensations originating in muscles, tendons and joints.
Cutaneous	Pertaining to the skin itself or the skin as a sense organ. Includes sensation of pressure, temperature, and pain.
Tactile	Pertaining to the cutaneous sense but more specifically the sensation of pressure rather than temperature or pain.
Force Feedback	Relating to the mechanical production of information sensed by the human kinaesthetic system.

The review presented in this chapter is important for the research described in the remainder of this thesis. One of the significant aspects of this thesis is its focus on the design of haptic cues for a variety of application areas, and this review is geared towards this end. The feedback described throughout the rest of thesis was created with an eye to the nature of the human haptic sensory system, an understanding of the capabilities of currently available display devices and bearing in mind the lessons that can be learnt from previous uses of haptic feedback in computer systems.

The review presented here also serves to justify the secondary aim of this thesis – to further design knowledge on the topic of haptic interfaces – as it highlights how much literature must be absorbed to gain an understanding of this topic. In order to create appropriate and valuable haptic feedback, a basic comprehension of the broad body of research discussed in this chapter is required. Furthermore, the understanding gained from this literature is no substitute for the kind of explicit design guidelines present in many other domains; at best it abstracted from real interaction tasks (in the case of the literature relating to psychophysics and hardware) or informal and rule of thumb (as that gained from the previous application domains). It is clear that the creation of concise and explicit design guidelines would facilitate the adoption of haptic feedback by designers and system developers and help ensure that it appears in future interfaces to computer systems. One of the goals of this thesis is to take steps towards this objective through the development of design guidelines relating to each of the domains studied in depth.

3. Haptic Augmentation of the Desktop

3.1. Introduction

This chapter reviews the literature relating to the haptic augmentation of Graphical User Interface widgets. It begins with a discussion of the motivations for adding touch to desktop interfaces borrowing from, and briefly covering, the related literature on audio augmentations. The subsequent literature review is split into two parts. The first, and more substantial, deals with the general literature and more specifically what is termed here as *single-target* interactions, while the second looks at *multi-target* interactions. The difference between these two areas is that the section on multi-target interactions is focused on the literature that explicitly considers the addition of haptics to groups of widgets while the section on single-target interactions makes no mention of this and examines either systems not backed up by empirical work or studies of widgets in isolation.

The addition of haptic feedback to GUI widgets is an important aspect of this thesis. It forms the first of the two exemplars chosen to illustrate the general claim that haptic cues relating to a user's representation (or avatar) in a computer system can improve a user's performance with (or experience of) that computer system. As a review of the literature in this area, this chapter provides a context in which the designs and evaluations of haptic widgets presented in the next chapter can be considered. It relates the current thinking on this topic, and represents the standpoint from which the empirical research described in the next chapter was conducted.

3.2. Motivations

The addition of non-visual information to interactions in GUIs is the focus of a growing body of research. While in recent years there have been a number of systems (reviewed in the next section) that have looked at haptic cues in this situation, there is a more established body of work examining non-speech audio augmentations. Consequently, much of the rationale for research into haptically enhanced desktops stems from this earlier audio work. This section describes the motivations for adding haptics to GUIs with an eye to the more established literature on audio augmentations.

Probably the most fundamental reason for augmenting a GUI with non-visual information is to increase user performance, the rationale being that the additional information will make tasks in the GUI easier. Typically, an increase in performance can occur either in terms of faster task completion times or lower error rates. A related aspect is the reduction of subjective workload [88], the amount of effort users feel they have to exert in order to perform a task. Brewster and colleagues [27, 28, 121] present a significant realisation of this motivation. They describe a series of studies investigating user performance with a variety of different sonically enhanced widgets including buttons and scrollbars. They report, in particular situations, improvements in task completion time, error rates and subjective workload. These results strongly support further work investigating the addition of non-visual information to GUIs, and are a strong motivational factor in the investigations of haptic widgets presented in this thesis. One significant difference between the audio and haptic modalities in the scenario of desktop interactions is that although audio feedback can notify users of the occurrence of errors, haptic feedback has the potential to go further and provide forces that physically prevent the occurrence of errors. This potentially beneficial property will be explored in this thesis.

Non-visual information can also be used to overcome problems of visual overload, a situation that can occur when so much information is being presented graphically that a user cannot simultaneously attend to all of it [27]. In this situation, the use of non-visual information can free up a user's sense of sight, allowing him or her to work more effectively. A significant example of visual overload relates to interactions with widgets that occur in the periphery of a user's attention. For

instance, in a task involving working with text, a user typically concentrates on studying the text and interacts with various widgets to perform operations (such as cut, copy, paste, save and search) in the sidelines, without fully focusing on them. This behaviour can lead to users experiencing problems (in terms of increased errors and task completion times) interacting with the widgets as their attention is directed elsewhere. Studies examining non-speech audio feedback have shown that it is possible to reduce the amount of visual attention required to manipulate widgets, which may alleviate this problem [27]. As another channel for the transmission of non-visual information there is the potential that haptic feedback could be used to gain similar benefits.

A related motivation for using non-visual feedback is to display information that is simply not presented in a GUI, often due to the additional visual clutter that would result from its inclusion. The literature on non-speech audio augmentations provides several significant examples of this. For instance, in one of the studies mentioned above Brewster *et al.* [28] observed that the visual presentation of a button in a GUI does not provide feedback that differentiates between all possible outcomes of an interaction. Critically, they isolated a situation in which failure to select a target results in identical visual feedback to a successful targeting operation. They then designed and evaluated audio feedback that distinguished between these two cases, and observed an increase in user performance they attributed to presence of this extra information about the system state. Another powerful example of the use of non-speech audio to transmit additional information about a system is Gaver's [73] audio icons. In this system graphical objects were associated with sampled sounds that reflected their properties. For instance, in one scenario files were associated a wooden sound, and this sound varied according to the size of the file. When a small file was selected a wooden sound that resembled one that might emanate from tapping a small wooden object was played, and conversely when a large file was selected, the sound of tapping a large wooden object was played. An important aspect of audio icons was that the audio feedback was carefully designed to leverage this kind of metaphor: there was typically a strong intuitive link between the meaning of a sound in the real world and its meaning in the virtual world of the GUI. While little evaluation of this system was presented, it is a compelling example, and

serves to highlight the possibilities afforded by using non-visual information to reveal hidden aspects of GUIs.

Finally, interfaces including non-visual information may be useful for impaired users. Specifically, the literature suggests that haptic interfaces have the potential to support interactions with both visually impaired users, and motion impaired users. The most striking example of an existing haptic interface for visually impaired users is Braille, the notation system that presents a haptic representation of written language, but the scope for the development of other displays seems large. For instance, work has begun examining the presentation of the much more complex structure of musical notation in a haptic form [39]. Understandably, haptic interfaces for motion impaired users focus less on information presentation and more on providing forces that support movement. The majority of existing examples come in the form of systems for rehabilitation. These provide haptic feedback designed to serve as a physiotherapeutic aid and include the Rutgers Ankle Master [53], an actuated foot pedal used to control a flying simulation and intended to help users recover from the motion impairments caused by a stroke, and several devices for the rehabilitation of the hand or wrist [30, 32]. While the examples described here do not directly relate to GUI tasks, they do highlight the applicability, and potential benefits, of haptic feedback to these user groups.

3.3. Single Target Haptic Augmentation

Empirical studies investigating the effects of haptic feedback on the performance of users engaged in typical GUI interaction tasks, such as target acquisition, began with simple augmentations of standard pointing devices in the early 1990's. In 1994 Akamatsu & Sate [3] developed a haptically enabled mouse with the ability to produce both tactile and force feedback. The tactile feedback was generated using a vibro-tactile technique. It was created by the motion of a pin positioned so as to slightly protrude through an area at the tip of the left mouse button, an area of the mouse over which a user's finger typically resides. The force feedback took the form of a simple electro-magnetic friction effect. This effect was created by embedding a solenoid in the housing of the mouse and using it in conjunction with an iron mouse mat. Activating the electromagnet caused the mouse to be attracted, and consequently to stick, to the surface of the mat. Using this mouse they conducted a

study investigating user performance in a target acquisition task. Each trial in their study consisted of participants moving from a preset starting point to a target, and then selecting that target. The study compared a control condition with a condition incorporating the use of haptically augmented targets. The haptic augmentation included both tactile and force display. The results indicated that the haptic condition led to small (in the order of 7-10%), but significant, reductions in task completion times, and Akamatsu & Sate concluded that haptic feedback has the potential to improve user performance in GUIs.

A follow up study was conducted by Akamatsu & MacKenzie [4] using the same multimodal mouse, but including two extra feedback conditions: one involving the display of only the tactile feedback, and one featuring only the force feedback. The results of this study essentially reinforce those of the first – that haptic feedback can exert a beneficial influence on user performance in GUI tasks – but are not so clear-cut. The sole presentation of tactile feedback was found to provide reductions in task completion time when compared to the control (5.6%), but could actually increase the error rate substantially (by up to 65%). Sole presentation of the force feedback, on the other hand, led to equivalent temporal performance to that observed in the control, but yielded reduced errors rates (by 12.1%). The combined force and tactile condition resulted in the quickest times (an improvement of 7.6% over the control), at the cost of a still substantial increase in the error rate (of 30.3%). The authors conclude that although haptic feedback appears to have the potential to provide performance improvements to users engaged in typical GUI tasks, it can also reduce performance. They suggest that the application of haptic feedback to these environments, if it is to be useful, will be a complex process.

Also in 1994, Engel *et al.* [65] built and investigated a trackball that incorporated 2 DOF force feedback. It had the ability to exert rotational forces independently around both x and y axes. They describe several studies conducted using this device. These included a maze navigation task in which the addition of force feedback cues representing the maze walls led to substantial reductions in task completion time (by 37%) and the rate of occurrence of boundary (or maze wall) crossing errors (by 76%). Acknowledging that the results from this maze task would not generalise well to all cursor activities, they also studied a target acquisition task. In this study

participants were required to select a randomly positioned button as rapidly as possible, and a control condition was contrasted with one in which attractive haptic forces were in place over the target. The results of this study reinforced those of the maze study; they demonstrated significantly reduced task completion times and error rates. The authors also note that the results, and their own anecdotal experience, appear to indicate that the time taken by users to learn how to use the trackball was dramatically reduced with the addition of force feedback.

Although the results cited in these studies seem promising, the devices that feature in them have now been superseded. More recently devices with a richer capacity for feedback have emerged, and in some cases have themselves been superseded. For instance, a commercial product, the Wingman Force Feedback Mouse [118], has also been developed. This is a mouse permanently mounted on a special base containing an assemblage of motors able to provide 2 DOF force feedback in a 2 inch square area. Drivers were provided for this device that integrated sophisticated haptic effects into Microsoft Windows. Buttons, icons and hyperlinks were enhanced with a snap to effect, which served to draw the cursor to their centre. Scrollbars and windows were bevelled into the surface of the workspace, providing haptic walls around their perimeter. Finally, isometric input mechanisms, which depend on the pressure a user exerts, were implemented in software and enabled for pressure sensitive selection and scrolling.

Some formal evaluation of basic haptic augmentations took place using this device. Hasser *et al.* [90] describe a target acquisition study comparing user performance in a condition in which targets were overlaid with an attractive force against those attained in a condition with no additional haptic feedback. They reported that the haptic feedback led to significant improvements in task completion time. Dennerlein *et al.* [51] continued this work, and examined performance in two more complex tasks. The first of these involved moving down a tunnel, and the second was a compound task consisting of moving down a similar tunnel, and selecting an item positioned at its far end. These tasks were intended to be analogous to menu interaction. In these studies a haptic condition featuring forces that pulled a user to the centre of the tunnel were contrasted with a condition with no additional feedback. The results revealed substantial reductions in task completion time (by 50%) in the

simple task of moving down the tunnel, and large, but reduced, improvements in the compound task of moving and selecting (of 25%). While this research provides positive evidence for the integration of haptic feedback into GUIs, it also highlights the complex nature of the influence that haptic cues can exert on user performance in this domain. Critically, in these studies, it is unclear why performance in the compound task is much reduced compared to the simpler task only involving movement. These unexpected effects underline the requirement for further study of this topic.

Rosenburg & Brave [162] briefly describe a pilot study involving a force feedback joystick and tasks involving interacting with buttons, menus and scroll bars. They explored the differences in performance achieved using what they termed active and passive haptic feedback. The active feedback involved the application of attractive haptic feedback over the interface targets (focused on a point in the case of the button, and on a line along the horizontal and vertical axes of the menu items and scrollbars respectively) while the passive feedback merely provided forces that resisted motion away from these critical areas. They found improvements in task performance time of up to 40% comparing their haptic augmentations against a control condition, but provide no formal analysis of their data. In general, they observed that the active feedback yields slightly improved results when compared to those achieved using the passive feedback. Although they do not formally examine it, one interesting aspect of Rosenburg & Brave's research is that they designed their system for use by motor-impaired users. Specifically they suggest the haptic feedback will be beneficial for users who experience difficulties in fine motor control. Such users can find working in a GUI very frustrating as it can be extremely challenging for them to successfully perform a basic targeting operation.

Recent research taking place at Cambridge University provides a more thorough examination of this topic. It includes studies characterising the movement capabilities of a group of users with different motor impairments [94] through to a number of investigations into the effects of different haptic feedback on this user group's performance in targeting tasks [108]. Using the Wingman force-feedback mouse described above, they have compared the performance attained by their group of impaired users when interacting with standard GUI targets against that achieved

when the targets are overlaid with what they term *gravity wells*, which are areas of attractive force, and also with a number of different types of friction, or resistive force. Their results are complex, due largely to the diverse motor skills of their subject group. For example, several of their users gain significant benefits (in the order of 50%) from both the gravity wells and the friction augmentations, while others only from the gravity, and yet others show little improvement from the addition of the haptic cues. Once again, these results highlight the fact that although haptic feedback has the potential to increase user performance in GUIs, its application is not simple, and performance improvements are not guaranteed. Indeed, this research suggests that it is more challenging to develop haptic augmentations that aid impaired users, than it is to create augmentations that help able-bodied users.

There have also been a number of systems developed where there has been little or no emphasis on evaluation. For instance, Ramstein *et al.* [158] produced the Pantograph, a ground-based two DOF force feedback device consisting of a surrogate that has free movement across a plane. Force feedback is exerted on this surrogate by motors through two exoskeletal arms. Several papers have been published discussing the use of the Pantograph in a variety of application areas including the kind of haptically augmented desktop discussed here [159]. In this desktop system the Pantograph was used to render two basic haptic objects: enclosures and frames. Each of these corresponded to different interaction objects in a GUI. Enclosures were haptically walled areas that exhibited a snap-to effect when a user moved over them, and provided resistance against any attempts to move off them. Enclosures were used to represent buttons, icons, and menu items. Frames resembled a thin groove and were presented over the borders of windows. They were designed to enable users to stay on the borders in order to manipulate, move and resize the windows. Unfortunately, no concrete evaluation of this system was presented.

Similarly, Zeleznik & Miller [135] extensively augmented the X Windows [169] desktop using the PHANToM [128]. Taking advantage of the PHANToM's three DOF force feedback, the surface of the desktop was represented as a horizontal plane, while icons formed dimples, or recessed areas and three-dimensional ridges separated menu items from one another. Windows were extensively augmented both

with haptic feedback and with pressure sensitive input mechanisms built on top of this feedback. For instance, users were able to adjust the position of a window by pressing into its surface and then moving laterally, much as objects are slid along a surface in a real environment. Collisions with other windows were also represented haptically, allowing users to neatly align them. Windows were also slightly dimpled and users could push against the lip of this recessed area in order to adjust the position of the window. Finally, users were able to raise and lower windows, in terms of their occlusion of one another, using vertical motions (motions orthogonal to the plane of the desktop). There is little evaluation of this system, save for a claim that in informal, observational tests of haptified drag and drop, users halved their task completion time, and become accustomed to the device very rapidly. Zeleznik & Miller acknowledge that this informal test is no substitute for rigorous user studies.

Finally, there have also been recent attempts to augment desktop interactions using tactile devices. Campbell *et al.* [38] describe a device they call the Tractile, a modified version of the isometric trackpoint device that is used for cursor control in many laptop computers. The functionality of this device resembles that of a small and very sensitive joystick. It typically appears as a cylindrical rubber protrusion positioned in the centre of the keyboard and is manipulated by a single finger. Campbell *et al.* added vibro-tactile feedback to this device through an electromagnetically activated pin mechanism situated underneath the rubber housing. They investigated user performance with this device in what they term steering tasks [1]. Steering tasks are operations in which a user has to move along a defined path, or tunnel, and are distinct from target acquisition tasks. Steering tasks have been likened to elementary GUI operations such as moving along a menu item in order to reach a submenu. Campbell *et al.*'s task involved participants moving along a visually defined horseshoe-shaped tunnel, under one of several conditions: a visual condition with no tactile feedback was compared to three conditions with additional tactile feedback. The first of these was a visual plus tactile condition in which visual dots on the screen matched haptically rendered bumps clustered along the line in the centre of the horseshoe's path. The second of these was an unconcerted condition in which visual dots were presented towards the borders of the horseshoe, while tactile stimuli were presented towards its centre. The final condition consisted of visually and haptically presented dots and bumps in a thin line on the borders of the

horseshoe. The results of this study revealed that the fastest task completion time was achieved in the visual and tactile condition, while the lowest error rate was found in the condition incorporating a narrow wall of bumps on the edge of the horseshoe. Unlike the other two tactile conditions the unconcerted condition resulted in no performance improvement over the visual condition. Campbell *et al.* suggest that this was due to the non-complementary nature of the visual and haptic cues presented in the unconcerted condition; the fact that the placing of the visual and haptic stimuli did not match. From this result they conclude that, in order to create effective haptic feedback, what you feel must be what you see. Their study also highlights the difficulties in creating useful haptic augmentations for desktop interactions, as the effects of haptic feedback on user performance are not simple or easily understood.

3.4. Multi Target Haptic Augmentation

The research reviewed thus far is mainly concerned with the presentation of single haptic targets – that is, individual haptic targets presented in an otherwise featureless haptic space – rather than the more realistic scenario incorporating multiple haptic targets presented simultaneously. This complex scenario is a more accurate portrayal of real interactions in GUIs as a typical interface currently involves the simultaneous presentation of tens, if not hundreds, of targets. In such a situation the influence exerted by haptic targets incidentally traversed by users as they move towards their desired destinations must be considered. It seems likely that the extraneous forces these widgets apply have the potential to alter the paths users wish to take, and consequently reduce their performance and subjective satisfaction.

One possible solution to this issue is to try to remove the unwanted haptic feedback by attempting to predict a user's desired destination, and applying the feedback only on this target. Such a calculation, if successful, would serve to reduce the complexity of the multi-target case to the simplicity of the single target case, and transfer the performance benefits gained there. However, as Dennerlein *et al.* [52] point out:

“...only enabling one force field is an unrealistic simulation for the implementation of force-feedback algorithms. If one confidently knew the desired target, why not then select that target automatically without using the mouse?”

According to this rationale, Dennerlein *et al.* [52] have begun to consider the implications of partially successful target prediction systems. They describe a study in which multiple targets are presented to users, and the number of haptic distracter targets active between a user and the destination target is manually controlled. They reason that adjusting the number of distracter targets simulates different accuracies of target prediction.

Their task involved the display of 13 evenly spaced circular targets arranged in a cross formation. Each trial consisted of one target becoming highlighted, which the user then had to select. They compared three haptic conditions against a control that featured no haptic feedback. The haptic feedback they used was in the form of an attractive force: a cursor moving over a target experienced forces drawing it to the target's centre. In the first haptic condition only the highlighted target was overlaid with forces. This simulated perfect target prediction and, therefore, the single target case. In the second haptic condition the highlighted target and one adjacent target were haptically active. This, they reasoned, might be the result of a relatively accurate target prediction algorithm. In the final haptic condition all 13 of the targets were haptically active. This condition represented a situation in which no target prediction was in use. Dennerlein *et al.*'s conclusions are mixed. As they use a similar set-up (in terms of the hardware used and forces generated) to Hasser [90], they replicate his results as regards objective measures in the condition resembling the single target case. They show a significant reduction in task completion time, relative to the control, in the condition in which a single target is haptically active. This objective gain is maintained, albeit at slightly reduced levels, in the other two haptic conditions. However, subjective measures (in Dennerlein *et al.*'s case designed to measure musculo-skeletal load) were more seriously effected. The haptic condition incorporating a single active haptic target led to significantly reduced musculo-skeletal load when compared to the control, but this advantage was not maintained in the other two haptic conditions. Indeed, the condition which rendered all targets as haptically active led to slight, and non-significant, increases in musculo-skeletal load when compared to the control condition.

The practicalities underpinning target prediction, however, seem more in doubt than the validity of the idea. Keuning & Houtsma [109, 110] describe several studies

investigating the accuracy of prediction of the final destination of a movement given its initial trajectory. They conclude that although the creation of an algorithm to perform such a task may be possible, the parameters that control it would vary substantially from direction to direction, device to device and user to user.

Munch & Dillmann [140] present a paper that supports this statement. They describe a complete system that provides not only haptic feedback in a GUI, but also a target prediction system that attempts to mediate the application of this feedback. The haptic feedback they use was based on that described by Akamatsu & Sate [3] - they developed a haptic mouse capable of generating vibro-tactile feedback through the use of an oscillating pin placed under a mouse button, and force feedback in the form of friction generated by the activation of a pair of solenoids over an iron mouse mat. Their target prediction system relied on both trajectory analysis and a dynamically generated model of application behaviour to determine user destination. An intrinsic element of this behaviour model was a substantial learning period for each specific combination of user and application. No formal evaluation of this system was presented, either in terms of the effectiveness of the target prediction system, or whether the haptic feedback generated was beneficial.

3.5. Summary

The pace of technological advancement in this field has been rapid, both in terms of the hardware produced and the software developed. For example, more recent projects to 'haptify' the desktop are not constrained to use the basic haptic effects described in the early studies by Akamatsu & Sato [3] and Engel *et al.* [65]. However, as technology has advanced there has been no corresponding progress in its evaluation. A number of systems have been created featuring a diverse range of haptic feedback, and there has been something of a sense that evaluation has taken a secondary role to implementation [135, 157]. A significant example of this is the lack of literature empirically and directly comparing within or across different haptic augmentations. For example, it is unclear from the literature at what strength it is appropriate to render an attractive force, or what height a haptic wall. Furthermore, it is even harder to discern whether a haptically walled area is superior to an area featuring an attractive force. This absence of comparative empirical data has led to a situation where it is extremely difficult to meaningfully assess and contrast the

different types of feedback available. Consequently, the haptic feedback used in many systems appears to have been designed on an *ad hoc* basis.

This problem is exacerbated by the fact that while the studies that have taken place have highlighted interesting aspects of the influence that haptic feedback can exert on the performance of GUI tasks, there is little real consistency in the results reported. There are no formal guidelines regarding what haptic feedback is appropriate in different situations, and this lack of consensus is reflected in the number of studies investigating essentially similar topics: numerous studies have focused on the augmentation of simple targeting tasks involving a single target. Indeed, the only firm conclusion that can be drawn from this research is that the addition of haptic cues to GUIs will have to be a complex process if it is to be an effective one.

Beyond this issue, the problems that haptic feedback may exert in multi-target situations must also be considered. The limited literature relating to target prediction suggests that although it may be an objectively effective solution to the potential problems of multi-target haptic interaction, it is also a costly and underdeveloped one. Differences between individuals, devices, and even applications may be enough to render such systems useless without substantial training times. More worryingly, the evidence suggests that partially successful systems may exert a damaging influence on subjective satisfaction. The presence of these significant disadvantages indicates that alternative solutions to these problems should be investigated.

Overall, the lack of concrete results relating to this topic presents significant problems for the application of haptic feedback in general situations – there is no simple methodology for adding haptic feedback to GUIs, and application developers may be forced to design their own solutions from the ground up. This problem is exacerbated by evidence indicating that arbitrary combinations of information presented in different modalities is ineffective, and can in fact lead to reductions in user performance [38]. These two facts combine to highlight the importance of a systematic empirical evaluation of the complex haptic augmentations of the desktop that current technologies are able to create. Without such research much time and effort may be wasted, and it is possible that we might even end up with haptically enhanced interfaces that are in fact harder to use than standard ones and haptics may

become relegated to the status of a gimmick. Consequently, it is important to ensure that as the complexity of the haptics that can be produced increases, the focus of HCI remains firmly on the user. There is a pressing need to evaluate, understand and disseminate in a palatable form how and where the complex haptic stimuli that can now be created should be applied.

3.6. Conclusions

This chapter reviewed the literature relating to the addition of haptic feedback to GUIs. It briefly detailed the motivations for such an addition, before discussing the previous literature. A crucial segregation in this review is between research that considers the potentially undesirable side effects of haptic feedback in multi-target scenarios, and that which does not consider this issue.

Within this thesis, the purpose of this chapter is to provide a context in which the research presented in the next chapter can be considered, to introduce the thinking behind it. This occurs throughout the review, but more specifically, it is achieved in the summary of the literature, where three critical weaknesses in the existing body of research are highlighted. The first of these is a general lack of comparative evaluation in systems featuring haptically augmented GUIs, which has led to a situation where it is difficult to determine what type of feedback it is most appropriate to use. The second of these relates to the problems apparent with the current approach to mediating the application of haptic feedback in situations incorporating multiple targets: it is inadequate, and new mechanisms need to be considered. Due in part to these two problems, the final significant weakness is that there is no methodology through which application developers or designers can integrate haptic cues into their systems; there are no guidelines that they can follow to simply plug in this functionality. These three facts motivate the research described in the next chapter.

4. Design, Implementation and Evaluation of Desktop Haptic Feedback

4.1. Introduction

In this chapter two studies that attempt to determine how haptic feedback can be added to Graphical User Interfaces are described and discussed. The first of these experiments looks at a simple target acquisition task. Its goal was to qualify what kind of haptic feedback is appropriate for general targeting tasks. This done, the second study examines a more complex situation involving multiple simultaneously active haptic targets. It attempts to address the question: does extraneous haptic feedback interfere with users performing targeting tasks, and, if so, is it possible to design feedback that does not possess this damaging influence? Some follow up work conducted to extend these ideas is then briefly described and discussed. This chapter closes with a set of design guidelines that attempt to support the creation of haptic widgets in general situations. They are drawn both from the research described here and also from the wider body of available literature, and are the first of their kind on this topic.

This work fits into the overall structure of this thesis by illustrating the performance improvements that can be gained using haptically augmented cursor interactions in a simple, well-established everyday scenario. Basic GUI interactions are the first example chosen to illustrate the general claim of this thesis: that haptic feedback relating to a user's avatar, or cursor, in a digital environment is beneficial and can

lead to increases in performance, or improvements in subjective experience. The guidelines presented at the end of the chapter fulfil the secondary goal of this thesis: to ensure that the design knowledge gained from the research described is presented in a form that easily is accessible to future designers and system developers.

4.2. Comparison of different forms of haptic feedback

4.2.1. Motivations

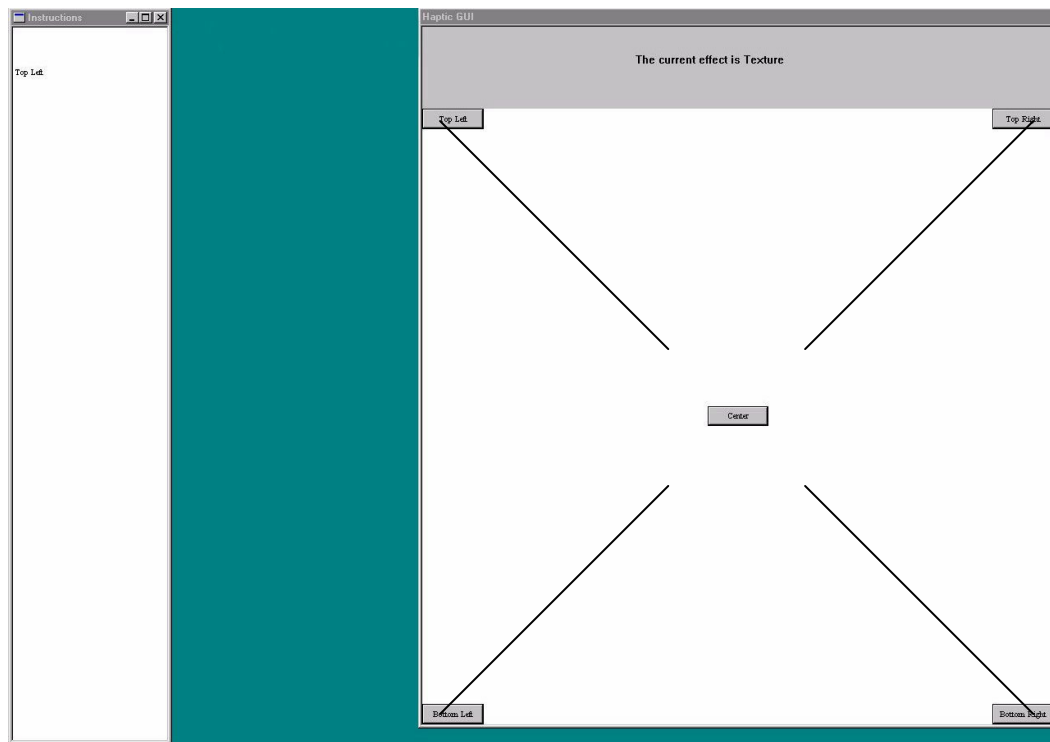
This section describes an initial experiment to evaluate the effects of a variety of haptic modifications in a simple targeting task using a sophisticated force feedback device: a PHANToM 1.0. The goal of this study was to distinguish between the performance levels obtained under different haptic augmentations. This choice of focus addresses a significant weakness observed in the previous literature: an inability to compare among the different forms of feedback employed by different researchers. The results of this study (which is the first to examine this particular topic) should provide detailed information as to the appropriateness (for haptically enabled targeting) of the various haptic effects that current devices are capable of producing. These data should inform the process involved in the creation of haptic widgets by clearly illustrating the positive and negative aspects of different haptic augmentations.

4.2.2. Experimental Task

The task chosen in this study was a simple target acquisition task; participants were required to seek to, and then select on-screen buttons. This task was chosen because it features prominently in the previous literature [3, 65] and also because it is a very elementary operation – it is both simple to perform and also perhaps the most fundamental and common cursor interaction.

Two factors were engineered into the task to make it more suitable for haptic augmentation. Firstly, it was felt that participants in the experiment should experience some visual distraction. This is not an unlikely circumstance in the typical operation of a GUI, particularly in the case of expert users. Expert users often concentrate on some central task and interact with graphical widgets in the periphery of their attention [27].

Figure 4.1. Screen shot of buttons experiment (annotated with lines indicating the path of the corner buttons).



Secondly, in this atmosphere of visual distraction, it seems likely that the haptic feedback will only really prove useful if the task encompasses some repetitive motion. Without some sort of repetitive motion the haptic task would rapidly dissolve into exhaustively searching the entire workspace for some haptically distinct area. This is clearly an inefficient strategy when compared to visually scanning the screen. Furthermore, repetitive motions are relatively common in desktop interactions [144] partly due to the fact that key widgets such as scrollbars and menus have relatively standardised screen positions (at least with respect to the content of the application they control).

To encompass these two factors two windows were placed on the screen at all times. One occupied the left-hand side of the screen and contained instructions as to the next target to seek. The other, larger, window was square and occupied the centre and right-hand side of the screen. It contained the experimental targets in the form of five buttons. One button was always positioned in the centre of the window. The other four were positioned one in each quadrant of the window, on the diagonals of the window. The position along the diagonals changed in the course of the experiment, but each button remained in a single quadrant of the screen throughout. This meant that each button remained in the same direction relative to the centre of

the window at all times. This was felt to provide an experimental task incorporating repetitive motions. Figure 4.1 is a screenshot illustrating the experimental environment, annotated with four diagonal lines indicating the paths that the corner buttons traverse.

Each of the buttons was labelled in accordance with its position on-screen. For instance “Top right” or “Bottom left”. The instructions in the left-hand window consisted of a list of these button names. In each trial the subject had to move over and select the button named at the bottom of the list. Target selection took place on the consecutive depression and release of the controller’s button over the target. Completion of this task caused a new button name to be appended to the list. When the window became filled, it was cleared again and the list began anew. This update mechanism was chosen to be more visually demanding than merely presenting the button names sequentially in an identical position.

To increase the repetitiveness of the motions required the centre button was named in every alternate trial. This led to the participants always moving along one of a very few paths, but continually experiencing variability in the distance they have to travel along these paths. This situation was one where it was felt that haptic feedback would afford performance increases while not being too far removed from an actual, realistic task.

4.2.3. Haptic Effects

Four different haptic effects were created: texture, friction, recess and gravity well. These effects are described in the subsequent sections. For each haptic effect three levels of magnitude were defined. This was thought necessary because while there is little literature comparing different haptic augmentations in GUIs there is none whatsoever comparing user performance on different versions of the same augmentation. For example, there is no empirical basis for choosing how deep to make a recess or dimple, or for determining the appropriate magnitude at which to present an attractive force. It was felt that in the absence of relevant information, it was an appropriate solution to present several different versions of each haptic augmentations varied along a salient scale. It should be noted that the primary goal of this study was to quantify the differences between the various haptic augmentations, and not to meaningfully distinguish between the levels of magnitude

within each of these augmentations. Instead, the motivation to include these levels was to take some initial measure of how user performance varied with different versions of the same haptic augmentation; to discover how sensitive users were to these variations as compared to their sensitivity between different haptic effects.


The magnitude of each augmentation used in the study was determined subjectively. They were validated in a pilot study conducted prior to the main experiment in which three participants were asked to rank the magnitude of the effects they experienced. These data, while present in insufficient quantities for analysis, was supportive of the validity of the chosen magnitude levels.

Texture

Texturing a button in a texture-less, flat workspace is a potential way of haptically signifying that the cursor is positioned over some interesting object. The texture used in this experiment was a set of ridges forming concentric circles centred on the middle of the target. This is illustrated in Figure 4.2.

This particular texture design was used because, when compared with other simple mathematical textures such as gratings, it was felt that this circular texture would go farthest to guaranteeing that users would encounter a number of ridges irrespective of both the direction from which they entered the target and the orientation of their path through it. It was also felt that a texture of circles might provide useful contextual information to a user. A motion perpendicular to any ridge on the surface of the button is guaranteed to move the cursor either directly towards, or away, from its centre. Furthermore, the curvature of the ridges provides an indication as to which direction is which.


Figure 4.2. 3D visualisation of the texture effect.



The textures were created through rotation of the normal vector of the surface of the workspace [178]. This is a commonly used mechanism for the generation of haptic textures [174, 195]. One consequence of this algorithm is that the strength of the textures (in terms of the lateral forces they generate) varies according to the force that a user exerts on the plane of the desktop. As a user exerts more force against the desktop plane, the strength of the normal force generated increases, which leads to a corresponding rise in the magnitude of the rotated force representing the texture. This is realistic in that if one were to push against a textured surface one would expect to experience more disturbance from that texture the harder it was that one pushed. However, it is not ideal, as it has the potential to form a confounding variable across participants. Each participant will rest against the plane with varying force, according to his or her preference, and will therefore feel the texture at different strengths. Consequently, although it is a less physically realistic model, it may have been advisable to implement the texture in some other way, such as with purely lateral forces. This mechanism of texture generation is illustrated in Figure 4.2.

The three levels of texture chosen varied the height of the ridges upward and the frequency of their spacing downward. A low level of texture consisted of a large number of ridges of small magnitude, while a high level of texture consisted of a smaller number of ridges that were of a larger magnitude. The distance between ridges on low magnitude was 1 mm, medium 1.5 mm and large 2 mm. The maximum angular rotation of the ridges ranged from 6 degrees to 12 degrees to 18 degrees. Although texture is extremely complex, and recent studies have suggested that the concept of roughness is poorly defined [130], for the purposes of this study these manipulations served to produce textures ranging from fine to coarse.

Figure 4.3. Cross-section of the texture effect with force orientation



Friction

The friction effect damped a user's velocity. This kind of haptic augmentation has appeared a number of times in the previous literature [3, 95], and, at least hypothetically, can influence user behaviour in several ways. Firstly (and similarly to the texture effect described above), a user should be able to feel the difference between the workspace as a whole and the augmented surface of a target. Secondly, users may experience an increase in performance simply from the fact that their velocity is damped when they are over key targets. For instance, users could develop a strategy whereby targets were approached more rapidly than normal, and the damping influence of the widget was used to snare them over the target and to prevent them from slipping off. A final benefit of augmenting a target with friction is that users may be less prone to make minor movement errors while resting over targets. The friction effect may damp unintentional motions into insignificance, meaning that users will only leave a target when they make explicit effort to do so.

The friction effect consisted of two components - a static and a dynamic friction. This is physically realistic [178]. Static friction refers to the force that must be overcome in order to set a stationary object into motion, while dynamic friction describes the resistance experienced by moving objects. The static friction was created using an attractive force, holding the cursor at its current location. It became active when users stopped moving, and inactive when a preset maximum force was overcome, or a preset distance travelled. It was implemented as a small attractive force, or gravity well. The dynamic friction was applied at all other times and was created using a force opposite to the user's last vector of motion – essentially a damping force.

Varying the magnitude of the friction effect was simply a matter of changing the strength of the force of the gravity well forming the static friction and by varying the magnitude of the vector applied during the simulation of dynamic friction. Both of these variables were increased as magnitude increased. Algorithmically the strongest friction was twice the strength of the weakest with the medium value equidistant between these two.

Recess

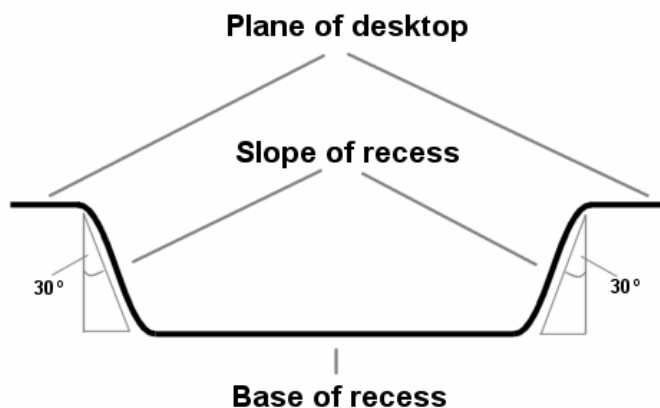
The recess effect was a three dimensional dimple in the plane of the workspace. Two illustrations of its geometry are shown in Figures 4.4 and 4.5. The sloped planes of the recess were slightly curved so as to minimise physical instabilities in the device during operation. The recess was similar to that implemented by other researchers [135, 159]. Like the friction effect, it also has the potential to serve three purposes: moving over (and falling into) or climbing out of the recess not only served to indicate that a target object was under the cursor, but also had the potential to support a user's targeting process. Participants, having moved into a recess, had to exert a degree of explicit effort to order to climb its walls and move away from it. This fact seemed likely to ensure that only purposeful attempts to leave the target would be successful.

To vary the magnitude of the recess augmentation, the angle of the slope of its walls, the final depth of its base, or both could be modified. In this instance, a decision was made to change the depth of the recess and to maintain the angle of the walls as a constant. The motivation for this was a feeling that the recess was essentially a hole,

Figure 4.4. 3D visualisation of the recess effect.



Figure 4.5. Cross-section of the recess effect.



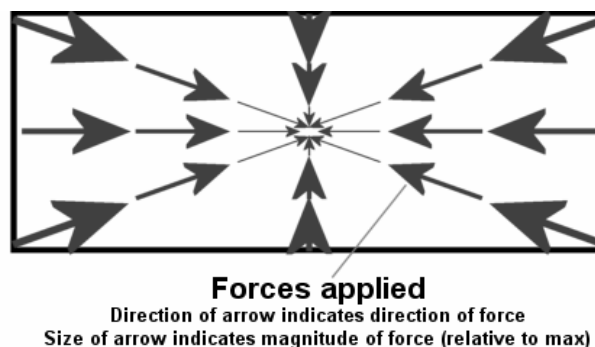
and that the depth of a hole is a more salient characteristic (with respect to magnitude) than the angle at which its sides slope. The depth of the recess ranged from low at 2 mm deep, through medium at 3 mm deep, to high at 4 mm deep. The angle of the slope remained constant at 30 degrees. Subjectively, a slope of this angle felt fairly steep.

Gravity Well

The gravity well was a ‘snap-to’ effect, sometimes called an attractive basin, and similar effects appear commonly in the previous literature on haptically augmented targeting [90]. When users moved over a target a constant force was applied pushing them towards the target’s centre. This force was tapered off around the very centre of the target, creating an area of softer force there. This was done for two reasons. Firstly, it minimised physical instabilities in the hardware, and secondly, it simply felt more natural to have some flexibility in the centre of the well. Figure 4.6 shows a plan view of a target augmented with a gravity well including arrows indicating force direction and magnitude.

As with the recess and friction effects, the gravity well has the potential to signify the presence of a target, in this case simply by the presentation of an attractive force. It may also provide a strong error reduction facility as it is essentially a homing force that attempts to ensure that the user remains on the central portion of a target. Hypothetically, it has the ability to capture users that move over it, and also to prevent small unintentional movements from causing the cursor to stray off it. A user must overpower the attractive force of the gravity well in order to leave the target.

Figure 4.6. Plan view of the gravity well effect with force direction and magnitude illustrated.



Varying the magnitude of the gravity was a simple matter of changing the maximum force it exerted. The three peak values chosen were 0.2 N, 0.4 N and 0.6 N. This was felt to encompass a range of forces from weak to strong.

4.2.4. System

The experiment was conducted under Windows NT SP3 on an Intel PII running at 300MHz. All software was written in C++ using Microsoft Visual Studio version 5. The graphical user interface presented to the user was generated using the standard Microsoft Foundation Classes and as such had the same look and feel as a typical application running on a Microsoft platform. The haptics code was created using the GHOST API supplied by SensAble Technologies. The force feedback was provided by a PHANToM 1.0 equipped with a pen stylus featuring a button. The operation of this button was linked to mouse click events. The workspace available to participants was restricted to a narrow vertical plane, 110 mm wide by 110 mm high and 2 mm deep. Motion along the x and y axes controlled cursor position. No explicit action was mapped to motion on the z-axis.

However, motion on the z-axis did play a significant role in the study as the feedback generated in both the texture and recess augmentations relied on the three dimensional position of the user. In the case of the recess effect this was due to the fact it formed a true three dimensional shape, while the lateral forces generated in the texture effect were reliant on the force a user exerted along the z axis. Both these effects could only be felt by a user resting on the back wall of the workspace. To ensure that participants were exposed to the haptic effects on the same basis throughout the course of the study the gravity well and friction effects were also made to function only when a user was against the back wall. Furthermore, to ensure that users were experiencing the haptic effects throughout the course of the study, cursor interactions were only enabled when users were against the back wall.

Graphically, the five target buttons were 75 pixels long by 25 pixels high on a 17 inch screen with a resolution of 1280 by 1024. This is a typical size for buttons or menu items on the screens of this size and resolution. Correspondingly, the haptic representations of the targets were approximately 10.5 millimetres long by 3.5 millimetres wide. Graphical cursor motions were confined to the window containing the targets which was 780 pixels square. The font used for all text in the study (the

instructions and labels on the buttons) was 8 point Times New Roman. A small size of font was used to increase the amount of visual attention that would be required for participants to read it.

4.2.5. Participants

There were sixteen participants. Four were female and twelve male. All were between the ages of eighteen and thirty. Most were computing students. All were regular and fluent computer users. Three users were left-handed and one was dyslexic. No subject had anything more than trivial previous exposure to haptic interfaces.

4.2.6. Experimental Design

The experiment followed a fully within subjects repeated measures design. Four different haptic effects were tested, each of these at three different magnitudes. This led to twelve haptic conditions. To balance this the Visual condition (which featured no additional haptic feedback) was also split into three separate sub-conditions, resulting a total of fifteen conditions. Each condition consisted of 40 button selections.

A repeated measures design was chosen for this study despite the nearly prohibitive number of conditions because of the novelty of haptic interfaces. Participants had not experienced haptic interfaces previously, and it seems likely that they would possess radically different levels of ability using such interfaces. A repeated measures design is the best way to factor out these individual differences.

Participants experienced the same stimuli in the same order for each condition. This fact heightened the already substantial potential influence of order effects. To remedy these effects the order that the conditions were presented to participants was carefully controlled. In this experiment a random or exhaustive allocation of condition orders to participants is clearly not a solution; the number of possible orders vastly exceeds experimental plausibility. A reasoned approach was used instead. The subject pool was split into eight groups. Table 4.1 shows the order in which each group experienced the conditions. Possible biases with this design are that the Visual condition is always in the central position and that different effect magnitudes are always grouped together.

Table 4.1. The eight different orders of condition presentation in buttons experiment.

Group Number										
1	2		3	4		5	6		7	8
Friction			Texture			Gravity			Recess	
Low	High		Low	High		Low	High		Low	High
Med	Med		Med	Med		Med	Med		Med	Med
High	Low		High	Low		High	Low		High	Low
Texture			Friction			Recess			Gravity	
High	Low		High	Low		High	High		Low	High
Med	Med		Med	Med		Med	Med		Med	Med
Low	High		Low	High		Low	Low		High	Low
Visual			Visual			Visual			Visual	
-----	-----		-----	-----		-----	-----		-----	-----
Gravity			Recess			Friction			Texture	
Low	High		Low	High		Low	Low		High	Low
Med	Med		Med	Med		Med	Med		Med	Med
High	Low		High	Low		High	High		Low	High
Recess			Gravity			Texture			Friction	
High	Low		High	Low		High	High		Low	High
Med	Med		Med	Med		Med	Med		Med	Med
Low	High		Low	High		Low	Low		High	Low

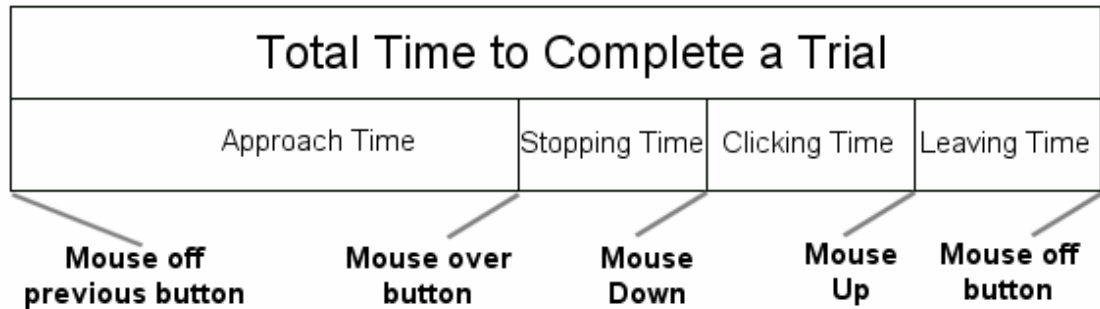
As participants had no previous exposure to haptic interfaces extensive training was thought to be essential to reduce the influence of practice effects. Users performed a training session that was a shortened version of the actual experiment. In the practice participants experienced all the conditions in the actual experiment but each condition contained half the number of trials. The order of the practice trials was not balanced – it remained the same for all subjects. This was a possible bias in the experiment and time constraints were the crucial factor in making this decision.

4.2.7. Measures

Timing Measures

Akamatsu & Sate [4] conducted a target acquisition study similar to the one described here. They were evaluating the effects of the tactile and force feedback generated by a simple haptic mouse and measured three distinct times for each target acquisition trial in the study: approach time, stopping time and clicking time. Approach time was the time taken for a user to reach a target from a pre-set starting position. In their study, an auditory stimulus signified the start of the clock (and the user motion). Stopping time was defined as the time it took users from when they moved over a target until they depressed the controller button, beginning to select the target. Clicking time was the span of time between a user depressing and releasing the controller button, the time required to actually select the target. This

Figure 4.7. Timing measures used in buttons study.



measurement scheme provided a great deal of detail as to the effects of haptic feedback on the precise details of user performance in target acquisition tasks, and consequently it is adopted here. However, for use in this study Akamatsu & Sate's [3] scheme is extended with an additional temporal measure: leaving time. Leaving time is defined to be the time between the successful selection of the target (the release of the controller button) and the user moving off the target. It is hoped that this measure will shed further light on user performance under haptic feedback. It is possible that it did not appear in Akamatsu & Sate's original scheme simply because each trial in their study was completed by the release of the controller button. They were concerned with user performance when selecting a target, but not with any subsequent effects that haptic feedback might exert. However, this new measurement scheme more accurately reflects real user experience; after a user has selected a target, the haptic feedback is still active and could affect performance in future tasks. This measurement scheme used in this study is illustrated in Figure 4.7.

The fact the experimental trials followed each other continuously, without an explicit break, exerts an influence on the measurement of the approach time. Without the occurrence of a definite event to indicate the start of each trial it is unclear at what point the approach time should begin to be measured. To resolve this issue, approach time was measured from the moment at which a participant last moved off a target.

Error measures

Four types of error were recorded, representing a comprehensive taxonomy. The error classifications were entitled wrong-target, slide-over, slip-off, and off-target. Wrong-target errors arose when participants simply selected the wrong target; one that was not currently named in the instruction window. A slide-over error was recorded when a subject moved over the correct target and then moved off it again without depressing the controller's button. This action represents non-optimal

behaviour (a chance has arisen to select the target, but this action did not occur), but is arguably a natural part of the targeting process. A slip-off error, defined by Brewster *et al.* [27], occurred when a user depressed the PHANToM's button over the appropriate target but then either purposely or accidentally moved off the target before releasing the button. This does not lead to a successful target selection, and is often confusing for users as it exhibits the same visual feedback as a correctly performed operation. The final error classification, off-target errors, was defined as when participants performed actions with the controller's button on the window area outside of the targets. This was a 'catch-all' category for remaining interactions that could occur through operation of the PHANToM's button. An example of such a situation is when a user's selection process misses a target entirely.

Subjective measures

There is evidence suggesting that quantitative measures of performance do not provide a full picture of the usability of an interface [202]. For instance, it is conceivable that users may be able to perform rapidly and accurately using an interface, but find it frustrating, tiring or mentally demanding. Questionnaires represent a mechanism for gaining access to this subjective data. Specifically, the concept of *workload*, defined by Hart & Wickens [89] as "...*the effort invested by the human operator into task performance...*", provides a framework for measuring this information. Consequently, a workload test was used in this evaluation to provide a more thorough and sensitive understanding of the influence of the haptic feedback.

The NASA Human Performance Research Group [81] analysed workload into six categories: mental demand, physical demand, time pressure, effort expended, performance level achieved and frustration experienced. These are defined as follows:

- Mental demand: the amount of mental and perceptual activity required by the task.
- Physical demand: the amount of physical activity required.
- Time pressure: the time pressure felt.

- Performance level achieved: the individual's perception of the degree of success.
- Effort expended: the degree of effort an individual invested.
- Frustration experienced: the amount of insecurity, discouragement, irritation and stress felt.

One way of assessing workload, known as the raw TLX method [35], was used in this study. It essentially involves gathering a measure of overall workload by having participants rate each of the six factors mentioned above, then simply calculating their average. For use in this study, the basic format was modified with the inclusion of an additional scale under the banner 'fatigue experienced'. This was an attempt to measure the physical discomfort participants experienced using the PHANToM for a protracted period. It seems likely that this is an important consideration in any situation where haptic feedback is applied to a user. Appendix A contains the materials used to gather TLX subjective workload throughout this thesis.

A questionnaire that gathered basic demographics, asked participants to order their preferences for each of the 5 feedback conditions, and then to rate their perceived performance under each of these conditions was also administered. A copy of this questionnaire is presented in Appendix B.

4.2.8. Procedure

All participants first performed a practice session. At this time written instructions were presented explaining the experimental task, and participants were free to ask any questions to clarify these. During the course of the practice participants were encouraged to explore the different haptic effects, in order to become accustomed to them. After every magnitude condition a window appeared requesting that they rest until they felt ready to continue.

After completion of the practice session participants immediately moved onto the experiment itself. This was essentially the same as the practice session except that participants were informed in the instructions that there was a prize of £30 for the fastest and most accurate participant. They were also required to fill out a NASA TLX subjective measures questionnaire [88] after they had completed each group of

three magnitude conditions that comprised an effect condition. In certain situations and depending on the forces being exerting the PHANToM can sometimes produce unwanted noise. To prevent any such noise from influencing participants' performance they were required to wear Sennheiser HD25 headphones throughout the experimental session. At the very end of the experiment participants were required to fill out the questionnaire gathering their demographics and preferences for each of the haptic conditions.

