

**UNIVERSITY OF GLASGOW**



**Investigating a Multimodal Solution for  
Improving Force Feedback Generated Textures**

**Marilyn Rose McGee**

Thesis submitted for the degree of Ph.D.  
Department of Computing Science

September 2002

# Abstract

---

This thesis empirically investigates the use of multimodal augmentation as an approach to improving the quantity and quality of force feedback generated textures. A review of haptic interaction issues indicated that haptic effects, such as texture, are being used successfully to convey rich information rather than simply to increase immersion or realism. This review also identified graphical user interaction as a potential application area that remains relatively unexplored and that might benefit from haptic interaction. An investigation was conducted as part of this research to empirically assess the effect on interaction of different haptic effects, including texture, used to augment a conventional desktop graphical user interface (GUI). The results from this work, as well as a review of texture simulation, confirm that texture as a haptic effect has the potential to enhance many applications but that effective simulation of this effect still needs investigation.

A set of experiments is presented that empirically investigate the effects of multimodally augmenting a simple model of force feedback texture with simple auditory texture cues. The results showed that judgments of perceived roughness of the force feedback and auditory textures varied as a function of the geometry of the model of the textures. Importantly judging the textures multimodally was also shown to have a significant effect on the perceived roughness of virtual surfaces but that this relationship may be more complex. The thesis shows that multimodally augmenting force feedback effects can serve to improve the quality of interaction available through current force feedback technology.

# Contents

---

<b>CHAPTER 1: INTRODUCTION .....</b>	<b>1</b>
1.1 THE EMERGENCE OF NEW HAPTIC INTERACTION ISSUES .....	1
1.2 HAPTIC INTERACTION IN GRAPHICAL USER INTERFACES.....	2
1.3 SIMULATING TEXTURE .....	3
1.4 FORCE FEEDBACK TEXTURES.....	3
1.5 THESIS AIMS .....	3
1.6 OVERVIEW OF STRUCTURE OF THESIS.....	4
<b>CHAPTER 2: HAPTIC INTERACTION .....</b>	<b>5</b>
2.1 HAPTIC TERMINOLOGY .....	5
2.2 THE HUMAN HAPTIC SYSTEM.....	7
2.2.1 <i>The cutaneous sense</i> .....	7
2.2.2 <i>The kinesthetic sense</i> .....	8
2.3 HAPTIC INFORMATION PROCESSING.....	8
2.3.1 <i>Unique attributes of the haptic sense</i> .....	9
2.3.1.1 <i>Bidirectionality</i> .....	9
2.3.1.2 <i>Multiparametered</i> .....	9
2.3.1.3 <i>Active versus passive perception</i> .....	9
2.3.1.4 <i>Spatial and temporal structure</i> .....	10
2.3.1.5 <i>Recall and association</i> .....	10
2.4 HAPTICS IN HUMAN COMPUTER INTERACTION .....	11
2.4.1 <i>Haptic technology</i> .....	11
2.4.1.1 <i>Tactile displays</i> .....	12
2.4.1.2 <i>Force feedback displays</i> .....	12
2.4.1.3 <i>The PHANToM force feedback device</i> .....	15
2.4.2 <i>Potential uses for haptics in human computer interaction</i> .....	17
2.5 NEW HAPTIC INTERACTION ISSUES .....	19

<b>CHAPTER 3: AN INVESTIGATION OF HAPTIC EFFECTS IN GRAPHICAL USER INTERFACES (GUIs)</b> .....	<b>20</b>
3.1 INTRODUCTION .....	20
3.1.1 <i>Motivation and Background</i> .....	20
3.1.2 <i>Related work by Oakley et. al. (2000)</i> .....	22
3.2 AN INVESTIGATION OF HAPTIC EFFECTS IN A SCROLL BAR TASK .....	22
3.2.1 <i>Device and Software</i> .....	23
3.2.2 <i>Haptic Effects</i> .....	23
3.2.3 <i>General Measures used in the Experiment</i> .....	24
3.2.4 <i>Participants</i> .....	24
3.2.5 <i>Design</i> .....	24
3.2.6 <i>Procedure</i> .....	25
3.2.7 <i>Hypotheses</i> .....	26
3.3 RESULTS FROM AN INVESTIGATION OF HAPTIC EFFECTS IN A SCROLL BAR TASK .....	27
3.3.1 <i>Timing results</i> .....	27
3.3.2 <i>Error results</i> .....	27
3.3.3 <i>Workload results</i> .....	27
3.4 GENERAL DISCUSSION OF GUI EXPERIMENT .....	28
3.5 CONCLUSIONS .....	30
<b>CHAPTER 4: SIMULATING TEXTURES</b> .....	<b>32</b>
4.1 REAL WORLD TEXTURES .....	32
4.1.1 <i>What are textures used for?</i> .....	32
4.1.2 <i>What defines a texture?</i> .....	32
4.1.3 <i>How are textures perceived?</i> .....	33
4.1.4 <i>The complex nature of texture</i> .....	34
4.1.5 <i>Roughness as the primary dimension of texture</i> .....	35
4.2 VIRTUAL TEXTURES .....	36
4.2.1 <i>What are virtual textures used for?</i> .....	36
4.2.2 <i>How are force feedback textures generated?</i> .....	37
4.2.3 <i>How are force feedback textures perceived?</i> .....	39
4.3 IMPROVING FORCE FEEDBACK GENERATED TEXTURES .....	41
4.3.1 <i>Problems encountered in exploration of force feedback texture</i> .....	41
4.3.2 <i>Brief introduction to multimodal augmentation</i> .....	42
4.4 SUMMARY .....	43
<b>CHAPTER 5: MULTIMODAL AUGMENTATION</b> .....	<b>44</b>
5.1 INTRODUCTION .....	44
5.2 PROCESSING MULTISENSORY INFORMATION IN THE REAL WORLD .....	44
5.2.1 <i>Intersensory interaction</i> .....	44
5.2.2 <i>Evidence of intersensory interaction in the real world</i> .....	46
5.2.3 <i>The effects of congruency</i> .....	47

5.3 THE USE OF MULTISENSORY INFORMATION IN HUMAN COMPUTER INTERACTION .....	47
5.3.1 Evidence of intersensory interaction in human computer interaction.....	48
5.3.1.1 Visual and haptic interaction .....	48
5.3.1.2 Audio and haptic interaction.....	49
5.3.2 Improving haptic user-computer interaction.....	49
5.4 MULTIMODAL AUGMENTATION OF FORCE FEEDBACK TEXTURES.....	50
<b>CHAPTER 6: MULTIMODAL ROUGHNESS EXPERIMENTS.....</b>	<b>52</b>
6.1 MOTIVATION AND GOALS.....	52
6.2 GENERAL EXPERIMENTAL OVERVIEW.....	53
6.2.1 General Design .....	53
6.2.2 Texture Stimuli.....	54
6.2.2.1 Haptic Virtual Textures .....	55
6.2.2.2 Auditory Textural Stimuli.....	56
6.2.2.3 Multimodal Textural Stimuli .....	56
6.2.3 Congruency and incongruency manipulation.....	57
6.2.4 Measuring Perceived Roughness .....	58
6.2.5 Device and software.....	59
6.2.6 The Interface and Virtual Workspace.....	59
6.3 UNIMODAL HAPTIC ROUGHNESS EXPERIMENT - EXP. 1(A).....	61
6.3.1 Experimental Design.....	61
6.3.2 Hypotheses .....	61
6.3.3 Results and discussion (Exp.1 (a)) .....	61
6.3.3.1 Effects of frequency on perceived roughness .....	61
6.3.3.2 Reliability of roughness judgments .....	63
6.3.3.3 Identical haptic stimuli .....	64
6.3.3.4 Different haptic stimuli.....	64
6.3.3.5 Two distinct notions of haptic virtual roughness .....	65
6.4 EXTENSION OF UNIMODAL HAPTIC ROUGHNESS EXPERIMENT - EXP. 1(B).....	65
6.4.1 Experimental Design.....	65
6.4.2 Hypotheses.....	67
6.4.3 Results and Discussion (Exp. 1(b)) .....	67
6.4.3.1 Effects of frequency on perceived roughness .....	67
6.4.3.1 Effects of frequency separation on perceived roughness .....	68
6.5 UNIMODAL AUDITORY ROUGHNESS EXPERIMENT - EXP. 2 .....	69
6.5.1 Experimental design (Exp.2) .....	70
6.5.2 Hypotheses.....	70
6.5.3 Results and discussion (Exp. 2) .....	71
6.5.3.1 Effects of frequency on perceived roughness.....	71
6.5.3.2 Effects of frequency separation on perceived roughness .....	72
6.5.3.3 Summary of auditory roughness results.....	73
6.5.3.3 Subjective auditory roughness results .....	74
6.6 MULTIMODAL ROUGHNESS EXPERIMENT - EXP. 3 .....	76
6.6.1 Experimental Design.....	77
6.6.2 Hypotheses.....	78
6.6.3 Results and discussion (Exp.3) .....	80
6.6.3.1 Effects of frequency on perceived roughness .....	80
6.6.3.2 Effects of frequency separation on perceived roughness.....	82

<b>CHAPTER 7: DISCUSSION OF EXPERIMENTAL RESULTS.....</b>	<b>85</b>
7.1 THE PERCEIVED ROUGHNESS OF FORCE FEEDBACK TEXTURES.....	85
7.1.1 <i>Two distinct notions of haptic virtual roughness.....</i>	85
7.1.2 <i>Reliability of the haptic roughness judgments.....</i>	86
7.1.3 <i>Judging identical unimodal haptic stimuli.....</i>	87
7.1.4 <i>Judging different haptic stimuli.....</i>	87
7.2 THE PERCEIVED ROUGHNESS OF AUDITORY TEXTURES.....	87
7.2.1 <i>The auditory function compared to the force feedback function.....</i>	88
7.3. THE EFFECTS OF MULTIMODALITY ON PERCEIVED ROUGHNESS.....	88
7.3.1 <i>Fusion of the unimodal texture percepts.....</i>	88
7.4 THE METHODOLOGY USED.....	89
7.4.1 <i>The nature of the device.....</i>	89
7.4.2 <i>The nature of exploration.....</i>	90
7.4.3 <i>The nature of the experimental paradigm.....</i>	90
7.4.4 <i>The model of texture used.....</i>	90
7.5 EMERGING GUIDELINES FOR FORCE FEEDBACK TEXTURE SIMULATION.....	91
7.5.1 <i>Perception of roughness via force feedback interaction is possible.....</i>	91
7.5.2 <i>Perception of roughness via simple auditory cue is possible.....</i>	91
7.5.3 <i>Multimodal augmentation can improve force feedback textures.....</i>	91
<b>CHAPTER 8: CONCLUSIONS AND FUTURE WORK.....</b>	<b>93</b>
8.1 CONTRIBUTIONS OF THIS THESIS.....	93
8.1.1 <i>Haptic characterisation.....</i>	93
8.1.2 <i>Empirical investigation of haptic effects in a graphical user interface.....</i>	93
8.1.3 <i>Empirical evaluation of a multimodal augmentation approach to improving force feedback texture.....</i>	93
8.1.3.1 <i>Unimodal haptic virtual roughness.....</i>	94
8.1.3.2 <i>Unimodal auditory virtual roughness.....</i>	94
8.1.3.3 <i>Multimodal virtual roughness.....</i>	94
8.2 SUGGESTIONS FOR FUTURE WORK.....	94
8.2.1 <i>Investigating other properties of the experience.....</i>	94
8.2.2 <i>Investigating other modalities.....</i>	95
8.2.3 <i>Other issues.....</i>	95
8.3 FINAL COMMENT.....	96

# Acknowledgments

---

I would like to thank both my first and second supervisors Philip Gray and Stephen Brewster. Together they have managed to provide far more than just the required supervisory feedback and guidance over the last 4 years. They have supported everything I have endeavored to achieve throughout my Ph.D. and nurtured all my ideas and methods even in my own moments of doubt when I needed brought back on track.

I would also like to thank my friends for constantly reminding me that I could do it and knowing exactly how to lift my spirits when all the hard work seemed too much. A large thank-you to Peter Lennon for keeping me smiling during the last 4 years and constantly reminding me what really matters in life.

And finally, a big thank you to my family, especially my wonderful mum and dad for never once expecting me to do anything extraordinary but always being quietly proud when they thought that I had. Everyone's constant support has meant that I could complete this thesis.

This research was supported under EPSRC project GR/M44866 and EPSRC studentship 98700418.

# Declarations

---

The empirical work presented in Chapter 3 on haptic effects in a graphical user interface was published in the Proceedings of ACM CHI 2000, The Hague. The first experiment from this publication was conducted by Ian Oakley in conjunction with my own, the second experiment reported in the paper. Oakley's work is reviewed in detail as a related work section in this thesis and the remaining experimental work presented in this Chapter is the authors own.

The remaining work has been published (as detailed in the references and appendix G) with co-authors, and supervisors, Philip Gray and Stephen Brewster. This thesis only exploits the work from these publications that are directly attributable to the author.

# Chapter 1: Introduction

---

This thesis makes a contribution to the development of haptically enhanced human-computer interfaces\* by bringing together an understanding of the physical attributes and capabilities of input/output devices with knowledge about the human perceptual and cognitive processes required for effective interaction with these devices. A specific haptic interaction issue (simulating texture) is tackled by considering both the nature of the primary interaction devices (a force-feedback device and audio output device), and the human perceptual and cognitive processes, (the processing of multisensory haptic and auditory information), required during interaction with the device for a given haptic task - judging the roughness of virtual textures.

## 1.1 The emergence of new haptic interaction issues


The field of computer haptics is not entirely new. The earliest haptic devices were teleoperation devices. The first teleoperator systems were developed at Argonne National Laboratories for chemical and nuclear handling [Goertz, 1964]. Examples of such systems have been both copied and adapted widely throughout the world. These early systems used wires, tapes and pulleys to couple a master robotic arm to a slave arm. The human operator exerts a force on the master, usually via a handle, and the slave copies the resulting motion of the master. Typical examples of such teleoperation are the handling of nuclear materials (dangerous activity), control of extremely small models (impossible activity) and space and underwater exploration (hazardous and expensive activity). The development of teleoperation since has been fast. Adaptation of technologies such as video and force feedback has made even more sophisticated virtual environments possible. The development of haptic devices in particular has been rapid.

The introduction of a wider variety of haptic devices, many of which can be placed next to a desktop computer has in turn greatly increased the number and variety of areas of computing which might benefit from intentional interaction via our haptic sense (Fig. 1.1). More and more applications have become viable over the last 20 years and therefore a newer generation of computer haptics has emerged. Teleoperation is no longer the only haptic application. Engineers and scientists manipulating robotic arms are no longer the only potential users of haptic devices. Importantly, haptic feedback can be used to convey a more complex, a more varied, and a more rich set of sensations and information to users of haptic technology.

---

\* Many of the terms used here relating to haptics are defined explicitly in the glossary (Appendix B) and also explained in greater detail in chapter 2.

With a better selection of high fidelity and commercially available devices emerging, and the potential applications exploding, there are a variety of new and important haptic interaction issues arising (Fig 1.1). These interaction issues need to be addressed if the field of computer haptics is to reach its full potential. This thesis addresses one such haptic interaction issue: the problem of cost-effective texture simulation.



**Figure 1.1: The emergence of new haptic interaction issues.**

With a better selection of high fidelity and commercially available devices emerging, and the potential applications exploding, there are a variety of new and important haptic interaction issues arising (Fig 1.1). These interaction issues need to be addressed if the field of computer haptics is to reach its full potential. This thesis addresses one such haptic interaction issue: the problem of cost-effective texture simulation.

## **1.2 Haptic interaction in graphical user interfaces**

Modern haptic devices are being designed with the desktop in mind. They are smaller, neater, more robust, and some are becoming affordable to mainstream computer users. Haptic interaction can provide much richer information than was previously available through simple contact cues to enhance immersion in teleoperation for example. With both these facts in mind, haptic effects have been added to common desktop interfaces to provide information regarding what button a user is over or what operation they are executing during desktop interaction.

## **1.3 Simulating texture**

Texture is used frequently in the real world to decide an object's identity and function when interacting with the environment. It also helps us tell the position and location of objects and assists when distinguishing between different objects. Texture information could be equally useful in virtual environments. Virtual objects can be made to appear more realistic for example when given a textured surface that approximates what we would expect to encounter in real life. Texture could also be used in virtual environments however as an extra channel through which to communicate meaningful information about virtual surfaces or objects. People often resort to their sense of touch (haptic modality) to extract textural information from real objects. It is likely then that texture will be an important attribute to display through force feedback interaction in haptic virtual environments. It is crucial therefore that the simulation of convincing texture via force feedback interaction be investigated.

## **1.4 Force feedback textures**

Although it is often assumed that most textural information is experienced cutaneously (via sensors in our skin), there is also a force-based nature to our texture perception (Katz, 1925). That is, the profile of a surface actually displaces our finger or hand as well as just stimulating sensors on the very surface of our skin. It is not unreasonable to assume therefore that force feedback devices could display textures to the user in the form of a series of forces meant to replicate the texture of a surface. Often these forces are much larger than those we would feel through our skin. Despite evidence to show that texture can be displayed and perceived through these devices, many attempts to simulate texture through force feedback interaction in fact results in users being perturbed from the textured objects (Chapter 3 for example). It would be helpful therefore if texture available through such devices could be improved such that realistically sized areas of a workspace could be textured in a useful and useable way. Virtual objects could then be classified by, or distinguished between, using textural properties of their surface.

## **1.5 Thesis aims**

Texture perception is complex and, consequently, it is difficult and costly to create genuinely realistic and effective simulations of texture using force-feedback input/output devices. It would be valuable, therefore, to find ways of improving simulated textures without having to pursue realism. The aim of this thesis is to investigate the potential of one such approach, that of the audio augmentation of haptic virtual textures generated via force-feedback devices as a means of increasing the quantity and quality of the information conveyed by a user's experience of such textures. The investigation is intended to provide results that can be utilised by designers of human-computer interfaces that include, or might include, force-feedback generated texture information.

## 1.6 Overview of structure of thesis

Chapter 2 presents a brief review of the range and nature of haptic technology currently available. There is a particular emphasis on devices falling into the category of force feedback (kinesthetic based) technology as opposed to tactile (cutaneous based) technology. The need for this classification is also discussed. Emerging haptic applications are discussed and new types of rich information that haptics can provide are highlighted.

Chapter 3 is an investigation of several haptic effects in a standard Graphical User Interface (GUI) environment. This is a recent area in which haptic interaction may benefit a range of users and applications. This experimental work shows (amongst other things) a need to look much closer at the ability of force feedback devices to convey *useful* texture information in the correct contexts.

In Chapter 4, the problem of simulating haptic texture is discussed in detail. The nature of real human texture perception is reviewed as well as previous attempts to simulate texture via haptic technology. This thesis argues that despite the scope for technology to advance, and despite the possibility of increased sophistication of models and algorithms for virtual textures, there may be more immediate and cost effective solutions to the problem of improving the simulating force feedback texture.

It is argued in Chapter 5, that the ability of human users to integrate multisensory information in a variety of sophisticated ways can be exploited in haptic interaction to produce multimodal textures with a possible increase to both their realism and their resolution. Multimodal augmentation of stimuli and multisensory integration of information will be reviewed in the context of human computer interaction in particular and proposed as a method of improving the simulation of texture via force feedback interaction devices.

Chapter 6 presents four experiments evaluating the perception of roughness of multimodal textures. The first experiment evaluates the perceived roughness of a set of unimodal haptic textures. The second evaluates the perceived roughness of an equivalent set of unimodal auditory textures. The third and final experiment evaluates the perceived roughness of a set of multimodal (haptic-auditory) textures in relation to the perceived roughness of the force feedback textures alone. The effects of the textures being multimodal are studied in this experiment as well as the effects of the multimodal textures being either congruent or incongruent.

Results from the multimodal roughness experiments are discussed in Chapter 7. The potential effects of combining haptic and auditory stimuli to convey information regarding virtual texture are discussed in terms of possible guidelines for designers of multimodal interfaces. In particular, the possibility of increasing the range and resolution of textures currently available through force feedback interaction is discussed.

Finally, Chapter 8 will summarise the main contributions that this thesis makes to the field of haptic and multimodal interaction. The possible shortcomings of the work carried out will also be reviewed and suggestions made for further work to resolve some of the issues discussed throughout the thesis.

# Chapter 2: Haptic Interaction

---

The word *haptic* generally refers to something that is associated with the sense of touch. More specifically *haptics* concerns the acquisition of information and/or the manipulation of objects through our sense of touch. *Computer haptics* is the field concerned with the techniques and processes associated with generating and displaying haptic stimuli to the human user (Srinivasan, 1989). *Haptic feedback* therefore is one of many available interaction media comparable to visual feedback (graphics) and auditory feedback (sound). It has similar properties to the more common interaction media of graphics and audio as well as unique properties that make it more (and perhaps less) suitable in certain contexts.

*Haptic interaction* is one of the most fundamental ways in which people perceive and effect changes in the real world around them. It is easy to realise the importance of our haptic sense in our everyday interactions in the non-computing sense. The haptic system can also be helpful when interacting with computers just as vision and audition are. *Multimodal Interfaces* exploit multiple sensory modalities in user interaction. It is quite possible that the integration of haptic input and/or output may be one of the major solutions to the problem of developing increasingly effective multimodal interfaces in computing.

Progress has been made in adding visual and auditory feedback to interfaces but haptic feedback has until recently largely been ignored. This sense has been neglected not because it is not useful but rather because it appears more complex than our other senses. In order to realise when and how haptics can help in interaction applications, knowledge of the human haptic sense is needed as well as knowledge of available haptic devices and an understanding of the nature of the resulting haptic user-computer interaction. This chapter presents a brief review of the human haptic system and how people use haptics in the real world. It also reviews the nature of current haptic devices and emerging haptic applications in order to discuss some of the new interaction issues arising in the field of computer haptics (Figure 1.1).

## 2.1 Haptic Terminology

It is common to hear people talking about their sense of 'touch' but less often to hear a reference to the haptic sense. The term 'haptic' is used more frequently with the increasing use of touch in computing. The terminology used to describe this sense is not as firmly established as that for the visual or auditory senses. Many different terms with many different definitions are still being used throughout the literature to describe haptic interaction and this makes it difficult to achieve a common understanding in the research. To rectify this a set of haptic definitions has been proposed (Table 2.1).

<b>Term</b>	<b>Definition</b>
Haptic	Relating to the sense of touch.
Proprioceptive	Relating to sensory information about the state of the body (including cutaneous, kinesthetic, and vestibular sensations).
Vestibular	Pertaining to the perception of head position, acceleration, and deceleration.
Kinesthetic	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 frequently the sensation of pressure rather than temperature or pain.
Force Feedback	Relating to the mechanical production of information sensed by the human kinesthetic system.

**Table 2.1: Haptic Definitions (First presented in Oakley, McGee, Brewster, & Gray, CHI 2000).**

These definitions are not entirely novel but rather aim to synthesise some of the meanings commonly presented in both the psychological and computing literature (e.g. Srinivasan, 1997; Lederman, 1979). This will at the very least provide this thesis, and other researchers, with a common vocabulary with which to study haptic interaction.

## 2.2 The human haptic system

This thesis defines the human haptic system to be the entire mechanical, sensory, motor and cognitive components of the body-brain system (Srinivasan & Basgodan, 1997). Haptics is often even more generally defined to be anything relating to the sense of touch. When people refer to 'touch' alone however it is not entirely clear if each of the components of our haptic sense are included in this definition. That is, the term 'touch' is frequently used to relate to an experience predominantly at the surface of the skin (cutaneous). This definition of touch neglects other important aspects of the haptic sense such as those concerned with muscle movement (kinesthetic) and body positioning (proprioceptive). The term 'haptic' should in fact contain each of these sub-components to make an entire haptic sensory system.

Under the umbrella term of haptic however, fall several significant distinctions. Perhaps most important of these for computer haptics is the division between cutaneous and kinesthetic information. There is some overlap between these two categories; critically both can convey the sensation of contact with an object. The distinction becomes important, however, when we attempt to describe emerging technology and the resulting interaction techniques (Section 2.4). Some devices are designed specifically to impinge on the cutaneous sense as a medium and others use primarily the kinesthetic sense for interaction. This distinction is explained in greater detail in the following sections.

### 2.2.1 The cutaneous sense

The skin is the largest of our sensory systems when measured in terms of area of the receptor surface (about 2 square meters). This makes the skin much more accessible than the eye and ear yet compared to vision and hearing, progress in skin research has been slow and the total number of basic principles emerging is still fairly small (Gibson, 1962). Stimulation of the skin informs the organism of what is directly adjacent to its own body. Skin sensitivity is especially acute in those parts of the body that are most relevant to exploring our immediate environment: the hands and fingers, and the lips and tongue for example.

The skin is a complex system mediating the sensory modalities of touch (mechanoreception), temperature (thermoreception), and pain (nociception) to make up the cutaneous sense. Mechanoreception refers to neural events and sensations that result from mechanical distortion or displacement of cutaneous tissue. It includes repetitive displacements, such as vibration, and a single displacement, such as touch or pressure.

Thermoreception refers to the ability to perceive temperature changes on the skin. The sensation of temperature is closely related to the temperature to which the skin has become adapted. Pain is more subjectively thought of as the sensation of hurt. The immediate stimulus for pain at the nerve ending is probably a chemical reaction brought about by damage to the surrounding tissues. It is considered by some as a separate modality, yet, by others it is considered as the product of excessive stimulation of any of the other sense modalities. What is clear is that the cutaneous sense itself can be split into these three separate senses any or all of which can be stimulated at a certain time to convey haptic sensation.

Although interacting with most haptic devices involves contact with the skin, this thesis is primarily concerned with force feedback devices, which rely predominantly on the kinesthetic sense, and therefore the cutaneous sense is not discussed in detail any further.

### 2.2.2 The kinesthetic sense

The vestibular and kinesthetic systems can be considered 'silent' systems. It is not easy to describe what is actually sensed through them. Without them, however, it is difficult to walk around and manipulate everyday items in our environment. It is very difficult to investigate the kinesthetic sense. The psychological data is substantially more limited than in the other senses, including the cutaneous sense. This lack of data is not a reflection of the relative importance of this system. In fact, this system provides us vital information for maintaining our normal, coordinated behaviour. Rather, it is difficult to design a study which is capable of isolating the cutaneous sense from the kinesthetic sense and so it is difficult to study purely and directly. It is very difficult for example to study our ability to detect changes in body position without involving the tactile system. Kinesthesia often involves the tactile sense by default. Many researchers have turned to physiological data to explore the kinesthetic system but this approach may be invasive and there still remains the issue of separating measurements from the muscles from measurements from the joints.

In addition to the kinesthetic *sensory* system, it can be said that the human haptic system also consists of the *motor system* that enables active exploration or manipulation of the environment and a *cognitive system* that can link sensations to perception and action. This thesis considers that all of the systems mentioned are included to some extent during haptic interaction.

## 2.3 Haptic information processing

As previously stated, haptic interaction cannot itself be investigated successfully if the human haptic system, and the devices with which the human user is intended to interact with, are not both well understood and compatible with each other. The visual system and more recently the auditory system have both been well studied in this respect. Haptics must be studied at the human information processing level if the advanced technology that is developing is to be of any real practical use. This section therefore presents a brief overview of some of the unique attributes of the human haptic sense that might make it a medium particularly suitable for certain types of human computer interaction. This understanding is only a step towards determining how haptic feedback might be synthesised and used successfully in interactive applications.

### 2.3.1 Unique attributes of the haptic sense

The haptic sense has certain attributes both at a physiological and psychological level that are unique to this sense and that might make it particularly useful or appropriate in certain interactive tasks. It makes sense that certain tasks, such as getting an overview of a virtual scene, may be better suited to our visual sense because of the ability of our eyes to gather spatial information effectively. Likewise, our auditory sense may be better suited to email reminder cues when we are visually pre-occupied with another task on our computer. There are cases where the nature of the sensory modality may suggest categories of tasks that each modality is best suited to. This section briefly presents some special or unique qualities of haptics and suggests how these might affect haptic interaction design in the way described.

#### 2.3.1.1 Bidirectionality

The haptic sense is a perceptual sense like vision and audition but is also closely tied and coordinated with motor functioning. Our environment isn't experienced purely passively through the haptic sense. We make decisions to move and reach out and grab objects of interest and move and manipulate them to achieve tasks or goals. Because of this our haptic sense is said to be *bi-directional*. That is, we both perceive and actuate via this sense. Much of what is perceived haptically relies on active exploration of space or objects for instance (Lederman and Klatzky, 1987).

#### 2.3.1.2 Multi-parametered

The haptic sense is *multi-parametered*. That is this modality is not a single one but instead is composed of several different sensations, which are conveyed by several different channels. It is the integration of cutaneous, kinesthetic, and proprioceptive information that makes our 'haptic system' complete (see table 2.1). Like vision and audition, the haptic modality has many qualitatively distinct components. Vision is used to perceive colour, depth, and motion for example. Touch is used to perceive force and pressure, temperature, and texture. Even texture, as a parameter, is itself multi-parametered. Textures can be hard, rough, sticky, and wet for example (textures are discussed in detail in Chapter 4). What is certain is that the qualitative variety in haptic sensation makes it challenging to classify and reproduce in computer simulations.

#### 2.3.1.3 Active versus passive perception

With vision and audition it could be argued that the eyes and ears are relatively passive in receiving stimuli compared to the haptic sense. The senses often receive information that we aren't actively seeking. For instance we may overhear someone talking about us at a party without intending to. Our sleep may be disturbed by noisy neighbours, and despite our efforts we cannot prevent our ears from hearing. With vision, it is more likely that we have to direct our attention somewhat to process information via our eyes. We at least have to look at the area of interest for example, albeit we are subjected to peripheral information in the scene.

Haptic interaction needs the most directed contact. An extreme case may be that we do in fact sense our clothes against our skin without intention but when we want to perceive an object in our environment via our haptic sense it is required that we reach out intentionally and explore it actively.

All this does not mean that there is less brain activity required to receive sounds or images. Rather it is simply that we are more limited in our ability to choose to attend to visual or auditory information. Sounds and images are frequently sensed that are of little interest or importance to our main task. These stimuli are sensed and then an active decision is made whether or to process them to a higher level. With haptic information however, a conscious decision often has to be made to actively reach out and make direct contact with the object we wish to sense.

An important distinction is often made in the literature therefore between active and passive haptic perception (e.g. Gibson, 1962). Active haptic perception involves intentional exploration of the object or surface via a conscious movement of our finger, hand, or other body part against the stationary object. Objects and materials can in fact also be haptically experienced passively however. That is, if the finger or hand is kept stationary and the object of interest is instead moved across or against *us* either mechanically or by another person. There are studies in the literature to show that both types of perception are in fact possible but that most certainly some haptic tasks might require active exploration (Gibson, 1962). Active exploration is at the very least what is most common in real world haptic perception. The importance of this distinction for this thesis is merely that the type of haptic perception being studied and referred to hereafter is active perception as defined by Gibson (1962) and explained here.

#### 2.3.1.4 Spatial and temporal structure

Haptic sensations have both a spatial (like visual) and temporal (like auditory) structure. That is, the information that is presented to us when we see something is generally regarded as being spread out spatially so that our eyes can receive a global picture. With sound however, the data generally arrives at our ears as a stream of data spread across time. The structure of haptic displays is not quite as definitive. Haptic information has both a spatial and a temporal structure.

#### 2.3.1.5 Recall and association

The haptic sense can be used for precise control and discrimination but it is less effective when compared to other senses in facilitating recall and association of absolute and relative resolutions (MacLean, 2000). That is, our haptic sense might enable us to discriminate between subtly different levels of a haptic property (such as hardness or roughness) but it is much more difficult to memorize and these different levels. As well as being able to discriminate between different wavelengths to detect colour in vision for example, it is also not difficult to recall the names for different colour hues. It is not clear whether such good recall and association exists for different haptic sensations.

## 2.4 Haptics in human computer interaction

*Haptic feedback* is defined here as computer control over the tactile or kinesthetic properties of a haptic interface. *Haptic interfaces* allow users to touch, feel, and manipulate objects simulated by virtual environments and teleoperator systems (Salisbury and Srinivasan, 1992). The last twenty years has brought significant research efforts in the fields of computer haptics. Simple, active haptic interfaces are now commercially available and research efforts on a range of even more sophisticated devices have intensified. The success of these devices depends on finding application tasks where haptics adds significant value rather than simply novelty. From a design viewpoint, a balance is required between the human haptic ability to sense haptic object properties, the fidelity of the interface device in delivering the appropriate mechanical signals, and the computational complexity in rendering the signals in real time. The nature of haptic devices is reviewed and the potential uses for haptics in computer interaction discussed.

### 2.4.1 Haptic technology

The keyboard, mouse, and trackball are familiar interaction devices that sense a user's hand movements. Although they apply forces on the user's hand upon contact and consequently provide tactile sensation, the forces are not under program control. A haptic device is defined here to provide position input like a mouse but also stimulate the sense of touch by supplying output to the user in the form of forces. Large forces are produced by "force feedback" devices and affect the finger and hand position and movement. Small-scale forces are produced by "tactile" devices and affect the skin surface by stretching and pushing it for example (Oakley *et. al.*, 2000).

Traditionally haptic devices were used in teleoperation (Goertz, 1964). These devices had to be custom ordered and were both complex and very expensive. The required materials such as carbon-fiber composites and rare earth magnets have since become available and affordable to commercial applications and even personal computers have the necessary power to calculate haptic interactions at the required 1000 Hz. The haptics market has recently seen the development of a new breed of lower cost, desktop peripherals and haptic hardware and software are now being developed with the specific goal of improving human computer interaction (Srinivasan & Basgodan, 1997).

Designing and building devices that provide effective haptic communication is made difficult by the lack of knowledge of the human haptic system. Understand the perceptual, motor, and cognitive abilities of the human user and making the physical design of the device compatible with these will improve the quality of haptic experience that these devices can provide. Given that our haptic system is composed primarily of the cutaneous and kinesthetic systems, haptic interaction devices can be classed most simply as either tactile or kinesthetic depending on the major system that they require for use. The definitions presented in this thesis assist in this classification process (Table 2.1).

### 2.4.1.1 Tactile displays

There are five main approaches to tactile display technology (Shimoga, 1993). These are visual, pneumatic, vibro-tactile, electro-tactile, and neuromuscular stimulations. In a visual display, the status of touch of the slave fingers is indicated by the appearance of an icon or via displaying the slave fingertip forces, digitally or graphically to the user. A pneumatics approach uses air jets, air pockets or inflatable bladders to provide touch feedback cues to the operator. Vibro-tactile approaches use vibrating pins, voice coils, or piezoelectric crystals to provide tickling sensation to the human operator's skin to signal the touch. The electro-tactile stimulation method provides electric pulses, of appropriate width and frequency. Finally, the neuromuscular stimulation approach provides the signals directly to the primary cortex of the operator's brain. The choice of the technology must be guided by such factors as the cost, complexity, weight, comfort, noise, power requirement, invasiveness and the extent of liability of the device.

A number of researchers have been seeking a suitable method of encoding language information for use by the skin. An early attempt was by Louis Braille in 1826 who was one of the first to devise a practical and widely accepted system for communicating language to the blind through the skin. None of the attempts so far has satisfied enough of the requirements of a communication system to substitute for hearing or sight. It is important to point out that the weaknesses in tactile devices to communicate information are not due to an inability to develop clever devices. Rather, the weakness is the gaps in knowledge of the functional characteristics of the skin as a sensory system and in the lack of systematic procedures when evaluating the effectiveness of the devices. This is also the case for force feedback interfaces and the kinesthetic sense. Tactile devices are not directly considered any further in this thesis. The focus will now shift therefore to force feedback devices and the discussions thereafter refer specifically to this force feedback interaction.


### 2.4.1.2 Force feedback displays

Force-feedback devices involve the mechanical production of information sensed by the human kinesthetic system. The underlying principle of force feedback devices originates from the belief that force and motion are arguably the most important haptic cues.

"Forces and motion imparted on/by our limbs and fingers contribute significant information about the spatial map of our environment." - Massie (1997).

Information about how an object moves in response to applied forces that arise when we attempt to move objects can provide useful cues to geometry, attributes, and events.

A force feedback system (Fig. 2.1) consists of interface hardware and a computation engine. A mechanical device acts as the interface between the physical and artificial worlds. The computation engine is a processor, which monitors the dynamic motions of the interface hardware through sensors in the mechanism and commands forces to the user by controlling actuators. The computation engine governs the forces felt by the user through control algorithms computed as a function of the sensor measurements. So by simulating the physics of the user's virtual environment forces can be computed in real time and sent to the actuators so that a user can feel them (Srinivasan & Basgodan, 1997).



**Fig. 2.1 Diagram of haptic rendering loop.**

There are three necessary criteria for an effective force feedback interface (Massie, 1997). Force feedback interactions with virtual environments must involve free motion, in which no physical contact is made with objects in the environment. That is, it should not exert external forces on user moving through virtual space. Secondly, solid virtual objects must feel stiff. The maximum obtainable stiffness depends on natural frequencies of the device and also on resolution of sensors and actuators and the servo rate. Most users can be convinced that a virtual surface with a stiffness of at least 20 Nt/cm represents a solid, immovable wall. Thirdly, virtual constraints must not be easily saturated. That is, a virtual wall should feel solid.

Force feedback systems can involve active or passive force feedback. Active force feedback controllers apply forces to the user by adding energy into the human-machine system. Passive force-feedback controllers apply forces to the user by removing energy from the human-machine system. For example, an active controller might use servomotors to generate feedback forces. The strength of the forces could be directly regulated by the computer, which regulates power to the motors. A passive controller might use energy dissipation elements such as a friction brake or a magnetic particle brake. These devices can not directly apply forces to the user; rather they can only apply resistance to the user's motion. The advantage of active force feedback control is that it is inherently general. When using active elements such as servomotors, the system can produce any general force sensation. The advantage of using passive force feedback control is that it is inherently stable and inherently safe for the user. This is because energy dissipation elements only resist motion but do not induce motion. The tradeoff between active and passive feedback is therefore a tradeoff between performance and safety.

Force feedback devices can be either body based or ground based. Body based devices include those that are flexible such as flexible gloves and suits worn by the user and exoskeletal such as jointed linkages affixed to user. These devices fit over and move with the limbs or fingers of the user. Because they are kinematically similar to the arm and hands that they monitor and simulate, they have the advantage of the widest range of unrestricted user motion. As position-measuring systems, body-based devices (gloves, suits, etc.) are relatively inexpensive and comfortable to use. Body based devices with rigid exoskeletons afford force display and slightly more accurate pose sensing, typically at the expense of greater bulk. Regardless of the exact mechanical design, providing force feedback with body based hand controllers remains a difficult problem, placing great demands on minimizing actuator size to make the control bandwidth of the device match human haptic capabilities.

Ground based devices include joysticks, mice, and steering wheels as well as tool-based / pen-based devices. Joysticks are probably the oldest of these technologies and were originally conceived to control light aircraft. Even the earliest of control sticks, connected by mechanical wires to the flight surfaces of aircraft, presented force information about loads on flight surfaces to pilots. Force reflecting joysticks are now commercially available in a wide range of prices and capabilities. Low cost devices (£50 - £1000) with two actuated degrees of freedom (dof) are targeted primarily toward video games (Microsoft Sidewinder, Immersion Impulse Stick, I-Force). Devices with more dof are produced in smaller quantities, generally have higher precision, and cost more (£1000 - £10,000). Joysticks with three actuated dof include the Immersion Impulse Engine 3000, and Cybernet PER-Force 3DOF. Force reflecting mice with 2 actuated dof are also commercially available at low cost (e.g. Immersion FeelIt mouse, around £70).

Force feedback devices have been used in a variety of applications. Notable applications of force feedback hand controllers to virtual environments include project GROPE (Brooks, Ouh-Young, & batter, 1990) which involves a simulator, the Argonne Mechanical Arm (ARM), and more recently the PHANToM, which were used successfully for force feedback during interactions with simulations of molecule docking. The MIT

Sandpaper system is a 3-dof joystick that is capable of displaying virtual textures (Minsky, Ouh-Young, Steele, Brooks, & Behensky, 1990). High performance hand controllers have also been developed by taking advantage of magnetic levitation technology (Salcudean, Wong, and Hollis, 1992). The 2-dof Pantograph has been developed for desktop applications, and the Freedom 7 has been developed for surgical simulation (Ramstein and Hayward, 1994). Per-Force hand controllers were developed in conjunction with NASA (Cybernet, 2000) and the PHANToM was developed at MIT (Massie and Salisbury, 1994).

Force feedback devices such as those described have been used as long as fifty years ago to give humans a sense of presence in remote or dangerous environments. The early developments in force-displaying haptic interfaces were driven by the needs of the nuclear energy industry and others for remote manipulation of materials (Sheridan, 1992). The force-reflecting teleoperator master arms in these applications were designed to communicate to the operator information about physically real tasks. Force feedback devices today are being used to convey a variety of information, in a variety of contexts, that exceeds the richness of simple contact information from remote sites conveyed during teleoperation.