4.2.9. Hypotheses

This experiment was exploratory in nature – it sought to investigate the differences between the four haptic effects against the context provided by the Visual condition, which featured no additional haptic feedback. Therefore, the experimental hypotheses were suitably unspecified and two-tailed: they simply suggested that differences would be observed in task completion time, number of errors and in the subjective data gathered.

4.2.10. Results

The raw data from this study are presented in Appendix C. Two basic comparisons were made during the analysis of all the temporal and error data. Firstly, each of the three different magnitude conditions within each haptic condition were compared against one another. Secondly, the data from each of the five sets of three magnitude conditions were averaged together and these data were used to compare the differences between the actual haptic effects. All group comparisons performed on the experimental data (including the subjective measures) were repeated measures multi-factorial ANOVA [48]. All subsequent pairwise comparisons were conducted with *post-hoc* t-tests [48], adjusted using Bonferroni confidence interval adjustments [93].

Temporal Analyses

Only one significant effect was detected when comparing the temporal results from the different magnitude conditions within each haptic augmentation. Within the three magnitude levels of the gravity well, there was a slight (60 ms) but highly significant ($F(3, 15)=10.834, p<0.001$) increase in the leaving time, the time measured from the moment the PHANToM's button is raised to the moment it moves off the target,

with increasing gravity well strength. This is illustrated in Figure 4.8. *Post-hoc* pairwise comparisons revealed that the low magnitude gravity well led to significantly shorter leaving times than the mid and high strength wells (respectively $p < 0.05$ and $p < 0.01$). The leaving time for the mid strength well also approached a significant improvement when compared to that achieved under the high strength well ($p = 0.062$).

Figure 4.8. Mean leaving time in different Gravity Well magnitude conditions in buttons study.




Figure 4.9. Mean total trial per trial in buttons study.




Figure 4.10. Mean approach time per trial in buttons study.

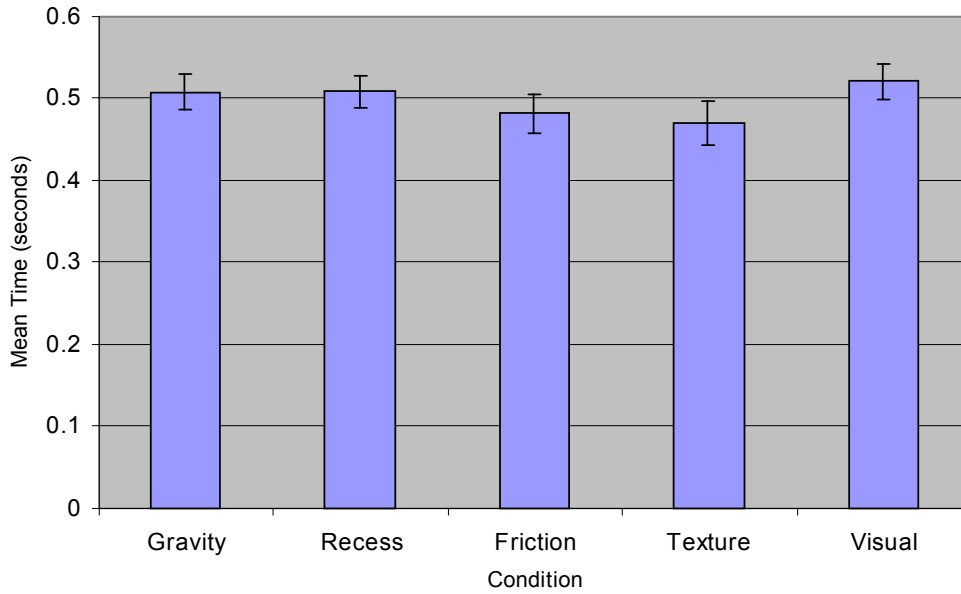



Figure 4.11. Mean stopping time per trial in buttons study.



Significant effects when comparing the averaged scores for each haptic effect were more commonplace. No significant effect was found for the overall time for each trial ($F(4, 15)=1.765, p=0.128$) but they were observed for each of the individual measures; the approach time ($F(4, 15)=5.108, p<0.01$), the stopping time ($F(4, 15)=5.750, p<0.01$), the clicking time ($F(4,15)=2.877, p<0.05$) and the leaving time ($F(4, 15)=9.668, p<0.01$). Figures 4.9, 4.10, 4.11, 4.12 and 4.13 illustrate total trial time and each of the individual timing measures. No order effects were found in an analysis of total task completion time ($F(4, 15) = 0.913, p=0.462$).

Figure 4.12. Mean clicking time per trial in buttons study.





Figure 4.13. Mean leaving time per trial in buttons study.



Despite these highly significant results *post-hoc* pair-wise t-tests on the means for these conditions yielded relatively few significant results. The gravity condition was found to yield a significantly slower leaving time than the Visual ($p < 0.01$), the recess ($p < 0.01$) and the texture conditions ($p < 0.01$), and a faster clicking time than the texture condition ($p < 0.05$). It is also worth noting that the difference between the best and worse performing effects (on the total trial time) was only 42 ms, a very short time.

Error Analyses

Once again few significant differences were detected among any of the magnitude conditions. As the strength of the gravity well increased participants made significantly fewer slide-over errors ($F(3, 15)=5.095, p<0.012$), although *post-hoc* comparisons did not reveal any significant differences between the means. The difference between the low and high magnitude gravity wells did approach

Figure 4.14. Mean total slide-over errors in different Gravity Well magnitude conditions in buttons study.

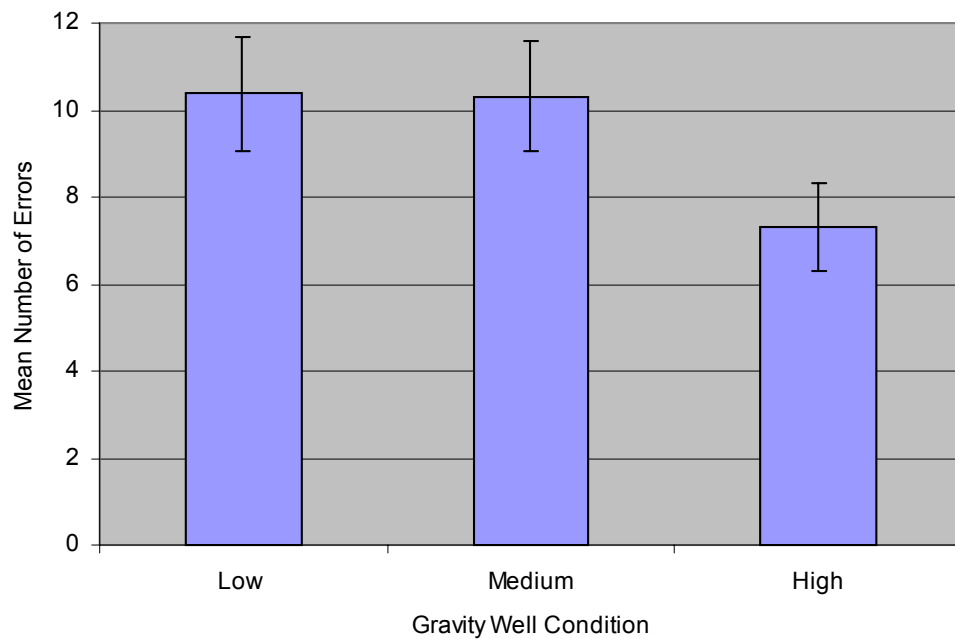


Figure 4.15. Mean total slide-over errors in different Texture magnitude conditions in buttons study.


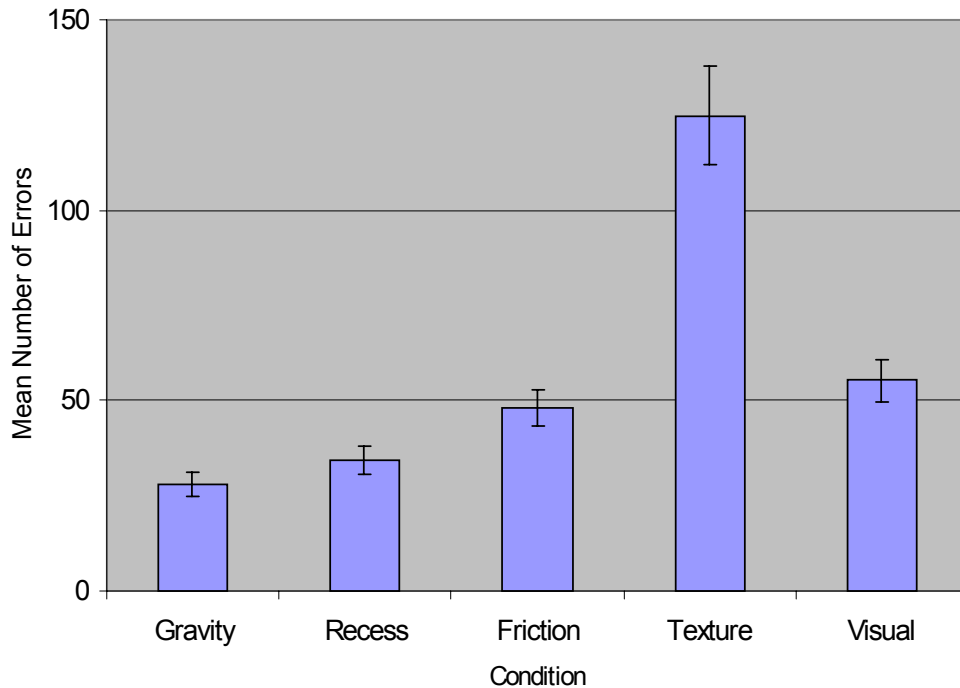


Figure 4.16. Mean total slide-over errors in buttons study.



significance ($p=0.056$). Conversely, as the magnitude of the texture increased participants made significantly more slide-over errors ($F(3, 15)=5.757, p<0.01$), and *post-hoc* comparisons bore this out. The low magnitude texture led to significantly fewer slide-over errors than the high magnitude texture ($p<0.05$). These data are illustrated in Figures 4.14 and 4.15.

Significant effects were found when comparing the averaged scores for each haptic effect for slide-over ($F(4, 15)=48.487, p<0.01$), slip-off ($F(4, 15)=20.81, p<0.01$) and off-target errors ($F(4, 15)=6.429, p<0.01$). No significant effects were found for participants attempting to select the wrong target ($F(4, 15)=0.315, p=0.867$). This was taken as positive evidence towards an absence of practice effects. Furthermore, order effects for the critical error classifications of slide-over ($F(4, 15) = 0.152, p=0.961$) and slip-off ($F(4, 15) = 0.123, p=0.974$) were not found. Figures 4.16, 4.17, 4.18 and 4.19 show these error results. An interesting aspect of these error data is that the graphs for all the slip-off and slide-over results possess very similar curves, suggesting a strong relationship between these error types.

Post-hoc analysis of the means generated from the slip-off data revealed that the results from the Gravity effect were significantly lower than those from the Friction ($p<0.01$), Texture ($p<0.01$) and Visual ($p<0.01$) conditions. The results from Recess

Figure 4.17. Mean total slip-off errors in buttons study.

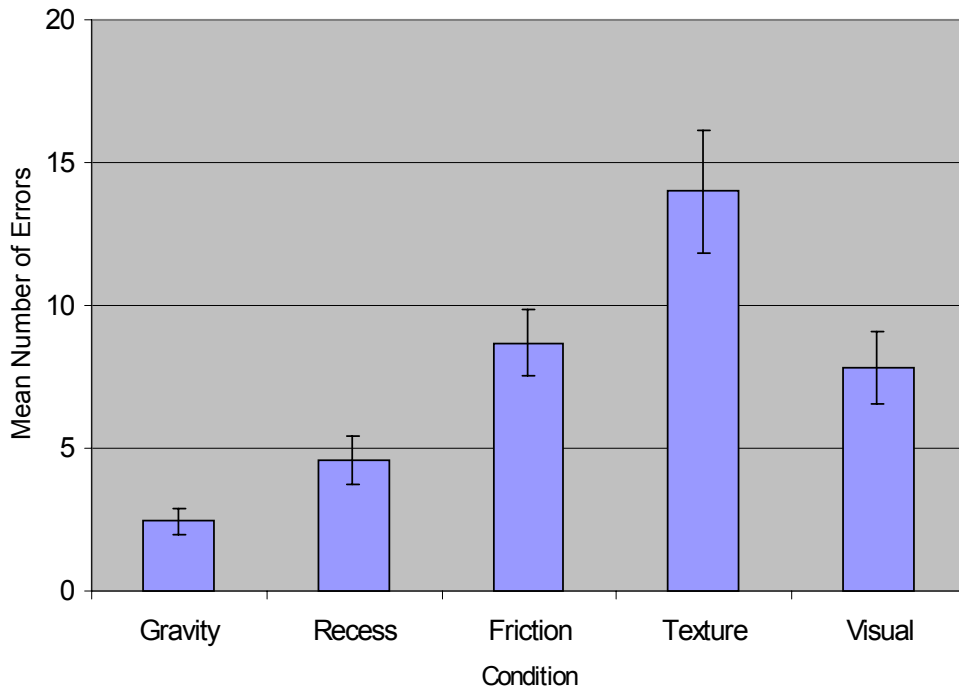
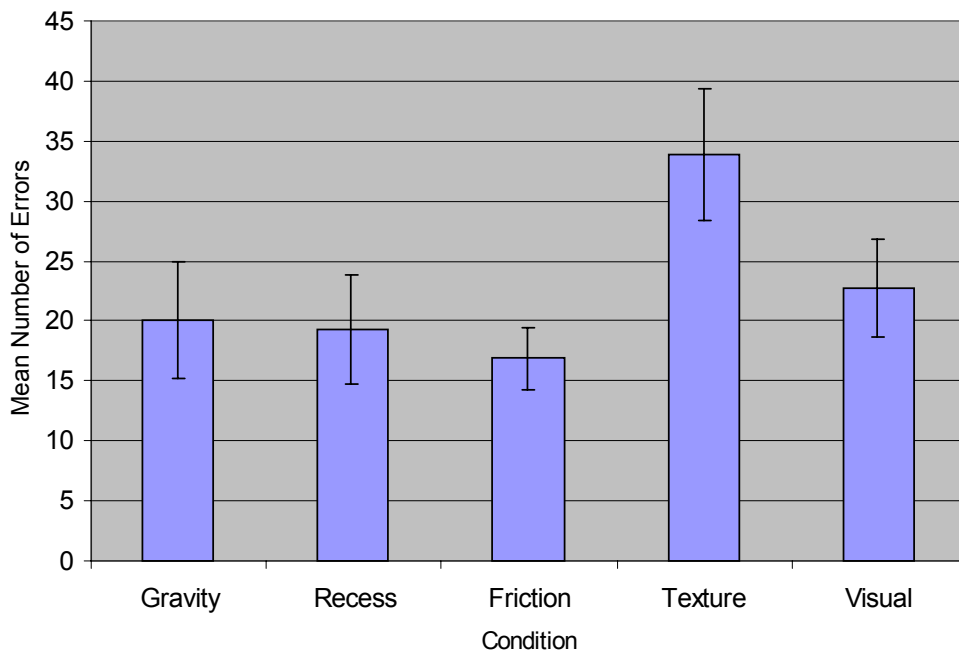



Figure 4.18. Mean total off-target errors in buttons study.



effect were also significantly improved when compared to the Friction ($p < 0.01$) and Texture ($p < 0.01$) effects. Finally, the Friction and Visual conditions yielded better results than the texture condition ($p < 0.01$, $p < 0.05$ respectively). Similar, *post-hoc* analysis of the slide-over results indicated that the Gravity and Recess effects performed better than the Friction, Texture and Visual conditions (all at $p < 0.01$ or

Figure 4.19. Mean total wrong-target errors in buttons study.



better). The Friction and Visual effects were also yielded improved results when compared to the Texture effect ($p < 0.01$). The analysis of the off-target data showed that Texture condition led to significantly more errors than all the other conditions (all at $p < 0.05$ or better) and that the Friction effect performed better than the Visual condition ($p < 0.05$). Tables 4.2, 4.3 and 4.4 present a more concise representation of the significance results observed in the pairwise comparisons of the error data.

Table 4.2. Results from pairwise comparisons of slip-offs in buttons experiment.

	Gravity	Recess	Friction	Texture	Visual
Gravity	-----	Not sig.	$p < 0.01$	$p < 0.01$	$p < 0.01$
Recess	-----	-----	$p < 0.01$	$p < 0.01$	Not sig.
Friction	-----	-----	-----	$p < 0.05$	Not sig.
Texture	-----	-----	-----	-----	$p < 0.05$

Table 4.3. Results from pairwise comparisons of slide-overs in buttons experiment.

	Gravity	Recess	Friction	Texture	Visual
Gravity	-----	Not sig.	$p < 0.01$	$p < 0.01$	$p < 0.01$
Recess	-----	-----	$p < 0.01$	$p < 0.01$	$p < 0.01$
Friction	-----	-----	-----	$p < 0.01$	Not sig.
Texture	-----	-----	-----	-----	$p < 0.01$

Table 4.4. Results from pairwise comparisons of off-targets in buttons experiment.

	Gravity	Recess	Friction	Texture	Visual
Gravity	-----	Not sig.	Not sig.	$p < 0.05$	Not sig.
Recess	-----	-----	Not sig.	$p < 0.05$	Not sig.
Friction	-----	-----	-----	$p < 0.01$	$p < 0.05$
Texture	-----	-----	-----	-----	$p < 0.01$

Subjective Measures

Figure 4.20 shows the TLX workload scores, scored out of 20 and adjusted so that higher scores consistently indicate higher workload. ANOVAs revealed that Overall Workload and all individual factors bar Time Pressure differed significantly. The details of these statistics are presented in Table 4.5. However, *post-hoc* pairwise comparisons between the conditions turned up relatively few concrete results. The Gravity condition was rated as yielding significantly lower Overall Workload than the Friction ($p<0.01$), Texture ($p<0.05$) and Visual ($p<0.05$) conditions, while the Recess and Visual conditions achieved significantly lower workload than the Texture condition (both at $p<0.05$). Turning to the individual scales, no significant differences were observed in the Mental Demand data. Participants rated the Gravity condition as significantly less Physically Demanding than the Texture condition ($p<0.05$), and also the Recess and Visual conditions as requiring less Effort Expended than the Texture condition ($p<0.01$ and $p<0.05$ respectively). Gravity achieved significantly better Performance Level Achieved than the Visual ($p<0.05$).

Figure 4.20. TLX workload results from buttons study.




Table 4.5. Results from ANOVA analysis of TLX results in buttons experiment.

	MD	PD	TP	EE	PLA	FE	Fat E	Overall
F value	5.033	2.803	2.471	5.654	5.191	10.693	4.138	26.449
P value	p<0.01	p<0.05	p=0.054	p<0.01	p<0.01	p<0.01	p<0.01	p<0.01


The Gravity, Recess and Friction conditions all resulted in significantly lower ratings of Frustration Experienced than the Texture condition (all $p<0.01$). Finally, the Recess condition led to significantly lower levels of Fatigue Experienced compared to the Texture condition ($p<0.05$).

The results of the post experiment questionnaire are shown in the Figure 4.21. Due to the unvalidated nature of this questionnaire no statistics were performed on these results. However, it is possible to observe a strong participant preference for the Gravity well and Recess conditions and a strong dislike for the Texture condition. Indifference is reserved for the Friction and Visual conditions. The results from this questionnaire appear to follow the same general trend as that found across the TLX subjective measures data.

4.2.11. Discussion

The timing results reveal few concrete significant differences. One reason for this may be that the approach time does not accurately represent the average time spent moving towards a target from the previous target. As mentioned in the measures section it instead represents the time from the immediately previous instance when a

Figure 4.21. Mean ratings from post-experiment questionnaire in buttons study.



user left a target, even if that target is the same target that they are moving over once again. Essentially, this means that if a user performs a slide-over or a slip-off the approach time is measured from the time of the occurrence of this error until he or she moves over the target once again. This is typically a very short span compared to the time taken to traverse the distance between two targets. Given the substantial differences in the occurrence of movement errors (such as slide-over) among the conditions, it seems likely that approach time is inversely related to the incident rates of these errors. This observation places serious doubts on the validity of the approach time measure used in this study.

Other than differences observed in the approach time, the most important timing differences are that participants took both longer to move off gravity wells of increasing strength, and to move off gravity wells when compared to performance under other conditions. This suggests that a weakness of the gravity well augmentation may be that it impedes users' performance as they attempt to leave a target. This factor has typically been overlooked by other researchers [3, 51] as target selection is usually viewed as completion of the experimental task. In a real scenario, however, users not only select their desired target, but also subsequently move off it in order to perform further interactions. The omission of this aspect of targeting from the previous literature casts doubt over its validity.

An analysis of the timing results based around this hypothesis yields an interesting trade-off. Ignoring the approach time, Figures 4.22 and 4.23 plot the selection time (the sum of the times to stop and click on the target) and the total time on target. They show that when considering only the selection time, the Gravity, and to some extent the Friction conditions, perform substantially better than they do when the total time on target is examined. Indeed, the Gravity condition yields the fastest selection time and one of the slowest total times on target. Formal analysis of these data reveals that both selection time and total time on target vary significantly (respectively $F(4, 15) = 6.860$, $p < 0.01$ and $F(4, 15) = 2.647$, $p < 0.05$). *Post-hoc* analysis indicates that the Gravity condition provides a faster selection than the Texture condition ($p < 0.05$), and a slower total time on target than the Recess condition ($p < 0.05$). The reduction in total time on target achieved by the Recess condition in comparison to the Friction condition approached significance ($p = 0.085$).

Figure 4.22. Mean selection time per trial in buttons study.

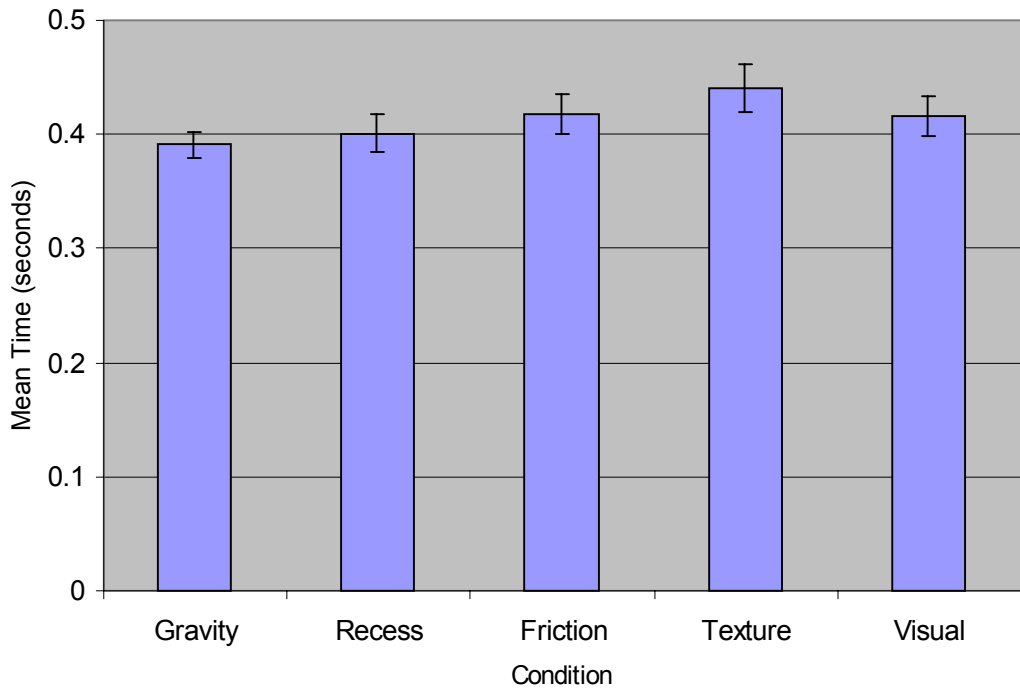



Figure 4.23. Mean total time on target in buttons study.



The difference is not surprising if the feedback that these effects generate is examined. Gravity wells exert a force that actively prevents a user from moving off a target. Therefore, this results in users spending a larger amount of time performing this action than they do in the other conditions. Similarly, the static component of the friction effect, which becomes active when a user is moving very slowly (such as

when they are performing a target selection) acts like a weak gravity well centred on the user's current position. The comparable nature of the feedback probably accounts for the similar effect it exerts on user performance.

This trade-off also suggests that different haptic augmentations may be suitable for different situations. In a scenario where targeting and selection, simply reaching and selecting a widget, are of paramount importance the gravity well may be appropriate. It has the potential to yield optimised times for this portion of the interaction. However, when considering more general purpose scenarios incorporating all stages of interaction with a target, the recess effect, which led to the lowest total time on target, may well be more suitable.

The error results from different magnitude conditions within the same effect revealed few differences. There was a reduction in number of slide over errors as gravity well strength increased. This suggests that the increasingly strong attractive forces make users less likely to overshoot their desired target. In other words, stronger gravity wells result in a significantly more effective snap-to behaviour, reinforcing the suggestion that gravity wells may be the most appropriate haptic augmentation if target selection (and not full interaction) is the primary goal. There was also an increase in slide over errors as texture strength increased, indicating that users found controlling their position more and more difficult as the magnitude of the texture increased.

More substantial significant effects were found comparing the error rates of the different haptic conditions with one another; indeed differences of several orders of magnitude were observed. The recess and gravity conditions resulted in the lowest occurrences of the slip-off and slide-over errors, while friction matched the performance of the Visual condition, and the texture effect resulted in far larger incidences of these errors. These two error types represent situations involving moving over the correct target and then moving off it, without successfully selecting it. It therefore seems reasonable to suggest that gravity well and recess effects strongly support targeting interactions, while the texture effect strongly hinders them. This evidence suggests that the texture effect substantially interferes with a user's ability to control their position accurately.

As a whole, the subjective data were generally in favour of the Gravity and Recess conditions, and firmly against the Texture condition. The Gravity and Recess conditions resulted in significantly less workload than the Texture condition on a number of factors, including Overall Workload. The Friction and Visual conditions each yielded a reduced workload when compared to the Texture condition on a single factor (respectively Frustration Experienced and Effort Expended). The only significant result from the TLX questionnaire not involving the Texture condition was an improvement in ratings of Performance Level Achieved in the Gravity condition over the Visual condition. This suggests that, subjectively, participants feel that they can gain performance benefits from this haptic augmentation, and indicates a willingness to accept the aid it offers in targeting tasks. The results from Fatigue Experienced, the additional term incorporated to the TLX questionnaire in this study, indicate that, at least in short term, haptic feedback does not lead to higher perceived levels of fatigue when compared to a Visual condition featuring no additional haptic feedback. This is an encouraging result.

Taking these results as a whole several conclusions can be drawn. Firstly, and most importantly, they demonstrate that the addition of haptic feedback to target acquisition tasks can either improve or reduce user performance substantially. The Gravity and Recess effects led to order of magnitude reductions in error numbers when compared to the standard Visual condition, while the Texture effect led to considerable increases. Supporting this claim is Akamatsu and Mackenzie's 1996 study [4] in which the presentation of tactile feedback on their haptic mouse increased the number of errors performed by participants, while force feedback reduced them. This conclusion supports the statement made in the previous chapter that applying haptic feedback to desktop user interfaces is far from straightforward, and that design guidelines are urgently needed to facilitate the consistent creation of beneficial haptic feedback.

Secondly, it seems clear that the addition of haptic feedback to widgets exerts its most dramatic and observable effects not by decreasing task completion time, but instead by reducing the occurrence of errors. While these two different metrics can access the same basic aspect of performance – for instance the occurrence of errors in the performance of a task typically results in an increase in the time taken to

complete that task – it was felt that in this evaluation more information about user behaviour could be gathered by measuring errors rather than focusing on task performance time. The measurement of slip-off and slide-over errors essentially performs the same function as measuring the total time it takes a user to approach a target prior to a successful selection. However, in this case it seems likely that the two error measurements provide a better window onto actual user behaviour than a simple single measurement of time.

Furthermore, supporting this suggestion are the results from both of the subjective measures (the TLX and the post experimental questionnaire). Both show a uniform trend towards favouring the gravity well and recess conditions, treating the friction and Visual conditions with ambivalence and reserving dislike for texture condition. The ratios of these results bears a striking resemblance to those of the slip-off and slide-over error results. This suggests that in this sort of simple, low level task where total completion time is very low, and the differences between a slow and a fast time are at the periphery of detection, that users may rely on error rates in their subjective assessments of performance. If this is the case, then for this kind of task, users may well regard a reduction in the number of errors to be more important than a decrease in the completion time.

A number of observations that go some way towards explaining the different levels of performance found for each haptic effect can also be made. For instance, it is possible that the reason the texture effect yielded such poor results across the whole range of measures was due to the fact that textural information, which is essentially tactile information, was being delivered through a force feedback device to the kinaesthetic sense. In the texture implementation used in this study, the perturbations were rendered by altering the orientation of the surface presented to the surrogate the user was holding. This has the potential to effect the position of the user's entire hand. However, in real environments, textural information is sensed by perturbations on the surface of the skin [194], not kinaesthetically, and is extremely unlikely to effect any large scale positional movements. Supporting this suggestion, some studies, such as that conducted by Campbell *et al.* [38], use tactile devices to present simple textural information and have reported that it led to improvements in user performance [4].

The performance of the friction effect matched that of the Visual condition in the majority of the measures. This indicates that users do not gain much advantage from the velocity damping it provides. However, from observing users interacting with the effect a second cause for this absence of improvement can be identified – an interaction between the influence the friction exerts and a user strategy for the effective completion of the task. Essentially, as users entered a target they began to slow their speed, achieving this more rapidly than they might expect due to effects of the dynamic friction. When they reached a given minimum velocity, the dynamic friction became inactive and the static friction became active, acting like a small gravity well centred on their current position. However, here the problem emerges. Users were not typically satisfied with their position on the target after this initial halt. Often they were on the very edge of the target and, although it is not essential, wanted to reposition themselves more centrally before selecting. The gravity well like force profile of the static friction prevented them from performing this kind of small adjusting manoeuvre. Users would then either waste time attempting to reposition themselves unnecessarily, or worse, accidentally fall off the button through this behaviour. While it is possible that this problem could be resolved through tweaking the parameters of the friction, possibly entirely removing one component or the other, it is an illustrative example of where user behaviour does not take full advantage of the supportive haptic forces. Two users adopted a strategy whereby they selected the target regardless of their position on it, and these participants expressed considerable preference for the friction in the post-experiment questionnaire.

Indeed, the adoption of strategies such as this undoubtedly exerts a considerable influence on user performance. Given that this study introduces not only the different haptic augmentations but also a novel cursor control device, it seems unlikely that users will attain their optimal level of performance. Indeed this assertion is supported by Jansson *et al.* [105] who conducted a study investigating the duration of practice effects in haptic tasks. They concluded that practice effects could take several hours to fully diminish.

4.2.12. Conclusions

This study looked at four different mechanisms for augmenting user interface targets with haptic feedback. These augmentations represent those studied, sometimes to confusing or contradictory ends, by other authors in this field [3, 65, 90]. By comparing these different haptic feedbacks against one another directly, the results of this experiment have highlighted interesting aspects of the behaviour elicited by the feedback that might not otherwise have come to light.

The results themselves support the observation that the addition of haptic feedback to GUIs is by no means a simple task. Haptic feedback has been shown to be able to significantly increase or decrease objective measures of performance by several orders of magnitude. Of the four effects studied, the Gravity and Recess effects consistently out-performed the un-augmented Visual condition, as well as the other two haptically enhanced conditions. However, this increase in performance is not without some cost. Examining the data closely indicates that users experience a measurable difficulty when leaving targets overlaid with the Gravity effect. In the simple case described here, where participants interact with an individual target, this effect may not be of critical importance – indeed the subjective measures suggest that users did not find anything amiss with the Gravity augmentation. However, in a more complex situation where users encounter many targets while making their way to their desired destination, this effect may be much more significant, and warrants further investigation.

4.3. Multi-Target Augmentation of a Menu System

4.3.1. Introduction and Motivations

The study described above, in conjunction with some of the previous literature [52, 109, 140] suggests that the haptic feedback that is appropriate in situations involving only single targets may not generalise well to those involving multiple targets. The reason for this is simple: the forces that aid targeting (such as those forming gravity wells) when in place over a user's desired destination seem likely to exert a damaging influence when they are over objects on the path to a user's desired destination. The forces on these extraneous targets are likely to perturb and disrupt a user's intended movements, potentially interfering with the targeting process.

One solution to this problem is target prediction, which involves reducing a situation involving multiple haptic targets to a situation involving only a single haptic target by predicting a user's destination, and applying forces only over that target. However, while this concept seems sound, there are substantial problems implementing a satisfactory system. Critically, the current research on prototype systems suggests that target prediction algorithms will require prohibitively long training times for each combination of user, device and application.

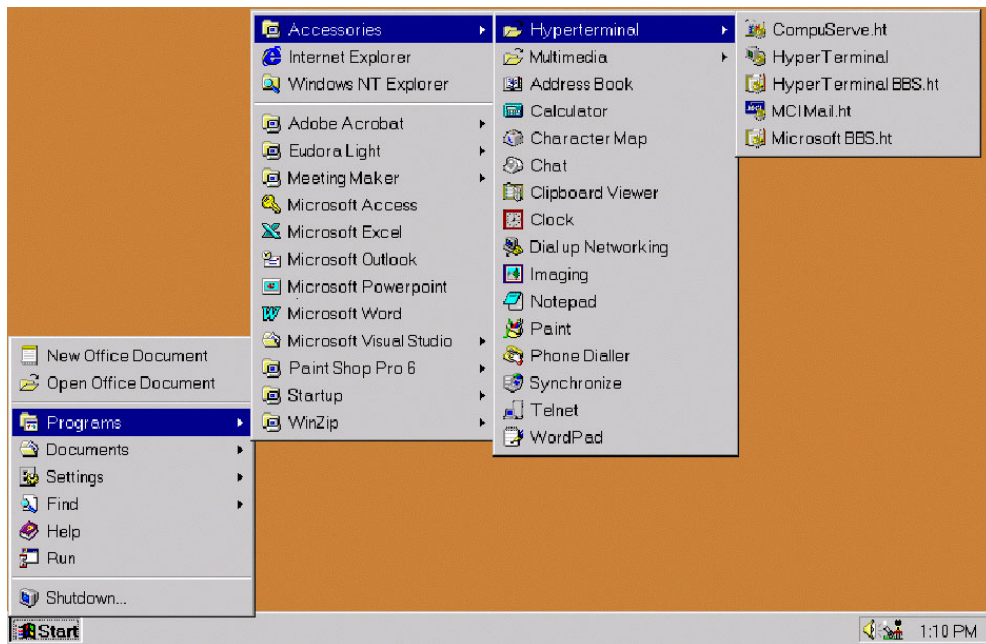
Given this limitation, this study proposes and examines a novel approach to this problem. It describes and evaluates a group of haptically enhanced widgets that has been designed to support a user's interactions by dynamically tailoring the forces presented based on directly measurable aspects of user behaviour. This differs from the target prediction approach, which essentially attempts to alter the applied feedback by determining intended destination, a complex and variable derivative of user behaviour. This solution, therefore, should not suffer from the problems inherent in target prediction systems, and has the potential to serve as a viable alternative.

4.3.2. Experimental Task

One task that may present specific challenges to a target prediction approach to the display of haptic feedback is menu selection. This is chiefly due to the large numbers and closely packed nature of the targets in a menu system. Very accurate target prediction would be required to avoid applying feedback over inappropriate targets. Furthermore, when a menu system includes hierarchical submenus, more than one targeting operation is required to complete a single interaction (as targets to bring up submenus must be selected prior to the final, desired, target), yielding haptic multi-target problems even with accurate target prediction for each level of the menu.

Several researchers have investigated issues relating to haptified menus. Dennerlein *et al.* [51] describe decreases in performance time in a haptically augmented condition in a task analogous to moving along a single menu item. Campbell *et al.* [38] demonstrate reductions in both task time and error rate with the use of tactile feedback in a steering task similar to menu interaction. These studies provide evidence that haptic feedback can support moving along the narrow tunnel of a

Figure 4.24. Screenshot of menu experiment.



single menu item. However, menus are always composed of groups of items, so the multi-target case must be considered.

Participants interacted with a menu system that had similar appearance and behaviour to the Start Menu (from Microsoft's Windows). It was activated by depressing the controller's button over a target (labelled "Start") that resided in the lower left corner of the screen, and consisted of 95 menu items arranged in 17 menus in a hierarchical structure. At its deepest the hierarchy extended to 4 sub-menus. The items in the menu system included the standard contents of the Start Menu in a typical Microsoft Windows install. This basic structure was extended using the menu items generated by commonly installed applications (such as Microsoft Office). Mimicking the behaviour of the actual Start Menu, menu items were selected by the release of the controller button. When this occurred the menu system collapsed back to its initial, inactive state. Clicking on the area outside of the start menu also led to its collapse back to its initial state.

Each trial in the study consisted of selecting an item from the menu. The current target and, as the cognitive aspects of searching the menu hierarchy are not relevant to this study, the path through the hierarchy to that target were displayed in a separate window directly beneath the experimental window. Successfully selecting the named target caused the name of a new target to be displayed. Figure 4.24 is a screenshot of the experimental window.

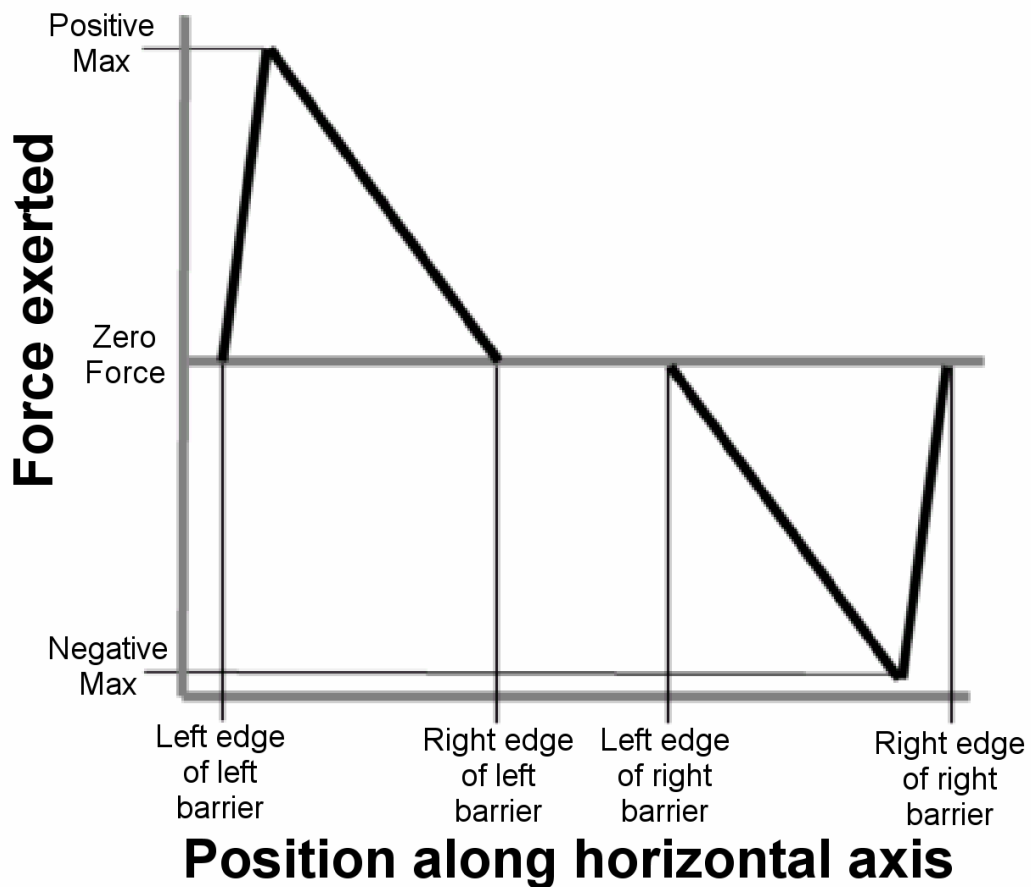
4.3.3. Haptic Feedback

Two types of haptic feedback featured in this study: Haptic barriers and Adjusted barriers. These were strongly related and are described below.

Haptic Barriers

Haptic barriers were simple two-dimensional haptic walls. To enable these barriers to reside adjacently they had a simple force profile ensuring that either side of the barrier returned to zero force. Four of these barriers arranged to enclose a rectangular area served to produce a single haptic target. The force profiles of two barriers are shown in Figure 4.25, orientated to serve as opposite walls of a target. One consequence of this implementation is that the corners of targets were subject to more substantial forces, as barriers in both x and y dimensions independently contributed force. This problem was partially resolved by capping the maximum exerted force to the maximum for a single barrier, but the corners of a target still

Figure 4.25. Force profile of two haptic barriers from menu study.




consisted of larger areas of the maximum force. This made diagonal motion more difficult than either horizontal or vertical motion. The maximum magnitude of the haptic barriers was 0.4 N. This algorithm had similar properties to those used in other studies of haptified targeting and steering tasks [51, 90]; moving over a target caused a user to be pulled into its centre, and leaving a target required overcoming the barrier forces surrounding it. It also closely resembled the recess and gravity effects that led to performance improvements in the button selection study described earlier in this chapter.

Adjusted Barriers

The Adjusted barriers were a modified version of the Haptic barriers. Two specific modifications were in place. The first of these comprised a reduction in the magnitude of all forces (to 0.08 N) when a user moved slowly (below 2 cm/s). This was designed to enable users to move from one menu item to an immediately adjacent one without being hindered by strong forces. The second modification was active when a user moved rapidly (above 2 cm/s), which in a menu system tends to occur either horizontally or vertically. It entailed a reduction in the force applied along each axis individually (to zero) in proportion to the user's speed along the opposite axis. This modification was designed such that, when a user was moving rapidly, the forces that opposed that motion would be reduced, while the ones that

Figure 4.26. Modification of force according to user direction in menu study.



supported it would be maintained - as a user moved horizontally in the menu, only the vertical forces that aided that motion would be presented, and *vice versa*. This would allow users to move across, or along, menu items at speed, gaining the benefits of forces supporting, without the cost of those obstructing, these actions. This concept is illustrated in Figure 4.26. These modifications were created around the two specific observations. The first of these was that there are two distinct phases to any targeting operation: an initial phase where users move rapidly to reach the general area of their desired destination, followed by a period of slower, more precise movements when they actually move over their target. The second observation was that, in order to interact with a menu, users have little choice as to what direction they move in – they must move vertically to reach their desired target in the current level of the menu, and horizontally to move onto a submenu. These two facts were exploited to produce a potentially more effective haptic augmentation of a menu system.

All transitions between different force magnitudes in the Adjusted barriers were gradual, so as not to disrupt users, but took place extremely rapidly. This was made possible due to the PHANToM's native 1000 Hz update rate.

4.3.4. System

The experiment was conducted under Windows NT SP3 on an Intel PIII running at 700MHz. All software was written in C++ using Microsoft Visual Studio version 6. The graphical user interface presented to the user was generated using bitmaps displayed through the Direct X API. This allowed the simple creation of a visually exact replica of the Start Menu. The haptics code was created using the GHOST API supplied by SensAble Technologies. The force feedback was provided by a PHANToM 1.0 [128] equipped with a pen stylus featuring a button. The operation of this button was linked to mouse click events. The workspace available to participants was restricted to a narrow vertical plane, 110 mm wide by 110 high mm and 2 mm deep. Motion along the x and y axes controlled cursor position. As the two haptic augmentations varied in only two dimensions (x and y), they were continually and uniformly presented to the user regardless of their position on the z-axis. Furthermore, no explicit action was mapped to motion on the z-axis.

4.3.5. Participants

There were eighteen participants in this study. Eleven were male, and seven female. All were between the ages of eighteen and forty-three. Most were computing students and all were regular and fluent computer users. Two participants were left-handed. No subject had anything more than trivial previous exposure to haptic interfaces.

4.3.6. Experimental Design

Three conditions were examined: Visual, Haptic, and Adjusted. The Visual condition did not incorporate any haptic feedback designed to improve targeting. The Haptic barriers were present over all targets in the Haptic condition and, similarly, the Adjusted barriers were present over all targets in the Adjusted condition. The experiment followed a repeated measures design and lasted approximately 1 hour. The eighteen participants were arranged into six order conditions each containing three subjects, resulting in a fully balanced experimental design. Each experimental condition consisted of the same set of 100 trials, presented in a random order. As the menu system extended to a depth of at most 4 sub-menus, 25 trials were drawn from each possible depth. Training on all conditions took place immediately prior to the experiment. The training conditions were presented in the same order as those in the experimental session and contained half the number of trials as the experimental conditions.

4.3.7. Measures

Timing Measures

Time to complete each trial was measured, taken from the time the Start button was pressed, until the selection of the correct menu item. If the Start button was pressed several times during the course of a single trial (as would be the case if the menu had collapsed back to its initial state due to the selection of an incorrect item) the most recent selection of the Start button was used as the start time for the trial.

Error measures

A variety of errors were also recorded. These followed a basically similar scheme to those in the buttons study described previously, adjusted for the different interactions present in menu operation. Firstly, the number of slide-over errors, where a

participant moved over the correct target, then moved off without progress towards the completion of the trial, were counted. Included in this total were both failures to select the named target, and failures to move on to an appropriate sub-menu when the opportunity occurred. This reflects the fact that menu interaction typically involves multiple targeting operations in order to reach a single target. Skip-ahead interactions, where a user took a shortcut from the current item to the next submenu by skipping over an adjacent item were also counted. These were collected in order to modify the number of slide-overs recorded, so that it more accurately reflected actual errors, and also to provide a measure of the conformity of a user's path to the ideal path through the menu. This led to an adjusted count of slide-over errors termed modified-slide-overs. Finally, the number of selections of wrong-items was counted. As menu selection is a single stage process (requiring only the release of the controller button over the desired target) slip-off errors could not occur. In this case, it seems likely that wrong-items errors will occur in the place of slip-offs.

Subjective measures

As in the buttons experiment described previously, subjective workload was recorded using a modified NASA TLX [88] questionnaire. The modification was in the form of a single additional term: fatigue experienced. A questionnaire was also administered at the end of the experiment. This simply collected basic demographics and is presented in Appendix D.

4.3.8. Procedure

All participants first performed a practice session. At this time written instructions were presented explaining the experimental task, and participants were free to ask any questions to clarify these. During the course of the practice participants were encouraged to explore the different haptic effects, to become accustomed to them. Between each condition they were instructed that they were free to rest for as long as they wanted.

After completion of the practice participants immediately moved onto the experimental session. This was similar to the practice session except that participants were required to fill out a NASA TLX subjective measures questionnaire [88] after

the completion of each condition. At the end of the entire study, participants filled out a brief questionnaire that gathered demographic information.

4.3.9. Hypotheses

The study investigating haptically enhanced buttons presented earlier in this chapter suggests that the Haptic condition may impair performance when compared to the Visual condition in situations where multiple haptic targets are presented simultaneously. This is due to the fact that haptic augmentations composed of attractive forces can lead to users experiencing difficulties as they attempt to leave a target. Such an event is likely to occur numerous times during the performance of any interaction in a complex multi-target environment, and it seems likely that this will impact negatively not only on task completion time, but also on error rate and subjective workload. The Adjusted condition has been designed to overcome these problems, and the critical hypothesis of this study is that it should yield improved task completion times, error rates and subjective workload when compared to the Visual or Haptic conditions.

4.3.10. Results

The raw data from this study are presented in Appendix E. All analyses of the data generated in this study were conducted using repeated measures ANOVA [48], followed by *post-hoc* pair-wise comparisons [48] using Bonferroni confidence interval adjustments [93].

Temporal Analyses

The results from the timing data are presented in Figure 4.27. An ANOVA revealed significant differences in this data ($F(2, 17)=11.815$, $p<0.01$) and the subsequent pairwise comparisons showed that users worked significantly slower in the Haptic condition than in both the Visual and Adjusted ones (respectively $p<0.05$ and $p<0.01$).

Error Analyses

The results from the error data are presented in Figures 4.28, 4.29, 4.30 and 4.31. ANOVA tests on all the error measures yielded significant results: slide-over ($F(2,17)=22.966$, $p<0.01$), wrong-target ($F=8.368$, $p<0.01$), skip-ahead ($F(2,$

Figure 4.27. Mean time per trial in menu study.





Figure 4.28. Mean total slide-over errors in menu study.



17)=65.520, $p<0.01$) and the derived modified-slide-over ($F(2, 17)=20.1$, $p<0.01$). Pairwise comparisons revealed that the Adjusted condition led to significantly less slide-over errors than both the Visual and Haptic conditions (both at $p<0.01$). The Adjusted and Haptic conditions produced significantly fewer occurrences of skip-ahead behaviour than the Visual condition (both at $p<0.01$). The Adjusted condition also led to significantly fewer modified-slide-over errors than either the Visual or

Haptic conditions (both at $p < 0.01$). Finally, the Adjusted and Haptic conditions produced significantly less wrong-target selections than the Visual condition (both at $p < 0.05$).

Figure 4.29. Mean total wrong-target errors in menu study.




Figure 4.30. Mean total skip-ahead actions in menu study.





Figure 4.31. Mean total modified-slide-over errors in menu study.



Subjective Measures

The results from the NASA TLX are presented in Figure 4.32. ANOVAs revealed significant differences in Overall Workload, and on all individual scales bar Performance Level Achieved. These figures are illustrated in Table 4.6. Pair-wise comparisons revealed that the Visual and Adjusted conditions received significantly lower ratings of Overall Workload than the Haptic condition (both at $p < 0.01$). Examining the individual scales, the Visual and Adjusted conditions were both rated as less taxing than the Haptic condition on Physical Demand (both at $p < 0.01$), Effort Expended (both at $p < 0.01$), Frustration Experienced ($p < 0.05$ and $p < 0.01$ respectively) and Fatigue Experienced (both at $p < 0.01$). The Adjusted condition also received a significantly lower rating of Mental Demand than the Haptic condition ($p < 0.01$). Finally, the Visual condition was rated significantly lower on the Time Pressure scale than the Haptic condition ($p < 0.05$). No significant differences were observed in the TLX results between the Visual and Adjusted conditions.

4.3.11. Discussion

Not all the experimental hypotheses were upheld. The Haptic condition did yield inferior performance on measures of time and workload when compared to the Visual condition, but still maintained equal or better levels of performance on the various error measures. The Adjusted condition did solve the problems discovered

Figure 4.32. TLX workload results from menu study.




Table 4.6. Results from ANOVA analysis of TLX results in menu experiment.

	MD	PD	TP	EE	PLA	FE	FatE	Overall
F value	9.570	24.227	4.574	14.483	1.798	7.845	13.318	17.699
P value	p<0.01	p<0.01	p=0.05	p<0.01	Not sig.	p<0.01	p<0.01	p<0.01

with the Haptic condition; it resulted in reduced task completion time and workload. However, it did not reduce all error counts. Furthermore, although it produced significantly fewer errors than the visual condition, it did not offer an improvement on task completion time and workload.

Essentially, the results indicate that the Adjusted condition combines the best aspects of both of the other conditions. Target selection errors are reduced to the same level as that observed in the Haptic condition, while speed and workload are not compromised when compared to the Visual condition. Users also found it easier to target in the Adjusted condition than in either of the other conditions, as indicated by the lower number of slide-over errors, while the power to ensure a user remains on a path is similar to that found in the Haptic condition, as indicated by the similar number of skip-ahead interactions.

A contributing factor as to why the reduction in error numbers achieved by the Adjusted condition (when compared to the Visual condition) did not lead to a corresponding reduction in task completion time, is that one critical category of errors – the wrong-target errors – were explicitly separated from the timing data: the clock was reset whenever one of these errors occurred. This manipulation, and the separation of these data, was intentional, as it was thought that this would provide more detailed information as to actual user experience: it makes it possible to record error rates, and to examine temporal performance that is not simply a function of these rates. However, if the time to recover from these errors had been included in task completion time, it seems likely the Adjusted condition would have attained the temporal improvement over the Visual condition the hypothesis of this study predicted.

The results of this study suggest that task completion time is increased in the Haptic condition because of the presence of interfering forces, and that workload is greater because of the disruptive nature of these forces. Observationally, the relatively large number of slide-over errors reported in the Haptic condition can be attributed to a problem users experienced at the closing stages of the targeting process. As they approached their final target they would accidentally stop on an adjacent target. They would then experience difficulties overcoming the strength of the haptic barrier between their position and their desired target. Eventually they would end up pushing sufficiently hard so as to not only move through the haptic barrier, but to also “pop through” the adjacent, desired target and move onto the target beyond that. They would then be presented with the same problem once again – moving onto an adjacent target. This process could repeat several times until the user managed to stop on the correct target.

Previous research on moving along haptified menu items [51] indicates that speed should increase with the additional feedback. It is possible that such an effect was not present here because users were unaware of the potential benefits of the Adjusted condition. Participants felt it was simply weaker than the Haptic condition and consequently took little advantage of the strong targeting forces provided. As informal evidence supporting this claim, after the experiment one participant spontaneously remarked that she was at a loss to understand the increased stability

she had observed in her movements under the Adjusted condition. This supports the previous claim (made in the study investigating haptically enhanced buttons) that, for users to gain the optimum performance using a haptically augmented GUI, appropriate strategies must be in use. It seems likely that more prolonged exposure, or explicit awareness of the beneficial forces, may lead to the generation of these strategies and therefore to measurable decreases in task completion time.

4.3.12. Conclusions

In conclusion, this study demonstrates that although directly applying single target haptic augmentations to a multi-target case is ineffective, careful consideration of the interactions being supported can transfer the benefits of haptic feedback in single target situations, at no cost. It proposes and validates a novel method of mediating the application of haptic feedback in situations incorporating multiple simultaneously active haptic targets. This method, based around altering the forces applied to a user according to directly measurable aspects of behaviour provides an alternative to target prediction systems that does not suffer from their problems. Unlike target prediction systems, the method described does not require complex algorithms to derive intended user behaviour, nor substantial training times. Consequently, although further validation of this technique with other groups of widgets is required, this study shows that it has the potential to serve as a viable alternative to target prediction.

4.4. Further Work on this Topic

Although the menu study described above provides a compelling example illustrating that appropriately adjusted haptic feedback can provide performance improvements in situations involving multiple simultaneously active haptic targets, it is questionable whether or not a designer or system builder would be able to apply the techniques it advocates to other widget groups. This one example does not clearly illustrate how to create appropriate feedback for use in more general situations. As the dissemination of a methodology for the creation of haptic enhancements for widgets is one of the goals of this chapter, this limitation was addressed through the design of haptic feedback to be used with other widget groups. Specifically, feedback was designed to support interactions with a toolbar, or other compact arrangement of very small targets. This example was chosen because

it represents a situation in which users are able to move in a large variety of directions in order to complete their tasks. Unlike the haptic cues used in the menu study, no simple rule about the likely directions users will move in can be used to create more effective feedback.

4.4.1. Haptic Feedback

The haptic feedback designed to support interactions with a toolbar was based on the same basic building blocks as that used in the menu study. Four haptic barriers were used to enclose each target. However, reflecting the smaller size of toolbar buttons compared to menu items, these barriers were subject to a lower maximum magnitude: 0.25 N.

As with the barriers in the menu study, a number of dynamic adjustments to the magnitude of the force applied were defined with the goal of better supporting a user's interactions. These adjustments partly differ to those used in the menu study, due to variations in behaviour and layout between the two widget groups, but remain derived from directly measurable aspects of user input. The following three rules were used to modify the maximum applied strength of the Haptic barriers:

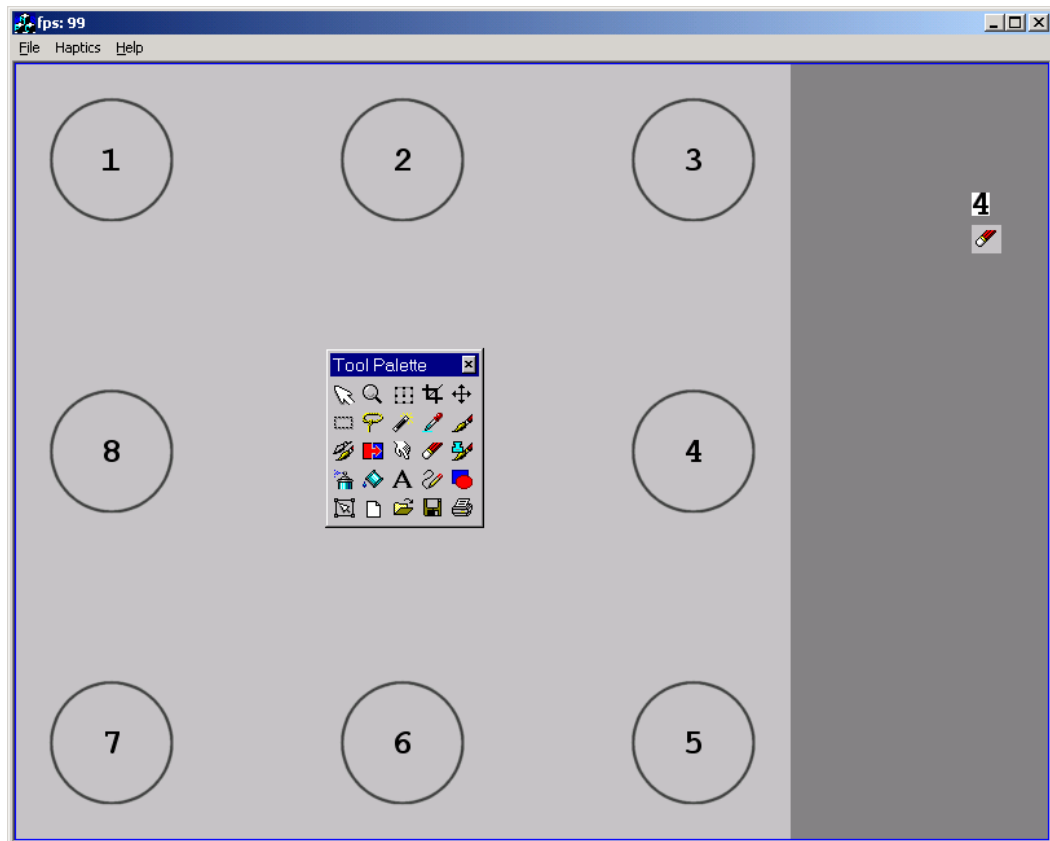
1. Reduce the maximum force applied if a user is moving slowly (beneath 2 cm per second) to a minimum of one third of its normal value. This alteration was taken directly from the menu study, and had a similar goal: it was designed to allow users to easily move onto targets adjacent to their current position.
2. If a user is moving rapidly (above 2 cm per second) and has only been on a button for a short time (less than 100 ms), reduce the maximum applied force by a factor of two. The goal of this modification was similar to that distinguishing between horizontal and vertical motion in the menu study. Essentially, it was designed to facilitate rapid unobstructed movements, by providing reduced forces when such motions are deemed to be taking place.
3. Increase the maximum force applied to three times its original amount if a user has begun to perform a click (by depressing the controller's button) and reduce the force back to normal levels when the click is completed (by releasing the button). The inclusion of this modification reflects the fact that

targets in a toolbar, unlike those in a menu, function like normal GUI buttons: selection is a two-phase process involving both the depression and release of the controller button over the target. This modification was included in order to increase the likelihood that a clicking action, once begun, would be successfully completed.

4.4.2. Evaluation of this feedback

A report by Adams [2]¹ (and published as Oakley *et al.* [147]) discusses the evaluation of this haptic feedback in some depth. As this work has a direct bearing on this thesis it is described in detail here. Adams's report examines the influence of this feedback on two different target selection tasks. The first of these focused on the widget group the feedback was designed for - the buttons in a small toolbar - while

Figure 4.33. Screenshot of toolbar study.



¹ This report formed an MSc IT thesis at Glasgow University, and Adams worked under the direct supervision of the author of this thesis.

the second looked at a randomly arranged set of much larger icons. Critically, the difference between these evaluations is that the toolbar study examined a highly structured and closely packed arrangement of small targets, much like the menu study, while the icons study investigated a more widely dispersed, unorganised and much larger group of targets. One of the goals of this second study was to contrast the effectiveness of the same haptic feedback with two different widgets groups. Screen shots from these studies are shown in Figures 4.33 and 4.34.

As with the menu study, each evaluation included three conditions: a Visual condition, a Haptic condition and an Adjusted condition. The Visual condition did not incorporate additional haptic feedback, while the Adjusted condition used the feedback described immediately above. The Haptic condition used the same arrangement of barriers as the Adjusted condition but without the presence of any dynamic modifications. In the study examining icons one alteration to the feedback described above was made: the maximum magnitude of the forces exerted by the haptic barriers was altered from 0.25 N to 0.65 N. This increase was a design decision and reflects both the larger size of the targets, and the increased area they occupied.

Figure 4.34. Screenshot of icons study.




Table 4.7. Data from toolbar study reported by Adams [2].

	Visual	Haptic	Adjusted
Mean Trial Time (seconds)	1.31	1.5	1.25
Mean Overall Workload	9.92	12.86	9.09
Mean Total Slip-offs	11.33	0.44	0.17
Mean Total Slide-overs	49.5	71.94	32.06
Mean Total Wrong-targets	4	3.89	2.72

Table 4.8. Statistics from toolbar study reported by Adams [2].

	Results of ANOVA	Results of <i>post-hoc</i> comparisons
Mean Trial Time (seconds)	F(2, 17) = 18.453, p<0.01	Visual and Adjusted less than Haptic (p<0.01)
Mean Overall Workload	F(2, 17) = 13.264, p<0.01	Visual and Adjusted less than Haptic (p<0.01)
Mean Total Slip-offs	F(2, 17) = 18.269, p<0.01	Adjusted and Haptic less than Visual (p<0.01)
Mean Total Slide-overs	F(2, 17) = 18.101, p<0.01	Adjusted less than Visual & Haptic (both at p<0.01), Visual less than Haptic (p<0.05)
Mean Total Wrong-targets	F(2, 17) = 0.581, not sig.	No significant differences

Measures

Both studies were subject to the range of subjective and objective measures used in the menu and button studies described in this thesis. Task completion time was gathered for each trial. Error measures included slide-overs, wrong-targets, and slip-offs. Slip-offs were recorded in these studies as the two stage selection process (the controller button must depress and release over the desired target) present in toolbar buttons and icons allowed them to occur. Subjective data were gathered using a NASA TLX modified to include a Fatigue Experienced term.

Results and Discussion

The results of the toolbar study strongly supported those reported in the menu study. Repeated measures ANOVA followed by *post-hoc* t-tests revealed that both workload and task completion time were increased in the Haptic condition when compared to either the Visual or Adjusted conditions. The Adjusted and Haptic conditions produced significantly fewer slip-offs than the Visual condition. The Adjusted condition also yielded significantly fewer slide-overs than the Visual and Haptic conditions. The raw data from this study is presented in Table 4.7, while Table 4.8 contains details of the statistics.

In short, the toolbar study reinforced the finding that appropriately adjusted haptic feedback can combine the favourable aspects of standard visual and haptic feedback.

Table 4.9. Data from icons study reported by Adams [2].

	Visual	Haptic	Adjusted
Mean Trial Time (seconds)	0.91	0.91	0.85
Mean Overall Workload	7.92	8.18	7.03
Mean Total Slip-offs	21.42	3.08	2.33
Mean Total Slide-overs	56	32.25	30.67
Mean Total Wrong-targets	2.08	4.67	3.17

Table 4.10. Statistics from icons study reported by Adams [2].

	Results of ANOVA	Results of <i>post-hoc</i> comparisons
Mean Trial Time (seconds)	F(2, 17) = 6.549, p<0.01	Adjusted less than Visual and Haptic (p<0.05)
Mean Overall Workload	F(2, 17) = 2.77, not sig.	No significant differences
Mean Total Slip-offs	F(2, 17) = 20.632, p<0.01	Adjusted and Haptic less than Visual (p<0.01)
Mean Total Slide-overs	F(2, 17) = 18.796, p<0.01	Adjusted and Haptic less than Visual (p<0.01)
Mean Total Wrong-targets	F(2, 17) = 7.058, p<0.01	Visual less than Haptic (p<0.01)

The Adjusted condition attained the low error count apparent in the Haptic condition and the rapid task completion time present in the Visual condition. The Haptic condition was more subjectively taxing than the other two conditions. These results support the claim that the standard haptic feedback that is effective in a single target situation leads to a degradation of performance when applied in a situation incorporating multiple targets. Like the menu study described earlier, it also demonstrates that appropriately adjusted haptic feedback can lead to performance benefits in these complex environments.

The icons study yielded data less directly coherent with the menu study. The subjective results indicated that the Adjusted condition yielded lower subjective workload on the scales of Effort Expended and Performance Level Achieved when compared to the Visual condition and reduced Frustration Experienced when compared to the Haptic condition. The Adjusted condition also resulted in significantly faster times than either the Visual or Haptic conditions and fewer slip-off and slide-over errors were present in the Adjusted and Haptic conditions when compared to the Visual condition. Finally, the Visual condition resulted in fewer wrong-target errors than the Haptic condition. Table 4.9 contains the raw data from this study, and the statistics are presented in Table 4.10.

To sum up, the Adjusted condition combined the fastest task completion times in the study with the lowest error rates, and showed a modest gain over the Visual and

Haptic conditions in subjective measures. However, the Haptic condition, while exhibiting the expected improvement in error rate compared to the Visual condition, did not produce the predicted performance reduction in terms of task completion time and subjective measures. As the same feedback was responsible for this performance hit in the toolbar and menu studies, this may indicate that target acquisition tasks relying on large, spatially separated targets are not sensitive indicators of performance. Techniques that improve performance in these situations may fail in more challenging scenarios and the ability to generalise from them should be questioned. A final interesting aspect of the results from this study is that the Visual condition resulted in fewer wrong-target errors than the Haptic condition. From observations, it seemed likely that this was a consequence of participants getting “snagged” on nearby distracter targets, when, without the haptic feedback, their velocity would have normally carried them over their desired target. It is also possible that participants were inadvertently using the haptic feedback as a trigger for button selection. As they moved over a target, the presence of the haptic feedback may have caused them to automatically select the target regardless of whether or not it was the correct target. This kind of automatic selection behaviour in response to haptic feedback has been previously observed [52].

Conclusions

The two studies conducted by Adams [2] and described in some detail here reinforce the precedent set by the menu study. They demonstrate that beneficial haptic feedback for use in complex multi-target scenarios can be produced through dynamic adjustment of the forces applied based around directly measurable aspects of user behaviour. In conjunction with the menu study they provide a compelling and relatively robust portfolio illustrating this claim, and the generation of such a set of examples was the primary goal of this research.

4.5. Design Guidelines for haptic widgets

4.5.1. Introduction

The studies described here indicate that the haptic augmentation of widgets in GUIs is a challenging problem. The first study explores the possibilities of haptic augmentation in a relatively simple task involving only a single haptically active

target. It highlights interesting and unexpected aspects of the influence of different haptic effects. The other studies described here examine haptic augmentations in situations featuring multiple simultaneously active haptic targets and indicate that augmentation with unadjusted attractive forces is at best not optimal, and at worst can reduce performance and subjective satisfaction. The fact that a number of papers have been published looking at the potential for target prediction to prevent these problems further supports this claim. However, the research reported here also suggests that appropriate haptic feedback – haptic feedback that provides performance improvements at no cost – can be created through various manipulations of standard haptic augmentations such as walled areas.

The complexity of augmenting widgets with haptic cues revealed both by the studies presented here and by the previous literature [3, 65, 90, 135] highlights a problem for the adoption of this kind of feedback by designers and software developers in the real world. It is possible to create systems that either improve or disrupt performance, and there are no guidelines suggesting how the former can be achieved and the latter avoided. To address this issue, this thesis presents an attempt to make explicit the process by which the performance enhancing haptic augmentations described here were created. This involves the presentation of the first set of guidelines for the design of effective haptic widgets to function in both single and multi target situations. Due to the very preliminary nature of these guidelines they are somewhat speculative. However, despite this, this thesis contends that they are both explanatory and informative.

4.5.2. Literature Review

Miller and Zeleznik [135] have previously presented three guiding principles to aid the creation of haptic widgets. Firstly, they suggest that haptic feedback should be used to reduce errors through guidance; to provide forces to support the motions that a user is undertaking. Secondly, they indicate that the forces applied should function as feedback; they should be based upon, but never control, a user's input. Finally, they state that any force feedback applied to a user should be overridable; a user should be able to pop through, or escape, from any haptically augmented area. The guidelines presented here share similar tenets to these principles. However, they try

to go further, to more precisely define the problems and solutions involved in adding haptic feedback to desktop widgets.

These haptic design guidelines are based around the idea that the force presented should support, and not oppose, a user's intent. This entails drawing a balance between allowing users to move where they want as freely as possible, and providing forces to improve targeting and reduce errors. An advantage of these guidelines is that they do not require target prediction, a currently immature technology. A disadvantage is that they do assume that an extremely dynamic simulation controls the haptic feedback. The guidelines rely on the rapid, smooth adjustment of force magnitude according to the current state of interaction, and this flexibility may be challenging to implement on some current consumer-level devices [118].

4.5.3. Methodology for Creation of Guidelines

User interface design guidelines are commonplace. The term encompasses a large number of very different artefacts ranging from the relatively abstract - such as the principles that underlie direct manipulation [173] - to the very specific - for instance, the details concerning creating conforming screen or dialogue layout on a given platform [7]. Consequently, there is little general formalism for the construction of design guidelines.

However, research has been conducted on the usefulness of design guidelines [92, 190], with the general conclusion that they are often difficult to interpret, and that practical examples illustrating how the concepts described should be instantiated are essential. For instance, Tetzlaff & Schwartz [190] report that the different participants in their study missed a wide variety of different critical concepts and details while attempting to create a system adhering to a set of interface design guidelines, and they point to the use of concrete examples as one way of rectifying this problem. Attempting to address similar issues, Henninger *et al.* [92] describe a system to augment the creation of design guidelines so that some of the wider context that the guidelines relate to is also captured. Baring this literature in mind, the guidelines detailed here attempt to describe general concepts, but to remain firmly rooted in practical examples.

4.5.4. Guidelines

Guiding Strategy

Haptic feedback has the potential to improve objective user performance in two ways: reducing the number of errors made, or decreasing task completion times. The research reported in this thesis suggests that it is more profitable to design haptic augmentations to achieve this first aim, to reduce errors. There are several reasons for this. Firstly, a reduction in errors can be linked to improvements in other metrics: it has been associated with decreases in task completion time [52], and some studies, including the buttons study presented in this chapter, have linked it to increased subjective satisfaction. Secondly, to gain an improvement in task completion time, users must adopt a movement strategy supported by the haptic augmentation, and there is no guarantee that this will occur. There are more assurances that forces to prevent errors will be successful. For instance, an attractive basin supports faster movement times, because targeting is simpler, and a reduction in errors, because a conscious effort is required to leave the target. However, although users may move towards the target more rapidly, this is a choice they make. The decrease in errors, on the other hand, is simply a property of the attractive forces applied over the target.

Choice of haptic augmentation

The buttons study presented in this thesis illustrates that widgets augmented with attractive basins or haptically walled areas typically provided the best performance improvements. However when designing the haptic feedback for a widget, it is also important to consider its shape, and the likely path a user will take over it. For instance, when using these standard targeting augmentations in conjunction with square or rectangular widgets, diagonal motion is more difficult than horizontal or vertical motion. These kinds of factors may have an impact on performance and subjective satisfaction, depending on the situation in which they are used. For instance, in the menu study described in this thesis, diagonal motion is extremely rare, and consequently this issue is unlikely to influence performance or satisfaction. On the other hand, in the toolbar study described above diagonal motion is much more common, and this issue has the potential to exert a more serious effect.

Interaction between force strength and widget size

The maximum strength used for any widget, or set of widgets, should be dependent on both the size of the widgets and density of the arrangement that they are presented in. As the toolbar study described above indicates, a dense arrangement of small widgets requires small forces, as large forces will severely hamper motion from one widget to an adjacent one. Also, motions over small, well packed widgets are likely to be slower, as only a short distance must be travelled. Consequently small forces are sufficient to aid targeting. Correspondingly, large, spatially separated widgets suit much stronger forces, as illustrated in the icons study described above. With the absence of nearby widgets, the presence of these stronger forces is less likely to be disruptive. Also users often approach large spatially distributed widgets at considerable speed. Thus, targeting benefits are likely to be maximised by increasing the strength of the forces applied, to match the increase in approach speeds.

Range of useful force magnitudes

The literature suggests that the maximum strengths of haptic targets should be in the range of 0.25 N (used in the toolbar study here) to 0.8 N (a large force used in the some studies of stand-alone widgets [51]). There is little research indicating whether these figures would be device dependent, and it is worth noting that for use in multi-target situations, feedback of these magnitudes will typically need to be adjusted as described later in these guidelines. Maximum applied strength is also likely to be highly dependent on individual differences. In any real system, it would be essential that maximum strength be user configurable. However, it seems likely that the general strength ratios between different sizes and densities of widgets would stay more or less the same across users; irrespective of the maximum strength a user chooses, the proportions between the magnitude of the forces applied over a large target, and of that applied over a small target should remain the same.

Exploit patterns of user behaviour

The haptic feedback present on a widget should capitalise on patterns of motion afforded by that widget. This is often related to the shape of the widget. In the study of haptically augmented menus described here, the fact that motion in a menu typically occurs either vertically or horizontally was exploited to provide only supportive forces. Given the similar shape of the widgets, this same manipulation has

the potential to apply to scrollbars. The scrollbar could exert strong targeting forces as a user moves along its length, but fade out these forces when a user attempts to move off it, in a direction perpendicular to its length. In a scrollbar it may also be appropriate to increase the strength of the targeting forces with increased speed along the scrollbar's length.

Exploit widget behaviour

Widget behaviour can also often be exploited to increase the effectiveness of a haptic augmentation at no cost. In both the toolbar and icons studies described above the Adjusted conditions incorporate haptic feedback designed to aid the completion of an action that has been begun. The strength of the haptic walls increases when a user begins to interact with a target, and reduces to normal levels on the completion of that action. This same strategy could be applied to any widget interaction that incorporates more than a single explicit stage. Beginning the interaction could trigger a change in the haptics, such as an increase in magnitude, designed to support the successful completion of the interaction.

Dynamic response to slow movement

Force should vary according to speed: slow motions require low forces. In situations incorporating densely packed widgets this is especially important. It is hard to traverse from one widget to an adjacent widget when opposed by even a relatively low force. Users often end up moving further than they intended, “popping through” the target widget onto one beyond it. They are then faced with the same task again – moving to an adjacent widget. This can lead to very frustrating interactions, and is arguably the biggest problem with multi-target haptic widgets. Varying the applied force such that slower motions are opposed by lower forces can overcome this problem, allowing users to move freely, while still providing a sufficiently strong force that accidental movements off a target are prevented. The strength of force required to support targeting clearly reduces in tandem with a reduction in the speed at which a user is moving, and to produce effective multi-target haptic augmentations it is essential to capitalise on this fact.

Dynamic response to rapid movement

Equally, an extremely rapid motion over a target typically indicates that it is not a user's final destination, and thus requires the application of low forces. Users do not want to be impeded by widgets that are nowhere near their final destination. Again this is especially important in situations where there is a high density of widgets. In these situations it is likely that users will traverse over numerous irrelevant widgets before reaching their desired targets. One mechanism to achieve this is that used in the toolbar and icon studies described above, where weaker forces are applied during the first 100 milliseconds that a user is over a widget. A disadvantage of this manipulation is that it may decrease the effectiveness of the behaviour observed by Dennerlein *et al.* [52] in which users throw themselves at speed towards a haptic target, relying on the forces it exerts to halt them. Reducing these forces for the first few moments that a user is over a target may make it less effective at capturing a rapidly moving user.

4.5.5. Conclusions from Guidelines

The guidelines described here serve several purposes. First and foremost, they represent the first detailed set of experimentally derived guidelines in this area: they describe the rules of thumb that led to the creation of the performance boosting widget sets described both in this thesis and by Adams [2]. Despite the growing body of research on this topic, there is no other formally presented set of design guidelines for the design of haptic widgets. They also bring together basic data published by a number of different researchers [51, 90] on the various haptically augmented widgets they have created. Consequently, they may function as a useful summary of this literature.

Finally, while this work focuses on augmenting standard GUIs, it is possible that the guidelines presented here may have further applicability. Many fish-tank VR systems [6] incorporate both haptic feedback and interface widgets, and the research described here will translate easily to these systems. Other uses may include haptic systems for scientific visualisation [60], or for visually impaired people [205]. In both these scenarios users are often required to explore complex arrangements of haptic targets. Applying the techniques outlined in these guidelines could make these tasks simpler and less demanding.

4.6. Conclusions

This chapter presents two studies investigating the use of haptic feedback to support user interactions with interface objects. The first study examines the application of haptic feedback to a simple target acquisition task, and has the goal of determining what sort of haptic feedback is appropriate for these tasks. The results indicate that haptically walled areas or wells of attractive force provide the greatest performance improvements. The second study builds on this work by applying the kind of feedback found to be successful in the first study to a more complex situation involving the simultaneous presentation of multiple targets. The results of this study indicate that directly transferring haptic augmentations that are effective in a single target situation to a multi-target situation is at best not optimal, and can lead to reductions in performance. However, this study also demonstrates that appropriately adjusted haptic feedback can result in performance improvements in multi-target situations. These findings are reinforced by follow-up work reported by Adams [2]. The final contribution of this section of this thesis is a set of design guidelines for the creation of haptically enhanced interface widgets. These guidelines are the first to be presented on this topic. They draw together the previous literature, and attempt to make explicit the process by which the empirically validated widgets described here were created.

In terms of the structure of this thesis, this chapter conclusively shows that the addition of haptic cues to a user's cursor can provide performance benefits in the everyday scenario of interactions in a GUI. It resolves some of the previous work reporting conflicting data on this topic [4, 65, 90], provides firm steps towards solving the significant problem of haptic augmentation in multi-target situations, and provides a framework under which other system builders can create feedback with similar beneficial properties. These achievements are important as this GUI scenario represents one of the specific situations chosen to demonstrate the general claim of this thesis: that the addition of haptic cues relating to a user's avatar in a computer system is beneficial. The research presented in this chapter strongly supports this claim, and the design guidelines presented achieve the secondary goal of preserving the insights collected in this chapter in a form suitable for use by future designers and system developers.

5. Haptics in Communication

5.1. Introduction

This chapter reviews the literature surrounding the role of haptic feedback in collaborative computer systems. It starts with a broad discussion of the potential motivations for adding haptic feedback to such systems, and then moves on to defining the specific class of collaborative systems that this section of this thesis is concerned with: shared editors. A description of some of the critical issues underpinning shared editors is followed by a discussion of the problems researchers have observed with these systems and some of the solutions that they have generated. Issues involved in evaluating collaborative systems are then discussed. Finally, there is an extensive review of the literature directly pertaining to touch communication. This includes previous systems that support communication through touch, but also the more theoretical aspects of this domain.

Haptic communication is the second case study, or exemplar, chosen to illustrate the general claim (and the crux of this thesis) that haptic cues pertaining to a user's representation (or avatar) in a computer system can improve a user's subjective experience of that computer system. This chapter fits into the structure of this thesis by providing a description of the context in which the haptic communication, and the evaluations of the haptic communication, were designed. It provides a detailed description and discussion of the background literature relevant to computer mediated communication in general and communication through touch specifically.

5.2. Motivations

With the advent of the Internet and globally connected computing, the desire to use this technology to allow distributed users to work together in structured ways has grown. In tandem with this growth, more and more work has become team-based. Consequently, systems to support collaborative work have many potential users and could provide substantial benefits. From a business perspective, work could be made more efficient and money saved by reducing the need to transport key employees to multiple sites. Expert help could also be brought in rapidly and with little disruption. Indeed, sophisticated tools for collaborating are becoming integrated into standard packages such as AutoCAD [9]. From an individual perspective, working from home could become a realistic possibility for workers in a far larger range of professions. Individual advice on troubleshooting computer problems is already available (at www.expertcity.com) in an interactive web-based collaborative system. The research area that investigates the role that computers can play in communication is known as Computer Supported Co-operative Work (CSCW), and applications that enable communication are known as groupware applications.

Despite these promising benefits, most groupware applications that attempt to support complex group activities, such as coordinating shared calendars, group project management, or the shared editing of documents, have failed. Organisational issues are an important factor contributing towards this, and Grudin [82] discusses the organisational impact of CSCW systems from the perspective of attempting to understand their consistent failure.

He isolates three significant organisational issues. Firstly, the disparity between who must do additional work to support the collaborative system and who gains benefit from the system. Grudin suggests that those who commission CSCW systems, individuals at a relatively high level of management, are likely to benefit from those systems at the expense of extra work from their subordinates. This leads to a situation where the system is viewed negatively by the vast bulk of its users, which has severe implications for its success. He suggests CSCW applications must benefit all users as equally as possible in order for them to become widely adopted. Secondly and thirdly, he describes two factors that contribute to a general lack of understanding of the requirements for a good CSCW system. Essentially, these two

points can be summarised by the statement that collaborative applications cannot be viewed in the same way as single user applications. To do so leads to a skewed perspective. He states that this distorted view is present among both those who develop CSCW systems, and those who evaluate them. He also provides four detailed case studies which provide a compelling illustration of his arguments.

However, another significant reason contributing to the general failure of groupware applications is that such systems do not support the richness of communication that is available to co-present users working in the real world [34]. In the limited environments currently available to them, users find it difficult to co-ordinate their work or discuss complex issues. In short, users find the majority of current groupware systems to be challenging and hard to use productively.

One reason for this is that these systems rely almost entirely on the visual channel for communication. Typically, a full duplex audio link transmitting speech is the only non-visual information provided. In contrast, in the real world we gain information of all sorts through all of our senses simultaneously. Audio enhancements of collaborative systems have been created [73] and, within their research context, demonstrated improvements in user communication and interaction. There has been little work, however, on the addition of haptic feedback to CSCW systems. It is possible that haptic cues would yield benefits in collaboration and communication scenarios. Exploring this suggestion, the next two chapters describe the background, design, implementation and limited evaluation of one use of such cues.

5.3. Literature Review

5.3.1. CSCW and Groupware

Communication is fundamental to human existence. An increasingly diverse set of sophisticated technologies to support communication – telegraph, telephone, e-mail, pagers, mobile telephones, World Wide Web (WWW) – have proven to be incredibly popular across different cultures and through decades. Computers offer the facility for further enriching communication in a host of ways.

Computer Supported Co-operative Work, the general discipline concerned with how computers can enable communication, encompasses a diverse spread of topics. These range from the psychological theories of group work, to how these theories are influenced by the mediation of communication by a computer system, to providing novel channels of communication for collaborators to utilise. It is important to consider that, in contrast to single user applications, the most important aspect of a groupware system is usually not what functionality it directly supports in whatever application area it is concerned with, but instead how it enables communication among its users. The functionality of the application is important, but regardless of this, it will only be successful if it adequately supports the communication needs of its groups of users.

Due to the ubiquitous nature of communication, groupware is not tied to any specific task domain or methodology. In Table 5.1 Johansen [106] provides a simple way of subdividing the range of groupware applications. He, and several others before him, categorise groupware along two simple axes – time and place. This leads to a description of distinctly different systems. In the top left cell of the table, participants are present in the same location at the same time. CSCW tools such as augmented whiteboards [156] attempt to aid peoples’ work by providing digital enhancements to a physically shared space. Typical enhancements are multiple pages, undo commands, the ability to save information for later use and provision for the creation of lists or structured diagrams. In the top right cell, users are working in the same place, typically on the same task but are present at different times. Groupware attempts to support these activities by providing version control systems or facilities for annotating shared objects. In the bottom right, users are neither communicating at the same time nor in the same place. This cell encompasses asynchronous communication systems such as email and voice mail. The bottom left cell deals with

Table 5.1. Johansen’s [106] classification of groupware applications.

	Same Time	Different Time
Same Place	Face to Face Meetings Augmented Whiteboards Meeting Rooms	Administration/Data Management Shared Files Version Control
Different Place	Remote Meetings Teleconferencing Shared editors	Co-ordination Email Voice Mail

remote simultaneous communication. It incorporates video conferencing and shared editors. Since the advent of the World Wide Web (and technologies such as webcams, video-enabled instant messaging and Microsoft Net Meeting [134]) video conferencing, while still subject to considerable interaction problems, has become relatively common. Shared editors are the aspect of CSCW that this section of this thesis is chiefly concerned with.

5.3.2. Shared Editors

Shared editors are tools that allow two or more users, situated at different locations, to work on the same document at the same time, as if they were in the same location. Typical scenarios for shared editors are engineers trying to produce or review design documents, distributed writers trying to author a document, or designers trying to sketch pictorial diagrams. Shared editors have therefore taken many forms; shared whiteboards, shared word processors and shared structured drawing tools have all been created [10, 80]. The defining tenet of these systems is that they allow users to interactively share the same information. As one user changes information, another can observe this change.

Shared editors usually feature a What You See Is What I See (WYSIWIS) interface or a variant on this, known as a relaxed WYSIWIS interface [181]. A WYSIWIS interface refers to an interface where all users of the system have the same view of the system at all times. As one user makes a change to an object in the system this change is reflected on the displays of all other users. In order to ensure that each user's view of the system remains the same, one constraint on WYSIWIS interfaces is that they either cannot feature scrollbars or alternatively they synchronise scrollbar positions among users. This synchronisation can lead to confusion among users and consequently many strict WYSIWIS systems simply do not feature a scrollable workspace, instead displaying only a single screen of information. This may make them inappropriate for some types of authoring – for instance, it is rare that any significant text document will fit on a single screen.

A relaxed WYSIWIS interface synchronises representations of the workspace as a whole, but allows collaborators to maintain different views on this workspace. This means that there is support for a workspace larger than a single screen of information

and that users can individually move to any part of this workspace. However, the information contained within the workspace is kept synchronised among the users.

Another distinction among shared editors is floor control policies. Floor control policies dictate who has access to the objects in a workspace at any given time and how this access is mediated. Numerous floor control policies exist [24, 170]. A simple policy is to nominate a single user to have access to the objects on the workspace. All other users can simply request that this user makes whatever changes they deem relevant. This approach is inefficient and inflexible – it forces a situation where only one user can make modifications to the shared data and it has been suggested that this single user could rapidly become a bottleneck [50]. Another floor control policy is based around the idea of a token. Only the user currently possessing the token can make changes to workspace. The token can either rotate systematically around the users, or alternatively the current holder of the token can pass it to whomever he or she sees fit. These types of policy are somewhat more flexible as they theoretically offer all users the opportunity to edit the shared document (although it is worth nothing that any specific user can suffer long, and possible arbitrary delays before realising this opportunity). These two floor control mechanisms are asynchronous – only one user can access the floor, the objects on the canvas, at any given time.

A synchronous floor control policy allows all users to access all objects on the canvas at any given time. This approach, while highly flexible and unconstrained, is accompanied by its share of new problems. Crucially, if two users attempt to modify the same object at the same time, differences between their local representations of that object can occur. To prevent this happening, it is common to lock objects that a user is editing [50]. Locking an object essentially means that it is made inaccessible to other users. This procedure ensures that all users are continually presented with the same shared data. However, the usage of object locks is not in itself simple. For instance, it is often not clear as to what granularity is it appropriate to lock objects. When considering a shared text editor is it best to lock data at the level of letters, words, sentences, or paragraphs? Furthermore, there is the issue of signifying locked objects to other users. Finally, there are implications for performance – the system must be capable of ensuring that two users do not access (and potentially alter) an

object simultaneously. There are several mechanisms for achieving this. Typical solutions involve brokering all requests at a central database or a server, or using a system of peer-to-peer messages [50]. Both these solutions can add lag to user actions, as they must first be validated on one or more remote machines. Lag can seriously influence user performance and subjective satisfaction [123].

There are advantages to synchronous floor control policies. Crucially synchronous editing does not constrain the activities of users of the system to any set pattern. The literature suggests that constraining user interactions to set patterns, which do not match their existing work patterns, is a major cause for the failure of groupware [82]. In a synchronous system, a group of users can choose to work on a single area of the workspace together or, in a relaxed WYSIWIS system, can move to completely separate areas to work in relative isolation. Synchronous systems support the development of work patterns to suit specific groups and tasks. Allowing synchronous activity also allows users to access the workspace as part of their communication – they can, for instance, write notes to one another, or draw illustrations to emphasise a point [17].

5.3.3. Problems with Shared Editors

Shared editors typically feature full-duplex audio and video links between collaborators. Despite the prevalence of these video-conferencing technologies, communication between collaborators still occurs in a deprived environment. A person is removed from communicating in the rich multi-sensory environment of the real world and required to work in a complex social setting through the primitive communicative medium of a window, or several windows, on a screen. A body of literature has attempted to classify, at a high level, the aspects of communication lacking in these distributed systems. This research defines the problems that the haptic communication described in the next chapter is designed to address, and consequently is reviewed here.

One of the most critical deprivations in collaborative environments is that of awareness [57, 155]. Gutwin *et al.* [84] define workspace awareness to include:

“...knowledge about who is in the workspace, where they are working, what they are doing and what they intend to do next.”

Similarly, in this thesis the term awareness is used to refer to the background, low-fidelity knowledge of the positions, actions and intentions of other people. In real world interactions, this information is gathered through casual glances at other workers, our peripheral vision, or through the sounds others make as they work. We accumulate awareness information from the world around us in a host of subtle and sophisticated ways and weave this rich tapestry of information into a background picture of what, and where, work is going on.

Coupled strongly to this concept of awareness are those of observed attention [104] and co-ordination [57]. Observed attention refers to the ability to know what another person is focusing on or referring to simply by observing his or her behaviour. This ability, characterised in the real world by the ability to see where someone is looking or pointing, makes talking about complex information simpler by providing a straightforward way of ensuring all participants are referring to the same object. Co-ordination is a product of awareness and is concerned with the higher level issues of providing a context for a group's activity as a whole. A co-ordinated group will not produce redundant work, nor will large amounts of time and effort be dedicated to the division of labour.

Information pertaining to gestures is also beneficial. Gestures in communication typically fall into one of two classes. Firstly, gestures to aid the flow of a conversation, for instance eye contact, and secondly bodily gestures, usually of the hands or arms, to illustrate, or reinforce, the information presented in a conversation. Eye contact is important in conversation not only because it aids token passing but also because it is the medium for the transmission of a large amount of important emotional content [126]. Tang & Minneman stress the importance of bodily gestures [189]. In observational studies of several group-drawing activities, they concluded that hand gestures are used regularly and productively in groups to:

“...act out sequences of events, refer to a locus of attention, or mediate their interaction....”

It is clear that gestural information of both kinds is important in communication.

Ishii *et al.* [103] provide a framework for explaining why these aspects of communication are difficult when mediated by a computer. They make a distinction

between the Shared Work Space (SWS) and the Inter-Personal Space (IPS). The SWS is the actual canvas that users are working on, be it a whiteboard on an office wall or a canvas in a multi-user editor. The IPS refers to the general communication that goes on between users in the same location, encompassing information pertaining to all the factors discussed above. Ishii *et al.* feel that while co-located users have easy access to both the SWS and IPS and can switch from observing one to the other quickly and easily, users mediating their work through a computer are either provided with no support for creating an IPS or are confronted with arbitrary *seams* between the SWS and IPS. A typical example of a seam is a video display being presented in a separate window or on a separate screen from the shared canvas. Ishii suggests that in a shared editor, users find that the seams between these two spaces make it difficult for them to switch from one to the other and reasons that removing these seams will facilitate collaboration by providing better access to the IPS. Ishii's solutions to the problems of a seamed workspace have typically been to use elaborate projection systems to overlay video of remote participants over the canvas that a user is working on [104].

Dourish & Bellotti [57] reinforce Ishii's ideas with a theory that culminates in similar recommendations. They discuss the awareness information provided by four shared editors, and conduct a study of one of these. The study involved ShrEdit [133], a textual editor, and a design task. Results were gathered through observation. After reviewing the systems, and their observations on the use of ShrEdit, they conclude that awareness information is vital to successful groupware systems but that there are several rules governing how and where it should be produced. Firstly, in accordance with Ishii, they suggest that it is vital that awareness information be presented within the shared workspace to enable users to easily choose when they wish to attend to it. They assert that awareness information contained outside of shared workspace will largely be of no use. Secondly, they state that awareness information must not be manually generated. The reasons for this are simple – users engaged in the process of generating awareness information are neither doing directly productive work nor liable to see any direct benefit from the work they are doing [82]. This makes it unlikely that they will be motivated to engage in explicitly generating this data. They term feedback that fits their mould *shared feedback*.

5.3.4. Solutions to these Problems

Many solutions to address these issues have been put forward. These are reviewed here as many generally applicable lessons can be learnt from the successes and failures of these techniques. Such information was critical in the design of the haptic communication presented in the next chapter. The solutions can be broadly classified as enhancements to the video channel, to the auditory channel, or simply as augmenting the semantic information displayed in the editor.

Video enhancements are by far the most common and are highly sophisticated. Many concentrate on allowing users to communicate using eye gaze, a facility not present in standard video links. Users can look directly at one another for token passing in conversation and in some systems can directly observe the focus of another user's attention [103]. Experimental studies have shown that patterns of dialogue use change (in terms of conversational token passing) when collaborators can maintain eye contact, but differ as to whether the effects of this change is positive [139] or negative [56].

Some systems have also enhanced video with views of the hands, to enable users to gesture and point as part of their interactions [104, 189]. These systems work by capturing video of the hands and projecting it over the remote user's workspace. Little more than observational evaluation – the statement that users found the gesturing capabilities useful and regularly engaged in it – of this method of communication has been provided. One obvious restriction of these gesturing video systems is that they are limited to a strict WYSIWIS architecture. If participants have different views of the same data then the gestures transmitted will be rendered meaningless.

Descriptions of audio enhancements are far rarer, the most influential being Gaver & *et al.*'s ARKola simulation [73]. In the ARKola simulation, two users had to maintain a set of machines in a virtual drink-bottling factory. They had to ensure that machines always had access to appropriate consumables, repair the machines when they broke and generally run a profitable drinks factory. Users were connected at all times by a full duplex speech link. The authors compared the behaviour of users of a purely visual version of this system with a version with added non-speech auditory feedback. This feedback indicated the status of the factory in an ambient way, much

like the running of a motor – all machines contributed sound to the simulation, continually changing their contribution according to their state. While there was no formal analysis of the system, observation of dialogues between the users led the authors to conclude that the presence of sound had positively influenced their interaction and facilitated a closer collaboration. They provide a compelling argument for this illustrated with some sampled sections of user dialogue. This evidence is reinforced in interview, where users stated that they found the sounds helpful in monitoring the state of the factory.

The RAVE system [72] used non-speech auditory cues to inform users about background system events, but presented little formal evaluation of this idea. Similarly, Beaudouin-Lafon & Karsenty [14] presented GroupDesign, a structured drawing system with the facility to associate an audio event with every graphical operation. In a different domain, Mynatt *et al.* described Audio Aura [141, 142], a lightweight audio augmented reality system. Users of this system continually wore wireless headphones and smart badges as they went through their day of work. The badges allowed a system of IR receivers to maintain the position of each user, and the headphones could receive sounds from an event server. As users moved around their workplace they were played a combination of abstract and explicit sounds relevant to their location. For instance, in the coffee room, sounds relating to the arrival of email were played. This system provided awareness information in a similar spatial context to that found in a shared editor.

Cohen describes two audio group awareness systems, ShareMon [45], and Out to Lunch [46]. ShareMon was a utility that allowed a user to monitor background file-sharing events by providing feedback when users logged on, logged off, and also at the beginning and end of an individual file share. ShareMon had three interfaces that were anecdotally compared to one another. It could present information graphically, as sampled non-speech sounds, and in Text To Speech (TTS). The conclusions of this study were that most users appeared to prefer the sound interface to the TTS and graphical ones, and that they found the sounds informative. On the other hand, some users referred to the sounds as intrusive and annoying. Out to Lunch attempted to represent the sounds of group activity, such as that to be found in a physically shared office, to a team of workers distributed around a building. The final system mapped

the mouse and keyboard activity of the group as a whole to a repeating chord and periodically played signature tunes for each user contributing to this activity. Although no formal evaluation of this system was presented, the authors speculate that people could use this information to gauge the level of general group activity, and also to discover who specifically was working.

It is worth noting that these audio enhancements fit natively into the idea of a seamless workspace – their presentation is easy to attend to as background information, and attending to it does not interfere with the main presentation of content. Furthermore, in accordance with Dourish and Bellotti’s framework, none of the feedback described is manually generated.

Enhancements of the semantic information present in a groupware interface typically involves including widgets that attempt to display visualisations pertaining to awareness, attention, co-ordination and gesturing. The simplest and most ubiquitous of these is the telepointer [80]. A telepointer is a local representation of a remote user’s cursor. They provide some awareness and attention information simply because they allow a user to locate other users and, possibly, to infer what they are doing. Telepointers also support basic gesturing; by moving your cursor rhythmically over an area of the screen you can hope to attract another user’s attention to some salient feature.


Numerous other widgets that can augment communication have been postulated. Baecker *et al* [11] in their description of Sasse, a textual shared editor, mention the inclusion of a vertical shared scrollbar. A shared scrollbar is an extra, non-interactive scrollbar that resides adjacent to the regular scrollbar and reflects the position of the viewports of the other users. They have little to say about users’ perception or use of this widget, save for the assertion that most users had little difficulty with it.

Gutwin *et al.* [84] present an evaluation of five awareness widgets in a groupware system considering only pairs of users. They examined miniature views, radar views, What You See Is What I Do (WYSIWID) views, a teleport facility, and finally multi-user scrollbars. The miniature view consisted of a small view of the entire workspace, complete with all artefacts on the workspace. A user gained information about the actions and location of the other user by observing these artefacts. The

radar view was an extension of the miniature view, augmenting it with more precise information about the location of the other user. The radar view contained small telepointers and shading, which indicated the extent of the other user's view on the workspace. The WYSIWID view provided a small view centred on the other user's cursor. The actions of the other user could be explicitly observed through this view, and objects seen in this view provided cues as to his or her location. The teleport facility allowed a user to depress the third button of his or her mouse and be temporarily transported to the location of the other user. Releasing this button would return the user to the original position. The shared scrollbar was similar to that in Baecker's system but consisted of both a horizontal and vertical scrollbar. Figure 5.1 is taken from Gutwin *et al's* paper and shows the widgets discussed. Gutwin *et al's* task involved formatting existing textual and graphical information into a newspaper style. Each pair of users participated in two sessions, each augmented with a single awareness widget. A control condition was also included. Results were collected in the form of observations, questionnaires, and an interview.

Gutwin's analysis indicated that users found the miniature and radar views most useful, but acknowledged that this was in part because they aided the formatting task as well as providing awareness information. Users were confused by the WYSIWID

Figure 5.1. Illustration of different awareness widgets (reprinted without permission from Gutwin *et al.* [84]).



view, partly due to a slow refresh rate and partly due to the fact that if they did see something relevant on the view they were often at a loss as to what to do about it – there was no way to directly influence the scene displayed on the WYSIWID view. Similarly, users found the two-dimensional shared scrollbar confusing, perhaps indicating that it is more suitable to a one-dimensional situation, as in Baecker's text editor. Finally, users tended to forget about the teleport facility, possibly because it had no graphical presence.

In contrast to the audio and the video enhancements of groupware, all of the widget augmentations (except the telepointer) form seams in the environment, according to Ishii's definitions and are presented outside the workspace, according to Dourish & Bellotti's work. To pay attention to the information in the awareness widgets, a user must explicitly direct attention away from the shared workspace and towards the widget. However they typically do comply with Dourish & Bellotti's second guideline, and do not rely on manually generated data.

5.3.5. Evaluation in CSCW

Evaluating any user interface is challenging and there is far from a single accepted procedure for doing this [131]. This situation is exacerbated in collaborative systems, where it is often not simply the interface that is being evaluated but instead the effectiveness of communication among users of the system. Due to this complexity, and the reliance of this thesis on evaluation, this next section reviews the literature relating to the evaluation of collaborative systems. Broadly, two basic approaches to evaluation exist in the literature: experimental science and ethnography.

Experimental science is typically based around laboratory experiments where participants sit through fixed tasks in a limited time frame. Typically, the experimenter manipulates a small number of variables, often only one, between two or more conditions. Great care is taken to ensure that the experiment is subject to no other variation. In this way, it is hoped, cause and effect can be established between changes in these variables and changes in systematically measured user performance and subjective measures [75].

Ethnography exists at the opposite pole. Ethnographic studies involve trained observers recording the activities of a group of users of a system over a protracted

length of time (often of several months) and in a real situation. In-depth interviews of participants are also usually conducted. Ethnographers analyse and study this rich wealth of data. The goal of ethnography is to draw conclusions about the use of a system in a genuine context, one subject to all the complexities of real life [86].

There is a great deal of debate between proponents of these two methodologies [138]. Experimental scientists criticise ethnography for being unrepeatable, potentially subject to the biases of the observers, and open to sampling problems. In return ethnographers criticise experimental science for not being ecologically valid, for destroying context rather than trying to understand it, and, despite all attempts to the contrary, being subject to uncontrolled, confounding influences.

In practically assessing groupware the choice between these two methodologies is often made by the context in which the research takes place. Experimental research is prevalent in systems where some new feature or channel of communication has been added to either an existing or relatively standard piece of groupware [73, 84, 103]. This style of research lends itself to experimental evaluation because it presents a simple pair of conditions – users working with the new augmentations and users working without. Evaluation in this way can also produce rapid results. Furthermore, ethnographic studies of such systems are difficult to conduct, as placing what is essentially a prototype system into a place of work is both undesirable, and unlikely. On the other hand ethnographic studies suit evaluations of more established groupware systems, often even suites of tools [152, 184, 191], for much the same reasons. Established and complete systems are more likely to be accepted into a workplace, and the longer-term nature of ethnographic evaluation is more likely to bear fruit.

In conclusion, both experimental science and ethnography have their place. Both feature regularly in the literature and they have very different features and characteristics. Each is suited to different styles and stages of research. It is important to note that every method of evaluation has its weaknesses and that reviews of evaluation techniques suggest that the full picture can only ever be revealed by the application of a variety of different methodologies [132]. As discussed in detail in section 5.2.7., the topic to be examined here – the role of haptic cues in structured communication or collaboration – is a relatively novel one, and

given its preliminary character, it seems appropriate to limit investigations to those falling under the scope of experimental science. Given this conclusion, a discussion of the evaluation techniques used in experimental studies of collaborative systems is presented below.

5.3.6. Techniques for Experimental Evaluation

Experimental science faces considerable challenges in evaluating groupware. As this approach is used extensively within this thesis, a brief review of the various empirical techniques that have been employed in the past is presented. This review informed the measures chosen for use in the studies reported in the next chapter.

Traditional objective methods for evaluating usability, such as task performance time, quality of output and efficiency [20] often show little variation in vastly different communicative environments, reflecting the resilience and versatility of human communication. Doherty-Sneddon *et al.* [56] make a distinction between problem solving tasks and design tasks, and suggest that significant differences in objective measures may be more likely to appear in studies of design tasks, possibly because design tasks require more communication and discussion as they possess no single correct solution.

Observation is a common method for evaluating groupware [73]. It basically entails watching and recording users interacting with a system and attempting to reason why they have performed certain actions, or what effect certain variations in the system are having on their behaviour. Dialogue analysis is a formal variation on observational evaluation and is commonly used in the more psychological accounts of groupware [56]. Dialogue analysis involves recording all utterances made by users, and then classifying these according to one of several theoretical models of dialogue structure. This classification is complex and to provide reliability the dialogue is usually completely analysed by two or more expert raters. The results of these separate analyses are then compared statistically. If they strongly correlate, then inter-rater reliability is achieved and the analysis is said to be valid. Once created, these formal accounts of dialogue can be reasoned upon. For instance in eye gaze experiments, dialogue analysis makes it simple to count the number of dialogues used to perform conversational token passing. Dialogue analysis is often used simply because it is more sensitive to changes in behaviour than other objective

measures. Unfortunately, dialogue analysis relies on the skill and time of expert raters and as such is often an impossible or uneconomical solution to the problem of evaluation.

Beyond these objective measures, several subjective experimental techniques have been applied to groupware. Questionnaires are a common way to evaluate groupware [16, 58] as they offer a repeatable mechanism for gaining subjective data. However, in groupware there appears to be something of a divergence as to how questionnaires are created and used, when compared to the larger discipline of usability as a whole. In the usability literature it is acknowledged that questionnaire design is far from simple, and that for a questionnaire to produce truly reliable results it must first undergo an extensive validation procedure [43]. Typically, questionnaires are designed to address issues, or gain opinions, on a number of different factors within a topic and each factor is represented in the questionnaire by a number of questions. The validation of a questionnaire involves its administration in a variety of settings and to a large number of users. This data is then taken and subjected to lengthy statistical analysis to attempt to find trends among the rating of items which correspond to the original factors. Validity is this demonstration of correspondence. Also the design of most questionnaires includes dummy or distracter items that serve the purpose of error checking by attempting to isolate users who are filling in the questionnaire inappropriately. For instance, it is common to check for participants filling in a questionnaire with complimentary rather than genuine responses. In contrast to this structured approach, questionnaire design in groupware evaluation appears to be *ad hoc*. Questionnaires are created for individual experiments and little detail is provided as to their contents. This has led to a situation in which, although questionnaires are regularly used for the evaluation of groupware, there are no established instruments for this purpose

Finally, interviews are often conducted to evaluate CSCW systems [11]. Interviews can be structured, with experimenters asking a set of fixed questions, or more typically unstructured, with the experimenter trying to draw out information about behaviours that had been observed while the users were working with the system. Interviews allow experimenters to target specific information in great detail but can be subject to significant biases, which can influence the quality of the results

gathered. For instance, confirmation bias [75] refers to the process by which people will seek to confirm a hypothesis that they already hold. In interviews, this manifests itself with what are essentially leading questions.