The success of the development of haptic devices relies both on low level psychophysics experiments to evaluate the human haptic system and on cognitive-perceptual research addressing the information processing capabilities of humans as they interact with such devices. It is this latter area of research that remains lacking in the growing field of computer haptics. Progress in haptics does not have to be limited by the reliance on development of new actuator hardware. Existing devices can be better matched to the human haptic system to improve the quality of haptic interaction.

#### 2.4.1.3 The PHANToM force feedback device

Force feedback devices as described in the last section include a new class of cursor-control peripherals. One such device is the PHANToM by SensAble Technology, Inc. (Figure 2.2 – see Massie, 1993 thesis for original descriptions of the device). The PHANToM is a commercial desktop interface with either 3 or 6 dof. It is a ground based, and pen-based based force feedback device. At the time of this writing, the price range for the PHANToM was about £10,000 - £35,000. The PHANToM uses a point force approach to haptic interaction. X, Y, and Z coordinates of the user's finger tip are tracked with optical encoders attached to three motors that control the X, Y, and Z forces exerted on the user's fingertip.

The PHANToM interface allows and measures motion along six dof and can exert controllable forces to the user along three of those dof. The torque from the motors is transmitted through a proprietary cable transmission to a stiff, lightweight aluminum linkage. At the end of this linkage is a passive, three dof gimbal attached to either a thimble or a stylus. Connection to the user through a thimble or stylus means that the PHANToM does not stimulate receptors in the human skin individually. Rather, pressure is distributed over the surface of the skin, with the aggregate result consisting of a force vector that users would mainly perceive through the strain at their finger tip muscles. When the force is transmitted through the PHANToM to the

gimbal, the force is effectively concentrated at the point where the axes of rotation coincide. This point was chosen to be inside the user's finger. Probing an object with the tip of a pencil gives some understanding of the basic mode of touch interaction that the PHANToM system uses.

The user pushes on the thimble or stylus to control the cursor and the thimble or stylus pushes back on the user to simulate physical encounters. To evoke the sensation of touching objects, the geometric, material, kinematic, and dynamic properties of the world to be represented are modeled. Computational methods (haptic rendering) must then be devised to determine the forces that result when the user interacts with the objects. So when the user moves the cursor into a graphical barrier, the phantom pushes back to prevent penetration. Because the Phantom impedes the cursor from passing, the user perceives the graphical barrier as being physically real. More advanced algorithms allow the Phantom to simulate not only hard surfaces but also springs, liquids, textures, vibrations and so on. In fact it can simulate anything else that can be represented mathematically.



**Figure 2.2: The PHANToM by SensAbleTechnology Inc.**

The user's finger position is located with respect to the virtual environment and collisions between the user's finger and the stored geometry of the virtual object are detected. A reaction force vector can then be calculated based on laws of physics and the appropriate force can be applied to the user's finger. The entire servo loop is then repeated (Fig 2.1). Some parallels can be drawn between haptic rendering and real time graphics rendering. Both require calculating surface normals across the geometry of an object. There are also differences in the rendering techniques. 30 -60 Hertz is sufficient for graphics because eyes cannot detect

motion quicker than this. Hands are fairly sensitive to vibrations even at 200 - 300 Hertz. To create convincing sensations of touch it has been found that the loop must occur at a very high rate of typically 1000 Hz or greater. The number of required calculations for each update with haptics is less however. A high resolution computer display has about 1000 by 1000, or 1 million pixels. Each of these pixels must be updated each frame. However, a point force haptic interface has the equivalent of 1 pixel to render each frame. Even standard personal computers have enough speed to execute the required haptic calculations.

The PHANToM is relatively inexpensive, portable, safe, and reliable. The fact that it is a force feedback cursor control device makes it leading edge technology and likely to be so for years to come. For these reasons, this device in particular has been used in this thesis to examine the haptic interaction issue.

## 2.4.2 Potential uses for haptics in human computer interaction

Until the last two decades, most of the work on haptic interfaces has been in the context of the teleoperation of remote robots. Teleoperation involves the transmission of control variables from the human operator to the robotic device, in response to the transmission of feedback variables from the robot to the human (Ellis *et al.*, 1995). There are many application tasks that may benefit from dynamic teleoperation. For example handling materials in hazardous environments (such as radioactive material or biological agents) or remotely operating in an environment in which safe human presence is expensive to achieve and maintain (such as outside a space station or on the ocean floor). More recently haptic hardware and software algorithms are being developed with the specific goal of actually improving human computer interaction in a variety of applications. Devices such as the PHANToM described in section 2.4.1.3 are accelerating the development of such applications. There has been a noticeable shift from touch being a purely robotic engineering issue to an entirely new and multidisciplinary field of computer haptics.

One of the most successful areas to apply force feedback technology is that of critical procedures training. People can learn to perform real-world tasks through the device. The Massachusetts Institute of Technology (MIT) and the Naval Air Warfare Command training Systems Division (NAWC/TSD) describe a virtual workbench for training electronic technicians for example. Perhaps the most successful application of this is in medical procedures training. Medical procedures (for example administering epidural anesthesia, palpating for cancerous lumps) are intrinsically haptic tasks. Haptic displays are required to simulate such tasks for training, because sensing of forces arising from tool-tissue interaction is critical for success. Many medical scenarios can be simulated ranging from routine tasks to risky, hard to perform procedures. Identical scenarios can easily be created and operators can 'undo' mistakes in training. Trainee's performance can also be objectively measured and evaluated. This is not intended to substitute the real thing, but practicing on the virtual could augment and enhance conventional training.

Innovations that have made it easier for most people to use computers have made it more difficult for the blind to use computers. In general, the visually impaired use text to interact with computers. This information

is temporal and not spatial like graphical information. Haptic devices allow blind people to take advantage of the spatial nature of today's computer applications. That is, today's computing environments are frequently windows based and very often exploit multi media by presenting graphs for mathematical data and maps for geographical data for example. Blind and visually impaired users can receive feedback regarding their interaction in the form of haptic cues rather than purely visual cues that are of very limited use to them (Rosenberg, 1994). Haptic interfaces have the potential to make virtual environments accessible to visually impaired users.

People with neuromotor disabilities may experience difficulties in the task of targeting small objects with traditional cursor control peripherals such as the mouse. That is, they do not have the same degree of fine motor control required to perfect the final stages of positioning a cursor accurately over a desired target. With force feedback interaction, users with neuromotor disabilities can use their spatial knowledge of where the target is along with some fine-tuning at the latter stages of targeting from force feedback to guide them into the target they were aiming for. Rosen and Aldestein (1981) have demonstrated that abnormal human tremor can be suppressed using force feedback hand controllers.

Another intrinsically haptic virtual environment task is testing the ease of manual assemble of complex mechanisms before they are manufactured. Designers are usually restricted to using 2D input devices and 2D graphic displays to design 3D parts. Using force feedback however the designer is able to actually feel mechanical constraints of virtual parts. Engineers can feel the characteristics of physical structures. They can feel tension, weight, or friction and manipulate their designs while getting a true "feel" for their structure. In computer aided design (CAD) a designer can also freely manipulate the mechanical components of an assembly in an immersive environment. In this way the designers can build virtual prototypes and adjust the ergonomics without wasting time or money on creating the prototype for users to test out in the field.

Haptic interfaces are still being used with the intent of improving a user's sense of presence but applications wider spread than teleoperation. It is fairly easy to reason why the entertainment industry would be willing to invest in force feedback technology. Haptic interfaces with two or fewer actuated degrees of freedom are now mass-produced for playing personal computer video games, making them relatively cheap, reliable, and easy to program. Although the complexity of the cues they can display is limited, they are surprisingly effective communicators. You can feel the recoil from your weapon, encounter turbulence in your flight simulation and walk into physical walls. If a joystick is vibrated when a player crosses a bridge (to simulate driving over planks) it can provide a landmark for navigation, and signal the vehicle's speed (vibration frequency) and weight (vibration amplitude). In fact, haptic feedback has been evident in sophisticated arcades for a long time and rumble packs are enhancing the gaming experience today in the same way that improvements in graphics might have done 5/10 years ago. Explicitly modeled force feedback that corresponds to the gaming interaction will provide a new sense of immersion otherwise not possible.

Haptic interfaces can reduce 'information clutter'. Unlike speakers and video monitors, haptic displays don't generally clutter a user's environment with unnecessary information. A good example of this property is a mobile phone set to vibrate rather than a ring tone. This haptic display provides only the right message ('you have a call'), to the right person (the owner), at the right time. This specificity is likely to become more important as embedded processors make more 'real world' objects intelligent and active. The same considerations suggest that haptic displays may reduce information clutter in virtual environments of increasing complexity. Haptic cues have also been developed to augment graphical user interfaces to windows operating systems, both Microsoft Windows (Immersion, 2000) and Linux/Unix (Miller and Zeleznik, 1999).

Haptic feedback can be used to add social context and/or augment the sense of shared environment enhancing the affect and communication available in an interface. Brave & Dahley (1997) introduced an approach for applying haptic feedback technology to interpersonal communication. Their 'inTouch' system provides a physical link between users separated by distance. This creates the illusion that the two people, separated by distance, are interacting with a shared physical object. The aim of the system is to enrich current real-time communication by opening a channel for expression through touch. This application differs from most others in that it does not attempt to create virtual objects with form, mass, and texture that can be felt through feedback from a haptically augmented device. Rather the idea is to create a physical link for expressing the movements or gestures of that person. Systems of this type are still at the prototype stage.

Commercial applications are becoming apparent as Internet purchasing becomes more secure and more popular. Natural physical materials are making way for the digital (e.g. online shopping). As such, opportunities for aesthetic indulgence such as stroking the fabric of an item of clothing before we buy it are becoming limited. Haptic feedback can bring aesthetic exploration into the digital world. Using haptic interaction, customers could potentially feel their products virtually before making the final decision to purchase.

In summary, there are a variety of new application areas in which haptic feedback is now being exploited. What is clear is that with more sophisticated devices and there are new haptic interaction issues emerging.

## **2.5 New haptic interaction issues**

The main contribution of this chapter is the characterisation of the human haptic system and the leading haptic technology in order that the capabilities of haptic devices can be better matched to the human haptic system. A set of haptic definitions was formally laid out to provide a common vocabulary with which to discuss haptic interaction research from this perspective. This chapter has highlighted the emergence of new haptic interaction issues involving richer and more complex haptic information than was previous available through teleoperation and simple force reflecting devices. Empirical research is now required if the potential of haptics in these new application areas is to be realised.

# Chapter 3: An Investigation of Haptic Effects in Graphical User Interfaces (GUIs)

---

## 3.1 Introduction

### 3.1.1 Motivation and background

Given that haptic devices are now commercially available, haptic interaction has become a potentially realistic solution to a variety of interaction design challenges as discussed in Chapter 2. This chapter reports an investigation of the use of haptic feedback as a way of reducing visual overload in the conventional desktop. The use of the PHANToM haptic device as a means of interacting with a conventional graphical user interface was investigated. Oakley *et al.* (2000)\* compared the effects of four different haptic augmentations on usability in a simple button-targeting task. The experiment reported in this thesis chapter involved a related but more ecologically-oriented searching and scrolling task. The studies were conducted at the same time in the University of Glasgow and followed a similar experimental design and analysis. Results from both Oakley's experiment and the experiment presented here indicated that the haptic effects employed in the GUI tasks did not improve users' performance in terms of task completion time. However, the number of errors made was significantly reduced. In addition, subjective workload measures showed that participants perceived many aspects of workload as significantly less with haptics. The results are described and the implications for the use of haptics in user interface design are discussed.

Desktop interfaces are becoming increasingly complex, and with this added complexity, problems are beginning to emerge. One such problem is information overload, where so much information is presented graphically that it becomes difficult to attend to all relevant parts (e.g. Brewster, 1997). Presenting information in other sensory modalities has the potential to lessen this problem. Attempts have been made to overcome information overload using non-speech sound during interactions such as button clicking and scrolling (Beaudouin-Lafon and Conversy, 1996) but there have been far fewer convincing empirical attempts to reduce overload by using haptic (or force feedback) technology.

---

\* The results of the experiment described in this chapter were first presented in a paper by Oakley, *McGee (the author)*, Brewster, and Gray (2000) at CHI, The Hague. Please see appendix G1 for the full paper. This chapter presents only the work from that paper which resulted directly from the work of the author carried out as part of this thesis.

Augmenting graphical user interfaces with haptic feedback is not entirely a new idea. Akamatsu and Sate (1994) developed a haptic mouse with the ability to produce what they termed ‘tactile feedback’, the ability to vibrate a user’s fingertip, and ‘force feedback’, a simple software controllable friction effect. Using this device they showed significantly decreased completion times in a targeting task offset by slightly increased error rates. In 1994 Engel *et al.* found improved speed and error rates in a generalised targeting task using a modified trackball with directional two degrees of freedom force feedback.

The devices used in these early studies have now been superseded. More advanced devices such as the Pantograph (Haptic Technologies Inc.), the FEELit mouse (Immersion Corp.), and the PHANToM (SensAble Technologies Inc.) have been developed. These devices have all been used to augment desktop interfaces. Ramstein *et al.* (1995) used the Pantograph to demonstrate performance increases in desktop interactions but provided little empirical evidence to support their claims. The FEELit mouse is a commercial product that offers users a haptically-enhanced desktop but there has been little evaluation of this device published (Rosenberg, 1997). Finally, the PHANToM has been used to create a haptically enhanced XWindows desktop (Miller and Zeleznik, 1998). Very little formal evaluation of this enhancement can be found in the literature.

The pace of technological advancement in this field is rapid, both in terms of the hardware produced and the software developed. Current projects to ‘haptify’ the desktop are not constrained to use the haptic effects described by Akamatsu and Engel. However, as technology has advanced there has been no corresponding progress in its evaluation. This disparity has led to a situation where there are no formal guidelines regarding what feedback is appropriate in different situations. This, along with evidence that shows arbitrary combinations of information presented to different senses is ineffective (Ramstein & Hayward, 1994; Ramstein *et al.*, 1996) leads to the conclusion that empirical evaluation of modern haptic augmentations of the desktop is urgently required if time and effort is not to be wasted. We might end up with haptically-enhanced interfaces that are in fact harder to use than standard ones and haptics may become just a gimmick, rather than the key improvement in interaction technology.

This chapter describes an experiment that empirically tests the use of haptics to augment targeting in the standard GUI. It is force feedback, and not tactile feedback that is evaluated in this work. In particular, it is an investigation similar to that of Oakley *et al.* (2000) but involving a more ecologically oriented task in which participants searched for and selected targets using haptic scrolling. Only some of the haptic effects tested in Oakley’s experiment were used to implement the haptic scroll bar. It was hypothesised that haptics will have a positive effect on performance. The experiment was not concerned with the influence of haptic distracters; it investigated haptic augmentation when there is guaranteed to be a clear path to target. The decision to adopt this approach reflects the preliminary nature of empirical research in this field.

### 3.1.2 Related work by Oakley *et. al.* (2000)

Oakley *et. al.* (2000) compared four different haptic effects (and a control of no haptic feedback) in a simple button targeting task. Each of the haptic effects was added to standard graphical buttons. This allowed the investigation of targeting (moving the cursor to the button) and mis-hitting errors (slipping-off the button when trying to press it). The experimental hypotheses were that differences would occur in task completion time, number of errors, and in the subjective data gathered. It was also predicted that the gravity well and recess would provide the largest reduction in errors, time and workload as they provided feedback that was highly appropriate to a simple targeting task.

Results from Oakley *et. al.* (2000) showed that significant effects were found when comparing the mean scores for each haptic effect for both slide over and slip off errors. *Post-hoc* analysis of the means showed that the most dramatic results were that participants in the gravity condition made significantly fewer errors of both sorts than in the control and that the converse was true of the texture condition – it caused significantly more errors than the control. Analysis of the temporal data was less conclusive. The total time taken to complete a trial was strongly biased by the number of errors made in each condition. Results revealed significant differences between the haptic effects. Gravity was significantly slower than recess. Oakley *et. al.* (2000) also found that the texture condition was significantly worse than the control across the whole board of subjective workload measures collected. The gravity condition consistently reduced workload and, in particular, achieved a significantly better score than the control in the performance level achieved category.

The hypotheses for the experiment presented here are described in more detail in the following sections but are based on similar hypotheses to the Oakley experiment discussed in this previous section.

## 3.2 An investigation of haptic effects in a scroll bar task

This experiment simulated a more ecologically realistic task than the simple button targeting task described by Oakley (2000). In this experimental task, reading was accompanied by scrolling through a document, selecting from the document, and returning to the scroll bar whilst still visually attending to the material being read. When users are required to scroll through a document it is the material in it that is of interest and not the scroll bar. Users want to concentrate on reading the material but often find themselves forced to move their visual attention to the scroll bar to ensure that the cursor is positioned appropriately to operate it. The time taken to make these frequent shifts in visual attention, and the frustration experienced by the need to do so, reduce the usability of the scroll bar. Problems associated with scrolling have been addressed previously (e.g. Brewster, 1997). Audio has been used to indicate the occurrence of errors to the user whereas haptics could also be used to prevent them altogether. Reducing the problem of slip off errors (while visually attending to a reading task) using force feedback technology has not yet been empirically evaluated.

### 3.2.1 Device and software

The device used in the experiment was the PHANToM 1.0 (see Figure 2.2 and Section 2.4.1.3 for details). It is a force feedback device (provides kinesthetic information as defined in Table 2.1) which, in the experiment, acted as a cursor control device in place of the traditional mouse. Optical sensors detect changes in the configuration of the PHANToM and the device uses mechanical actuators to apply forces back to the user calculated from this positional information and the stored algorithmic models of the haptic effects used. To operate the device users hold a stylus. The graphical interface to the experiment was generated using standard Microsoft Windows widgets and these performed in exactly the same way as standard widgets. The workspace was a box 160 mm wide x 160 mm high x 2 mm deep. The haptic effects were present only on the back wall of the workspace.

### 3.2.2 Haptic effects

The haptic effects used were based exactly on those models of haptic effect used in Oakley's experiments (texture, gravity, recess, and friction; see Oakley *et al.*, 2000). These built on and added to the effects used in previous studies in the force feedback literature (e.g. Akamatsu & Sate, 1994). The effects were all aimed at improving targeting and reducing problems of mis-hitting or slipping off interface widgets. Two of these effects were then used in the experiment presented here in this chapter. These were chosen as they were found to be the most suitable for the type of targeting task scrolling involves.

The effects used in the scroll bar experiment were:

**Recess:** The recess effect was a hole in the back of the workspace, with a depth of 2 mm and edges sloped at 45°. This effect also features strongly in previous literature (Miller and Zeleznick, 1998; Ramstein, 1995). A diagram of the geometry of a recess is presented in Figure 3.1. A recess could potentially provide useful feedback by the simple fact that to leave it, the wall at the edge must be climbed. This may make it harder to accidentally slip-off a button (a problem noted by Brewster *et al.*, 1995).



**Figure 3.1: Diagram of the geometry of haptic recess effect.**

*Gravity Well:* The gravity well was a ‘snap-to’ effect. When users moved over a button a constant force of 0.5N was applied that pushed them towards the button’s centre. This force tapered off around the very centre so that the user could rest in the centre. The gravity well promised the same benefits as the recess – a reduction in errors through the simple mechanism of preventing a user from accidentally slipping off a button.

The scroll bar in the thesis investigation was composed of these two effects and will be described in more detail later.

### 3.2.3 General measures used in the experiment

In order to get a full range of quantitative and qualitative results, time, error rates, and subjective workload measures were used. The performance measures were therefore (a) mean time per trial (secs.), (b) mean number of movements on/off scroll bar (including all required movements), and (c) subjective workload ratings. Time was measured from when the user activated the send button at the end of the previous trial until the send button was activated at the end of the current trial (see procedure section 3.2.6).

The subjective workload measurement was a modified version of the NASA Task Load Index (TLX) (Hart & Wickens, 1990). NASA reduced workload to six factors: mental demand, physical demand, time pressure, effort expended, performance level achieved, and frustration experienced. A seventh factor: fatigue was also added. One potential problem with force feedback devices is the physical strain placed on the user. By adding this factor it would be possible to find out if haptic effects caused any additional perceived fatigue. Participants filled-in workload charts after each condition in both experiments.


### 3.2.4 Participants

Twenty participants were used in this experiment (one was female and the remaining nineteen male). All were between the ages of seventeen and twenty-seven. Most participants were first-year computing science students from the University of Glasgow. All were regular and fluent computer users. All users were right-handed. Participants had nothing more than trivial previous exposure to the PHANToM and none had participated in Oakley's associated experiment.

### 3.2.5 Design

The experiment used a within-subjects repeated-measures design. Each participant experienced both a visual-only condition (visual) and a visual - haptic condition (haptic). The visual condition used a standard graphical scroll bar only. In the haptic condition, this same scroll bar was overlaid with haptic effects (see Figure 3.1 and 3.2).

Recess and gravity well were chosen as these were the most effective in Oakley *et al.*, 2000.




**Figure 3.2: The experimental scroll bar. Haptic effects from visual-haptic condition are indicated.**

The up and down arrow buttons used gravity wells. These acted as a haptic indication that the user was in the appropriate place to press the button successfully. The rest of the scrolling area used a recess effect that allowed the user to 'fall into' the slider area. A diagram of these effects is shown in both Figures 3.1 and 3.2. The haptic feedback allowed the user to reserve his/her visual attention for the primary task, as being over the widget was indicated through touch. The order of the presentation of the conditions was counterbalanced to evenly distribute the effects of practice and fatigue. Training was given to each participant in each condition prior to the experiment.

### 3.2.6 Procedure

Figure 3.3 shows the interface to the task. Participants had to read a four-digit numerical code from the instruction window. They then had to scroll vertically through a large file of codes (presented in the data window) to find the target code, highlight the code (either by double clicking on it or dragging across it), and press a button to send this code to the target window. The widgets operated as in standard desktop applications. The data window contained the same list of 2000 randomly generated but numerically ordered codes in each condition. Forty codes had to be entered in each condition. The list was formatted such that there were three columns of codes, simulating a standard document read from left to right and from top to

bottom. The highlight operation was included to force the user off the scroll bar. This ensured repeated targeting of the scroll bar. The experiment's duration was typically 40 minutes.



**Figure 3.3: The interface used in the Scroll Bar Experiment. Top left window is the instruction window, the bottom left is the target window, the large window to the right is the data window and in the centre is the send button.**

Expert GUI users are more than likely to encounter tasks where their visual attention is required in this manner. The large quantity of codes forces the need to scroll and monitor where in the file the user is currently. The numerical ordering of the codes however reduces the mental demand of processing each single code while still placing reasonably high demand on the visual search element of the task. Users often are required to concentrate on some central task and interact with graphical widgets in the periphery of their attention in this way (Brewster, 1997). Tasks involving scrolling in particular require the user to utilise their peripheral vision to monitor feedback from the scroll bar.

### 3.2.7 Hypotheses

It was hypothesised that when the scroll bar was haptically-enhanced, the participants would (a) take significantly less time to complete the task; (b) move on and off the scroll bar significantly less; and (c) perceive the workload during the task as significantly less.

## 3.7 Results

### 3.7.1 Timing results

Table 3.1 shows the timing and movement on/off scroll bar results. Paired T-tests established that haptic feedback did not significantly reduce the average trial time as predicted ( $T_{19} = 0.46$ ,  $p < 0.32$ ).

Mean Trial Time (secs.)		No. times on/off scroll bar	
Visual	Haptic	Visual	Haptic
11.7251	11.9668	107	97
SD=2.77	SD=2.84	SD=25	SD=22


**Table 3.1: Timing and movement results from Experiment.**

### 3.7.2 Error results

*Movement on/off scroll bar:* Paired T-tests showed that participants in the haptic condition moved on and off the scroll bar area significantly less than in the visual condition ( $T_{19} = 2.37$ ,  $p < 0.05$ ).

### 3.7.3 Workload results

Figure 3.4 shows the workload scores. Paired T-tests were carried out on the visual versus haptic conditions for each of the categories. Mental demand was not significantly less in the haptic condition as expected. Both the effort and frustration ratings were significantly reduced in the haptic condition (Effort:  $T_{19} = 2.80$ ,  $p < 0.01$ , Frustration:  $T_{19} = 2.04$ ,  $p < 0.05$ ). There was no significant difference in fatigue experienced. The hypothesis that the haptic condition would reduce workload is therefore confirmed in part.



**Figure 3.4: NASA TLX workload ratings for GUI experiment.**

### 3.8 General discussion

The timing results indicate that the haptic effects did not reduce the time taken for the task as hypothesised. There were also no real differences between the effects – only 42 ms between the best and worst effects (recess and gravity) in Oakley *et al.* (2000). The explicit separation of the error data from the timing data is no doubt a contributing factor to the lack of temporal variations across conditions. However, one potential reason for the lack of time reduction is that, in all of the effects used, participants had to exert more force to overcome the haptic effects. In the control condition they could just slide over the interface with no obstacles. In the haptic conditions they had to climb out of recesses, overcome gravity forces applied, etc. For participants to produce the forces required to do this could have taken them more time.

Further work is needed on the haptic effects themselves and the types of desktop tasks that would benefit most from them. It may have been that the haptic effects chosen were inappropriate either for reducing time or for the tasks chosen. Other previous work has claimed a significant reduction in performance times (Miller and Zeleznick, 1998; Rosenberg, 1997). The present work suggests that things are not so clear-cut and care must be taken when using haptics to try to reduce performance times.

The error results were more conclusive. Oakley *et al.* (2000) showed a significant reduction in the number of errors produced across the different haptic conditions (where gravity and recess caused the fewest errors and texture the most). The movement results presented in this chapter confirm this. Results showed a significant reduction in the number of times a participant moved on/off the scroll bar in the haptic condition. This showed that the haptic recess aided participants in remaining on target, demonstrating that haptics can provide a significant practical benefit for interaction. The haptic groove placed over the scroll bar allowed users to scroll up and down without slipping off. They could do this without looking at the bar as once the cursor was in the groove it would stay there. To move out of the recess they had to lift off the scroll bar and it was difficult to do this by mistake, as it required a physical effort. Gravity and recess therefore are the most effective for targeting tasks (which are important for using many standard GUI widgets, for example hitting a button, selecting a menu item or dragging the scrollbar thumb) in the sense that they made it very hard to slip off a target once on it.

Texture on the other hand only indicated that the cursor was over a target, and did not constrain users to the target, which was one of the reasons it was less effective in this case (Oakley *et al.*, 2000). Texture also had the problem that it could potentially perturb users' movements, making it hard for them to stay on target. This resulted from the kinesthetic force feedback device used here. We often use cutaneous stimulation to feel much of the richness of fine-grained texture in the real world (Lederman *et al.*, 1996). A kinesthetic device can only simulate gross textures, requiring larger forces, which then make it harder for users to move precisely. The PHANTOM is very effective at simulating gravity and recess effects as these require movement and so are kinesthetic tasks. There are no devices, as yet, which combine both tactile and kinesthetic force feedback. This is an important distinction arising again in the empirical work here and is explored in much greater detail in the remainder of the thesis.

The results show that interface designers must be aware of the facilities of the devices they are using in order to generate haptic effects that will improve usability. This might seem obvious, but this area is in its infancy and new devices are appearing all the time, each having different functionality to the last.

The subjective workload measures reported are also important. Papers concerning other haptically-enhanced desktops have not presented any such data. In developing multimodal interfaces (ones that use multiple sensory modalities) it is very important to consider what effects they have on users' workload. Users may perform tasks well and quickly and yet find them frustrating and requiring more effort to complete than they would expect. This dissociation between behavioral measures and subjective experience has been addressed in studies of workload. Hart and Wickens (1990) suggest that cognitive resources are required for a task and there is a finite amount of these. As a task becomes more difficult, the same level of performance can only be achieved by the investment of more resources. Just measuring time or error rates does not give the whole picture of the usability of a haptic device. Workload is particularly important in this area as we know little yet of the effects on cognitive/attentional resources of using such devices.

The experiment presented in this chapter showed the effect of haptics in a more realistic situation. In this case there was a significant reduction in effort and frustration – the fact that it was easy to stay on the scroll bar due to the recess effect made the task much less effortful (the reduction in the number of movements on/off the scroll bar confirms this). We had expected that this might also lead to reductions in other categories (e.g. mental demand) but these showed no significant reductions. This suggests that we need further studies of workload to learn more about the affect of haptics in desktop interactions.

One other area that was investigated was fatigue. Using a device that requires the user to apply force could cause fatigue. It is important to investigate this if force feedback devices are to be used in desktop situations (where people might use the interfaces for long periods of time). Results from Oakley *et al.* (2000) showed that gravity and recess effects did not cause any more fatigue than the control condition. On the other hand, texture caused significantly more fatigue than the control. This is likely to be for the reasons as discussed above – to simulate texture with a kinesthetic device required larger forces to be applied and these, in turn, required the users to exert larger forces to overcome them. Results from the current experiment again showed no increase in fatigue with the use of gravity well and recess effects. This research shows that appropriate haptic effects used correctly may have no impact on fatigue, but used incorrectly may significantly increase it. This is only a first step in investigating this problem and further work is needed to ensure that we can design haptic interfaces to avoid fatigue.

Oakley *et al.* (2000) showed that the different effects had markedly different levels of workload. Gravity well and recess came out best, indicating that they were effective at reducing error rates and decreasing workload. This suggests that they are very robust and can be successfully used in haptic interfaces of the type described here. Texture came out the worst in terms of workload, suggesting that, in general, it is hard to do effectively with force feedback. It is also becoming apparent from the GUI work that texture may not in fact be entirely suitable for the type of desktop targeting tasks described. This may be due to the nature of the texture effect not constraining the user to a desired target and in fact sometimes perturbing users from the area of interest. An important conclusion from this work therefore is that texture may in fact be more suitable in other contexts where constraint is not the main type of haptic interaction.

### **3.9 Conclusions**

This research has shown that haptics may have some benefits in graphical user interfaces. Reductions in the number of errors made and subjective workload experienced can be gained. It has also been shown that the haptic effects used must be matched to the capabilities of the device – trying to simulate effects not supported by the device in use can have serious negative effects on all aspects of usability. As technology progresses it is easy to focus on what benefits new equipment may afford whilst forgetting to measure the benefits actually produced. Recent work on haptically-enhanced desktops has been firmly orientated towards implementation and the experiments described here begin to redress the balance. These empirical findings provide a firm foundation for future researchers to build on and some basic principles for developers to use.

Chapter 3 presented an empirical evaluation of different haptic effects in a standard GUI scrolling task and discussed their effect on performance time, errors, and subjective workload. Results from the experiments showed that haptic effects do not necessarily improve performance in terms of time. This may be due to the extra force required to overcome certain haptic effects used. The number of errors made was reduced with the addition of haptics, however. The reduction in slip-off errors may be particularly useful in tasks such as scrolling where constraining the user to the object of interest may be beneficial. In addition, subjective workload ratings showed that the amount of effort and frustration experienced was significantly reduced by the addition of haptics in both experiments. These measures also showed that the haptic conditions did not cause any additional perceived fatigue to the users.

It is important to continue evaluating the effects of haptics on performance as different effects may be more or less suitable for different tasks. A particularly strong observation of this nature was that the gravity and recess effects seem particularly well suited to the GUI experiments presented here. On the other hand, the texture effect used appeared to be least useful in the GUI context and in fact often perturbed users from the small area of a widget such as a button. Due to the underlying nature of force feedback textures, this effect may not in fact be at all suitable for constraining users to desired objects as in the experiments discussed. It may on the other hand be extremely useful in contexts where constraint to an object is not the primary factor. How force feedback based texture effects can in fact be used appropriately and effectively is the matter of discussion for the remainder of the thesis.

# Chapter 4: Simulating textures

---

## 4.1 Real world textures

### 4.1.1 What are textures used for?

In human sensing and manipulation of everyday objects, the perception of surface *texture* is fundamental to accurate identification of objects (Katz, 1925/1989). When exploring an object in our environment we extract various types of information from it to help us identify what it is and what use it might be to us. As well as using the size and shape of an object for example, we often have to resort to the *texture* of the object to determine whether that object is glass, metal, or wood for example. The orientation and placement of textured elements on a surface can also be used to infer the orientation of an object in our environment or add to depth perception in a 3-dimensional scene. Texture is used primarily however for extracting the material properties such as roughness, stickiness, and so on from an object (see Figure 4.1).

### 4.1.2 What defines a texture?


The definition of texture in Webster's Online Dictionary is "the distinctive or identifying character or characteristics of something". This definition appears to be based on the work of David Katz, whose seminal work *The World of Touch* (1925/1989) set the agenda for much of the work on haptic texture perception that has appeared since. This description appears somewhat vague however and does not give any indication of what a textured object looks or feels like when compared to a non-textured object. One of Katz' central concepts was a distinction between two types of surface properties, *qualities* and *identifying characteristics*. *Qualities* are defined as properties on which any haptic surface can be rated such as the hardness or roughness of an object. *Identifying characteristics* on the other hand related to the overall feel of a surface such as 'leatheriness' of leather or the 'rubberiness' of rubber and so on.

Different definitions of texture appear in different contexts and therefore texture does not have a single definition. In vision, the texture of a surface is defined by how coherently it reflects light. A perfectly smooth surface such as a mirror reflects light uniformly, for example, whereas other surfaces have a rougher texture, made up of tiny reflecting surfaces set at different angles. Texture in the context of food science on the otherhand is a measure of the mechanical behaviour of foods measured by sensory or physical means. In soil geology, texture is defined as the relative proportion of sand, silt and clay of the dominant soil for each soil map polygon. Texture in botany is described as the appearance of a plant in terms of coarseness or fineness, roughness or smoothness, heaviness or lightness, denseness or thinness. What is clear is that there is no one simple underlying property of texture with which to form a simple single definition of texture. In fact, as the botany definition illustrates, defining texture may involve a set

of dimensions such as rough-smooth, hard-soft, slippery-sticky and so on. The multidimensional nature of texture, and the conflict in its underlying parameters, results in the simulation of texture being a complex design problem. The uncertainty as to the main parameters of texture leads to an uncertainty in the parameters to model and vary for virtual textures.

### 4.1.3 How are textures perceived?

It could be argued that vision alone is best at extracting information on global shape and structural cues about objects. Haptics alone on the other hand is extremely useful at extracting material cues and texture or compliance of an object. A common way to explore surface features on objects in our environment is to drag a finger or other probe (such as a pencil) across them. As the features on the surface become smaller and more closely spaced they become indistinguishable as individual elements, and are perceived as *texture* on the objects' surface. This definition serves as a useful perceptual definition in that it describes how it is that people come to encounter 'texture' when exploring a surface. Lederman's well-documented Exploratory Procedures (Figure 4.1) detail how objects and surfaces are explored haptically. Contour following is required for extracting shape for example while pressure and lateral motion are required for extracting texture.



**Figure 4.1: Exploratory Procedures taken from Lederman (1974).**

The goals of texture perception studies are normally to identify aspects of the surface geometry and physics that govern textural percepts. These studies have highlighted further the complex nature of texture.

#### 4.1.4 The complex nature of texture

Although some aspects of surface texture have been studied, an overall understanding of haptic texture perception is far from complete. Little is known about its dimensionality, its physical determinants, or its neural mechanisms. What is known is that something cannot simply be textured, or not textured. It is Katz' *identifying characteristics* that people often use to describe an object's texture subjectively. It is the *qualities* of a texture however that are scalable and are most often used in the scientific study of haptic texture perception. That is, we have to be able to distinguish between different textures and place them along appropriate dimensions just as say light stimuli can be placed along an intensity dimension. The exact number and nature of these underlying dimensions for texture however are not completely known or understood.

It is likely that texture is in fact a multidimensional concept. There is no doubt that roughness has the psychometric properties expected of a continuous dimension (Lederman and Taylor, 1972; Stevens and Harris, 1962). Hardness although less thoroughly studied also appears to be a scalar dimension of texture (Klatzky, Lederman, and Reed, 1987). Beyond these two dimensions, despite agreement that there may indeed be more, very little is understood. A multidimensional scaling analysis by Hollins *et al.* (1993) examined the possibility that other tactile dimensions of texture might exist. This study found that roughness-smoothness and hardness-softness were the most robust dimensions but that the other dimensions that might exist are slippery-sticky, flat-bumpy, and warm-cool. Until the exact number and nature of the dimensions underlying texture are understood, the problem of simulating haptic virtual textures is not entirely straightforward.

Chapter 2 highlighted the important distinction between cutaneous and kinesthetic perception. This is an extremely significant distinction in the perception of textures. When we experience textures in the real world we would most often drag our finger across the surface and judge the texture from the experience directly at our fingertip. It is unclear however as to the exact roles of vibration (cutaneous) and force (kinesthetic) cues in real texture perception and so this again causes uncertainty in the generation or simulation of virtual textures. Early studies by Katz (1925/1989) of textural sensations arising from moving a stylus across sandpapers led him to suggest that texture perception is based on vibratory cues transmitted through the stylus to the fingers. An implication of this work was that texture perception might be based on vibratory cues even when the fingers were in direct contact with the surface to be sensed. Such cues would result from repeated cutaneous stimulation of tactile mechanoreceptors as the stylus moves across the surface.


Lederman *et al.* (1974) on the other hand suggest that texture perception is mediated by force cues created by spatial geometry of the surface, not by vibratory cues generated by the repeated and regular stimulation of mechanoreceptors as the finger moves across a surface. It may also be possible to affect surface texture perception using vibratory cues, as suggested by Katz. In fact it is most likely that both kinds of cues could be involved, depending on the task to be executed.

#### 4.1.5 Roughness as the primary dimension of texture

Most psychophysical research on texture has focused on one underlying dimension, that of surface roughness. Surface roughness is one of texture's most prominent perceptual attributes and is used to guide texture discrimination. Roughness describes something that to the sight or touch has inequalities, as projections or ridges. The primary physical determinants of perceived roughness are themselves not entirely clear however. It is not clear for example whether a texture must be abrasive to be rough or whether 2-dimensional raised dot patterns can also be rough. Neither is it clear whether precisely controlled linear gratings are rough or whether textures must be irregular in nature to be perceived as rough. In actual fact, all these types of surfaces have been successfully judged as rough in the literature. The texture model that is used for the experimental work in a regular sine wave pattern and will be described in more detail in the sections that follow.


Much of the research on texture perception in the real world is of a psychophysical nature. The study of psychophysics of real textures started with work by Stevens in 1962. He applied his classic magnitude estimation technique to the perception of roughness. Participants experience a range of stimuli with different physical characteristics. For example, "roughness" stimuli using pieces of sandpaper with varying grit size which participants rub their fingers across. As previously mentioned, it is also common to use regular sine wave or square wave linear gratings as roughness stimuli. In this sense it is the variation in any of the physical geometry of the wave that in turn causes the variation in perceived roughness of that stimulus.

The exact contribution of the various geometry variables is not entirely clear either however. Using magnitude estimation techniques, Lederman *et al.* (1974, 1979) have shown that for engraved linear metal gratings with rectangular waveforms (Figure 4.2) the groove width between the ridges exerted the strongest effect on perceived roughness. That is they found roughness to be a monotonically increasing function of the groove width of the plates (Figure 4.3). These results were for perception of real physical gratings explored via the bare fingertip.



**Figure 4.2: Large-scale diagram of the physical profile of Lederman type metal gratings.**

Increases in ridge width however tended to decrease perceived roughness but with a more modest effect than groove width (Figure 4.3). They also comment however that it is highly unlikely that roughness is a function of a single physical dimension. Variables that did not affect the perception of roughness in these experiments included scanning velocity and spatial frequency.



**Figure 4.3: Diagram of approximate relationship between groove width and perceived roughness (Lederman *et al.*, 1974).**

The large body of work on real texture perception by Lederman *et. al.* Highlights the complex nature of measuring the roughness of textured surfaces. What it has confirmed however is that roughness continues to be one of the primary determinants of texture and therefore the most studied in the real texture perception literature. It seems reasonable therefore to begin with the dimension of roughness when exploring the generation and perception of virtual textures.

## **4.2 Virtual textures**

### **4.2.1 What are virtual textures used for?**

Haptic virtual textures might be an effective means of representing certain information about virtual objects explored through touch. Section 4.1 explained how the physical properties of textures might be very complex and difficult to reproduce for virtual textures. In a virtual world, haptic information can both increase the sense of realism of an object as well as convey informational content regarding what the object is, where it is, what it is for and so on. Textures also have important aesthetic properties that can play an important role in consumer-oriented design, marketing, and selling of products.

In virtual human and veterinary medicine, textures might be used in simulation. Distinctive textures of a tissue might indicate different stages of a condition or particular abnormalities in a patient, for example. Haptic training environments, which are becoming increasingly common aids for medical and veterinary training, would benefit from being able to convey different textures effectively. In many of the physical sciences soils, fossils, plants, and other materials all have surfaces of which their texture might carry significant information. The texture of a soil might indicate where the soil came from or details about the soil's composition. Details of virtual fossils might be explored haptically revealing small-scale information.

Using texture in the multimodal visualization of scientific data would allow regions to be *haptically textured* to indicate distinct areas of interest in the workspace. Data sets could be categorised by texture rather than colour in bar charts and pie charts for instance. Textures allow a larger amount of information to be available in haptic visualization techniques. Visually impaired people in particular would benefit from haptic visualization techniques. Visually impaired people are deprived of a great deal of information when exploring virtual scenes through only their tactile sense because it is difficult to perceive 3D aspects of 2D tactile pictures (Jansson, 1988). If textures could be conveyed effectively with such devices then the textures could be used to convey rich tactile pictures rather than the purely 2D aspects of tactile pictures currently available to visually impaired users.