In conclusion, a variety of different evaluation techniques have been applied to experimental studies of groupware, and each technique has its advantages and disadvantages. Furthermore it seems likely that different techniques will be effective in different situations. The design of the evaluations of collaborative systems described within this thesis bears these issues in mind, and draws on the review presented here to direct the selection of the specific measures chosen in each case.

5.3.7. Haptics in Communication

Haptic communication is a relatively unexplored area of research. This section considers the potential of the sense of touch for communication from several different perspectives before reviewing existing systems that support haptic communication. From a cultural standpoint, there is a common perception that touch is a deeply personal and emotional sense. An illustration of this assertion is the integration of our emotional reactions to the world and tactile terminology. We ‘feel’, and have experiences that are described using terms related to the sense of touch, such as ‘tender’, ‘rough’ and ‘touching’. The sense of touch is also one of the earliest avenues for communication in infants, and many researchers [70, 87] have linked effective and supportive tactile communication to healthy emotional and social development.

The psychological literature concerning social tactile communication can be divided broadly into two camps, one predating the other. The earlier portion of the literature focuses on the role of touch as a means to express intimacy [127], while the later, mainly stemming from an influential 1973 study by Henley [91], is concerned with touch as a mediator of power and status. Henley asserts that touch is a status privilege, and that high status individuals are more likely to initiate touch with low status individuals, than vice versa. Henley also recorded that men touch women more than women touch men and claimed that this reflects the relative status of each gender in society. Much of the subsequent literature has remained focused on this assertion [66, 85, 125], and tried to prove, disprove or merely highlight gender and status related differences in tactile communication. It is also worth noting that

substantial cultural differences seem likely to affect haptic communication. The conclusions that can be drawn from this literature are that communication through touch is often loaded with emotional content, and can be an important channel for social interactions.

Beyond these issues, a number of touch-mediated languages have been created, primarily to support communication with deaf-blind individuals. These languages demonstrate that, regardless of the social implications, tactile communication can be both rich and expressive. Fingerspelling is a tactile language based around hand contact. The pressure and movements of one communicator's hand are sensed and interpreted by the hand of other communicator. Tadoma [186] refers to a communication method that allows the direct perception of speech through touch. To sense the speech, an individual places his or her thumb on the speaker's lips and fingers on the speaker's throat. Much of the expressiveness of speech can be maintained in Tadoma.

There is little literature specifically concerning Computer Mediated Communication (CMC) and touch. However, general observations on the differences between communication in the real world and that mediated by a computer may well generalise to interactions through touch. While the psychological literature suggests that communication through touch is extremely personal, and therefore potentially an unsuitable modality for work orientated communication, some of the CMC literature indicates that people feel less bound by such social conventions during computer mediated interactions. Sproull and Kiesler [177] state:

“People who interact via computer are isolated from social rules and feel less subject to criticism and control. This sense of privacy makes them feel less inhibited in their relations with others.”

A powerful example of this kind of behaviour, where normally strict social rules are flouted, can be observed in text based CMC systems, such as Internet Relay Chat (IRC), or Multi-User Dungeons (MUD). It is commonplace for users of such systems to adopt false identities, often of the opposite gender [127]. Another example comes from network conferencing. Dubrovsky *et al.* [59] conducted an experiment indicating that in network conferences, as opposed to real world conferences, status

played a reduced role in determining which participants expressed their views. The mediation of communication through a computer had reduced the influence of these social cues.

Computing science has also ventured into the realm of tactile communication. Systems that support communication through touch have been designed, and typically the social elements of such systems are the driving force behind the research. Brave & Dahley [25] state:

“Touch is a fundamental aspect of interpersonal communication. Whether a greeting handshake, an encouraging pat on the back, or a comforting hug, physical contact is a basic means through which people achieve a sense of connection, indicate intention, and express emotion.”

The majority of computer science research dealing with haptic communication has reflected this statement and focused on connecting people and supporting interpersonal relationships.

Perhaps the first communicative haptic environment was Telephonic Arm Wrestling [200], an art exhibit consisting of a pair of spatially separated robot arms allowing two remote users to arm wrestle with one another. The force that one user would exert on one robot arm would be applied to the other and *vice versa*. Several devices have been developed on a similar theme. The shaker in Feather, Scent and Shaker [183] allowed users to shake a device in their hand and have this represented as vibration in another user’s coupled device. The shaker contained an electromagnet and movement of one user’s device induced a current in its coil that was then transmitted to the other user’s device to produce vibration. The authors suggested their device would encourage: *“...light-hearted play amongst friends...”*

The Bed [55] describes an attempt to create a distributed bed to connect remote partners using a variety of abstract non-explicit feedback. For instance, the breathing of the remote partner was sampled, analysed, and then locally represented as the gentle vibration of a pillow. inTouch, [25, 26] (illustrated in Figure 5.2) was a pair of coupled devices each consisting of three rollers. Rotating a roller in one device caused a similar movement in the other device, allowing users to push against one another, and to play. inTouch provided relatively rich feedback, as each roller could

Figure 5.2. The inTouch haptic communication device (reprinted without permission from www.media.mit.edu/tangible).



be manipulated, either clockwise or anticlockwise, independently of the others. Each of these systems is concerned with the use of haptic feedback to support interpersonal communication, but provides little evidence to justify this interest as they are characterised by a lack of reported evaluation.

However, the most detailed and informative account relating to a device in this area is reported by Fogg *et al.* [67]. They describe HandJive, another pair of coupled devices, in this case created as a toy that supports people's desire to fidget when listening to group presentations such as lectures. The most useful aspect of this research, however, is the fact that it exposes the iterative design process that led to the final construction of the devices. This design process involved the production of a number of prototypes, each subject to some user testing and observational analysis. The results from these intermediary steps, and a final, more substantial, observational study led the researchers to conclude that users found communication through touch compelling and enjoyable, but, if possible, tended to use it to compete with one another. These concerns were reflected in the design of the final device, which consisted of a pair of cylinders, jointed together at the centre. Each cylinder could rotate around this joint to lock into one of five discrete positions (including straight). A change in position of the device was reflected in the other coupled device. The construction of HandJive is illustrated in Figure 5.3. To avoid competition over the position of the cylinders each of a pair of users could control the motion only on a single, and opposite, axis. Only collaboratively were they able to move the device through its full range of motion. The researchers suggested that two users could co-operatively construct "dances", transmit rudimentary messages, or play simple games using the device.

Figure 5.3. Illustration of handJive haptic communication device (reprinted without permission from Fogg *et al.* [67]).

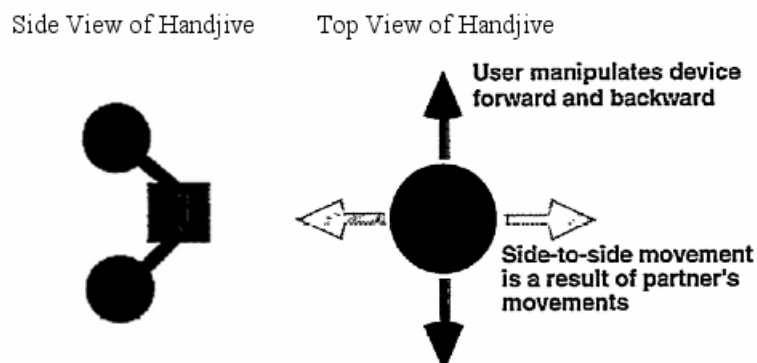
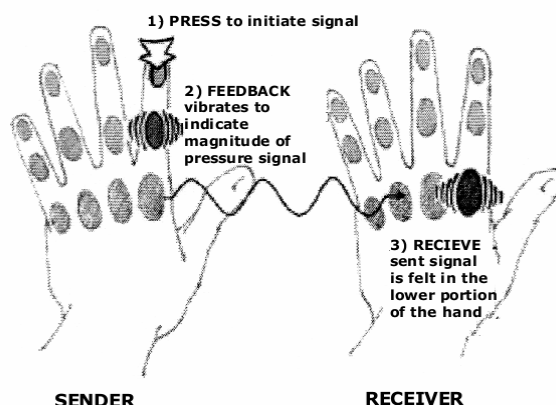


Figure 5.4. The ComTouch haptic communication device (reprinted without permission from Chang *et al.* [40]).




Chang *et al.* [40] also describe the design of a haptic device to support communication between two individuals. Each of their devices consisted of ten vibro-tactile actuators and five pressure sensitive pads. The device was designed to be handheld (the driving concept was to integrate this technology into a mobile telephone) and arranged such that two actuators and one pad were situated under each digit of a user's hand. As a user pushed against a pressure pad, a vibration with a magnitude dependant on the magnitude of the pressure exerted was displayed both directly on the user's device (on one of the actuators under the digit exerting the pressure), and also on a coupled device (on the other actuator under the same digit). This is illustrated in Figure 5.4 for a single digit. The goal of this research was to investigate how people used this tactile communication as an accompaniment to an audio conversation. Several studies were conducted. In situations when users were discouraged from speaking, the authors observed that users could develop simple tactile languages (and indeed, they compared these to Morse code). When users

could speak freely they tended to use the touch communication to emphasise points, coordinate turn-taking and express their presence by playfully mimicking one another's tactile communications.

It is possible that haptics can have more impact than simply acting as a conduit for interpersonal communication. Durlach & Slater [61] speculate that the sense of touch may be vital to the sense of presence that users perceive in Collaborative Virtual Environments (CVEs). They reason that the ability to feel objects or other users would enhance feelings of interaction and direct manipulation which have been linked with an increased sense of presence. They also refer to touch not being a “distance sense” – if we feel something, then it must be close to us, making a simulation more compelling. Finally, they suggest that users are unused to receiving illusions of touch and are continually bombarded with artificial visual and auditory stimuli, and therefore haptic simulations are more likely to draw users in and increase their subjective experiences of presence. This last effect would obviously hold only while haptic simulations are a rarity.

In line with this reasoning, there have been several attempts to add haptic feedback to CVEs. In a companion paper to the one described above, Basdogan *et al.* [13] discuss the effects of haptic feedback on the cooperative performance of physical task in a virtual environment. They demonstrate that both objective performance, and

Figure 5.5. Illustration of interface for haptic communication (reprinted without permission from Basdogan *et al.* [13]).



a loosely defined concept termed the “sense of togetherness” achieved among participants can be increased through the addition of haptic feedback to a virtual simulation of co-operatively steering a ring along a wire. In their task, pairs of users would move a ring along a wire by positioning themselves at opposite sides of the ring, pushing towards each other, and then moving laterally. Essentially, they would act as if one were the thumb and one the forefinger of a hand performing a gripping action. This behaviour is pictured in Figure 5.5. While the results from this experiment were statistically significant, they were over a small sample of users and those relating to the “sense of togetherness” were based on an unvalidated questionnaire. The authors admit that this work is non-conclusive and ongoing.

Sallnas *et al.* [166] report a study with a similar theme. Their task also involved the cooperative manipulation of virtual objects. Specifically pairs of participants used a cooperative gripping action (analogous to that in Ho *et al.*'s study) to position groups of randomly dispersed blocks into specific three-dimensional configurations. Sallnas *et al.* [166] were interested in the abilities of the haptic feedback to affect both objective performance and users' perceived virtual and social presence – the extent to which they felt immersed in the virtual environment and the amount they felt socially present with the remote user. They showed significant improvements in objective measures, and questionnaire measures showed that the addition of haptic feedback significantly increased virtual, but not social, presence.

Noma & Miyasato [145] also describe a virtual environment object manipulation system. Their system includes the facility for multiple haptic devices to simultaneously affect the orientation and position of an object. They describe a pilot study consisting of three subjects in which each subject controlled two devices and attempted to relocate and reorient an object in each of three conditions. In the first condition, users held an object that corresponded to the virtual object, in the second they experienced haptic feedback designed to resolve discrepancies between the positions of their hands and in the third had access to only graphical information. The results showed that task times were lowest in the condition incorporating the real object. The task time was also significantly lower with the haptic feedback than without. Noma & Miyasato believe that the results can be generalised to the case where different users have hold of each haptic device. This assumption is perhaps

questionable and the three subjects in their study do not form a particularly convincing subject pool. A further caveat regarding the objective measures (such as task completion time) used in all three of these CVE based studies is that the demonstration that haptic feedback increases performance in physical tasks is not particularly surprising. Indeed, it is perhaps more surprising that users managed to achieve the manipulations without the haptic feedback at all.

Finally on this topic, Alhalabi & Horiguchi [5] describe a system that enables two users of a CVE to haptically shake hands with one another. In their system each user is represented by an arm shaped avatar. When the hands of these avatars collide, the movements made by one user are haptically presented to the other user, enabling the remote equivalent of a handshake. Although they do not perform a comparative study, they evaluate the quality of this feedback using a simple custom built questionnaire. They conclude that users of this system find the feedback both intuitive and satisfying. A crucial unifying point about the CVE based studies described here is that they share a concern over the lack of support for social interaction in collaborative virtual environments and view touch as a valid mechanism for increasing it.

5.3.8. Summary of Haptics in Communication

To briefly summarise, the psychological literature suggests that touch may be a modality that is too personal to use for the communication of anything but the most intimate of messages. However, a mounting body of CMC and computer science research takes an opposing stance. Firstly, it is far from clear that computer mediated communication through touch shares the same highly emotionally loaded connotations apparent in real world tactile communication. It may be the case that virtual communication through touch is acceptable in situations where it would be inappropriate in the real world. Furthermore, while some computer science research has looked at intimate communication, as highlighted by the psychological literature, more has focused on communication for other purposes, such as entertainment, supporting social interactions and bonding, or to perform co-operative tasks in virtual environments. The perspective of this body of work is that haptic communication has the potential to improve the social aspects of CMC, and that such an improvement is desirable. Viewing this diverse literature as a whole, this thesis

adopts the position that the current understanding of the social or interpersonal aspects of communication through touch supports the investigation of its use as a channel for communication among the users of collaborative systems. Overall, the literature suggests that the positive effects of mediated haptic communication - such as increased levels of engagement and presence and a heightened sense of togetherness - seem likely to outweigh any potentially undesirable side-effects relating to the intimate and personal nature of real world communication through touch.

5.4. Conclusions

The literature relating to communication through touch is reviewed in this chapter. It began with a general introduction to the field of CSCW, but rapidly focused in on one particular area within that field: synchronous collaboration. Specifically, it has been concerned with shared editors. These are tools that support the simultaneous authoring of documents by groups of physically separated users. A brief description of the various classifications of shared editors is followed by a more in-depth discussion of the communication problems that have been observed in these systems. Solutions that have been developed to address these issues are then discussed at some length. The role played by this section of the literature review in this thesis is to precisely define a problem area for the design of the feedback presented in the next chapter.

This chapter also examines evaluation methodologies in CSCW research. Broadly, there are two forms of evaluation: experimental, and typically lab-based, investigations, and ethnographic studies of systems in real environments. While both have their merits, the preliminary nature of the topic to be examined in this thesis – the integration of a novel modality for communication into an existing application domain – led to the conclusion that experimental evaluations would be more appropriate in this instance. Following this decision is a review of the different techniques that have been applied to evaluate collaborative systems. This section of the literature review serves to provide the rationale for the basic nature of the evaluations presented in the next chapter, and informs the choice of the specific measures chosen in each case.

The final topic examined in this chapter is communication through touch. This review begins with the psychological literature, then moves on to discuss the various systems that support haptic communication. These are examined in some depth, and despite the relative small quantity of concrete evaluation found, draws the conclusion that there is a strong argument supporting the investigation of touch as a modality for communication.

Overall, the role of this chapter within this thesis is to provide a context and a background within which to consider the haptic communication designed, described and evaluated in the next chapter. It achieves this by precisely defining the application area to be studied, considering the possibilities for evaluation and discussing the current research on touch in communication.

6. Design, Implementation and Evaluation of Haptic Communication

6.1. Introduction

This chapter describes in detail, and goes some way towards evaluating, one possible use of haptic feedback in shared editors: haptically enhanced telepointers. Given the context provided by the literature review presented in the previous chapter, this section starts by describing the high level design of the haptic communication. Following this are details of how it was implemented. Two evaluations are then presented. The first of these is a general evaluation of pairs of users performing a complex design task either with or without the haptic communication. It was intended to provide a broad range of information about how the communication was used – how user behaviour altered with the presence of the haptic communication – and to gather data on participants' perceptions of it. The second study was more focused. It examined a single aspect of the haptic communication, and explored a hypothesis that was isolated (through observation) in the first study. The chapter concludes with a set of design guidelines generated from the two studies presented in this thesis, and the previous literature on this topic. Reflecting the infancy of research in this area, the goal of these guidelines is not to precisely describe how to create effective haptic communication. Instead they seek to present the design issues relevant to this topic, and to map out the motivations that are driving research in this area. These guidelines are the first of their kind on this topic.

The central claim of this thesis is that the addition of haptic cues to cursor interactions can yield performance benefits, or improvements in subjective experience. This general statement is to be demonstrated through the thorough investigation of two specific examples. This chapter forms an empirical examination of the second example chosen. It details investigations into the influence of haptic feedback relating to a user's representation or avatar in the domain of collaborative systems. These evaluations, and therefore this chapter's contribution, are critical to this thesis. A secondary goal of this thesis is to further design knowledge on the topics examined; to ensure that future researchers or system developers can directly benefit from the work described here. The design guidelines presented in this chapter serve this purpose. They are intended to concisely summarise the available literature on haptic communication and collaboration in a form that allows easy consumption by designers, researchers and other interested parties.

6.2. Design of Haptic Communication

Given the discussion (presented in the previous chapter) of the problems users experience with shared editors – such as awareness, attention, co-ordination and gesturing – and the suggestion that haptic feedback may be able to at least partially address these issues, here the design of a novel mechanism for computer mediated haptic communication is presented. It entails enabling haptic cursor interactions between collaborators in synchronous shared editors. In this design telepointers are transformed from being a simple graphical representation of position to physical avatars, tangible bodies, in the virtual space that can haptically influence one another. A total of five specific haptic interactions between these avatars have been designed.

6.2.1. Push

Firstly, the telepointers can push one another around the workspace. As one cursor encroaches on another, both can feel a force pushing them apart, or if one cursor intersects another at speed then the other cursor will be pushed away. This effect could be used as a warning, for instance if a user was about to perform some disastrous action another user might attempt to push him or her aside in order to prevent this. Another potential use would be to catch another user's attention, the remote equivalent of a tap to the shoulder. This interaction is reminiscent of others in

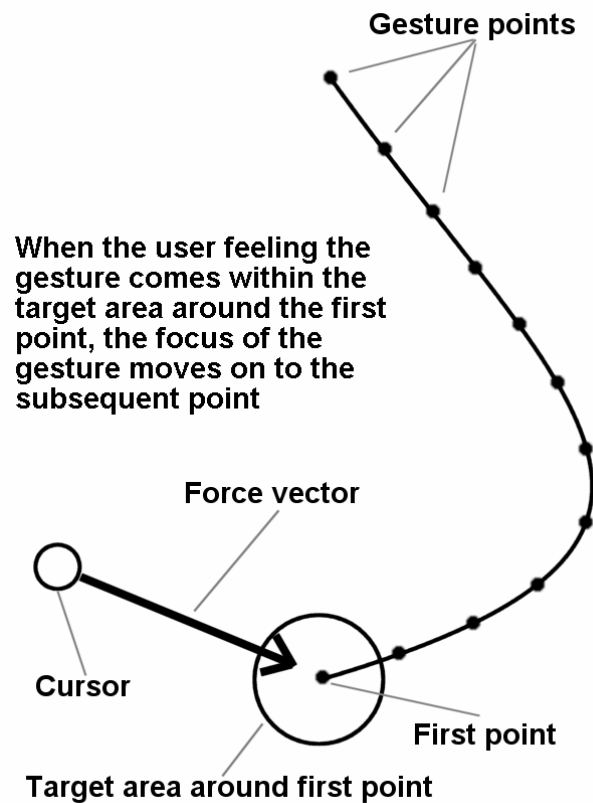
the literature – for instance both the arm-wrestling simulation [200] and inTouch [25] are basically mechanisms that allow physically separated users to push against one another. In this instance, however, the pushing simulation is much more complex, as it is embedded within the context of a spatial workspace – to push a user you must first locate that user, and as you push them he or she can retreat away from you. Currently, the push effect is implemented with each cursor represented as a frictionless sphere. A consequence of this is that it is difficult for cursors to push each other uniformly; they tend to slip and slide off each other. A more complex haptic simulation, including friction or possibly even an attractive force between cursors engaged in a push interaction, might prove more effective at providing a consistent sensation of pushing or being pushed.

6.2.2. Gesture

The second haptic communication designed extends the technique of gesturing with telepointers by allowing a telepointer to haptically take hold of another by moving over it and depressing a controller button. Once held, subsequent movements are played back haptically to the other cursor until the button is released. This operation has the effect of grabbing a pointer and then making it follow your path. While this is far from directly analogous to how gestures are created and perceived by co-located individuals, it does considerably extend and make concrete the basic gesturing function of telepointers. One can firmly and interactively transmit a complex spatial pattern to a remote user, without words.

There were some problems in implementing this gesture effect. The basic algorithm involved storing key points along the path of the user creating the gesture, based upon the distance from the current point to the previous key point. This distance was small (5 mm), to maintain the fine detail of the gesture. When the gesture begins, an attractive force towards the first point in the gesture is applied to the user receiving the gesture. The magnitude of this force increases with the range from the user to the point. When the user comes within a certain target range of the point, the focus of the gesture moves on to the subsequent key point. Again, to maintain the fine detail of the gesture this target range was kept small: 1 cm. This procedure iterates for all the points in the gesture. This is summed up in Figure 6.1.

Figure 6.1. Visualisation of haptic gesturing algorithm.



However, it was observed that while using this algorithm, users experienced difficulties – they became lost and unable to follow the gesture. This was attributed to the fact that the forces of attraction used are relatively weak and become weaker as a user approaches a target area, making it difficult to locate these areas. There were several possible solutions to this problem. As larger forces had been mapped to greater distances, simply increasing the magnitude of the forces applied when users became close to a point was not an appropriate solution. Nor was increasing the size of the range at which a user is said to have reached a point, as doing this would reduce the fidelity (the resolution at which the gesture was rendered) of the gesture, and small perturbations on the path would not be perceived. Furthermore, it was felt that it would be easier for users to detect changes in the direction of a force rather than just in its magnitude.

To achieve these goals the gestures were smoothed. As time went by without the user reaching the currently active key point in the gesture, the target area around that point would expand. Eventually it would encompass the user, at which stage the simulation would turn its attention to the subsequent key point in the gesture, with a

small active range once more. Moving the simulation along the path of the gesture even while the user remains stationary means that the magnitude and direction of the force applied to the user will continually change. A further consequence of this is that if a person ignores the forces from a gesture then eventually all they will feel is a force to the last point of the gesture – the details would have been smoothed away. This algorithm has the benefits of initially presenting the user with an accurate representation of the gesture and then gradually reducing its resolution. In this reduction of resolution, it also ensures that a user is presented with vectors of varying magnitude and direction while remaining on the gesture's path. The algorithm also only reduces resolution as it needs to – if a person begins to follow the gesture closely again, after losing it for a short time, the resolution will increase once more. A temporal aspect to the gesture is also added. If one ignores the gesture for long, it will slowly lose detail and eventually vanish.

Finally, this gesture effect was further enhanced to factor in the speed of the user making the gesture. The force applied to the user receiving the gesture varies according to the speed at which the person recording the gesture is moving, above a certain minimum. This allows users to highlight or emphasise certain parts of a gesture by varying their speed.

6.2.3. Proximity

The third interaction between the telepointers was designed to provide some simple awareness information (according to Gutwin *et al.*'s [84] definition of workspace awareness). The viscosity or resistance to movement present in the workspace was increased when another user drew near a user's position. Alternatively, if a user was stationary when another approached, a small vibration was applied. This provided a haptic proximity sense analogous to the physical sensation of presence perceived when close to another. While the information content of this effect is low, for instance it will not help a user determine who is approaching, nor from what direction they hail, it is hoped to have the advantage of being obvious while remaining unintrusive.

6.2.4. Locate and Grab

The remaining two interactions are focused towards the awareness and co-ordination problem of being unable to locate other users in the workspace. The work described earlier in this thesis has shown that haptics can provide benefits in targeting tasks. Finding other users in a canvas is fundamentally a targeting task. In accordance with this, a locate effect was implemented which allowed users to activate a homing force on their cursor which would tug them towards another user. The force profile of this effect is similar to that of the gravity well haptic widget augmentation that proved to be effective for targeting tasks in the buttons experiment described in Chapter 4. In this instance, however, another user is the focus of the targeting force and it is applied at two distinct levels. Initially a small force is applied, which allows a user to determine in what direction another user is located. After a brief time, this force is increased to actually move the user towards the other user's position. This effect is somewhat similar to the functionality of the teleport interaction described by Gutwin *et al* [84]. The locate effect differs from Gutwin's teleport in that it does not instantaneously transport a user anywhere, nor does it provide a method of returning to the original position. While both the initial tug to indicate direction and the fact that a user is physically pulled along may prevent the disorientation observed in users of Gutwin *et al.*'s teleport effect, it might be beneficial to include functionality that allows users to return to their original position after locating another user. This locate effect also suffers the disadvantage that, unlike the push and gesture effects, there is no inherent mechanism for specifying which other user you wish to interact with. Some external dialogue would have to be invoked in systems containing more than two users.

The final interaction is an inverse version of the locate effect. This grab interaction allows users to turn on a homing force that pulls all other users in the workspace towards their position. This allows a user to request other users to come to some location in the document without being burdened by having to describe that location. The locate and grab effects were designed to facilitate easier navigation and co-ordination between users in the workspace.

6.2.5. Discussion of Haptic Cursor Communication

An advantage to these haptic interactions is that they comply with Ishii's recommendations for a seamless workspace [102] and also with Dourish & Bellotti's [57] principles of shared feedback – that feedback should be presented in same space as content and that it should not be manually generated. A disadvantage of these interactions is that, save for the proximity and push effects, all are required to be explicitly invoked by a user and all share a common invisibility, much like the teleport effect described by Gutwin *et al.* [84]. Gutwin's teleport was an awareness widget with no visual representation and participants used it infrequently. In interviews, it was revealed that they had simply forgotten about it. Similarly, it is possible then that users will not use the haptic effects for the same reason, although the act of manipulating a haptic device may serve to remind them.

A final consideration in the design of this haptic communication was how intrusive it could be. A user working on a diagram, for instance, would probably not appreciate the application of arbitrary forces by other users. The push, gesture, and grab interactions allow a user to haptically influence another user with intrusive forces and the grab interaction in particular does this without any associated visual feedback. Modes are a potential solution to this problem. Three modes are suggested – working, communication and observation. In the working mode a user can interact with the canvas and can create content, but cannot be haptically influenced by another user. In the communication mode, users cannot interact with the canvas but have access to the haptic communication. In the observation mode, users can neither communicate haptically nor access the canvas. In situations involving a two-dimensional canvas and three-dimensional haptic device these three modes could be mapped to the z-axis of the device. Users would adjust their position on the canvas using movements on the x and y axes, and use their position on the z axis to mediate their mode. Closest to the canvas would be the working mode, beyond that the communication mode and, furthest away, the observation mode. This mapping supports the physical metaphor of the canvas. You must be on the canvas to work, near the canvas to interact with other workers and when far from the canvas, you can simply watch. Haptic barriers could be placed between the spaces representing these modes to ensure that users did not accidentally slip from one to another, while

altering the graphical cursor representing the user could serve to indicate which mode a user was currently operating in.

6.3. Generic Implementation Details

6.3.1. Haptic Simulation

The haptic simulation was implemented in C++ under Windows NT and for the PHANToM (from SensAble Technologies) using the propriety GHOST API. The PHANToM's used in these studies were three DOF force feedback devices. One was a model 1.0, and the other a 1.5. Participants manipulated them by holding a pen-like surrogate featuring a single button.

6.3.2. Implementation of Networked Haptics

Haptic feedback in general requires high update rates, typically of 500 Hz or more [137]. Most modern networks, for instance Ethernet or the Internet, do not provide a quality of service at anywhere near this level. Several researchers are working to investigate and address this issue [201]. The goal of this research, however, is to investigate the potential, for a user, of collaborative haptics. Consequently, a high performance, loss-less transmission medium was required. To achieve this the refresh rate requirements for the haptic cursor communication were analysed.

In general, high update rates are required to support the production of stable stiff objects [137] and, in the haptic cursor communication, only the push effect involves the production of an object at all – the other cursor. The proximity effect does not require any such representation, merely modifying the viscosity of the workspace based on the distance to other cursors, and as such should be fairly insensitive to low update rates. Changes in the viscosity from one millisecond to another are relatively unimportant and should fall well below the perceptual threshold. The locate, grab and pull interactions all regard the other user's position as the target of a homing force and as such, when the range to the target is small, so is the magnitude of the targeting force. Lower forces can be adequately represented by lower update rates [137]. Conversely when the distance to the target is large, the targeting force is also relatively large. However, in this instance, the distance that the target must move to cause a substantial change in the targeting force is also large. These factors all

combine to make the haptic cursor communication relatively tolerant of low update rates.

Furthermore, an analysis of the total network demands of haptic cursor communication leads to the conclusion that even assuming a high update rate, a low total bandwidth is sufficient. This is because the only remote information of significance to the simulation is the position and state of other cursors – for instance whether they are engaged in a gesture. This can be expressed using a very small number of bytes. For instance, the current implementation requires the transmission of only 16 bytes per PHANToM per update: 4 bytes (representing a 32 bit floating point number) for each of the three axes, and 4 bytes of status information (representing a currently not fully utilised set of 32 separate 1 bit flags).

Finally, to further simplify development, some of the requirements of a collaborative system were relaxed. Firstly, it was decided that the system would be restricted to consisting of only two users and secondly that the implementation of the communication itself would be through a dedicated streaming connection rather than across a network. Both of these choices are commonly seen in groupware research [73, 104] and typically reflect the intention of the work - to evaluate the potential of the communication for a user, not to investigate underlying network performance.

With these analyses in mind, the haptic communication was implemented running over serial cables and providing a dialogue between two machines at 100 Hz. This implementation suffers from the restrictions of being simple for two users, but more complex for more, and of being confined to machines positioned only a few metres from one another. Due to the update-tolerant nature of the cursor communication, this refresh rate provides a subjective experience that is both stable and of an acceptable quality.

6.4. Initial Exploratory Study

6.4.1. Introduction and Motivations

Given the novel design of the haptic communication, and the novel context of use it was created for (a work scenario in a shared editor) the initial study investigating the cursor communication was chosen to be exploratory in character. The principal goal

of the study was not to precisely quantify objective performance improvements, but instead to provide a broad overview of as many aspects of users' perceptions and uses of the novel feedback as possible. A secondary goal was to use the results of this exploratory study to generate more defined and precise hypotheses for more focused studies in the future. These considerations are reflected in all aspects of the design of the study from the choice of experimental task, through to the assessment and measurement procedures chosen.

6.4.2. Experimental Task

The task chosen for this preliminary evaluation of the haptic communication was a CASE (Computer Aided Software Engineering) task. There were a number of reasons for this choice. Firstly, CASE is an established domain for collaborative tools [78]. Secondly, suitable problems, ideal solutions, and semi-expert users are easy to find in an academic computing science environment. Thirdly, as CASE tasks are, in essence, design tasks they lack a single solution and make good candidates for yielding complex collaborative activity [56]. The specific task chosen was to ask pairs of participants to read a hypothetical problem statement describing a paper based business system, and then collaboratively design a set of Unified Modeling Language (UML) [182] diagrams that describe a computer system that would replace it. This problem was taken from the syllabus of a university level software engineering class, and appears in Appendix F.

6.4.3. Haptic Feedback

The focus of this study was a two dimensional desktop editing task. In order to provide a standard pointer interface the PHANToM was given control of the cursor: movements made in the x and y axes with the PHANToM corresponded to x and y motions of the cursor on the screen. However, unlike standard pointing controllers the PHANToM is an absolute position device. A specific physical position of the device corresponds to a specific virtual position on screen. To accommodate this, the range of the cursor while under control of the PHANToM was restricted to the active experimental window, and stiff walls lined the edges of this area, preventing users from further motion. Operation of the PHANToM's button was mapped to a left mouse click.

According to the design of the haptic communication presented earlier in this chapter, movements in the z axis should control the interaction mode that the user is in (either working, communicating or observing). However, in light of the preliminary nature of this study, this use of modes was not implemented, and users were continually able to make modifications to canvas objects and to haptically communicate with one another. It was felt that while the modes would be an essential component of the haptic communication in a real world scenario, their use in this exploratory scenario might impede evaluation of the haptic feedback. Therefore, the haptic workspace users experienced was a narrow vertical plane. The dimensions of haptic workspace were 110mm by 110 mm by 20mm.

The five different haptic communications described in earlier in this chapter were all used in this study. Of these five, the activation of the proximity and push effects was solely mediated by position: they were turned on or off based around the relative locations of each user. The gesture, locate and grab effects, on the other hand, all required an explicit command to initiate and halt. All commands were mapped to the PHANToM's single button. To gesture, a user moved over another user and depressed the button; to end the gesture the button was released. This control mechanism was designed to support the metaphor of taking hold of, and pulling along, another user. The locate effect was activated by pulling backwards from the workspace and depressing the button. Releasing the button stopped the effect. Finally the grab effect was initiated by pulling backwards from the workspace, performing a double click and keeping the button down after the second click. Once more, releasing the button stopped the effect. All other buttons presses were passed to the system as normal mouse events. The rationale for using the PHANToM's button in this way was to ensure that the haptics could only be activated while the PHANToM was being held.

6.4.4. Collaborative Editor

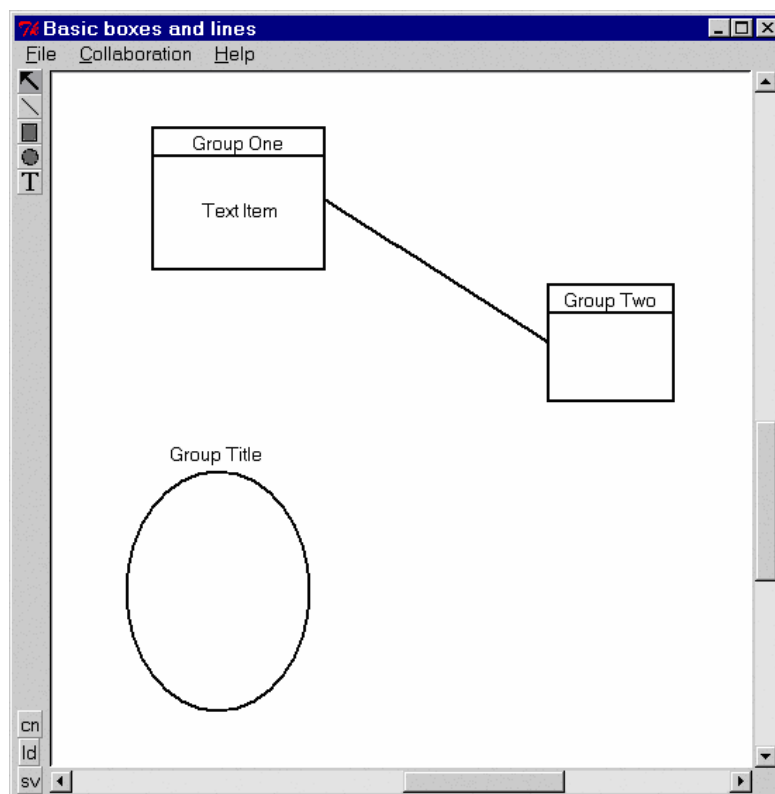
In order to build a CASE system that would enable the evaluation of the haptic communication, GroupKit [161], a high level tcl/tk based groupware toolkit (developed by Saul Greenburg and GroupLab at the University of Calgary) was employed. GroupKit provides support for the management of shared information,

including creation, manipulation and mediation of access, and basic groupware aids such as telepointers. It is a commonly used tool within the groupware community.

The product of coupling the haptic communication with GroupKit is CHASE (Collaborative Haptics And Structured Editing), a synchronous, Relaxed What You See Is What I See (RWYSIWIS) structured drawing tool, pictured in Figure 6.2. It provides telepointers and allows users to simultaneously work on a large canvas while each maintaining a separate view of it. CHASE allows users to create and edit four types of object: text items, rectangular groups, oval groups, and links. Text items can be placed in group objects and links can be made between them. All items can be freely moved, edited, resized and otherwise manipulated. This editor supports the creation of basic diagrams from a number of CASE methodologies, including UML.

One consequence of the combination of the haptic cursor communication and a collaborative RWYSIWIS workspace is that users can feel forces that attempt to move them outside of their current view of the workspace. An example of such a

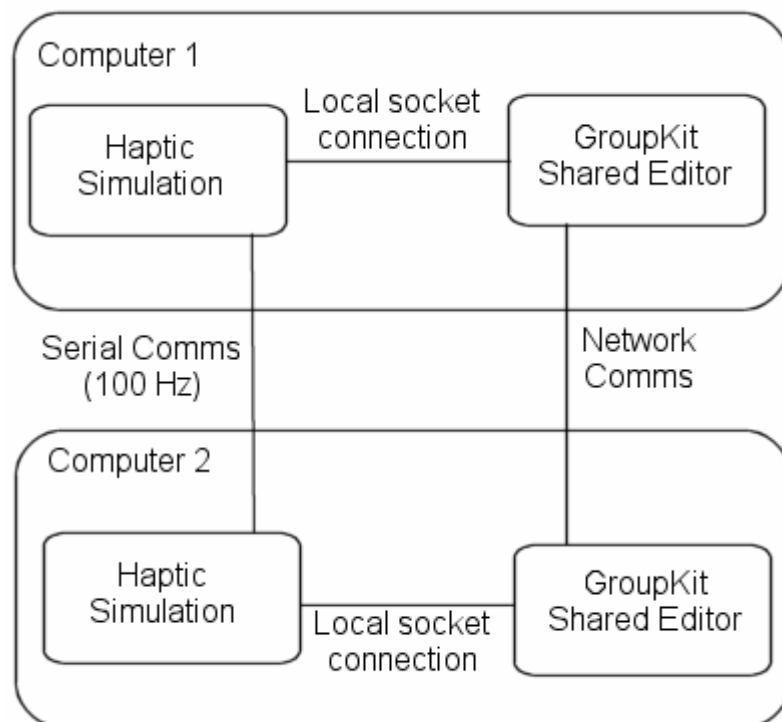
Figure 6.2. Screenshot of CHASE.



situation would be if one user activated the locate effect when the other user was on a distant (and currently not in view) part of the canvas. To resolve these forces intelligibly, the haptic workspace was restricted to only allow cursor movement within the window of the GroupKit application; walls were presented at the edge of the CHASE window. When users pushed into these walls, as they would if forces were pulling them off the workspace, the workspace scrolled in the direction of their push, gradually changing their view until their target was on screen. The adoption of this solution had the consequence of providing an unusual mechanism for scrolling – simply pushing into the walls of the workspace.

The haptic simulation communicated with CHASE through mouse events and also through a local socket connection. The mouse events formed the main part of the interface, and functioned simply as a consequence of using the PHANToM as a cursor control device. The socket connection was used to communicate all other information, such as the scroll position of the CHASE workspace, and to transmit scrolling events when the PHANToM pushed into a wall. A diagram of this architecture is shown in Figure 6.3.

Figure 6.3. Architecture of the CHASE system.



6.4.5 Participants

Sixteen users, all computing science students, in eight pairs, participated in the experiment. This group consisted of 9 men and 7 women, and had an average age of 25. The participants came from a relatively diverse range of cultural backgrounds: 10 were British, 1 French, 1 German, 1 Finnish, 1 Greek, 1 Indian, and 1 Azeri. None had previously used a shared editor, nor had anything more than trivial experience with haptic interfaces.

6.4.6. Experimental Design

The experiment was designed to be between subjects: four pairs, forming the Haptic condition, solved a problem with the aid of the collaborative haptic feedback; four pairs, the Visual condition, worked without it. Both sets used the PHANToM in the same restricted workspace, and had access to the same novel scrolling technique. This simple between subjects design was chosen primarily because of the complexity of the task. In order to gain the richest data on the participants' collaboration, they were being asked to solve a complex (and potentially lengthy) design problem, and it was felt that this process could be mentally taxing, and that subjects would be unwilling to solve two such problems.

6.4.7. Procedure

For simplicity participants were seated in the same room, separated by a partition. They could not see one another, and no video link was provided. They were able, and encouraged, to talk freely. A disadvantage of this simple set-up is that audio information from the environment, such as the sounds generated by key presses was present. This kind of information is not typically passed between individuals working at different locations, and its presence may have influenced user behaviour in the study.

Participants in both conditions went through an extensive training phase. A manual to CHASE, and in the case of the Haptic condition, to the haptic cursor communication was provided and each feature explained and demonstrated. This manual is presented in Appendix G. Participants were then required to copy a printed UML diagram into CHASE. At this stage it was clearly explained that the example diagrams consisted of acceptable UML constructs. The recording equipment was

then switched on and subjects then began to solve the UML problem. There was no time limit imposed on the task; subjects were instructed to stop when they believed the problem had been solved satisfactorily. This typically took an hour. After the completion of this task subjects filled in the subjective measures, described in the next section.

6.4.8. Measures

Observation

Users were observed, with both audio and video recorded, throughout the experiment. No formal dialogue analysis of this data is currently planned. One user's screen was captured with a screen grabber, the other by a video camera positioned behind the user. The second arrangement, while providing an inferior view of the screen, often actually provided more information as to the problems users were experiencing. Users would often make gestures or point when discussing critical items, and these movements were recorded by the camera pointed towards the screen. Audio was recorded by a single microphone positioned on the partition separating the two users.

Subjective measures

Four questionnaires were administered at the end of the experiment. Each questionnaire assessed a different aspect of subjective user experience. It was hoped that by gathering this broad range of data it would be possible to shed light on the influence of the haptic communication on widely different aspects of user experience and opinion. The four areas examined were workload, usability, presence and sense of togetherness or collaboration. These specific areas were chosen as each holds some direct bearing on the system. A measure of workload assesses how taxing participants found an experimental task, while a usability measure reveals practical problems with the system – does it support its users in achieving the tasks they are set? Given that the literature on haptic communication suggests it can positively influence levels of presence, and possibly a sense of togetherness or social presence, questionnaires to measure these factors are also important, and therefore included. Basic demographics were also gathered as an integral part of the presence questionnaire described in some detail below.

Once again, the workload questionnaire used was the NASA TLX [88]. As in the previous studies reported here, an additional term – Fatigue Experienced – was included in this questionnaire, as it seems likely that it is a key factor to assess with regard to haptic interfaces. The usability questionnaire administered was QUIS [117], a relatively standard questionnaire for assessing the usability of computer systems. This questionnaire consists of 27 items. The total score over all items gives an overall usability rating, but the results can also be decomposed into the following three factors: System Usefulness, Information Quality and Interface Quality. System Usefulness refers to whether or not the system appropriately supports the tasks it was designed for. Information Quality attempts to measure users' perceptions of the quality of the feedback presented in the system (for instance, whether or not the error messages, or screen layout, were well designed and useful). Finally, Interface Quality assesses whether or not users found the system pleasant and easy to use.

The presence questionnaire used was the ITC Presence Questionnaire [116], a recently created but well validated instrument. This questionnaire consists of 44 items, arranged into four factors entitled Spatial Presence, Engagement, Naturalness and Negative Effects. Spatial Presence refers to a user's sense of being in the virtual space the system represents. Engagement is an expression of how involving or enjoyable users found the system. Naturalness attempts to measure how interacting with the system compares with interactions and experiences in the real world. Finally, Negative Effects captures how physically disorientating a system is. An example of the effects this factors tries to measure would be the motion sickness, or dizziness that some users suffer from when exposed to technologies such as immersive visual displays.

The final questionnaire administered was created specifically for this experiment and was designed to assess collaboration. It consisted of ten items, which were rated on seven-point Likert type [63] scales ranging from strongly disagree to strongly agree. The questionnaire addresses five collaboration issues and each is measured by two questions. In each pair of questions, one elicits a higher rating for greater levels of collaboration, one a lower. The five issues addressed are: Communication Achieved, Location, System Support for Communication, Team Working and Awareness. Communication Achieved was intended to refer how easy it was, in general, to

communicate with the other user. Location looked at how simple it was to find the other user on the canvas. System Support for Communication tried to tap into user's perceptions of how well the system supported their desire to communicate. Team Working assessed whether or not participants felt they worked together. Finally, Awareness attempted to record how conscious of the other users actions participants felt they were. The ten questions are shown in Table 6.1. They are presented in the order they appeared when administered to participants. The table has also been augmented with details as to which factor each question is assessing. The use of this custom questionnaire, rather than a previously validated tool, reflects the lack of an accepted measurement instrument for assessing collaboration.

Objective Measures

Objective measures in the form of the time it took pairs of users to complete their diagrams, and final models they generated were also gathered, but little weight has been attached to them. This is due to the well-reported insensitivity of objective data gathered from collaborative tasks [56]. Consequently, no analysis of this data is currently planned.

6.4.9. Hypotheses

Given the discussion of the complexities of evaluation in collaborative systems (presented in the previous chapter), the preliminary state of this work, and the desire for this study to provide a broad spread of data in order to generate hypothesis for future studies, the experimental hypotheses are necessarily somewhat vague. They predict that the haptic communication will improve the quality of users' subjective

Table 6.1. Questions from custom collaborative questionnaire.

Question	Factor
There were times when I was unable to communicate effectively with the other user	Communication Achieved
It was easy to find the other user	Location
I found the communication in this system to be effective	System Support for Communication
I worked alone	Team working
The other user and I coordinated our actions together	Awareness
Communicating with the other user was simple	Communication Achieved
I often found it hard to locate the other user	Location
I was not aware of the activities of the other user	Awareness
The system did not support my desire to communicate	System Support for Communication
The software made it easy to work as a team	Team working

experience (as measured by questionnaire), and also that their observable behaviour will alter to take advantage of the novel functionality provided. They suggest that the haptic communication will be visibly and positively employed, and lead to decreased ratings of workload, and increased ratings of usability, presence and collaboration.

6.4.10. Results

Observations

General Observational Results

There was substantial variation in the use of the haptic communication. One pair of users in the Haptic condition did not use the communication at all, while others embraced it, using the various effects regularly. This may reflect the fact that touch is a very personal sense. Individual differences and social factors may well exert a strong influence over the adoption and use of this kind of communication. Furthermore it seems likely that communication of this sort breaks new social ground. Most users prefixed use of the haptic communication with a verbal warning, even when the communication would not affect the other user, as with the locate effect. Users have not previously been able to communicate in this way and appear unsure what protocols should mediate their interaction.

Users also appeared to find the interface to the haptic communication difficult. This could be attributed to the fact that the communication was completely controlled through the haptic device; it was simply overloaded. Users had the most problems initiating the locate and grab effects. These effects required them to operate the controller's button while maintaining a position against the back wall of the workspace, a relatively complex action with little interactive feedback to support it.

There was also a trade-off between perceiving the haptic feedback and performing the task. Participants had to type in order to complete the task, and when involved in this activity were unable to hold the PHANTOM, and therefore unable to experience the haptic feedback. However, as the task was not solely reliant on content creation, but also on discussion (during which times the participants typically held their haptic device) this effect did not overwhelm the observational results of the study. A final crucial point is that the majority of the users appeared to find the haptic communication engaging and helpful, rather than annoying or intrusive.

Observations of Locate and Grab Effects

The majority of the participants in the haptic condition immediately understood the purpose and applicability of the locate and grab effects. They used the locate effect regularly and the grab effect more infrequently. Reflecting the fact that users found it easier to find one another, pairs in the Haptic condition tended to use much more of the available canvas space and created diagrams that spread over a larger area than pairs in the Visual condition. Diagrams in the Visual condition tended to be very compact, with different sections abutting one another.


Participants in the Visual condition also used many different techniques for finding one another or specific objects, which were mainly absent from the Haptic condition. They would describe their position by references to the position of the scrollbars at their location, or by naming nearby objects or diagrams. One participant in the visual condition occasionally instigated a more extreme solution. Upon finding the other user's telepointer he would endeavour to maintain a view on it for as long as possible through rapid scrolling – giving the appearance of pursuing the other user.

One pair in the Visual condition did produce a diagram that occupied a large and diverse area of the screen. They suffered from considerable confusion while discussing where to begin new elements of their solution, and in keeping track of one another when they came to review their diagrams. Such confusion was far less evident in the Haptic condition.

Observations of Gesture Effect

The haptic gesture was used infrequently. Graphical telepointer gestures, such as when a shape is described by simply moving along its contours, or more commonly, when an object is indicated by moving over it, were far more prevalent. Contributing to this is the fact that there is a time cost associated with the haptic gesture when compared to visual gestures. A haptic gesture involves moving over another user, taking hold of that user, then moving back to the item of interest. A graphical gesture, on the other hand, simply involves the final stage of this process: moving over the item of interest. The haptic gesture also provides little enhancement of the most common use of a graphical gesture: indicating a single object.

Figure 6.4. TLX workload results from CHASE study.

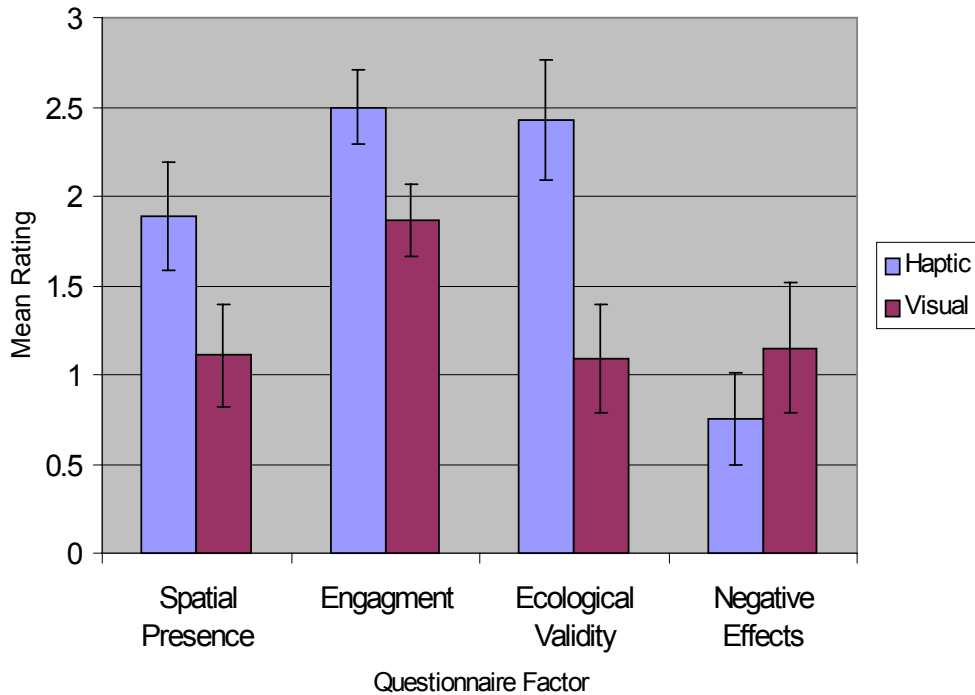


However, the haptic gesture was used in more complex situations. One pair of participants used the haptic gesture to indicate several objects, spread over an area of the workspace too large to fit on a single visually presented screen, and therefore too large to be effectively displayed with a graphical gesture. This may indicate that the haptic gesture is useful for illustrating complex sets of data. Another pair of participants used the gesture in an entirely novel way - one user selected a group object that was being discussed and the other user then began to gesture to the first user, steering her (and consequently the object) towards what she thought was an appropriate location.

Observations of Proximity and Push Effects

Observing any direct usage of the proximity and push effects was challenging, as neither requires any explicit action to initiate, nor causes an observable motion. However the two effects may have combined to increase a feeling of presence in the workspace. Users appeared to be more confident about using graphical gestures in the Haptic condition, which may have stemmed from an increased sense of presence brought about by the more tangible representation provided by haptic communication.

Figure 6.5. ITC Presence questionnaire results from CHASE study.