The aesthetic properties of textures mean they are also bound to play an important role in any the design of products, and the marketing and selling of products in the multimedia world. In particular, arts and textile industries expanding their work to be available via the Internet would benefit from tools that allow them to convey the feel of their products to the user (Moody *et al.*, 2001). The simulation of online clothes catalogues would certainly benefit if the potential buyer could select their preferred material online.

#### 4.2.2 How are force feedback textures generated?

Texture generation and mapping has received considerable attention in graphics (Ekert *et al.*, 1994). With the increasing use of haptics in a variety of applications, it is becoming more and more important to explore the haptic representation of texture. There have been two main computational approaches to the generation of texture. The Stochastic approach (Fritz and Barner) assumes a spatial structure in a random field in which each texture element is calculated statistically according to its surrounding texture elements. Stochastic input parameters are derived by analysing actual force data. This techniques has been used by Green for the display of soils (Green, 1998).

In the Structured approach, the spatial structure of texture is emphasised. A texture is composed of a primitive pattern that is repeated throughout the texture. This thesis explores a simple structured approach to the generation of textures. It is based on a model where a function (regular sine wave) determines an approximation to a textured surface. This type of texture although not based exactly on reality is sufficient to convey the sensation of texture and roughness (e.g. Minsky, 1990).

Siira and Pai (1996) define haptic texture to be all those effects that are not explicitly accounted for by traditional rigid body contact normal (constraint) and lateral (friction) forces (Fig 4.4). This is a definition intended to guide the generation of haptic texture. It by no means uses a full understanding of real texture perception to display texture. Rather it is intended to approximate the feel of texture through interaction with a force feedback device. A person moves the probe of the force feedback device across a textured surface and the appropriate forces can be generated and felt providing surface texture based on the forces that are encountered. When the user exerts a force which is normal to the textured surface then they will either feel a stiff or rigid surface or a vibrotactile texture if texture is present. If a user exerts a force in a direction lateral to the textured surface they will feel either friction or surface texture if texture is present. It is *surface texture* as defined by Siira and Pai (1996) that is discussed in the experimental work of this thesis. That is, the user drags a probe, exerting some force, across a textured surface in a lateral direction to experience texture.



**Fig. 4.4: Definition of haptic texture according to Siira and Pai (1990).**

In the study of real texture perception, researchers are usually concerned with the surface geometry and physics that govern surface roughness. This is also the case in virtual texture perception. Currently, there are a number of researchers looking at ways to approximate the geometric and physical properties of real textures in virtual environments. Green and Salisbury (1997) for example developed a method to acquire data about the physical properties of textures with the help of the PHANToM by registering lateral forces and the z position of the endpoint of a stylus when it is moved over the surface. These are not exact copies of the actual physical surfaces but are intended to be sufficient for accurate perception of the textures.

It is crucial to consider the possibility that exact physical modeling of real textures (given how complex these are to begin with) may not result in virtual textures that generate the conceptual representation intended by their real counterpart. So, for example, playing back the forces recorded by dragging a devices probe across real corduroy may not in fact produce perception of corduroy when the virtual counterpart is conveyed to the user through the device. This makes the modeling of virtual textures more complex than other physical percepts in that the exact physical parameters required for the creation of realistic textures are not yet fully known. The physical experience of a virtual texture then is affected both by the capabilities of the device and the algorithms used to generate the textures.

### 4.2.3 How are force feedback textures perceived?

Much of the research on haptic display of textures focuses on the algorithms and hardware necessary to convey a convincing simulation of a haptic texture or percept. Much less work has been done on the actual perceptual responses to these simulated textures or percepts. It is these perceptual responses that have the ability to affect haptic interaction and so it is perceptual responses that this thesis deals with primarily. In designing haptic interfaces, designers should not assume that the virtual world will be perceived in exactly the same ways as the real world (Colwell, 1998), particularly given the current limitations of haptic devices which often use probes and one point of contact on a virtual surface (Lederman, 1974). Relationships between perception and simulated properties of virtual surfaces are even less well defined than their real counterparts. The existing work on these relationships will be summarised in this section.

Early perceptual and physiological studies argued purely for a spatial code in texture perception. Direct exploration with the fingerpad in particular involves spatially distributed cues providing the main percept for subjects' judgments of surface properties (Lederman and Taylor, 1972). In remote environments, however, spatially distributed cues on the fingerpad do not correspond to the surface geometry at the distal point of the interface, rather they correspond to the geometry of the probe itself. The user is forced to adopt vibrational cues transmitted via the probe (Kontarinis and Howe, 1993). Wellman and Howe (1995) have experimented with a system that combines force and vibrotactile feedback in conveying information about surfaces. They argued that relatively high frequency vibrations play an important role in conveying information about surface stiffness and texture in a way that simple force reflection does not. Katz (1989) described how the proximal stimulus that permits textural distinctions is the pattern of vibrations transmitted to the fingers holding the stylus, which constitutes a purely temporal cue. This temporal nature to texture perception may be true given the relative success of force feedback conveyed texture.

Three-dimensional (3D) haptic interfaces that provide point-source force feedback, such as the PHANTOM (used in this work), are also able to simulate surface texture with a surprising degree of fidelity. Even 2D joysticks can convey convincing texture. Minsky *et al.* (1990) used control algorithms with a joystick and spring forces based on a local gradient to simulate fine grained surfaces. In 1995, Minsky investigated both display algorithms and perceptual characterizations of these types of textures. The lateral force gradient algorithm was used to create textured surfaces, and perceptual judgments of roughness and other textural characteristics were elicited from subjects as the surface force characteristics and spatial geometry were varied. It was found that in judging the roughness of textures, variations in force accounted for over 96% of the variability in roughness estimations, whereas spatial frequency played little or no role. Virtual texture perception may indeed have both a spatial and temporal nature, and given the nature of both tactile and force feedback devices it seems reasonable that both devices are capable of generating texture percepts to some extent. This thesis focuses on the use of a probe (on the force feedback device) to explore and perceive force feedback generated textures.

Weisenberger & Krier (1997) conducted investigations designed to assess the roles of force alone, and vibration alone, in transmitting textural cues. One haptic interface provided *force* feedback to the fingertip via the PHANToM and the other provided *vibratory* feedback to the fingertip. They aimed to determine whether subjects could reliably identify surfaces whose edges and interior detail varied in frequency, intensity, and spatial density, and to map the range over which such judgments could be performed. With the vibratory display, subjects could differentiate surfaces that differed in vibratory frequency, intensity, and spatial density. For the force display, they obtained minimum values for differences in spatial frequency and amplitude that permitted discrimination between pairs of surfaces.

Weisenberger and Krier were attempting to determine the parameters that define the useable workspace for texture simulation using force feedback. Conventional haptic interfaces can not accurately reproduce trajectories to the small scales typical of haptic textures. Weisenberger and Krier (2000) did find that the threshold for textures improves as the amplitude of the grating is increased. It is also important to point out however that larger amplitudes may in turn cause you to fall out of or off of textures (see results from Chapter 3; Weisenberger and Krier, 2000). In addition, higher amplitudes, with most currently employed models, have been found to lead to greater instability of the PHANToM. The work presented in Chapter 6 therefore examines ways to improve the range of textures available through force feedback while maintaining stability with force feedback generated textures. Overall, Weisenberger and Krier confirmed that both vibrotactile and force feedback can be used to generate distinguishable surface features and that for manipulation tasks that require surface differentiation, it may not be necessary to create completely realistic simulations. Simple, regular models of texture such as sine waves are sufficient to elicit textural percepts.

West and Cutkosky (1997) demonstrated the use of haptic displays to produce geometric surface features at the sub-millimeter level. A stylus was connected to a haptic interface and configured such that users could explore real and virtual surfaces using the same apparatus. Surfaces consisted of sinusoidal profiles and subjects were asked to explore the surfaces and count the number of waves detected. Tests were also conducted using the subjects' fingertips instead of the stylus for physical surface exploration. Subjects' perceptions of sinusoidal features on virtual and physical walls were found to be qualitatively similar. Despite the limited stiffness and bandwidth of the haptic interface the ability of humans to accurately count virtual and physical waves was also similar with both real and virtual interaction. They also investigated the scales at which individual peak or valley features give way to general sensations of roughness or smoothness when dragging a stylus over the manufactured surfaces and virtual surfaces.

West and Cutkosky found that subjects performed better when using fingertips than when using a stylus for all but the highest frequencies tested. With higher frequencies the stylus can fall into troughs that are too small for the fingertip. Stylus exploration however was not good at trials with lowest amplitudes. They concluded that to obtain better performance with features of low amplitude and high spatial frequency the bandwidth of the device must be improved.

Jansson *et al.* (1998) conducted a study in which sighted, blindfolded students explored and made magnitude estimations of the roughness of both real and virtual sandpapers with four different degrees of physical coarseness. They used the Impulse Engine to present the sandpapers. Their judgments were made on haptics alone, the auditory information being masked. The virtual sandpapers were simulated by modeling normal and tangential forces recorded during exploration of real sandpapers. It was found that the perceived roughness was greater for virtual sandpapers. They found a highly significant relationship between perceived roughness and spatial period with wider grooves leading to increasing roughness.

Colwell *et al.* (1998) used a magnitude estimation technique to assess the roughness of ten grooved textures. Regression analysis was used to determine how much of the variation in the sensation of the textures could be accounted for by the variations in the groove width. Overall, they found a highly significant relationship between the perception of virtual texture and its simulated physical characteristics. 3 out of 9 blind participants tested in this study perceived narrower grooves as rougher than wider grooves. For the other 6, they perceived the wider grooves to be rougher than the narrower grooves. Only 5 out of 13 sighted participants showed a significant relationship between perception of virtual texture and its simulated physical characteristics. Colwell *et al.* concluded that it could not be assumed that physical variations in roughness of virtual textures could necessarily be easily detected or discriminated from one another. In fact, virtual textures may not be perceived in the same way as their real counterparts.

## **4.3 Improving force feedback generated textures**

### **4.3.1 Problems encountered in our exploration of force feedback texture**

It has been illustrated that the physical properties of textures are very complex and that this may partly be affecting the successful simulation of virtual textures. One of the major constraints for successful haptic texture is the generation of a sufficiently realistic texture given hard constraints on computational costs. It has been shown however that force feedback devices are capable of producing successful textural percepts using simple linear force models. It is also clear that texture can be used to convey information about an object rather than simply increase its realism. The exact nature of the underlying physical geometry of the most successful virtual percepts is still not known however. It is important to continue research on the physical geometry of virtual textures and particularly important that the perceptual responses to these stimuli be evaluated.

A force feedback interface relies on the sensing, calculating, and actuating of forces on the users' finger(s), hand, or body. The forces are large-scale compared to the small-scale forces produced by tactile devices that stimulate the skin to convey texture. Methods to convey texture purely with force feedback devices are a matter of ongoing research and are achieving varying success (e.g. Colwell *et al.*, 1998; Minsky, 1990; Lederman; Jansson *et al.*, 1998). It has been found however that the range and resolution of the available

textures is under question. In Chapter 3 for example it was found that the geometrically gross textures used perturbed users' movements, making it hard for them to stay on a desktop target that was textured.

It has been suggested that for these to improve, the bandwidth of the device must increase. The remainder of this thesis explores a means to improve force feedback generated textures without relying on more sophisticated software or hardware. A more immediate and cost effective solution might be to use other modalities to improve the perception and cognition of force feedback textures.

### 4.3.2 Brief introduction to multimodal augmentation

Taylor *et al.* (1973) state that texture perception potentially involves the coordinated action of a number of sensory systems - cutaneous, kinesthetic, visual, and auditory. It is likely that we use each of these sources of information for texture perception. An impression of roughness in particular is easily obtained by the visual, haptic, and auditory senses and so the visual sense (or any other for that matter) does not take a precedence over the others. That is, there is no strict 'hierarchy' to the senses for the case of texture or roughness perception. One is not necessarily more 'significant' than the others (Lederman and Abott, 1981).

Katz (1989) stated that people are highly skilled in using such touch-produced sounds to identify the material of various objects and Lederman (1979) reported that people are capable of making texture discriminations when cutaneous, kinesthetic, and visual cues are eliminated. In fact, Lederman and Taylor (1972) showed that subjects could not easily differentiate the roughness of regularly grooved surfaces in which the uncut portion between the grooves (the land) was the only aspect of the surface to be varied. They suggested that sounds produced when touching a surface could perhaps serve as an additional source of information in such cases.

Simple auditory cues could therefore potentially be used to increase the magnitude of a sensation such as roughness (for example) such that the mechanical effort requirements of the haptic device are reduced or the maximum stiffness or roughness that can be displayed are in fact increased (Ruspini and Khatib, 1998). The use of multiple modalities, in this case haptic and auditory, is therefore one way to increase the available resolution during haptic texture perception without increasing the computational costs significantly.

As an example, Heller (1982) investigated the interaction of the visual and haptic senses in the perception of surfaces. It was found that vision and touch have similar levels of accuracy in the perception of roughness. It was also found that bimodal perception was superior for texture judgments. It was suggested that vision might aid the perception of roughness by allowing an active explorer to guide their hand movements in a more efficient manner. The addition of auditory cues to force based textures might also have positive effects on interaction with those textures. The exact nature of the effects of this multimodal augmentation to current force feedback generated textures is discussed in the next chapter.

## 4.4 Summary

This chapter examined the nature of both real textures and virtual textures and the possibility of improving the range and resolution of force feedback based textures available via the current haptic technology. The nature of real textures was reviewed in order to understand the ways in which force feedback textures might be generated and to explore the possible limitations of force feedback devices in displaying virtual textures. It was shown that textures have been found to be perceptible through predominantly the kinesthetic sense (Katz, Lederman) and force based textures have been displayed convincingly to the kinesthetic sense (e.g. Minsky). Given the larger scale of force based textures compared to cutaneous textures, it has also been shown that users may be perturbed off small textured areas (Chapter 3). It appears then that although force feedback devices are perfectly capable of conveying textures, the exact range and resolution of textures available is an important research question that this thesis addresses.

One way to improve the quantity and quality of force feedback generated textures has now been briefly introduced. That is, our ability to process and integrate information from multiple sensory modalities simultaneously can augment the potential information available from any sensory modality individually. Chapter 5 will now discuss such a multimodal augmentation approach in the context of improving force feedback textures.

# Chapter 5: Multimodal augmentation

---

## 5.1 Introduction

This chapter introduces the concept of multimodal augmentation of force feedback generated textures. Beginning with multisensory information processing, the chapter proceeds to progressively narrow the focus of attention, identifying issues relevant to the design of the set of experiments reported in Chapter 6.

## 5.2 Processing multisensory information in the real world

Research in cognitive psychology has revealed that the brain seems to involuntarily recognise links between information presented via complementary modalities. Ryan (1940) for example, in his survey of work concerning intersensory relationships, notes:

"It is a commonly observed fact that most objects of our everyday lives are perceived by means of two or more sensory modalities working in cooperation."


Multiple sensory cues are in fact not merely *integrated* successfully, but also *transformed*, producing a reaction larger than that which would be expected from a sum of its parts (Stein and Meredith, 1993). Stein and Meredith suggest that this is indeed a fundamental basis for perception. This effect is the underlying idea behind the effect described throughout the thesis as multimodal augmentation.

### 5.2.1 Intersensory Interaction

Intersensory interaction is the modification of responses to stimulation in one modality by concurrent or juxtaposed stimulation in another. Conclusions from sensory interaction studies remain somewhat unclear. Senses appear to both interact and function independently. When interactions do appear it is often difficult to decide whether the positive results actually reflect an effect of stimulating one sense on processes of another, or instead an effect on the subjects' expectations of a multisensory stimulus and hence their judgements. A crucial factor is the objective relationship between the stimuli.

Meaningful related stimuli have a definite effect. The voice of a good ventriloquist, for example, sounds displaced in space, away from the ventriloquist's mouth, which does not move, towards the dummy's, which does. Of course, the success of the sensory interaction depends on our expectation of a voice originating from a moving mouth rather than a stationary one. Intersensory interaction has *not* occurred if the addition of a second sensory modality does not change the nature of the perception by the first modality. Often, of course,

perception *does* change when information is available to a second modality and, in such a case it is said that intersensory interaction has taken place.



**Figure 5.1: Illustration of the senses as an integrated system of information providers (Stein and Meredith, 1993).**

*Sensory substitution* research investigates the potential for one sensory stimulus to replace the information normally obtained through another sensory system. Sensory substitution has been used successfully in virtual environments where absent force feedback is replaced by analogous visual or aural displays for example. Information normally experienced through one of our sensory modalities can be experienced via another modality with approximately equal effect. This proves particularly useful when one modality is unavailable or restricted due to sensory disability or contexts, such as flight where one modality (vision) is already overloaded. Sensory substitution is not the type of intersensory interaction that will be discussed any further in this thesis.

The other primary method of intersensory interaction research is *Cross Modal interaction*. Cross-modal interactions are perceptual illusions in which users may use sensory cues in one modality to 'fill in' missing components of a perceptual experience. Psychophysics produced the method called "cross-modality matching" (Stevens, 1959). These studies show that people are capable of making cross-modal intensity matches quickly, easily, and reliably. Children as young as 5 years old for example can readily match brightness to loudness (Bond & Stevens, 1969). Synesthesia is the most extreme form of cross modal matching. Synesthetic people say that sensory events of one modality take on perceptual qualities or characteristics normally deemed appropriate to another. To many synesthetic people, loud sounds are bright; soft sounds are dim or dark (Marks, 1975). The extent of the possible additive effects of additional sensory

stimuli may vary. People may or may not automatically utilize the most effective sensory information available to them.

Combined sensory stimulation involves two or more modalities producing complementary and/or additive effects not achievable through the single sensory stimulus. This type of sensory interaction is the basis of the multimodal augmentation approach discussed throughout the thesis.

### 5.2.2 Evidence of intersensory interaction in the real world

The linking of auditory and visual stimuli has been demonstrated in many experimental studies. Spelke (1979), for example, presented two films simultaneously to infants, each showing an animal bouncing at a different rate, accompanied by a single soundtrack. Sounds were arbitrarily paired with visual objects, so that the infants could only respond to intermodal relationships through an awareness of the temporal synchronisation of stimuli. The infants were found to direct 64% of their looking time towards the film that was synchronised with the soundtrack.

Lipscomb and Kendall (1994) proposed a hypothetical model of film music perception. They suggested that if the relationship of the aural and visual accent structures is consonant, and if the associations of the musical style are judged appropriate to the visual stimulus, attention will be focused on a composite of the two stimuli, rather than on either stimulus in isolation. If this model is at all accurate then modalities will only fully complement each other if the coupling is judged to be 'appropriate'. Exactly how this subjective measure of 'appropriateness' should be measured is not clear from the literature, however.

Krumhansl and Schenk (1997) explored the interrelation of the structures of the musical and visual elements of audiovisual material. These stimuli were presented to subjects either in a composite form or singly. Subjects were asked to indicate the occurrence of structurally important events. A strong similarity was found across all conditions, with the composite condition having an additive relationship with the single-mode conditions. This suggests both that musical and visual elements are capable of conveying similar structural information, and that similarities in the structure of the two stimuli may assist the resolution of potentially ambiguous interpretations of the structure of one stimulus.

The studies reviewed in this section show that *temporal synchrony* of stimulus, *appropriateness* of stimulus combination and *ambiguity* of the judgment all might play a part in multisensory or multimodal processing. Stein and Meredith (1993) also suggested that characteristics such as spatial and temporal information are extracted from both aurally and visually presented stimuli and that "the *coincidence* or *covariance* of these strengthens the effect of the combined stimulus".

### 5.2.3 The effect of congruency

When deciding whether a light is dim or bright, people respond faster and more accurately when the light is accompanied by an informationally irrelevant sound that "matches" the brightness (Stevens & Marks, 1965). Responses to a dim light would then be faster and more accurate when it is accompanied by a soft or low pitched sound, and responses to a bright light would be faster and more accurate when it is accompanied by a loud or high-pitched sound. This phenomenon is known as the 'congruence effect'. Disentangling the contributions of low-level sensory components and higher-level semantic components to congruence interactions should contribute greatly to the improved design of multimodal interfaces.

When a multisensory percept is processed, the level at which the combined stimuli are integrated into a whole piece of information can vary. Sensory stimuli can be ignored or attended to; used alongside each other to reinforce a mental representation generated by one of the stimuli; combined in an additive way to produce a greater effect than either single stimuli; or processed against each other in competition to produce perceptual or cognitive conflict. Each of these conditions is a possibility when multisensory percepts are used in multimodal interaction.

Redundant information occurs when multiple pieces of information provide essentially the same content. The effects of redundant information can be to make the internal representation of the information stronger. This may result in more confidence in a perceptual judgment of a stimulus or a greater sense of realism regarding a percept. It is often difficult to establish whether the multiple pieces of redundant information are in fact being attended to and being processed or whether certain representations are being disregarded altogether. If cues in modality A are in fact redundant with cues in modality B then with or without the second modality information, absolute perceptual estimates of a stimulus should be essentially the same (Lederman, 1979).

Conflicting information might disrupt the information flow altogether such that the informational content is lost or the communication of the information fails. On the other hand, the information could become distorted rather than lost if multiple pieces of information are conflicting. The effects of this conflict can be very dangerous to interaction as the communication of information takes place as if everything were normal. The fact is that the information might have changed its meaning altogether by the time it is processed. If information from modality A conflicts with information from modality B (and assuming they are both attended to) then perceptual judgments regarding the stimulus might be different from those made when unimodal cues are present alone (Lederman, 1979).

## 5.3 The use of multisensory information in human computer interaction

Computer interfaces often provide incomplete or impoverished sensory cues when compared to the physical environment. Even multimodal interfaces engage only a fraction of human sensory bandwidth (Barfield *et al.*, 1995). Detailed study of both intrasensory illusions and intersensory interactions is important because in many cases the presence of an illusion or sensory enhancement might allow the system design to be simplified and, therefore, to increase its cost-effectiveness (Durlach & Mavor, 1994). Designers could use the

stimulation from displays for which they have greater control of fidelity (such as visual or aural displays) to augment the experience of stimulation for which the designer might have less control or fidelity (such as haptic displays). Exaggerating or distorting the representations of physical phenomena may create desirable cognitive effects.

Using more than one modality *appropriately* when interacting with a computer provides greater scope for computer interface design. A *multimodal interface* involves the use of more than one of our sensorimotor channels (for example vision, hearing, and touch) in the communication process between a computer and its users. Multimodal systems represent a developing area for computing that thrives on novel input and output technologies becoming available in mainstream computing. Such technologies allow for more complex and more powerful displays that have the potential to reach broader populations, encourage a wider range of applications, and improve the interaction experience altogether.

Despite the availability of technologies which allow multimodal interfaces to be implemented (at a realistic cost) there is a lack of applied knowledge on how our senses interact when using them. In computer environments the designer can control the presentation of multisensory information. It is possible then that there are techniques or methods that could compensate for limits in the performance of interface hardware. Research into these possibilities will allow the expansion of the range of perceptual experiences in a variety of multimodal interaction environments. Little work exists on the systematic study of how the output display mode could be better designed to coincide more closely with human information processing capabilities during multimodal interaction. This thesis adds to this area of research.

### 5.3.1 Evidence of intersensory interaction in human computer interaction

#### 5.3.1.1 Visual and haptic interaction

Srinivasan, Beauregard, and Brock (1996) studied the impact of visually presented spatial cues on the perception of stiffness in virtual environments. They used a series of psychophysical experiments designed to measure human performance in the discrimination of the stiffness of two virtual springs. The relationship between visual information on spring deformation and actual spring deformation were systematically varied. Results showed that graphically manipulated visual information could give rise to compelling haptic illusions about mechanical properties such as stiffness of objects in Virtual Environments.

More recently, Biocca *et al.* (2001) found that users who manipulated the visual analog of a physical force (a virtual spring) reported haptic sensations of 'physical resistance' even though the interface included no haptic displays. They suggest that presence may derive from the process of multimodal integration. They suggest that this perceptual phenomenon might be used to improve user experiences with multimodal interfaces, specifically by supporting limited sensory displays (such as haptic displays) with appropriate synesthetic stimulation to other sensory modalities (such as visual and auditory analogs of haptic forces).

Durfee *et al.* (1997) examined whether the addition of a visual display can trick users into believing that a soft wall is stiffer. They found that the effect of having equal or unequal haptic springs had a highly significant effect on the mean percentage of errors that followed the visual display. Subjects were more likely to make errors in determining the relationship between two springs when each spring had an equal haptic stiffness as compared to when there was a 20% difference in spring stiffness. When subjects did make errors, the errors were more likely to follow the visual display when the haptic springs had a 20% difference in stiffness as compared to when they were the same. Errors tend to increasingly follow visual cues as the amount of disparity between the haptic and the visual cues increases. Visual cues were more effective in influencing human perception of stiffness when a subject's level of interaction with the environment was limited. So in some cases, visual cues can be used in virtual environments to compensate for deficiencies in haptic displays.

### 5.3.1.2 Audio and haptic interaction

DiFranco, Beauregrad, and Srinivasan (1997) studied whether the auditory modality can similarly influence haptic perception. They conducted psychophysical experiments examining the influence of sound on the haptic perception of stiffness. Subjects used the PHANToM to feel the stiffness of various virtual surfaces. DiFranco *et al.* (1997) showed that in the absence of haptic cues, subjects ranked by auditory cues alone without any disbelief or perception of unnaturalness in the stimuli. The use of auditory cues then can potentially be used to overcome the limitations of representing rigidity with force feedback device.

### 5.3.2 Improving haptic user-computer interaction

One of the major barriers in the advancement of haptic computer interaction has been the inability of the current haptic display technologies to provide both the range and bandwidth of forces necessary to exploit the sensory and motor capabilities of the human haptic system. This results in the fidelity of many haptic virtual objects or percepts still being disappointingly low. Even as hardware becomes more sophisticated, knowledge of our human haptic system is still essential to detail the design specifications of haptic interfaces. A successful haptic interface should represent a good match between the human haptic system and the hardware used for sensing and displaying the haptic information. Since interface designers can control the presentation of multisensory information, it is possible that certain techniques could compensate for limits in the performance of haptic interface hardware. Researching these possibilities will allow the expansion of the range of haptic experiences in a variety of multimodal interaction environments.

There are specific interaction issues emerging from the increasing use of haptic interfaces that could potentially be solved using careful addition of multimodality. The representation of rigid walls for example has been a particular limitation of force feedback interfaces. Rosenberg (1994) describes the representation of a virtual rigid surface as the most frequently distorted sensory percept attempted by virtual force reflection systems. Virtual rigid surfaces are often described as "mushy", "sticky", or "bouncy", for example (Rosenberg, 1994). This is despite the fact that the ability to produce a convincing rigid wall is considered a

primary requirement of any general-purpose haptic interface (Jex, 1991). The display of a convincing virtual percept should not, however, be limited to the physical modalities.

It would be beneficial to know the extent to which peoples' perceptions are affected by such multimodal percepts. In doing so, ways in which to manipulate what the user perceives at the interface could be established. Such multimodal augmentation effects may allow designers to produce interfaces which have less than optimal unimodal capabilities. In particular, these effects could be used to overcome limitations of the device, the user, the environment, or the task. This knowledge could also be used to avoid coupling percepts that result in perceptual or cognitive conflict and which in turn might adversely affect the processing of that information during interaction. This can lead to significant savings in computational requirements. When two or more single percepts are combined however, the resulting multimodal percept may become a weaker, stronger, or altogether different percept. It is important therefore to compare the accuracy and/or reliability of multimodal, as compared to unimodal, judgments in multimodal augmentation research.

If only one sensory modality receives information about an event, perception of that event is generally unambiguous. In many situations however, information about an event is received by two (or more) modalities, and a question arises about the nature of the interaction between those modalities in the perception of such an event. Experimental situations may be designed to allow only one, or more than one, modality to receive information, and these situations can be used to study questions of intersensory interaction. It is useful to know whether the contributions of multiple redundant sources of information are additive (the more the better with diminishing returns) or substitutive (one simply confirms the other).

Much of the research into such multisensory integration in HCI is of a psychophysical level. Perceptual thresholds for example are examined in such studies. In 1994, Rosenberg discussed haptic-audio registration and reported that a delay as high as 100ms can exist between the presentation of haptic and audio sensations in the display of a virtual rigid surface before users notice any perceptual distortion. His work addressed the registration of audio and haptic cues in parallel and the temporal requirements that must be met for effective coupling of audio cues with haptic percepts. Temporal synchronisation is an important issue in haptic-audio multimodal interaction. Equally important, however, are the semantic considerations when combining percepts of different modalities. That is, how can the users conceptual understanding of a percept be altered through multimodal interaction? There has been little work to date that deals specifically with the semantics of associating sound and haptics in virtual environments.

## **5.4 Multimodal augmentation of force feedback textures**

Multimodal augmentation is defined as the combined sensory stimulation of more than one modality producing complementary and additive effects not achievable through the single sensory stimulus. It is a form of intersensory interaction defined here to apply to multimodal human computer interaction specifically and implies an *additive* effect of the sensory interaction which can influence the multimodal interaction either positively or negatively.

It has been demonstrated that textures can be generated and perceived via force feedback (Chapter 3 and 4) but that there are perceptual and interaction issues concerned with doing so successfully (Chapter 3 and 4). This chapter has demonstrated that multisensory stimuli can be combined to produce interaction effects that may augment the perceptual experience of the stimuli. In particular, audio cues have been used successfully to increase the magnitude of a haptic sensation such as rigidity such that the mechanical effort requirements of the haptic device are reduced or the maximum stiffness that can be displayed were in fact increased. The use of audio to augment a force feedback percept such as texture therefore is the subjects of the following chapter.

There has been little empirical work so far to examine if the haptic and auditory senses might interact in a similar way to that of the visual and auditory senses as detailed in the literature (section 5.2.2). In particular, there has been little systematic evaluation of how such effects (if they exist) can impact real human computer interaction issues such as the problem of improving force feedback textures being investigated in this thesis. The particular interaction being explored is that of multimodal (haptic - auditory) augmentation of these force feedback textures. The issue of the congruency of these multimodal stimuli (as discussed in section 5.2.3) is also a matter of concern in the multimodal texture experiments presented in Chapter 6.

# Chapter 6: Multimodal roughness experiments

---

## 6.1 Motivation and goals

It would be beneficial to know the extent to which we can affect peoples' perception by coupling haptic (force feedback) stimuli with stimuli in another modality (e.g. auditory) in a systematic way. In doing so we could establish ways in which to manipulate what the user will perceive at the interface, perhaps to overcome limitations of the device or the user for example. The addition of auditory information to force feedback virtual surfaces might increase the range and/or resolution of textures available to the designer for instance. This type of knowledge could be used to avoid coupling percepts that result in perceptual or cognitive conflict, which in turn might adversely affect the processing of that information. As has been highlighted, a design issue that might benefit from this type of solution is the problem of improving the quantity and/or quality of textures (or other haptic properties) available through current force feedback technology.

Information processed by multiple modalities which produces *conflicting* information in some way may cause the resulting multimodal percept to become distorted or completely lost. This should be avoided if multimodality is to be used in a positive way. Alternatively, people might process only one modality of information from the many available to them in a multimodal percept. This might be particularly true when both stimuli are intended to convey the same information for example. The modality, or modalities employed may depend on factors such as physical/perceptual ability, personal preference, or the nature of the task. A percept composed of multiple modalities might combine to give more than the sum of the individual parts. Two unimodal percepts, when combined, could produce an additive effect not possible with either unimodal percept alone. Such complementary pairings might act to increase the quality and/or quantity of information available through a haptic interface for example.

In the following experiments, sets of either unimodal haptic, unimodal auditory, or multimodal (haptic-auditory) textures were rated by participants to establish how rough each stimulus was perceived to be relative to each of the other stimuli. This allowed a set of both haptic and auditory textures to be perceptually classified along the dimension of roughness. This sheds light on the relationships that exist between the geometry of the textures and their relative perceived roughness. The haptic and auditory stimuli were then combined to produce multimodal haptic-audio roughness percepts for the final experiment. The combined textures were either congruent or incongruent in terms of the information each modality conveyed regarding the number of ridges/bumps on the virtual surface. Resulting multimodal percepts might provide *redundant*, *complementary*, or even *conflicting* haptic-audio information. The possible effects of these different levels of congruency are therefore discussed. It is hypothesised that multimodal textures, and in particular, the incongruent multimodal textures will provide complementary information. This complementarity could be used to increase the range and/or resolution of textures available through force feedback interaction.

## 6.2 General experimental overview

The same basic experimental paradigm was used throughout all the multimodal roughness experiments presented in this chapter. An overview of the design and strategy for the set of experiments is presented in this section. Specific details and results from each study are presented in turn after each experiment in this chapter. A general discussion for the set of experiments is presented in Chapter 7.

### 6.2.1 General design

The experimental work presented in this chapter consists of the following studies:

- 1(a):** Unimodal haptic roughness experiment
- 1(b):** Extension of unimodal haptic roughness experiment
- 2:** Unimodal auditory roughness experiment
- 3:** Multimodal haptic-audio roughness experiment

The general design for all four experiments was a fully counter-balanced within groups design. All participants were exposed to all conditions within the experiment. Participants only took part in one of the above experiments but within this experiment experienced all possible pairs of texture stimuli. The presentation of the textural stimuli and of the conditions within each experiment was counterbalanced in all experiments.

#### ***Main Independent Variable (all experiments):***

Spatial Frequency of the texture (number of cycles per 30mm).

Described in more detail in section 6.2.2.

#### ***Independent Variable (multimodal experiment):***

Congruency of the haptic-auditory texture pairing (congruent Vs Incongruent)

Described in more detail in section 6.2.3.

#### ***Dependent Variable (all experiments):***


Perceived roughness.

Operationalised in these experiments as the number of times each texture was judged as the roughest of a pair.

Described in more detail in section 6.2.4.

## 6.2.2 Texture stimuli (IV for all experiments)

Throughout this work the model of texture used is simple and deterministic as described in Chapter 4 (see section 4.2.2). Textures were generated as sinusoidal waves or gratings to approximate the feel of a regular but bumpy/ridged surface (see Fig 6.1). The resulting texture profiles depend therefore on the amplitude and frequency of the sinusoidal waves. The only physical parameter to be varied in the model of texture being used is the spatial frequency. The textures had fixed amplitude of 0.5mm and variable spatial frequency (cycles per 30mm). Spatial frequency is defined throughout as the number of cycles of wave per fixed 30mm virtual workspace (see section 6.2.6). The frequencies used varied from 5 to 45 cycles per 30mm virtual workspace. Higher frequencies were more tightly packed waves and lower frequencies were more loosely packed waves. The result of these textures was a bump felt at the peak of each wave.




**Figure 6.1: Diagrammatic view of the profile of the texture.**

This may be a simplified model of texture and it is by no means a definitive one (see chapter 4 for further discussion). A regular texture with only one manipulated parameter was chosen to allow the effects of multimodality and congruency to be the primary experimental factor. More complex models of texture would have increased the chances of these effects being obscured. Moreover, it has been shown that such models of texture are sufficient to convey the sensation of roughness (e.g. Minsky, 1990; Lederman, 1974). If multimodal augmentation is effective for simple texture implementations such as a regular sine wave then this increases the attractiveness of such an approach for haptic interface designers.

### 6.2.2.1 Haptic (Force Feedback) virtual textures

Haptic textures were generated as the sinusoidal waves described in the previous section. They were presented to subjects as located on a rectangular patch on the back wall of the PHANToM's experimental workspace (see section 6.2.6). As the user drags the stylus of the PHANToM in a lateral motion against the back wall of the PHANToM workspace, the probe that they are holding is perturbed by the resulting forces. The forces that are returned are generated from the physical geometry of the sine wave. The effect is the sensation of moving a probe across a bumpy surface. Figure 6.2 shows a simple representation of the resulting haptic forces.



**Figure 6.2: Indication of forces resulting from amplitude and frequency of texture.**

The forces are in fact modeled as a point contact in the z-direction and calculated and actuated using a simple linear force model (Hooke's Law) where:

$$\mathbf{F} = \mathbf{k} \times \mathbf{PenDist}$$

$\mathbf{F}$  is the force applied back to the user.


$\mathbf{k}$  is a constant chosen to make a surface feel both hard and stable.

$\mathbf{PenDist}$  is the distance penetrated by the stylus/probe.

### 6.2.2.2 Auditory textural stimuli

Auditory textures were generated from the same sinusoidal waves. The resulting profile therefore still depended on the amplitude (fixed) and frequency (varied) of the waves. The result of dragging the probe along these textures however was a single MIDI note (percussion organ) generated from and heard at the peak of each wave (see Fig. 6.3 for diagrammatic view).

Although the PHANToM device was still used to explore the virtual auditory space, no experimental forces were experienced through the device in the unimodal auditory conditions. Only the sound was heard. The effect is a note heard to indicate each bump along the virtual texture. This abstract notion of auditory textural cue was chosen as a simple starting point for this work. The auditory cues in this work semantically match the haptic cues in the sense that both are indicating movement over a bump or ridge on the surface. More realistic sounds could be digitised or in fact recorded but an important first step is exploring the multimodal effects of the simplest auditory cues. Participants experienced the auditory cues via headphones in order to mask the motor-generated sounds from the device as far as possible.



**Figure 6.3: Diagrammatic view of the profile of the multimodal texture for the congruent condition.**

### 6.2.2.3 Multimodal textural stimuli

Multimodal textures were generated from the same sinusoidal waves as described in the unimodal conditions. The difference in this condition is that there is the opportunity for one of the textures being explored to be multimodal rather than purely auditory or haptic. The result of dragging the PHANToM pen across these textures was a single MIDI note generated from and heard at the peak of every wave *as well* as the haptic forces. The bumps of each wave were therefore both audible and touchable in this condition (Fig. 6.3). In this way, multimodal textures could be compared to each of the force feedback textures.

### 6.2.3 Congruency and incongruency manipulation (IV for multimodal experiment)

In the multimodal roughness experiment there is an additional variable manipulation. The multimodal stimuli can be either congruent or incongruent. Congruency is a complex notion and is discussed in more detail in chapter 5. It is important to clarify what is meant by congruency in the context of this work. In the multimodal experiment congruency and incongruency are two separate states rather than a continuum. The congruent and incongruent conditions are determined by the information provided by each modality relating to the number of bumps/ridges encountered on a virtual surface. That is, if the number of contact sounds matches the number of contact forces then they are defined as congruent. Incongruency occurs when the number of contact sounds *does not* match the number of contact forces.

The haptic and auditory bumps in the congruent condition are displayed at the same place along the texture. They can be said to coincide. The haptic and auditory bumps in the incongruent condition are not displayed at the same place along the texture. The frequency of auditory cues is 120% of the frequency of the haptic cues. An incongruent multimodal texture with haptic frequency of 10 would have an auditory frequency of 12. An incongruent multimodal texture with haptic frequency of 30 would have an auditory frequency of 36 and so on. This type of incongruency could of course be explored further by making the incongruency bi-directional. That is, the auditory frequency could be *less than* the haptic frequency. In addition, various magnitudes of incongruency could be selected. The increase of 20% was chosen so as to displace the auditory cues from the haptic cues enough that they were obviously mismatched or incongruent without causing any cognitive disrupt that would make the textures appear unbelievable as rough surfaces.

A short pilot study was conducted to ensure that users perceived the incongruent auditory and haptic cues to be generated from the same stimulus. Displaying the MIDI note and the force as a result of the tip of the PHANToM reaching the designated point on the wave appeared to result in a convincing percept according to the participants' subjective responses in the pilot study.

As previously explained, neither haptic nor auditory textures are designed to model physically accurate or optimum representations of a rough surface. Rather, they are designed to give feedback approximate to that obtained when real textures are explored. In this way, the actual effects of experiencing such feedback multimodally as opposed to unimodally can be explored. It is the combined effect of experiencing the virtual textures via both the haptic and auditory senses simultaneously that is the subject of the final multimodal roughness experiment.

## 6.2.4 Measuring Perceived Roughness

As established in Chapter 4, the primary determinant of texture, and certainly the most researched, is roughness. For this reason, roughness was chosen as the dimension of texture to be explored. The methodology used however could potentially be used to explore either other modalities or other dimensions of texture. Likewise it could be modified to investigate multimodal augmentation of other haptic properties altogether such as stiffness or friction. The following experiments investigate the effects of the variation in the texture model on perceived roughness of the texture stimuli.

Although surface roughness is one of texture's most prominent perceptual attributes (Katz, 1989), the precise physical determinants of roughness are not exactly clear (Lederman, 1974). Because there is still debate over the actual parameters that determine roughness, users' *perception* of virtual roughness (regardless of the underlying physical model) is an increasingly important issue in virtual haptic interaction. It is exactly this perception that is being investigated here. It is hypothesised that a relationship will exist between frequency of texture and its relative perceived roughness. It is not assumed however that this relationship will necessarily be a simple monotonic one.