Subjective Measures

The raw data from these measures are presented in Appendix H. All analyses were conducted using two sample between groups t-tests [48]. Figure 6.4 shows the TLX results. Overall Workload did not significantly change between the conditions. The Haptic condition, however, was significantly more Physically Demanding than the Visual ($p < 0.05$), and the difference in the Frustration Experienced factor approached significance ($p < 0.065$). Figure 6.5 illustrates the results from the ITC presence questionnaire. The Haptic condition yielded significantly greater subjective ratings of Spatial Presence ($p < 0.05$), Engagment ($p < 0.05$) and Naturalness ($p < 0.5$). The results from the QUIS usability questionnaire are presented in Figure 6.6. Overall Usability was significantly improved in the Haptic condition ($p < 0.05$) as were the individual factors of System Usefulness ($p < 0.01$) and Interface Quality ($p < 0.05$). Results from the custom questionnaire are shown in Figure 6.7 (adjusted so that higher values consistently indicate increased collaboration). While this questionnaire was developed simply for this experiment, and as such little trust can be placed in the validity of the data that it produces, it is worth noting the unanimously rated superiority of the Haptic condition. Indeed, a comparison of the means from this questionnaire indicates that this superiority attains significance ($p < 0.01$).

Figure 6.6. CS usability questionnaire results from CHASE study.





Figure 6.7. Custom collaborative questionnaire results from CHASE study.



6.4.11. Discussion

The information gained from the questionnaires supports that gained through observation. With regard to the TLX workload questionnaire, it is possible that both the trend towards increased Frustration Experienced in the Haptic condition, and also

the significant increase in Physical Demand, are attributable to the observed difficulty users had with the interface to the haptic communication. Most users found the haptic communication hard to invoke, and this is an area that requires substantial improvement. The use of input mechanisms such as pie menus [37] may enable the creation of a more usable interface to the communication (by reducing the load on the PHANTOM's single button) while still ensuring that users can only activate the haptic feedback when they are holding their haptic device.

The results from the presence questionnaire also tally well with the observational data. In keeping with the observation that the haptic communication may increase a user's sense of presence, the Haptic condition yielded significantly greater subjective ratings of Spatial Presence and Naturalness. This indicates that the haptic feedback resulted in participants feeling not only more present with regard to their interactions in the canvas, but also that it became more like interacting in the real world. Furthermore, supporting the observation that users found the haptic communication appealing and engaging, it achieved significantly higher ratings than the Visual condition in the Engagement category. Taken as whole, the results from this questionnaire strongly suggest that the haptic feedback substantially improves user experience.

The data from the QUIS usability questionnaire are also strongly in favour of the haptic feedback. An improvement in Overall Usability indicates that in general participants in the Haptic condition found that the system easier to use. The improvements observed in the System Usefulness and Interface Quality factors decompose this finding into an improvement in the functionality of the system (the haptic feedback led to participants rating the system as better supporting the collaborative task) and in the ease of use of the system. This reflects the observations that users appeared to work more effectively in the Haptic condition. Given the problems that participants appeared to encounter with the interface to the haptic communication, it is perhaps unsurprising that it did not lead to an improvement in Information Quality. This factor represents the suitability of the information presented by the system as regards its own state.

Finally, the data from the custom questionnaire are also supportive of the suggestion that the haptic feedback improves user experience. While, due to the fact that the

questionnaire has not undergone a validation process, little weight can be attached to the formal analysis of the data, the general trend shown is compelling. The Haptic condition rates higher than the purely Visual condition on all questions. This informally suggests that the haptic feedback provides a broad spectrum of benefits to communication ranging from the specific (such as those isolated using the Location factor in the questionnaire), to the more general (for instance, the Communication Achieved factor).

Taken as a whole the questionnaire data powerfully supports the experimental hypothesis that the haptic communication would be beneficial to pairs of users working together in a shared editor. Although some aspects of subjective workload are increased by the presence of the haptic feedback, these can be attributed to poor design in the interface to the communication, and not in the communication itself. The usability, presence and custom collaboration questionnaires all yielded largely positive results on the haptic feedback. This success on a diverse range of subjective measures is strong evidence for the usefulness and value of the haptic communication.

The observational results of this experiment were valuable in clarifying and reinforcing the data from the questionnaires. They consistently explain, and provide a context for the more concrete data revealed in the subjective measures. However, they were also useful in suggesting future avenues for research. The observations contain sufficient detail about how participants used the haptic communication to allow speculation on novel hypotheses that might be interesting to examine in future studies. For instance, the gesture effect appeared useful when indicating complex information, such as gesturing to encompass a variety of objects, or a complex shape. An evaluation of this technique in such situations may yield interesting results. The locate (and to a lesser extent the grab) effect was well received by subjects and it may be interesting to compare this navigation and coordination technique to other, possibly visual, aids (such as those described by Gutwin *et al.* [84]). Finally, a direct comparison of the effect of the proximity and push effects on subjective ratings of presence might also prove informative.

6.4.12. Conclusions

In conclusion, both positive and negative aspects of the haptic communication were unearthed in this study. On the negative side, users experienced problems with the interface to the haptic communication, and there appears to be substantial variations in user adoption of the communication, probably due to the personal nature of the sense of touch. On the other hand, the majority of the participants appeared to find the haptic communication engaging and used it frequently. It also significantly increased subjective ratings of presence, improved usability and appeared to facilitate collaboration.

However, the goals of this study were to shed light on how participants perceived and used the haptic cursor communication and to generate hypotheses for further, more tightly focused, work. Within these parameters, it has been a wholehearted success. Taking both the largely complementary observations and questionnaire data, it has provided a broad assessment of the haptic communication and highlighted several interesting avenues for future research.

6.5. Haptic Gesturing Study

6.5.1. Introduction and Motivations

In the previous study, the observations of participants using the haptic communication led to the generation of a number of hypotheses for further study. In this section one of these hypotheses is examined in some depth – the contention that the haptic gesture effect used in the previous study aids the transmission of complex spatial data concerning shape is investigated. There are several reasons why this particular effect was chosen for detailed study over the others. Firstly, it is the most complex of the interactions in the haptic communication, not only algorithmically, but also in terms of what it enables two users to communicate to one another. This in itself makes it an interesting prospect for future study. Secondly, several novel and unexpected behaviours were observed when participants used this communication, suggesting that a closer study is likely to yield surprising and informative data. Finally, the kind of haptic guidance that the gesture effect represents has applicability in other domains (such as physical training [77]), and therefore any

results gained from a study may well generalise to a wider area of research. These factors combine to make this effect an intriguing topic for close examination.

6.5.2. Literature Review

Bodily gestures, or gestures of the hands and arms are a fundamental part of communication. They are typically used to illustrate, or reinforce, information presented in conversation. Graham and Argyle [77] empirically investigated the role of gestures in the communication of information concerning shape. They studied pairs of participants. One participant described a shape to the other, whose task was to draw it as accurately as possible. The communication took place either with or without the aid of hand gestures. The images produced were then systematically rated for similarity to the originals. Analyses of these ratings demonstrated that communication through hand gestures had significantly improved the quality of the images. This indicates that gestures can form a vital communication channel for conveying information about shape.

As mentioned in the previous chapter, gestures also have a role in group communication and work. In an observational study of paper-based group work Tang & Minneman [189] conclude that bodily gestures are used regularly and productively in groups to:

“...act out sequences of events, refer to a locus of attention, or mediate their interaction....”

Consequently, gestural information is often integrated into synchronous computer mediated collaborative systems. A common mechanism for achieving this is through the use of telepointers, a communication aid described in the section 5.2.4 of the previous chapter. However, the richness of telepointer gestures is low, when compared to hand gestures in the real world. Several systems have addressed this issue. Ishii *et al.* [103, 104] have presented a substantial body of work describing the design of ClearBoard, a collaborative system that overlays a vertically presented shared document or canvas with a view of a remote collaborator. The experience is analogous to facing another user through a sheet of glass. Tang and Minneman [189] describe VideoDraw, a similar system. It also involves the projection of a video overlay onto a shared workspace. Tang and Minneman’s system is horizontally

presented and focuses explicitly on the display of the hands. Both of these systems allow users to observe the gestures that another user is making, much as they would in the real world. Unfortunately, neither of these systems has been subjected to a thorough evaluation – in Tang and Minneman’s case there is little more than the statement that users reported that the gesturing capabilities were useful, and that they regularly engaged in them. Both systems are also subject to several restrictions. Firstly they are inherently limited to situations containing only pairs of users, and secondly they are confined (at least in their current reported designs) to a strict What You See Is What I See (WYSIWIS) architecture. This means that participants are required to have the same view of the shared canvas at all times. Different views of the canvas would render the gestures transmitted in these systems meaningless. Standard telepointers suffer from neither of these problems.

This experiment investigates the use of haptic guidance as a mechanism for improving the communication of gestures in collaborative systems. Several researchers have preliminarily investigated the effects of similar haptic guidance from the perspective of physical training. Yokokohji *et al.* [203] discuss how haptic and visual information can be effectively combined to teach students specific object manipulations. They consider combining graphical display with playback of the direct forces experienced by an expert, or of the motion taken by an expert (similar to the haptic gesture described earlier in this chapter). They conducted a preliminary study involving learning a specific manipulation of a cube in a virtual environment, and concluded that display of the motion was more effective than the display of force. Sakuma *et al.* [164] describe a system designed to study the influence that haptic guidance exerts on learning the skilled physical task of shodo, or Japanese calligraphy. An expert calligrapher haptically and repeatedly described a pair of Japanese characters to novices. The researchers then performed an informal evaluation of the quality of the characters produced by the novices with and without this haptic feedback. They concluded that haptic feedback had improved the quality of the novices’ calligraphy.

However, while physical training is one possible use for haptic guidance, it is not the only one. This study is concerned with using haptics to improve the communication between users in synchronous collaborative systems. It seeks to investigate whether

haptic guidance, or gesturing, can improve performance in cognitive, rather than physical, tasks. Correspondingly, the haptic gesturing is applied to the image-based experimental paradigm used by Graham and Argyle [77]. Telepointer gestures in the standard visual case, the haptic-only case, and the combined case which incorporates feedback in both of these modalities, are compared and contrasted against one another. The hypotheses of this study are that although the sole presentation of the haptic information may well reduce task performance and user satisfaction when compared to the visual presentation of this information, the combination of the haptic and visual feedback will be additive, leading to increased levels of performance and satisfaction.

6.5.3. Haptic Feedback

As mentioned in the introduction to this section, the experiment was conducted under Windows NT and used two PHANToMs from SensAble Technologies [128], each equipped with a pen stylus featuring a button, to provide the haptic interface. The haptic gesture or guidance used in this study is that described in the design section at the start of this chapter, and also that used in the previous study. In this study the haptic workspace was restricted to a narrow vertical plane 110mm by 82.5 mm by 20mm. This corresponded to a graphical range of 640 by 480 pixels. Motion along the x and y axes controlled cursor position. No action was mapped to motion on the z axis. Depression of the PHANToM's button was used to indicate one of a number of events, as discussed in the section describing the experimental task below.

6.5.4. Materials

The materials consisted of forty-five black and white line drawings. Twelve of these drawings were reproduced from the low codability images illustrated in Graham and Argyle [77]. These images have been exposed to a substantial validation process designed to produce images that are difficult to describe. The other 33 images were Chinese characters, or kanji, which were manually adjusted to contain only lines of a uniform width. These kanji were chosen on a basis of their distinctiveness from the wider set of available Chinese characters, and were subject to no validation process. The images were split into two groups. Fifteen images were used in the practice session and the remaining thirty in the experimental session. The ratio of kanji to Graham and Argyle's images was maintained for both these groups. These groups

remained fixed for all participants. However, within the practice and experimental sessions all images were presented randomly. The images used in this study are presented in Appendix I.

6.5.5. Task

There were two distinct roles in the experiment: encoder and decoder. The line drawings were presented to the encoder, whose task was to describe them to the decoder. The decoder had to reproduce these images. Decoders were able to draw black lines through the depression of a button on the PHANTOM's stylus, and erase lines through the same mechanism, combined with a modifier key. The encoders were not able to see any representation of the decoders' drawings. Participants were located in the same room, separated by a partition. Although they could not see one another, they were able, and encouraged, to talk freely. The encoder could also gesture (through telepointers) to the decoder. Gesturing was activated by the depression of the PHANTOM's button and could not interrupt the decoder's drawing activity. An hourglass cursor was used to signify that the decoder was unavailable for gesture interaction. Equally the decoder could not interrupt a gesture by beginning to draw, nor draw while a gesture was taking place. This prevented the decoder from simply tracing the shapes demonstrated by the encoder. Each trial was subject to a time limit of eighty seconds. The decoder was able to end the trial at any time, by pressing the space bar. A progress bar on the right of the experimental window kept both participants aware of the time remaining in each trial.

Figure 6.8. Encoder view in gesture study

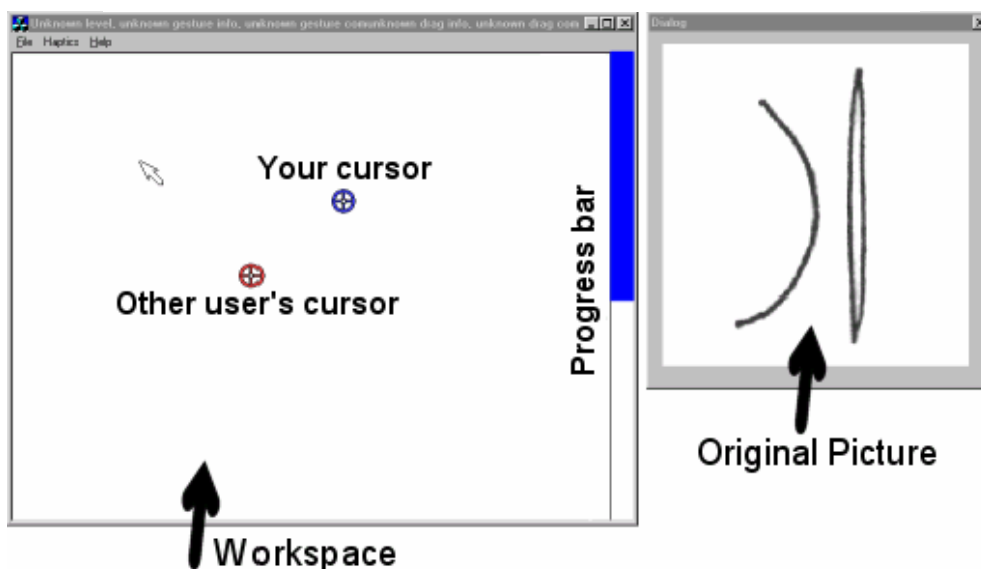
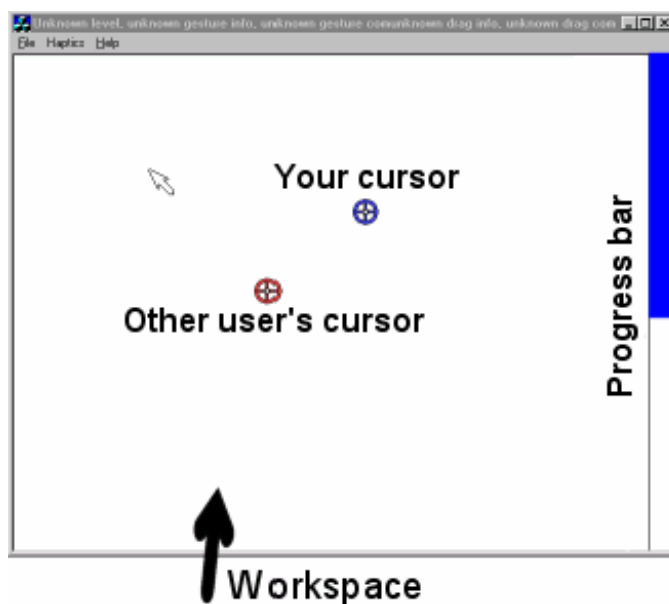


Figure 6.9. Decoder's view in gesture study.



The decoder experienced three different conditions: Visual, Haptic, and Combined. Each incorporated a different mechanism for gesturing with telepointers. In the Visual condition graphical gestures were available. When the encoder gestured, standard graphical telepointers appeared. Correspondingly, in the Haptic condition, haptic gestures were available. These are described in detail in the section of this chapter detailing the design of the haptic communication. Haptic gestures did not provide the decoder with any visual information – while a haptic gesture was underway, neither the local nor the remote cursor was drawn. The Combined condition incorporated feedback in both modalities. Decoders were able to both see and feel the shape described by the encoders. Encoders experienced all gestures in the standard graphical form only. Figures 6.8 and 6.9 illustrate the encoder and decoder interfaces to the experiment during a Visual gesture, while the instructions appear in Appendix J.

6.5.6. Participants

The experiment involved thirty-six participants, organized into eighteen pairs. They were aged between eighteen and thirty-one. Twelve were female, while twenty-four were male. Thirty-one were native English speakers and the other five were fluent. Four pairs of participants did not know each other prior to the experiment. The majority of the participants were computing science students, the rest were

experienced computer users. No participant had more than trivial previous experience with haptic interfaces, and all were unfamiliar with kanji.

6.5.7. Experimental Design

The experiment featured a repeated measures, fully balanced design. The decoder in each pair of participants experienced all three of the communication conditions (Visual, Haptic and Combined). To ensure there were no practice or order effects, there were six order conditions, each completed by three pairs of participants.

6.5.8. Measures

Objective Measures

Objective measures were taken in the form of task completion time. Completion time was manually controlled by the decoder (by pressing the space bar) and was consequently subject to some noise: participants exhibited a tendency to pause to appreciate finished artwork. Objective data in the form of quality was also gathered. The images produced by the encoders were captured and rated for similarity to the original images by three independent raters, none of whom were aware of the hypothesis of the experiment.

Subjective Measures

Demographic data was collected using the questionnaire shown in Appendix K. Subjective measures in the form of NASA TLX [88] were also gathered. Each decoder completed a TLX questionnaire for each gesture condition, while encoders simply completed one for the entire study. Once again, an extra factor, in the form of Fatigue Experienced, was added to the TLX questionnaires.

Observation

The entire experiment was also observed and recorded to video. A single microphone recorded all audio. A screen grabber captured the contents of the encoder's screen, including the original images and all cursor movements involved in gestures. A video camera was trained on the decoder's screen and captured the production of the image. This footage was intended to support observational results and no detailed analysis is currently planned.

6.5.9. Procedure

Before the experiment began, participants were allowed to choose which role they adopted, and remained in that role throughout the experiment. Practice took place immediately before the experimental session and consisted of fifteen images, five presented in each condition. The experiment was twice as long as the practice and contained ten image presentations in each condition. The order of conditions was always the same in the practice and experimental sessions. Participants were provided with comprehensive instructions as to their tasks, and the first few practice trials were spent guiding them through any problems they experienced. Both participants were made aware of the details of the three gesturing conditions, and were always kept up to date as to which condition was currently underway. Participants were paid £5 for taking part in the experiment and competed for a shared prize of £50 that was awarded to the fastest, most accurate pair. This provided some incentive for participants to try their hardest.

6.5.10. Hypotheses

The hypotheses of this study are that the haptic condition will yield poorer results than the visual condition on the scales of quality of images produced, subjective experience and task completion time, while the combined condition will produce superior results. It is suggested that although haptic gesturing alone may not provide as much information as visual gesturing, the combination of information in both of these modalities will be constructive, and result in more effective communication.

6.5.11. Results


Incomplete Data

Due to the self-regulated nature of the timing, there were three trials (attributable to two pairs) which ended accidentally after a short period of time. All three trials were in the Combined condition. The data from these trials was not included in any analysis.

Objective Measures

The timing data are presented in Figure 6.10. The raw figures are in Appendix L. Analyses of timing data were conducted using repeated measures ANOVA [48] followed by *post-hoc* t-tests [48], using Bonferroni confidence interval adjustments

Figure 6.10. Mean time per trial in gesture study.



[93]. The ANOVA revealed significant results ($F(2,17)=13.505$, $p<0.01$) which the t-tests clarified into the Haptic condition yielding significantly higher task completion times than either the Visual or Combined conditions (both $p<0.01$). There was no significant difference in completion time between the Visual and Combined conditions.

Subjective Measures

Results from the TLX questionnaire are presented in Figure 6.11, adjusted so that higher ratings consistently indicate higher workload. Like the timing data, all analyses of the subjective measures were conducted using repeated measures ANOVA and *post-hoc* t-tests, using Bonferroni confidence interval adjustments. Table 6.2 summarises the data from the ANOVAs. The t-tests showed that Overall workload was significantly higher in the Haptic condition than in the Visual and Combined conditions (both $p<0.001$). The Haptic condition was rated significantly more taxing than the Visual and Combined conditions in all individual scales except Physical Demand, Fatigue Experienced, and in the case of the Visual condition, Effort Expended (all $p<0.05$). There were no significant differences between the Visual and Combined conditions in any aspect of the subjective measures.

Figure 6.11. TLX workload results from gesture study.




Table 6.2. Results from ANOVA analysis of TLX results in gesture experiment.

	MD	PD	TP	EE	PLA	FE	FatE	Overall
F value	12.013	0.258	10.669	2.341	9.232	9.752	2.012	7.988
P value	P<0.01	Not sig.	P=0.01	Not sig.	Not sig.	P<0.01	Not sig.	P<0.01

Correlation of Image Ratings

Three otherwise uninvolved participants were paid £50 each to rate the images produced by the encoders for similarity to the originals. One rater failed to rate one image. It was assigned the average of the ratings it had received from the other two raters. These rating sets were then inter-correlated using the Pearson product moment coefficient of correlation [48]. Each set correlated significantly (at $p < 0.01$) with the others. The coefficients of correlation are shown in Table 6.3. Appendix L contains the image ratings, and Appendix M, a sample of the materials sent to each rater.

Analysis of Image Ratings

The average rating assigned by the three raters was calculated. The three truncated trials in the Combined condition were assigned a rating equal to the average of the

Table 6.3. Rater Coefficients of Correlation in gesture experiment.


	Rater 1	Rater 2	Rater 3
Rater 1	-----	0.534	0.604
Rater 2	0.534	-----	0.492
Rater 3	0.604	0.492	-----

remaining one hundred and seventy-seven ratings in that condition. These data were then scrutinised with a repeated measures ANOVA [48] and *post-hoc* pair-wise comparisons using t-tests [48] with Bonferroni confidence interval adjustments [93]. This process revealed significant differences in the data ($F(2,17)=22.456$, $p<0.001$). The t-tests showed this to be attributable to the fact that the Haptic condition led to significantly poorer images, when compared to both the Visual and Combined conditions ($p<0.01$). No significant difference was found between the Combined and Visual conditions. Figure 6.12 contains the average rating for each condition.

Observations

Both voice and gestural communication appeared essential for this task. Encoders often described shapes vocally and used the gestures to reinforce this. While some descriptions were structured (for instance “it looks like a backwards ‘C’”) it was common for the voice communication to be very vague. A typical example would be for the encoder to state “Straight along and then curves like this”, while the gestural

Figure 6.12. Mean image rating per condition in gesture study.



communication would be used to describe both the orientation of the straight portion of the stroke and the details of the curvature at the end. This behaviour occurred regularly and productively throughout all three experimental conditions.

Participants developed many strategies for maximising the usefulness of the gestural communication. Decoders in the Haptic and Combined conditions were constrained to follow the path taken by the encoder, and this was a commonly adopted technique in the Visual condition. One problem that emerges from this action is that it can be difficult to precisely locate the start point of the gesture once again. Several strategies were used to resolve this. Some pairs developed a solution whereby they marked the start point of each gesture with a dot. They did this by decomposing each actual gesture into two discrete gesture interactions. The first gesture interaction marked the start point, while the second described the shape. One pair developed this further and simply used the gestures to indicate key points that could be connected by straight lines. Another common strategy in the visual condition was to hover over the start point of a gesture for the duration of that gesture. Cursor position was used to signify the start point of the gesture.

Several versions of a strategy whereby the decoders were able to use gestures to communicate with the encoders were also present. In its simplest form decoders would mark the end point of a recently drawn line as the encoder began a gesture. This would ensure that a subsequent gesture began from an appropriate shared frame of reference. A more sophisticated form of this inverted communication was present in a few trials in the Visual condition. An encoder would initiate a gesture as normal, but instead, the decoder would perform the gesture – describing the shape that had already been drawn.

Most variation in strategy came within the haptic condition. One decoder drew the shapes backward, in order to maintain an appropriate position for the start of the next gesture. Another pair developed a technique of rapidly initiating and halting gestures in order to provide some visual representation of key points on the path. Several encoders returned the decoders to the start position of a gesture, often leading to confusion on the part of the decoder as to the location of the end point of the line to be drawn. Often the decoder would draw the line looping back on itself.

This misunderstanding of the end points of a gesture cropped up in several other ways. Many of the images included short lines, and these were often illustrated by rapidly moving up and down their length. In the haptic condition this led to confusion as to which end the line began and ended at. This became a problem when a subsequent line was said to start at either the beginning or end of the previous line, and the encoder and decoder had inverted representations as to the orientation of the line.

Participants also experienced great difficulty with the voice communication. Most encoders regularly confused their left and right (and to a lesser extent up and down, or backwards and forwards, if referring to a shape like a 'c') in speech, often while engaged in a directionally accurate gesture. It was encouraging to note that, while it was not universal, decoders in all conditions regularly questioned this conflict of information.

In cases of lesser conflict between information in different modalities (such as the origin of a new line, referenced to two different places on a previous line), it was unclear from simple observation whether decoders put more stock in what was said, or what was gestured. All decoders appeared to favour one or the other representation at different times. However, decoders did seem more likely to rely on visual information over voice, than haptic information over voice. This may be attributable to the novelty of communication in the haptic modality.

6.5.12. Discussion

The hypotheses that the Haptic condition would be significantly more subjectively demanding, take longer and result in poorer images than either the Visual or Combined conditions have been upheld. However, the Combined condition did not yield images of a significantly higher quality, nor result in a subjectively improved or more rapid experience than that attained in the Visual condition.

Several factors may have contributed to this fact. Haptic gestures inherently take time. A user is constrained to move along a path and, upon completion of the gesture, has to reorient his or her self before beginning another task. Only one subject used the potentially efficient strategy of drawing backwards from the completion point of a gesture. Haptic gesturing also forces a user to adopt a

particular strategy. In the Visual condition, it was common to follow the path described by the encoder, but it was not universal. Some participants did not adopt this strategy and even among those who did, it was typically not used for very short lines. No such choice was available in either of the haptically enabled conditions.

Training and experience may also be highly influential in this experiment. Beyond the simple fact that participants were unused to this kind of haptic feedback, the encoders were entirely naïve about the gestural communication. Although they were made aware of the three different types, and when each was in use, they never experienced the haptic gesturing, only the standard graphical gesturing. This may have led them to understand and pursue communication strategies that are effective only through visual means, and not through haptic ones. An example of this is gesturing by moving rapidly backwards and forwards along a short line. This leads to confusing haptic gestures, but is a commonly used and effective strategy in graphical gestures.

Although the primary hypothesis of this study - that gestures in the Combined condition would yield improved results when compared to purely visual gestures - was not shown, it does not rule out the usefulness of this kind of haptic feedback on collaborative canvases. Here the direct communication of information was investigated, but this kind of feedback may have other uses. In a canvas larger than a single window, a Combined gesture could be used to ensure that another user sees all elements of a gesture; it could be used to provide influence over (or guarantees as to the focus of) another user's attention. Understanding what another user is attending to is a common problem in collaborative systems [104]. It could also be used simply as a coordination tool. One user could use the gesture to transport another to some item of common interest.

The results from the Haptic and Combined conditions have some bearing on the physical training application domain briefly described in the introduction to this study. In the subjective measures results, the lack of significant increases in Physical Demand and Fatigue Experienced suggests that haptic guidance may well be suitable for the prolonged use that would occur in training simulations.

The data from the Haptic condition may also be relevant to interfaces for visually impaired users. Haptic feedback has the potential to provide effective computer interfaces and data visualisations for visually impaired people. Diverse research investigating different aspects of how this might be achieved is currently underway [47, 205]. While no study has yet investigated the impact of the kind of guidance described here on interfaces for the visually impaired, it is an obvious path down which to tread. The observations of participants in the Haptic condition – the details of strategies they developed, and the problems they encountered – could serve to direct the generation of more sophisticated haptic guidance techniques to support visually impaired users.

An interesting follow-up to this study would be to look at voice-only communication, to provide a baseline against which the other conditions could be measured. Similarly it might also be interesting to re-run each of the conditions described here without voice communication. Provided some mechanism was introduced to allow subjects to perform basic token passing communication, to indicate when an action had been completed, the tasks described here should be possible. A study such as this could investigate in greater detail the role of mixed modality communication of this sort. For instance it could address the effect that conflicting information in different modalities has on performance. Although voice communication was vital in the experiment described here, it was also responsible for a lot of wasted time, and inappropriate drawings. Encoders often confused their left and right in speech, but not graphically or haptically. This dissonance typically resulted in either a time consuming discussion as to the appropriate direction, or an incorrect drawing. The reduction in performance that would be expected from the loss of voice communication might be compensated for by the lack of cross-modal interference.

6.5.13. Conclusions

This study compared the effectiveness of standard telepointer gestures with purely haptic gestures and with gestures incorporating feedback in both modalities, in a task involving the communication of complex information about shape. As expected, the haptic gestures were less effective than the visual gestures, while unexpectedly the combined haptic and visual gestures did not improve on the results gained from the

purely visual gestures. While this study did not show the hypothesised improvement in performance using the combined gestures, the experiment did highlight interesting aspects of this kind of interaction.

In contrast to the majority of previous work on this, or related, topics [164, 203], the rigorous nature of the study described here means that the results have some general applicability. Due to the structured combination of subjective and objective information gathered, it seems likely that the insights gained can be applied to a variety of domains including interfaces for visually impaired users and virtual reality training simulations. They also have the potential to serve as lessons towards a more effective system of haptically enhanced gesturing.

6.6. Guidelines for Haptic Communication and Collaboration

The studies described here, in conjunction with some of the previous literature, suggest that haptic feedback has the potential to positively influence user experience, and, in some cases, task performance [166], in collaborative environments. However, the systems that exist have been created in a very *ad hoc* manner. Different researchers have built systems examining very different aspects of haptic communication, and functioning in very different contexts [13, 26, 40]. Due to the emerging nature of this field, little of this work currently possesses a great deal of depth. One consequence of the somewhat shallow but broad quality of the research in this area is that it can be challenging to gain an overall understanding of the issues involved in, and potential benefits of, the integration of haptic stimuli into communicative or collaborative systems. The lack of a concise summary of the conclusions of this body of work forms a barrier to advancing our understanding of this new and promising research topic.

To address this issue, the next section of this thesis presents a set of design guidelines to support the integration of haptic feedback into collaborative and communicative environments that draws together all the available literature relevant to this field. Reflecting the embryonic stage of research in this area, the goal of these guidelines is not to precisely define how to create effective communication through touch, but instead to summarise the generally applicable contributions from the

existing research and to describe a design space in which haptic communication has the potential to flourish. To this end, the guidelines cover a variety of topics ranging from very practical issues regarding what kind of haptic feedback might function feasibly over a network, through to grounded speculation as to the potential benefits of this kind of communication. These guidelines are intended to map out the significant issues relating to haptic communication in order to provide future designers considering the creation of novel forms of haptic messages with a framework describing the problems that must be overcome and the benefits that can be achieved.

6.6.1. Design Considerations for Haptic Communication

Latency

Perhaps the most significant practical challenge relating to haptic communication is problems with latency. The bi-directional transmission of real-time haptic communication (such as that required for an interaction like shaking hands) would need to be extremely rapid (in the order of 500 updates per second [137]) in order to provide a stable and realistic sensation. Such a high level of performance and quality of service is beyond current local area network technologies such as Ethernet [122], and far beyond that available over wide area networks such as the Internet. This issue, therefore, must be considered in the design of any haptic communication. Beyond the possibility that technological innovations (either in network protocols [122] or transmission algorithms for real-time haptic data [18, 201]) will solve this problem, there are a number of design solutions that can address this issue. In some contexts, one potential solution is simply to use a dedicated physical connection between two haptic devices. This solution works well in a relatively straightforward research context (as illustrated in this chapter and in the research of other groups [13, 166]) and may also have applicability in the area of interpersonal communication. Given the discussion presented earlier (in Section 6.3.2) indicating that although latency requirements are demanding, the overall transmission bandwidth required for haptic communication can be relatively low, it is possible to imagine a scenario where haptic data could be modulated onto an existing circuit switched network such as the telephone system [122]. The latency in such systems is typically restricted only by the physical properties of the carrier (for instance the speed of light in a fibre optic cable), and is therefore negligible in many cases. Accepting this fact, it is

possible to envisage a system that includes a real-time haptic channel for communication accompanying a traditional voice call, and it seems likely that this would make an interesting and feasible real-world scenario for haptic interpersonal communication.

Another potential solution, employed to some extent in the haptic communication described in this chapter, is to design feedback that operates effectively over higher latency systems. In this chapter one way this effect was achieved was by building feedback that had a reduced sensitivity to the remote environment. This is best illustrated by the proximity effect, which altered the magnitude of a non-disruptive haptic cue depending on some aspect of the remote system. Another method to reduce the latency demands is to consider situations which operate with uni-directional feedback. A good example of this technique is the gesture effect presented in this thesis. In this effect, the user making the gesture experienced cues that did not depend on the remote environment, while the data the other user perceived were buffered (in the form of the gesture path) and then experienced at a pace that depended on his or her input. This fact made the gesture effect relatively insensitive to update rate, and a similar technique has also been used by other researchers examining this kind of haptic communication [164].

A final solution addressing the problem of latency would be to explore the potential of haptic communication in asynchronous scenarios: situations where communication delays are not important, as the communication is not simultaneously bi-directional. Essentially, asynchronous communication involves one or more self-contained messages being sent between users. Existing and established examples of this kind of communication include letters, faxes, voicemail, emails, and instant messages. There are a number of compelling reasons why this might make an interesting topic to examine with respect to haptics. Firstly, asynchronous communication is extremely commonplace, and this pervasive nature suggests that it is a format that we find useful, meaningful and convenient. Secondly, and extending this point, asynchronous communication is already popular in a number of different modalities. Textual messages are most common, but audio messages are extremely established and visual messages in the form of mobile picture messaging [143] are an emerging technology. Given the general adoption and

interest in asynchronous communication of many forms, and in several different modalities, it seems likely that an investigation of the potential of asynchronous haptic messages could yield interesting and worthwhile results, while entirely avoiding problems with the latency of the underlying transmission medium.

Physicality of communication

As observed in the first study presented here, haptic communication requires physical attention: users must be holding or touching the haptic device in order to experience the feedback. In some situations, for instance where users have a larger task to complete, this activity competes for the limited resources of a user's body. They are either experiencing the communication, or completing the task. This fact seems likely to influence the kind of scenarios in which haptic communication is effective. For example, in the gesture study described in this thesis, the task was tightly coupled to the haptic device, and consequently participants experienced the vast majority of the haptic feedback. In contrast, in the first study presented in this chapter, users sometimes failed to experience the haptic communication as a requirement of the task was to type at the keyboard, away from the haptic device. The other user would transmit haptic information, and it would not be received due to the fact that the user was occupied elsewhere.

Much of the existing research on this topic has implicitly acknowledged this fact and examined scenarios that involve interaction solely with a haptic device. For instance, the majority of the other work examining collaborative tasks in virtual environments [13, 166] has focused exclusively on physical manipulations of virtual objects conducted using a haptic device, and therefore avoided conflicts in physical attention. Equally, the work on interpersonal communication has typically featured stand-alone devices. No other form of communication (such as typing) is supported in conjunction with the haptic communication, and consequently there is no competition for physical attention.

This issue seems likely to exert a significant effect on the kind of scenarios that are suitable for the addition of haptic communication. While haptic communication should not interfere with audio and visual communication, its successful application does appear to preclude physical activities that take place away from the haptic device. Such activities include other forms of physical input such as mouse or

keyboard use and designers need to be aware of this fact, and try to integrate haptic communication seamlessly and fully into the input and output aspects of the devices and scenarios that they are considering.

One interesting possibility for future examination may be wearable haptic devices [187, 193]. These devices are becoming richer in terms of the information they can transmit, and have the potential to sidestep the issue of physical attention, as they can transmit haptic information regardless of a user's current activity. Indeed, although there have been few serious examinations of these devices in communication scenarios, this fact is often an important motivation supporting and justifying their design.

6.6.2. Benefits Conveyed by Haptic Communication

Social Presence

The subjective data and observations from the first collaborative study presented in this chapter provide evidence that haptic communication can strongly and positively influence the social aspects of a collaborative task. They suggest that haptic feedback can increase levels of presence, and make users feel more engaged with one another. These data are backed up by the interests of, and results reported by, a number of other authors [13, 25, 67, 166]. Examined as a whole, this research presents a compelling case for the beneficial social aspects of remote communication through touch, and it seems likely this motivation will be a persistent force driving future research into this topic. Haptic communication appears to support strong interpersonal connections between remote individuals, a desirable feature that seems starkly lacking in current communication and collaboration systems.

Spatial Communication

The focus of three forms of the haptic communication presented in this thesis is the transmission of simple spatial information in the form of directional cues. The grab and locate effects allow users to identify their own or another's location, while the gesture effect supports the transmission of a relatively complex spatial pattern. The similar focus of these designs represents a fundamental observation about the kind of kinaesthetic cues that our haptic system is capable of perceiving: they are inherently directional. We have a strong and deep-rooted comprehension of directional

information (whether something is up or down, left or right) embedded within our haptic perceptual system. This inherent directional understanding can be used to create comprehension of more complex spatial information. Some evidence suggests that kinaesthetic spatial representations have a greater influence on our cognitive processes than visual ones [185] and, as observed in the gesture study presented in this thesis, it can be challenging to accurately convey spatial information through speech. Given this assertion - the native suitability of haptic cues to spatial, or directional, information - it seems likely that further work examining the haptic transmission of spatial data will yield interesting results, and prove valuable. This topic seems a likely and profitable candidate for further exploration within the domain of haptic communication.

Shared Physical Manipulation

Shared physical manipulation appears in a number of studies of haptic communication. This research has examined two main scenarios: the collaborative completion of a physical task [13, 166], and training scenarios [164]. The studies examining physical task completion typically report highly significant increases in objective performance, while, reflecting the more complex nature of a skill acquisition scenario, the research in training has yielded less conclusive results. Both these areas, however, seem extremely promising avenues for future research; the blue-sky benefits they suggest could be achieved seem substantial. The justification for remote training is relatively straight-forward: many skilled physical tasks, such as those involved in such disparate domains as complex surgery and a variety of forms of art, take many years to master and it seems likely that new forms of haptically enhanced tuition could speed up this process. Collaborative completion of a physical task, such as moving or manipulating objects, is applicable to a number of real world scenarios. While most studies have thus far examined these systems at a fairly conceptual level, looking at the kind of interactions that can be supported, this kind of shared physical manipulation could be invaluable in complex tele-manipulation scenarios such as tele-surgery. This suggestion is perhaps best illustrated through an example: if two people are working together in a remote environment then haptic feedback mediating their interactions with each other seems likely to be just as important as haptic feedback relating to the physical aspects of the remote environment. In short, although work is just beginning to appear examining haptic

feedback in tasks involving collaborative manipulation, the potential benefits that could be achieved in this area seem substantial, and this topic warrants further investigation.

6.6.3. Conclusions from Guidelines

The guidelines presented here do not form a cookbook for the design of haptic communication. Instead, and reflecting the infancy of this topic, they attempt to sum up the significant design problems influencing, and the potential benefits motivating current research into haptic communication. The goal of this activity is to concisely present all aspects of the current understanding of this topic so that researchers and designers of novel systems within this field can effectively and meaningfully work within the context established by the existing work. The goal of these guidelines is let the creators of future systems easily build on the work that has taken place previously, through the structured presentation of its main tenets.

6.7. Conclusions From Haptic Communication

This chapter has presented the design, implementation and a limited evaluation of a novel form of haptic communication based around augmenting the abilities of telepointers in a collaborative environment. Its final contribution is a set of design guidelines, the first presented on this topic, summarising the insights gained from the two studies, and from the previous literature.

Two evaluations were presented in this chapter. The first was a general evaluation featuring all aspects of the designed communication and involved two users collaborating on a complex design task. The subjective data gathered revealed that the haptic feedback could provide substantial improvements in user experience, and observations highlighted a number of interesting possibilities for further investigations. One of these possibilities was followed up in the second study, which looked at haptic gesturing in some depth. This study gathered a broad range of objective and subjective data, and although the hypotheses were not fully confirmed, its results were valuable, and applicable to a number of other domains.

Importantly, while the evaluation of the haptic communication was not fully comprehensive, the chapter has shown that the addition of haptic feedback to a user's

cursor in a collaborative scenario can lead to significant improvements in user experience as measured by established, validated and reliable questionnaires. It also reports detailed observations on users' perceptions and uses of the haptic communication. These represent a valuable contribution as they capture some of the context of use surrounding the communication that is absent in the purely numerical data gathered using objective measures of performance (such as task completion time) or questionnaires. The design guidelines presented at the end of this chapter further this achievement. They distil the available research on this topic into a form that is suitable for future researchers and designers to use as a starting point for the generation of novel haptic communication; they describe the major design issues that influence this area, and the motivations that drive it. They should prove a valuable reference for creation of future systems that support haptic communication.

This work fits into the overall structure of this thesis by demonstrating the second exemplar of the general hypothesis that haptic feedback applied to a user's cursor (or avatar in the virtual world) can provide qualitative and quantitative benefits. This chapter has examined a haptically augmented cursor in a collaborative scenario, and illustrated the benefits that this can bring in terms of subjectively assessed user experience and qualitatively observed alterations in coordination and work patterns. The design guidelines presented contribute towards the secondary aim of this thesis: to provide design knowledge on the specific topics studied. They ensure that the insights gained from the construction of the haptic communication and the empirical studies described are easily available to other researchers and designers.

7. Conclusions

7.1. Introduction

This final chapter presents an overview of the work described in this thesis. It is split into four sections. The first two sections describe a summary of the research conducted, and a discussion of the limitations of this work. In both cases the discussion is split into that which relates to the main claims of this thesis, and then into two more specific categories each covering one of the two exemplars explored in depth: desktop GUIs, and collaborative systems. The third section discusses where future work might extend the ideas this thesis puts forward, while the fourth section presents a summary of its overall contributions. Again, this final section is split into three parts, one covering the central claim and each of the two domains studied.

7.2. Summary of research

The central claim of this thesis is that the addition of haptic cues relating a user's representation, typically a cursor, in a computer system can provide both increases in objective performance and improvements in subjective experience. This claim was investigated through empirical study in two specific, and very different, domains. The results of this work were positive; the predicted improvements were observed in both cases, thus demonstrating the general claim. A further goal of this thesis was to develop design guidelines to support the adoption of haptic feedback in novel scenarios within the domains studied. These were also created, albeit at different levels of detail, achieving this second aim. The research conducted within each domain is briefly summarised in the following two sections.

7.2.1. Haptics in desktop user interfaces

Chapter 4 contains the research relating to the integration of haptic feedback into desktop GUIs. It begins by describing a study investigating the influence of four different forms of haptic feedback situated over a target in a simple selection task. Its goal was to contrast these different haptic cues against one another and against a visual condition featuring no additional haptic feedback. This kind of directly comparative study had not previously been conducted. The four different forms of feedback examined were gravity well (essentially an attractive force), recess (a bevelled area in the plane of the desktop), friction (a damping force) and texture (a pattern of concentric circles). The results of this study indicated that users achieved the best performance (in the form of the lowest error rates) when using the gravity well and recess conditions and the worst when using the texture condition. An important conclusion from this study is that haptic feedback has the potential to both substantially increase and decrease objective performance in this scenario.

The detailed measures used in this study also supported a hypothesis postulated by a number of other researchers [52, 140]: that haptic feedback that restricts users' movements has the potential to reduce their performance. The most significant example of such a situation is when multiple haptically active targets are displayed simultaneously. Haptic feedback relating to users' incidental movements over targets that they are not interested in has the potential to interfere with their intended actions. As the display of multiple targets is an essential part of desktop GUIs, the next section of this chapter turned to this topic. A novel mechanism for mediating haptic feedback in situations involving multiple targets was proposed. It involved the dynamic modification of the presented haptic cues based around directly measurable aspects of user behaviour (such as cursor velocity or direction). The haptic feedback for a menu widget that leveraged this principle was then designed and evaluated. This study compared a visual condition, a standard haptic condition (featuring feedback similar to that used in the previous study) and a condition that incorporated dynamically adjusting feedback. Its results confirmed the disruptive nature of standard haptic feedback when groups of widgets are considered (in terms of decreased temporal performance and increased subjective workload) and that the dynamic rendering approach proposed in this thesis can successfully counter this. The dynamic condition maintained the low error rates observed in the haptic

condition, at no cost in terms of time or subjective measures when compared to the visual condition. The concept of dynamic haptic feedback was then further explored through the design of haptic cues for another widget group: toolbars. This feedback was evaluated in toolbars, and in a related study on icons, by an MSc student [2], and these results provided further evidence supporting the validity of this idea.

The final section of this chapter attempts to distil the information gained from the studies reported, from the associated design process that led to the creation of feedback used, and from the previous literature on this topic in order to create a set of design guidelines that support the creation of haptic augmentations for use in desktop scenarios. These guidelines are firmly rooted in practical examples, and cover a range of issues relevant to the creation of haptics for use with GUI widgets ranging from the kind and strength of feedback to use, through to issues regarding how best to capitalise on the shape of a widget when defining a dynamically adjusting force profile. These guidelines should provide future system developers with a basis for creating haptic feedback that helps, and not hinders, users engaged in GUI tasks.

7.2.2. Haptics in collaborative systems

Chapter 6 presents the work conducted in the domain of collaborative systems. It begins by describing the design of five different interactions, each allowing a pair of users to haptically communicate with one another through the medium of shared cursors. The different interactions are termed push, gesture, locate, grab and proximity. Push allows two cursors to physically collide and push one another, as if they were objects in the real world. Gesture lets one user trace a complex shape, which is transmitted and haptically displayed to the other user in real time. Respectively, locate and grab allow users to activate a homing force leading them to another user, and to do the opposite, to activate a force summoning other users to their position. Finally, proximity entails the alteration of the viscosity of movement in the workspace according to the distance between users, providing an unobtrusive awareness mechanism.

The first of the two studies presented in this chapter examined the influence of the haptic cursor communication in a high-level design task conducted in a shared graphical editor. The goal of this study was to gain an overall understanding of how

the haptic communication was perceived and used by participants. It had a relatively simple structure. Pairs of participants were requested to solve a CASE problem collaboratively, and two conditions were compared. The first featured all aspects of the haptic communication, while the second featured none. The results of this study came in the form of a bank of questionnaire data, and from observations. Analysis of the questionnaire results revealed that, despite difficulties with the interface to the communication, users found the feedback compelling and engaging and felt that it made the collaborative task easier to complete. The observations reinforced this sentiment, illustrating how the haptic feedback had allowed users to work together more effectively. The observations also revealed a number of novel behaviours, which appeared to warrant further, and more focused, investigation.

The second study reported in Chapter 6 continued this thread of investigation and detailed a focused experiment examining one particular use of the gesture communication. It was concerned with the ability of the haptic gesture as a tool of explanation; it addressed the question of whether or not haptic cues could be used to increase the effectiveness of techniques to remotely convey information regarding shape. This study compared three conditions in which one user described a shape to another, whose task was to draw it. Each condition featured a voice channel for communication, and either purely visual gestures (in the form of the motion of a cursor), purely haptic gestures (described in detail at the beginning of Chapter 6) or the combination of these two. The main hypothesis of this study was that the combined condition would allow users to constructively incorporate the information perceived in these two modalities, leading to qualitatively improved, or more rapidly attained results. However, the objective measures used – task completion time and an independently gathered measure of the similarity of the copied images to the originals – did not confirm this hypothesis. Unsurprisingly, the purely haptic condition yielded lower performance than the other two conditions, but the combined condition did not offer an improvement over the solely visual condition. Observation of users in this study suggested several reasons why this might be the case, and these represent valuable insights for the designers of any future system for haptic gesturing.

The final section in this chapter attempts to collate the work on haptic communication and collaboration into a set of design guidelines. The focus of these guidelines differs somewhat from the focus of those presented on the topic of haptics in GUIs; the goal in this instance is to present the critical large-scale design issues relevant to this topic, and to map out the unique potential this modality possesses in communication or collaboration scenarios. The reason for this difference stems from the emerging and embryonic nature of this topic. Currently, it is a largely unexplored field, and the work that has taken place is composed of relatively novel but shallow contributions. In this situation, the guidelines adopt the entirely appropriate role of attempting to map out the available space in a concise and explicable form.

7.3. Limitations of this work

The crux of this thesis is essentially a proof by example; a general claim was instantiated in two domains, and through a limited but rigorous empirical investigation, shown to be true in these situations. This was taken as sufficient to illustrate the validity of the general claim. However, while these two examples, partly due to their disparate nature, do provide a compelling case supporting the general claim, an investigation of further examples would strengthen this. Unfortunately, such investigations are beyond the scope of this thesis, and this represents a weakness of this work that can only be addressed by further research on this topic. One domain which may prove an interesting and informative subject for future investigations is drawing or graphical modelling tools, where the mouse is often used to perform complex and varied operations, or as a device for artistic expression. Some existing modelling software [42] incorporates related haptic feedback as central part of its interface but an investigation from the perspective of augmenting the cursor with haptic cues to support user actions may yield novel designs.

The limitations of the work conducted in each of the two domains are described in the following two sections. Where necessary, these refer to the more detailed discussions presented in the Chapters 4 and 6, that fully describe this work.

7.3.1. Haptics in desktop user interfaces

The primary limitation with the work on haptics in GUIs relates to whether or not it can be generalised to novel scenarios. This work explores the addition of feedback in only a handful of very specific situations. Two studies, investigating button and menu widgets, are described in detail, as is the design of another form of feedback for use with toolbars. The influence of this feedback is examined in both toolbars and icons by Adams [2]. From these three designs and four studies, and from an understanding of the previous literature on this topic, this thesis draws up a set of design guidelines intended to cover the creation of haptic feedback to accompany any widget. Although these guidelines are general-purpose in character, it is clear that they would be strengthened and extended through the creation of further designs and empirical investigations. Interesting interface components to examine would include complex widgets such as colour selection tools, or different forms of GUI interaction, such as drag-and-drop.

One further important unexamined aspect of this work is the generalisability of the haptic cues studied to different display devices; all the research described in this thesis was conducted using the same haptic device, the PHANToM [128]. The developing nature of haptic hardware means that there are substantial differences between the quality of feedback that can be rendered, and the form factor adopted, in currently available haptic devices. In terms of quality, the feedback studied in this thesis is near the high end of the available spectrum, a beneficial property as it seems likely that future devices will attain or supersede this level of rendering performance. Therefore, the feedback examined here seems likely to become more applicable, rather than less. A more serious concern relates to the form factor of the PHANToM. As with many currently available haptic devices [22, 69] the PHANToM is designed around the pretext that a user must hold and interact through a surrogate that can be moved freely in 3 dimensional space. This free floating interaction is substantially different from current mouse interaction, where users work with a device that rests on a physical surface, typically the top of a desk, and often support their arm on that surface. This form factor may add stability, and increase the fidelity of user movements. Although the guidelines presented here attempt to unify research conducted using different display devices, there has been little research directly comparing the same haptic augmentations on different hardware, and research into

this essentially ergonomic issue is an important next step for this topic. If the haptic cues isolated as beneficial in this thesis can be shown to be effective on a variety of different hardware platforms, this will add weight to the argument supporting the integration of these cues into GUIs. On a somewhat promising note, Yu *et al.* [204] compared user performance on two different haptic devices (the PHANToM [128] and the Wingman Force Feedback Mouse [118]) in a graph exploration task, and found little variation between the two platforms. This result suggests that the value of haptic cues may be relatively independent of the form factor of the display device, and is supportive of the general applicability of the research presented in this thesis. However, further work is required to extend and explore this finding.

More focused issues also affect the buttons study presented at the beginning of Chapter 4. Perhaps most significantly, there are a number of issues with the experimental design. Although this study was purposefully constructed to examine the effects of many different forms and strengths of haptic feedback in a simple targeting task, it was arguably over-complex. For instance, too many conditions existed for the construction of a balanced experimental design (with any reasonable number of participants) and so a compromise was created that presented the conditions in a precisely controlled order that attempted to minimise practice effects. A statistical analysis suggested this approach was successful, but the conclusions of the study may have been clearer if it had featured a simpler design, comparing fewer conditions. Furthermore, this study also presented each form of haptic feedback at 3 different magnitudes, then averaged this data when comparing between different forms. While few significant differences were observed within the magnitude conditions, suggesting this was an acceptable approach, it seems likely the conclusions of this study would have been made firmer if it had possessed a simpler design featuring feedback at only a single magnitude.