To obtain a measure of perceived roughness for the sets of texture stimuli participants made a fixed choice response regarding a series of pairs of the textured surfaces. The response method was a modified forced choice paradigm. Participants could decide that the roughest surface was on the left patch of texture presented or on the right patch or both patches of texture could be judged as the same roughness (see Fig. 6.4). This methodology was chosen as it was felt that magnitude estimation would not allow additional response categories to be incorporated. That is, this thesis was concerned not purely with the magnitude of the set of textures but also with the reliability with which participants decided the relative roughness of the set of textures. This design would also allow further measures such as a 'confidence in response' rating to be gathered in future work.

The resulting measure is a count of the number of times each separate surface is judged as rougher, less rough than, and the same roughness as each of the other surfaces. Using these measures the texture stimuli could then be placed along the roughness dimension to determine the likelihood that any frequency of texture would be judged as rougher than any of the others. In this way, a roughness profile for the set of textures could be plotted for each of the conditions and the effect of the modality of the judgment on the relative perceived roughness ratings examined.


## 6.2.5 The device and software

A PHANToM 1.0 force feedback device by SensAble Technologies was used to create the haptic virtual surfaces (see Fig. 2.2). For a full description of this device refer to Chapter 2. A standard PC with a Haptic Interface card was used to control the haptic interaction. The GHOST software development toolkit was used and the haptic texture effects were programmed in C++. The visual interface was constructed using Microsoft Foundation Class software. The auditory cues were generated and displayed using MIDI, a soundblaster card and standard headphones.

The PHANToM device was beside the computer where the mouse would normally sit. Headphones were used in all conditions to mask sounds from the PHANToM. The keyboard and mouse were removed from direct contact, as these were not needed for any exploration or response during the experiments. Subjects interact with the PHANToM by holding the pen-like stylus in all the experiments. By scraping this stylus/probe back and forth across the textured area the appropriate forces and/or sounds can be calculated from the positional information of the tip of the probe and the stored algorithmic models of the textured surface with which the user is interacting. The stylus switch on the probe of the PHANToM is used as a button to select each trial response a participant has to make.

## 6.2.6 The Interface and Virtual Workspace


A dialogue box (Figure 6.4) appeared on the screen throughout the experiment and provided a visual indication of where the virtual textures 'existed' as well as an interface for making their responses. Participants were instructed to drag the probe of the device over each of the indicated textured surfaces (the left and right square patches on the screen) and make a judgment on the roughness of the pair of textures. After each texture comparison, the participant clicked on the appropriate radio button in the centre of the dialogue box and clicked the 'next' button to start the next trial.



**Figure 6.4: The interface for roughness comparisons.**

The textures were presented in pairs as rectangular patches on the back wall of the workspace (see Fig. 6.5). Although the interface to the textures is 2 dimensional, the virtual workspace is in fact 3 dimensional. As previously mentioned, the participants use the PHANToM to explore the textures in all conditions, even when no forces are displayed. As such, participants were instructed to scrape the probe of the PHANToM back and forth across the stimulus surface to form an impression of how rough the surface seems to them. The participant was then asked to make a judgment regarding their comparison of the two surfaces. They made their response by clicking the appropriate button on the screen with the stylus switch on the probe of the PHANToM. When textures were haptic, users would feel the appropriate forces actuated back via the probe of the PHANToM. When textures were auditory, the virtual workspace that is explored coincides with that of the haptic workspace but the appropriate auditory cues are heard via the headphones.

Clicking the button labeled 'next' presented the next pair of surfaces. When the participant had completed all the trials they were given a message indicating that they had finished the experiment and a summary file for their responses was automatically stored for that participant.



**Figure 6.5: Diagrammatic view of the virtual workspace and interface.**

Participants compared each texture to itself and to each of the others twice (in a random order). Participants were allowed to explore each of the textures during that trial for as often as they liked and could switch between exploring the one of the left to exploring the one on the right as often as they liked to compare the two textures. They were instructed however that it was their initial response to the textures that mattered most and that there were not right or wrong answers for each of the trials. Participants made their response by clicking the switch on the probe of the PHANToM to select the response that reflected their roughness judgment for each trial.

A training session identical to the experiment but with less trials allowed the participants to become familiar with the device and the interface. Importantly, it also allowed them to adopt an exploration strategy for experiencing the textures comfortably and successfully.

## 6.3 Unimodal haptic roughness experiment (Exp. 1(a))

### 6.3.1 Experimental design

This was a within subjects (N= 12) design using the PHANToM device and the experimental software and set up as described in Section 6.2. The texture stimuli in this study were the unimodal haptic textures as described in Section 6.2. No auditory cues were presented in this experiment. Participants made a series of roughness judgments on a set of these force feedback generated textures. The textures had fixed amplitude of 0.5mm and one of 6 frequencies (cycles per fixed 30mm length of surface) - 10, 15, 20, 25, 30, or 35. The 12 participants compared each texture to itself and to each of the others twice (in a random order) resulting in 42 trials that lasted an average of 35 minutes. This incorporated comparisons with all possible frequencies and provided repetition of each comparison to test for response reliability. On each trial the user rated either one of the textures (left or right) as roughest or both the textures as the same roughness. The proportion of times each texture is rated as rougher than, as well as less rough than, and the same as, each of the other textures was then determined. This produced a dependent variable of perceived roughness for the range of textures compared.

### 6.3.2 Hypotheses

**Hypothesis (A):** The frequency of the texture will have a significant effect on the proportion of times that texture is rated as rougher than each of the others.

**Hypothesis (B):** The frequency of the texture plotted against perceived roughness scores will not necessarily produce a monotonic mapping from frequency of texture to perceived roughness. This will further reflect the complex nature of the concept of roughness.

### 6.3.3 Results and discussion (Exp.1 (a))

#### 6.3.3.1 Effects of frequency on perceived roughness

A one way ANOVA showed that the frequency of the texture had a significant effect on perceived roughness. For the main analyses, perceived roughness is operationalised as the number of times (or the likelihood) that a texture is judged as the roughest of a pair. As there were three response categories however, this could be expressed as either: the likelihood that a texture was judged as least rough of a pair, the most rough of a pair, or the likelihood that the texture pair are judged as the same roughness. Using this first measure gives an indication of the magnitude of roughness for each of the texture frequencies, however, and this allows a more direct comparison with previous work using magnitude estimations.

Effect of frequency of texture on:

Number of times texture was judged as	Most rough of the pair -	(F=9.73, p<0.01)
Number of times texture was judged as	Least rough of the pair -	(F=16.287, p<0.01)
Number of times the textures were judged as	The same roughness -	(F=7.632, p<0.01)


In summary then, the number of times each (frequency of) texture was judged as *roughest* was measured as the overall roughness score (Table 6.1). As noted above, the frequency of the texture was shown to have a significant effect on perceived roughness. That is, there was a significant effect of frequency on the number of times a texture was judged as the roughest of a pair (F=9.73, p<0.01). The number of times each (frequency of) haptic texture was judged as roughest tended to increase as the frequency of the texture increased (see Figure 6.6).

Frequency of texture	10	15	20	25	30	35
Perceived roughness score	24	18	38	35	61	69

**Table 6.1: Effect of frequency on perceived roughness.**

The graph in Fig. 6.6 shows the general trend for increased frequency (for this range) to lead to increased perceived roughness. Pairwise comparisons showed that the greater frequency had the significantly greater perceived roughness score with the exception of the pairs 10-15 and 20-25. There was no significant difference between the perceived roughness scores for frequencies 20 and 25. A frequency of 15 on the other hand had a perceived roughness score significantly less than that of frequency 10. It is likely however that the range used in the experiment is only a sample from a more complex function. In fact, the graph shown may not be part of a simple monotonically increasing function at all. Instead it may be part of a quadratic function of perceived roughness as suggested by Lederman *et al.* (1974). It may be likely that there is more than one maximum roughness generated from the set of frequencies.

In fact, it is likely that as the frequency of the texture goes below 10, the surface becomes a series of distinct bumps or waves rather than a unified texture. At the other end of the range, frequencies somewhere beyond 35 will become almost smooth again as the force profile becomes essentially flat. This would explain the apparent monotonic function in fact being a sampled portion of a quadratic (or more complex) curve for example. The range of frequencies tested would have to be extended at either side to determine whether this general increasing trend holds or whether a more complex (quadratic for example) function is revealed.



**Figure 6.6: Graph of effect of texture frequency on overall perceived roughness score.**

### 6.3.3.2 Reliability of roughness judgments

Participants made relative roughness judgements on every possible pair of texture frequencies twice. As well as evaluating the effect of frequency on the type of judgment made, the reliability of the judgment made for each texture pair could also be gathered. It was found that 61% of the time, subjects were in fact consistent in their responses for the judgments overall. That is, if they judged the pair as identical on first presentation they would judge them as identical on the second occasion and likewise for judging one or other of the textures as the roughest of the pair.

Only 7% of the overall judgments were in fact conflicting, or unreliable between trials (chance level 33%). That is, the subjects responded that a texture was roughest on one presentation and least rough on the other presentation. This is a small percentage of the trials and so it can be asserted from the double presentation of trials that the ratings were reliable measures of perceived roughness.

32% of the judgments were unreliable but in a manner more easily explained by indecision. That is, 32% of the trials resulted in participants' judgments changing from one being the roughest texture to both the textures being the same roughness. This could be due to the variation allowed, throughout the experiment, in factors such as force applied and speed of exploration. On one presentation of a pair, the exploration technique may lead to perceptually indiscriminable textures whereas on another presentation the exploration may produce distinct textures.

Overall, repetition of the trials has shown 61% consistency in roughness judgments between trials for any given texture pair. A very small percentage of the trials caused conflicting judgments. This at least validates the roughness measures used in the experiments to an extent.

### 6.3.3.3 Identical haptic stimuli

Textures with equal frequency were not reliably judged as the same roughness (accuracy: 50%-87.5%, mean: 64%). The rate at which physically identical textures were not in fact judged as the same perceived roughness was significantly higher than chance. This shows a strong effect of interaction on the perceptual-cognitive result of the stimuli. This could be explained by the freedom with which the participants were allowed to explore the stimuli. Many experiments have restricted the speed at which the textures are probed and/or the hand force applied during exploration. This allows these factors to be disentangled from the effects of the variation in the physical stimulus (spatial frequency). This also eliminates natural exploratory behaviour however something, which these experiments aimed to retain at the expense of a fully controlled psychophysical experiment.

Some of the variation in perceived roughness of identical textures arises then from the interaction between the probe and the physical model of texture such that different perceived profiles may arise from haptically identical profiles. Lower frequencies were more subject to these variations in perceptual differences. This is perhaps due to the interaction between probe size and texture-profile size; lower frequencies being more easily/significantly affected by differences in hand force and exploration speed.

### 6.3.3.4 Different haptic stimuli

Looking at the number of times different textures are rated incorrectly as the same perceived roughness showed the effect of frequency separation on the likelihood that the textures are judged as the same even when the textures are physically different (in terms of frequency).

When textures separated by a frequency of 5 are compared, 68% of the responses are that the texture pair are the same roughness. Interestingly, in the cases where physically identical textures are compared, 64% of the time the textures are correctly rated as the same perceived roughness. A frequency separation of 5 cycles therefore was not sufficient to significantly separate the perceived level of roughness for the haptic textures used. As frequency differences increased however, participants found it increasingly easy to decide whether the textures felt the same or different. Textures separated by frequency of 10 were rated the same roughness on 29% (below chance – 33%) of the judgments and those separated by 15 were rated as the same only 17% of the time.

Texture pairs deviating from the mean percentage judged as the same from their frequency difference group are the texture pair 25-30 and the texture pair 20-35. The texture pair 25 and 30 were judged as the same 92% of the time. This figure is higher than any other texture pair, including all those texture pairs where the textures had the identical frequencies. The texture pair 20-35 on the other hand were judged the same roughness only 8% of the time. This is significantly lower than the other texture pairs separated by the same frequency.

Although increased frequency differences made it increasingly easy to decide whether the textures felt the same or different it also made it increasingly difficult to decide which of the two was in fact the roughest. This might be caused by the range of stimuli generating two distinct notions of roughness. Frequencies of 10 and 15 were perceived as "bumpy" or "corrugated roughness" whereas frequencies from 20-35 were perceived as "sharp" or "sandpaper roughness".

### 6.3.3.5 Two distinct notions of haptic virtual roughness

Participant comments suggest that the lower frequency of 10 was considered very rough "like corrugated material". The higher frequencies of 30 and 35 however were also labeled as very rough but "like sandpaper". It is possible then that two frequencies from opposite ends of the scale can be perceived as equal in roughness magnitude but from different roughness scales. Experiment 1b (described next) extended the range of textures (5 to 45 cycles) to evaluate whether the increasing frequency leading to increasing perceived roughness relationship still held beyond the range used in Experiment 1a and whether bimodal peak roughness points emerged. This follow up study also evaluated our suggestion from Experiment 1(a) that comparing two textures from either end of the frequency range would increase the likelihood that they would be judged as different but also increase the likelihood that they would not be able to compare the textures on the same roughness scale.

## 6.4 Extension of unimodal haptic roughness experiment - Exp. 1(b)

### 6.4.1 Experimental design

This experiment involved a similar paradigm to that of experiment 1(a).

In addition however it involved:

- (1) Extending the range of frequencies used
- (2) Examining the possibility that there were two distinct notions of roughness emerging from the range of textural stimuli used.

(1) Extended range of frequencies

In Experiment 1(a) participants rated 6 textures (10, 15, 20, 25, 30, and 35). In this experiment however, subjects rated 3 additional textures. These were extended either side of the previous boundaries to include 5 cycles at the lower frequency (loosely packed wave) end and 40 and 45 cycles at the higher frequency (tightly packed wave) end. A pilot study showed that frequencies lower than 5 were difficult to perceive as textures at all. In fact, frequencies below this were subjectively judged as separate unidentifiable object. When pilot subjects were asked to evaluate frequencies beyond 45, their subjective feedback suggested that the texture became 'more like a vibration than a texture'. For these reasons, the range of frequencies was set as 5, 10, 15, 20, 25, 30, 35, 40, 45 for this experiment. This would allow us to plot an extended profile for the function of perceived roughness relative to frequency of texture. This would allow upper and lower force feedback texture boundaries to be included and determine their effect on the overall function for perceived roughness.

(2) Two distinct notions of roughness

Participants in Experiment 1(a) could rate the textures as the same, the one on the left as rougher, or the one on the right as rougher. Participants in this experiment were given the same options but with the additional response option of rating the textures as *not comparable on the same roughness scale*.

As in Experiment 1(a), this set of responses allowed us to evaluate:

- (a) Whether the participant perceived the two textures as the same or as different in terms of roughness.
- (b) The number of times each texture was rated as the roughest of the pair.

In addition however, the added response in Experiment 1(b) allowed evaluation of:

- (c) The textures participants felt were different but not comparable along the same roughness scale.

This was added as it was observed in experiment 1(a) that people often perceived a haptic difference but that they could not decide easily which one was in fact rougher. This additional response category might be used then when the two textures being compared are considered to be from two separate categories of roughness.

## 6.4.2 Hypotheses


**H1** - Increasing frequency of haptic texture (or number of bumps) will lead to an increase in the perceived roughness of the texture (replicating Experiment 1(a)).

**H2** - There will be a bimodal function of roughness with a frequency from either end of the scale being perceived as the roughest of the set. That is, textures compared from either end of the frequency range are more likely to be rated as 'not comparable on the same roughness scale' than textures compared within the high range or within the low range.

## 6.4.3 Results and Discussion (Exp. 1(b))

### 6.4.3.1 Effects of frequency on perceived roughness

An ANOVA showed that the likelihood that a texture was judged as roughest of any texture pair is significantly affected by the texture of the frequency being compared ( $F=2.05$ ,  $p<0.01$ ). That is, pairwise comparisons showed that as the frequency of the texture increased, the general trend was for the overall perceived roughness of that texture to increase (Figure 6.7). This confirms what was found in Experiment 1A and is now true for the extended range of frequencies used (5 to 45 cycles). It should be noted however that there is a significant dip at a frequency of 15 to a lower perceived roughness than a frequency of 10.




**Figure 6.7: Effect of frequency of texture on perceived roughness; Experiment 1 (b) - extended frequency range.**

### 6.4.3.2 Effects of frequency separation on perceived roughness

The effect of frequency separation of a texture pair was again analysed as in Experiment 1(a). This time, the fourth category of "can't judge the textures on the same roughness scale" was of particular importance. It was hypothesised that the larger the frequency separation of the texture pair, the more likely it was that this response would category would be used. Figure 6.8 shows the results from this analysis.

It can be seen from Figure 6.8 that the additional response category was used infrequently. Pairwise comparisons showed that the only significant rise in the number of times this response was used was from a frequency separation of 30 to the frequency separation of 35 and 40. When the frequency separations were at their largest, the subjects were more likely to use the response that they could not compare the frequencies along the same roughness scale. The hypothesis that textures from extreme ranges of the frequency scale will be more difficult to rate on the same roughness scale is confirmed in part. It is difficult to hold these results as anything more than suggestive at this stage however as it was also found that this response category was used unpredictably. It was used significantly less overall than the response categories "rougher than", "less rough than", and "same roughness as". It was also used unpredictably between the participants.



**Figure 6.8: Effect of frequency separation of texture pair on the perceived roughness rating.**

Some subjects opted not to use this response category at all. This response category was perhaps not a valid measure for capturing the dual notion of roughness present perhaps present in the range of stimuli after all. This response was dropped from the experimental paradigm hereafter and the original three responses were retained for the unimodal auditory and the multimodal experiments.

## 6.5 Unimodal auditory roughness experiment (Exp.2)

The results of the haptic studies suggest that larger frequency differences lead to more easily distinguishable textures but also to difficulties in using the dimension of roughness in comparing textures. It has also been shown that large textures have been found to throw users off of some textured areas. This thesis has suggested that the addition of audio information to such force feedback textures might ameliorate some of these problems. It has been proposed throughout that the combined (multisensory) presentation of haptic and audio textural information could increase the range and/or resolution of textures available to the designer without disturbing interaction through force feedback devices.

The main concern for the unimodal auditory experiment was to explore the effects of frequency of an arbitrary sound or note (based on the underlying texture model) on the perceived roughness of the auditory texture. This perceptual classification process is necessary before the haptic and auditory stimuli are combined. If something is not known of the perceptual ratings of the auditory stimuli (now that we know something about the haptic stimuli), then some artifact of the auditory stimulus chosen might obscure the effects of multimodality.

Katz (1925) stated that people are highly skilled in using sounds to identify the material of various objects. Whether they can also use sound to differentiate the roughness of surfaces has been a matter under investigation. Lederman (1979) evaluated the role of touch-produced sounds in judging surface roughness and found that subjects were capable of judging roughness on the basis of sound alone. Auditory judgments were found to be similar, but not identical to corresponding haptic touch judgments. When both sources of information were available, subjects tended to use the tactile cues. Roughness then, can be aurally estimated just as loudness can. Taylor *et al.* (1973) stated that texture perception potentially involves the coordinated action of a number of sensory systems (cutaneous, kinesthetic, visual, and auditory). The questions are, when and how are each of these sources used either separately or together for virtual texture perception.

Lederman and Taylor (1972) found that subjects could not easily differentiate the roughness of regularly grooved surfaces in which the uncut portion between the grooves (the land) was the only aspect of the surface to be varied. It is suggested that sounds produced when touching a surface in this instance could serve as an additional source of information. That is, when a decision is difficult or compromised then perhaps the additional (auditory) modality will be more fully utilised in the multimodal judgment of roughness.

If auditory cues complement haptic cues in a redundant manner, or if they are ignored altogether, then with or without the auditory information, estimates of roughness should be essentially the same. If sounds provide information that is incongruent in some way with the haptic cues (and if they are attended to) then estimates of surface roughness made when both cues are available might certainly be different from those made when haptic cues are present alone. If auditory cues complement haptic cues in an additive way then likewise the estimates of surface roughness made multimodally might be different from those made purely haptically.

### 6.5.1 Experimental design (Exp.2)

The experimental procedure and methodology were equivalent to that of the two haptic roughness experiments. The independent variable is still the frequency of the texture. This becomes however the number of *notes heard* per fixed texture patch rather than the number of *bumps felt*. The dependent measure is still the relative perceived roughness in the form of the number of times any texture is judged as the roughest of a pair. There were 10 participants (7 female and 3 male) none of which could be considered experts with the device. Four participants had never used the PHANToM before. Six participants had used the PHANToM previously for completely unrelated tasks. All the subjects were given a short verbal introduction to the PHANToM device and were exposed to a training condition that was identical in nature to the experimental phase.

The auditory textural stimuli were as described in section 6.2. A distinct MIDI note was played at an identical point on each wave peak. That is, the audio note sounded at each peak of the wave as though the noise was actually caused by contact with a physical bump. No haptic forces were present in this condition. The resulting experience is that of hearing contact with a rough surface via a probe rather than feeling the roughness of the surface. Although some effort was made to approach realism in order to make the task cognitively achievable it was never a main aim to simulate any particular material or object. Rather, the objective was to identify if the surface property of roughness could be simulated through auditory stimuli alone. A particular aim was to identify if subjective roughness comparisons could be made of these purely auditory stimuli which are based on a simple, sine wave based texture model

There were 9 different texture frequencies (5 cycles per 30mm to 45 cycles, in steps of 5 cycles). Each was paired with each of the others to produce 36 different texture pairs. Each comparison was made twice by each subject to result in 72 trials altogether. Each subject experienced all 72 trials in a random order.

### 6.5.2 Hypotheses

**H1** - The frequency of the texture will have an effect on the number of times a texture is judged as the roughest of a pair.

This would confirm that a relationship between frequency of texture and the perceived roughness of the texture also exists in the auditory modality for the stimuli used.

**H2** – Larger frequency separations will reduce the likelihood that the textures are judged as the same roughness.

This would show that the resolution of the perceived roughness is improved as the frequency difference between the pair of auditory textures being compared is increased.

**H3** – Larger frequency separations will increase the likelihood that the higher frequency texture is judged as the roughest.

This would confirm that increasing the frequency separation between auditory textures to be compared increases the likelihood that the relationship between frequency and perceived roughness holds.


**H4** – Frequency separation will have no significant effect on the number of times the lower frequency texture is judged as roughest (chance level).

This will show that the likelihood that a lower frequency texture is judged as the most rough (inverse to the hypothesised relationship) is not affected by the frequency difference between the texture pair.

### 6.5.3 Results and discussion (Exp. 2)

#### 6.5.3.1 Effects of frequency on perceived roughness

Figure 6.9 shows the results from the auditory roughness experiment. An increasing function similar to that found for the haptic condition was found. Increasing frequency of the auditory texture stimuli led to increasing likelihood that the texture would be judged as roughest of any pair. The first hypothesis (H1) is therefore confirmed. People rate the auditory textures as increasingly rough as their frequency increases.



**Figure 6.9: Effect of frequency of texture on mean number of times a texture is judged as roughest of a pair.**

### 6.5.3.2 Effects of frequency separation on perceived roughness

Table 6.2 shows the likelihood of each response for each frequency separation. The number of times each response is made (for each frequency difference) is divided by the total number of presentations (across subjects) of texture pairs with that frequency difference. There are more instances of textures being separated by 5 than by 35 for example. In addition, there are no instances where the textures are separated by 45 as when 45 is compared to its furthest away texture (5) the frequency separation is only 40.

There are 10 subjects. There are 2 trials per texture pair. This makes 20 responses per texture pair. There are more texture pairs separated by 5 cycles than 10 cycles and more of separation 10 than 15 and so on. For example there are 160 trials in total that include pairs separated by 5 cycles. There are only 140 trials that include pairs separated by 10 cycles, 120 for pairs separated by 15 cycles and so on.

Response	Frequency difference								
	5	10	15	20	25	30	35	40	45
% highest rougher	15	28	38	53	63	71	71	63	-
% lowest rougher	10	24	27	28	34	28	29	29	-
% same	75	48	35	19	3	1	0	8	-

**Table 6.2: Effect of frequency separation (between a texture pair) on perceived roughness response.**

If each response were equally likely all of the time (chance level) then each of the responses would be made 33% of the time. Significantly above or below this would mean that there is an effect of frequency on the subjective roughness judgment. The results confirm that a relationship exists between the frequency of the auditory texture and its perceived roughness.

### 6.5.3.3 Summary of auditory roughness results

- (1) The lower frequency texture of the pair is judged as roughest at or below chance level for all frequency differences.

When 5 cycles and even 10 cycles separate the texture pair the number of times the lower frequency is judged as roughest is significantly below chance. This is due to the greater likelihood that these pairs are judged as equal roughness. That is, when judging the perceived roughness of the auditory textures, the likelihood that the lowest frequency texture is judged as roughest is merely chance level, except when the textures are separated by only 5 or 10 cycles, when it is unlikely that the lower frequency will be judged as more rough at all.

- (2) The higher frequency texture of the pair is judged as the roughest of a pair increasingly reliably as frequency separation increases.

When auditory textures are only separated by 5 cycles the number of times the highest frequency is judged as the roughest is significantly below chance, again due to the significant number of times the 'same' response is used. This confirms again that unimodally the frequency separation of 5 is not sufficient to discriminate two textures in terms of roughness magnitude.

When the auditory textures are separated by 10 or 15 cycles, the number of times the highest frequency is judged as the roughest is increasing but still only slightly above chance level. Many of the responses at this separation are still in the 'same' category. When the frequency separation reaches 20, 50% of the responses reflect the higher frequency texture as being the roughest of the pair. This effect increases but tails off as towards the high end of the range. As the frequency difference goes beyond 30 the added effect of increasing frequency separation disappears. This highlights that simply increasing the frequency separation (if possible given the auditory perceptual thresholds) does not increase the discriminability without limit.

- (3) The number of times texture pairs are judged as the same roughness is a decreasing function of frequency separation between the textures.

Texture pairs are judged as the same roughness 75% of the time at frequency separation of 5. They are judged as the same almost 50% of the time even when the frequency separation increases to 10 cycles. This reaches chance level at frequency separations of 15 but decreases rapidly below chance and towards 0 as the frequency separation increases.

#### 6.5.3.4 Subjective results


A post experimental questionnaire was given to the participants after the unimodal auditory study to examine how people found the task of judging the roughness of purely auditory and considerably arbitrary stimulus (see Appendix E3).

When asked how easy subjects found it to tell that any two textures seemed different, 1 person responded that he found it quite hard and the remaining participants found it either average or quite easy. When asked how easy subjects found it to choose which surface was the roughest, the subjects again rated it as about average but with a slight tendency towards hard. This suggested that the task was perceived overall as fairly average to perform and that if anything, it was slightly harder to make the relative roughness estimate than it was purely to decide if they felt the same or different. This simply reflects the decision process of comparing the two percepts mentally and deciding which has the greater (roughness) magnitude.

When asked how many different textures were used to make up the set the responses were varied. Over half of the subjects thought that either 4 or 5 textures were used (around half the actual number). Only 2 people grossly underestimated the number of textures used (2/3) and only two overestimated the number used (10-12). Only 1 subject responded with the correct number of textures. The fact that people underestimated the number of textures used by around half is understandable given that those close together (separated by 5 cycles) were often rated as seeming exactly the same.

When subjects were asked to draw what the textures might have looked like, 25% constructed a matrix of regularly spaced dots to represent their view of the textures. One person drew a matrix of dots but spaced irregularly. 33% drew a series of vertical lines to represent the temporal occurrence of the note being played. One of these people clarified this by drawing a closely spaced series of lines and a more widely spaced series of lines to indicate the different scales of roughness. 25% drew horizontal lines but with peaks and troughs (waveform). All of these subjects gave an example of a closely spaced wave and a more loosely spaced wave to indicate how the textures varied. One subject suggested that height of the bumps (amplitude) varied both within a texture and between the textures. One final person drew the texture as a grid of intersecting horizontal and vertical lines. This most certainly highlights the variation in mental models of the textures. It could be that the notion of roughness is completely unattached to the physical nature of the stimulus and more directly to the excitation that occurs from interaction with the stimulus.

Subjects were also asked how many different notions of roughness they thought they experienced. This was intended to be distinct from the number of different stimuli and rather aimed to capture whether the subjects felt some of the textures belonged in separate categories of roughness as suggested in the unimodal haptic roughness experiments. Only 1 subject said that he had only 1 notion of roughness. 33% claimed to have experienced 3 notions of roughness and 1 subject claimed to have experienced as many as 5. It was felt that those with high numbers might have interpreted ‘notions of roughness’ to relate to the number of different stimulus. This is disproved by the fact that the numbers do not match the numbers in the question that directly asked for an estimate of the number of different stimuli. Importantly however, 50% said they had 2 notions of roughness which is what we might expect since previous observations from our roughness experiments have indicated that their may be (at least) 2 kinds of roughness - sandpaper roughness (high frequencies) and corrugated roughness (low frequencies).



**Figure 6.10: Diagram representing the different mental models of texture sketched by participants.**

## 6.6 Multimodal roughness experiment (Experiment. 3)

The unimodal experiments presented in the previous sections examined the perceived roughness of a set of force feedback generated textures (conveyed via a PHANTOM device) and a set of auditory textures based on these. These studies have confirmed that we can and do use both haptic and auditory information when making relative perceived roughness judgments of unimodal virtual textures. This work has also highlighted the possible perceptual limitations involved in reliably and confidently judging the relative roughness of a set of force feedback based textures given the gross nature of many force feedback generated textures and the small amount of workspace often available for textured surfaces or objects.

As Chapter 5 has highlighted, one possible way to increase the range and resolution of these force feedback textures is to combine them with the auditory texture stimuli and evaluate if and how they combine in a positive way. The multisensory percepts might combine for example to make the textures more easily discriminable from one another in terms of roughness. They may also combine to extend the perceptual boundaries possible in terms of perceived roughness. Or rather than alter the actual value of roughness perceived, the multisensory percepts might increase the reliability or confidence of the roughness judgments.

It would be beneficial to know the extent to which we can affect peoples' perception by coupling haptic (force feedback) stimuli with stimuli in another modality (e.g. auditory) in a systematic way as suggested. In doing so we could establish ways in which to manipulate what the user will perceive at the interface, perhaps to overcome limitations of the device for example. In this instance, the addition of auditory information to force feedback virtual surfaces might increase the range and/or resolution of textures available to the designer. Likewise, this information could be used to avoid coupling percepts that result in perceptual or cognitive conflict and which in turn might adversely affect the processing of that information.

Previous unimodal studies of the perceived roughness of a set of force feedback generated textures have shown some possible limitations in reliable roughness discrimination. It was found for example that participants did not always judge identical textures as the same roughness. Nor did they necessarily judge adjacent textures in a set as reliably different in terms of roughness. The suggested multimodal approach offers a cost-effective solution to overcoming the possible perceptual limitations of the currently available devices and texture models. Such a solution exploits the human ability to combine and integrate information from multiple sensory modalities into a fused, meaningful and whole percept. It is hypothesised that presenting combined haptic and audio percepts of roughness will increase the reliability with which people can make comparative roughness judgements of force feedback textures.

The work presented in this final experiment empirically investigates the effects of adding auditory textural cues to the existing haptic textures. It is hypothesised that the existence of an additional cue (in the auditory modality) will change the response patterns of users when asked to judge relative perceived roughness of the set of force feedback (now multimodal) textures.

## 6.6.1 Experimental Design

The experimental device, system, software, set-up, interface and general procedure were identical to the previous experiments (see section 6.2 for details). The specific details in this multimodal experiment are presented in the following section.

Computing Science students with no prior experience of the PHANToM participated in the experiment. No participants reported any auditory or haptic sensory abnormalities that might affect their performance. All participants experienced all texture comparisons in all conditions. The order in which the modality conditions were experienced and the texture comparisons presented within each condition were counterbalanced.

A within subjects (N=18) design was used with two independent variables:

**Frequency of texture:** Identical to spatial frequency manipulations in previously described unimodal experiments. There are 6 levels of Frequency in this experiment – 10, 15, 20, 25, 30, 35 cycles per 30mm. These boundaries were selected due to observations from the previous work in this chapter. Participants previously commented that textures of 5 cycles felt more like individual bumps than texture elements and that those of 40 and 45 were more of a smooth buzzing vibration when compared with the other more 'corrugated' or 'jagged' textures. The perceived roughness scores also confirmed this. The range of frequencies sampled came therefore from the monotonically increasing section of the function found in the unimodal haptic roughness experiments.

**Modality of judgment:** There are 3 conditions – Haptic (H), Multimodal Congruent (MMC), and Multimodal Incongruent (MMI). Each participant experienced all three conditions. The order of the conditions was counterbalanced across the subject.

All conditions involved comparing a unimodal haptic texture with another texture. The 'other' texture could be (1) also unimodal haptic, (2) multimodal congruent, or (3) multimodal incongruent. The definitions below explain further the notion of multimodality as well as the notion of congruency versus incongruency used in this experiment (also see section 6.2).

### Haptic Condition (H)

A haptic texture is compared against another haptic texture. No auditory stimuli are presented at all.

### Multimodal Congruent Condition (MMC)

Every haptic texture is compared against every multimodal texture. The haptic frequency and auditory frequency are numerically identical. That is the number of haptic bumps felt matches the number of auditory bumps heard.

## Multimodal incongruent (MMI)

Every haptic texture is compared against every multimodal texture. The auditory frequency is 120% of the haptic frequency. That is, the number of auditory bumps is 20% greater than the number of haptic bumps.

## Relative perceived roughness measurement

The dependent measure was the relative perceived roughness rating of each texture. This rating was gathered as a count of the number of times each of three possible responses was used: texture is the roughest of a pair; texture is least rough of a pair; and texture is the same roughness as the other texture in the pair as described in section 6.2). The effect of texture frequency on perceived roughness rating was evaluated as well as the effect of the modality of the judgments on those perceived roughness ratings.

The same modified forced choice paradigm was used to allow users to rate the perceived roughness of any two textures. Participants could rate the texture on the left as rougher, the texture on the right as rougher, or both as the same roughness as in the unimodal experiments. The *same* option was included to examine how reliably two physically identical stimuli are perceived as the same roughness in addition to how rough each different (frequency of texture) is rated compared with each of the others.

## 6.6.2 Hypotheses

**H1:** Increasing haptic frequency will lead to increased perceived roughness in all modality conditions.

**H2:** The modality of the judgment will have an effect on the number of times haptically identical textures are judged as the same.

**H3:** The modality of the judgement will affect the likelihood that different textures are successfully judged as different.

**H4:** The incongruency of the multimodal textures will have an effect on the perceived roughness judgements.

*Conflict:* If information processed by multiple modalities produces conflicting information in some way then the resulting multimodal percept may become distorted or completely lost in the process. Alternatively, the judgment of the multimodal percept might change in some unpredictable way. If participants are unable to easily make roughness judgments in the multimodal conditions when they could do so easily in both the unimodal conditions then this may be indicative of conflict. If the roughness profiles are distorted by the multimodal conditions then this may also be a sign of conflict occurring. Conflict may occur at the perceptual level and/or the cognitive level. That is, hearing a stimulus when feeling a stimulus might alter the perceptual experience altogether such that the haptic and auditory information have different perceptual qualities now that they are experienced together. On the other hand, the haptic and auditory perceptual experiences may be the same as the unimodal experience but the cognitive process of integrating them into a meaningful whole might be problematic resulting in conflict. Perhaps in this case, one of the modalities might be disregarded altogether when forced to make a judgment.

*Redundancy:* People might process only one modality of information from the many available to them in a multimodal percept. This might be particularly true of the congruent condition where both haptic and auditory stimuli are intended to convey the same information (the number of bumps on the texture). The modality employed may depend on physical/perceptual ability, personal preference, or the nature of the task at the time. If we are more used to using either our haptic or our auditory sense for roughness judgments then that might affect which modality we opt to disregard in this situation. The true effects of providing redundant information are in actual fact somewhat difficult to predict. Redundant information might increase the mental representation of the information. This may in turn lead to increased confidence or reliability of judgments without necessarily altering the content of the information. If the audio stimulus and haptic stimulus are congruent and redundant then with or without the auditory information, perceptual judgments of a virtual surface will be essentially the same.

*Complementarity:* A percept composed of multiple modalities might combine to in fact give more than the sum of the individual parts. That is, two unimodal percepts, when combined, produce some additive effect not possible with either unimodal percept alone. Such complementary pairings of haptic and audio stimuli might act to increase the quality and/or quantity of information available through a haptic-audio interface.

If the auditory stimulus and haptic stimulus are incongruent but complementary then multimodal (haptic-audio) judgments of roughness might move along the roughness dimension in the direction predicted by the direction of the incongruency. That is, when an audio and haptic stimulus are combined such that the audio stimulus is more rough than the haptic stimulus then the multimodal judgment of roughness is moved along the roughness dimension in the direction of increasing roughness. Likewise, when an audio stimulus and haptic stimulus are combined such that the audio stimulus is less rough than the haptic stimulus then the multimodal judgment of roughness is moved along the roughness dimension in the direction of decreasing roughness.

### 6.6.3 Results and discussion (Exp.3)


A 2 factor fully crossed factorial ANOVA was used to determine the effects of (1) *Frequency* of the haptic texture (*6 levels*) on perceived roughness and (2) *Modality* of the judgement (*3 levels*) on perceived roughness as well as (3) the interaction between the two factors *Frequency* and *Modality*.

#### 6.6.3.1 Effects of frequency on perceived roughness

Results from the ANOVA show that there was a significant effect of haptic frequency of the texture on the number of times a texture was judged as the roughest of any pair across all the texture comparisons ( $F_{5,85} = 16.22$ ;  $p < 0.01$ ). Pairwise comparisons showed that increasing frequency leads to increased perceived roughness for the range of textures compared (Fig 6.11).

**H1** is therefore confirmed - increasing haptic frequency leads to increased perceived roughness in all three modality conditions.

There was no significant effect of modality of judgement on the number of times any frequency of texture was judged as the roughest of any pair (see Figure 6.11). In addition, there was no significant interaction effect between the frequency factor and modality factor.



**Figure 6.11: Effect of frequency of texture on likelihood that texture will be judged as roughest of any texture pair.**

The *perceived roughness* rating being considered in this evaluation was the overall likelihood that a texture will be judged as the roughest of *any* texture pair collapsed across the entire set of comparisons for each texture. This work has found that the range used invoked two possibly separate notions of roughness: that of 'corrugated roughness' at the lower levels of 10 and 15 cycles and that of 'sandpaper roughness' at the higher frequencies of 30 and 35. It is possible therefore that the modality of the judgment might have significant effects when the range of frequencies is analysed in finer detail.


The increasing function found in all modality conditions replicates the function found for unimodal haptic textures in our previous study. It further confirms that people can successfully judge the relative roughness of a set of simple sinusoidal textures. It shows that, for the model of texture used and the range of frequencies sampled, varying the frequency (number of waves per texture patch) is sufficient to enable people to rate the relative roughness of the set of textures. This is true regardless of the modality of the comparison. This does not, however, alter the likelihood that roughness has more defining parameters than frequency alone (wave cycles per patch) as defined in this experiment.

### Identical haptic stimuli

The number of times a haptically identical pair of textures is judged as the same perceived roughness approaches 100% only in the haptic condition when the textures being compared both have a low frequency of 10. Even then, this likelihood is only 83%. Higher frequencies in the haptic condition have an even lower likelihood (mean = 50%) of being judged as perceptually the same in terms of perceived roughness (see figure 6.11).

In both multimodal conditions, the likelihood that haptically identical textures are perceived as the same roughness is significantly lower than in the haptic condition. This confirms the **H2** that the modality of the judgment will have an effect on the perceived roughness ratings (see figure 6.12).

The likelihood of haptically identical textures being judged as the same roughness decreases in the multimodal congruent condition and decreases further in the multimodal incongruent condition. This confirms both **H2** that multimodality will have an effect on the roughness judgments and **H4** that the incongruent within the multimodality will have an effect. There are frequencies at which the likelihood of the same response is equal regardless of condition. On the other hand, there are also frequencies at which the likelihood of same responses is dramatically different across conditions. The responses at individual frequencies may therefore need further exploration.



**Figure 6.12:** Effect of frequency of texture on likelihood that a texture will be judged the same as any other texture.

### 6.6.3.2 Effects of frequency separation on perceived roughness

#### Different haptic stimuli

In our previous studies, a frequency separation of 5 cycles was not sufficient for participants to be able to decide that the textures were different in terms of roughness any more than they would for haptically identical textures. It is possible that making the decision at this resolution in the multimodal conditions would be different.

Fig.6.13 shows the likelihood that two textures are judged the different at every possible frequency separation

(or resolution) including the cases where they are haptically identical. When the textures are haptically identical, the modality of the judgement has an effect on the likelihood that these textures are judged as different. **H3** is therefore confirmed.

In the haptic condition, a frequency separation of zero means that the same physical stimuli were presented. We might expect likelihood to approach zero for the probability of these identical textures being judged as perceptually different. In fact, the likelihood that identical stimuli are judged as different in the haptic condition is around chance level (33.3%) showing that people do not necessarily perceive physically identical force feedback stimuli as the same. This is perhaps not alarming given the freedom participants have to use as little or as much force in their exploration as well as their own exploration speed which in turn could vary within and between the trials. It does confirm that in practice, there is a strong chance of the haptic interaction affecting the perceptual and cognitive textural experience.