Furthermore, as discussed in some length in Chapter 4, the timing measures used in the buttons study were not optimal. Specifically, the approach time measure did not accurately capture the time it took users to move from one target to another, and instead varied inversely according to the number of slip-off errors users performed. Consequently, it did not make a satisfactory temporal measure; on average, as error

count increased, it decreased. A future study that accurately measures this variable may illustrate more substantial temporal differences in performance.

7.3.2. Haptics in collaborative systems

The work on haptics in collaborative systems is primarily limited by a lack of objective evidence illustrating the benefits that this thesis maintains it can provide. The observations and subjective data reported in the first collaborative study are compelling, but are not backed up by the objective data gathered in the second study examining the haptic gesture. This work would clearly be improved if supportive objective data were to be unearthed. However, as mentioned in Chapter 5, it is worth noting that it is notoriously difficult to gather valid objective results in studies of CSCW systems; this is a problem that affects the topic as a whole, and not only the work presented in this thesis.

One potential route through which objective data could be gathered is through further investigations of the hypotheses generated by the observations presented in the first collaborative study. Although a number of hypotheses were generated, covering all aspects of the haptic communication, only one of these - the use the haptic gesture to convey complex shape - was explored empirically, and focused investigations into the others may reveal the desired objective data.

Each of the studies also possesses specific limitations. The limitations of the first collaborative study are primarily concerned with the validity of measures used within it. The goal of evaluation was to assess user opinion of, and subjective experience using, the haptic feedback. However, it was conducted with a relatively small pool of participants; of the 16 users in the study, only 8 actually experienced the haptic communication. Furthermore, while the measures used were appropriate for the goals of the study, they could have been strengthened with further techniques. Given sufficient time and expertise, a transcription and dialogue analysis of the speech between the participants may have provided a more concrete representation of the differences in collaborative practice reported in the observations. Some recent studies [21] have also used trust games, such as Prisoners' Dilemma, as an objective measure of the social relationship attained between participants in a collaborative task. The use of these more quantifiable measures may have increased the validity of the results of this study, and provided objective data supporting its claims.

The limitations of the second study examining the haptic gesture are less concerned with the measures used, and more with the selection of the experimental task. While the shape description task chosen is relevant to gesturing in the real world, which is often used to describe shapes or objects, it may not be appropriate in the case of a shared editor. A recent study by Gutwin *et al.* [83] compared regular telepointer gestures (similar to the graphical gesture studied in this thesis) against a condition in which the telepointers had been augmented with a gradually fading line (or trail) indicating their recent path. The tasks examined in this study were much simpler than the task used in the gesture study described in this thesis and included indicating a group of existing objects by encircling them and highlighting a line connecting two objects. The goal of Gutwin *et al.*'s manipulation was to overcome the effects of network delay on the perception of telepointer gestures, and their study firmly indicated that this was achieved through the addition of cursor trails. While the goals of this study are not exactly coincident with those of the gesture study in this thesis, the measures employed have considerable bearing. Gutwin *et al.*'s study attempts to measure performance on the kinds of tasks that users perform in shared editing environments, and not on more abstract tasks, such as that examined in the gesture study. It is possible, therefore, that an investigation of the haptic gesture relying on measures more akin to Gutwin *et al.*'s would reveal subtler, and critically, quantifiable differences in performance. Such measures may prove to be a more sensitive indicator of real performance in tasks involving gesturing in shared editors.

7.4. Future work

This section does not comment on future investigations related to the studies conducted in this thesis, as this is covered in the specific chapters where this work is described and, to some extent, in the limitations section above. Instead, it discusses the more general issues that this thesis has raised and that appear to warrant further attention.

7.4.1. Guidelines for the addition of haptic cues to user interfaces

An important aspect of this thesis is the development of design guidelines to support the addition of haptic cues to a user interface. Design guidelines represent a valuable

and practical first step towards enabling system developers to integrate haptic cues into their applications in a way that is appropriate: a way that provides users with greater benefits than costs. The guidelines presented in this thesis attempt to achieve this, but, unfortunately, may not be not widely applicable. They seem unlikely to generalise to novel scenarios as they are inherently domain specific. They provide relevant insights within their specific topics, but may have less to offer if applied to other scenarios. This fact can be illustrated by simply comparing the two sets of guidelines (presented in Chapters 4 and 6) against one another. They are radically different, and each has little to offer the other domain.

An overall, and relatively long-term, goal for future study would therefore be the generation of design guidelines for haptic feedback that are broadly relevant: that have bearing on haptic interfaces in general. It would seem likely that such a set of guidelines would not be as directly informative as those presented in this thesis; they would not attempt to cover the intricacies within any specific topic. However, they would have the beneficial property of conveying an overall sense of what haptic feedback is able to provide users and how interfaces that successfully leverage the potential of this modality can be created. The generation of such a set of general-purpose design guidelines could only be achieved through the creation of further guidelines for specific domains, a worthy goal in itself. When guidelines for a sufficient number of domains have been created, their contents could be generalised into a set of generally applicable meta-guidelines.

7.5. Summary of contribution

This section summarises the contributions of this thesis to the field of haptic interface design. Even considering the limitations described above, this thesis represents a substantial advancement. It proposes a novel design approach for the addition of haptic feedback to computer interfaces – that of considering the relevance of haptic cues from the perspective of a user’s representation in a system – then moves on to instantiate and validate this idea in two distinctly different domains. This validation serves to demonstrate the general claim. Beyond this illustration, this thesis also contributes significantly to the body of design knowledge relating to the integration of haptic cues into user interfaces. It presents the first design guidelines

covering this topic for the two domains studied. The contributions of this thesis will now be discussed according to each of the two areas studied.

7.5.1. Haptics in desktop user interfaces

The research presented in this thesis represents a significant advancement to the literature relating to the addition of haptic cues to desktop user interfaces. The first study described compares the performance supported by different forms of haptic feedback in place over individual targets. No other empirical research has directly compared different haptic cues in this way, and this work provides firm evidence demonstrating what kind of feedback is appropriate for a targeting task. This issue was critical, as there had been many different examinations of this task using a variety of forms of feedback [3, 90] and primarily stemming from the absence of the kind of comparative data that this study has provided. This study conclusively showed what kind of haptic feedback is effective in a simple targeting task.

The second study directly addressed the issue of whether or not haptic feedback can provide benefits in situations incorporating multiple simultaneously active haptic targets; situations where users will experience haptic cues that are not directly relevant to their current task, and may hinder their performance. The contributions of this study were twofold. Firstly, it empirically confirmed that haptic feedback can reduce user performance in situations involving multiple targets, a possibility that had been the subject of some speculation within the research community. Secondly, it proposed and evaluated a novel solution to this problem based around the dynamic alteration of the forces rendered according to directly measurable aspects of user behaviour. This idea represents a significant contribution as other potential solutions to this problem, such as target prediction, currently suffer from unresolved algorithmic and usability problems. The approach introduced here does not suffer from these disadvantages.

The final contribution on this topic was a further design exercise around the idea of dynamic alteration of haptic cues (evaluated by Adams [2]), and the development of a set of design guidelines that support the creation of haptic widgets. These guidelines are detailed and general purpose. They represent a substantial contribution to this field as they are the first guidelines to be created dealing with this topic and

are firmly grounded on empirical evidence. They enable system developers to integrate useful haptic cues into their new projects and products.

7.5.2. Haptics in collaborative systems

The research presented in this thesis on haptic feedback in collaborative systems has an important bearing on this emerging field. Its first significant contribution is one of design; it presents a novel perspective for haptic communication. This perspective, based around the cursors of two users interacting directly with one another, differs from much of the previous literature which has focused on interactions through the mediation of a shared virtual object [13, 166], and may have much to offer users. One significant advantage of this approach is that, as it natively focuses on communication between users, explorations within its framework have the potential to reveal a much richer set of user-to-user interactions than other approaches to haptic communication. This suggestion is supported by the diversity of the feedback designed and presented at the beginning of Chapter 6.

The empirical examination of the haptic communication also makes several key contributions. The first of these is that the subjective data from the first collaborative study strongly supports the claim that haptic cues can increase levels of virtual and social presence among users. Although other research supporting this assertion is now available, this study was one of the first systematic empirical examinations of this topic. As such it is a significant contribution. The observational data from both communication studies also represents an important contribution. This information describes, at a the level of an informed observer, many of the nuances and unexpected uses of the haptic communication. This kind of data is a regular output of CSCW research, and reflects the complexity of the topic. The observations provide significant in-depth insights into the real use of the haptic communication, and are something that future researchers and designers could benefit from. Due to the more focused nature of the second collaborative study, the observations of users in the gesture experiment are especially relevant; they reveal a great deal about the mechanisms by which complex telepointer gestures are conveyed between users, especially (but not only) with regard to haptic communication. These observations seem likely to be important in the development of gesturing systems in a number of

other domains, such as systems for visually impaired users, or systems for physical training.

The final contribution from this domain is the development of a set of design guidelines on the topic of haptic communication. Reflecting the infancy of this research area, these guidelines do not seek to make firm recommendations as to how haptic communication should be created, or what its characteristics should be. Instead, their goal is to map out the practical design issues that currently constrain research into haptic communication, and to summarise the motivations for studying this area. These guidelines are a valuable contribution as they are the first to be created in this domain, and also as they serve to define the space available for work in this field. They provide a concise summary of the salient points of this small but disparate body of research, in a form suitable for consumption by future researchers and designers.

7.6. Conclusions

In conclusion, this thesis set out to illustrate a general claim – that haptic cues relating to a user’s cursor can lead to performance benefits. This was achieved and, in illustrating this claim, two domains were explored and rigorous empirical evidence was gathered. Design guidelines serving to further the adoption or integration of this kind of feedback into the research or system development communities were also presented. These guidelines are the first of their kind on their respective topics. This thesis has provided a novel, empirically grounded, approach for addition of haptic cues to user interfaces and, through the design guidelines, provided a mechanism whereby this can be adopted, and hopefully generalised, by system developers. This is an important first step to achieve the integration of haptic cues into user interfaces so that the general population can enjoy the benefits that this feedback can convey.

8. References

1. Accot, J. and Zhai, S. *Performance Evaluation of Input Devices in Trajectory-based Tasks: An Application of Steering Law*. In Proceedings of *ACM CHI'99*. 1999. Pittsburgh, PA: ACM Press: pp. 466-472.
2. Adams, A., *Evaluating Haptic Forces in a Multi-Target Environment*, Work submitted for Masters of Science in *Department of Computing Science*, University of Glasgow: Glasgow, 2002.
3. Akamatsu, M. and Sate, S., *A multi-modal mouse with tactile and force feedback*. *International Journal of Human-Computer Studies*, 1994(40): pp. 443-453,.
4. Akamatsu, M. and MacKenzie, S., *Movement Characteristics using a Mouse with Tactile and Force Feedback*. *International Journal of Human Computer Studies*, 1996(45): pp. 483-493.
5. Alhalabi, M.O. and Horiguchi, S. *Tele-Handshake: A Cooperative Shared Haptic Virtual Environment*. In Proceedings of *EuroHaptics 2001*. 2001. Birmingham, UK: University of Birmingham: pp. 60-64.
6. Anderson, T. and Breckenridge, A. *A Graphical and Haptic user Interface for Creating Graphical, Haptic User Interfaces*. In Proceedings of *Fourth Phantom User's Group Workshop*. 1999. Cambridge, MA: pp. 48-51.

7. Apple, *Human Interface Guidelines: The Apple Desktop Interface*. 1987, Reading, MA: Publishing Company, Inc.
8. Ashley, S., *Simplifying Controls*, in *Automotive Engineering International*. 2001. pp. 123-126.
9. AutoDesk. *AutoDesk Home Page, 2002*, <http://www.autodesk.com/>
10. Baecker, R., Glass, G., Mitchell, A., and Posner, I. *Sasse: The Collaborative Editor*. In Proceedings of *ACM CHI'94*. 1994. Boston, MA: ACM Press: pp. 459-460.
11. Baecker, R., Nastos, D., Posner, I., and Mawby, K. *The User-Centred Iterative Design of Collaborative Writing Software*. In Proceedings of *ACM CHI'93*. 1993. Amsterdam, NL: ACM Press: pp. 399-405.
12. Barrett, J. and Krueger, H., *Performance effects of reduced proprioceptive feedback on touch typists and casual users in a typing task*. Behaviour and Information Technology, 1994. 13(6): pp. 373-381.
13. Basdogan, C., Ho, C., Srinivasan, M.A., and Slater, M., *An experimental study on the role of touch in shared virtual environments*. ACM Transactions of Computer-Human Interaction, 2000. 7(4): pp. 443-460.
14. Beaudouin-Lafon, M. and Karsenty, A. *Transparency and awareness in a real-time groupware system*. In Proceedings of *ACM UIST'92*. 1992. Monterey, CA: ACM Press: pp. 171-180.
15. Beauregard, G.L. and Srinivasan, M.A., *Sensorimotor interactions in the haptic perception of virtual objects*. 1996, RLE TR-607, MIT: Cambridge, MA.
16. Begole, J., Rosson, M.B., and Shaffer, C.A., *Flexible collaboration transparency: supporting worker independence in replicated application-sharing system*. ACM Transactions of Computer-Human Interaction, 1999. 6(2): pp. 95-132.

17. Bekker, M., Olson, J., and Olson, G. *Analysis of Gestures in Face-to-Dace Design Teams Provides Guidance for How to use Groupware in Design*. In Proceedings of *Designing Interactive Systems (ACM DIS'95)*. 1995. Ann Arbor, MI: ACM Press: pp. 157-166.
18. Belghit, I., Hennion, B., and Guerraz, A. *Predictive Algorithms for Distant Touching*. In Proceedings of *EuroHaptics 2002*. 2002. Edinburgh, UK: University of Edinburgh: pp. 61-65.
19. Berkelman, P.J., Butler, Z.J., and Hollis, R.L. *Design of a Hemispherical Magnetic Levitation Haptic Interface Device*. In Proceedings of *ASME IMECE*. 1996. Atlanta: pp. 483-488.
20. Bevan, N., *Measuring Usability as Quality of Use*. *Software Quality Journal*, 1995. 4: pp. 115-150.
21. Bos, N., Olson, J., Gergle, D., G., O., and Wright, Z. *Effects of four computer-mediated communications channels on trust development*. In Proceedings of *ACM CHI'02*. 2002. Minneapolis, MI: ACM Press: pp. 135-140.
22. Bouguila, L., Ishii, M., and Sato, M., *Scaleable SPIDAR: A Haptic Interface for Human Scale Virtual Environments*. *Haptic Human Computer Interaction*, 2001. Springer-Verlag LNCS 2058.
23. Bouzit, M., Popescu, G., Burdea, G., and Boian, R. *The Rutgers Master II-ND Force Feedback Glove*. In Proceedings of *IEEE VR2002, 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. 2002. Orlando, FL: IEEE: pp. 145-152.
24. Boyd, J. *Floor Control Policies in Mutli-User Applications*. In Proceedings of *ACM CHI'93*. 1993. Amsterdam, NL: ACM Press: pp. 107-108.
25. Brave, S. and Dahley, A. *inTouch: A Medium for Haptic Interpersonal Communication*. In Proceedings of *ACM CHI '97*. 1997. Atlanta, GA: ACM Press: pp. 22-27.

26. Brave, S., Ishii, H., and Dahley, A. *Tangible Interfaces for Remote Collaboration and Communication*. In Proceedings of ACM CSCW'98. 1998. Seattle, WA: ACM Press: pp. 169-178.
27. Brewster, S.A., *Using Non-Speech Sound to Overcome Information Overload*. Displays, Special issue on multimedia displays, 1995. 17: pp. 179-189.
28. Brewster, S.A., Wright, P.C., Dix, A.J., and Edwards, A.D.N. *The Sonic Enhancement of Graphical Buttons*. In Proceedings of Interact '95. 1995. Lillehammer, Norway: Chapman and Hall, London: pp. 43-48.
29. Brooks, F.P.J., Ouh-Young, M., Batter, J.J., and Kilpatrick, P.J. *Project GROPE - Haptic Displays for Scientific Visualization*. In Proceedings of SIGGRAPH '90. 1990. Dallas, TX: ACM Press: pp. 177-185.
30. Brown, P., Jones, D., Singh, S., and Rosen, J. *The Exoskeleton Glove for Control of Paralyzed Hands*. In Proceedings of IEEE International Conference on Robotics and Automation. 1993. Atlanta, GA: IEEE: pp. 642-647.
31. Burdea, G., *Force and Touch Feedback for Virtual Reality*. 1996, New York, NY: John Wiley and Sons.
32. Burdea, G., Dishpande, B., Liu, B., Langrana, N., and Gomez, D., *A Virtual Reality Based System for Hand Rehabilitation and Diagnosis*. Presence: Teleoperators and Virtual Environments, 1996. 6(2): pp. 229-240.
33. Burgin, J., Stephens, B., Vahora, F., Temkin, B., Marcy, W., P.J., G., and Krummel, T.M. *Haptic Rendering of Volumetric Soft-Bodies Objects*. In Proceedings of Third Phantom User's Group Workshop. 1998. Cambridge, MA
34. Buxton, W., *There's More to Interaction than Meets the Eye: Some Issues in Manual Input*, in *User Centered System Design: New Perspectives on Human-Computer Interaction*, D.A. Norman and S.W. Draper, Eds. 1986, Lawrence Erlbaum Associates: Hillsdale, New Jersey. pp. 317-337.

35. Byers, J.C., Bittner, A.C., and Hill, S.G., *Traditional and raw task load index (TLX) correlations: Are paired comparisons necessary?*, in *Advances in industrial ergonomics*, A. Mital, Ed. 1989: Taylor & Francis. pp. 481-485.
36. Cadiere, G.B., Himpens, J., Vertruyen, M., and Favretti, F., *The world first obesity surgery performed by a surgeon at a distance*. *Obesity Surgery*, 1999. 9(2): pp. 206-209.
37. Callahan, J., Hopkins, D., Weiser, M., and Shneiderman, B. *An empirical comparison of pie vs. linear menus*. In *Proceedings of ACM CHI'98*. 1998. Los Angeles, CA: ACM Press: pp. 95-100.
38. Campbell, C.S., Zhai, S., May, K.W., and Maglio, P.P. *What you Feel Must be What you See*. In *Proceedings of Interact'99*. 1999. Edinburgh, UK: IOS Press: pp. 383-390.
39. Challis, B.P. and Edwards, A.D.N., *Design Principles for Tactile Interaction*. *Haptic Human-Computer Interaction*, 2001. Springer-Verlag LNCS 2058: pp. 17-24.
40. Chang, A., O'Modhrain, S., Gunther, J.R., Gunther, E., and Ishii, H. *ComTouch: Design of a Vibrotactile Communication Device*. In *Proceedings of Designing Interactive Systems (ACM DIS '02)*. 2002. London, UK: ACM Press: pp. 312-320.
41. Chen, E. *Six Degree-of-Freedom Haptic System for Desktop Virtual Prototyping Applications*. In *Proceedings of First International Workshop on Virtual Reality and Prototyping*. 1999. Laval France: pp. 97-106.
42. Cheshire, D.G., Evans, M.A., and Dean, C.J. *Haptic Modelling - An Alternative to Industrial Design Methodology*. In *Proceedings of EuroHaptics 2001*. 2001. Birmingham, UK: University of Birmingham: pp. 124-128.
43. Chin, J.P., Diehl, V.A., and Norman, K.L. *Development of an Instrument Measuring User Satisfaction of the Human-Computer Interface*. In

- Proceedings of *ACM CHI'88*. 1988. Washington DC: ACM Press: pp. 213-218.
44. Chodack, J. and P., S. *Spacesuit glove thermal micrometeoroid garment protection versus human factors design parameters*. In Proceedings of *SAE Conference on Space Station and Advanced EVA*. 1991: Society of Automotive Engineers: pp. 67-83.
 45. Cohen, J., *Monitoring Background Activities*. Auditory Display: Sonification, Audification, and Auditory Interfaces, Proceedings of the First International Conference on Auditory Display, 1992: pp. 499-531.
 46. Cohen, J. *Out to Lunch: Further Adventures Monitoring Background Activity*. In Proceedings of *The Second International Conference on Auditory Display*. 1994. Santa Fe, NW: Addison Wesley: pp. 15-20.
 47. Colwell, C., Petrie, H., Kornbrot, D., Hardwick, A., and Furner, S. *Haptic Virtual Reality for Blind Computer Users*. In Proceedings of *ACM ASSETS '98*. 1998. Marina del Rey, CA: ACM Press: pp. 92-99.
 48. Coolican, H., *Research Methods and Statistics in Psychology*. 1994, London, UK: Hodder and Stoughton.
 49. Cooper, B. *Driving on the Surface of Mars using the Rover Control Workstation*. In Proceedings of *SpaceOps'98*. 1998. Tokyo, Japan: NASDA
 50. Coulouris, G., Dollimore, J., and Kindberg, T., *Distributed systems: concepts and design*. 3rd edition. 2000, Boston, MA: Addison-Wesley.
 51. Dennerlein, J.T., Martin, D.B., and Hasser, C. *Force-feedback improves performance for steering and combined steering-targeting tasks*. In Proceedings of *ACM CHI'2000*. 2000. The Hague, NL: ACM Press: pp. 423-429.
 52. Dennerlein, J.T. and Yang, M.C., *Haptic force-feedback devices for the office computer: Performance and musculoskeletal loading issues*. *Human Factors*, 2001. 43(2): pp. 278-286.

53. Deutsch, J., Latonio, J., Burdea, G., and Boian, R., *Post-Stroke Rehabilitation with the Rutgers Ankle System - A case study*. Presence: Teleoperators and Virtual Environments, 2001. 10(4): pp. 416-430.
54. DiFranco, D., Beauregard, G.L., and Srinivasan, M.A. *The effect of auditory cues on the haptic perception of stiffness in virtual environments*. In Proceedings of *ASME Dynamic Systems and Control Division*. 1997. Dallas, TX: ASME: pp. 17 - 21.
55. Dodge, C. *The Bed: A Medium for Intimate Communication*. In Proceedings of *ACM CHI'97*. 1997. Atlanta, GA: ACM Press: pp. 371-372.
56. Doherty-Sneddon, G., Anderson, A., O'Malley, C., Langton, S., Garrod, S., and Bruce, V., *Face-to-Face and Video-Mediated Communication: A Comparison of Dialog Structure and Task Performance*. Journal of Experimental Psychology: Applied, 1997. 3(2): pp. 105-125.
57. Dourish, P. and Bellotti, V. *Awareness and Coordination in Shared Workspaces*. In Proceedings of *ACM CSCW'92*. 1992. Toronto, Canada: ACM Press: pp. 107-114.
58. Dourish, P., Bellotti, V., Mackay, W., and Ma, C. *Information and context: lessons from the study of two shared information systems*. In Proceedings of *ACM Conference on Organisational Computing Systems (COOCS'93)*. 1993. Milpetas, CA: ACM Press: pp. 42-51.
59. Dubrovsky, V.J., Kiesler, S., and Sethna, B.N., *The Equalization Phenomenon: status effects in computer mediated and face-to-face decision making groups*. Human Computer Interaction, 1991. 6: pp. 119-146.
60. Durbeck, L.J.K., Macias, N.J., Weinstein, D.M., Johnson, C.R., and Hollerbach, J.M. *SCIRun Haptic Display for Scientific Visualisation*. In Proceedings of *Third Phantom User's Group Workshop*. 1998. Cambridge, MA

61. Durlach, N. and Slater, M., *Presence in Shared Virtual Environments and Virtual Togetherness*. Presence: Teleoperators and Virtual Environments, 2000. 9(2): pp. 214-217.
62. Durlach, N.I., Delhorne, L.A., Wong, A., Ko, W.Y., Rabinowitz, W.M., and Hollerbach, J., *Manual discrimination and identification of length by the finger-span method*. Perception and Psychophysics, 1989. 46(1): pp. 29-38.
63. Eagly, A.H. and Chaiken, S., *The Psychology of Attitudes*. 1993, Orlando, FL: Harcourt Brace.
64. Ellis, R.E., Ismaeil, O.M., and Lipsett, M.G., *Design and Evaluation of a High Performance Haptic Interface*. Robotica, 1996. 14: pp. 321-327.
65. Engel, F.L., Goossens, P., and Haakma, R., *Improved Efficiency through I- and E-Feedback: A Trackball with Contextual Force Feedback*. International Journal of Human-Computer Studies, 1994. 41(6): pp. 949-974,.
66. Fisher, J.D., Rytting, M., and Heslin, R., *Hands Touching Hands: Affective and Evaluative Effects of an Interpersonal Touch*. Sociometry, 1976. 39(4): pp. 416-421.
67. Fogg, B.J., Cutler, L.D., Arnold, P., and Eisbach, C. *HandJive: A Device for Interpersonal Haptic Entertainment*. In Proceedings of ACM CHI'98. 1998. Los Angeles, CA: ACM Press: pp. 57-64.
68. Fong, T. and Thorpe, C., *Vehicle Teleoperation Interfaces*. Autonomous Robots, 2001. 11(1): pp. 9-18.
69. Force Dimension. *Delta Haptic Device*, 2002, <http://www.forcedimension.com/products/>
70. Frank, L.K., *Tactile Communication*. Genetic Psychology Monographs, 1957. 56(2): pp. 209-255.
71. Fulke, E., *Reading Braille*, in *Tactile Perception: A Sourcebook*, W. Schiff and E. Foulke, Eds. 1982, Cambridge University Press: Cambridge, UK.

72. Gaver, W., Moran, T., MacLean, A., Lovstand, L., Dourish, P., Carter, K., and Buxton, W. *Realizing a Video Environment: EuroParc's RAVE System*. In Proceedings of *ACM CHI'92*. 1992. Monterey, CA: ACM Press: pp. 27-35.
73. Gaver, W., Smith, R.B., and O'Shea, T. *Effective Sounds in Complex Systems: the Arkola Simulation*. In Proceedings of *ACM CHI'91*. 1991. New Orleans, LO: ACM Press: pp. 85-90.
74. Gibson, J.J., *Observations on Active Touch*. *Psychological Review*, 1962. 69(6): pp. 477-491.
75. Gleitman, H., *Psychology*. 1991, New York: Norton & Company.
76. Gorman, P.J., Lieser, J.D., Murray, W.B., Haluck, R.S., and Krummel, T.M. *Assessment and validation of a force feedback virtual reality based surgical simulator*. In Proceedings of *Third Phantom User's Group Workshop*. 1998. Cambridge, MA
77. Graham, J.A. and Argyle, M., *A cross-cultural study of the communication of extra-verbal meaning by gestures*. *International Journal of Psychology*, 1975. 10(1): pp. 57-67.
78. Graham, T.C.N., Stewart, H.D., Ryman, A.G., Kopae, A.R., and Rasouli, R. *A World-Wide-Web Architecture for Collaborative Software Design*. In Proceedings of *Software Technology and Engineering Practice (STEP'99)*. 1999. Pittsburgh, PA: IEEE Press
79. Gray, H., *Anatomy of the Human Body*. 20th edition, Ed W.H. Lewis. 1918, Philadelphia: Lea & Febiger.
80. Greenberg, S., Roseman, M., and Webster, D. *Issues and Experiences Designing and Implementing Two Group Drawing Tools*. In Proceedings of *Hawaii International Conference on System Sciences*. 1992. Kuwahi, Hawaii: IEEE Press: pp. 149-150.
81. NASA Ames Research Centre. *Task Load Index (NASA-TLX) v1.0 computerised version, 1987*.

82. Grudin, J. *Why CSCW applications fail: problems in the design and evaluation of organization of organizational interfaces*. In Proceedings of ACM CSCW'88. 1998. Portland, OR: ACM Press: pp. 65-84.
83. Gutwin, C. and Penner, R. *Improving Interpretation of Remote Gestures with Telepointer Traces*. In Proceedings of ACM CSCW'02. 2002. New Orleans, LO: ACM Press: pp. 49-57.
84. Gutwin, C., Roseman, M., and Greenberg, S. *A Usability Study of Awareness Widgets in a Shared Workspace Groupware System*. In Proceedings of ACM CSCW'96. 1996. Boston, MA: ACM Press: pp. 258-267.
85. Hall, J.A. and E.M., V., *More "Touching" Observations: New Insights on Men, and Interpersonal Touch*. Journal of Personality and Social Psychology, 1990. 59(6): pp. 1155-1162.
86. Hammersley, M. and Atkinson, P., *Ethnography: Principles in Practice*. 1983, London, UK: Routledge.
87. Harlow, H.F., *The Nature of Love*. American Psychology, 1958. 13: pp. 673-685.
88. Hart, S. and Staveland, L., *Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research*, in *Human mental workload*, N. Meshkati, Ed. 1988: Amsterdam, NL. pp. 139-183.
89. Hart, S.G. and Wickens, C., *Workload assessment and prediction*, in *MANPRINT, an approach to systems integration*, H.R. Booher, Ed. 1990, Van Nostrand Reinhold.: New York. pp. 257-296.
90. Hasser, C., Goldenburg, A., Martin, K., and Rosenberg, L. *User Performance in a GUI Pointing Task with a Low-Cost Force-Feedback Computer Mouse*. In Proceedings of *Seventh Annual Symposium on Haptic Interfaces*. 1998. Anaheim
91. Henley, N.M., *Status and Sex: Some Touching Observations*. Bulletin of the Psychonomic Society, 1973. 2: pp. 91-93.

92. Henninger, S., Haynes, K., and Reith, M.W. *A Framework for Developing Experience-Based Usability Guidelines*. In Proceedings of *Designing Interactive Systems (ACM DIS'95)*. 1995. Ann Arbor, MI: ACM Press: pp. 43-53.
93. Holland, B.S. and Copenhaver, M., *Improved Bonferroni-type multiple testing procedures*. Psychological Bulletin, 1998. 104: pp. 145-149.
94. Hwang, F. *A study of cursor trajectories of motion-impaired users*. In Proceedings of *ACM CHI'02*. 2002. Minneapolis, MI: ACM Press: pp. 842-843.
95. Hwang, F., Keates, S., Langdon, P., Clarkson, P.J., and Robinson, P. *Perception and Haptics Towards More Accessible Computers for Motion Impaired Users*. In Proceedings of *Perceptual User Interfaces (ACM PUI'01)*. 2001. Orlando, FL: ACM Press: pp. 1-9.
96. Immersion Corporation. *3D Technologies CyberForce, 2002*, <http://www.immersion.com/products/3d/interaction/cyberforce.shtml>
97. Immersion Corporation. *3D Technologies CyberGrasp, 2002*, <http://www.immersion.com/products/3d/interaction/cybergrasp.shtml>
98. Immersion Corporation. *Gaming Joysticks, 2002*, <http://www.immersion.com/products/ce/gamingjoystick.shtml>
99. Immersion Corporation. *Gaming Mice, 2002*, <http://www.immersion.com/products/ce/gamingmice.shtml>
100. Immersion Corporation. *Immersion Medical - Laparoscopy, 2002*, http://www.immersion.com/products/medical/laparoscopy_hardware.shtml
101. *Pulse Data International - Notetakers, 2002*, <http://www.brailnote.com>
102. Ishii, H. *TeamWorkStation: Towards a Seamless Shared Workspace*. In Proceedings of *ACM CSCW'90*. 1990. Los Angeles, CA: ACM Press: pp. 13-26.

103. Ishii, H., Koyashi, M., and Grudin, J. *Integration of Inter-Personal Space and Shared Workspace: ClearBoard Design and Experiments*. In Proceedings of *ACM CHI'92*. 1992. Monterey, CA: ACM Press: pp. 33-42.
104. Ishii, H. and Koyayashi, M. *ClearBoard: A Seamless Medium for Shared Drawing and Conversation with Eye Contact*. In Proceedings of *ACM CHI'92*. 1992. Monterey, CA: ACM Press: pp. 525-532.
105. Jansson, G., *Can the Efficiency of a Haptic Display be Increased by Short Term Practice at Exploration?* Haptic Human Computer Interaction, 2001. Springer-Verlag LNCS 2058: pp. 88-97.
106. Johansen, R., Martin, A., Mittman, R., and Saffo, P., *Leading Business Teams*. 1991, Reading, MA: Addison-Wesley.
107. Johnson, K.O. and Phillips, J.R., *Tactile spatial resolution - I. Two point discrimination, gap detection, grating resolution and letter recognition*. *Journal of Neurophysiology*, 1981. 46(6): pp. 1177-1191.
108. Keates, S., Hwang, F., Langdon, P., Clarkson, P.J., and Robinson, P. *Cursor measures for motion-impaired computer users*. In Proceedings of *ACM Assets*. 2002. Edinburgh, UK: ACM Press: pp. 135-142.
109. Keuning, H. and Houtsma, A.J.M. *Cursor Displacement and Velocity Profiles for Targets in Various Locations*. In Proceedings of *Eurohaptics 2001*. 2001. Birmingham, UK: University of Birmingham: pp. 108-112.
110. Keuning, H. and Houtsma, A.J.M., *Cursor Trajectory Analysis*. Haptic Human Computer Interaction, 2001. Springer-Verlag LNCS 2058: pp. 127-134.
111. Klatzky, R.L., Lederman, S.J., and Metzger, V., *Identifying Objects by Touch: An 'Expert System'*. *Perception and Psychophysics*, 1985. 37: pp. 299-302.

112. LaMotte, R.H. and Srinivasan, M.A. *Surface microgeometry: Neural encoding and perception*. In *Proceedings of Information Processing in the Somatosensory System*. 1991: Macmillan Press
113. Lederman, S.J. and Klatzky, R.L., *Hand Movements: A Window into Haptic Object Recognition*. *Cognitive Psychology*, 1987. 19: pp. 342-368.
114. Lederman, S.J. and Klatzky, R.L., *Extracting Object Properties through Haptic Exploration*. *Acta Psychologica*, 1993. 84: pp. 29-40.
115. Lederman, S.J., Klatzky, R.L., and Barber, P.O., *Spatial and movement-based heuristics for encoding pattern information through touch*. *Journal of Experimental Psychology: General*, 1985. 114: pp. 33-49.
116. Lessiter, J., Freeman, J., Keogh, E., and Davido, J., *A Cross-Media Presence Questionnaire: The ITC Sense of Presence Inventory*. *Presence: Teleoperators and Virtual Environments*, 2001. 10(3): pp. 282-297.
117. Lewis, J.R., *IBM Computer Usability Satisfaction Questionnaires: Psychometric Evaluation and Instructions for Use*. *International Journal of Human-Computer Interaction*, 1995. 7(1): pp. 57-75.
118. Logitech. *WingMan Force Feedback Mouse, 2000*, <http://www.logitech.com/cf/products/productoverview.cfm/30>
119. Loomis, J.M., *An investigation of tactile hyperacuity*. *Sensory Processes*, 1979. 3: pp. 289-302.
120. Loomis, J.M., *Distal attribution and presence*. *Presence: Teleoperators and Virtual Environments*, 1992. 1(1): pp. 113-119.
121. Lumsden, J., Brewster, S.A., Crease, M., and Gray, P.D. *Guidelines for audio-enhancement of graphical user interface widgets*. In *Proceedings of British HCI 2002*. 2002. London, UK: BCS: pp. 6-9.
122. Mackenzie, L.M., *Communications and Networks*. 1998: McGraw-Hill.

123. MacKenzie, S. and Ware, C. *Lag as a determinant of human performance in interactive systems*. In Proceedings of ACM CHI'93. 1993. Amsterdam, NL: ACM Press: pp. 488-293.
124. MacLean, K.E., Shaver, M.J., and Pai, D.K. *Handheld Haptics: A USB Media Controller with Force Sensing*. In Proceedings of IEEE VR2002, 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. 2002. Orlando, FL: IEEE: pp. 311-319.
125. Major, B., Schmidlin, A.M., and Williams, L., *Gender Patterns in Social Touch: The Impact of Setting and Age*. Journal of Personality and Social Psychology, 1990. 58(4): pp. 634-643.
126. Mantei, M.M., Baecker, R.M., Sellen, A.J., Buxton, W.A.S., and Milligan, T. *Experiences in the Use of a Media Space*. In Proceedings of ACM CHI'91. 1991. New Orleans, LO: ACM Press: pp. 203-208.
127. Mantovani, G., *Virtual Reality as a Communication Environment: Consensual Hallucination, Fiction and Possible Selves*. Human Relations, 1995. 48(6): pp. 669-683.
128. Massie, T.H., *Initial haptic explorations with the PHANToM: Virtual touch through pointer interaction*, Work submitted for Masters of Science, Massachusetts Institute of Technology: Boston, MA, 1996.
129. McGee, M.R., *The Cognitive Impact of Multimodal Interfaces*. 1999, Technical Report for first year of Ph.D. thesis, Department of Computing Science, University of Glasgow.: Glasgow.
130. McGee, M.R., Gray, P.D., and Brewster, S.A. *Mixed Feelings: Multimodal Perception of Virtual Roughness*. In Proceedings of Eurohaptics 2002. 2002. Edinburgh, UK: University of Edinburgh: pp. 47-52.
131. McGrath, J., *Methods for the Study of Groups*, in *Readings in Groupware and Computer Supported Cooperative Work*, R. Baecker, Ed. 1993, Morgan Kaufmann: San Francisco, CA. pp. 200-204.

132. McGrath, J., *Methodology Matters*, in *Human-Computer Interaction: Towards the Year 2000*, R. Baecker, Ed. 1995, Morgan Kaufmann: San Francisco, CA. pp. 152-169.
133. McGuffin, L. and Olson, G., *ShrEdit: A Shared Electronic Workspace*. 1992, Cognitive Science and Machine Intelligence Laboratory, University of Michigan: Michigan.
134. *Netmeeting Home Page, 2002*,
<http://www.microsoft.com/windows/netmeeting/>
135. Miller, T. and Zeleznik, R. *An Insidious Haptic Invasion: Adding Force Feedback to the X Desktop*. In *Proceedings of ACM UIST'98*. 1998. San Francisco, CA: ACM Press: pp. 59-64.
136. Miller, T. and Zeleznik, R. *The design of 3D haptic widgets*. In *Proceedings of Symposium on Interactive 3D graphics*. 1999. Atlanta, GA: ACM Press
137. Minsky, M., Ouh-Young, M., Teele, O., Brooks, F.P., and Behensky, M., *Feeling and Seeing: Issues in Force Display*. Special Issue on the 1990 symposium of 3D graphics, 1990. 24(2): pp. 235-243.
138. Monk, A., Nardi, B., Gilbert, N., Mantei, M., and McCarthy, J. *Mixing oil and water?: Ethnography versus experimental psychology in the study of computer-mediated communication*. In *Proceedings of ACM CHI'93*. 1993. Amsterdam, NL: ACM Press: pp. 3-6.
139. Monk, A.F. and Watts, L. *A Poor Quality Video Links Affects Speech but not Gaze*. In *Proceedings of ACM CHI'95*. 1995. Denver, CO: ACM Press: pp. 274-275.
140. Munch, S. and Dillmann, R. *Haptic Output in Multimodal User Interfaces*. In *Proceedings of ACM IUI'97*. 1997. Orlando, FL: ACM Press: pp. 105-112.
141. Mynatt, E.D., Back, M., Want, R., Baer, M., and Ellis, J.B. *Designing audio aura*. In *Proceedings of ACM CHI'98*. 1998. Los Angeles, CA: ACM Press: pp. 566-573.

142. Mynatt, E.D., Back, M., Want, R., and Frederick, R. *Audio aura: light-weight audio augmented reality*. In Proceedings of *ACM UIST'97*. 1997. Banff, Alberta, Canada: ACM Press: pp. 211-212.
143. Nelson, D. and Muthukrishnan, S. *Design issues in multimedia messaging for next generation wireless systems*. In Proceedings of *Second ACM International Workshop on Data Engineering for Wireless and Mobile Access*. 2001. Santa Barbara, CA: ACM Press: pp. 98-103.
144. Newell, A. and Rosenbloom, P.S., *Mechanisms of Skill Acquisition and the Law of Practice*, in *Cognitive Skills and their Acquisition*, J.R. Anderson, Ed. 1980, Lawrence Erlbaum & Associates.: Hillsdale, N.J.
145. Noma, H. and Miyasato, T. *Cooperative Object Manipulation in Virtual using Virtual Physics*. In Proceedings of *ASME Annual Winter Meeting, Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. 1997. Dallas, TX: pp. 101-106.
146. Norman, D.A., *The Design of Everyday Things*. 1998, Cambridge, MA: MIT Press.
147. Oakley, I., Adams, A., Brewster, S.A., and Gray, P.D. *Guidelines for the Design of Haptic Widgets*. In Proceedings of *BCS HCI'2002*. 2002. London, UK: Springer: pp. 195-212.
148. Oakley, I., Brewster, S.A., and Gray, P.D. *Can you feel the force? An Investigation of Haptic Collaboration in Shared Editors*. In Proceedings of *EuroHaptics 2001*. 2001. Birmingham, UK: University of Birmingham: pp. 54-59.
149. Oakley, I., Brewster, S.A., and Gray, P.D. *Communicating with Feeling*. In Proceedings of *Haptic Human Computer Interaction*. 2001. Glasgow, UK: Springer-Verlag: pp. 61-68.
150. Oakley, I., Brewster, S.A., and Gray, P.D. *Solving Multi-Target Haptic Problems in Menu Interaction*. In Proceedings of *ACM CHI'01*. 2001. Seattle, WA: ACM Press: pp. 357-358.

151. Oakley, I., McGee, M.R., Brewster, S.A., and Gray, P.D. *Putting the feel into 'look and feel'*. In Proceedings of *ACM CHI'2000*. 2000. The Hague, NL: ACM Press: pp. 415-422.
152. Olson, J.S. and Teasley, S. *Groupware in the Wild: Lessons Learned from a Year of Virtual Collocation*. In Proceedings of *ACM CSCW'96*. 1996. Boston, MA: ACM Press: pp. 419-427.
153. Pallas-Areny, R. and Webster, J.G., *Sensors and Signal Conditioning*. 2000, New York, NY: John Wiley and Sons.
154. Partrick, N., *Design, Construction, and Testing of a Fingertip Tactile Display for Interaction with Virtual and Remote Environments*, Work submitted for Masters of Science in Department of Mechanical Engineering, MIT: Cambridge, MA, 1990.
155. Pedersen, E.R. and Sokoler, T. *AROMA: abstract representation of presence supporting mutual awareness*. In Proceedings of *ACM CHI'97*. 1997. Atlanta, GA: ACM Press: pp. 51-58.
156. Pederson, E.R., McCall, K., Moran, T.P., and Halasz, F.G. *Tivoli: An Electronic Whiteboard for Informal Workgroup Meetings*. In Proceedings of *ACM CHI'93*. 1993. Amsterdam, NL: ACM Press: pp. 391-398.
157. Ramstein, C. *A Multimodal User Interface System with Force Feedback and Physical Models*. In Proceedings of *Interact'95*. 1994. Lillehammer, Norway: Chapman and Hall, London
158. Ramstein, C. and Hayward, V. *The PANTOGRAPH: A Large Workspace Haptic Device for a Multi-Modal Human-Computer Interaction*. In Proceedings of *ACM CHI'94*. 1994. Boston, MA: ACM Press: pp. 57-58.
159. Ramstein, C., Martial, O., Dufresne, A., Carignan, M., Chassé, P., and Mabileau, P. *Touching and hearing GUI's: design issues for the PC-Access system*. In Proceedings of *ACM ASSETS '96*. 1996. Vancouver, BC: ACM Press: pp. 2-9.

160. Raymaekers, C. and Coninx, K. *Menu Interactions in a Desktop Haptic Environment*. In Proceedings of *EuroHaptics 2001*. 2001. Birmingham, UK: University of Birmingham: pp. 49-53.
161. Roseman, M. and Greenberg, S., *Building Real Time Groupware with GroupKit, A Groupware Toolkit*. ACM Transactions of Computer-Human Interaction, 1996. 3(1): pp. 66-106.
162. Rosenberg, L. and Brave, S. *Using Force Feedback to Enhance Human Performance in Graphical User Interfaces*. In Proceedings of *ACM CHI'96*. 1996. Vancouver, BC: ACM Press: pp. 291-292.
163. Russo, M. and Tadros, A. *Controlling Dissipative Magnetic Particle Brakes in Force Reflective Devices*. In Proceedings of *Advances in Robotics*. 1992. New York, NY: ASME
164. Sakuma, M., Masamori, S., and Harada, T. *A System for Japanese Calligraphy Lesson with Force Feedback on Internet*. In Proceedings of *ASME Dynamic Systems and Control Division*. 1999: ASME: pp. 429-434.
165. Salisbury, J.K. and Srinivasan, M.A., *Phantom-Based Haptic Interaction with Virtual Objects*. IEEE Computer Graphics and Applications, 1997. 17(5).
166. Sallnas, E., Rasmus-Grohn, K., and Sjostrom, C., *Supporting Presence in Collaborative Environments by Haptic Force Feedback*. ACM Transactions on CHI, 2000. 7(4): pp. 461-476.
167. Sanders, M.S. and McCormick, E.J., *Human Factors in Engineering and Design*. 1987, New York: McGraw-Hill Book Co.
168. Sato, K., Igarashi, E., and Kimura, M. *Development of a Non-constrained Master-arm with Vibrotactile feedback device*. In Proceedings of *International Conference on Advanced Robotics*. 1991. Pisa, Italy: pp. 334-338.

169. Scheifler, R.W. and Gettys, J., *X Window System*. 3rd edition. 1992, Burlington, MA: Digital Press.
170. Shen, H. and Dewan, P. *Access Control for Collaborative Environments*. In Proceedings of *ACM CSCW'92*. 1992. Toronto, Canada: ACM Press: pp. 51-58.
171. Shennib, H., Bastawisy, A., McLoughlin, J., and Moll, F., *Robotic computer-assisted telemanipulation enhances coronary artery bypass*. *Journal of Thoracic Cardiovascular Surgery*, 1999. 117(2): pp. 310-316.
172. Shimoga, K.B. *A Survey of Perceptual Feedback Issues in Dexterous Telemanipulation: Part 2. Finger Touch Feedback*. In Proceedings of *IEEE Virtual Reality Annual International Symposium (VRAIS'93)*. 1993. Seattle, WA: IEEE: pp. 271-279.
173. Shneiderman, B. *Direct Manipulation for Comprehensible, Predictable, and Controllable User Interfaces*. In Proceedings of *ACM IUI'97*. 1997. Orlando, FL: ACM Press: pp. 33-39.
174. Siira, J.O. and Pai, D.K. *Fast haptic textures*. In Proceedings of *ACM CHI'96*. 1996. Vancouver, BC: ACM Press: pp. 231-232.
175. Silvermintz, L., *Sensor Application to Functional Neuro-Muscular Stimulation*, in *Tactile Sensors for Robotics and Medicine*, J.G. Webster, Ed. 1988, John Wiley and Sons: New York, NY. pp. 309-330.
176. Snibbe, S., MacLean, K.E., Shaw, R., Roderick, J., Verplank, W.L., and Scheeff, M. *Haptic Techniques for Media Control*. In Proceedings of *ACM UIST'01*. 2001. Orlando, FL: ACM Press: pp. 199-208.
177. Sproull, L. and Keisler, S., *Connections: New ways of Working in Networked Organisations*. 1991, Cambridge, MA: MIT Press.
178. Srinivasan, M.A. and Basdogan, C., *Haptics in Virtual Environments: Taxomy, Research Status and Challenges*. *Computers and Graphics - Special Issue on Haptic Displays in Virtual Environments*, 1997. 21(4).

179. Srinivasan, M.A., Beauregard, G.L., and Brock, D.L. *The impact of visual information on haptic perception of stiffness in virtual environments*. In Proceedings of *ASME Dynamic Systems and Control Division*. 1996. Atlanta, GA: ASME: pp. 555-559.
180. Srinivasan, M.A. and Chen, J.S. *Human performance in controlling normal forces of contact with rigid objects*. In Proceedings of *Advances in Robotics, Mechatronics, and Haptic Interfaces*. 1993. New Orleans: ASME
181. Stefik, M., Bobrow, D.G., Foster, G., Lanning, S., and Tatar, D., *WYSIWIS revised: early experiences with multiuser interfaces*. *ACM Transactions on Information Systems (TOIS)*, 1987. 5(2): pp. 147-167.
182. Stevens, P. and Pooley, R., *Using UML*. 1999: Addison-Wesley.
183. Strong, R. and Gaver, B. *Feather, Scent & Shaker*. In Proceedings of *ACM CSCW'96*. 1996. Boston, MA: ACM Press: pp. 363-364.
184. Suchman, L.A. and Trigg, R.H., *Understanding Practice: Video as a Medium for Reflection and Design*, in *Readings in Groupware and Computer Supported Cooperative Work*, R. Baecker, Ed. 1991, Morgan Kaufmann: San Francisco, CA. pp. 205-217.
185. Tan, D.S., Pausch, R., Stefanucci, J.K., and Proffitt, D.R. *Kinesthetic cues aid spatial memory*. In Proceedings of *ACM CHI'02*. 2002. Minneapolis, MI: ACM Press: pp. 806-807.
186. Tan, H.Z., Durlach, N., Rabinowitz, W., and Reed, C.M. *Information Transmission with a Multi-finger Tactual Display*. In Proceedings of *International Sensory Aid Conference*. 1996
187. Tan, H.Z. and Pentland, A. *Tactual displays for wearable computing*. In Proceedings of *First International Symposium on Wearable Computers*. 1997. Cambridge, MA: pp. 84-89.
188. Tan, H.Z., Srinivasan, M.A., Eberman, B., and Cheng, B. *Human factors for the design of force-reflecting haptic interfaces*. In Proceedings of *ASME*

- Dynamic Systems and Control Division*. 1994. Chicago, IL: ASME: pp. 353-359.
189. Tang, J.C. and Minneman, S.L., *VideoDraw: A Video Interface for Collaborative Drawing*. ACM Transactions on Information Systems, 1991. 9(2): pp. 170-184.
 190. Tetzlaff, L. and Schwartz, D.R. *The Use of Guidelines in User Interface Design*. In Proceedings of *ACM CHI'91*. 1991. New Orleans, LO: ACM Press: pp. 329-333.
 191. Twidale, M., Randall, D., and Bentley, R. *Situated Evaluation for Cooperative Systems*. In Proceedings of *ACM CSCW'94*. 1994. Chapel Hill, NC: ACM Press: pp. 441-452.
 192. Van der Linde, R.Q., Lammertse, P., Frederiksen, E., and Ruiters, B. *The Haptic Master, a New High Performance Haptic Interface*. In Proceedings of *EuroHaptics 2002*. 2002. Edinburgh, UK: University of Edinburgh: pp. 1-5.
 193. van Erp, J.B.F., *Tactile Navigation Display*. Haptic Human Computer Interaction, 2001. Springer-Verlag LNCS 2058: pp. 165-173.
 194. Verrillo, R.T. and Gescheider, G.A., *Perception via the Sense of Touch*, in *Tactile Aids for the Hearing Impaired*, I. Summers, Ed. 1992, Whurr Publishers.
 195. Wall, S.A. and Harwin, W.S., *Interaction of Visual and Haptic Information in Simulated Environments: Texture Perception*. Haptic Human Computer Interaction, 2001. Springer-Verlag LNCS 2058: pp. 39-44.
 196. Wang, X., Seet, G.G.L., Lau, M.W.S., Low, E., and Tan, K.C. *Exploiting Force Feedback in Pilot Training and Control of an Underwater Robotics Vehicle: an Implementation in LabVIEW*. In Proceedings of *OCEANS 2000 MTS/IEEE Conference and Exhibition*. 2000. Providence, RI: IEEE: pp. 2037-2042.

197. Wanger, L. *Haptically Enhanced Molecular Modeling : a case study*. In Proceedings of *Third PHANTOM Users Group Workshop*. 1998. Cambridge, MA
198. Webster, M., *Webster's Ninth New Collegiate Dictionary*. 1985, Springfield, MA: Merriam-Webster Inc.
199. Welch, R.B. and Warren, D.H., *Intersensory interactions*, in *Handbook of Perception and Human Performance*, K. Boff, L. Kaufman, and J. Thomas, Eds. 1986, Wiley and Sons.
200. White, N. and Back, D. *Telephonic Arm Wrestling, 1986*, <http://www.bmts.com/~normill/artpage.html>
201. Wilson, J.P., Kline-Schoder, R.J., Kenton, M.A., and Hogan, N. *Algorithms for Network-Based Force Feedback*. In Proceedings of *PHANTOM's User Group (PUG'99)*. 1999. Boston, MA
202. Wright, P.C. and Monk, A.F. *Evaluation for design*. In Proceedings of *People and Computers*. 1989. Cambridge: Cambridge University Press: pp. 345-358.
203. Yokokohji, Y., Hollis, R., Kanade, T., Henmi, K., and Yoshikawa, T. *Toward Machine Mediated Training of Motor Skills - Skill Transfer from Human to Human via Virtual Environment*. In Proceedings of *5th IEEE International Workshop on Robot and Human Communication*. 1996: IEEE: pp. 32-37.
204. Yu, W. and Brewster, S.A. *Comparing Two Haptic Interfaces for Multimodal Graph Rendering*. In Proceedings of *IEEE VR2002, 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. 2002. Orlando, FL: IEEE
205. Yu, W., Ramloll, R., and Brewster, S.A., *Haptic graphs for blind computer users*. *Haptic Human-Computer Interaction*, 2001. Springer-Verlag LNCS 2058: pp. 41-51.

206. Zachmann, G. *Minimal hierarchical collision detection*. In Proceedings of *Virtual Reality Software and Technology*. 2002. Hong Kong, China: ACM Press: pp. 121-128.
207. Zhai, S., Smith, B.A., and Selker, T. *Improving Browsing Performance: A study of four input devices for scrolling and pointing tasks*. In Proceedings of *Interact'97*. 1997. Sydney, Australia: IOS Press: pp. 286-292.
208. Zhu, H., *Electrotactile Stimulation*, in *Tactile Sensors for Robotics and Medicine*, J.G. Webster, Ed. 1988, John Wiley and Sons: New York, NY. pp. 341-354.

Appendix A. Materials for gathering TLX data

This appendix contains the definitions and forms handed out to experimental participants relating to the TLX questionnaire used in the studies described in this thesis. The definitions are taken from [ref].

Table A.1. Workload descriptions given to subjects when filling in the workload charts.

The descriptions are taken from [111]

Rating scale definitions		
Title	Endpoints	Description
Mental Demand	Low/High	How much mental, visual and haptic activity was required? (e.g. thinking, deciding calculating, looking, feeling listening, cross-monitoring, scanning, searching)
Physical Demand	Low/High	How much physical activity was required? (e.g. pushing, pulling, turning, controlling)
Time Pressure	Low/High	How much time pressure did you feel because of the rate at which things occurred? (e.g. slow, leisurely, rapid, frantic)
Effort Expended	Low/High	How hard did you work (mentally and physically) to accomplish your level of performance?
Performance Level Achieved	Poor/Good	How successful do you think you were in accomplishing the mission goals?
Frustration Experienced	Low/High	How successful do you think you were in accomplishing the mission goals?
Fatigue Experienced	Low/High	How much fatigue did you experience? (e.g., tiredness, strain, discomfort)

Table A.2. Form given out to subjects to capture TLX workload data.

Name : _____ Condition Completed : _____


Mental Demand




Physical Demand




Time Pressure



Effort Expended




Performance Level Achieved



Frustration Experienced



Fatigue Experienced



Appendix B. Questionnaire Administered in Buttons Study

This appendix contains a copy of the questionnaire administered after the first study described in Chapter 4 investigating haptically augmented buttons.

Questionnaire

Age: _____

Gender: M / F

Nationality: _____

Rate your level of computer experience (tick one):

None

Basic

Intermediate

Expert

Rate your order of preference for the five conditions (1 best, 5 worst):

Texture..... —

Friction..... —

Recess..... —

Gravity..... —

Visual..... —

Do not fill out this section

Condition:

Appendix C. Raw Data from Buttons Study

This appendix contains a copy of the raw data gathered during the first study described in Chapter 4 investigating haptically augmented buttons. It includes all subjective and objective data.

Table C.1. TLX data from gravity well condition.

	Mental Demand	Physical Demand	Time Pressure	Effort Expended	Performance Level Achieved	Frustration Experienced	Fatigue Experienced
Participant 1	5	3	7	7	15	3	3
Participant 2	3	10	4	12	14	10	14
Participant 3	10	10	8	12	4	6	8
Participant 4	0	0	2	0	2	2	2
Participant 5	7	3	4	3	9	3	2
Participant 6	3	8	4	3	6	1	1
Participant 7	8	10	16	8	4	9	10
Participant 8	12	12	13	11	4	9	11
Participant 9	15	13	11	11	10	18	17
Participant 10	1	3	3	3	5	3	4
Participant 11	3	5	7	7	13	4	5
Participant 12	7	11	8	10	10	8	8
Participant 13	6	6	6	5	6	5	3
Participant 14	2	4	6	6	6	1	2
Participant 15	7	14	11	14	6	6	14
Participant 16	5	9	9	10	12	9	10

Table C.2. TLX data from recess condition.

	Mental Demand	Physical Demand	Time Pressure	Effort Expended	Performance Level Achieved	Frustration Experienced	Fatigue Experienced
Participant 1	5	3	7	3	13	5	3
Participant 2	2	4	2	6	13	4	4
Participant 3	5	10	8	9	4	4	6
Participant 4	2	2	4	2	2	0	2
Participant 5	4	4	4	4	8	4	4
Participant 6	3	8	3	5	6	2	2
Participant 7	9	11	16	9	10	9	9
Participant 8	14	14	14	13	4	8	12
Participant 9	16	16	14	13	10	18	19
Participant 10	3	3	4	5	9	5	4
Participant 11	3	5	7	4	18	4	5
Participant 12	6	12	7	9	10	6	11
Participant 13	10	12	10	10	8	5	6
Participant 14	2	4	4	4	7	2	3
Participant 15	7	14	12	14	8	8	17
Participant 16	6	10	10	12	12	13	8

Table C.3. TLX data from friction condition.

	Mental Demand	Physical Demand	Time Pressure	Effort Expended	Performance Level Achieved	Frustration Experienced	Fatigue Experienced
Participant 1	3	3	5	3	13	5	4
Participant 2	8	7	0	8	10	8	6
Participant 3	8	11	9	10	5	7	6
Participant 4	6	6	6	6	8	6	2
Participant 5	10	8	4	9	8	10	4
Participant 6	3	7	6	6	6	2	2
Participant 7	10	18	17	17	10	10	16
Participant 8	14	14	14	15	4	8	13
Participant 9	16	17	19	19	15	20	20
Participant 10	4	4	4	4	10	2	6
Participant 11	3	10	10	10	14	4	8
Participant 12	10	10	8	10	10	9	11
Participant 13	6	6	5	5	8	6	6
Participant 14	2	2	4	1	5	1	2
Participant 15	8	16	12	16	10	10	16
Participant 16	6	8	10	9	10	11	10

Table C.4. TLX data from texture condition.

	Mental Demand	Physical Demand	Time Pressure	Effort Expended	Performance Level Achieved	Frustration Experienced	Fatigue Experienced
Participant 1	3	3	5	4	12	4	4
Participant 2	11	9	2	12	18	16	12
Participant 3	9	11	9	11	4	9	9
Participant 4	15	6	12	12	16	16	2
Participant 5	8	10	4	12	10	14	4
Participant 6	4	10	10	10	7	3	2
Participant 7	11	10	17	13	7	16	12
Participant 8	14	16	14	18	6	8	14
Participant 9	15	17	16	18	13	20	20
Participant 10	3	3	4	4	12	8	6
Participant 11	3	8	10	8	16	5	6
Participant 12	10	13	8	14	11	14	12
Participant 13	14	12	12	13	15	14	14
Participant 14	2	2	4	3	8	6	5
Participant 15	12	16	12	16	14	14	17
Participant 16	11	12	10	10	12	16	11

Table C.5. TLX data from visual condition.

	Mental Demand	Physical Demand	Time Pressure	Effort Expended	Performance Level Achieved	Frustration Experienced	Fatigue Experienced
Participant 1	5	3	5	4	15	5	5
Participant 2	12	12	2	12	18	20	10
Participant 3	5	9	8	8	4	3	5
Participant 4	4	4	4	4	4	2	2
Participant 5	6	8	6	8	11	8	4
Participant 6	2	2	2	2	6	1	1
Participant 7	9	9	17	9	5	9	9
Participant 8	18	17	16	17	7	17	15
Participant 9	18	18	15	15	11	19	19
Participant 10	2	1	2	4	10	3	4
Participant 11	3	6	9	6	14	3	6
Participant 12	11	13	8	14	13	10	10
Participant 13	12	10	12	10	10	10	9
Participant 14	2	0	4	1	6	2	1
Participant 15	5	17	12	16	12	12	18
Participant 16	6	8	9	9	11	9	10

Table C.6. Participant Preference for each condition.

	Gravity Well	Recess	Friction	Texture	Visual
Participant 1	2.000	4.000	3.000	5.000	1.000
Participant 2	2.000	4.000	1.000	5.000	3.000
Participant 3	1.000	2.000	3.000	5.000	4.000
Participant 4	1.000	2.000	4.000	5.000	3.000
Participant 5	1.000	2.000	3.000	5.000	4.000
Participant 6	1.000	2.000	4.000	5.000	3.000
Participant 7	1.000	5.000	2.000	3.000	4.000
Participant 8	1.000	2.000	4.000	5.000	3.000
Participant 9	3.000	2.000	5.000	4.000	1.000
Participant 10	2.000	1.000	5.000	4.000	3.000
Participant 11	1.000	2.000	4.000	5.000	3.000
Participant 12	2.000	1.000	4.000	5.000	3.000
Participant 13	2.000	1.000	3.000	4.000	5.000
Participant 14	2.000	3.000	1.000	5.000	4.000
Participant 15	1.000	2.000	3.000	5.000	4.000
Participant 16	1.000	3.000	2.000	5.000	4.000

Table C.7. Mean approach and stopping time in gravity well condition.

	Approach Time (ms)				Stopping Time (ms)		
	Gravity: low	Gravity: medium	Gravity: high		Gravity: low	Gravity: medium	Gravity: high
Participant 1	.571	.446	.400		.201	.190	.200
Participant 2	.546	.400	.354		.244	.247	.175
Participant 3	.469	.479	.510		.165	.221	.146
Participant 4	.495	.680	.731		.169	.205	.223
Participant 5	.532	.434	.483		.235	.211	.211
Participant 6	.719	.711	.684		.275	.294	.355
Participant 7	.492	.531	.416		.225	.266	.242
Participant 8	.486	.435	.346		.198	.196	.189
Participant 9	.571	.550	.543		.251	.205	.191
Participant 10	.432	.375	.371		.219	.207	.214
Participant 11	.469	.483	.478		.219	.217	.225
Participant 12	.630	.675	.715		.278	.283	.247
Participant 13	.495	.480	.465		.224	.218	.235
Participant 14	.455	.458	.463		.223	.222	.245
Participant 15	.522	.437	.441		.260	.229	.220
Participant 16	.548	.491	.490		.332	.306	.316

Table C.8. Mean approach and stopping time in recess condition.

	Approach Time (ms)				Stopping Time (ms)		
	Recess: low	Recess: medium	Recess: high		Recess: low	Recess: medium	Recess: high
Participant 1	.433	.479	.543		.183	.226	.211
Participant 2	.522	.478	.433		.238	.213	.235
Participant 3	.474	.548	.501		.158	.137	.123
Participant 4	.615	.525	.641		.194	.181	.199
Participant 5	.489	.500	.462		.242	.222	.224
Participant 6	.552	.659	.618		.373	.345	.336
Participant 7	.412	.615	.425		.223	.234	.238
Participant 8	.445	.389	.402		.158	.178	.166
Participant 9	.561	.557	.610		.223	.195	.213
Participant 10	.379	.334	.369		.212	.202	.229
Participant 11	.519	.503	.467		.214	.231	.211
Participant 12	.640	.722	.674		.268	.251	.295
Participant 13	.539	.563	.455		.215	.212	.266
Participant 14	.531	.518	.454		.314	.272	.255
Participant 15	.465	.501	.512		.246	.279	.265
Participant 16	.479	.441	.460		.321	.308	.269

Table C.9. Mean approach and stopping time in friction condition.

	Approach Time (ms)				Stopping Time (ms)		
	Friction: low	Friction: medium	Friction: high		Friction: low	Friction: medium	Friction: high
Participant 1	.464	.472	.439		.278	.221	.189
Participant 2	.446	.429	.456		.300	.286	.215
Participant 3	.540	.475	.489		.178	.198	.162
Participant 4	.587	.585	.650		.246	.266	.213
Participant 5	.395	.355	.459		.244	.204	.210
Participant 6	.643	.689	.696		.347	.323	.369
Participant 7	.457	.452	.401		.209	.245	.246
Participant 8	.427	.520	.408		.212	.193	.186
Participant 9	.475	.443	.467		.175	.208	.234
Participant 10	.322	.294	.317		.208	.182	.182
Participant 11	.460	.527	.513		.158	.201	.213
Participant 12	.616	.680	.616		.320	.349	.312
Participant 13	.431	.514	.479		.257	.226	.277
Participant 14	.419	.470	.472		.267	.275	.270
Participant 15	.449	.469	.391		.268	.288	.298
Participant 16	.486	.413	.461		.316	.327	.334

Table C.10. Mean approach and stopping time in texture condition.

	Approach Time (ms)				Stopping Time (ms)		
	Texture: low	Texture: medium	Texture: high		Texture: low	Texture: medium	Texture: high
Participant 1	.434	.433	.514		.221	.198	.237
Participant 2	.462	.462	.319		.315	.349	.416
Participant 3	.407	.398	.384		.174	.158	.145
Participant 4	.613	.641	.637		.263	.231	.237
Participant 5	.415	.364	.374		.254	.250	.265
Participant 6	.687	.608	.657		.343	.353	.313
Participant 7	.411	.434	.454		.241	.279	.379
Participant 8	.431	.548	.522		.220	.212	.243
Participant 9	.418	.326	.432		.246	.232	.243
Participant 10	.281	.261	.335		.202	.185	.170
Participant 11	.546	.511	.477		.182	.171	.179
Participant 12	.640	.666	.717		.362	.330	.385
Participant 13	.345	.449	.466		.284	.294	.250
Participant 14	.456	.416	.478		.371	.321	.303
Participant 15	.493	.565	.479		.253	.284	.283
Participant 16	.395	.407	.387		.398	.283	.250

Table C.11. Mean approach and stopping time in visual condition.

	Approach Time (ms)				Stopping Time (ms)		
	Visual	Visual	Visual		Visual	Visual	Visual
Participant 1	.473	.501	.482		.217	.216	.175
Participant 2	.517	.494	.579		.250	.285	.199
Participant 3	.560	.496	.438		.186	.152	.190
Participant 4	.545	.640	.670		.234	.255	.236
Participant 5	.500	.477	.510		.247	.248	.233
Participant 6	.683	.598	.688		.351	.355	.337
Participant 7	.485	.479	.392		.232	.281	.291
Participant 8	.448	.507	.431		.170	.187	.200
Participant 9	.516	.516	.613		.204	.203	.199
Participant 10	.388	.295	.334		.178	.188	.180
Participant 11	.525	.643	.536		.206	.219	.215
Participant 12	.650	.562	.700		.251	.290	.307
Participant 13	.464	.530	.557		.266	.215	.215
Participant 14	.482	.375	.396		.281	.334	.267
Participant 15	.609	.550	.649		.248	.324	.280
Participant 16	.545	.504	.465		.380	.356	.364

Table C.12. Mean clicking and leaving time in gravity well condition.

	Clicking Time (ms)				Leaving Time (ms)		
	Gravity: low	Gravity: medium	Gravity: high		Gravity: low	Gravity: low	Gravity: medium
Participant 1	.211	.192	.234		.281	.211	.192
Participant 2	.201	.207	.170		.331	.201	.207
Participant 3	.158	.155	.161		.279	.158	.155
Participant 4	.179	.215	.204		.340	.179	.215
Participant 5	.174	.189	.185		.373	.174	.189
Participant 6	.123	.123	.122		.241	.123	.123
Participant 7	.147	.140	.118		.210	.147	.140
Participant 8	.111	.111	.104		.328	.111	.111
Participant 9	.137	.127	.135		.260	.137	.127
Participant 10	.126	.137	.141		.429	.126	.137
Participant 11	.164	.165	.167		.356	.164	.165
Participant 12	.178	.178	.219		.390	.178	.178
Participant 13	.227	.178	.212		.299	.227	.178
Participant 14	.145	.145	.154		.256	.145	.145
Participant 15	.120	.119	.132		.253	.120	.119
Participant 16	.153	.158	.148		.261	.153	.158

Table C.13. Mean cicking and leaving time in recess condition.

	Clicking Time (ms)				Leaving Time (ms)		
	Recess: low	Recess: medium	Recess: high		Recess: low	Recess: medium	Recess: high
Participant 1	.272	.314	.243		.195	.192	.259
Participant 2	.172	.191	.196		.304	.256	.323
Participant 3	.163	.174	.174		.255	.250	.219
Participant 4	.198	.206	.210		.301	.287	.271
Participant 5	.185	.178	.184		.321	.310	.280
Participant 6	.138	.142	.134		.256	.267	.257
Participant 7	.146	.147	.145		.220	.159	.227
Participant 8	.092	.110	.115		.325	.282	.305
Participant 9	.116	.123	.127		.202	.205	.225
Participant 10	.132	.152	.130		.394	.389	.395
Participant 11	.162	.180	.194		.373	.359	.373
Participant 12	.183	.256	.190		.405	.414	.341
Participant 13	.190	.188	.202		.275	.232	.235
Participant 14	.141	.133	.152		.291	.263	.273
Participant 15	.124	.141	.128		.233	.215	.239
Participant 16	.176	.154	.153		.279	.322	.269

Table C.14. Mean clicking and leaving time in friction condition.

	Clicking Time (ms)				Leaving Time (ms)		
	Friction: low	Friction: medium	Friction: high		Friction: low	Friction: medium	Friction: high
Participant 1	.234	.230	.238		.325	.266	.312
Participant 2	.194	.225	.236		.380	.338	.338
Participant 3	.175	.164	.174		.279	.206	.243
Participant 4	.210	.208	.191		.426	.340	.291
Participant 5	.212	.212	.210		.276	.229	.300
Participant 6	.164	.162	.147		.254	.298	.314
Participant 7	.146	.149	.151		.361	.343	.253
Participant 8	.129	.114	.133		.308	.295	.351
Participant 9	.111	.129	.142		.178	.162	.191
Participant 10	.129	.135	.156		.383	.329	.449
Participant 11	.136	.162	.161		.373	.387	.435
Participant 12	.193	.182	.195		.448	.356	.385
Participant 13	.156	.193	.197		.363	.302	.308
Participant 14	.162	.152	.157		.259	.205	.253
Participant 15	.131	.132	.131		.228	.252	.271
Participant 16	.159	.178	.167		.280	.271	.352

Table C.15. Mean clicking and leaving time in texture condition.

	Clicking Time (ms)				Leaving Time (ms)		
	Texture: low	Texture: medium	Texture: high		Texture: low	Texture: medium	Texture: high
Participant 1	.212	.244	.214		.232	.280	.234
Participant 2	.240	.230	.222		.244	.238	.278
Participant 3	.176	.159	.182		.198	.179	.175
Participant 4	.210	.209	.195		.436	.447	.393
Participant 5	.195	.191	.189		.268	.216	.190
Participant 6	.174	.160	.163		.225	.220	.203
Participant 7	.144	.157	.167		.169	.209	.270
Participant 8	.113	.111	.100		.246	.200	.253
Participant 9	.148	.144	.143		.220	.211	.237
Participant 10	.140	.139	.132		.393	.341	.398
Participant 11	.153	.147	.155		.399	.329	.363
Participant 12	.192	.188	.219		.443	.404	.391
Participant 13	.211	.225	.193		.321	.278	.236
Participant 14	.166	.157	.168		.196	.221	.181
Participant 15	.138	.145	.132		.173	.206	.189
Participant 16	.213	.197	.184		.271	.247	.199

Table C.16. Mean clicking and leaving time in visual condition.

	Clicking Time (ms)				Leaving Time (ms)		
	Visual	Visual	Visual		Visual	Visual	Visual
Participant 1	.216	.219	.214		.217	.187	.236
Participant 2	.220	.195	.219		.246	.285	.268
Participant 3	.169	.173	.171		.269	.243	.231
Participant 4	.196	.225	.191		.432	.442	.343
Participant 5	.184	.186	.192		.272	.346	.277
Participant 6	.140	.154	.144		.235	.272	.248
Participant 7	.148	.152	.124		.254	.308	.289
Participant 8	.116	.120	.105		.249	.303	.242
Participant 9	.143	.151	.158		.213	.224	.191
Participant 10	.127	.137	.130		.347	.326	.289
Participant 11	.153	.175	.161		.334	.347	.334
Participant 12	.172	.179	.174		.401	.328	.320
Participant 13	.208	.209	.192		.293	.251	.250
Participant 14	.168	.149	.159		.243	.234	.195
Participant 15	.129	.129	.128		.190	.202	.219
Participant 16	.205	.180	.189		.257	.264	.261

Table C.17. Total slip-off and slide-over errors in gravity well condition.

	Slip-off				Slide-over		
	Gravity: low	Gravity: medium	Gravity: high		Gravity: low	Gravity: low	Gravity: medium
Participant 1	.00	.00	3.00		5.00	.00	.00
Participant 2	2.00	.00	.00		9.00	2.00	.00
Participant 3	2.00	1.00	1.00		16.00	2.00	1.00
Participant 4	.00	.00	1.00		9.00	.00	.00
Participant 5	1.00	2.00	.00		7.00	1.00	2.00
Participant 6	1.00	.00	.00		8.00	1.00	.00
Participant 7	.00	1.00	1.00		19.00	.00	1.00
Participant 8	4.00	2.00	.00		13.00	4.00	2.00
Participant 9	3.00	1.00	1.00		22.00	3.00	1.00
Participant 10	1.00	.00	.00		14.00	1.00	.00
Participant 11	1.00	.00	.00		10.00	1.00	.00
Participant 12	.00	.00	.00		4.00	.00	.00
Participant 13	2.00	1.00	.00		4.00	2.00	1.00
Participant 14	.00	1.00	.00		12.00	.00	1.00
Participant 15	.00	.00	1.00		7.00	.00	.00
Participant 16	1.00	2.00	2.00		7.00	1.00	2.00

Table C.18. Total slip-off and slide-over errors in recess condition.

	Slip-off				Slide-over		
	Recess: low	Recess: medium	Recess: high		Recess: low	Recess: medium	Recess: high
Participant 1	2.00	.00	3.00		12.00	12.00	9.00
Participant 2	1.00	.00	.00		5.00	1.00	5.00
Participant 3	3.00	5.00	4.00		22.00	14.00	24.00
Participant 4	2.00	.00	4.00		18.00	16.00	7.00
Participant 5	4.00	1.00	6.00		17.00	10.00	10.00
Participant 6	1.00	1.00	1.00		12.00	11.00	10.00
Participant 7	1.00	.00	1.00		23.00	10.00	17.00
Participant 8	.00	1.00	2.00		6.00	14.00	14.00
Participant 9	2.00	3.00	2.00		13.00	15.00	21.00
Participant 10	.00	1.00	2.00		18.00	18.00	18.00
Participant 11	4.00	1.00	1.00		8.00	11.00	13.00
Participant 12	1.00	.00	.00		5.00	1.00	4.00
Participant 13	1.00	2.00	1.00		5.00	5.00	12.00
Participant 14	.00	.00	.00		3.00	4.00	7.00
Participant 15	2.00	1.00	.00		18.00	12.00	14.00
Participant 16	5.00	1.00	.00		8.00	7.00	8.00

Table C.19. Total slip-off and slide-over errors in friction condition.

	Slip-off				Slide-over		
	Friction: low	Friction: medium	Friction: high		Friction: low	Friction: medium	Friction: high
Participant 1	1.00	.00	2.00		13.00	14.00	12.00
Participant 2	.00	.00	1.00		15.00	14.00	7.00
Participant 3	7.00	7.00	4.00		20.00	19.00	25.00
Participant 4	4.00	2.00	1.00		20.00	7.00	14.00
Participant 5	8.00	2.00	2.00		19.00	29.00	10.00
Participant 6	3.00	2.00	3.00		15.00	7.00	11.00
Participant 7	1.00	2.00	5.00		27.00	22.00	45.00
Participant 8	7.00	2.00	6.00		18.00	18.00	16.00
Participant 9	8.00	6.00	1.00		21.00	22.00	24.00
Participant 10	1.00	2.00	.00		18.00	30.00	18.00
Participant 11	1.00	4.00	4.00		12.00	5.00	6.00
Participant 12	1.00	3.00	2.00		4.00	8.00	8.00
Participant 13	2.00	3.00	.00		13.00	12.00	12.00
Participant 14	7.00	3.00	1.00		18.00	9.00	15.00
Participant 15	3.00	3.00	4.00		15.00	16.00	18.00
Participant 16	.00	6.00	2.00		13.00	18.00	19.00

Table C.20. Total slip-off and slide-over errors in texture condition.

	Slip-off				Slide-over		
	Texture: low	Texture: medium	Texture: high		Texture: low	Texture: medium	Texture: high
Participant 1	1.00	1.00	4.00		13.00	18.00	26.00
Participant 2	1.00	2.00	3.00		26.00	25.00	50.00
Participant 3	12.00	9.00	15.00		49.00	55.00	78.00
Participant 4	2.00	4.00	3.00		23.00	23.00	17.00
Participant 5	4.00	1.00	2.00		41.00	44.00	44.00
Participant 6	5.00	3.00	4.00		59.00	46.00	74.00
Participant 7	5.00	7.00	7.00		51.00	70.00	64.00
Participant 8	5.00	10.00	4.00		31.00	26.00	51.00
Participant 9	9.00	8.00	11.00		67.00	54.00	60.00
Participant 10	.00	4.00	5.00		37.00	43.00	28.00
Participant 11	.00	2.00	2.00		12.00	10.00	36.00
Participant 12	4.00	4.00	4.00		12.00	40.00	27.00
Participant 13	7.00	3.00	4.00		34.00	45.00	30.00
Participant 14	4.00	6.00	3.00		35.00	39.00	37.00
Participant 15	2.00	4.00	4.00		27.00	37.00	52.00
Participant 16	7.00	5.00	8.00		66.00	57.00	107.00

Table C.21. Total slip-off and slide-over errors in visual condition.

	Slip-off				Slide-over		
	Visual	Visual	Visual		Visual	Visual	Visual
Participant 1	1.00	3.00	2.00		20.00	13.00	17.00
Participant 2	2.00	2.00	2.00		10.00	10.00	8.00
Participant 3	4.00	5.00	10.00		20.00	25.00	32.00
Participant 4	.00	1.00	1.00		13.00	12.00	18.00
Participant 5	2.00	2.00	2.00		23.00	18.00	19.00
Participant 6	1.00	1.00	.00		19.00	18.00	16.00
Participant 7	5.00	3.00	1.00		33.00	31.00	43.00
Participant 8	.00	8.00	.00		24.00	21.00	15.00
Participant 9	4.00	11.00	2.00		30.00	32.00	15.00
Participant 10	2.00	.00	1.00		25.00	30.00	32.00
Participant 11	3.00	6.00	4.00		14.00	7.00	17.00
Participant 12	2.00	.00	1.00		9.00	14.00	6.00
Participant 13	1.00	1.00	5.00		12.00	6.00	14.00
Participant 14	4.00	3.00	3.00		17.00	20.00	23.00
Participant 15	3.00	3.00	2.00		11.00	18.00	11.00
Participant 16	4.00	1.00	1.00		16.00	15.00	12.00

Table C.22. Total off-target and wrong-target errors in gravity well condition.

	Off-target				Wrong-target		
	Gravity: low	Gravity: medium	Gravity: high		Gravity: low	Gravity: medium	Gravity: high
Participant 1	8.00	1.00	7.00		.00	1.00	.00
Participant 2	2.00	6.00	2.00		3.00	.00	.00
Participant 3	15.00	37.00	11.00		.00	.00	.00
Participant 4	8.00	32.00	20.00		1.00	1.00	.00
Participant 5	5.00	15.00	13.00		.00	2.00	.00
Participant 6	3.00	1.00	1.00		1.00	1.00	1.00
Participant 7	6.00	7.00	3.00		2.00	4.00	3.00
Participant 8	12.00	10.00	7.00		1.00	2.00	1.00
Participant 9	17.00	13.00	3.00		1.00	.00	.00
Participant 10	11.00	6.00	6.00		3.00	.00	2.00
Participant 11	1.00	6.00	1.00		.00	.00	.00
Participant 12	2.00	.00	4.00		.00	1.00	.00
Participant 13	1.00	4.00	1.00		.00	.00	.00
Participant 14	.00	.00	.00		.00	.00	1.00
Participant 15	1.00	3.00	4.00		.00	.00	2.00
Participant 16	1.00	4.00	.00		1.00	.00	1.00

Table C.23. Total off-target and wrong-target errors in recess condition.

	Off-target				Wrong-target		
	Recess: low	Recess: medium	Recess: high		Recess: low	Recess: medium	Recess: high
Participant 1	3.00	10.00	10.00		3.00	.00	1.00
Participant 2	1.00	2.00	2.00		1.00	.00	.00
Participant 3	19.00	14.00	24.00		.00	.00	.00
Participant 4	23.00	24.00	12.00		.00	1.00	.00
Participant 5	15.00	1.00	5.00		.00	2.00	1.00
Participant 6	1.00	4.00	2.00		1.00	.00	.00
Participant 7	5.00	4.00	10.00		2.00	5.00	1.00
Participant 8	3.00	6.00	9.00		.00	1.00	1.00
Participant 9	6.00	15.00	18.00		3.00	.00	.00
Participant 10	10.00	6.00	8.00		.00	.00	1.00
Participant 11	.00	9.00	1.00		.00	2.00	1.00
Participant 12	.00	2.00	.00		.00	.00	.00
Participant 13	.00	.00	3.00		.00	.00	.00
Participant 14	2.00	.00	.00		1.00	1.00	.00
Participant 15	8.00	2.00	7.00		3.00	.00	.00
Participant 16	2.00	.00	.00		2.00	1.00	.00

Table C.24. Total off-target and wrong-target errors in friction condition.

	Off-target				Wrong-target		
	Friction: low	Friction: medium	Friction: high		Friction: low	Friction: medium	Friction: high
Participant 1	4.00	4.00	.00		.00	.00	.00
Participant 2	2.00	5.00	2.00		1.00	1.00	1.00
Participant 3	16.00	13.00	13.00		.00	1.00	.00
Participant 4	6.00	4.00	10.00		.00	2.00	1.00
Participant 5	6.00	12.00	7.00		1.00	.00	.00
Participant 6	6.00	1.00	4.00		.00	.00	.00
Participant 7	12.00	4.00	12.00		.00	1.00	3.00
Participant 8	4.00	6.00	4.00		1.00	2.00	.00
Participant 9	6.00	10.00	7.00		4.00	2.00	4.00
Participant 10	9.00	11.00	9.00		1.00	3.00	1.00
Participant 11	12.00	1.00	6.00		.00	.00	.00
Participant 12	4.00	2.00	2.00		1.00	.00	.00
Participant 13	1.00	3.00	5.00		.00	.00	.00
Participant 14	4.00	4.00	1.00		1.00	.00	1.00
Participant 15	6.00	.00	1.00		.00	.00	.00
Participant 16	3.00	2.00	4.00		1.00	.00	3.00

Table C.25. Total off-target and wrong-target errors in texture condition.

	Off-target				Wrong-target		
	Texture: low	Texture: medium	Texture: high		Texture: low	Texture: medium	Texture: high
Participant 1	2.00	1.00	4.00		1.00	.00	.00
Participant 2	.00	5.00	3.00		1.00	1.00	.00
Participant 3	30.00	27.00	40.00		2.00	.00	.00
Participant 4	7.00	6.00	10.00		.00	.00	.00
Participant 5	13.00	10.00	8.00		.00	1.00	.00
Participant 6	10.00	16.00	7.00		1.00	1.00	.00
Participant 7	9.00	10.00	22.00		.00	3.00	3.00
Participant 8	14.00	15.00	20.00		.00	.00	2.00
Participant 9	23.00	18.00	9.00		4.00	1.00	1.00
Participant 10	13.00	9.00	15.00		3.00	1.00	.00
Participant 11	7.00	3.00	18.00		.00	.00	.00
Participant 12	1.00	7.00	3.00		.00	.00	2.00
Participant 13	12.00	14.00	6.00		.00	.00	.00
Participant 14	7.00	6.00	7.00		1.00	2.00	1.00
Participant 15	4.00	5.00	12.00		.00	.00	.00
Participant 16	6.00	12.00	36.00		2.00	2.00	1.00

Table C.26. Total off-target and wrong-target errors in visual condition.

	Off-target				Wrong-target		
		Visual Average			Visual	Visual	Visual
Participant 1		15.00			.00	1.00	1.00
Participant 2		.00			.00	.00	1.00
Participant 3		62.00			.00	1.00	.00
Participant 4		16.00			.00	.00	1.00
Participant 5		19.00			.00	.00	.00
Participant 6		7.00			.00	.00	1.00
Participant 7		36.00			1.00	4.00	2.00
Participant 8		39.00			.00	1.00	.00
Participant 9		43.00			4.00	1.00	5.00
Participant 10		38.00			8.00	1.00	5.00
Participant 11		17.00			1.00	.00	.00
Participant 12		11.00			.00	1.00	.00
Participant 13		13.00			.00	.00	.00
Participant 14		8.00			.00	1.00	1.00
Participant 15		21.00			.00	.00	.00
Participant 16		19.00			3.00	1.00	.00

Appendix D. Questionnaire Administered in Menu Study

This appendix contains a copy of the demographics questionnaire administered after the second study described in Chapter 4 investigating a haptically augmented menu system.

Questionnaire

Age: _____

Gender: M / F

Nationality: _____

Rate your level of computer experience (tick one):

None

Basic

Intermediate

Expert

Do not fill out this section

Condition:

Appendix E. Raw Data from Menu Study

This appendix contains a copy of the raw data gathered during the second study described in Chapter 4 investigating a haptically augmented menu system. It includes all subjective and objective data.

Table E.1. Results from TLX in visual condition.

	Mental Demand	Physical Demand	Time Pressure	Effort Expended	Performance Level Achieved	Frustration Experienced	Fatigue Experienced
Participant 1	7	3	6	6	6	6	2
Participant 2	4	14	2	12	2	7	18
Participant 3	13	14	10	10	10	10	11
Participant 4	0	4	0	3	0	1	2
Participant 5	4	3	0	3	2	2	2
Participant 6	1	4	2	3	0	0	3
Participant 7	8	6	7	11	7	6	5
Participant 8	11	8	10	9	11	7	6
Participant 9	7	10	12	11	10	10	11
Participant 10	13	10	12	13	9	10	12
Participant 11	14	14	10	14	14	14	14
Participant 12	1	4	4	6	14	2	8
Participant 13	6	6	7	10	5	8	7
Participant 14	12	11	9	14	12	16	11
Participant 15	8	8	5	12	11	17	18
Participant 16	19	14	0	13	10	12	16
Participant 17	14	10	10	14	14	6	10
Participant 18	8	7	10	13	7	10	6

Table E.2. Results from TLX in adjusted condition.

	Mental Demand	Physical Demand	Time Pressure	Effort Expended	Performance Level Achieved	Frustration Experienced	Fatigue Experienced
Participant 1	3	3	3	3	4	4	7
Participant 2	4	17	2	15	4	1	17
Participant 3	9	7	8	9	6	2	7
Participant 4	0	3	0	2	0	6	6
Participant 5	9	11	0	12	8	10	12
Participant 6	0	4	2	5	20	1	0
Participant 7	11	7	7	11	8	7	7
Participant 8	15	12	12	10	9	9	10
Participant 9	6	9	9	10	12	11	12
Participant 10	13	6	12	12	7	6	7
Participant 11	14	14	10	14	10	10	16
Participant 12	2	2	3	4	11	1	4
Participant 13	2	8	13	10	13	17	12
Participant 14	6	4	10	12	5	3	13
Participant 15	7	3	5	11	13	10	17
Participant 16	7	14	0	14	2	2	17
Participant 17	6	2	18	6	6	2	6
Participant 18	6	10	7	13	8	12	10

Table E.3. Results from TLX in haptic condition.

	Mental Demand	Physical Demand	Time Pressure	Effort Expended	Performance Level Achieved	Frustration Experienced	Fatigue Experienced
Participant 1	13	15	15	13	13	13	15
Participant 2	15	15	3	14	7	3	17
Participant 3	17	18	14	18	14	20	20
Participant 4	1	10	3	6	0	6	10
Participant 5	15	20	1	20	7	18	18
Participant 6	0	6	5	6	20	0	3
Participant 7	13	7	8	12	8	6	8
Participant 8	16	16	13	14	6	17	16
Participant 9	9	14	12	12	14	11	14
Participant 10	14	15	12	15	2	13	14
Participant 11	14	20	10	14	10	10	18
Participant 12	6	13	11	11	12	1	10
Participant 13	14	20	15	15	17	20	20
Participant 14	11	9	7	15	7	12	13
Participant 15	7	12	4	13	18	15	15
Participant 16	18	19	0	19	10	19	17
Participant 17	6	18	14	18	14	18	14
Participant 18	12	12	8	13	7	13	12

Table E.4. Objective measures from visual condition.

	Mean Time per trial (seconds)	Average total number of Slide-Over errors	Average total number of Wrong-Target errors	Average total number of Skip-ahead errors
Participant 1	3.38148	67	4	17
Participant 2	3.45469	82	11	18
Participant 3	3.52834	129	2	19
Participant 4	3.04663	102	0	19
Participant 5	4.03138	107	18	22
Participant 6	2.71028	96	4	33
Participant 7	2.81584	88	31	22
Participant 8	3.77855	138	4	20
Participant 9	3.73801	92	4	18
Participant 10	3.13461	94	7	5
Participant 11	3.30843	129	10	40
Participant 12	3.83154	115	2	10
Participant 13	2.96406	107	4	27
Participant 14	2.76817	129	3	21
Participant 15	3.7043	113	11	46
Participant 16	3.83242	169	8	42
Participant 17	2.97087	88	2	40
Participant 18	3.76445	254	9	72

Table E.5. Objective measures from adjusted condition.

	Mean Time per trial (seconds)	Average total number of Slide-Over errors	Average total number of Wrong-Target errors	Average total number of Skip-ahead errors
Participant 1	3.61453	45	5	3
Participant 2	3.4951	47	8	1
Participant 3	3.19426	83	2	1
Participant 4	3.22742	66	1	1
Participant 5	3.38274	52	4	0
Participant 6	2.71418	41	0	1
Participant 7	3.05917	50	9	1
Participant 8	4.00689	70	0	5
Participant 9	3.85509	67	4	3
Participant 10	3.26747	76	1	1
Participant 11	2.93482	42	7	2
Participant 12	3.73686	81	0	0
Participant 13	3.25097	49	0	5
Participant 14	2.49422	91	3	11
Participant 15	3.44952	47	2	3
Participant 16	3.352	100	8	3
Participant 17	2.86259	28	3	11
Participant 18	3.60827	149	7	31

Table E.6. Objective measures from haptic condition.

	Mean Time per trial (seconds)	Average total number of Slide-Over errors	Average total number of Wrong-Target errors	Average total number of Skip-ahead errors
Participant 1	4.08488	65	1	0
Participant 2	3.68987	71	3	1
Participant 3	3.57783	118	4	4
Participant 4	3.51895	77	0	0
Participant 5	4.11415	123	5	1
Participant 6	2.95162	80	1	1
Participant 7	3.46771	133	5	0
Participant 8	4.01944	94	1	7
Participant 9	4.40135	148	5	0
Participant 10	4.04535	179	0	0
Participant 11	3.7267	115	3	0
Participant 12	4.03268	166	1	1
Participant 13	3.19391	94	0	12
Participant 14	2.93047	121	3	6
Participant 15	5.44716	150	2	0
Participant 16	3.39367	109	6	10
Participant 17	3.10265	33	2	3
Participant 18	3.71103	158	6	27

Appendix F. Problem used in Shared Editor Study

This appendix contains a copy of the software engineering problem used in the first study presented in Chapter 6 looking at all aspects of the haptic cursor communication.

Instructions

A description of an existing manual (paper-based) system is given below. Using the collaborative editor, you are asked work in pairs to begin the process of designing a computerised system using an object-oriented approach. All work should be done in the editor. Please leave final versions for each of the deliverables requested clearly labelled in the editor.

Develop an outline object model, using the noun identification technique (or any other appropriate method). Draw a class diagram summarising the classes you have chosen and the associations between them.

Identify the use cases for the system; draw a diagram showing the actors, use cases and their interactions.

Develop two of your use cases into more detail and show how they interact with the classes you have chosen for 1, above. Add more detail to the class diagram (or modify it) as required.

If you are in any doubt about the meaning of any part of the description then state any assumptions you have made in order to complete your solution.

Description of Existing System

A club runs a video hire service for its members. There is a library of videos shelved in alphabetical order of title. There is no restriction on the number of videos that a member may have on loan at any time and videos are requested for return only if they are required by another member. Multiple copies of popular titles are available.

In order to borrow a video a member selects it from the shelf, takes it to the issue desk and gives the video, together with their club membership card, to the librarian. The librarian takes the video ID card from the video sleeve and adds the membership number to the video ID card. The librarian places the video ID card into the loans file, returns the club membership card to the member and the member leaves with the video.

When a member returns a video, they give the video back to the librarian, who finds the video ID card in the loans file and places the card in the video sleeve before returning the video to the library shelves.

To reserve a video that is on loan, a member asks the librarian to reserve the video with a given title and leaves their membership number. The librarian finds a relevant video ID card in the loans file and adds the member's number on the reservation column on the video card. The librarian also makes a note of the member who has the video on loan and completes a return request card addressed to that member. The address is taken from the library copy of the club membership list. When the reserved video is returned the librarian puts the video on the reserved shelf and completes a reservation ready card addressed to the member who requested the reservation, again taking the address from the membership list. The member requiring the video can then pick it up from the library and have it issued in the usual manner.

Librarians are often asked questions by members about the availability of videos and to assist them in answering these queries the librarians use title and subject indexes, in addition to inspecting the loans file. A video may be indexed under several different subject categories.

New videos are purchased by a sub-committee of the club and passed to the librarian for inclusion in the library. (The selection, ordering and payment for videos is outside the scope of the current investigation.). On receipt of a new video, a librarian has to make out a new video ID card for insertion into the video sleeve and create the relevant index entries. Librarians are also responsible for updating the membership list using information received from the Club Secretary.

The customer wants to develop a computerised information system that will record the details of all the videos in the library; keep accurate records of loans and reservations; support the process of recalling videos; and generally support the work of the librarians.

Appendix G. Instructions and User Manual from Shared Editor Study

This appendix contains a copy of the instructions and user manual that were handed out to participants in the first study presented in Chapter 6 looking at all aspects of the haptic cursor communication. The final page of this appendix detailing information as to the operation of the haptic communication was presented only to participants in the Haptic condition, and not those in the Visual condition, who would not experience this feedback.

CHASE – Minimal Manual

Chase (Collaborative Haptics and Structured Editing) is a structured collaborative editor. It allows two or more users to simultaneously work on the same structured diagram, for instance a software-engineering diagram.

Features

Telepointers

Other users in the editor are graphically represented by hollow round cursors. These cursors can be used to locate the other user and as a means of communication – by hovering over a location you can hope to draw attention to it.

Tools

To select a tool, click on its icon. The currently highlighted tool is displayed with a darker background. After an object tool (text, box, circle, line) has been used once the selected tool resets to the arrow. This behaviour can be overridden by double clicking on the desired tool. The tool background will then go black and the tool may be used multiple times.

Arrow

The arrow tool selects objects. All objects can be selected by left clicking on them. Furthermore, text objects can be added to an existing selection (of other text objects) by holding down the shift key and left clicking on them. Finally, all the text objects in an area can be selected by drawing a box around them.

A selected object appears with black square ‘handles’ in the corners. Adjusting these handles causes the object to be resized. Text objects cannot be resized.

To delete an object, select it and press the delete key.

Text

The text tool allows you to place and edit text objects. To place a text object, select the text tool and click on the canvas. To edit a text object simply click on it with the text tool. To finish editing a text object click elsewhere on the canvas. The ‘enter’ or ‘return’ keys will not finish editing a text object, they will simply bring you onto the next line.

While editing a text object another users telepointer can disappear behind it. Consequently, making long, multi-line, text strings is not recommended.

Box Group

Group objects are composed of a container and a title. The box container is a rectangular shape. After you create a group object you will be left editing its title. A group object can be selected by clicking on any of the lines that it is composed of, or its title. Text objects can be placed within group objects. A text object within a group object will be moved when that group object is moved and deleted when that group is deleted. Membership of a group is determined whenever a group is selected, meaning that one group can effectively steal objects from another by being placed in an overlapping position and then moved. Only one group object can be selected at a time.

Circle Group

The circle group has a similar behaviour to the box group

Line

Lines connect groups together. As a point in a line is placed it will jump towards the nearest group. A line object must connect two separate group objects, and makes no distinction between circle and box groups. Lines can be edited, or moved but will always snap to the nearest available shape. Lines maintain their connections when the groups they connect move.

Known Issues

If the groups on either end of a line are moved simultaneously, the behaviour of the line can become erratic.

Telepointers are available only when the other user they signify is present on the canvas. If the other user moves off their canvas, their telepointer will disappear.

Two users editing the same text object at the same time can cause unexpected results.

Collaborative Haptics

There are five ways to communicate haptically. It is worth noting that any force presented by the PHANToM can be ignored. To gain full benefit of the haptics, resisting the applied forces is often not the best approach. As several of the haptic effects are performed by pulling backwards on the PHANToM, to enter data in the editor it is best to rest lightly on the front of the workspace.

Proximity

This effect is passive – it is always on. As two users approach one another, each feel an increasing viscosity, or stickiness, in the workspace. Alternatively, if a user is still when another approaches they will feel a slight vibration. This effect can be used to monitor the proximity of other users.

Push

The push effect enables users to push one another about the workspace – it turns the telepointers into physical objects. If one cursor knocks into another, both should feel a force pushing them away. This effect could be used to catch the attention of another user.

Locate

The locate effect allows a user to activate a homing force towards the other user in the environment. To activate it, the PHANToM is pulled backwards and its button is depressed. Initially a small force is presented which indicates the direction of the user. This rapidly changes into a more substantial force to guide you to this location. To end the effect, release the button.

Grab

The grab effect is an inverted version of the locate effect. As a user activates it all other users in the workspace feel a homing force to the first user's location. This effect is initiated by pulling back on the PHANToM and performing a double click, leaving the button down after the second click. It is ended by releasing the button. For the duration of this effect the user activating it will experience increased viscosity against their movements

Gesture

The gesture effect allows one user to guide another through a complex path. To initiate a gesture one user moves over another and depresses the PHANToM's button. The other user will now feel forces constraining them to precisely follow the path of the first user. The gesture is ended by releasing the PHANToM's button. The user performing the gesture will experience increased viscosity in the workspace while they have the other user in tow.

Appendix H. Raw Data from CHASE Study

This appendix contains a copy of the raw data gathered during the first study described in Chapter 6 looking at all aspects of the haptic cursor communication. It includes all subjective data.

Table H.1. Results from custom collaborative questionnaire in haptic condition.

	Participant 1	Participant 2	Participant 3	Participant 4	Participant 5	Participant 6	Participant 7	Participant 8
Q1	2	4	6	4	4	5	3	5
Q2	4	6	4	6	6	5	4	5
Q3	4	5	4	5	6	5	4	5
Q4	6	6	6	6	5	5	5	6
Q5	4	6	5	5	5	6	4	6
Q6	4	6	4	4	4	5	3	6
Q7	4	6	5	5	6	2	4	6
Q8	5	6	6	5	6	6	4	6
Q9	5	1	2	6	5	6	4	5
Q10	5	5	5	6	6	5	4	5

Table H.2. Results from custom collaborative questionnaire in visual condition.

	Participant 9	Participant 10	Participant 11	Participant 12	Participant 13	Participant 14	Participant 15	Participant 16
Q1	3	5	4	3	2	6	5	0
Q2	2	3	1	2	2	3	1	1
Q3	2	4	4	3	3	3	4	1
Q4	4	5	4	3	6	6	4	5
Q5	4	4	6	4	5	6	3	2
Q6	4	4	5	3	4	6	3	0
Q7	2	5	4	3	5	3	0	1
Q8	4	5	5	3	6	5	4	5
Q9	3	4	6	3	4	4	3	1
Q10	4	4	6	5	5	5	4	1

Table H.3. Results from TLX in haptic condition.

	Participant 1	Participant 2	Participant 3	Participant 4	Participant 5	Participant 6	Participant 7	Participant 8
Mental Demand	13	18	17	13	10	4	12	14
Physical Demand	13	2	17	14	6	6	12	6
Time Pressure	8	6	3	6	6	1	6	6
Effort Expended	12	10	8	17	2	9	9	14
Performance Level Achieved	5	14	12	18	6	9	13	14
Frustration Experienced	11	14	18	17	19	9	12	14
Fatigue Experienced	11	0	2	1	6	0	16	6

Table H.4. Results from TLX in visual condition.

	Participant 9	Participant 10	Participant 11	Participant 12	Participant 13	Participant 14	Participant 15	Participant 16
Mental Demand	20	6	13	14	12	15	12	15
Physical Demand	2	7	2	2	1	10	3	11
Time Pressure	0	13	2	0	6	2	3	3
Effort Expended	16	18	12	10	10	14	13	12
Performance Level Achieved	0	18	11	0	9	18	10	3
Frustration Experienced	0	18	5	6	12	16	13	12
Fatigue Experienced	20	15	13	2	2	0	3	13

Table H.5. Results from ITC presence questionnaire in haptic condition.

	Participant 1	Participant 2	Participant 3	Participant 4	Participant 5	Participant 6	Participant 7	Participant 8
Age	25	21	28	25	23	26	23	26
Gender	M	M	F	F	F	M	F	F
Occupation	Student	Researcher	Researcher	Student	Student	Student	Student	Student
Nationality	British	British	German	Indian	Finnish/British	British	Greek	British
Experience	3	4	4	4	4	4	4	3
Games	4	2	1	1	1	2	2	2
Weekly TV	1	1	1	2	1	1	1	3
Education	6	6	6	6	6	6	6	6
TV Size	1	2	1	1	2	1	2	2
TV Know	2	2	2	1	3	2	3	3
Stereo Use	n	n	y	y	n	y	n	n
VR Use	n	y	y	n	n	n	n	y
3D Know	1	3	2	1	2	3	2	2
VR Know	1	3	2	1	2	3	2	3
a1	2	2	1	2	1	0	4	2
a2	1	1	1	1	0	2	0	0
a3	3	0	2	3	2	0	1	3
a4	3	3	2	1	1	2	4	2
a5	4	2	3	3	3	2	3	4
a6	3	3	3	3	3	4	4	4
b1	3	2	1	1	3	2	4	4
b2	3	4	3	3	3	4	3	4
b3	2	2	1	3	2	0	4	1
b4	3	2	3	3	3	4	3	4
b5	4	3	3	3	2	2	3	3
b6	3	4	2	1	3	2	4	3
b7	1	2	1	1	3	1	1	0
b8	3	3	3	3	3	2	4	3
b9	3	2	1	1	1	0	1	4
b10	2	0	3	1	3	0	3	0
b11	4	4	3	3	2	4	3	4
b12	4	4	3	3	3	2	3	4
b13	3	3	2	3	3	4	3	3
b14	0	0	0	1	0	0	1	0
b15	4	3	0	1	1	2	3	3
b16	3	2	1	3	2	2	3	3
b17	4	0	3	3	3	4	4	3
b18	3	2	3	1	1	0	1	3
b19	4	3	4	3	4	4	4	4
b20	1	2	2	1	1	0	1	3
b21	0	2	1	1	2	2	1	0
b22	0	0	0	1	1	0	0	1
b23	3	3	0	1	2	0	0	3
b24	1	0	0	1	0	0	0	1
b25	4	2	0	1	1	0	1	3
b26	0	0	0	1	0	0	0	0
b27	3	4	1	1	3	4	0	3
b28	3	3	1	1	3	4	2	3
b29	0	1	0	1	0	0	0	0
b30	3	2	0	1	0	0	1	1
b31	3	3	1	2	1	2	3	2
b32	3	4	3	3	1	4	4	4

Table H.6. Results from ITC presence questionnaire in visual condition.