Figure 6.13: Effect of frequency separation on likelihood that two textures are judged as different perceived roughness - Experiment 3).

The likelihood of multimodal textures with identical haptic frequencies being judged as different is significantly higher than in the unimodal haptic condition. This shows that the additional auditory stimulus has an effect and that the incongruency in turn has an effect on the likelihood that haptically identical stimuli are judged as different. This would suggest that the auditory stimuli are in fact attended to and incorporated into the roughness judgement at this level. As the frequency separation increases beyond 5 cycles the likelihood of the textures being judged as different increases towards 100% very rapidly. A frequency separation of 15 cycles (or more) in the range of frequencies and workspace used is sufficient to elicit reliable difference judgments between the two textures.

More work is needed on both the absolute and relative perceived roughness of force feedback textures at the perceptual level as well as in the context of a texture dependant task. Many applications require virtual objects to have realistic surface properties and therefore simulating convincing texture is important. Many haptic tasks may, in fact, require these surfaces or objects to be discriminable in terms of their relative roughness or to make them classifiable according to their perceived roughness. It is necessary therefore for haptic research to continue to perceptually classify the range and resolution of roughness (and other texture dimensions) available through current technology.

# Chapter 7: Discussion of Experimental Results

---

## 7.1 The perceived roughness of force feedback textures

The frequency of the haptic textures had a significant effect on perceived roughness. The general trend of this function was found to be an increasing function of perceived roughness as the frequency of the texture increased. The quantitative results suggested that perceived roughness was rated along a scale dependant on the spatial frequency of the texture. It was thought that the function found was perhaps part of a more complex function of perceived texture and therefore an extended range of stimuli were also tested in an attempt to reveal this. Qualitative results suggested the existence of two perceptually distinctive notions of roughness yet the quantitative results don't confirm this.

Lederman *et al.* (1999) obtained a quadratic function for perceived roughness via a probe. The results from this thesis however show an approximately increasing function. It is suggested by Lederman *et al.* that an increasing function of perceived roughness might be achieved if the speed of exploration is high. It may be that the free exploration allowed during the experiments in this thesis naturally induced such a high speed of exploration and thus resulted in the monotonically increasing function.

### 7.1.1 Two distinct notions of haptic virtual roughness

Participant comments suggested that the lower frequency of 10 was considered very rough "like corrugated material". The higher frequencies of 30 and 35 however were also labeled as very rough but "like sandpaper". It is possible then that two frequencies from opposite ends of the scale can be perceived as equal in roughness magnitude but from different roughness scales. Experiment 1(b) extended the range of textures (5 to 45 cycles) and showed that the increasing frequency leading to increasing perceived roughness relationship still held beyond the range used in Experiment 1(a). Bimodal peak roughness points did not emerge as might have been expected. Comparing two textures from either end of the frequency range increased the likelihood that they would be judged as different but also increase the likelihood that they would not be able to compare the textures on the same roughness scale.

The larger the frequency separation of the texture pair, the more likely it was that the response of "can't judge the textures on the same roughness scale" would be used. When the frequency separations were at their largest, the subjects were more likely to use this response. The hypothesis that textures from extreme ranges of the frequency scale will be more difficult to rate on the same roughness scale is confirmed in part. It is difficult to hold these results as anything more than suggestive at this stage, however, as it was also found that

this response category was used unpredictably. It was used significantly less overall than the response categories "rougher than", "less rough than", and " same roughness as". It was also used unpredictably between the participants. Some subjects opted not to use this response category at all. It is possible that people under experimental conditions would feel they had to make a decision and not choose this fourth response as they felt that this was 'opting out'. This response category was perhaps not a valid measure for capturing the dual notion of roughness perhaps present in the range of stimuli. This response was dropped from the experimental paradigm thereafter and the original three responses were retained for the unimodal auditory and the multimodal experiments.

### 7.1.2 Reliability of the haptic roughness judgments

Subjects were consistent in their judgments on average 61% of the time. That is, if they judged a pair as identical (or rougher) on first presentation they would judge them as identical (or rougher) on the second occasion. It might be hypothesised that it would be preferred if the judgments were reliable nearer 100% of the time. As such, it could be said that presenting two stimuli at any time would result in the same roughness judgment being made each time. This is unlikely to occur however given the fact that each time the two textures are encountered either between trials or between subjects, there is a possibility that the exploration of both those textures is altered. Haptic interface designers should be aware that users presented with two haptic textures on any two occasions are likely to have the same relative perceived roughness of the pair on both occasions approximately 61% of the time.

It was only a small percentage of the time that the judgments were in fact conflicting, or unreliable between trials (well below the level of chance at 33%). That is, the subjects responded that a texture was roughest on one presentation and least rough on the other presentation. This is only a small percentage of the trials and so it can be asserted from the double presentation of trials that the ratings were fairly reliable measures of perceived roughness. Some of the conflict in the between trial judgments could be attributable to the dual notion of roughness. Subjects comparing textures from extreme ends of the stimuli range may perceive both textures as equally rough in magnitude although different perceptually. This might result in them choosing to assign one as roughest first time around and the other as roughest the second time around to reflect this. It is impossible to say for certain whether this strategy was used in the experiment conducted.

Almost a third of the judgments were unreliable but in a manner more easily explained by simple indecision. That is, 32% of the trials resulted in participants' judgments changing from one being the roughest texture to both the textures being the same roughness. This again could be due to the variation allowed, throughout the experiment, in factors such as force applied and speed of exploration. On one presentation of a pair, the exploration technique may lead to perceptually indiscriminable textures whereas on another presentation the exploration may produce distinct textures.

### 7.1.3 Judging identical unimodal haptic stimuli

Unimodal haptic textures with equal frequencies were not reliably judged as the same roughness (mean 64% correct). The rate at which physically identical textures were not in fact judged as the same perceived roughness was significantly higher than chance. This shows a strong effect of interaction on the perceptual-cognitive result of the stimuli. This could be explained by the freedom with which the participants were allowed to explore the stimuli. Many experiments have restricted the speed at which the textures are probed and/or the hand force applied during exploration. This allows these factors to be disentangled from the effects of the variation in the physical stimulus (spatial frequency). This also eliminates natural exploratory behaviour however, something, which these experiments aimed to retain at the expense of a fully controlled psychophysical experiment. Some of the variation in perceived roughness of identical textures arises from the interaction between the probe and the physical model of texture such that different perceived profiles may arise from haptically identical profiles. Lower frequencies were more subject to these variations in perceptual differences. This is perhaps due to the interaction between probe size and texture-profile size; lower frequencies being more easily/significantly affected by differences in hand force and exploration speed.

### 7.1.4 Judging different haptic stimuli

Looking at the number of times different textures are rated the same perceived roughness showed the effect of frequency separation on the likelihood that the textures are judged as the same even when the textures are physically different. When textures separated by a frequency of 5 are compared, 68% of the responses are that the texture pair is the same roughness. A frequency separation of 5 cycles therefore was not sufficient to significantly separate the perceived level of roughness for the haptic textures used. As frequency differences increase however, participants found it increasingly easy to decide whether the textures felt the same or different. Textures separated by frequency of 10 for example were rated the same roughness only 29% (below chance - 33%) of the judgments and those separated by 15 were rated as the same only 17% of the time. The texture pair 25 and 30 was judged as the same 92% of the time. This figure is higher than any other texture pair, including all those texture pairs where the textures had the identical frequencies. The texture pair 20-35 on the other hand was judged the same roughness only 8% of the time. This is significantly lower than the other texture pairs separated by the same frequency. It could be that for the frequencies tested there are different perceptual effects depending on the range of frequencies under examination.

## 7.2 The perceived roughness of auditory texture

The main purpose of the unimodal audio experiment was to explore the effects of frequency of an arbitrary sound or note on the perceived roughness of the auditory texture as in the haptic experiment. This perceptual classification process was necessary before the haptic and auditory stimuli were combined. If we do not know something of the perceptual ratings of the auditory stimuli (now that we know something about the haptic stimuli) then some artifact of the auditory stimulus chosen might have obscured the effects of multimodality.

## 7.2.1 The auditory function compared to the force feedback function

Lederman (1979) evaluated the role of touch-produced sounds in judging surface roughness and found that subjects were capable of judging roughness on the basis of sound alone. It is supposed then that roughness can be aurally estimated just as, for example, loudness. This thesis examined whether this would be true of simple MIDI note stimuli alone as a cue to roughness.

The higher frequency texture of a pair in the haptic experiments was judged as the roughest of a pair increasingly reliably as the frequency separation increased. This was also true in the auditory experiment. When auditory textures were only separated by 5 cycles the auditory judgments were likely to be rated as the same perceived roughness just as in the haptic modality. This confirms again that unimodally the frequency separation of 5 is not sufficient to discriminate two textures in terms of roughness magnitude. When the frequency separation reached 20, 50% of the responses reflect the higher frequency texture as being the roughest of the pair. This effect increased but tailed off as towards the high end of the range. As the frequency difference went beyond 30 the added effect of increasing frequency separation disappeared. This highlights that simply increasing the frequency separation (if possible given the auditory perceptual thresholds) does not increase the discriminability without limit.

Overall, a similar function was found for the auditory modality. On individual examination of each user's perceived roughness profile, however, it was observed that a subset of the participants had opposite functions of perceived roughness. That is, increasing spatial frequency of the texture led to decreased perceived roughness. More work is needed in the auditory modality to determine what the underlying parameters for roughness are to say anything concrete about this observation.

## 7.3 The effects of multimodality on perceived roughness

The function found in all modality conditions replicates the general increasing function found for unimodal haptic and auditory textures. It further confirms that people can successfully judge the relative roughness of a set of simple sinusoidal textures. It shows that, for the model of texture used and the range of frequencies sampled, varying the frequency (number of waves per texture patch) is sufficient to enable people to rate the relative roughness of the set of textures. This is true regardless of the modality of the comparison.

### 7.3.1 Fusion of the unimodal texture percepts

One important observation from the multimodal experiments was that the two unimodal texture percepts fused into a meaningful percept. That is, the haptic and auditory information was believed to have come from the same stimulus - that is one virtual texture with a sound and a feeling rather than two separate stimuli. None of the main hypotheses of this work aimed to directly examine the fusion of the multimodal percept. In fact, it was merely desirable that the two sources were fused enough to allow the user to make multimodal roughness judgments without too much cognitive effort. From pilot studies and qualitative responses it was found that both the haptic and auditory information were attended to and subjects felt that both sources helped

them make their relative roughness judgments. Intersensory discrepancy work might shed more light on whether the haptic or auditory information becomes more dominant under certain conditions or for specific tasks. What the multimodal work does show, however, is that judgments made multimodally differed from those made unimodally in either condition. This might be a result of the multimodal augmentation taking place.

In both multimodal conditions (i.e. haptic texture compared with haptic-auditory texture), the likelihood that *haptically* identical textures (i.e. both being compared have equal haptic spatial frequency) are perceived as the same roughness is significantly lower than in the unimodal haptic condition. That is, people are likely to rate a multimodal texture as different from a haptic texture more so than they would for two purely haptic textures despite the actual haptic spatial frequency being identical in for both textures in both comparisons. This shows that the additional cue of auditory information relating to the number of bumps on the surface makes the texture more discriminable from its purely haptic equivalent. This demonstrates that the auditory cue is attended to in these multimodal comparisons of haptically identical stimuli.

The likelihood of haptically identical textures being judged as the same roughness decreases in the multimodal congruent condition and decreases further in the multimodal incongruent condition. This confirms that multimodality will have an effect on the roughness judgments and also that the incongruence within the multimodality will have an effect. There are frequencies at which the likelihood of the same response is equal regardless of condition. On the other hand, there are also frequencies at which the likelihood of same responses is dramatically different across conditions. The responses at individual frequencies may therefore need further exploration before the effects of multimodality can be used constructively to augment force feedback textures.

## **7.4 The methodology used**

### **7.4.1 The nature of the device**

The force feedback device used in this work was arguably the device with the highest fidelity and resolution on the market at the time of research and had been for a number of years. Although haptic technology research is developing, it seems that the nature of the main qualities of the PHANToM force feedback used will remain in any successful future force based technology. It seems reasonable to be confident that the results from this work will be transferable to similar interaction devices.

The probe attachment used in this work simulates exploration of texture by dragging a pencil across a surface for example. Auditory interaction with textures is most common and most pronounced via probe interaction as described. It was also felt that using the stylus attachment on this device was a reasonable decision because much haptic interaction work takes place via this method of exploration. It is reasonable, however, to suggest that similar tasks using the finger/thimble attachment should be compared with the findings from this thesis.

## 7.4.2 The nature of exploration

It was important in this thesis that the exploration of the textures was not unnaturally restricted, that is there were no time restrictions, or restrictions on the force that is applied or the speed used to explore the virtual objects. Existing studies on both real and virtual texture perception have controlled for many of these free interaction variables in order to achieve more psychologically sound results. While this is required to fully understand the underlying nature of textures, it is also now valid to study texture exploration and perception under conditions more typical of human computer interaction tasks. People may perceive and understand textures one way under the controlled conditions of a laboratory while their experience completely alters when asked to make similar textural judgments on their own, in their own time, and under their own exploratory conditions. The work in this thesis, although conducted under controlled empirical conditions, allowed for this natural exploration to occur in order that the observations might still hold true when such tasks are expected in real haptic interaction applications.

## 7.4.3 The nature of the experimental paradigm

Real texture perception studies most frequently use the magnitude estimation technique to determine functions for perceived roughness of sets of textural stimuli. Much of the work done so far on virtual texture perception has also employed this technique. While this technique is a commonly accepted strong psychophysical method, this technique is perhaps most often employed so as to make the virtual findings comparable with the existing knowledge of real texture perception. The forced choice paradigm is also a commonly accepted psychological method used in perceptual and cognitive psychology. A modified forced choice paradigm was chosen as it was felt it was better suited to the information processing and interaction involved in the study. Although the thesis explores the perception of systematically varied stimuli, it was also intended to capture the decision processes at a higher level than the purely psychophysical. Pure psychophysical understanding was perhaps sacrificed in this experimental work but was replaced by a methodology which better suited the high level processing expected in haptic interaction tasks.

## 7.4.4 The model of texture used

The experimental work conducted employed a simple linear force model of force feedback texture. Of course it is possible that more accurate or realistic representations of haptic virtual texture may be possible as the nature of texture becomes better understood and the body of research on texture simulation advances. What this work has shown, however, is that a simple model of texture is sufficient to convey information regarding the relative roughness of force feedback textures. This thesis confirms that for haptic interaction tasks that require surface differentiation, it may not be necessary to create realistic simulations.

## **7.5 Emerging guidelines for force feedback texture simulation**

The previous sections identified potential effects that the results might have on interaction design. This section summarises these findings in a form suitable for user interaction designers.

### **7.5.1 Perception of roughness via force feedback interaction is possible**

Overall the results show that force feedback textures can be used to convey information regarding the relative perceived roughness of textures. The resolution and discriminability of these textures may not provide a large enough set of distinct textural qualities, however. One design step to avoid is to simply increase the physical geometry of the force feedback textures to make them more distinct. This has been found to make force feedback textures unusable in some circumstances as the users haptic cursor simply gets thrown off the textured areas because of the large forces required. Designers should work therefore within the useable range of force feedback stimuli and be aware of the limitations of force based textures.

Given that an increasing function for perceived roughness was achieved along the single physical geometrical property of spatial frequency, this scale could be used to map a set of distinct and ordinal values to the roughness property. As such, designers could use texture as an information provider in a variety of haptic applications described.

### **7.5.2 Perception of roughness via simple auditory cues is possible**

Auditory texture, based on the same texture model can also be rated successfully along a roughness dimension. Designers should be aware therefore that the auditory modality might be equally successful at conveying virtual textural information. This is despite the relatively abstract and arbitrary nature of the auditory stimuli used. That is, simple auditory notes played in succession can reliably be perceived as texture. This is an extremely important finding given the low cost, and low computational requirements for such a display. Although more realistic, more sophisticated generation of sound to convey texture may improve the realism of virtual texture, computationally simple sounds may be sufficient to convey information about the texture of an object.

### **7.5.3 Multimodal augmentation can improve force feedback textures**

This thesis has confirmed that multimodally augmenting simple force feedback textures with similarly simple auditory textural cues can significantly affect the perceived roughness of those textures. In particular, multimodal augmentation of force feedback textures increases the accuracy of, and reduces the threshold of, relative perceived roughness judgments of the textures.

The many applications that would benefit from the use of haptic texture (discussed in Chapter 2) now have the potential to display an even wider range and resolution of textural information if designers choose to present them multimodally. What is also clear however is that the way in which we integrate multimodal information is complex and that the effects may depend on many factors of the interface and interaction. Now that it is clear that force feedback textures augmented with auditory information can have a significant effect on interaction, more work is required to determine the exact nature of this relationship.

# Chapter 8: Conclusions and future work

---

This chapter summarises the work described in this thesis and the results that have been achieved. It also discusses some of the limitations of the work and how these might be overcome. This chapter then suggests possible areas for future research in the multimodal augmentation of force feedback textures, both to overcome limitations of the work reported here and to address new questions. It concludes by stating the main contributions of the thesis to the area of multimodal and haptic user interfaces.

## 8.1 Contributions of this thesis

### 8.1.1 Haptic characterisation

At the highest level of description, this thesis is concerned with finding ways of improving modern haptic interaction involving force feedback interaction. A problem that has to be overcome in order for this to be achieved is the lack of understanding of the human haptic system as well as a lack of appreciation for the capabilities of force feedback devices and how these might influence the resulting haptic interaction. There was no clear conceptual framework in place in order for this type of research to be carried out. A major contribution of this thesis is a set of haptic definitions that characterise haptics from the perspective of both the haptic device and the human haptic system. This now provides a framework for describing haptic interaction.

### 8.1.2 Empirical investigation of haptic effects in a Graphical User Interface

An aim of this thesis was to evaluate the use of haptic effects to convey information to the user, or to guide user interaction, rather than to increase immersion or make simulations more realistic. The main problem in this research area is that it is not straightforward to achieve this. Despite previous implementations of haptic effects in the conventional desktop, no empirical investigations of such use of haptics exist. A major contribution of this thesis was to provide such an investigation in order that the potential of haptic effects to enhance conventional graphical user interaction was explored. This investigation showed that it is indeed difficult to use haptic effects to guide user interaction to assist graphical user interaction. This provided further motivation for the study of haptic effects to convey information.

### 8.1.3 Empirical evaluation of a multimodal augmentation approach to improving force feedback textures

This thesis aimed to study the potential of a specific haptic effect to convey useful information to the user. The haptic effect of force feedback texture was highlighted as a percept that would be useful in many of the current applications, yet is proving difficult to simulate successfully at low cost with the available interaction devices. A main problem in this area has been that much of the work on texture perception and generation has

been of a psychophysical nature. Although this is required to obtain a pure understanding of the nature of real and virtual textures, it is also necessary to achieve an applied understanding of haptic texture simulation in order that user guidelines for effective interaction with virtual textures can be established. This thesis provides such an applied approach and does so for the auditory augmentation of force feedback texture.

#### 8.1.3.1 Unimodal haptic virtual roughness

This thesis has shown that a simple linear force sine wave based model of texture can convey the sensation of roughness. A major contribution is the confirmation of this under a free exploration, forced choice paradigm. This adds to the research on virtual roughness perception by showing that the effects found exist when haptic exploration of the textures is close to that which might be required in a real haptic texture exploration task.

#### 8.1.3.2 Unimodal auditory virtual roughness

Simple MIDI notes generated and played briefly at the peak of the sine wave texture model were sufficient to convey the sensation of roughness to users. An increasing function similar to that found for the haptic textures was revealed for the auditory modality. Despite previous research showing that auditory cues alone might be sufficient to convey roughness, this research has shown that this is true not only for physical textures but also for simple models of virtual auditory texture. Such cost effective auditory textures can now be used to couple with haptic textures to convey effective multimodal textures.

#### 8.1.3.3 Multimodal virtual roughness

This thesis has shown an effect of multimodality for the haptic-auditory roughness judgment task studied. That is, roughness judgements are different when made purely haptically and multimodally. It has also been shown that this effect may be complex and needs more research. The experiments also revealed an effect of congruency. That is the roughness judgments were affected by whether the multimodal texture was congruent or incongruent in the experiment. The effects of incongruency should also be further investigated, however. One simple notion of incongruency was chosen and there may be many other forms capable of affecting the perception of multimodal textures. A certain amount of incongruency might improve resolution and reliability of judgments, although a high level of incongruency is likely to result in detrimental effects to both the user judgments and the usability of the textures.

## 8.2 Suggestions for future work

### 8.2.1 Investigating other properties of the experience

The task examined in this thesis was that of rating roughness of virtual textures. Future work could adopt a similar approach to that used in this thesis to explore other haptic percepts, such as rigidity, the simulation of which is affected by physical limitations of force feedback hardware. Although some work has been done on

the addition of auditory cues to haptic percepts of rigidity or hardness, none have focused primarily on the multimodal augmentation approach.

Given the complexity of texture revealed in this thesis, it is also reasonable to assume that the work done could be replicated for dimensions of texture other than roughness. This might include the hardness of a surface as described above or it may include other dimensions of texture not yet fully understood in the texture perception literature such as 'stickiness' or 'leatheriness'.

The models of texture used in this thesis were simple linear sine wave based models. The sounds used were also simple and arbitrary in that they were not chosen to represent any real material properties. In spite of this, people were happy to rate all the stimuli in terms of their roughness and commented that the stimuli felt like textures. Subjects also commented that they were unsure as to the exact material they believed the textures to be made of.

### 8.2.2 Investigating other modalities

The modalities used in this experiment were the haptic modality, force feedback in particular, and the auditory modality. Given that texture is also perceived through our visual sense it would be beneficial to investigate the effects of repeating this work with the third modality of vision included.

It would be interesting if the experimental work could be replicated with digitized or sampled haptic and auditory textures rather than the simple representation of texture cues used in this thesis. This may increase the realism of the textures but it would also be interesting to discover whether the roughness discrimination judgments would improve with more sophisticated texture simulation.

### 8.2.3 Other issues

Although the effects of multimodality were examined, the issue of dominance of any one modality was not addressed. The thesis demonstrates that both the haptic and auditory cues are used in some way when multimodal judgments are made. It is not clear, however, whether either modality would be dominant for any particular texture judgment tasks.

These experiments did not collect measurements of confidence or time taken to make judgments. This might be an important factor to include in future work. It may be that unimodal haptic and auditory discriminations of roughness are possible but difficult under some circumstances. Multimodal textures may reduce the cognitive load or processing time involved in such judgments even when there is no direct effect on the magnitude judgments of the textures. Multimodal judgments might result in a reduction in time taken to make the roughness judgments or an increased confidence in the roughness judgments overall.

### 8.3 Final comment

This thesis makes a major contribution in two related research areas. The successful generation of haptic percepts is an important area for research given the increase in haptic interaction devices and applications that are considering haptic interaction as a means to improve the user experience. Despite the ability of haptic technology to progress at a fast rate, it is equally important that the possible limitations of such devices is considered from a human information processing perspective. That is, psychology research and human computer interaction research should continue to examine the human haptic system in order to understand the processes required to perceived and interact with haptic information conveyed via a computer. This thesis embraces exactly this approach and considers the possible limitations of current force feedback devices when used to represent textures. This thesis explored the nature of texture itself and the nature of force feedback generation of texture. A major contribution is additional evidence that simple haptic models of force feedback texture can be used to convey the sensation of roughness.

A second area of research that benefits from the results from this thesis is multimodal interaction research. The force feedback texture problem has been investigated by multimodally augmenting the force feedback textures with auditory textures. A major contribution of this experimental work is evidence that multimodal textures are perceived differently than unimodal textures. This suggests that multimodal augmentation by adding auditory texture cues to force feedback textures works. The exact effects of the multimodal augmentation can now be the main concern for future research given the existence of multimodal effects. In addition, the experimental work revealed an effect of the congruency of the multimodal textures. This finding should encourage research into multimodal augmentation of haptic interfaces to include congruency as a factor that might affect the multimodal interaction.

Overall, force feedback texture simulation has been highlighted as a concern for haptic human-computer interfaces. It has been shown that haptic simulation of this type of percept may be improved by multimodally augmenting the haptic percept. It has also been demonstrated that haptic technology and haptic interfaces can be improved by considering more closely the nature of haptic perception rather than purely focusing on the hardware and software issues. The quantity and quality of virtual textures available through one modality (i.e. force feedback) can be increased by the appropriate addition of texture information to another modality (i.e. audio). An empirical evaluation of the perceived roughness of haptic, auditory, and multimodal (haptic-auditory) virtual textures has confirmed that this multimodal augmentation is indeed a potential approach for improving force feedback textures.

# Appendices: Contents

---

<b>Appendix A:</b> References	98-105
<b>Appendix B:</b> Glossary	106-107
<b>Appendix C:</b> Instructions, questionnaires, and raw data for GUI Experiment (Chapter 3)	108-120
C1: Consent and Introduction form	109
C2: Instruction Sheets	110-115
C3: Debrief form	116
C4: NASA TLX Rating Scale Definitions	117
C5: NASA TLX Rating Scale form	118
C6: Raw Data Collection	119-121
<b>Appendix D:</b> Instructions, questionnaires, and raw data for Haptic Roughness Experiment (Chapter 6)	122-136
D1: Consent and Introduction form	123
D2: Instruction sheets	124-127
D3: Post Experimental Questionnaire	128
D4: Raw Data Collection – Haptic Exp. 1A	129-131
D5: Raw Data Collection – Haptic Exp. 1B	132-136
<b>Appendix E:</b> Instructions, questionnaires, and raw data for Auditory Roughness Experiment (Chapter 6)	137-150
E1: Consent and Introduction form	138
E2: Instruction sheets	139-142
E3: Post Experimental Questionnaire	143-144
E4: Raw Data Collection	145-150
<b>Appendix F:</b> Instructions, questionnaires, and raw data for Multimodal Roughness Experiment (Chapter 6)	151-161
F1: Consent and Introduction form	152
F2: Instruction sheets	153-155
F3: Post Experimental Questionnaire	156-157
F4: Raw Data Collection	158-160
<b>Appendix G:</b> Paper and Publications resulting from thesis	161-192
G1: Putting the ‘Feel into Look and Feel’, <i>CHI 2000</i> .	162-169
G2: The Effective Combination of Haptic and Auditory Textural Information, <i>Lecture Notes in Computing Science, 2001</i> .	170-180
G3: The Effective Combination of Haptic and Auditory Textural Information, <i>Proceedings of Eurohaptics, University of Birmingham, 2001</i> .	181-185
G4: Feeling Rough: Multimodal Perception of Virtual Roughness, <i>Poster presentation, CHI 2001</i> .	186-187
G5: Mixed Feelings: Multimodal Perception of Virtual Roughness, <i>Proceedings of Eurohaptics, University of Edinburgh, 2002</i> .	188-191

## Appendix A: References

---

- [1] Akamatsu, M. & Sate, S. (1994). A Multi-modal mouse with tactile and force feedback, *International Journal of Human Computer Studies*, pp. 73-80.
- [2] Aldridge, R.J., Carr, K., England, R., Meech, J.F. & Solomonides, T. (1996). Getting a Grasp on Virtual Reality. *CHI 96*, Short Papers (April 13-18), pp. 229-330.
- [3] Balbo, S., Coutaz, J. & Salber, D. (1993). Towards Automatic Evaluation of Multimodal User Interfaces. *Intelligent User Interfaces*, pp. 201-208.
- [4] Baron-Cohen, S., Burt, L., Smith-Laittan, F., Harrison, J. & Bolton, P. (1996). Synaesthesia: prevalence and familiarity. *Perception*, 25, pp. 1073-1079.
- [5] Beaudouin-Lafon, M. & Conversy, S. (1996). Auditory Illusions for Audio Feedback. *CHI 96*, Short Papers (April 13-18), pp. 299-300.
- [6] Bellik, Y. & Burger, D. (1994). Multimodal Interfaces: New Solutions to the Problem of Computer Accessibility for the Blind. *CHI 94*, Short Papers(April 24-28), pp. 267-268.
- [7] Ben-Artzi & Marks, L.E. (1995). Visual-auditory interaction in speeded classification: Role of stimulus difference. *Perception & Psychophysics*, 57(8), pp. 1151-1162.
- [8] Biocca, F., Kim, J. & Choi, Y. (2001). Visual Touch In Virtual Environments: An Exploratory Study of Presence, Multimodal Interfaces, and Cross-Modal Sensory Illusions. *Presence*, 10(3), pp. 247-265.
- [9] Boring, E.G. (1942). *Sensation and perception in the history of experimental psychology*. New York: Appleton-Century.
- [10] Bowman, T. (1997). VR Meets Physical Therapy. *Communications of the ACM*, 40(8), pp. 59-60.
- [11] Boyd, C. & Darken, R. (1996). Psychological Issues of Virtual Environment Interfaces. *CHI 96*, Workshops (April 13-18), pp. 426.
- [12] Boyd, C. (1995). Human and Machine Dimensions of 3D Interfaces for Virtual Environments. *CHI 95*, Doctoral Consortium (May 7-11), pp. 41-42.
- [13] Brave, S. & Dahley, A. (1997). inTouch: A Medium for Haptic Interpersonal Communication. *CHI '97*, March, pp. 22-27.
- [14] Brewster, S., Wright, C. & Edwards, A.D.N. (1995). Parallel earcons: reducing the length of audio messages. *International Journal of Human-Computer Interaction*, 43, pp. 153-175.
- [15] Carissa, J. & Sathian, K. (2001). Temporal Cues Contribute to Tactile Perception of Roughness. *Journal of Neuroscience*, 21(14), pp. 5289-5296.
- [16] Colwell, C., Petrie, H. & Kornbrot, D. (1998). Haptic Virtual Reality for Blind Computer Users. *In Proceedings of ASSETS '98*, April, Los Angeles.
- [17] Conner, C.E. & Johnson, K.O. (1992). Neural coding of tactile texture: Comparison of spatial and temporal mechanisms for roughness perception. *Journal of Neuroscience*, 12, pp. 3414-3426.

- [18] Connor, C.E., Hsiao, S.S., Phillips, J.R. & Johnson, K.O. (1990). Tactile roughness: neural codes that account for psychophysical magnitude estimates. *Journal of Neuroscience*, 10, pp. 3823-3836.
- [19] Coutaz, J., Salber, D., Carraux, E. & Portolan, N. (1996). NEIMO, A Multiworkstation Usability Lab for Observing and Analysing Multimodal Interaction. *CHI 96*, Videos (April 13-18), pp. 402-403.
- [20] DiFranco, D.E., Beauregard, G.L. & Srinivasan, M.A. (1997). The Effect of Auditory Cues on the Haptic Perception of Stiffness in Virtual Environments. *Proceedings of the ASME Dynamic Systems and Control Division*, 61, pp. 17-22.
- [21] Dillon, P., Moody, W., Bartlett, R., Scully, P., Morgan, R. & James, C. (2001). Sensing the fabric: to simulate sensation through sensory evaluation and in response to standard acceptable properties of specific materials when viewed as a digital image. *Lecture Notes in Computing Science*, 2058 (Springer-Verlag, Berlin, Germany), pp. 205-217.
- [22] Driver, J. & Spence, C. (1998). Crossmodal links in spatial attention. *Proceedings of the Royal Society of London*, Series B(353), pp. 1-13.
- [23] Durfee, W.K., Hendrix, C.M., Cheng, P. & Varughese, G. (1997). Influence of Haptic and Visual Displays on the Estimation of Virtual Environment Stiffness. *Proceedings of the ASME Dynamic Systems and Control Division*, 61, pp. 139-144.
- [24] Easton, R.D., Srinivas, K., Greene, J. (1997). Do Vision and Haptics Share Common Representations? Implicit and Explicit Memory Within and Between Modalities. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(1), pp. 153-163.
- [25] Edwards, A., Edwards, A.D.N. & Mynatt, E.D. (1994). Enabling Technology for Users with Special Needs. *CHI 94*, Tutorials (April 24-28), pp. 405-406.
- [26] Ekman, G., Hosman, J. & Lindstrom, B. (1965). Roughness, Smoothness, and Preference: A study of Quantitative Relations in Individual Subjects, pp. 18-26.
- [27] Ellis, R.E., Ismaeil, O.M. & Lipsett, M.G. (1996). Design and Evaluation of a High Performance Haptic Interface. *Robotica*, 14, pp. 321-327.
- [28] Fasse, E.D. & Hogan, N. (1994). Quantitative Measurement of Haptic Perception. *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, pp. 3199-3204.
- [29] Freides, D. (1974). Human Information Processing and Sensory Modality. *Psychological Bulletin*, 81(5), pp. 284-310.
- [30] Fritz, J.P. & Barner, K.E. (1996). Stochastic models for haptic texture. *Proceedings of the International Symposium on Intelligent Systems and Advanced Manufacturing*, 3901, pp. 34-44.
- [31] Gibson, J.J. (1962). Observations on Active Touch. *Psychological Review*, 69(6), pp. 477-491.
- [32] Gillan, D. & Cooke, N.J. (1994). Methods of Perceptual and cognitive Research Applied to Interface Design and Testing. *CHI 94*, April 24- 28 (Tutorials), pp. 367-368.
- [33] Goertz, R.C. (1964). Manipulator systems development at ANL, *Proceedings of the 12th Conference on Remote Systems Technology, ANS. Teleoperation Systems*. pp. 117-136.
- [34] Grabowski, N.A. & Barner, K.E. (1998). Data Visualization Methods for the Blind Using Force Feedback and Sonification. *Conference on Telemanipulator and Telepresence Technologies V*, 3524 (November), pp. 131-139.

- [35] Green, D.F. (1998) *Haptic Simulation of Naturally Occuring Textures and Soil Properties*. Master of Science, Engineering, Massachusetts Institute of Technology.
- [36] Haanpaa, D.P. & Roston, G.P. (1997). An Advanced Haptic System For Improving Man-Machine Interfaces. *Computers and Graphics*, 21(4), pp. 443-449.
- [37] Hardwick, A., Rush, J., Furner, S. & Seton, J. (1996). Feeling it as well as Seeing it - Haptic Display within Gestural HCI for Multimedia Telematic Services. *Proceedings of Gesture Workshop '96*, ED. P.A.Harling, & A.D.N.Edwards, pp. 105-116.
- [38] Heller, M.A. (1982). Visual and Tactual Texture Perception: Intersensory Cooperation. *Perception and Psychophysics*, 31(4), pp. 339-344.
- [39] Hill, S.G., Iavecchia, H.P., Byers, J.C., Bittner, A.C., Zaklad, A.L., and Christ, R.E. (1992). Comparison of Four Subjective Workload Rating Scales. *Human Factors*, 34(4), pp. 429-439.
- [40] Hollins, M., Bensmaia, S. & Risner, R. (1998). The duplex theory of tactile texture perception. *Proceedings of the International Society for Psychophysics* (4), pp. 115-120.
- [41] Hollins, M., Faldowski, R., Rao, S. & Young, F. (1993). Perceptual dimensions of tactile surface texture: A multidimensional scaling analysis. *Perception and Psychophysics*, 54(6), pp. 697-705.
- [42] Jacob, R.J.K. & Sibert, L.E. (1992). The Perceptual Structure of Multidimensional Input Device Selection. *CHI 92*, May 3-7, pp. 211-218.
- [43] Jansson, G. (1999). Can a haptic display rendering virtual 3D objects be useful for people with visual impairment? *Journal of Visual Impairment and Blindness*, submitted Jan 1999.
- [44] Jansson, G., Fanger, J., Konig, H, Billberger, K. (1998). Visually Impaired Person's use of the PHANToM for Information about texture and 3D form of Virtual Objects. *Proceedings of the Third PHANToM Users Group Workshop*, Cambridge, MA, MIT.
- [45] Jansson, G., Petrie, H., Colwell, C., Kornbrot, D., Fanger, J., Konig, H., Billberger, K., Hardwick, A., and Furner, S. (1999). Haptic Virtual Environments for Blind People: Exploratory Experiments with Two devices. *International Journal of Virtual Reality* 4(1).
- [46] Jeong, W. (2001). Exploratory User Study of Haptic and Auditory Display for Multimodal Geographical Information Systems. *Extended Abstracts of Conference on Human Factors in Computer Systems*, April 2001, pp. 73-74.
- [47] Johnson, K.O. & Lamb, G.D. (1981). Neural mechanisms of spatial tactile discrimination: Neural patterns evoked by Braille-like dot patterns in the monkey. *Journal of Physiology*, 310, pp. 117-144.
- [48] Jones, B. & O'Neil, S. (1985). Combining Vision and Touch in Texture Perception. *Perception and Psychophysics*, 37(1), pp. 66-72.
- [49] Jones, F.N., Singer, D. & Twelker, P.A. (1962). Interactions among the Somesthetic Senses in Judgments of Subjective Magnitude. *Journal of Experimental Psychology*, 64(2), pp. 105-109.
- [50] Katz, D. (1925/1989). *The world of touch*. New Jersey: Erlbaum.
- [51] Keyson, D.V. (1997). Dynamic cursor gain and tactual feedback in the capture of cursor movements. *Ergonomics*, 40(12), pp. 1287-1298.

- [52] Klatzky, Lederman & Metzger (1985). Identifying Objects by Touch: An 'Expert System'. *Perception and Psychophysics*, 37(4), pp. 299-302.
- [53] Klatzky, R.L. & Lederman, S.J. (1999). Tactile roughness perception with a rigid link interposed between skin and surface. *Perception and Psychophysics*, 61(4), pp. 591-607.
- [54] Klatzky, R.L. & Lederman, S.J. (1995). Identifying objects from a haptic glance. *Perception and Psychophysics*, 57(8), pp. 1111-1123.
- [55] Knecht, S., Kunesch, E., Schnitzler, A. (1996). Parallel and serial processing of haptic information in man: Effects of parietal lesions on sensorimotor hand function. *Neuropsychologia*, 34(7), pp. 669-687.
- [56] Lamb, G.D. (1983). Tactile discrimination of textured surfaces: Psychophysical performance measurements in humans. *Journal of Physiology*, 338, pp. 551-565.
- [57] Lederman, S.J. & Abbott, S.G. (1981). Texture Perception: Studies of Intersensory Organization Using a Discrepancy Paradigm, and Visual Versus Tactual Psychophysics. *Journal of Experimental Psychology: Human Perception and Performance*, 7(4), pp. 902-915.
- [58] Lederman, S.J. & Taylor, M.M. (1972). Fingertip force, surface geometry, and the perception of roughness by active touch. *Perception and Psychophysics*, 12, pp. 401-408.
- [59] Lederman, S.J. (1974). Tactile roughness of grooved surfaces: The touching process and effects of macro- and microsurface structure. *Perception and Psychophysics*, 16, pp. 385-395.
- [60] Lederman, S.J. (1979). Auditory Texture Perception. *Perception*, 8, pp. 93-103.
- [61] Lederman, S.J. (1983). Tactual roughness perception: Spatial and temporal determinants. *Canadian Journal of Psychology*, 37, pp. 498-511.
- [62] Lederman, S.J., Klatzky, R.L., Hamilton, C.L. & Ramsay, G.I. (1999). Perceiving Roughness via a Rigid Probe: Psychophysical Effects of Exploration Speed and Mode of Touch. *Haptics-e The Electronic Journal of Haptic Research* (<http://www.haptics-e.org>), 1(1), pp. 1-20.
- [63] Lederman, S.J., Loomis, J.M. & Williams, D.A. (1982). The role of vibration in the tactual perception of roughness. *Perception and Psychophysics*, 32, pp. 109-116.
- [64] Lederman, S.J., Summers, C., Klatzky, R.L. (1996). Cognitive salience of haptic object properties: Role of modality-encoding bias. *Perception*, 25, pp. 983-998.
- [65] Lederman, S.J., Thorne, G. & Jones, B. (1986). Perception of Texture by Vision and Touch: Multidimensionality and Intersensory Integration. *Journal of Experimental Psychology: Human Perception and Performance*, 12(2), pp.169-180.
- [66] MacLean, K.E. (2000). Designing with Haptic Feedback. *Proceedings of IEEE Robotics and Automation*, San Francisco, CA.(April 22-28).
- [67] Mark, W.R., Randolph, S.C., Finch, M., Van Verth, J.M., Taylor, R.M. (1996). Adding Force Feedback to Graphics Systems: Issues and Solutions. *Computer Graphics: Proceedings of SIGGRAPH '96*, (August 1996) pp. 447-452.
- [68] Marks, L.E. (1982). Bright Sneezes and Dark Coughs, Loud Sunlight and Soft Moonlight. *Journal of Experimental Psychology: Human Perception and Performance*, 8(2), pp. 177-192

- [69] Mason, A.H. (2001). Sensory Information for Collaborative Interaction in Augmented Environments. *Conference on Human Factors in Computer Systems, Extended Abstracts (Doctoral Consortium)*, pp. 77-78.
- [70] Massie, Thomas H., (1993). Design of a Three Degree of Freedom Force-Reflecting Haptic Interface. *SB thesis, MIT EECS Department*. May, 1993.
- [71] McGee, M.R. (2000). The Effective Combination of Haptic and Audio Textural Information. *Proceedings of Haptics in HCI Workshop*, University of Glasgow.
- [72] McGee, M.R., Gray, P.D. & Brewster, S. (2001). Feeling Rough: Multimodal Perception of Virtual Roughness. *In Proceedings of Eurohaptic Workshop*, University of Birmingham .
- [73] Menne, J.M., Menne, J.W. (1972). The Relative Efficiency of Bimodal Presentation as an Aid to Learning. *AV Communication Review*, 20(2), pp. 170-180.
- [74] Miller, T. & Zeleznik., R. (1998). An Insidious Haptic Invasion: Adding Force Feedback to the X Desktop. *UIST ACM*.
- [75] Miner, N., Gillespie, B. & Caudell, T. (1996). Examining the Influence of Audio and Visual Stimuli on a Haptic Interface. *Proceedings of the 1996 IMAGE Conference*, Scottsdale, Arizona, 23-28 June (23-28 June), pp. 1-8.
- [76] Minsky, M. & Lederman, S.J. (1996). Simulated Haptic Textures: Roughness. *Proceedings of the ASME Dynamic Systems and Control Division*, 58, pp. 421-426.
- [77] Minsky, M., Ouh-Young, M., Steele, O. & Brooks, F.P. (1990). Feeling and Seeing: Issues in Force Display. *Computer Graphics*, 24(2), pp. 235-243.
- [78] Mulligan, R.M. & Shaw, M.L. (1980). Multimodal signal detection: Independent decision Vs. integration. *Perception and Psychophysics*, 28(5), pp. 471-478.
- [79] Munch, S. & Dillman., R. (1997). Haptic Output in Multimodal User Interfaces. *Proceedings of the 1997 international conference on Intelligent user interfaces*, pp. 105-112.
- [80] Nigay, L. & Coutaz, J. (1993). A Design Space for Multimodal Systems: Concurrent Processing and Data Fusion. *INTERCHI 93*, April 24-29, pp. 172-178.
- [81] Nygren, T.E. (1991). Psychometric Properties of Subjective Workload Measurement Techniques: Implications for their use in the Assessment of Perceived Mental Workload. *Human Factors*, 33(1), pp. 17-33.
- [82] Oakley, I., McGee, M.R., Brewster, S. & Gray, P. (2000). Putting the Feel in 'Look and Feel'. In *Proceedings of CHI 2000* (The Hague, Netherlands), ACM Press, pp. 415-422.
- [83] Ogi, T. & Hirose, M. (1996). Effect of Multisensory Integration for Scientific Data Sensualization. *Japan Society of Mechanical Engineers international Journal*, 39(2), pp. 411-417.
- [84] Oviatt, S. (1999). Mutual Disambiguation of Recognition Errors in a Multimodal Architecture. *CHI 99*, papers (May 15-20), pp. 576-583.
- [85] Oviatt, S., DeAngeli, A. & Kuhn, K. (1997). Integration and Synchronization of Input Modes during Multimodal Human-Computer Interaction. *CHI 97*, Papers(March 22-27), pp. 415-422.

- [86] Owen (1996). Redundant Sources of Visual Information for Active Control: Additive or Confirmatory? *IMAGE Conference* (Arizona, June, 1996), pp. 1-6.
- [87] Pavani, F., Spence, C. & Driver, J. (2000). Visual Capture of Touch: Out of Body Experiences with Rubber Gloves. *Psychological Science*, 11(5), pp. 353-359.
- [88] Penn, P., Petrie, H., Colwell, C., Kornbrot, D., Furner, S. & Hardwick, A. (2001). The haptic perception of texture in virtual environments: an investigation with two devices. *Lecture Notes in Computing Science, 2058* (Springer-Verlag, Berlin, Germany), pp. 25-30.
- [89] Perry, L.D.S., Smith, C.M., and Yang, S. (1995). An Investigation of Current Virtual Reality Interfaces. *Crossroads - ACM electronic publication*, 3(3), pp. 1-11.
- [90] Posner, M.I., Nissen, M.J. & Klein, R.M. (1976). Visual dominance: an information-processing account of its origins and significance. *Psychological Review*, 83, pp. 157-171.
- [91] Quek, F.K.H. & Petro, M.C. (1993). Human-Machine Perceptual Cooperation. *Interchi '93* (24-29 April), pp. 123-130.
- [92] Ramachandran, V.S., Ramachandran, D. & Cobb, S. (1995). Touching the Phantom Limb. *Nature*, 377(6549), pp. 489-490.
- [93] Repperger, D.W., Hass, M.W. (1997). Design of a haptic stick interface as a pilot's assistant in a high turbulence task environment. *Perceptual and Motor Skills*, 85, pp. 1139-1154.
- [94] Richard, Burdea, Gomez & Coiffet (1994). A Comparison of Haptic, Visual, and Auditive Force Feedback for Deformable Virtual Objects. *Proceedings of the 4th International Conference on Artificial Reality and Telesistance* (July), pp. 49-62.
- [95] Robles-De-La-Torre, G. & Hayward, V. (2001). Force can overcome object geometry in the perception of shape through active touch. *Nature*, 412(July), pp. 445-448.
- [96] Rock, I. & Victor, J. (1964). Vision and touch: An experimentally created conflict between the two senses. *Science*, 143, pp. 594-596.
- [97] Rosenberg, L. (1996). Using Force Feedback to Enhance Human Performance in Graphical User Interfaces. *CHI 96*, April 13-18, pp. 291-292.
- [98] Rosenberg, L.B. (1994). Design of a Virtual Rigid Surface: Haptic/Audio Registration. *CHI'94, Conference Companion*, Boston, MA., USA.(April 24-28), pp. 257-258.
- [99] Ruspini, D. & Khatib, O. (1998). Acoustic Cues for Haptic Rendering Systems. *PHANToM Users Group, Oct 1998*, pp. 1-4.
- [100] Ryan, T.A. (1940). Interrelations of the sensory systems in perception. *Psychological Bulletin*, 37, pp. 659-698.
- [101] Sakaguchi, Y. & Nakano, K. (1994). Haptic Recognition System with Sensory Integration and Attentional Perception. *Multisensor Fusion and Integration for Intelligent Systems*. (Oct. 2-5), pp. 288-291.
- [102] Salber, D. (1994). Observing Users in Multimodal Interaction. *CHI 94, Interactive Posters* (April 24-28), pp. 103-104.

- [103] Salzman, M.C., Dede, C., McGlynn, D. & Loftin, R.B. (1996). ScienceSpace: lessons for Designing Immersive Virtual Realities. *CHI, Interactive Posters* (April 13-18), pp. 89-90.
- [104] Sherrick, C.E. (1970). Temporal Ordering of Events in Haptic Space. *IEEE Transactions on Man-Machine Systems*, 11(1), pp. 25-29.
- [105] Shimoga, K.B. (1993). A Survey of Perceptual Feedback Issues in Dexterous Telemanipulation: Part 2. Finger Touch Feedback. *Proceedings IEEE International Symposium on Virtual Reality*, Seattle, WA, Sept 18-22.
- [106] Siira, J.O. & Pai, D.K. (1996). Fast Haptic Textures. *CHI 96*, April 13-18, 1996 (Short Papers), pp. 231-232.
- [107] Simpson, C. & Prusak, L. (1995). Troubles with Information Overload - Moving from Quantity to Quality in Information Provision. *International Journal of Information Management*, 15(6), pp. 413-425.
- [108] Sivak, M. (1996). The information that drivers use: is it indeed 90% visual? *Perception*, 25, pp. 1081-1089.
- [109] Srinivas, K., Greene, A.J., and Easton, R.D. (1997). Implicit and Explicit memory for Haptically Experienced Two-Dimensional Patterns. *Psychological Science*, 8(3), pp. 243-246.
- [110] Srinivasan, M. & Basgodan, C. (1997). Haptics in Virtual Environments: Taxonomy, Research Status, and Challenges. *Computers and Graphics*, 21(4), pp. 393-404.
- [111] Stein, B.E. & Meredith, M.A. (1993). *The Merging of the Senses*. Cambridge, MA: MIT Press.
- [112] Stevens, S.S. & Harris, J.R. (1962). The Scaling of Subjective Roughness and Smoothness. *Journal of Experimental Psychology*, 64(5), pp. 489-494.
- [113] Strickland, D., Hodges, L., North, M., & Wehorst, S. (1997). Overcoming Phobias by Virtual Exposure. *Communications of the ACM*, 40(8).
- [114] Sutcliffe, A. & Faraday, P. (1994). Systematic Design for Task Related Multimedia Interfaces. *Information and Software Technology*, 36(4), pp. 225-234.
- [115] Tan, H.Z., Srinivasan, M.A., Eberman, B., and Cheng, B. (1994). Human Factors For the Design of Force-Reflecting Haptic Interfaces. *Dynamic Systems and Control*, 55(1), pp. 353-359.
- [116] Taylor, M.M. & Lederman, S.J. (1975). Tactile roughness of grooved surfaces: A model and the effect of friction. *Perception and Psychophysics*, 17, pp. 23-36.
- [117] Tierney, K.J., De Largy, P., and Bracken, M. (1995). Formation of an Equivalence Class Incorporating Haptic Stimuli. *The Psychological Record*, 45, pp. 431-437.
- [118] Walk, R.D. & Pick, H.L. (1981). *Intersensory Perception and Sensory Integration*. New York: Plenum.
- [119] Wall, S.A. & Harwin, W.S. (2001). Interaction of visual and haptic information in simulated environments: texture perception. *Lecture Notes in Computing Science*, 2058 (Springer-Verlag, Berlin, Germany), pp. 108-117.
- [120] Ward, L.M. (1982). Mixed-Modality Psychophysical Scaling: Sequential Dependencies and other Properties. *Perception and Psychophysics*, 31(1), pp. 53-62.

- [121] Ware, C. (1996). Multimedia Output Devices and Techniques. *ACM Computing Surveys*, 28(1), pp. 181-182.
- [122] Warren, McCarthy & Welch (1983). Discrepancy and Nondiscrepancy Methods of Assessing Visual-Auditory Interaction. *Perception and Psychophysics*, 33(5), pp. 413-419.
- [123] Weisenberger, J.M. & Krier, M.J. (1997). Haptic Perception of Simulated Surface Textures via Vibratory and Force Feedback Displays. *Proceedings of the ASME Dynamic Systems and Control Division*, 61, pp. 55-60.
- [124] Weisenberger, J.M., Krier, M.J. & Rinker, M.A. (2000). Judging the Orientation of Sinusoidal and Square-wave Gratings Presented via 2-DOF and 3-DOF Haptic Interfaces. *Haptics-e*, 1(4), pp. 1-20.
- [125] Welch, R. & Warren, D. (1981). Immediate perceptual response to intersensory discordance. *Psychological Bulletin*, 88, pp. 638-667.
- [126] Wenzel, E.M. & Gaver, W.W. (1993). Perceptual vs. Hardware Performance in Advanced Acoustic Interface Design. *Interchi 93, Panel* (April 24-29), pp. 363-366.
- [127] West, A.M. & Cutkosky, M.R. (1997). Detection of Real and Virtual Fine Surface Features with a Haptic Interface and Stylus. *Proceedings of the ASME Dynamic Systems and Control Division*, 61, pp. 159-166.
- [128] Wickens, C.D., Sandry, D.L., Vidulich, M. (1983). Compatibility and Resource Competition between Modalities of Input, Central Processing, and Output. *Human Factors*, 25 (2), pp. 227-248.
- [129] Wu, J., Morita, S. & Kawamura, S. (1996). Human Sensory Fusion on Visual and Tactile Sensing for Virtual Reality. *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, April 1996, pp. 2365-2370.
- [130] Zhai, S., Smith, B.A., Multi-Stream Input: An Experimental Study of Document Scrolling Methods, *IBM Systems Journal*, Vol 38, No.4, p 642-651.

# Appendix B: Glossary of Terms

---

<b>Audio:</b>	Sound. The medium we use our auditory modality (or sense of hearing) to process.
<b>Auditory Cognition:</b>	The study of psychological effects derived by the auditory system.
<b>Channel:</b>	A means of communication or expression: as a path along which information passes.
<b>Computer haptics:</b>	The field concerned with the techniques and processes associated with generating, displaying and processing of haptic stimuli to or by the human user (Srinivasan).
<b>Congruent:</b>	Generally meaning superposable so as to be coincident throughout.
<b>Cutaneous:</b>	Pertaining to the skin itself or the skin as a sense organ. Includes sensation of pressure, temperature, and pain.
<b>Degrees of Freedom:</b>	The number of independent parameters required to specify the position and orientation of an object. Used to classify force feedback devices.
<b>End Effector:</b>	An interface mechanism attached to the distal point of a robot manipulator or haptic interface.
<b>Exploratory Procedure:</b>	A stereotypical series of hand movements used unconsciously to extract information regarding object properties (Lederman).
<b>Force Feedback:</b>	Relating to the mechanical production of information sensed by the human kinesthetic system. Used specifically to describe a class of haptic device.
<b>Haptic:</b>	From Greek <i>haptikos</i> meaning to grasp or touch. Relating to the sense of touch. Dependent on feeling by touch.
<b>Haptically:</b>	(Adverb) Relating to something done through the sense of touch.
<b>Haptics:</b>	The study of touch.
<b>Integration:</b>	The act or process of an instance of integrating coordination of mental processes into a normal effective percept or concept in the individual's environment. Used specifically to describe the processing of multisensory information.
<b>Interaction:</b>	Mutual or reciprocal action or influence. Used specifically to describe how multisensory pieces of information influence one another.
<b>Kinesthetic:</b>	Meaning the feeling of motion. Relating to sensations originating in muscles, tendons, and joints.
<b>Mechanoreceptor:</b>	A nerve ending in the human skin sensitive to mechanical stimulus such as stretch, pressure, or vibration.
<b>Modality:</b>	A sensory system, a sense. Usually qualified to specify the sense intended.
<b>Multi-modal:</b>	Having or involving several modes, modalities, or maxima. Specifically used to mean a computer interface requiring the intended and primary use of more than one sensory mode for interaction

<b>Multisensory:</b>	Relating to or involving several physiological senses
<b>Perception:</b>	Awareness of the elements of environment through physical sensation.
<b>PHANToM:</b>	Personal HAPTic iNTERface Mechanism. A specific force feedback device made by SensAble Technologies Inc. used in the experimental work in this thesis.
<b>Proprioception:</b>	Relating to sensory information about the state of the body (including cutaneous, kinesthetic, and vestibular sensations).
<b>Roughness (rough):</b>	A primary dimension used to sense, perceive, and mentally judge the texture of a surface or object.
<b>Simulation:</b>	Attempting to predict aspects of the behaviour of some system by creating an approximate (mathematical) model of it.
<b>Spatial:</b>	Relating to, occupying, or having the character of space.
<b>Skin:</b>	The external covering of the body. It is an extremely complex and vital organ involved in the mediation of the cutaneous senses.
<b>Tactile:</b>	Pertaining to the cutaneous sense but most frequently the sensation of pressure rather than temperature or pain.
<b>Tactile Device:</b>	A device that uses the tactile senses to apply stimulation to the user.
<b>Teleoperation:</b>	The act of controlling a device, typically a robot manipulator, from a remote location.
<b>Temporal:</b>	Of or relating to time as distinguished from space.
<b>Texture:</b>	A measure of the variation of the intensity of a surface, quantifying properties such as smoothness and roughness amongst other things.
<b>Touch:</b>	Loosely and generally, the contact of some object with the body or the sensory experience which accompanies such contact.
<b>Usability:</b>	The effectiveness, efficiency, and satisfaction with which users can achieve tasks in a particular environment of a product.
<b>Vestibular:</b>	Pertaining to the perception of head position, acceleration, and deceleration.
<b>Virtual Environment:</b>	A world or situation simulated on a computer. Can incorporate haptic, graphical or auditory cues.
<b>Virtual Reality:</b>	A human-computer interface in which the computer creates a sensory-immersing environment that interactively responds to and is controlled by the user.
<b>Workspace:</b>	The volume through which a robot manipulator, haptic interface or other framework is free to move.

## **Appendix C: Instructions, questionnaires, and raw data for GUI experiment**

---

# Scroll Bar Experiment

Marilyn Rose McGee  
Department of Computing Science  
University of Glasgow

Thank you for agreeing to take part in this experiment.

The task you are about to perform is a simple and common computing task. You should attempt to complete the task as quickly as possible but with as few errors as possible.

You should read the instructions very carefully.

If you have difficulty understanding any part of the experiment please let the experimenter know immediately.

Please fill out the details below.

Name:  Age:  Gender:

How frequently do you use computers?

- Less than 1 hour per week
- Between 1 hour and 10 hours a week
- More than 10 hours per week

Which of these do you most often use (please tick only one option):

- a mouse to operate a scroll bar
- page up/down keys on the keyboard to operate a scroll bar

When you use a scroll bar, which of the following do you perform most often (you may tick more than one option):

- up/down arrows at either end of the scroll bar
- the thumb wheel in the middle of the scroll bar
- the area either side of the thumb wheel


Please tell the experimenter you have reached the end of this page.

**C1: Consent and Introduction form**

# Instructions

You will be asked to perform this experiment in two almost identical sections. After each section you will be given a break and then asked to complete a short workload assessment.

## The Interface



There are three windows on the screen.

The largest window is the Data window. It is similar to a text window that is found in most word processor packages (e.g. a word document). This contains the data file that you will be required to search.

The small window at the top left is your Instruction window. It will tell you which 4-digit number code you must find in the large data window.

The small window at the bottom left is the Target window. This is where you should send the target code that you have just found and highlighted in the data window.

## The Data File

In the Data window, you will be presented with a data file containing a very large list (2000) of 4 digit number codes. The codes increase in value from the start of the file to the end (i.e. they are ordered). There are no codes starting with a zero. That is, the codes range from 1000 to 9999. There are three of these codes per line. A section of the file might therefore be:

1003	1009	1014
1023	1026	1031
1039	1047	1052

## The Task

You must search through the codes until you find the target code indicated to you by the instruction in the small window at your left.


e.g. 1031

When you find the target code, it must be highlighted. This is achieved by dragging horizontally from left to right across the entire code or double clicking on the correct number code. Remember that the switch on the stylus acts as the left mouse button.

When the correct code has been highlighted, you should send this code into the Target window by pressing the << button. This will send the highlighted code to the appropriate window to form the new smaller list of target codes.

When this has been completed, the instruction window will change to tell you to find the next code to be found.

## The Scroll Bar



You are restricted to using the vertical scroll bar on the right of the data window to move around the document. That is, you can only scroll up and down (not left and right) and the page up/down arrows on the keyboard cannot be used, as it is the scroll bar actions that we are interested in evaluating. You can use the scroll bar in the same way you would expect to use most common scroll bars. The only difference here is that you are controlling the cursor with a force feedback device called the Phantom.

## The Force Feedback device

You will probably be used to using a mouse to control a cursor. Particularly in tasks that involve scrolling, you will be familiar with positioning the cursor (a) directly over the thumb on the scroll bar or (b) over the up and down arrow buttons or (c) in the scrolling area above or below the thumb. Instead of moving the mouse from left to right and up and down across the mouse mat you will move the stylus (pen) in 3D space and use the silver switch on the pen as if it were the left mouse button.

### Completing the task

When you have performed the required operation on all 50 codes, a message will appear saying that you have reached the end. You should then let the experimenter know that you have finished.

You will then be asked to take a short break and fill in the workload questionnaire given to you.

Workload has been defined as "the effort invested by the human operator into task performance" (Hart and Wickens, 1990). You will be asked to fill in a set of 6 ratings defined by NASA Human Performance Group. A seventh was added to allow you to rate the level of frustration experienced during the task. You will be given an explanation of the 7 factors before you should fill in these subjective ratings.

After you have rated your subjective workload you will repeat the experiment under the other condition.

Thank you again for agreeing to take part.

## Visual Condition

In this condition, when you are over the up/down arrow buttons, you should know because you will see that the cursor is over them graphically. In addition, when you are in the rest of the scroll bar area you will see that your cursor is over the scroll bar graphically.

As in any normal scrolling operation, you need to be aware of where you are on the scroll bar to be able to operate it most effectively. In this condition, the feedback you require to tell you where you are is presented visually.

Please try to read, find, and highlight the code as quickly as possible while trying to ensure you are always finding and selecting the correct code each time.

## Haptic Condition

In this condition, when you are over the up/down arrow buttons, you should know because you will see that the cursor is over them graphically AND you will feel your cursor 'dropping into' a hole in the button. In addition, when you are in the rest of the scroll bar area you will see that your cursor is over the scroll bar graphically AND you will feel your cursor falling into a trough in the scroll bar area.

As in any normal scrolling operation, you need to be aware of where you are on the scroll bar to be able to operate it most effectively. In this condition, the feedback you require to tell you where you are is presented both visually and haptically.

Please try to complete the task as quickly as possible while trying to ensure you are selecting the correct code each time.

## Training Session

This is a short training session to allow you to become familiar with the task.

You should perform this in the same way as you have been told to perform the actual experiment although you will not be expected to fill out workload ratings at this stage.

You will experience both conditions for an equal amount of time and you will be told which one you are about to experience.

If you have any questions during or after this training please let the experimenter know before you progress to the experiment.

## Debrief

Thank you again for taking time to complete this experiment.

This is part of research that is looking at the possible effects of haptically enhancing the standard graphical user interface.

At the moment, we see feedback from our interaction and more recently we have begun to hear feedback from our interaction. Force feedback technology such as the Phantom allows us to add haptic effects to standard widgets such as scroll bars, buttons, and menus and to study their effect on interaction.

If you are interested in this sort of work then you are welcome to mail me with specific questions. Alternatively you can visit my web site and follow some of the links from there.

Marilyn Rose McGee

Mail - [mcgeemr@dcs.gla.ac.uk](mailto:mcgeemr@dcs.gla.ac.uk)

<http://www.dcs.gla.ac.uk/~mcgeemr>

## Rating Scale Definitions

Title	Endpoints	Description
<b>Mental Demand</b>	<i>Low/High</i>	How much mental, visual and haptic activity was required? (e.g. thinking, deciding, calculating, looking, feeling, exploring)
<b>Physical Demand</b>	<i>Low/High</i>	How much physical activity was required? (e.g. pushing, pulling, turning, controlling)
<b>Time Demand</b>	<i>Low/High</i>	How much time pressure did you feel because of the rate at which things occurred? (e.g. slowly, leisurely, rapid, frantic)
<b>Effort Expended</b>	<i>Low/High</i>	How hard did you work (mentally and physically) to accomplish your level of performance?
<b>Performance Level Achieved</b>	<i>Poor/Good</i>	How successful do you think you were in doing the task set by the experimenter? How satisfied were you with your performance? Don't just think of your success in terms of pressing buttons, but how you felt you performed.
<b>Frustration Experienced</b>	<i>Low/High</i>	How much frustration did you experience? (e.g. were you relaxed, content, stressed, irritated, discouraged?)
<b>Fatigue Experienced</b>	<i>Low/High</i>	Did you find the condition physically tiring or straining?


Name :

Condition Completed :


**Mental Demand**




**Physical Demand**



**Time Pressure**



**Effort Expended**




**Performance Level Achieved**



**Frustration Experienced**



**Fatigue Experienced**



participant	average time per code (seconds)		no.times on/off scroll bar		no.times on/off <i>adjusted</i>	
	<b>haptic</b>	<b>visual</b>	<b>haptic</b>	<b>visual</b>	<b>haptic</b>	<b>visual</b>
1	11.614475	12.0246	75	121	35	81
2	10.807825	13.39765	86	80	46	40
3	8.830075	7.888675	82	84	42	44
4	17.06525	15.148825	73	89	33	49
5	10.482825	14.6289	108	92	68	52
6	17.105475	15.37655	116	104	76	64
7	9.51095	9.73945	87	96	47	56
8	10.408575	12.704675	73	125	33	85
9	8.6547	11.119525	91	98	51	58
10	9.7828	9.621075	87	108	47	68
11	10.138675	9.142175	104	119	64	79
12	10.74415	8.8457	109	96	69	56
13	14.85235	8.6082	85	91	45	51
14	10.0574	8.17695	83	75	43	35
15	14.197675	12.492575	151	158	111	118
16	17.105475	15.37655	110	104	70	64
17	13.894925	13.457025	132	174	92	134
18	10.7336	10.283975	130	136	90	96
19	13.838675	16.7289	70	103	30	63
20	9.51095	9.73945	87	96	47	56
<b>mean</b>	<b>haptic</b> 11.9668413	<b>visual</b> 11.725071	<b>haptic</b> 96.95	<b>visual</b> 107.45	<b>haptic</b> 56.95	<b>visual</b> 67.45

**C6.1: Raw Data Collection – performance time and errors**

Workload Factor	mental demand		physical demand		time pressure		effort expended		visual
	condition	haptic	visual	haptic	visual	haptic	visual	haptic	
		5	6	6	7	3	4	10	16
		7	7	12	15	14	9	13	16
		10	12	9	8	10	12	9	10
		18	18	12	5	12	14	14	19
		5	5	10	14	5	6	8	10
		15	18	4	8	13	17	15	16
		7	4	7	12	4	2	7	10
		8	15	9	11	6	20	11	17
		15	16	11	7	8	9	9	8
		16	17	11	9	16	16	16	17
		14	14	13	14	18	17	14	17
		12	7	6	9	11	8	11	7
		16	12	4	16	6	4	9	5
		9	9	9	13	11	11	12	13
		5	6	2	4	8	6	5	7
		8	12	12	13	10	14	12	12
		5	15	14	13	7	13	9	16
		2	4	6	2	8	4	2	8
		8	7	15	16	18	15	19	17
		2	2	3	9	4	9	5	9
Mean Score		9.35	10.3	8.75	10.25	9.6	10.5	10.5	12.5

**C6.2: Raw Data Collection – Subjective workload ratings (1/2)**

Workload Factor	frustration experienced		fatigue experienced		performance Level achieved	
	haptic	visual	haptic	visual	haptic	visual
condition	10	12	16	18	16	14
	14	15	6	5	5	5
	4	6	6	6	5	6
	15	14	12	11	7	12
	4	3	11	16	4	5
	7	11	11	12	2	4
	4	7	8	12	6	4
	6	12	12	17	13	7
	13	16	13	5	5	7
	6	12	10	8	2	4
	14	17	8	16	8	5
	13	9	8	7	7	6
	13	8	4	4	3	2
	7	11	12	14	8	11
	7	8	11	10	5	12
	13	17	15	16	13	13
	8	15	12	15	4	6
	14	10	2	10	8	6
	16	15	10	13	8	5
	8	9	3	2	6	9
Mean Score	9.8	11.35	9.5	10.85	6.75	7.15

**C6.3: Raw Data Collection – Subjective workload ratings (2/2)**

# **Appendix D: Instructions, questionnaires, and raw data for haptic roughness experiment**

---

# Judging the Roughness of Surfaces

## Introduction

Thank you for taking part in this experiment.

If at any time you do not understand what is being asked of you, or you need to stop for any reason at all then please let me know.

This is not a test - it is simply a means of collecting data on how rough each of the textures presented seems to different people.

Please fill in the following details and then let the experimenter know you are ready to continue.

Name:	_____
Age (optional):	_____
Sex:	Male / Female
To your knowledge, do you have normal sense of touch?	Yes / No
If no, then please give details below	
_____	
_____	
_____	
Have you ever used this force feedback device (the PHANToM) before?	Yes / No

**D1: Consent and Introduction form**

# Judging the Roughness of Surfaces

## Instructions

Please read the following instructions very carefully.

You can ask the experimenter to clarify anything that doesn't make sense as you are reading them.

The following experiment looks at how people perceive textures. In particular, I am interested in how *rough* we judge different surfaces to be.

You will be given some practice trials that will allow you to become familiar with the task and the way the textures will be presented to you before the actual experiment begins.

There are no right or wrong answers, only your own judgment of how rough the surface seems. Your initial feeling is most important so try not to spend too much time worrying over each trial.

When you have made your decision you simply check the box(es) that correspond(s) to your decision for that trial. Clicking the button labeled "Next" will begin the next trial and instructions will appear at the end to tell you when you have completed the experiment.

The next couple of pages explain the device and the interface in more detail.

## **Haptic Textures**

The textures that you can touch are experienced by *dragging* the pen-like probe of the PHANToM device back and forth across the virtual surfaces. The textured area is outlined visually by a box but the actual texture of the surface cannot be seen - only felt. The texture is a patch on the back wall of the screen so you might need to apply a small force forward to actually feel the texture. The device will provide feedback to you to let you know how rough the surface you are dragging across is.

## **Making your response**

Sometimes the texture on the left will feel roughest. Sometimes the texture on the right will feel roughest. Sometimes the textures will feel exactly the same, and sometimes you may not be able to compare the textures on the same roughness scale. Each response is equally valid.

# The PHANToM Device

The device you will be using to explore the textures AND make all your responses is called the PHANToM.

It is similar to a mouse in that you can enter input to the computer via the device.

This device can also send feedback that you can *feel* in the form of output from the computer back to you.


The forces sent back to you simulate how it feels to interact with virtual objects by touching them.

Hold the pen like probe on the PHANToM as you would hold a pen. You will have a chance to make sure you are comfortable with this before you are asked to do anything.

On the end of the probe, there is a silver switch that you can press just like you would press the left button on a mouse. Pressing this switch while over a target/button will select that target/button.

If you have any difficulties during the experiment then please let the experimenter know.

# Experimental Interface



(1) You can experience the texture by dragging the pen like probe of the PHANToM across the rectangular patches labeled 'left' and 'right'.

(2) Using the PHANToM probe, positioning the cursor over a button or check box and clicking the silver switch on the end of the probe will select that button.

(3) A selected checkbox will turn black.

(4) Pressing the button labeled 'next' will present the next two textures to be compared.

(5) Let the experimenter know when the experiment ends and no more textures are being presented to you.

# Additional Questions

(1) On average, how easy did you find it to tell that any two textures felt different?

Very easy    quite easy    about average    quite hard    very hard

(2) How many *different* textures do you think you might have been used to make up the whole set?

-----

(3) On average, how easy did you find it to choose which surface you thought was roughest?

Very easy    quite easy    about average    quite hard    very hard

(4) Can you draw below what you think the textures might have looked like?



**Haptic Perception of Virtual Roughness  
Experiment 1(a)**

**RAW DATA**

**Paired comparisons  
of 6 virtual textures**

<b>Participant</b>		10,10	10,15	10,20	10,25	10,30	10,35	15,15	15,20
1	s	R	R	R	R	R	R	R	R
	R	L	R	R	L	R	R	R	s
2	s	s	s	R	R	R	R	R	s
	s	s	L	R	L	R	s	R	R
3	s	s	s	s	R	R	s	R	R
	R	R	R	R	R	R	s	R	R
4	L	L	L	R	R	R	R	R	s
	s	L	s	L	L	L	R	s	s
5	s	s	R	R	R	R	L	s	L
	L	R	R	R	R	R	L	L	L
6	L	s	R	s	R	R	L	s	s
	L	L	L	L	R	R	s	s	R
7	L	s	s	R	R	R	s	R	s
	s	s	R	R	R	R	s	s	L
8	s	s	L	L	L	L	R	L	R
	s	R	L	L	L	L	s	R	R
9	s	s	s	s	R	R	s	R	R
	R	R	R	R	R	R	s	R	R
10	s	s	R	R	R	R	L	s	L
	L	R	R	R	R	R	L	L	R
11	L	s	s	R	R	R	s	R	s
	s	s	R	R	R	R	s	s	s
12	L	s	R	s	R	R	L	s	s
	L	L	L	L	R	R	s	s	s
texture on right rated as:	<b>more rough</b>	3	6	12	15	19	21	6	9
	<b>less rough</b>	9	5	6	5	5	3	6	3
	<b>same roughness</b>	12	13	6	4	0	0	12	12

### Frequency Pair

	15,25	15,30	15,35	20,20	20,25	20,30	20,35	25,25	25,30	25,35	30,30	30,35	35,35
1	R	R	s	s	s	R	R	R	R	s	s	s	R
2	R	R	R	R	s	R	R	R	s	s	s	s	R
3	R	R	R	R	s	s	R	L	s	R	s	s	R
4	R	R	R	R	s	s	R	R	R	s	s	s	s
5	s	R	R	s	R	R	R	s	s	R	s	s	R
6	s	R	R	R	s	R	R	s	s	s	s	s	s
7	L	R	R	s	R	L	R	L	s	L	s	R	s
8	L	R	L	s	s	s	L	s	s	R	s	s	s
9	R	R	R	L	s	R	R	s	s	R	L	s	s
10	R	R	R	s	L	L	L	s	s	s	s	L	L
11	L	R	R	s	L	s	L	s	s	L	s	s	s
12	s	R	R	s	s	s	s	L	s	R	s	s	L
13	R	R	R	s	s	R	R	s	s	R	s	s	s
14	R	R	R	L	s	R	R	s	s	R	s	s	s
15	s	L	L	s	L	L	L	s	s	L	R	L	s
16	R	L	L	s	s	L	L	s	s	s	s	L	S
17	s	R	R	s	R	R	R	s	s	R	s	s	R
18	s	R	R	R	s	R	R	s	s	s	s	s	s
19	R	R	R	L	s	R	R	s	s	R	L	s	s
20	R	R	R	s	L	L	L	s	s	s	s	L	L
21	R	R	R	s	s	R	R	s	s	R	s	s	s
22	R	R	R	L	s	R	R	s	s	R	s	s	s
23	L	R	R	s	L	s	L	s	s	L	s	s	s
24	s	R	R	s	s	s	s	L	s	R	s	s	L
>	13	22	20	5	3	12	15	3	2	12	1	1	5
<	4	2	3	4	5	5	7	4	0	4	2	4	4
<>	7	0	1	15	16	7	2	17	22	8	21	19	15

**Roughness  
Score**

(no. of times  
texture rated as  
roughest of pair)

<b>F= 10</b>	<b>F= 15</b>	<b>F= 20</b>	<b>F= 25</b>	<b>F= 30</b>	<b>F= 35</b>
<b>24</b>	<b>18</b>	<b>38</b>	<b>35</b>	<b>59</b>	<b>69</b>

**Haptic Perception of  
Virtual Roughness  
Experiment  
1(B)**

**Raw Data**

	frequency pair	5 10	5 15	5 20	5 25	5 30	5 35	5 40	5 45
<b>Participant</b>	1	0	0	0	0	0	0	0	0
		0	0	0	0	0	0	0	0
	2	0	3	3	3	3	0	3	3
		3	3	3	0	0	3	3	3
	3	0	0	0	0	0	0	0	0
		0	0	0	0	0	0	0	0
	4	2	2	2	2	2	2	2	2
		2	2	2	2	2	2	2	2
	5	0	0	0	0	0	3	0	0
		0	3	0	0	3	3	3	0
	6	0	0	0	0	1	0	0	0
		0	0	0	0	0	0	0	0
	7	0	0	0	0	0	0	0	0
		0	0	0	0	0	0	0	0
	8	0	0	0	0	0	0	0	0
		0	0	0	0	0	0	0	0
	9	0	0	0	0	0	3	0	0
		0	3	0	0	3	3	3	0
	10	0	3	3	3	3	0	3	3
		3	3	3	0	0	3	3	3

**Haptic Perception of  
Virtual Roughness  
Experiment  
1(B)**

**Raw Data**

	frequency pair	10 15	10 20	10 25	10 30	10 35	10 40	10 45
Participant 1	1	1	0	0	0	0	0	0
		0	1	0	0	0	0	0
2		0	0	0	3	2	0	3
		0	0	3	0	0	2	0
3		1	0	0	0	0	1	0
		1	0	0	0	0	0	0
4		2	0	3	0	2	2	2
		2	2	2	0	0	0	2
5		1	0	0	0	0	0	0
		0	0	1	0	0	0	0
6		2	0	1	0	0	0	0
		2	2	0	0	0	0	0
7		1	0	0	0	0	0	0
		0	1	0	0	0	0	0
8		1	0	0	0	0	1	0
		1	0	0	0	0	0	0
9		1	0	0	0	0	0	0
		0	0	1	0	0	0	0
10		0	0	0	3	2	0	3
		0	0	3	0	0	2	0

**Haptic Perception of  
Virtual Roughness  
Experiment  
1(B)**

**Raw Data**

	frequency pair	15 20	15 25	15 30	15 35	15 40	15 45
<b>Participant</b>	1	0	0	0	0	0	0
	2	1	0	0	0	0	0
3	3	0	0	0	0	3	0
	1	1	1	0	0	0	0
4	1	1	0	0	2	0	2
	1	0	0	0	0	0	0
5	1	0	0	0	0	0	0
	1	1	0	0	0	0	0
6	1	2	2	0	0	0	0
	0	2	0	1	0	0	0
7	0	0	0	0	0	0	0
	2	1	0	0	0	0	0
8	1	1	0	0	0	0	0
	1	1	0	0	0	0	0
9	1	0	0	0	0	0	0
	1	1	0	0	0	0	0
10	0	0	0	0	0	0	0
	3	0	0	0	0	3	0

**Haptic Perception of  
Virtual Roughness  
Experiment  
1(B)**

**Raw Data**

	frequency pair	20 25	20 30	20 35	20 40	20 45
<b>Participant</b>	1	1	0	0	0	0
		1	1	0	0	0
	2	0	0	0	0	2
		1	1	2	0	2
	3	1	0	0	0	0
		1	1	0	0	0
	4	2	0	0	0	0
		0	2	0	0	0
	5	1	1	0	0	0
		1	0	0	2	0
	6	0	1	1	0	0
		1	1	1	0	0
	7	1	0	0	0	0
		1	1	0	0	0
	8	1	0	0	0	0
		1	1	0	0	0
	9	1	1	0	0	0
		1	0	0	2	0
	10	0	0	0	0	2
		1	1	2	0	2

**Haptic Perception of  
Virtual Roughness  
Experiment  
1(B)**

**Raw Data**

	frequency pair	30 35	30 40	30 45	35 40	35 45	40 45
<b>Participant</b>	1	1	0	0	1	1	1
		1	0	0	1	1	1
2	1	1	1	1	1	1	1
		0	0	2	1	1	2
3	1	1	1	0	2	1	1
		1	1	1	1	1	1
4	1	1	1	2	1	1	0
		3	1	0	1	1	0
5	1	1	1	0	2	1	2
		1	1	0	1	1	1
6	1	1	2	1	1	0	2
		2	1	2	0	0	1
7	1	1	0	0	1	1	1
		1	0	0	1	1	1
8	1	1	1	0	2	1	1
		1	1	1	1	1	1
9	1	1	1	0	2	1	2
		1	1	0	1	1	1
10	1	1	1	1	1	1	1
		0	0	2	1	1	2

## **Appendix E: Instructions, questionnaires, and raw data for auditory roughness experiment**

---

# Judging the Roughness of Surfaces

## Introduction

Thank you for taking part in this experiment.

If at any time you do not understand what is being asked of you, or you need to stop for any reason at all then please let me know.

This is not a test - it is simply a means of collecting data on how rough each of the textures presented seems to different people.

Please fill in the following details and then let the experimenter know you are ready to continue.

Name:	_____
Age (optional):	_____
Sex:	Male / Female
To your knowledge, do you have normal sense of hearing? Yes / No	
If no, then please give details below	
_____	
_____	
_____	
Have you ever used this force feedback device (the PHANToM) before?	
Yes / No	

**E1: Consent and Introduction Form**

# Judging the Roughness of Surfaces

## Instructions

Please read the following instructions very carefully.

You can ask the experimenter to clarify anything that doesn't make sense as you are reading them.

The following experiment looks at how people perceive textures. In particular, I am interested in how *rough* we judge different surfaces to be.

You will be given some practice trials that will allow you to become familiar with the task and the way the textures will be presented to you before the actual experiment begins.

There are no right or wrong answers, only your own judgment of how rough the surface seems. Your initial feeling is most important so try not to spend too much time worrying over each trial.

When you have made your decision you simply check the box(es) that correspond(s) to your decision for that trial. Clicking the button labeled "Next" will begin the next trial and instructions will appear at the end to tell you when you have completed the experiment.

The next couple of pages explain the device and the interface in more detail.

## **Audio Textures**

The textures that you can hear are experienced by *dragging* the pen-like probe of the PHANToM device back and forth across the virtual surfaces. The textured area is outlined visually by a box but the actual texture of the surface cannot be seen or felt - only heard. The texture is a patch on the back wall of the screen so you might need to apply a small force forward to actually hear the texture. The device will provide feedback to you to let you know how rough the surface you are dragging across is.

## **Making your response**

Sometimes the texture on the left will feel roughest. Sometimes the texture on the right will feel roughest. Sometimes the textures will feel exactly the same. Each response is equally valid. Please don't expect to be wrong or right - just try to make a quick initial judgment of the surfaces.

## The PHANToM Device

The device you will be using to explore the textures and make all your responses is called the PHANToM.

It is similar to a mouse in that you can enter input to the computer via the device.


This device can also be programmed to send feedback that you can *feel* in the form of output from the computer back to you.

Hold the pen like probe on the PHANToM as you would hold a pen. You will have a chance to make sure you are comfortable with this before you are asked to do anything.

On the end of the probe, there is a silver switch that you can press just like you would press the left button on a mouse. Pressing this switch while over a target/button will select that target/button.

If you have any difficulties during the experiment then please let the experimenter know.