	Participant 9	Participant 10	Participant 11	Participant 12	Participant 13	Participant 14	Participant 15	Participant 16
Age	22	40	27	29	28	23	22	26
Gender	M	M	F	M	M	M	M	F
Occupation	Student	Student	Student	Student	Student	Student	Student	Student
Nationality	British	British	French	British	British	Azeri	British	British
Experience	4	3	3	3	3	4	3	3
Games	4	1	1	3	2	2	3	3
Weekly TV	4	1	1	2	1	1	2	1
Education	6	6	6	6	6	6	6	6
TV Size	1	1	2	2	1	2	1	1
TV Know	1	2	2	1	3	3	2	2
Stereo Use	n	y	n	n	y	n	n	y
VR Use	n	n	n	n	n	n	y	n
3D Know	1	2	1	1	3	1	2	2
VR Know	1	1	1	1	3	2	2	2
a1	1	3	2	0	0	3	0	0
a2	1	0	0	0	0	0	1	0
a3	0	1	2	0	0	3	0	1
a4	1	3	2	0	0	3	0	1
a5	2	3	2	2	3	3	1	4
a6	1	3	3	2	2	4	1	1
b1	2	3	3	1	4	4	3	1
b2	3	2	3	2	2	4	3	1
b3	2	3	2	3	4	4	3	4
b4	3	3	3	3	4	3	3	2
b5	3	1	3	2	3	3	3	0
b6	1	3	2	2	1	3	3	1
b7	1	1	2	1	0	0	0	0
b8	2	3	3	2	1	3	1	0
b9	1	2	2	1	0	0	0	1
b10	2	0	3	3	4	0	3	3
b11	2	3	2	1	1	2	1	3
b12	3	3	2	1	0	2	3	3
b13	1	1	2	1	0	2	1	3
b14	0	0	0	0	0	2	1	3
b15	1	1	0	0	0	2	0	0
b16	0	2	2	2	1	2	1	1
b17	3	3	3	2	4	0	1	3
b18	1	1	1	2	0	2	0	1
b19	4	3	4	3	4	2	4	4
b20	1	0	1	0	0	2	0	2
b21	1	0	1	2	3	1	3	3
b22	0	0	0	0	0	0	0	0
b23	0	0	0	0	0	0	3	0
b24	0	0	0	0	0	0	0	0
b25	1	0	1	0	2	0	0	0
b26	0	0	0	0	0	2	1	3
b27	2	0	2	0	0	1	4	1
b28	2	0	2	0	0	0	4	0
b29	0	0	0	0	0	1	0	1
b30	1	0	2	0	0	1	0	3
b31	1	0	2	2	0	2	0	0
b32	2	2	3	2	0	2	1	1

Table H.7. Results from usability questionnaire in haptic condition.

	Participant 1	Participant 2	Participant 3	Participant 4	Participant 5	Participant 6	Participant 7	Participant 8
Q1	6	5	4	4	5	5	6	6
Q2	4	5	4	4	4	6	5	5
Q3	5	6	4	4	5	6	5	6
Q4	4	5	4	4	4	6	4	5
Q5	6	3	4	4	4	4	4	5
Q6	4	4	4	4	4	6	6	6
Q7	6	4	5	4	5	6	5	4
Q8	6	4	4	4	4	6	6	6
Q9	1	3	1	2	1	0	2	1
Q10	4	3	2	2	4	2	4	4
Q11	1	3	2	4	4	0	3	3
Q12	1	2	2	2	3	2	6	5
Q13	4	5	4	4	4	4	5	4
Q14	2	4	4	4	4	6	6	3
Q15	6	5	5	4	4	4	4	3
Q16	4	5	4	2	4	6	5	4
Q17	4	6	4	2	4	5	4	6
Q18	5	4	4	2	2	4	1	6
Q19	5	5	4	2	5	5	5	6

Table H.8. Results from usability questionnaire in visual condition.

	Participant 9	Participant 10	Participant 11	Participant 12	Participant 13	Participant 14	Participant 15	Participant 16
Q1	4	5	4	1	4	2	5	3
Q2	3	5	4	2	5	1	5	3
Q3	2	5	2	1	4	4	6	3
Q4	2	5	2	0	4	4	6	3
Q5	2	5	2	0	4	4	4	3
Q6	5	5	4	0	6	4	6	3
Q7	6	6	4	5	6	4	6	3
Q8	5	5	2	1	4	4	5	2
Q9	3	0	0	3	2	1	0	0
Q10	6	6	3	2	4	4	5	4
Q11	0	0	1	3	3	1	3	3
Q12	3	5	2	3	3	2	5	3
Q13	5	5	5	4	3	4	5	3
Q14	5	5	4	4	3	4	5	3
Q15	5	5	4	0	3	5	5	3
Q16	4	4	4	0	3	4	4	3
Q17	4	5	2	0	4	4	5	3
Q18	1	5	0	2	4	4	2	1
Q19	4	6	4	0	5	4	4	3

Appendix I. Images Used in Gesture Study

This appendix contains a copy of the images used in the second study described in Chapter 6 investigating communication through the haptic gesture. The first page contains the images that were used in the practice session, the second page those that were used in the experimental session.

Figure I.1. Images used in practice session is gesture study.


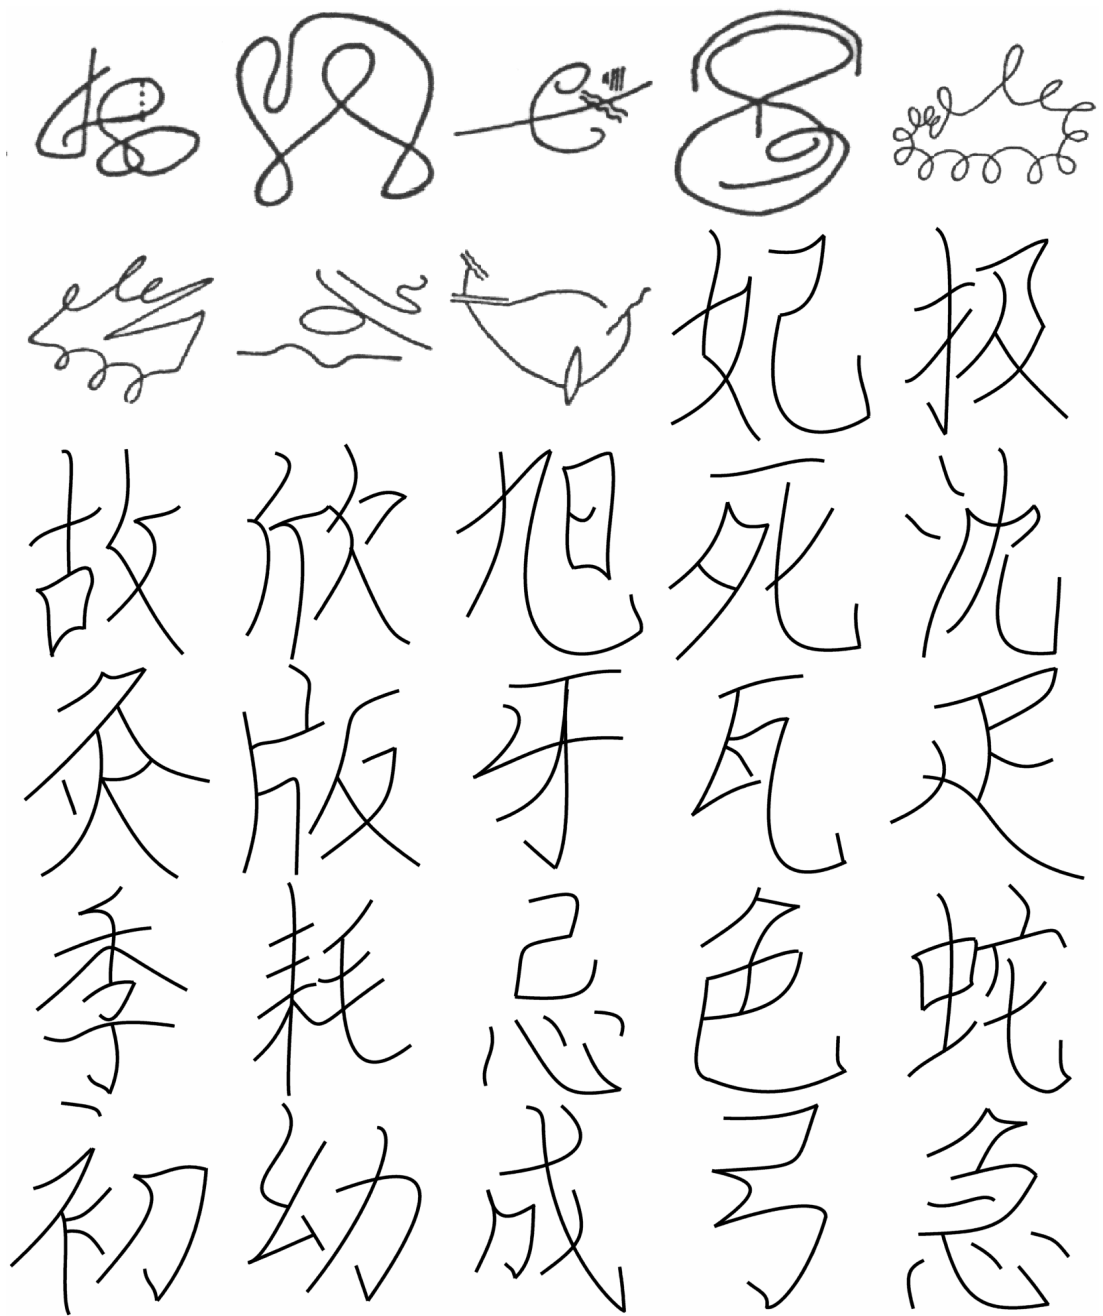


Figure I.2. Images used in experimental session is gesture study.



Appendix J. Instructions from Gesture Study

This appendix contains a copy of the instructions given to participants in the second study described in Chapter 6 investigating communication through the haptic gesture. The first page contains the instructions given to the encoders, while the second contains those given to the decoders. The third page is a copy of the instructions given to the image raters.

Instructions

In this experiment you will use the phantom as a cursor control device, in a narrow vertical plane.

Your task will be to describe shapes presented to you to another user, who attempts to draw these shapes. You will not be able to see the other user's drawings but can talk freely to them at all times. You can also gesture to the other user. When gesturing you will be able to see the other users cursor as well as your own and will also feel a slight resistance to your motions.

The other user will experience three types of gesture.

Graphical gestures allow them to see your position

Touch gestures provide forces that push the other user along the path of your cursor, but provide no visual information – for the duration of the gesture the other user can see neither your cursor nor their own.

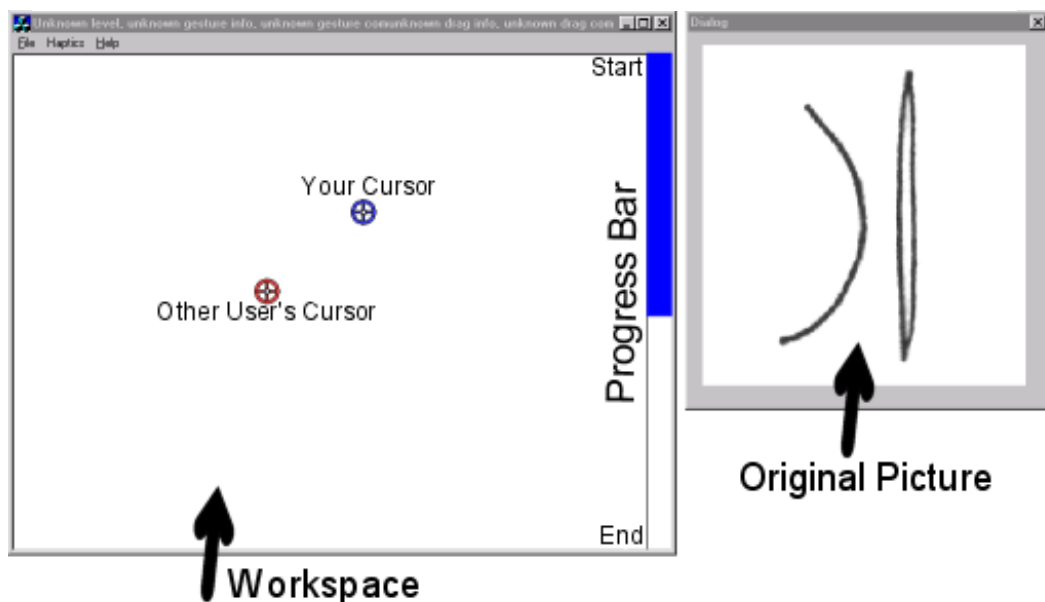
Combined gestures incorporate touch and graphical gestures. The other user can both see and feel the gesture.

The other user cannot draw while you are gesturing. Equally you cannot interrupt them while they are drawing, doing so results in your cursor changing to an hourglass. You can use this to determine when the other user stops drawing. To activate a gesture, depress the phantom's button. To end a gesture release it.

Each trial is timed and lasts a maximum of 80 seconds. Time remaining is indicated by a progress bar on the right. The other user controls the start of each trial and can also end a trial at any time. Please try to work as quickly as possible during the trials but feel free to rest between them.

You will be given an opportunity to practice both with the phantom and with the other user. The practice will involve the other user experiencing each type of gesture. The order that they are presented in will remain the same in the actual experiment. The practice consists of 15 trials, the experiment 30.

Figure J.1. Annotated encoder view in gesture study.



There is a £50 shared prize for the fastest, most accurate subjects

Instructions

In this experiment you will use the phantom as a cursor control device, in a narrow vertical plane.

Your task will be to draw, as accurately as possible, shapes described to you by another user. The other user will be unable to see your drawings but you can talk freely to them at all times. The other user can also gesture to you. You will experience three types of gesture.

Graphical gestures allow you to see the others user position

Touch gestures provide forces that push you along the path of the other user's cursor, but provide no visual information – for the duration of the gesture you will be not be able to see your cursor, or the other user's.

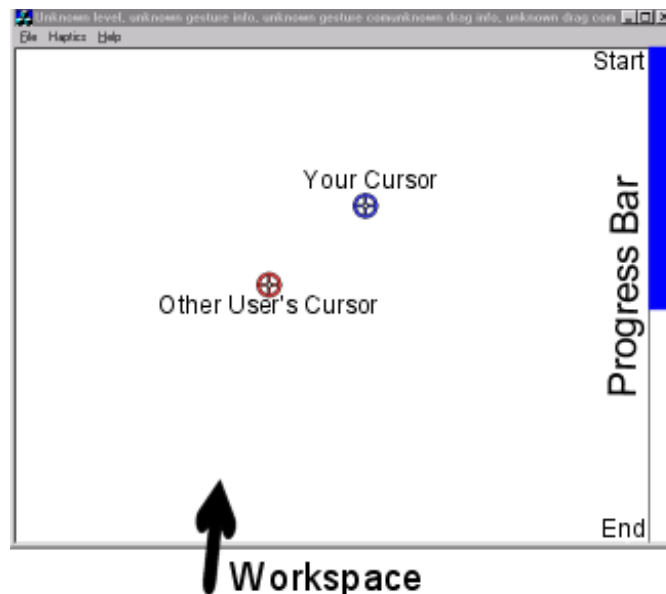
Combined gestures incorporate touch and graphical gestures. You can both see and feel the gesture.

You cannot draw while the other user is gesturing. Equally they cannot interrupt you while you are drawing. To draw depress the PHANToM's button. To stop drawing release it. To erase hold down shift as you begin to draw.

Each trial is timed and lasts a maximum of 80 seconds. Time remaining is indicated by a progress bar on the right. You control the start of each trial and can also end a trial at any time. Trials are started and stopped by pressing the space bar. Please try to work as quickly and accurately as possible during the trials but feel free to rest between trials.

You will be given an opportunity to practice both with the phantom and with the other user. The practice will involve each type of gesture. The order that they are presented in will remain the same in the actual experiment. The practice consists of 15 trials, the experiment 30.

Figure J.2. Annotated decoder view in gesture study.




There is a £50 shared prize for the fastest, most accurate subjects

Instructions for raters

Your task is to rate copies of an image, in terms of similarity, to the original image. Similarity is taken to include all aspects of the image, such as the number of lines in each image, the size of these lines, their position, curvature, and the accuracy of intersections with other lines. Be careful not to skew your ratings according to the most attractive images.

There are 30 original images, each presented alone on a page. Attached to each original are 18 copies of that image, printed 9 to a page. Rate these copies by marking on them. Rating should take place on the following scale:

Figure J.3. Scale used in image rating process.



Your ratings need not include the whole range of the scale for any given set of images.

Some of the images include arrow cursors and cursors in the shape of red and blue circles. These cursors should **not** be considered in the rating that you give the image.

Please attempt to ensure that your ratings remain consistent throughout the set of images.

Raters will be paid £50 on return of a complete set of rated images.

Appendix K. Questionnaire Administered in Gesture Study

This appendix contains a copy of the demographics questionnaire administered after the second study described in Chapter 6 investigating communication through the haptic gesture.

Questionnaire

Age: _____

Gender: M / F

Nationality: _____

Did you know the other user (prior to the experiment): Y / N

Rate your level of computer experience (tick one):

None

Basic

Intermediate

Expert

Do not fill out this section

Condition:

Appendix L. Raw Data from Gesture Study

This appendix contains a copy of the raw data gathered during the second study described in Chapter 6 looking at communication through the haptic gesture. It includes all subjective and objective data.

Table L.1. TLX results from decoder visual condition.

	Mental Demand	Physical Demand	Time Pressure	Effort Expended	Performance Level Achieved	Frustration Experienced	Fatigue Experienced
Decoder 1	14	7	17	14	3	10	14
Decoder 2	5	10	4	9	9	5	9
Decoder 3	8	13	14	12	4	12	12
Decoder 4	14	7	17	17	6	3	7
Decoder 5	16	8	12	15	8	14	0
Decoder 6	12	5	4	11	11	6	8
Decoder 7	11	3	8	5	7	7	3
Decoder 8	6	6	10	8	16	6	10
Decoder 9	2	2	9	2	5	13	7
Decoder 10	16	12	12	16	6	14	14
Decoder 11	13	17	10	13	12	9	7
Decoder 12	11	5	12	12	12	10	7
Decoder 13	12	10	10	14	4	8	10
Decoder 14	14	16	10	14	8	4	4
Decoder 15	15	4	16	13	5	5	11
Decoder 16	8	10	18	9	3	19	13
Decoder 17	14	16	17	14	15	11	12
Decoder 18	16	17	17	16	5	17	19

Table L.2. TLX results from decoder haptic condition.

	Mental Demand	Physical Demand	Time Pressure	Effort Expended	Performance Level Achieved	Frustration Experienced	Fatigue Experienced
Decoder 1	14	7	17	14	3	10	14
Decoder 2	7	7	14	9	3	15	6
Decoder 3	16	16	19	16	2	14	15
Decoder 4	17	5	17	16	4	7	10
Decoder 5	17	10	13	16	5	15	4
Decoder 6	14	9	7	10	6	13	11
Decoder 7	13	7	16	7	2	11	4
Decoder 8	15	6	18	10	8	16	14
Decoder 9	10	5	9	2	2	17	9
Decoder 10	15	11	17	16	6	15	11
Decoder 11	16	13	17	16	8	12	10
Decoder 12	15	5	12	10	11	11	4
Decoder 13	16	14	12	18	2	14	16
Decoder 14	17	18	13	18	11	3	5
Decoder 15	13	5	16	14	11	3	13
Decoder 16	8	8	17	15	3	19	17
Decoder 17	18	15	17	17	8	16	13
Decoder 18	18	14	17	18	4	16	16

Table L.3. TLX results from decoder combined condition.

	Mental Demand	Physical Demand	Time Pressure	Effort Expended	Performance Level Achieved	Frustration Experienced	Fatigue Experienced
Decoder 1	14	7	17	14	3	10	14
Decoder 2	9	6	6	7	10	11	7
Decoder 3	8	16	14	14	8	11	13
Decoder 4	11	8	16	15	7	2	5
Decoder 5	14	8	10	14	9	12	0
Decoder 6	11	5	7	11	9	6	8
Decoder 7	11	11	11	12	6	4	9
Decoder 8	14	14	14	17	13	16	16
Decoder 9	6	7	10	5	1	15	11
Decoder 10	10	10	10	9	9	6	7
Decoder 11	12	18	12	15	12	10	7
Decoder 12	13	7	13	11	11	10	10
Decoder 13	8	6	8	10	12	4	8
Decoder 14	8	10	12	12	15	2	0
Decoder 15	13	3	16	13	8	3	11
Decoder 16	8	6	15	8	2	20	16
Decoder 17	13	7	17	15	14	11	11
Decoder 18	14	14	12	14	10	12	16

Table L.4. TLX results from encoder condition.

	Mental Demand	Physical Demand	Time Pressure	Effort Expended	Performance Level Achieved	Frustration Experienced	Fatigue Experienced
Encoder 1	17	10	19	13	10	12	14
Encoder 2	19	12	15	19	5	12	10
Encoder 3	16	2	16	4	10	10	4
Encoder 4	14	3	16	4	16	4	2
Encoder 5	16	12	15	15	10	12	15
Encoder 6	16	2	16	18	15	10	8
Encoder 7	17	3	15	13	8	10	3
Encoder 8	15	3	15	15	13	14	16
Encoder 9	17	10	20	18	10	18	4
Encoder 10	15	8	13	16	6	5	13
Encoder 11	15	10	19	16	10	12	3
Encoder 12	18	4	5	17	14	10	12
Encoder 13	16	4	16	16	10	18	4
Encoder 14	15	17	7	12	13	15	15
Encoder 15	11	15	11	15	10	6	16
Encoder 16	16	7	14	17	11	13	10
Encoder 17	16	12	14	15	11	10	7
Encoder 18	7	1	9	7	8	7	7

Table L.5. Mean trial times (seconds) in each condition.

	Visual	Haptic	Combined
Participants 1	59.0519	66.9162	62.456
Participants 2	62.2245	68.195	58.3489
Participants 3	59.2803	72.3452	68.2491
Participants 4	67.4981	76.9836	68.9972
Participants 5	76.0684	80.3144	72.6093
Participants 6	58.9395	66.0177	63.34878
Participants 7	61.7908	66.9142	71.0432
Participants 8	59.8431	75.097	69.2054
Participants 9	64.2775	75.9932	69.7223
Participants 10	77.7448	78.4657	71.1211
Participants 11	62.0871	75.5757	63.8087
Participants 12	65.984	75.2682	71.5829
Participants 13	48.8876	56.3921	60.3969
Participants 14	54.0081	70.5785	67.386
Participants 15	76.4811	76.6964	75.022
Participants 16	68.4944	66.4185	68.1679
Participants 17	70.257	66.8962	67.651
Participants 18	67.2256	73.9321	56.9758

Table L.6. Similarity ratings (from three raters) for first 90 encoder images.

Condition	R 1	R 2	R 3		Condition	R 1	R2	R 3
Visual	6	4	7		Haptic	2	3	1
Haptic	5	3	3		Haptic	6	6	2
Visual	6	7	8		Visual	6	8	7
Haptic	3	4	4		Haptic	4	2	2
Haptic	5	4	4		Haptic	4	4	1
Haptic	5	6	3		Visual	4	5	5
Combined	6	9	6		Visual	5	7	4
Visual	3	3	7		Visual	5	8	4
Haptic	7	6	7		Haptic	1	4	1
Combined	2	7	3		Visual	6	8	7
Combined	8	7	7		Haptic	4	7	4
Combined	4	3	2		Combined	3	5	5
Combined	2	4	3		Visual	5	8	7
Visual	4	6	4		Combined	4	4	2
Haptic	4	2	2		Visual	5	6	5
Visual	7	7	5		Combined	6	2	2
Haptic	5	3	7		Visual	5	7	3
Haptic	2	5	5		Combined	5	7	6
Combined	6	6	3		Combined	7	8	4
Combined	4	3	3		Combined	5	2	4
Haptic	7	10	6		Haptic	5	6	4
Combined	6	8	9		Combined	0	0	0
Combined	7	7	8		Combined	5	3	6
Visual	7	8	7		Haptic	3	2	1
Combined	3	4	2		Haptic	6	5	3
Visual	4	4	5		Haptic	2	5	5
Haptic	4	4	2		Visual	5	5	5
Combined	4	4	3		Visual	4	7	6
Haptic	4	3	2		Combined	4	4	6
Haptic	5	3	3		Haptic	4	1	2
Haptic	4	1	3		Visual	5	7	5
Haptic	2	2	1		Haptic	4	5	5
Haptic	4	6	4		Combined	3	3	5
Visual	5	6	8		Visual	5	4	4
Haptic	3	3	4		Visual	5	8	6
Combined	1	4	1		Visual	2	7	4
Haptic	5	7	3		Haptic	5	2	4
Combined	5	9	7		Visual	6	4	8
Combined	5	4	5		Haptic	1	2	1
Combined	6	8	7		Haptic	2	2	1
Haptic	4	3	5		Combined	4	3	2
Combined	5	7	5		Visual	3	3	5
Combined	4	3	2		Combined	5	5	3
Haptic	5	4	6		Combined	6	7	4
Combined	4	6	3		Haptic	4	4	4

Table L.7. Similarity ratings (from three raters) for second 90 encoder images.

Condition	R 1	R 2	R 3		Condition	R 1	R2	R 3
Haptic	4	3	1		Combined	4	5	3
Combined	5	4	7		Haptic	2	3	1
Haptic	3	6	4		Haptic	6	4	5
Combined	4	7	1		Haptic	4	2	5
Visual	5	7	2		Visual	7	4	8
Haptic	6	5	4		Haptic	6	4	7
Combined	3	3	1		Haptic	1	2	1
Haptic	6	4	7		Combined	6	8	4
Combined	6	4	2		Combined	6	8	7
Visual	2	2	3		Haptic	4	5	4
Visual	1	7	4		Haptic	5	5	5
Visual	3	5	3		Haptic	2	5	3
Visual	5	3	1		Combined	5	5	2
Visual	5	3	4		Visual	5	4	4
Combined	8	8	9		Visual	6	7	4
Combined	5	5	5		Combined	3	4	1
Haptic	3	6	2		Haptic	2	2	1
Visual	7	8	9		Visual	5	6	6
Combined	4	3	3		Visual	5	8	6
Combined	4	6	6		Visual	6	3	5
Visual	2	3	3		Haptic	2	2	1
Combined	5	3	6		Visual	5	8	2
Combined	3	2	2		Visual	5	4	5
Haptic	4	5	4		Visual	3	2	2
Combined	3	2	2		Visual	6	7	6
Haptic	5	3	3		Visual	5	8	7
Haptic	4	2	3		Visual	7	4	7
Combined	5	4	5		Visual	6	5	6
Visual	4	9	6		Visual	4	2	4
Combined	4	5	3		Visual	5	8	5
Visual	3	3	2		Combined	4	7	4
Haptic	3	6	2		Combined	6	5	6
Visual	2	1	1		Visual	5	4	4
Haptic	4	4	4		Combined	5	7	3
Combined	5	1	6		Combined	7	10	4
Haptic	6	8	6		Haptic	3	3	4
Combined	3	3	4		Combined	5	6	4
Visual	4	4	4		Visual	5	7	2
Haptic	3	4	2		Combined	5	2	2
Visual	4	4	5		Combined	8	7	4
Combined	6	8	7		Haptic	4	2	3
Visual	6	6	4		Haptic	5	7	3
Haptic	5	3	2		Visual	5	6	5
Combined	5	4	3		Visual	5	7	5
Visual	7	9	9		Haptic	5	5	4

Table L.8. Similarity ratings (from three raters) for third 90 encoder images.

Condition	R 1	R 2	R 3	Condition	R 1	R2	R 3
Visual	2	5	3	Visual	4	5	2
Haptic	2	5	4	Visual	5	3	2
Visual	5	7	5	Haptic	4	4	4
Combined	4	4	4	Visual	5	5	4
Haptic	4	3	1	Haptic	5	4	3
Haptic	4	5	4	Combined	7	3	8
Combined	4	4	4	Combined	7	5	4
Combined	1	2	1	Visual	6	3	7
Visual	4	4	4	Haptic	4	4	3
Visual	7	8	4	Combined	4	3	1
Visual	6	5	4	Visual	6	7	9
Combined	5	7	4	Haptic	3	2	2
Visual	2	2	2	Visual	8	7	9
Combined	3	2	4	Combined	5	6	5
Combined	6	7	5	Combined	5	8	3
Haptic	3	5	4	Visual	3	4	4
Combined	6	6	5	Combined	4	3	4
Haptic	4	2	3	Visual	1	3	1
Haptic	4	3	7	Haptic	1	4	2
Visual	6	6	3	Visual	6	4	5
Haptic	4	4	3	Haptic	4	2	5
Visual	5	4	5	Combined	6	7	6
Combined	4	4	3	Haptic	4	3	5
Visual	2	5	1	Visual	5	6	4
Haptic	5	3	2	Haptic	4	3	4
Haptic	5	4	3	Combined	4	2	2
Visual	8	7	6	Visual	8	10	8
Haptic	8	9	8	Combined	5	5	3
Haptic	5	8	5	Haptic	6	4	2
Visual	4	5	4	Haptic	2	1	1
Combined	5	9	2	Combined	6	8	8
Visual	4	9	3	Haptic	4	5	5
Visual	6	7	4	Visual	3	4	4
Combined	8	9	5	Combined	5	6	4
Visual	2	4	2	Visual	5	4	4
Visual	6	9	5	Haptic	1	1	1
Visual	4	6	3	Visual	5	4	3
Visual	4	6	7	Haptic	2	3	4
Visual	5	6	2	Combined	3	4	2
Visual	4	3	4	Haptic	6	7	7
Combined	4	3	4	Haptic	3	2	2
Combined	6	7	3	Haptic	6	6	5
Haptic	3	2	2	Combined	6	7	5
Combined	2	3	1	Combined	6	4	3
Combined	5	7	5	Visual	3	8	3

Table L.9. Similarity ratings (from three raters) for fourth 90 encoder images.

Condition	R 1	R 2	R 3		Condition	R 1	R2	R 3
Haptic	4	4	5		Visual	6	7	5
Visual	5	3	4		Combined	5	4	3
Combined	4	4	4		Combined	5	6	4
Combined	5	3	3		Haptic	3	3	2
Haptic	4	6	4		Haptic	4	5	2
Visual	5	3	2		Haptic	5	3	5
Visual	3	3	1		Haptic	5	2	4
Combined	3	3	2		Visual	5	3	3
Combined	0	0	0		Haptic	3	7	8
Haptic	5	6	4		Combined	3	7	4
Haptic	4	7	7		Combined	2	4	1
Haptic	5	2	4		Haptic	4	4	3
Haptic	5	7	4		Combined	3	4	2
Visual	6	4	8		Combined	0	0	0
Combined	3	7	7		Visual	4	2	2
Combined	6	6	4		Visual	3	2	2
Haptic	3	1	1		Combined	7	8	2
Haptic	2	3	2		Haptic	6	7	6
Combined	7	7	2		Visual	5	7	4
Visual	4	6	2		Combined	3	2	1
Haptic	3	3	3		Haptic	3	3	6
Haptic	4	2	4		Visual	4	5	3
Combined	6	8	7		Visual	3	4	1
Haptic	4	2	3		Combined	6	5	4
Haptic	1	4	1		Visual	7	5	6
Haptic	1	3	1		Visual	3	4	3
Visual	6	3	3		Visual	5	4	2
Combined	8	8	9		Haptic	4	5	3
Haptic	6	7	4		Combined	6	6	7
Haptic	3	2	2		Visual	2	2	1
Haptic	6	4	3		Visual	5	4	2
Visual	7	8	6		Combined	3	3	2
Visual	7	7	6		Visual	7	9	6
Visual	6	8	6		Haptic	6	7	6
Combined	2	2	2		Haptic	4	3	4
Combined	3	6	3		Combined	4	7	4
Visual	6	7	4		Combined	7	7	7
Haptic	5	5	3		Combined	3	7	3
Haptic	4	3	5		Combined	5	6	3
Visual	4	3	7		Visual	2	4	2
Combined	5	7	6		Haptic	3	5	2
Combined	5	7	4		Visual	6	7	7
Visual	6	6	5		Visual	2	4	1
Combined	5	4	8		Haptic	3	7	4
Visual	4	3	4		Visual	3	3	3

Table L.10. Similarity ratings (from three raters) for fifth 90 encoder images.

Condition	R 1	R 2	R 3	Condition	R 1	R2	R 3
Combined	3	6	3	Visual	5	6	6
Combined	6	5	6	Visual	5	6	5
Visual	5	7	5	Combined	2	9	3
Combined	4	5	4	Haptic	7	4	4
Visual	6	5	5	Visual	5	3	3
Visual	6	9	4	Visual	5	3	3
Visual	4	8	3	Visual	4	5	3
Haptic	2	2	1	Combined	1	1	1
Haptic	5	6	2	Haptic	3	6	3
Haptic	3	2	1	Haptic	1	4	1
Combined	6	7	8	Combined	5	6	2
Combined	6	8	6	Combined	3	3	1
Haptic	5	3	3	Visual	6	7	5
Combined	5	4	3	Haptic	8	9	9
Haptic	4	7	2	Combined	5	2	4
Haptic	3	3	2	Visual	4	6	3
Combined	4	4	5	Combined	4	3	4
Combined	6	8	5	Combined	6	9	6
Combined	5	7	7	Combined	5	5	7
Visual	5	2	5	Haptic	2	2	2
Visual	8	10	9	Combined	4	2	3
Combined	5	8	3	Visual	7	5	6
Visual	7	7	6	Visual	6	7	6
Haptic	4	3	4	Visual	6	5	3
Combined	6	4	5	Combined	5	8	5
Combined	4	8	4	Combined	6	7	5
Visual	4	2	3	Haptic	4	6	3
Combined	5	8	6	Visual	4	4	4
Combined	4	5	4	Haptic	2	2	3
Haptic	4	5	5	Haptic	1	6	3
Combined	3	3	3	Haptic	5	6	5
Visual	6	3	3	Haptic	5	7	3
Visual	6	5	4	Haptic	4	7	6
Haptic	6	6	4	Combined	4	6	3
Haptic	6	5	7	Combined	5	8	3
Haptic	3	2	3	Haptic	4	6	4
Combined	4	7	4	Haptic	7	7	9
Visual	5	4	3	Haptic	4	2	9
Combined	5	4	4	Haptic	8	5	6
Combined	1	4	1	Haptic	6	8	8
Visual	4	4	3	Visual	3	1	2
Haptic	4	7	5	Combined	5	8	3
Visual	6	7	6	Visual	5	6	5
Haptic	3	6	5	Visual	4	4	6
Haptic	2	6	5	Combined	5	4	2

Table L.11. Similarity ratings (from three raters) for sixth 90 encoder images.

Condition	R 1	R 2	R 3		Condition	R 1	R2	R 3
Visual	4	4	2		Visual	6	8	5
Visual	6	6	7		Combined	6	6	8
Visual	6	6	4		Visual	4	5	4
Combined	4	5	3		Haptic	7	7	5
Visual	5	7	2		Combined	3	2	3
Combined	6	4	4		Combined	7	8	8
Visual	7	6	3		Combined	4	8	2
Visual	7	8	5		Combined	4	4	6
Visual	3	3	2		Visual	5	4	8
Combined	5	5	6		Combined	3	4	4
Visual	7	5	4		Haptic	3	3	3
Haptic	4	5	3		Combined	4	5	3
Haptic	6	7	7		Visual	4	2	4
Haptic	4	6	3		Haptic	6	7	5
Combined	3	2	2		Visual	5	7	8
Combined	6	6	6		Haptic	5	6	4
Haptic	2	2	4		Visual	4	3	4
Haptic	5	7	3		Visual	5	4	5
Haptic	6	5	6		Haptic	6	7	2
Haptic	4	2	4		Visual	5	9	6
Combined	5	9	8		Haptic	2	2	2
Visual	4	5	2		Combined	3	5	4
Combined	5	7	7		Combined	3	6	3
Combined	4	2	3		Combined	5	3	3
Combined	2	4	2		Haptic	2	6	4
Haptic	4	2	6		Combined	4	2	4
Visual	3	7	4		Haptic	6	4	4
Combined	5	6	8		Combined	4	6	4
Visual	5	6	5		Visual	5	7	4
Haptic	5	4	4		Combined	8	3	3
Combined	6	8	5		Visual	3	4	1
Haptic	7	7	4		Haptic	3	2	3
Haptic	5	5	4		Visual	3	3	4
Haptic	4	2	3		Visual	5	8	4
Visual	5	7	4		Combined	3	4	2
Haptic	3	7	2		Visual	4	2	4
Combined	5	4	5		Visual	7	7	7
Haptic	4	4	4		Visual	6	7	5
Visual	5	7	4		Haptic	4	3	3
Haptic	2	2	2		Combined	5	9	8
Visual	5	6	7		Haptic	3	6	4
Haptic	2	2	2		Visual	3	3	4
Combined	8	9	9		Combined	7	6	8
Haptic	3	3	5		Combined	4	6	7
Haptic	4	5	4		Visual	6	6	4

Appendix M. Sample of Images Used in Rating Process in Gesture Study

This appendix contains a sample of the images used in the rating process described in the second study in Chapter 6 investigating communication through the haptic gesture. The images presented here have been through the rating process. All data relating to 2 specific source images are presented, each occupying 3 pages. In each case, the a large version of the source image is followed by 2 pages each containing 9 of the copied images. The printed code visible in the top-left corner of each copied image served to uniquely identify it, while the written numbers in the bottom right were added by the rater.

Figure M.1. Original image (K42) used in rating process.



K42

Figure M.2. First nine copies used in rating process for image K42.

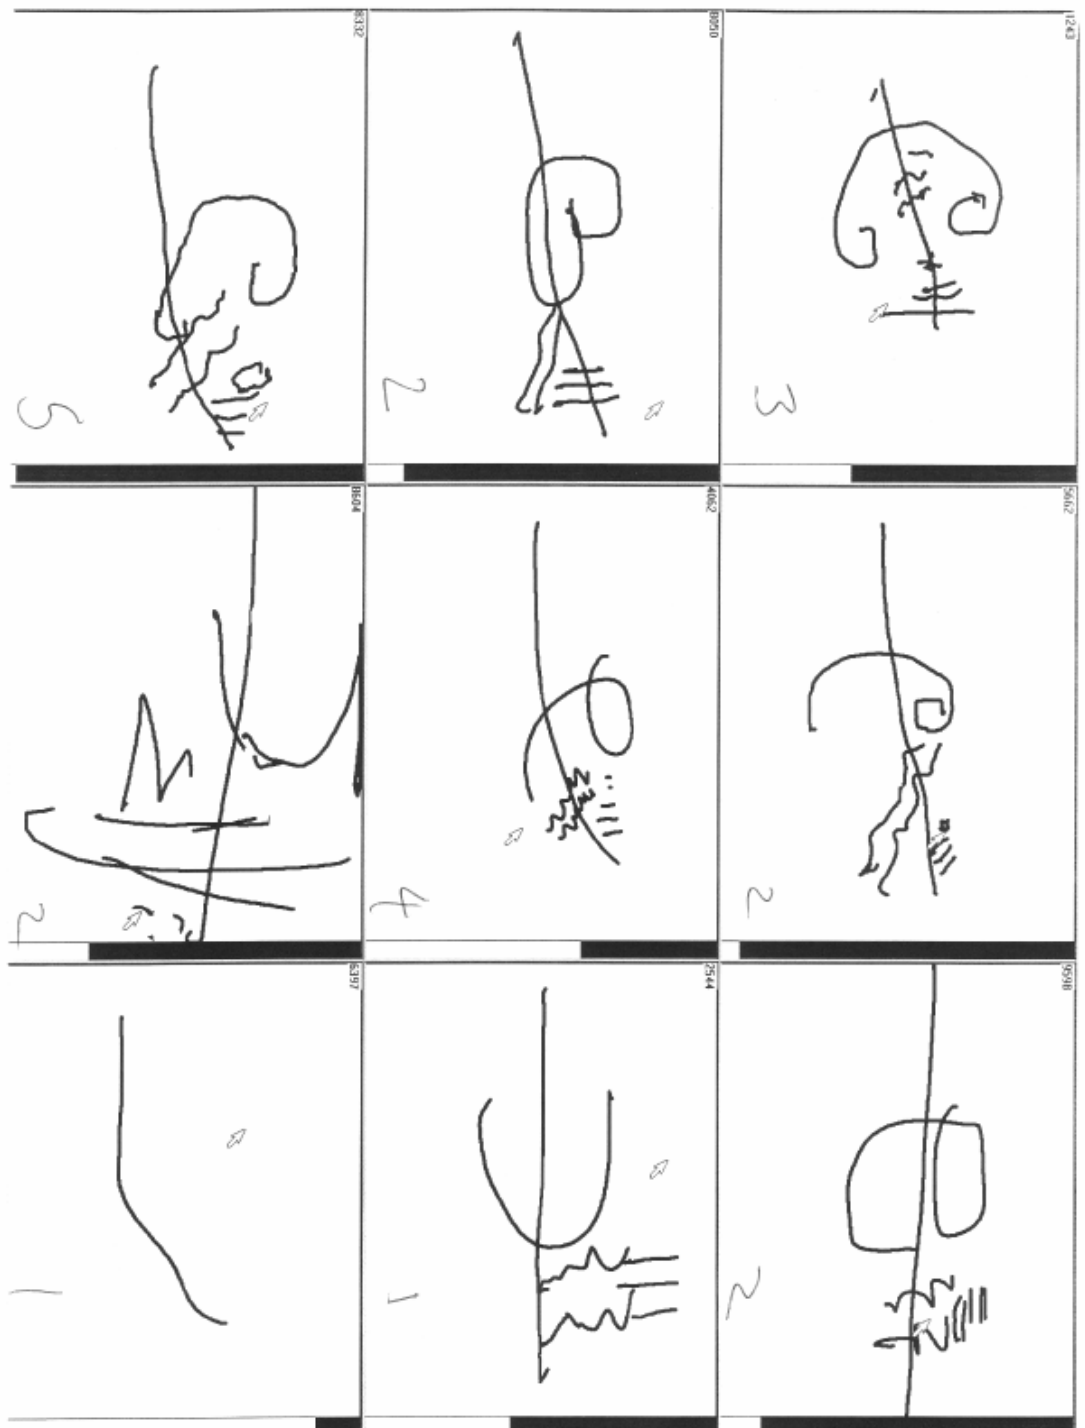


Figure M.3. Second nine copies used in rating process for image K42.

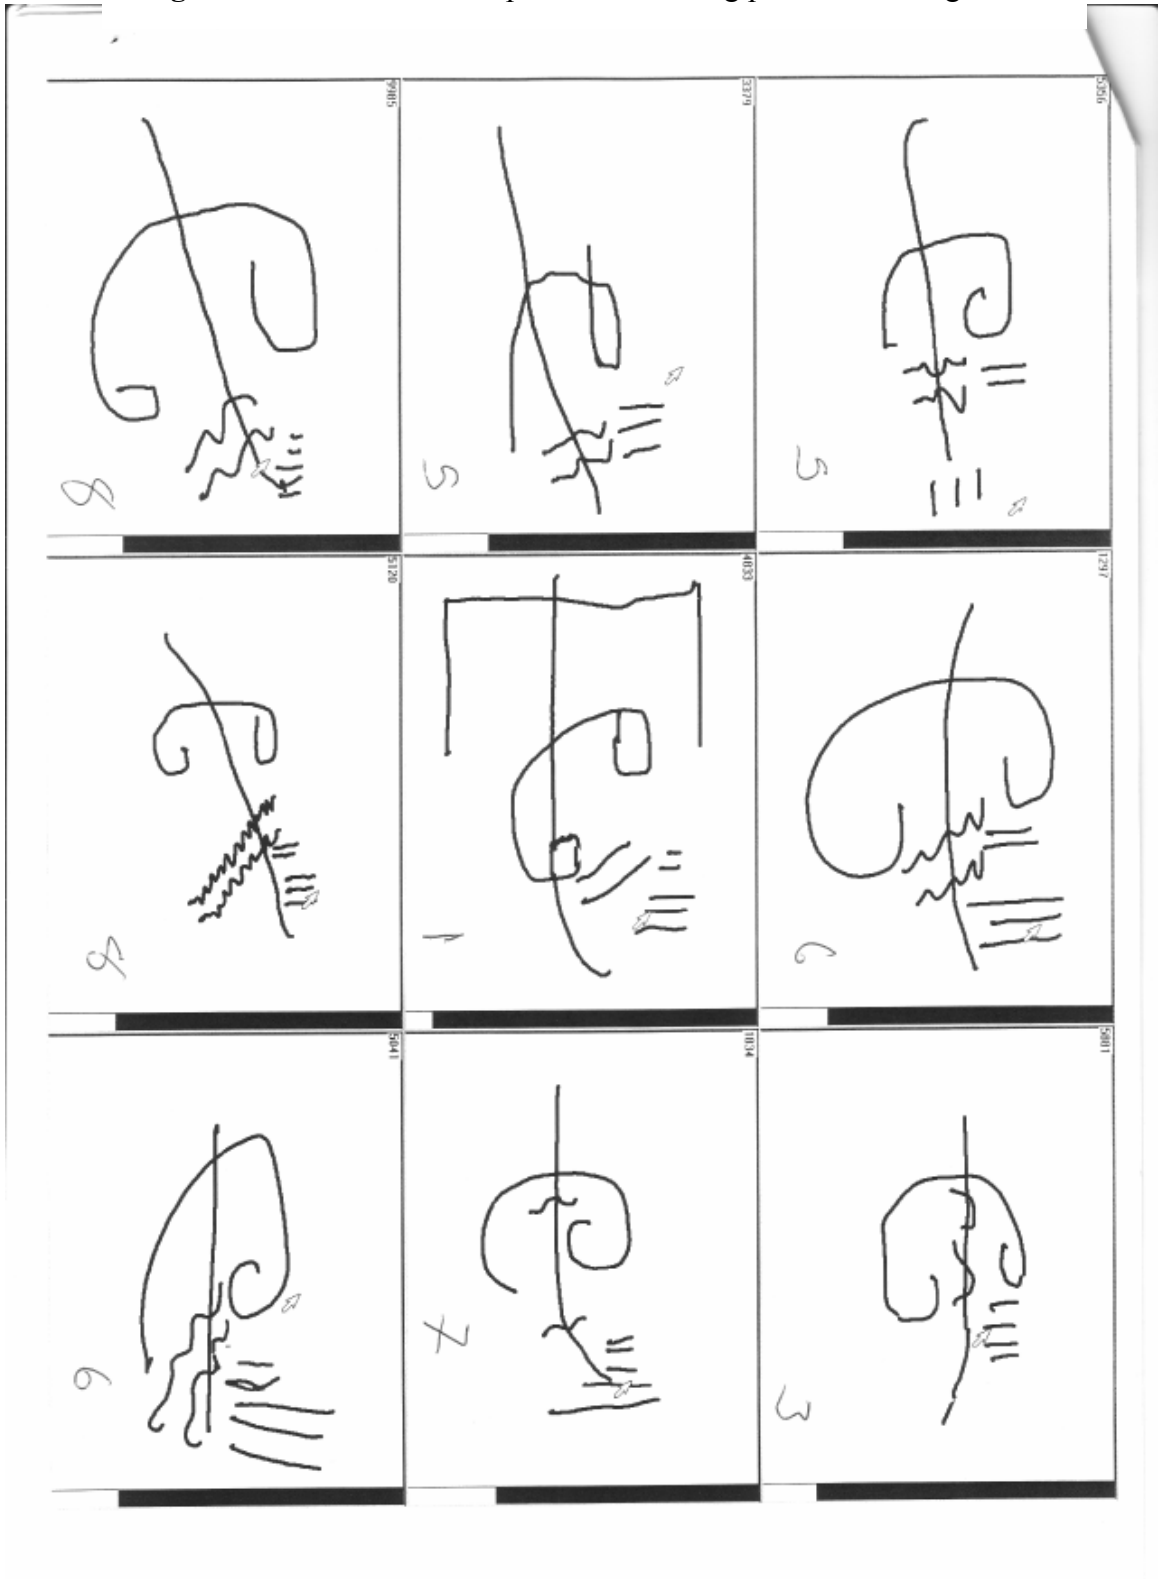


Figure M.4. Original image (K11) used in rating process.



Figure M.5. First nine copies used in rating process for image K11.

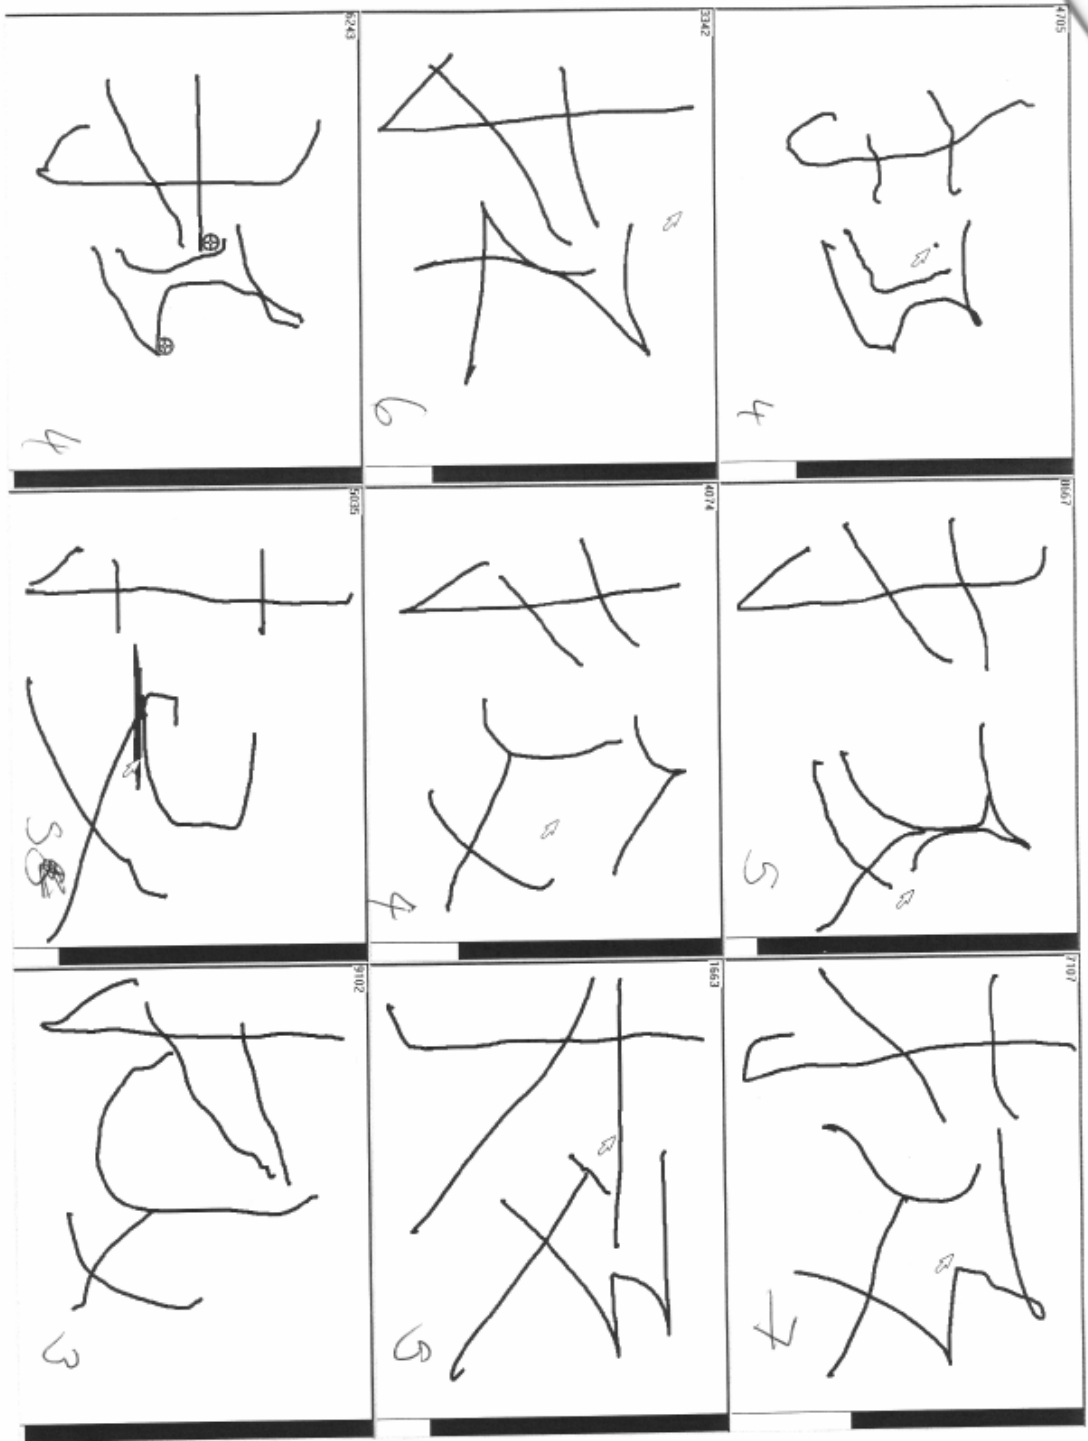


Figure M.6. Second nine copies used in rating process for image K11.