# Experimental Interface



(1) You can experience the texture by dragging the pen like probe of the PHANToM across the rectangular patches labeled 'left' and 'right'.

(2) Using the PHANToM probe, positioning the cursor over a button or check box and clicking the silver switch on the end of the probe will select that button.

(3) A selected checkbox will turn black.

(4) Pressing the button labeled 'next' will present the next two textures to be compared.

(5) Let the experimenter know when the experiment ends and no more textures are being presented to you.

## E2.4: Instructions Sheets (4/4)

# Additional Questions

Please circle the appropriate response:

(1) On average, how easy did you find it to tell that any two textures felt different?

Very easy    quite easy    about average    quite hard    very hard

(2) How many *different* textures do you think might have been used to make up the whole set?

-----

(3) On average, how easy did you find it to choose which surface you thought was roughest?

Very easy    quite easy    about average    quite hard    very hard

(4) Can you draw below what you think the textures might have looked like if you could see them?



**PTO**



**Audio Roughness Experiment  
Raw Data**

Participant	5 10	5 15	5 20	5 25	5 30	5 35	5 40	5 45
1	0 0							
2	0 0							
3	0 1	0 0						
4	0 0							
5	1 0	0 0						
6	0 0							
7	0 0	1 1						
8	2 2	0 2	2 2	2 2	2 0	2 0	2 0	0 2
9	2 2							
10	0 0	1 0	0 0	0 0	0 0	0 0	0 0	2 2
11	2 2							
12	0 1	0 0						
<b>Right Rougher</b>	15	18	18	18	18	18	18	15
<b>Left Rougher</b>	6	5	6	6	5	5	5	7
<b>Same</b>	3	1	0	0	0	0	0	2
<b>agreement</b>	9	10	12	12	11	11	11	11
<b>discrepant</b>	0	1	0	0	1	1	1	1
<b>indecision</b>	3	1	0	0	0	0	0	0

**Audio Roughness Experiment  
Raw Data**

Participant	10 15	10 20	10 25	10 30	10 35	10 40	10 45
1	1 1	0 0	0 0	0 0	0 0	0 0	0 0
2	1 1	0 1	0 0	0 0	0 0	0 0	0 0
3	1 0	0 2	0 2	2 0	2 1	0 0	0 2
4	1 0	0 0	0 2	0 0	0 0	0 0	0 0
5	1 0	0 0	0 0	0 0	0 0	0 0	0 0
6	1 0	0 1	0 0	0 0	0 0	0 0	0 0
7	1 1	1 0	0 0	0 0	0 0	0 0	0 0
8	1 1	2 1	2 2	2 2	0 2	2 2	2 2
9	2 0	2 2	2 2	2 2	2 2	0 2	2 2
10	1 0	2 2	0 2	2 0	2 2	2 2	2 2
11	1 0	2 2	2 2	1 2	2 2	0 2	2 2
12	1 1	0 0	0 0	0 0	0 0	0 0	0 0
<b>Right Rougher</b>	7	12	15	16	15	18	15
<b>Left Rougher</b>	1	8	9	7	8	6	9
<b>Same</b>	16	4	0	1	1	0	0
<b>agreement</b>	5	7	9	9	10	10	11
<b>discrepant</b>	1	1	3	2	1	2	1
<b>indecision</b>	6	4	0	1	1	0	0

**Audio Roughness Experiment  
Raw Data**

Participant	15 20	15 25	15 30	15 35	15 40	15 45
1	1 1	0 0	1 1	1 0	0 0	0 0
2	1 1	1 0	0 1	0 1	2 0	0 0
3	2 1	2 2	2 1	2 0	0 2	2 1
4	1 1	0 0	0 1	2 0	0 0	0 2
5	1 1	0 1	1 0	0 0	0 0	0 0
6	1 1	0 0	0 1	0 0	2 0	0 0
7	0 1	0 1	0 1	0 0	0 0	0 0
8	1 1	1 1	1 2	0 2	2 2	0 2
9	0 1	2 2	2 2	2 1	2 2	2 2
10	2 1	2 2	1 1	0 2	2 2	2 2
11	0 1	2 2	1 2	2 0	2 2	2 2
12	1 1	1 1	1 0	0 1	0 2	0 0
<b>Right Rougher</b>	3	9	6	14	12	14
<b>Left Rougher</b>	2	8	5	6	12	9
<b>Same</b>	19	7	13	4	0	1
<b>agreement</b>	7	9	3	4	8	9
<b>discrepant</b>	0	0	0	5	4	2
<b>indecision</b>	5	3	9	3	0	1

**Audio Roughness Experiment  
Raw Data**

Participant	20 25	20 30	20 35	20 40	20 45
1	1 1	1 1	0 1	0 0	0 0
2	1 1	1 1	1 1	1 1	0 0
3	2 1	0 2	2 1	0 0	2 0
4	0 1	1 1	0 0	0 1	0 0
5	1 1	1 1	1 1	0 0	0 1
6	1 1	1 1	0 1	1 1	0 0
7	1 1	0 2	0 1	0 1	0 0
8	1 1	1 1	2 0	2 2	2 1
9	1 1	2 2	2 2	2 2	2 2
10	2 1	1 2	2 2	2 2	2 2
11	1 1	2 2	2 2	2 2	2 2
12	1 1	0 1	0 1	0 0	0 0
<b>Right Rougher</b>	1	3	7	10	14
<b>Left Rougher</b>	2	7	8	8	8
<b>Same</b>	21	14	9	6	2
<b>agreement</b>	9	8	6	10	9
<b>discrepant</b>	0	2	1	0	1
<b>indecision</b>	3	2	5	2	2

**E4.4 Raw Data Collection – Audio Exp. - (4/6)**

**Audio Roughness Experiment  
Raw Data**

Participant	25 30	25 35	25 40	25 45	30 35	30 40	30 45
1	1 1	1 1	0 1	0 1	1 1	1 1	1 1
2	1 1						
3	2 1	0 2	1 1	1 1	1 1	2 1	2 1
4	1 1	1 1	1 1	0 2	1 1	1 2	0 0
5	1 1						
6	1 1						
7	1 1	1 0	0 0	0 0	1 2	1 1	0 0
8	1 1	1 2	1 0	0 1	1 1	1 1	1 2
9	2 1	1 2	2 2	2 2	1 1	2 2	2 2
10	1 1	1 1	1 2	2 1	1 1	1 1	1 1
11	2 1	1 2	2 2	2 2	1 1	2 2	2 2
12	1 1	0 1	0 1	2 1	1 1	1 1	1 1
<b>Right Rougher</b>	0	3	5	5	0	0	4
<b>Left Rougher</b>	3	4	5	7	1	6	6
<b>Same</b>	21	17	14	12	23	18	14
<b>agreement</b>	9	6	8	8	11	10	10
<b>discrepant</b>	0	1	0	1	0	0	0
<b>indecision</b>	3	5	4	4	1	2	2

**Audio Roughness Experiment  
Raw Data**

Participant	35 40	35 45	40 45
1	1 1	1 1	1 1
2	1 1	1 1	1 1
3	1 1	1 1	1 1
4	1 1	1 1	1 1
5	1 1	1 1	1 1
6	1 1	1 1	1 1
7	1 1	0 1	1 1
8	1 1	1 0	1 2
9	2 0	1 1	1 2
10	1 1	2 1	1 1
11	2 0	2 1	1 2
12	1 1	1 1	1 1
<b>Right Rougher</b>	2	2	0
<b>Left Rougher</b>	2	2	3
<b>Same</b>	20	20	21
<b>agreement</b>	10	8	9
<b>discrepant</b>	2	0	0
<b>indecision</b>	0	4	3

## **Appendix F: Instructions, questionnaires, and raw data for multimodal roughness experiment**

---

# Judging the Roughness of Virtual Surfaces

## Introduction

Thank you for taking part in this experiment.

If at any time you do not understand what is being asked of you, or you need to stop for any reason at all then please let the experimenter know immediately.

This is not a test of your ability, it is simply a means of collecting data on how rough virtual surfaces seem to different people. Different people might have different responses.

Please fill in the following details and then let the experimenter know you are ready to continue.

Name:	_____
Age (optional):	_____
Sex:	Male / Female
To your knowledge, do you have normal sense of touch?	Yes / No
If no, then please give details below	
_____	
_____	
To your knowledge, do you have normal sense of hearing?	Yes / No
If no, then please give details below	
_____	
_____	
Have you ever used this force feedback device (the PHANToM) before?	
	Yes / No

**F1: Consent and Introduction form**

# Judging the Roughness of Virtual Surfaces

## Instructions

Please read the following instructions very carefully.

You can ask the experimenter to clarify anything that isn't clear at any time as you are reading them.

You will also get a short training session that should make everything much clearer.

The following experiment looks at how people perceive virtual surfaces. In particular, we are looking at how *rough* we judge these different surfaces to be.

There are no right or wrong answers, only your own judgment of how rough the surface seems. Your initial feeling is most important so try not to spend too much time worrying over each trial. You may decide that one of the surfaces seems rougher or you may decide that they both feel the same in terms of roughness.

When you have made your decision you simply check the box that corresponds to your decision for that trial. Clicking the button labeled "Next" will begin the next trial and instructions will appear at the end to tell you when you have completed the experiment.

The next couple of pages explain the device you will be using to explore the surfaces and the experimental interface in more detail.

## **Exploring the Virtual Surfaces (textured patches)**

The surfaces can be explored by *dragging* the pen-like probe of the PHANToM device back and forth across the virtual surfaces. The textured area is outlined visually by a box as a guide but the actual texture on the surface cannot be seen - only felt and/or heard.

The texture (or surface) is a patch on the back wall of the screen so you might need to apply a small force forward to actually feel or hear the texture. There will always be some form of surface present. The device will provide feedback to you to let you know how rough the surface you are dragging across is.

## **Making your response**

Sometimes the surface on the left will seem roughest. Sometimes the surface on the right will seem roughest. Sometimes the surfaces will seem exactly the same in terms of their roughness. Each response is equally valid. Different people may even have different responses so don't worry about being right or wrong. Just try your best to make an accurate judgement as quickly as you can so that you are reasonably happy with your judgment.

# The PHANToM Device

The device you will be using to explore the surfaces AND make your responses is called the PHANToM.

It is similar to a mouse in that you can enter input to the computer via the device.

This device can also send feedback that you can *feel* in the form of output from the computer back to you.


The forces sent back to you simulate how it feels to interact with virtual objects by touching them.

Hold the pen-like probe on the PHANToM as you would hold a pen. You will have a chance to make sure you are comfortable with this before you are asked to do anything.

On the end of the probe, there is a silver switch that you can press just like you would press the left button on a mouse. Pressing this switch while over a target/button will select that target/button.

If you have any difficulties during the experiment then please let the experimenter know.

# Experimental Interface



(1) You can experience the texture by dragging the pen like probe of the PHANToM across the rectangular patches labeled 'left' and 'right'.

(2) Using the PHANToM probe, positioning the cursor over a button or check box and clicking the silver switch on the end of the probe will select that button.

(3) A selected checkbox will turn black.

(4) Pressing the button labeled 'next' will present the next two textures to be compared.

(5) Let the experimenter know when the experiment ends and no more textures are being presented to you.

# Additional Questions

**Please circle the appropriate response and provide as much detail as possible when a written answer is required:**

(1) On average, how easy did you find it to tell that any two textures *felt* different?

Very easy    quite easy    about average    quite hard    very hard

(2) On average, how easy did you find it to tell that any two textures *sounded* different?

Very easy    quite easy    about average    quite hard    very hard

(3) On average, did you feel that the sound of a surface or the feeling of a surface was more important in making your roughness judgement?

Sound / Feeling

Provide explanation of necessary:

---

---


---

(4) How confident were you about making your decisions?

extremely unsure    fairly unsure    indifferent    fairly sure    extremely sure

**PTO**

(5) Can you draw below what you think the textures might have looked like if you could see them?



Finally - You were asked to judge the *roughness* of virtual textures.

(6) Can you describe below (using pictures, words, or both) what 'roughness' means for you?

---

---

---

---

---

---

---

---

---

---

You can let the experimenter know that you are finished.

Thank you again for taking part.

**F3.2: Post Experimental Questionnaire (2/2)**

## Multimodal Roughness Results

Between conditions

Number of times texture judged as roughest of pair

### ROUGHER THAN

Subject	HAPTIC					
	10	15	20	25	30	35
1	0	2	4	6	5	5
2	2	3	1	4	5	6
3	0	1	4	5	7	9
4	4	4	5	4	5	5
5	0	2	5	6	5	5
6	1	3	3	5	7	7
7	0	1	2	4	5	8
8	0	2	4	6	7	7
9	0	2	4	5	8	8
10	6	2	5	2	1	2
11	3	1	3	4	5	7
12	0	4	2	6	7	7
13	0	2	4	6	7	8
14	0	0	3	4	7	8
15	2	1	2	3	5	7
16	0	5	4	4	8	9
17	9	5	4	3	1	0
18	2	3	1	4	5	5
<b>Total</b>	29	43	60	81	100	113
<b>Mean</b>	1.6	2.4	3.3	4.5	5.6	6.3
<b>%</b>	16	24	33	45	56	63

F4.1: Raw Data Collection – Multimodal Exp. - (1/3)

Subject	MM-C					
	10	15	20	25	30	35
1	1	1	1	3	6	4
2	5	2	2	4	2	4
3	1	3	2	6	8	9
4	1	2	3	6	8	7
5	1	1	1	3	6	4
6	1	2	3	4	5	3
7	0	0	2	4	5	6
8	1	3	2	6	8	9
9	1	2	5	5	7	6
10	4	4	3	5	5	0
11	0	1	3	0	6	7
12	0	3	0	3	4	3
13	0	4	7	8	8	10
14	1	1	3	5	6	9
15	1	1	3	5	6	8
16	0	4	7	8	8	10
17	7	6	2	1	1	0
18	4	2	2	4	2	4
<b>Total</b>	29	42	51	80	101	103
<b>Mean</b>	1.6	2.3	2.8	4.4	5.6	5.7
<b>%</b>	16	23	28	44	56	57

**F4.2: Raw Data Collection – Multimodal Exp. - (2/3)**

### Multimodal Roughness Results

Between conditions      Within Subjects

Number of times texture judged as roughest of pair

### ROUGHER THAN

Subject	MMI					
	10	15	20	25	30	35
1	0	0	3	2	7	5
2	4	4	2	3	1	5
3	0	2	4	5	6	10
4	1	4	7	8	9	5
5	0	0	3	2	7	4
6	0	0	1	2	3	3
7	0	1	1	4	6	5
8	0	2	4	5	6	10
9	4	2	6	4	5	6
10	2	6	3	4	2	2
11	2	1	2	6	6	5
12	3	2	0	4	1	2
13	0	1	5	7	5	9
14	0	1	3	3	5	8
15	0	1	3	3	5	8
16	0	1	5	7	5	9
17	9	5	1	1	1	1
18	4	4	2	3	1	5
<b>Total</b>	29	37	55	73	81	102
<b>Mean</b>	1.6	2.06	3.06	4.1	4.5	5.7
<b>%</b>	16	20.6	30.6	41	45	57

F4.3: Raw Data Collection – Multimodal Exp. - (3/3)

# **Appendix G: Papers and publications resulting from thesis**

---

# Putting the Feel in ‘Look and Feel’

Ian Oakley, Marilyn Rose McGee, Stephen Brewster and Philip Gray

Glasgow Interactive Systems Group

Department of Computing Science

University of Glasgow, Glasgow, G12 8QQ, UK

Email: {io, mcgeemr, stephen, pdg}@dcs.gla.ac.uk Web: www.dcs.gla.ac.uk/~stephen

## ABSTRACT

Haptic devices are now commercially available and thus touch has become a potentially realistic solution to a variety of interaction design challenges. We report on an investigation of the use of touch as a way of reducing visual overload in the conventional desktop. In a two-phase study, we investigated the use of the PHANToM haptic device as a means of interacting with a conventional graphical user interface. The first experiment compared the effects of four different haptic augmentations on usability in a simple targeting task. The second experiment involved a more ecologically-oriented searching and scrolling task. Results indicated that the haptic effects did not improve users performance in terms of task completion time. However, the number of errors made was significantly reduced. Subjective workload measures showed that participants perceived many aspects of workload as significantly less with haptics. The results are described and the implications for the use of haptics in user interface design are discussed.

## Keywords

Haptics, force feedback, multimodal interaction.

## INTRODUCTION

Desktop interfaces are becoming increasingly complex, and with this added complexity, problems are beginning to emerge. One such problem is information overload, where so much information is presented graphically that it becomes difficult to attend to all relevant parts [4]. Presenting information in other sensory modalities has the potential to lessen this problem. Attempts have been made to overcome information overload using non-speech sound during interactions such as button clicking and scrolling [3, 5] but there have been no convincing empirical attempts to reduce overload by using haptic (or force feedback) technology. This new technology allows users to *feel* their interfaces and has the potential to radically change the way we use computers in the future. We will be able to use our powerful sense of touch as an alternative mechanism to

send and receive information in computer interfaces.

Augmenting graphical user interfaces (GUIs) with haptic feedback is not a new idea. In 1994 Akamatsu and Sate [1] developed a haptic mouse with the ability to produce what they termed ‘tactile feedback’, the ability to vibrate a user’s fingertip, and ‘force feedback’, a simple software controllable friction effect. Using this device they showed significantly decreased completion times in a targeting task offset by slightly increased error rates. Engel *et al.* [7] found improved speed and error rates in a generalised targeting task using a modified trackball with directional two degrees of freedom force feedback.

The devices used in these early studies have now been superseded. More advanced devices such as the Pantograph (Haptic Technologies Inc.), the FEELit mouse (Immersion Corp.), and the PHANToM (SensAble Technologies Inc.) have been developed. These devices have all been used to augment desktop interfaces. Ramstein *et al.* [11] used the Pantograph to demonstrate performance increases in desktop interactions but provided little empirical evidence to support their claims. The FEELit mouse is a commercial product that offers users a haptically-enhanced desktop but there has been little evaluation of this device published [14]. Finally, the PHANToM has been used to create a haptically enhanced XWindows desktop [10]. No formal evaluation of this enhancement can be found in the literature.

The pace of technological advancement in this field is rapid, both in terms of the hardware produced and the software developed. Current projects to ‘haptify’ the desktop are not constrained to use the haptic effects described by Akamatsu and Engel. However, as technology has advanced there has been no corresponding progress in its evaluation. This disparity has led to a situation where there are no formal guidelines regarding what feedback is appropriate in different situations. This, along with evidence that shows arbitrary combinations of information presented to different senses is ineffective [12, 13], leads to the conclusion that empirical evaluation of modern haptic augmentations of the desktop is urgently required if much time and effort is not to be wasted. We might even end up with haptically-enhanced interfaces that are in fact harder to use than standard ones and haptics may become just a

gimmick, rather than the key improvement in interaction technology that we believe it to be.

### Haptic Terminology

Many different terms with many different definitions are used throughout the literature to describe haptic interaction. One reason for this is that the area is in its infancy. To rectify this problem we propose a set of haptic definitions that should prove useful for further research in this area.

The word ‘haptic’ has grown in popularity with the advent of touch in computing. We define the human haptic system to consist of the entire sensory, motor and cognitive components of the body-brain system. It is therefore closest to our understood meaning of proprioceptive (see Table 1). We define haptics therefore to be anything relating to the sense of touch. Under this umbrella term, however, fall several significant distinctions. Most important of these is the division between cutaneous and kinesthetic information (see Table 1). There is some overlap between these two categories; critically both can convey the sensation of contact with an object. The distinction becomes important however when we attempt to describe the emerging technology. In brief, a haptic device provides position input like a mouse but also stimulates the sense of touch by applying output to the user in the form of forces. Tactile devices affect the skin surface by stretching it or pulling it, for example. Force feedback devices affect the finger, hand, or body position and movement. Using these definitions (summarised in Table 1), devices can be categorised and understood by the sensory system that they primarily affect.

Term	Definition
Haptic	Relating to the sense of touch.
Proprioceptive	Relating to sensory information about the state of the body (including cutaneous, kinesthetic, and vestibular sensations).
Vestibular	Pertaining to the perception of head position, acceleration, and deceleration.
Kinesthetic	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 kinesthetic system.

Table 1: Definitions of Terminology.

### EXPERIMENTAL OVERVIEW

This paper describes two experiments that empirically test the use of haptics to augment targeting in the standard GUI. It is force feedback, and not tactile feedback that is evaluated in this work. Experiment 1 compared user performance with haptically-enhanced buttons using four

different haptic effects in a simple targeting task. Experiment 2 involved a more ecologically oriented task in which participants searched for and selected targets using haptic scrolling. We hypothesise that in both experiments haptics will have a positive effect on performance.

Neither of the experiments described is concerned with the influence of haptic distracters; both investigate haptic augmentation when there is guaranteed to be a clear path to target. The decision to adopt this approach reflects the preliminary nature of empirical research in this field.

### Device and Software

The device used in both experiments is the PHANToM 1.0 (see Figure 1). It is a force feedback device (provides kinesthetic information as defined in Table 1) which, in the experiments, acted as a cursor control device in place of the traditional mouse.

Optical sensors detect changes in the configuration of the PHANToM. The device uses mechanical actuators to apply forces back to the user calculated from this positional information and the stored algorithmic models of the objects with which the user is currently interacting. To operate the device users hold a stylus.

The graphical interface was generated using standard (MFC) widgets and these performed in exactly the same way as standard widgets. The workspace was a box 160 mm wide x 160 mm high x 2 mm deep. The haptic effects were present only on the back wall of the workspace.

### Haptic Effects

Four haptic effects were used in the experiments. These built on and added to the effects used in previous studies. The effects were all aimed at improving targeting and reducing problems of mis-hitting or slipping off interface widgets. The effects used were:

*Texture:* Texturing a button in a texture-less, flat workspace is a potential way of haptically signifying that the cursor is positioned over an interesting object. The texture implemented here formed a set of concentric circles 7.5 mm apart and centred around the middle of the target. The



Figure 1: The Phantom 3D force feedback device from SensAble Technologies. The stylus shown has a button that can be used for performing the mouse clicks in the experiments reported.

texture was created by vector rotation (force perturbation) [15] and the maximum rotation applied was  $12^\circ$ . A visual representation is shown in Figure 2. This texture pattern was used because it was felt that it would maximise the possibility that users would encounter ridges irrespective of the direction they began from or travelled in.



Figure 2: Diagram of the geometry of haptic texture effect.

**Friction:** The friction effect damped a user's velocity. Haptically-enhanced interfaces that use a friction effect are common in previous literature. This is partly because they can be produced with simple hardware – for instance with an electromagnet placed in the base of a mouse [1, 2] – and partly because it seems advantageous to provide feedback that causes a user to stop when over an interesting target. The friction effect used here was realistically modelled with both a static and a dynamic component. The static component restricted users to a point until they attained an escape velocity. The dynamic component attempted to slow them whilst they were in motion.



Figure 3: Diagram of the geometry of haptic recess effect.

**Recess:** The recess effect was a hole in the back of the workspace, with a depth of 2 mm and edges sloped at  $45^\circ$ . This effect also features strongly in previous literature [10, 11]. A diagram of the geometry of a recess is presented in Figure 3. A recess could potentially provide useful feedback by the simple fact that to leave it, the wall at the edge must be climbed. This may make it harder to accidentally slip-off a button (a problem noted by Brewster *et al.* [5]).

**Gravity Well:** The gravity well was a 'snap-to' effect. When users moved over a button a constant force of 0.5 N was applied that pushed them towards the button's centre. This force tapered off around the very centre so that the user could rest in the centre. The gravity well promised the same benefits as the recess – a reduction in errors through the simple mechanism of preventing a user from accidentally slipping off a button.

#### General Measures Used in the Experiments

In order to get a full range of quantitative and qualitative results, time, error rates, and subjective workload measures were used in both of the experiments. The subjective

workload measurement was a modified version of the NASA Task Load Index (TLX) [8]. NASA reduce workload to six factors: mental demand, physical demand, time pressure, effort expended, performance level achieved, and frustration experienced. We added a seventh factor: fatigue. One potential problem with force feedback devices is the physical strain placed on the user. By adding this factor it would be possible to find out if haptic effects caused any additional perceived fatigue. Participants filled-in workload charts after each condition in both experiments.

#### EXPERIMENT 1

In the first experiment the haptic effects were compared to investigate which was the most effective. To do this we added each of the haptic effects to standard graphical buttons. This allowed us to investigate targeting (moving the cursor to the button) and mis-hitting errors (slipping-off the button when trying to press it).

#### Hypotheses

Experiment 1 was an exploratory experiment – we wanted to investigate the differences between the different haptic effects and a control condition. Therefore, the experimental hypotheses were that differences would occur in task completion time, number of errors and in the subjective data gathered. We predicted that the gravity well and recess would provide the largest reduction in errors, time and workload as they provided feedback that was highly appropriate to a simple targeting task.

#### Participants

There were sixteen participants. Four were female and twelve were male. All were between the ages of eighteen and thirty. Most were computing students from the University of Glasgow. All were regular and fluent computer users. Three users were left-handed and one was dyslexic. None had anything more than trivial previous exposure to the PHANToM.

#### Design

The experiment followed a within-subjects repeated-measures design. Each participant underwent each of the four haptic conditions, each encompassing one of the effects described above, and a control condition. The control condition used the PHANToM device but no haptic effects were applied – in essence the device worked like a normal mouse. The order of the presentation of the conditions was counterbalanced to evenly distribute the effects of practice and fatigue. Participants were randomly allocated to conditions. Training was given in each condition in a session immediately prior to the experiment. Each condition in the training session constituted 60 button presses and in the experimental session 120 presses. The experiment's duration was typically 45 minutes.

#### Task

A simple button pressing/targeting task was used. This task was chosen because it featured prominently in the previous literature [1, 2, 7] and also because it is a very elementary

operation – it is both simple to perform and also perhaps the most fundamental cursor operation.

Two factors were engineered into the task to make it more suitable for haptic augmentation. Firstly, it was felt that participants 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. They concentrate on some central task and interact with graphical widgets in the periphery of their attention [4]. Secondly, in this atmosphere of visual distraction, we assert that the haptic feedback will only really prove useful if the task encompasses some repetitive motion. Without such motions 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. Repetitive motions are also common in desktop interactions (moving to menu bars, clicking buttons, etc.).




Figure 4: The interface used in Experiment 1.

To encompass these two factors two windows were placed on the screen at all times (see Figure 4). One, the instruction window, occupied the left-hand side of the screen and contained instructions as to the next target to seek. The other, larger, window occupied the centre and right-hand side of the screen and contained the targets in the form of five buttons. One button was always positioned in the centre, the other four were positioned one in each quadrant of the window, on the diagonals of the window. The position along the diagonals was changed in the course of the experiment, but each button remained in a single quadrant of the window throughout. This meant that each button remained in the same direction relative to the centre of the window at all times. The buttons moved along the diagonals to prevent users employing a purely mechanical repetition. To ensure users moved along only a few trajectories to reach each of the buttons, every second button press was the centre button. The buttons were labelled in accordance with their positions on screen, for instance “top right” or “bottom left”. The instruction

window indicated the next target button, on successfully pressing the named button, a new name was presented.

### Measures

Data were gathered from all button presses in the experiment. The performance measures were (a) mean time per trial (secs.), (b) mean number of errors, and (c) subjective workload ratings. Times were measured at four stages: time to find target button; time to move onto target button; time to press target button; and time to move off target button. Errors were measured as when a participant moved over a button but failed to press it. There were two categories: the first was where the user simply slid over the button, arguably as a part of the normal targeting process. The second, more serious error is known as a ‘slip-off’ [4]. This occurs when a user presses the mouse down over a button but moves off it before releasing the mouse, thus not selecting it. The feedback for this is the same as for a successful mouse click. An error of this type can go unnoticed for some time and cause considerable confusion.

### Results from Experiment 1

The error data are presented in Figures 5 and 6. Results were analysed using ANOVA tests. Significant effects were found when comparing the mean scores for each haptic effect for both slide over ( $F_{4,15} = 48.487$ ,  $p < 0.001$ ) and slip off ( $F_{4,15} = 20.81$ ,  $p < 0.001$ ) errors. Order effects for both slide over ( $F_{4,15} = 0.152$ ,  $p = 0.961$ ) and slip off ( $F_{4,15} = 0.123$ ,  $p = 0.974$ ) errors were not found.

	Gravity	Recess	Friction	Texture	Control
Gravity	-----	Not sig	$p < 0.001$	$p < 0.001$	$p < 0.002$
Recess	-----	-----	$p < 0.01$	$p < 0.003$	Not sig
Friction	-----	-----	-----	$p < 0.04$	Not sig
Texture	-----	-----	-----	-----	$p < 0.016$

Table 2: Analysis of slip-off errors in Experiment 1.

	Gravity	Recess	Friction	Texture	Control
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 3: Analysis of slide-over errors in Experiment 1.

A summary of the results revealed by *post-hoc* analysis of the means (using Bonferroni confidence interval adjustments) is shown in Tables 2 and 3. The most dramatic results were that participants in the gravity condition made significantly fewer errors of both sorts than in the control and that the converse was true of the texture condition – it caused significantly more errors than the control.

Analysis of the temporal data was less conclusive; the total time taken to complete a trial was strongly biased by the number of errors made in each condition. It was felt that this invalidated it as a measure – it would merely be a reflection of the number of errors in each condition.




Figure 5: Slide over errors in Experiment 1.




Figure 6: Slip-off errors in Experiment 1.




Figure 7: Total time on button in Experiment 1.

Instead, the total time on a button during a successful trial was analysed (see Figure 7). An ANOVA revealed significant differences between effects. Subsequent pairwise comparisons (using Bonferroni adjustments) revealed that gravity was significantly slower than recess ( $p < 0.05$ ). It is also worth noting that the difference between the best and worst performing effects was only 42 ms, a very short time. No order effects were found in this temporal analysis ( $F_{4,15} = 0.913$ ,  $p = 0.462$ ).

To validate analysing time and errors separately we ran a Pearson correlation. The timing results did not correlate with the slide over ( $r = 0.0$ ,  $p < 1.0$ ) or slip off ( $r = 0.019$ ,




Figure 8: Workload results from Experiment 1.

$p < 0.976$ ) errors. The two error results strongly correlated with one another ( $r = 0.938$ ,  $p < 0.018$ ).

Figure 8 shows the TLX workload scores (scored out of 20). The texture condition was significantly worse than the control across the whole board of measures. The gravity condition consistently reduced workload and, in particular, achieved a significantly better score than the control in the performance level achieved category ( $p < 0.018$ ).

## EXPERIMENT 2

This experiment simulated a more realistic task where reading was accompanied by scrolling through a document, selecting from the document, and returning to the scroll bar whilst still visually attending to the material being read. When users are required to scroll through a document it is the material in it that is of interest and not the scroll bar. Users want to concentrate on reading the material but often find themselves forced to move their visual attention to the scroll bar to ensure that the cursor is positioned appropriately to operate it. The time taken to make these frequent shifts in visual attention, and the frustration experienced by the need to do so, reduce the usability of the scroll bar. Problems associated with scrolling have been addressed previously [e.g. 4, 16]. Reducing these problems using force feedback technology has not yet been empirically evaluated.

## Hypotheses

It was hypothesised that when the scroll bar was haptically-enhanced, the participants would (a) take significantly less time to complete the task; (b) move on and off the scroll bar significantly less; and (c) perceive the workload during the task as significantly less.

## Participants

Twenty new participants were used: one was female and the remaining nineteen male. All were between the ages of seventeen and twenty-seven. Most participants were first-

year computing science students from the University of Glasgow. All were regular and fluent computer users. All users were right-handed. Participants had nothing more than trivial previous exposure to the PHANToM.

### Design

The experiment again used a within-subjects repeated-measures design. Each participant underwent both a visual-only condition (visual) and a visual and haptic condition (haptic). The visual condition used a standard graphical scroll bar only. In the haptic condition, this scroll bar was overlaid with haptic effects (recess and gravity well were chosen as these were the most effective in Experiment 1). The up and down arrow buttons used gravity wells. These acted as a haptic indication that the user was in the appropriate place to press the button successfully. The rest of the scrolling area used a recess effect that allowed the user to 'fall into' the slider area. Therefore, the haptic feedback allowed the user to reserve his/her visual attention for the primary task, as being over the widget was indicated through touch. The order of the presentation of the conditions was counterbalanced to evenly distribute the effects of practice and fatigue. Training was given to each participant in each condition prior to the experiment.




Figure 9: The interface used in Experiment 2. The top left window is the instruction window, the bottom left is the target window, the large window to the right is the data window and in the centre is the send button.

### Procedure

Figure 9 shows the interface to the task. Participants had to read a four-digit numerical code from the instruction window. They then had to scroll vertically through a large file of codes (presented in the data window) to find the target code, highlight the code (either by double clicking on it or dragging across it), and press a button to send this code to the target window. The widgets operated as in standard desktop applications. The data window contained the same list of 2000 randomly generated but numerically ordered codes in each condition. Forty codes had to be entered in each condition. The list was formatted such that there were three columns of codes, simulating a standard document read from left to right and from top to bottom. The highlight operation was included to force the user off the scroll bar. This ensured repeated targeting of the scroll bar. The experiment's duration was typically 40 minutes.

### Measures

The performance measures were (a) mean time per trial (secs.), (b) mean number of movements on/off scroll bar (including all required movements), and (c) workload ratings. Time was measured from when the user activated the send button at the end of the previous trial until the send button was activated at the end of the current trial. Subjective ratings were collected as before.

### Results from Experiment 2

**Timing results:** Table 4 shows the timing and movement on/off scroll bar results. Paired T-tests established that haptic feedback did not significantly reduce the average trial time as predicted ( $T_{19} = 0.46, p < 0.32$ ).

Mean Trial Time (secs.)		No. times on/off scroll bar	
Visual	Haptic	Visual	Haptic
11.7251	11.9668	107	97
SD=2.77	SD=2.84	SD=25	SD=22

Table 4: Timing and movement results from Experiment 2.

**Movement on/off scroll bar:** Paired T-tests showed that participants in the haptic condition moved on and off the scroll bar area significantly less than in the visual condition ( $T_{19} = 2.37, p < 0.05$ ).

**Workload Results:** Figure 10 shows the workload scores. Paired T-tests were carried out on the visual versus haptic conditions for each of the categories. Mental demand was not significantly less in the haptic condition as expected. Both the effort and frustration ratings were significantly reduced in the haptic condition (Effort:  $T_{19} = 2.80, p < 0.01$ , Frustration:  $T_{19} = 2.04, p < 0.05$ ). There was no significant difference in fatigue experienced. The hypothesis that the haptic condition would reduce workload is therefore confirmed in part.

### GENERAL DISCUSSION

The timing results from the two studies indicate that the haptic effects added to the buttons and scroll bar did not




Figure 10: Workload results from Experiment 2.

reduce the time taken for either task, as hypothesised. There were also no real differences between the effects – only 42 ms between the best and worst effects (recess and gravity) in Experiment 1. The explicit separation of the error data from the timing data is no doubt a contributing factor to the lack of temporal variations across conditions. However, we suggest that one potential reason for the lack of time reduction is that, in all of the effects used, participants had to exert more force to overcome the haptic effects. In the control condition they could just slide over the interface with no obstacles, in the haptic conditions they had to climb out of recesses, overcome gravity forces applied, etc. For participants to produce the forces required to do this could have taken them more time.

Further work is needed on the haptic effects themselves and the types of desktop tasks that would benefit most from them. It may have been that the haptic effects chosen were inappropriate either for reducing time or for the tasks chosen for these experiments. Other previous work has claimed a significant reduction in performance times [10, 14]. The present work suggests that things are not so clear-cut and care must be taken when using haptics to try to reduce performance times.

The error results were more conclusive. Experiment 1 showed a significant reduction in the number of errors produced across the different haptic conditions (where gravity and recess caused the fewest errors and texture the most). Gravity and recess were the most effective for targeting tasks (which are important for using many standard GUI widgets, for example hitting a button, selecting a menu item or dragging the scrollbar thumb) in the sense that they made it very hard to slip off a target once on it; participants could not just knock the pointer off the target, they had to make an explicit movement to leave. Texture only indicated that the cursor was over a target, and did not constrain users to the target, which was one of the reasons it was less effective in this case. Texture also had the problem that it could potentially perturb users' movements, making it hard for them to stay on target. This resulted from the kinesthetic force feedback device used here. We use cutaneous stimulation to feel much of the richness of fine-grained texture in the real world [9]. A kinesthetic device can only simulate gross textures, requiring larger forces, which then make it harder for users to move precisely. Texture is much more suitable to production by tactile devices such as the *Tractile* from Campbell *et al.* [6]. The PHANToM, on the other hand, is very effective at simulating gravity and recess effects as these require movement and so are kinesthetic tasks. There are no devices, as yet, which combine both tactile and kinesthetic force feedback.

Haptic devices are now reaching the desktop. For example, the FEELit Mouse [14] adds low cost haptic effects to the standard graphical interface. Our results show that interface designers must be aware of the facilities of the devices they

are using in order to generate haptic effects that will improve usability. This might seem obvious, but this area is in its infancy and new devices are appearing all the time, each having different functionality to the last.

The movement results from Experiment 2 showed a significant reduction in the number of times a participant moved on/off the scroll bar in the haptic condition. This showed that the haptic recess aided participants in remaining on target, demonstrating that haptics can provide a significant practical benefit for interaction. The haptic groove placed over the scroll bar allowed users to scroll up and down without slipping off. They could do this without looking at the bar as once the cursor was in the groove it would stay there. To move out of the recess they had to lift off the scroll bar and it was difficult to do this by mistake as it required a conscious effort.

The subjective workload measures taken across both experiments are important. Papers concerning other haptically-enhanced desktops have not presented any such data. In developing multimodal interfaces (ones that use multiple sensory modalities) it is very important to consider what effects they have on users' workload. Users may perform tasks well and quickly and yet find them frustrating and requiring more effort to complete than they would expect. This dissociation between behavioral measures and subjective experience has been addressed in studies of workload. Hart and Wickens [8] suggest that cognitive resources are required for a task and there is a finite amount of these. As a task becomes more difficult, the same level of performance can only be achieved by the investment of more resources. Just measuring time or error rates does not give the whole picture of the usability of a haptic device. Workload is particularly important in this area as we know little yet of the effects on cognitive/attentional resources of using such devices.

Experiment 1 showed that the different effects had markedly different levels of workload. Gravity well and recess came out best, indicating that they were effective at reducing error rates and decreasing workload. This suggests that they are very robust and can be successfully used in haptic interfaces of the type described here. Texture came out the worst in terms of workload, suggesting that, in general, it is hard to do effectively with the device used here. Experiment 2 showed the effect of haptics in a more realistic situation. In this case there was a significant reduction in effort and frustration – the fact that it was easy to stay on the scroll bar due to the recess effect made the task much less effortful (the reduction in the number of movements on/off the scroll bar confirms this). We had expected that this might also lead to reductions in other categories (e.g. mental demand) but these showed no significant reductions. This suggests that we need further studies of workload to learn more about the affect of haptics in desktop interactions.

One other area that we investigated was fatigue. Using a device that requires the user to apply force could cause fatigue. It is important to investigate this if force feedback devices are to be used in desktop situations (where people might use the interfaces for long periods of time). Results from Experiment 1 showed that gravity and recess effects did not cause any more fatigue than the control condition. On the other hand, texture caused significantly more fatigue than the control. This is likely to be for the reasons as discussed above – to simulate texture with a kinesthetic device required larger forces to be applied and these, in turn, required the users to exert larger forces to overcome them. Experiment 2 again showed no increase in fatigue with the use of gravity well and recess effects. This research shows that appropriate haptic effects used correctly may have no impact on fatigue, but used incorrectly may significantly increase it. This is only a first step in investigating this problem and further work is needed to ensure that we can design haptic interfaces to avoid fatigue

### CONCLUSIONS

Our research has shown that haptics may have some benefits in graphical user interfaces. Reductions in the number of errors made and subjective workload experienced can be gained. We have also shown that the haptic effects used must be matched to the capabilities of the device – trying to simulate effects not supported by the device in use can have serious negative effects on all aspects of usability. As technology progresses it is easy to focus on what benefits new equipment may afford whilst forgetting to measure the benefits actually produced. Recent work on haptically-enhanced desktops has been firmly orientated towards implementation and the experiments described here begin to redress the balance. Our empirical findings provide a firm foundation for future researchers to build on and some basic principles for developers to use.

### ACKNOWLEDGEMENTS

This research was supported under EPSRC project GR/L79212 and EPSRC studentship 98700418. Thanks must also go to the SHEFC REVELATION Project, SensAble Technologies and Virtual Presence Ltd.

### REFERENCES

1. Akamatsu, M. & Sate, S. (1994). A multi-modal mouse with tactile and force feedback. *International Journal of Human-Computer Studies* (40), 443-453.
2. Akamatsu, M. & MacKenzie, S. (1996). Movement Characteristics using a Mouse with Tactile and Force Feedback. *International Journal of Human Computer Studies* (45), 483-493.
3. Beaudouin-Lafon, M. & Conversy, S. (1996). Auditory illusions for audio feedback. In *ACM CHI'96 Conference Companion* (Vancouver, Canada) ACM Press, Addison-Wesley, 299-300.

4. Brewster, S.A. (1997). Using Non-Speech Sound to Overcome Information Overload. *Displays*, 17,179-189.
5. Brewster, S.A. (1998). The design of sonically-enhanced widgets. *Interacting with Computers*, 11(2), 211-235.
6. Campbell, C.S., Zhai, S., May, K.W. & Maglo, P. (1999). What you feel must be what you see: adding tactile feedback to the trackpoint. In *IFIP Interact'99*, (Edinburgh, UK), IOS Press, 383-390.
7. Engel, F.L., Goossens, P. & Haakma, R. (1994). Improved Efficiency through I- and E-Feedback: A Trackball with Contextual Force Feedback. *International Journal of Human-Computer Studies*, 41(6), 949-974.
8. Hart, S.G. & Wickens, C. (1990). Workload assessment and prediction. MANPRINT, an approach to systems integration, 257-296, Van Nostrand Reinhold.
9. Lederman, S.J., Summers, C. & Klatzky, R.L. (1996). Cognitive salience of haptic object properties: Role of modality-encoding bias. *Perception*, 25, 983-998.
10. Miller, T. & Zeleznik, R. (1998). An Insidious Haptic Invasion: Adding Force Feedback to the X Desktop. In *ACM UIST'98*, (San Francisco, CA) ACM Press, 59-64.
11. Ramstein, C. (1995). A Multimodal User Interface System with Force Feedback and Physical Models. In *IFIP Interact'95*, (Lillehammer, Norway) Chapman & Hall, 157-162.
12. Ramstein, C. & Hayward, V. (1994). The Pantograph: A Large Workspace Haptic Device for Multi-Modal Human-Computer Interaction. In *Summary Proceedings of ACM CHI'94*, (Boston, MA) ACM Press, Addison-Wesley, 57-58.
13. Ramstein, C., Martial, O., Dufresne, A., Carignan, M., Chassé, P. & Mabileau, P. (1996). Touching and hearing GUIs: Design issues for the PC-access system. In *Proceedings of ACM Assets'96*, (Vancouver, Canada), ACM Press, 2-10.
14. Rosenberg, L.B. (1997). FEELit mouse: Adding a realistic sense of FEEL to the computing experience. <http://www.force-feedback.com/feelit>
15. Srinivasan, M.A. & Basdogan, C. (1997). Haptics in Virtual Environments: Taxonomy, Research Status, and Challenges. *Computers and Graphics*, 21(4), 393-404.
16. Zhai, S., Smith, B.A., & Selker, T. (1997). Improving Browsing Performance: A study of Four Input Devices for Scrolling and Pointing Tasks. In *Proceedings of Interact 97*, (Sydney, Aus) Chapman & Hall, 286-293.

# The Effective Combination of Haptic and Auditory Textural Information

Marilyn Rose McGee, Phil Gray, Stephen Brewster

Multimodal Interaction Group  
Glasgow Interactive Systems Group  
Department of Computing Science,  
University of Glasgow

Email (first author) - mcgeemr@dcs.gla.ac.uk  
Tel - (0141) 330 3541  
Internet - <http://www.dcs.gla.ac.uk/~mcgeemr>

**Abstract.** With the increasing availability and quality of auditory and haptic means of interaction, it is not unusual to incorporate many modalities in interfaces rather than the purely visual. The user can be powerfully affected however when information presented in different modalities are combined to become multimodal. Providing interface designers with the means to implement haptic-audio interfaces might result in adverse effects to interaction unless they are also equipped with structured knowledge on how to select effective combinations of such information. This work introduces 'Integration of Information' as one important dimension of haptic-audio interaction and explores its effects in the context of multimodal texture perception. The range and resolution of available textures through force feedback interaction is a design consideration that might benefit from the addition of audio. This work looks at the effect of combining auditory and haptic textures on people's judgment of the roughness of a virtual surface. The combined haptic-audio percepts will vary in terms of how congruent they are in the information they convey regarding the frequency of bumps or ridges on the virtual surface. Three levels of integration (*conflicting*, *redundant*, or *complementary*) are described and their possible implications discussed in terms of enhancing texture perception with force-feedback devices.

**Keywords.** Haptic, audio, force-feedback, texture perception, multimodal information processing, intersensory integration.

## Introduction

### Motivations

Multimodal Interfaces involve the use of multiple human modalities in the interaction (input, output, or both) between the human user and the computer. Haptic-audio interfaces therefore involve the use of both haptic and audio means of interaction (see Table 1. for definitions). In particular, the term haptic-audio interfaces is used here to refer to the communication of certain information to the user through an interface using a combined haptic and audio representation of this information rather than a single modality representation. The advances in both haptic and audio technology have resulted in such haptic-audio interfaces becoming increasingly realistic to

implement in a wide range of applications yet we have little organized knowledge on how best to design them. This work contributes to a body of knowledge on how to effectively combine haptic and auditory information.

The way we integrate information from different sensory modalities is complex (Wickens *et al*, 1983) and can seriously contribute to the quality of interaction in multimodal interfaces. The term 'integration of information' is used to refer to the information processing involved in combining two (or more) different modalities presented together to convey the same piece of information. Two modalities can be combined and the resulting multimodal percept may be a weaker, stronger, or altogether different percept. The effects of combining haptic and audio information must therefore be systematically explored to realize the potential of haptic-audio interfaces as well as to avoid creating interfaces that afford poor interaction.

There are specific interaction issues emerging from the increasing use of haptic interfaces, which could potentially be solved using careful addition of audio. One such interaction issue is that of haptically representing texture. In particular, force feedback devices are being used to convey texture by perturbing the user's hand or finger movements kinesthetically rather than cutaneously as with tactile devices (e.g. Lederman, 1999; West and Cutkosky, 1997). This often relies on much larger forces than those typically experienced on the skin during real texture perception (Katz, 1989). We have found in our previous work that such gross textures can perturb the users' movements so much that the ability to stay on the textured surface is adversely affected (Oakley *et al* 2000).

## **Goals**

This work discusses and empirically evaluates the dimension of 'Integration of Information' in the specific context of haptic-audio texture perception. The goals of the ongoing work are to: (a) explore the effects of combining haptic and audio information at varying levels of integration and (b) determine the potential benefits of using haptic-audio percepts of texture to overcome the limitations of presenting texture through force feedback alone.

## **Previous Research**

Within multimodal research, there have been distinct areas of specialized interest emerging. It has become clear from the research that exploring how our sense modalities behave in interaction should allow us to choose appropriate combinations of modalities according to the devices being used, the population of users, the environment, and the nature of the task.

Much of the work to date has focused on coordinating multimodal input for example (e.g. Oviatt, 1997), or the coordination of multimodal output for a specialized population such as visually impaired or physically disabled users (e.g. Mynatt, 1997; Stevens *et al*, 1997). Less work exists on the systematic study of how the combination of multimodal output of information could be better designed to coincide more closely with human information processing capabilities during multimodal interaction. In addition, little work exists on matching these information-processing capabilities to the nature of the interaction device(s) being used.

Visual displays have dominated interface research in the past but more recently auditory displays have been developed and tested (e.g. Brewster, 1997; Mynatt, 1997). With the lack of touch in interfaces now being strongly challenged, haptic technologies have also emerged at a rapid rate (Srinivasan, 1997). With the visual, auditory, and haptic channels (see Table 1. for definitions) all now technically available, multimodal interfaces can reach wider populations, increase the potential realism of displays, and generally increase the quantity and quality of information we can convey through the interface.

In human sensing and manipulation of everyday objects, the perception of surface texture is fundamental to accurate identification of an object (Katz, 1989). In a virtual world also, haptic texture information can both increase the sense of realism of an object as well as convey informational content regarding what the object is, where it is, what it is for and so on (Jansson *et al*, 1998).

Textures might be used in human and veterinary virtual medicine to assist in diagnosis of certain conditions. The texture of a tissue might indicate how well scarred tissue is healing for example. Using texture in the visualization of data could allow areas of interest to be 'textured' in the same way as colours are used in graphical visualization. Different textures could indicate different keys on a graph or chart for example. Being able to discriminate between various virtual textures in the textile industry might also prove beneficial. With an increasing number of customers shopping online for a variety of products, being able to convey different textures of objects will become crucial. For a variety of reasons it is desirable to be able to represent textures as effectively as possible in virtual environments.

There has been considerable previous work investigating the perceptual aspects of real surface textures. Lederman *et al*. (1974) suggest that texture perception is mediated by force cues created by spatial geometry of the surface. It is also possible that surface texture perception uses vibratory cues generated by the repeated and regular stimulation of mechanoreceptive afferents as the finger is moved across a surface (Katz, 1989). In fact, it is possible that both kinds of cues are involved, depending on the task to be executed (Weisenberger and Krier, 1997). Far less is known about the perceptual response to virtual surfaces. The physical properties of textures are very complex and are proving difficult to reproduce for virtual textures. For example, is a rough surface characterized

by irregular or regular surface elements? What effect does inter-element spacing have on perceived roughness? Representing texture with force feedback devices in particular has proved problematic.

Force feedback devices detect changes in the device's configuration and then use mechanical actuators to apply appropriately calculated forces back to the user. Importantly, the interaction relies on kinesthetic information being conveyed to the user rather than cutaneous information (see table 1). These devices often simulate textures with larger forces than those experienced in real texture perception. In our previous work for example we found that the gross textures implemented perturbed users' movements making it hard for them to stay on a desktop target (Oakley *et al.*, 2000).

Haptic	Relating to the sense of touch.
Kinesthetic	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 kinesthetic system.

**Table 1:** Definitions (Oakley, McGee, Brewster and Gray, CHI 2000)

It could perhaps be argued that texture is more suitable to production by tactile devices. Despite the early perceptual and physiological arguments for a spatial code to texture, three-dimensional force feedback interfaces are able to simulate surface texture (Weisenberger and Krier, 1997). It is the degree of fidelity and realism achievable with such devices that is of primary interest. The interaction issue then is how to overcome any limitations of using force feedback devices alone to represent texture.

The display of a convincing haptic percept such as texture should not necessarily be limited to the haptic modalities. Audio and visual cues can be associated with a haptic display to contribute to the realism or informational content of the display (Rosenberg, 1994). The current work investigates the conditions under which audio cues do and do not enhance force feedback based texture perception.

## Current Work

It would be beneficial to know the extent to which we can affect peoples' perception by coupling auditory and haptic percepts in a systematic way. In doing so we can establish ways in which to manipulate what the user will perceive at the interface. In particular, we could use this information to overcome limitations of a device. For instance, the addition of audio information to force feedback virtual surfaces might increase the range and/or resolution of textures available to the designer. Likewise, this information could be used to avoid coupling percepts that result in perceptual or cognitive conflict and which in turn might adversely affect the processing of that information.

In the current work, haptic and auditory textures will be rated by a group of participants to establish how rough each stimuli is in terms of each of the other stimuli. This will result in a set of haptic and audio textures identifiable along the dimension of increasing roughness. These haptic and audio stimuli can then be combined to produce multimodal haptic-audio roughness percepts in the main study. The combined textures will be either congruent or incongruent in terms of the information each modality conveys regarding the number of ridges/bumps on the virtual surface. Resulting multimodal percepts might provide *redundant*, *complementary*, or *conflicting* haptic-audio information. The effects of the different levels of congruency and resulting levels of integration of the information will be discussed.

## Device

The PHANToM 1.0 force feedback device by SensAble Technologies will be used to create the haptic virtual surfaces (see Fig. 1). Force feedback devices have optical sensors that detect changes in the device's configuration. The device then uses mechanical actuators to apply forces back to the user calculated from the positional information and the stored algorithmic models of the objects with which the user is interacting. The interaction relies on kinesthetic information being conveyed to the user rather than cutaneous information (see table 1).



**Fig 1:** The Phantom 3D force feedback device from SensAble Technologies.

Subjects interact with the device by holding a pen-like stylus attached to a passive gimbal. The user is instructed to scrape the probe of the PHANToM back and forth across the textured area to produce the haptic and/or auditory feedback regarding the roughness of the surface. The stylus switch on the probe of the PHANToM is used to select any response a participant has to make.

### **Haptic and Auditory Textures**

Neither haptic nor auditory textures are designed to necessarily model physically accurate or optimum representations of a rough surface. Rather, they are designed to give feedback approximate to that obtained when real textures are explored. In this way, the actual effects of experiencing such feedback multimodally as opposed to unimodally can be explored.

The haptic textures are generated as sinusoidal gratings on a rectangular patch on the back wall of the workspace. Forces are modeled as a point contact in the z-direction. The resulting profile depends on the amplitude and frequency of the 'wave'. The haptic textures will have a fixed amplitude of 0.5mm and frequency (cycles per fixed length of surface) can have one of 6 values - 10, 15, 20, 25, 30, or 35 cycles.

The auditory textures will consist of a sound played to indicate contact with a ridge/bump on the haptic virtual surface. The number of contact sounds can be matched to the number of ridges/bumps experienced haptically (congruent) or provide more or less contact sounds than there are haptic bumps/ridges (incongruent). The exact effect of this congruency/incongruency on the perceived level of roughness of a percept is the subject of investigation.

### **Manipulating Congruency**

Congruency/Incongruency are determined by the information provided by each modality relating to the number of bumps/ridges encountered on a virtual surface. If the number of contact sounds matches the number of haptic bumps/ridges then they are defined as congruent. Incongruency occurs when the number contact sounds does not match the number of haptic bumps/ridges.

Incongruency however has directionality. Audio information might indicate more or less bumps/ridges than the haptic information. In this case, the incongruency could act to move the level of perceived roughness of a surface up or down the roughness dimension. The direction of incongruency will depend on how frequency of the haptic bumps/ridges, and frequency of contact sounds, unimodally map to level of perceived roughness.

### **Measuring Perceived Roughness**


Surface roughness is one of texture's most prominent perceptual attributes. The precise physical determinants of roughness however are not exactly clear (e.g. Lederman, 1974). Because there is still debate over the actual parameters that determine roughness, users' *perception* of virtual roughness (regardless of the underlying physical model) is an increasingly important issue in virtual haptic interaction.

Participants will make a fixed choice response regarding a pair of surfaces. The roughest surface can be on the left, the right, or they can be judged as the same roughness. The proportion of times a surface is judged as rougher than each of the other surfaces can be obtained and the surfaces can then be placed along the roughness dimension.

### **Task and Procedure**

The haptic-audio surfaces will be presented in pairs as rectangular patches on the back wall of the workspace (see Fig. 2). Participants will be instructed to scrape the probe of the PHANToM back and forth across the stimulus surface to form an impression of how rough the surface seems to them. They will be asked to try to maintain the same speed throughout the experiment. The participant will then be asked to make a judgment regarding their comparison of the two surfaces. They make their response by clicking the appropriate button on the screen with the stylus switch on the probe of the PHANToM.

Clicking the button labeled 'next' will present the next pair of surfaces. When the participant has completed all the trials they will be given a message indicating that they are finished the experiment and a summary file for their responses will automatically be stored for that participant.



**Fig. 2:** Diagrammatic view of interface

## Hypotheses and Implications

### Integration of Information

Haptic-audio percepts of texture may reduce, increase, or completely alter the informational content of the percept being conveyed multimodally. The exact effects of the haptic-audio coupling will depend on the level at which the information is integrated. The level at which the multimodal information is integrated will depend, in part, on the level of congruency between the haptic and audio stimuli.

Participants will experience congruent and incongruent pairings of haptic and audio textures. The level of integration of these combinations can be *conflicting*, *redundant*, or *complementary*, each of which has the potential to affect perception and resulting interaction in different ways.

**H1 - Conflict:** If information processed by multiple modalities attempts to convey conflicting information in some way then the resulting multimodal percept may become distorted or completely lost in the process. Alternatively, the judgment of the multimodal percept might change in some unpredictable way.

If the audio stimulus and haptic stimulus are incongruent and conflicting then multimodal (haptic-audio) judgments of roughness will move along the roughness dimension but in the opposite direction predicted by the direction of the incongruency.

**H2 - Redundancy:** People might process only one modality of information from the many available to them in a multimodal percept. The modality employed may depend on physical/perceptual ability, personal preference, or the nature of the task for example. The actual effects of providing redundant information are somewhat difficult to predict. Redundant information might increase the mental representation of the information. This may in turn lead to increased confidence or reliability of judgments without necessarily altering the content of the information.

If the audio stimulus and haptic stimulus are congruent and redundant then with or without the auditory information, perceptual judgments of a virtual surface will be essentially the same. That is, the unimodal (haptic) and multimodal (haptic-audio) judgments of roughness will be at the same level along the roughness dimension.

**H3 - Complementarity:** A percept composed of multiple modalities might combine to in fact give more than the sum of the individual parts. That is, two unimodal percepts, when combined, produce some additive effect not possible with either unimodal percept alone. Such complementary pairings of haptic and audio stimuli might act to increase the quality and/or quantity of information available through a haptic-audio interface.

If the audio stimulus and haptic stimulus are incongruent but complementary then multimodal (haptic-audio) judgments of roughness will move along the roughness dimension in the direction predicted by the direction of the incongruency. That is, when an audio and haptic stimulus are combined such that the audio stimulus is more rough than the haptic stimulus then the multimodal judgment of roughness is moved along the roughness dimension in the direction of increasing roughness. Likewise, when an audio stimulus and haptic stimulus are combined such that the audio stimulus is less rough than the haptic stimulus then the multimodal judgment of roughness is moved along the roughness dimension in the direction of decreasing roughness.

## Future Work

Perceptual judgments of the unimodal stimuli are currently being gathered in preparation for combining them to produce the haptic-audio percepts. The next stage of the work will be to combine the haptic and audio textures to produce the congruent and incongruent multimodal percepts. This work will shed light on the ability of audio stimuli to alter the effect of haptic virtual stimuli and the different levels at which the haptic-audio percepts are integrated.

Work is underway to conduct an applied experiment of haptic-audio integration during force feedback texture perception. Veterinary simulation and visualization for the blind are being considered as possible applications areas. Results from the current study will serve to provide predictions regarding the effects of coupling haptic and audio information in a more applied example of force-feedback texture perception. Future work will also include a more in depth exploration of the levels at which we integrate haptic and audio information and how such organised knowledge would aid interface designers in the effective combination of haptic and audio information.

## Acknowledgments

This research is supported under EPSRC project GR/L79212 and EPSRC studentship 98700418. Thanks also go to the SHEFC REVELATION Project, SensAble Technologies and Virtual Presence Ltd.

## References

1. Brewster, S.A. (1997). Using Non-Speech Sound to Overcome Information Overload. *Displays*, 17, pp. 179-189.
2. Jansson, G., Fanger, J., Konig, H, Billberger, K. (1998), Visually Impaired Person's use of the PHANToM for Information about texture and 3D form of Virtual Objects, *Proceedings of the Third PHANToM Users Group Workshop*, Cambridge, MA: Massachusetts Institute of Technology.
3. Katz, D. (1989) *The World of Touch*, (Translated by Krueger, L.E.), Erlbaum, Hillsdale, NJ. Original work published in 1925.
4. Lederman, S.J. (1974). Tactile roughness of grooved surfaces: the touching process and effects of macro- and microsurface structure, *Perception and Psychophysics*, 16, 2, pp. 385-395.
5. Lederman, S.J., Klatzky, R.L., Hamilton, C.L., Ramsay, G.I. (1999). Perceiving Roughness via a Rigid Probe: Psychophysical Effects of Exploration Speed and Mode of Touch, *Haptic e-journal*, 1,1, (<http://www.haptics-e.org>).

6. Mynatt, E.D. (1997) Transforming graphical interfaces into auditory interfaces for blind users. *Human-Computer Interaction*, 12, PP. 7-45.
7. Oakley, I., McGee, M.R., Brewster, S. & Gray, P. (2000) Putting the Feel in 'Look and Feel', In *Proceedings of ACM CHI 2000*, The Hague, ACM Press, Addison-Wesley, pp.415-422.
8. Oviatt, S.L. (1997) Multimodal interactive maps: Designing for human performance. *Human -Computer Interaction*, 12, pp. 131-185.
9. Rosenberg, L.B. (1994), Design of a Virtual Rigid Surface: Haptic/Audio Registration, *CHI' 94*, April 24-28, pp. 257-258.
10. Srinivasan, M.A. & Basdogan, C. (1997), Haptics in Virtual Environments: Taxonomy, Research Status, and Challenges, *Computers and Graphics*, 21, 4, pp. 23-34.
11. Stevens, R.D., Edwards, A.D.N., & Harling, P.A. (1997). Access to mathematics for visually disabled students through multimodal interaction. *Human-Computer Interaction*, 12, pp. 47-92.
12. West, A.M. & Cutkosky, M.R. (1997), Detection of Real and Virtual Fine Surface Features with a Haptic Interface and Stylus, *Proceedings of the ASME Dynamic Systems and Control Division*, 61, pp. 159-166.
13. Wickens, C.D., Sandry, D.L., Vidulich, M. (1983), Compatibility and Resource Competition between Modalities of Input Central Processing, and Output, *Human Factors*, 25, 2, pp. 227-248.
14. Weisenberger, J.M., Krier, M.J. (1997). Haptic Perception of Simulated Surface Textures via Vibratory and Force Feedback Displays, *Proceedings of the ASME, Dynamic Systems and Control Division*, 61, pp.55-60.

# Feeling Rough: Multimodal Perception of Virtual Roughness

Marilyn Rose McGee, Philip Gray, Stephen Brewster

Glasgow Interactive Systems Group  
Department of Computing Science  
University of Glasgow, Glasgow, G12 8QQ, UK

*Tel: (0141) 330 3541*

*Email: {mcgeemr, pdg, stephen}@dcs.gla.ac.uk*

*Web: <http://www.dcs.gla.ac.uk/~mcgeemr>*

## Abstract

*The texture of a real or virtual surface can both increase the sense of realism of an object as well as convey information about an object's identity, type, location, function, and so on. It is crucial therefore that interface designers know the range of textural information available to them through current interaction devices in virtual environments. We have examined roughness perception of a set of force feedback generated textures (conveyed via a PHANToM device) in order to better understand the range and resolution of textural information available through such interaction. We propose that the addition of audio stimuli will increase further the potential for conveying more varied and realistic texture percepts through force feedback interaction. We are currently examining roughness perception of a set of auditory stimuli and will use both sets of unimodal results to test the potential benefits of combining haptic and auditory textural stimuli.*

## Keywords

Haptic, auditory, force feedback, texture, roughness perception, multisensory, multimodal interaction.

## Introduction

Despite the increasing prevalence of haptics in today's computing environments, the effective representation of such information is still a relatively new design problem for human computer interaction research. Force feedback interfaces in particular pose a variety of design questions such as what can and cannot be communicated convincingly via such devices.

In human sensing and manipulation of everyday objects, the perception of surface texture is fundamental to accurate identification of an object [5]. In a virtual world also, texture information can both increase the sense of

realism of an object as well as convey information about what the object is, where it is, and what it is for [4]. Through force feedback interaction in particular we can provide textural information that we can literally feel through the haptic modality. Given that it is often argued that touch is the 'reality sense' [2], being able to feel the texture of a virtual object should surely lead to increased realism of the object.

Previous work investigating the perception of real surface textures has shown that the physical properties of textures are complex and that an overall understanding of texture perception remains somewhat elusive [e.g. 3,4,5]. Textures are therefore proving difficult to reproduce successfully in virtual environments. It has been accepted however that roughness (along with hardness) is certainly one of the primary properties of a surface used to identify and classify an object. We have chosen therefore to focus our research initially on this dimension of roughness of virtual surfaces.

Simulating textures with force feedback devices in particular has proved an interesting research problem. Force feedback devices convey texture by actuating kinesthetic forces on the users' finger, hand, or body. This type of interaction relies on forces created through kinesthetic movement or displacement of the device and user limbs or joints while much of the texture perception we are used to comes through tactile stimulation of the mechanoreceptors on or just below the surface of our skin [5]. We have found in our previous work that such 'gross' or large textures can perturb the users' movements so much that the ability to stay on the textured surface is adversely affected [7]. More careful design of such force feedback based textures is required if these devices are to reach their full potential.

High fidelity force feedback devices (such as the PHANToM) are becoming increasingly realistic

interaction tools in a variety of applications where the texture of a virtual surface may be of great importance. Medical research for example can exploit such interaction in surgical and diagnostic simulations where the texture of tissue or organs may provide crucial information or feedback to the surgeon during a procedure.

Force feedback interaction is also improving the ability to design and prototype a variety of products ranging from the commercial (e.g. cars) to artistic and historical artifacts (e.g. sculptures and jewelry). E-Commerce will also benefit in that companies can provide their customers with a close representation of the feel of their products before they buy. The textile and fashion industry in particular could anticipate increased online sales if the texture of the clothing could be felt before purchasing [1].

It is crucial therefore that interaction designers know the potential range of textural information available through each modality and indeed each device available to them. With the increasing prevalence of force feedback interaction, it is particularly important to establish this for the haptic modality and high end of the range devices such as the PHANToM (SensAble Technologies).

Past research suggests that texture representation is possible through force feedback interaction but that the ideal solution is yet to be found. This is due in part to the mismatch between real texture perception (which involves both cutaneous and kinesthetic sensation) and virtual texture perception (which normally relies on either cutaneous sensation through tactile devices or kinesthetic sensation through force feedback devices). The problem could potentially be solved by advancing the currently available devices in order that the devices better suit real texture perception [e.g. 8]. This hardware-based solution is inevitable as the technology advances.

Another solution may be to improve the physical and mathematical modeling of real textures to produce the optimum algorithms for generating realistic virtual textures. This method currently has mixed results, as it cannot be assumed that the virtual exploration of texture matches that of real texture perception. Exact physical modeling therefore may be pointless if the interaction used to experience the texture differs significantly from that assumed by the physical model.

## Proposed Solution

Our approach offers a cost-effective approach that makes use of the currently available devices and even the simplest physical models of texture. We propose a multimodal solution that exploits the human ability to combine and integrate information from multiple sensory modalities into a fused and meaningful and whole percept.

We hypothesise that presenting combined haptic and audio percepts of roughness will increase the quantity and quality of textural information available through force feedback interaction alone.

## Overview of Experiments

The current work involved: (1) the evaluation of the effect of texture frequency on perceived roughness of a set of force feedback generated textures, (2) a follow up study extending the range of frequencies used and examining the possibility that there were two distinct notions of roughness emerging from the range of textural stimuli used, and (3) the evaluation of the effect of texture frequency on perceived roughness of auditory textures created from the profiles of the force feedback textures. This perceptual classification process will serve as a basis from which to test the eventual effects of systematically combining the haptic and auditory texture stimuli.

## The Force Feedback Device

The PHANToM 1.0 force feedback device by SensAble Technologies (Figure 1) was used to generate the virtual textures. Optical sensors detect changes in the device's configuration and mechanical actuators apply forces back to the user. Users interact with the device by holding a pen-like stylus attached to a passive gimbal on the device.



Figure 1: The PHANToM 3D force feedback device from SensAble Technologies.

By scraping this stylus/probe back and forth across the textured area the appropriate forces can be calculated from the positional information of the tip of the probe and the stored algorithmic models of the textured surface with which the user is interacting.

## Haptic Textures

Haptic textures were generated as sinusoidal waves or gratings on a rectangular patch on the back wall of the workspace. Figure 2 shows a diagrammatic view of the profile of a texture and the forces generated as a result of this profile.




Figure 2: (a) diagrammatic view of the profile of the texture; (b) indication of forces resulting from amplitude and frequency of texture wave.

The resulting texture profiles depend therefore on the amplitude and frequency of the sinusoidal waves. The textures had fixed amplitude of 0.5mm and variable frequency (cycles per 30mm). The frequencies used varied from 5 – 45. Higher frequencies were more tightly packed waves and lower frequencies were more loosely packed waves. The result of these textures was a bump felt at the peak of each wave.

## Auditory Textures

Auditory textures were generated from the same sinusoidal waves on a rectangular patch on the back wall of the workspace. The resulting profile still depended on the amplitude and frequency of the waves. The result of these textures was a single MIDI note generated from and heard at the peak of each wave. No experimental forces were experienced through the device.

## Roughness Comparisons

Participants in Experiment 1 could rate the textures as the same, the one on the left as rougher, or the one on the right as rougher. Participants in Experiment 2 were given the same options but with the additional response option of rating the textures as not comparable on the same roughness scale. This set of responses allowed us to evaluate (a) whether the participant perceived the two textures as the same or as different in terms of roughness, and (b) the number of times each texture was rated as the roughest of the pair.

In addition, the added response in Experiment 2 allowed us to evaluate (c) which textures participants felt

were different but not comparable along the same roughness scale. This was added as it was observed in experiment 1 that people often perceived a haptic difference but that they could not decide easily which one was in fact rougher.

## Procedure

Participants (Experiment 1, N=12; Experiment 2, N=10, Experiment 3, N= 12) were instructed to drag the probe of the device over each of the indicated textured surfaces and make a judgment on the roughness of the pair of textures. Participants compared each texture to itself and to each of the others twice (in a random order). In experiment 1, subjects compared 6 textures. In experiment 2 the frequency range was extended to include 9 textures.




Figure 3: Interface for roughness comparisons

Participants were allowed to explore each of the textures during that trial for as often as they liked and could switch between exploring the one of the left to exploring the one on the right as often as they liked to compare the two textures. They were instructed however that it was their initial response to the textures that mattered most and that there were not necessarily right or wrong answers for each of the trials. Participants made their response by clicking the switch on the probe of the PHANTOM to select the response that reflected their roughness judgment for each trial.

A training session identical to the experiment but with less trials allowed the participants to become familiar with the device and the interface. Importantly, it also allowed them to adopt an exploration strategy for experiencing the textures comfortably and successfully.

## Hypotheses

**Independent Variable:** frequency of texture (cycles per 30mm).

**Dependent Variable:** Perceived roughness, operationalised as the number of times each texture was judged as the roughest of the pair.

**Exp. 1:** The frequency of the haptic texture (or number of bumps) will have an effect on the perceived roughness of the texture.

**Exp. 2-a:** Increasing frequency of haptic texture (or number of bumps) will lead to an increase in the perceived roughness of the texture.

**Exp. 2-b:** There may be a bimodal function of roughness with a frequency from either end of the scale being perceived as the roughest of the set.

**Exp. 2-c:** Textures compared from either end of the frequency range are more likely to be rated as not comparable on the same roughness scale than textures compared within the high range or within the low range.

**Exp. 3:** The effect of frequency of audio texture (or number of notes) will have an effect on the perceived roughness of the virtual texture.

## Haptic Results (Experiment 1)

### Effects of Frequency on Perceived Roughness

The frequency of the texture was shown to have a significant effect on perceived roughness. That is, there was a significant effect of frequency on the number of times a texture was judged as the roughest of a pair ( $F=9.73$ ,  $p<0.01$ ). The number of times each (frequency of) haptic texture was judged as roughest tended to increase as the frequency of the texture increased (see Figure 4).




Figure 4: Effect of frequency on perceived roughness.

It is likely however that the range used in the experiment is only a sample from a more complex function. In fact, the graph shown may not be part of a simple monotonically increasing function at all. Instead it

may be part of a quadratic function of perceived roughness as suggested by people such as Lederman *et al.* [6]. At the very least, it may be likely that there is more than one maximum roughness generated from the set of frequencies.

### Two distinct notions of haptic virtual roughness?

Participant comments began to suggest that the lower frequency of 10 was considered very rough 'like corrugated material'. The higher frequencies of 30 and 35 however were also labeled as very rough but 'like sandpaper'. It is possible then that 2 frequencies from opposite ends of the scale can be perceived as equal in roughness magnitude but from different roughness scales.

Experiment 2 extended the range of textures (5-45 cycles) to evaluate whether the increasing frequency leading to increasing perceived roughness relationship still held beyond the range used in Experiment 1 and whether the bimodal peak roughness points emerged. This follow up study also evaluated our suggestion from Experiment 1 that comparing two textures from either end of the frequency range would increase the likelihood that they would be judged as different but also increase the likelihood that they would not be able to compare the textures on the same roughness scale. Final results from this evaluation will be presented at the workshop.

### Identical Haptic Stimuli

Textures with equal frequency were judged as the same roughness on an average of 64% of the trials. It appeared that higher identical frequencies were more likely than lower identical frequencies to be successfully judged as the same. This could perhaps due to the interaction between probe size and texture-profile size - lower frequencies being more susceptible to differences in hand force and exploration speed. Further statistical analysis of exp.1 and exp.2 will investigate this hypothesis further.

### Different Haptic Stimuli

A frequency separation of 5 cycles was not sufficient to significantly separate the perceived level of roughness for the haptic textures used. That is, textures separated by a frequency difference of 5 cycles were often judged as the same roughness. As frequency differences increased participants found it increasingly easy to decide whether the textures felt the same or different but increasingly difficult to decide which of the two was in fact the roughest. These results will be discussed in more detail at the workshop.

### Auditory Roughness (Experiment 3)

The audio virtual roughness experiment is currently underway and the effects of frequency of notes on perceived roughness of the audio textures are being evaluated using the same experimental paradigm. The MIDI instrument being used is piano although this will be compared to other instruments in the future. Our main concern for the initial audio experiment was purely to explore the effects of frequency of an arbitrary sound or note on the perceived roughness of the auditory texture. Results from the auditory roughness experiment will also be presented at the workshop.

### Future Work

The results of the haptic studies suggest that larger frequency differences lead to more easily distinguishable textures but also to difficulties in using the dimension of roughness in comparing textures. Large textures have also been found to throw users off of some textured areas [7]. The addition of audio information to such force feedback textures might ameliorate some of these restrictions.

We propose that the combined (multisensory) presentation of haptic and audio textural information will increase the range and/or resolution of textures available to the designer without disturbing interaction through force feedback devices. Results from the unimodal haptic and audio studies will be presented at the workshop. Our future multimodal (haptic – audio) experiment will also be discussed in more detail.

### Acknowledgments

This research is supported under EPSRC project GR/M44866 and EPSRC studentship 98700418. Thanks also go to the SHEFC REVELATION Project, SensAble Technologies and Virtual Presence Ltd.

### References

1. Dillon, P., Moody, W., Bartlett, R., Scully, P. & Morgan, R., (2000), Simulation of Tactile Sensation through Sensory Evaluation of Textiles when viewed as a Digital Image, *Proceedings of 1<sup>st</sup> International Workshop on Haptic Human-Computer Interaction*, University of Glasgow, UK, 31<sup>st</sup> August-1<sup>st</sup> September, pp. 63-68.
2. Gibson, J.J. (1962), Observations of Active Touch, *Psychological Review*, 69, 6, pp477-491.
3. Hollins, M., Faldowski, R., Rao, S., Young, F. (1993), Perceptual Dimensions of Tactile Surface Texture: A Multidimensional Scaling Analysis, *Perception and Psychophysics*, 54, 6, pp. 697-705.
4. Jansson, G., Fanger, J., Konig, H, Billberger, K. (1998), Visually Impaired Person's use of the PHANToM for Information about texture and 3D form of Virtual Objects, *Proceedings of the Third PHANToM Users Group Workshop*, Cambridge, MA: Massachusetts Institute of Technology.
5. Katz, D. (1989) *The World of Touch*, (Translated by Krueger, L.E.), Erlbaum, Hillsdale, NJ, Original work published in 1925.
6. Lederman, S.J. (1974). Tactile roughness of grooved surfaces: the touching process and effects of macro- and microsurface structure, *Perception and Psychophysics*, 16, 2, pp. 385-395.
7. Oakley, I., McGee, M.R., Brewster, S. & Gray, P. (2000) Putting the Feel in 'Look and Feel', *In Proceedings of ACM CHI 2000*, The Hague, ACM Press, Addison-Wesley, pp.415-422.
8. Wall, S.A. & Harwin, W.S., (2000), Interaction of Visual and Haptic Information in Simulated Environments: Texture Perception, *Proceedings of 1<sup>st</sup> Workshop on Haptic Human Computer Interaction*, 31<sup>st</sup> August - 1<sup>st</sup> September, Glasgow, Scotland, pp. 39-44.

# Haptic Perception of Virtual Roughness

Marilyn Rose McGee, Philip Gray, Stephen Brewster

Glasgow Interactive Systems Group  
Department of Computing Science  
University of Glasgow, Glasgow, G12 8QQ, UK

Tel: (0141) 330 3541

Email: {mcgeemr, pdg, stephen}@dcs.gla.ac.uk

Web: <http://www.dcs.gla.ac.uk/~mcgeemr>

## ABSTRACT

The texture of a virtual surface can both increase the sense of realism of an object as well as convey information about object identity, type, location, function, and so on. It is crucial therefore that interface designers know the range of textural information available through the haptic modality in virtual environments. The current study involves participants making roughness judgments on pairs of haptic textures experienced through a force-feedback device. The effect of texture frequency on roughness perception is analysed. The potential range and resolution of textural information available through force-feedback interaction are discussed.

## Keywords

Haptics, force-feedback, texture perception.

## INTRODUCTION

Despite the increasing prevalence of haptics in today's computing environments, the effective representation of haptic information is still a relatively new design problem for human computer interaction research. Force feedback interfaces in particular pose a variety of design questions. For example, what can and cannot be communicated convincingly via such devices?

The perception of surface texture is a specific design issue in haptic environments. In human sensing and manipulation of everyday objects, the perception of surface texture is fundamental to accurate identification of an object (Katz, 1989). In a virtual world also, haptic texture information can both increase the sense of realism of an object as well as convey information about what the object is, where it is, and what it is for (Jansson *et al.*, 1998).

There has been considerable previous work investigating the perception of real surface textures (e.g. Lederman *et al.*, 1974; Katz, 1989). The physical properties of textures are

complex and difficult to reproduce for virtual textures. Little is known about the perceptual response to virtual surfaces. Representing textures with force feedback devices in particular has proved problematic.

Force-feedback devices convey texture by actuating kinesthetic forces on the users' finger, hand, or body. This often relies on much larger forces than those typically experienced on the skin during real texture perception (Katz, 1989). We have found in our previous work that such gross textures can perturb the users' movements so much that the ability to stay on the textured surface is adversely affected (Oakley *et al.* 2000).

Force feedback devices are nonetheless becoming increasingly realistic interaction tools in a variety of applications where texture perception may be of importance. It is crucial therefore that designers know the range of textural information available through the haptic modality in virtual environments. The current study investigates the effects of frequency of texture on the relative perceived roughness of a set of force feedback generated textures.

## CURRENT EXPERIMENT

The current study involved participants making a series of roughness judgments on a set of force feedback generated textures explored via the PHANToM force feedback device. The user can rate one of the textures as roughest or both the textures as the same roughness. In this way, the proportion of times each texture is rated as rougher than each of the other textures can be determined.

**Hypothesis (A):** The frequency of the texture will have an effect on the proportion of times that texture is rated as rougher than each of the others.

**Hypothesis (B):** The frequency of the texture plotted against perceived roughness score will not produce a monotonic mapping from frequency of texture to perceived roughness. This will further reflect the complex nature of the concept of roughness.

The PHANToM 1.0 force feedback device (by SensAble Technologies) was used to create the haptic virtual surfaces.

Optical sensors detect changes in the device's configuration and mechanical actuators apply forces back to the user. Users interact with the device by holding a pen-like stylus attached to a passive gimbal on the device. By scraping this stylus/probe back and forth across the textured area the appropriate forces can be calculated from the positional information of the tip of the probe and the stored algorithmic models of the textured surface with which the user is interacting.

Six haptic textures were generated as (a series of) sinusoidal ridges on a rectangular patch on the back wall of the workspace. The resulting profile depended on the amplitude and frequency of the ridges. The textures had fixed amplitude of 0.5mm and one of 6 frequencies (cycles per fixed length of surface) - 10, 15, 20, 25, 30, or 35.

12 participants compared each texture to itself and to each of the others twice (in a random order) resulting in 42 trials that lasted an average of 35 minutes. Participants made their response by clicking the stylus switch on the probe of the PHANToM to select the button that reflects their roughness judgment for each trial.

## RESULTS AND DISCUSSION

The results highlight the complex nature of the concept of roughness as well as providing some guidelines as to how to present perceptually distinct virtual roughness through force feedback interaction.

### Effects of Frequency on Perceived Roughness

The number of times each (frequency of) texture was judged as roughest was measured as an overall roughness score (Table 1).

Frequency of texture	10	15	20	25	30	35
Perceived roughness score	24	18	38	35	61	69

Table 1. Effect of frequency on perceived roughness.

With the exception of a trough at frequency of 15 increasing frequency (for this range) leads to increased perceived roughness. It is likely however that the range used is only a sample from a probable quadratic function of perceived roughness (see Lederman et. al., 1974). In fact, it is likely that as the frequency of the texture goes below 10, the surface becomes a series of distinct bumps or waves rather than a unified texture. On the other side of the range, frequencies somewhere beyond 35 will become almost smooth again as the force profile becomes essentially flat.

### Same-Same Judgments

Textures with equal frequency were not reliably judged as the same roughness (accuracy: 50%-87.5%, mean: 64%). Lower frequencies were more subject to variations in perceptual differences. This is perhaps due to the interaction between probe size and texture-profile size;

lower frequencies being more easily affected by differences in hand force and exploration speed.

### Same-Different Judgments

A frequency separation of 5 was not sufficient to significantly separate the perceived level of roughness. For larger frequency differences, participants found it easy to decide whether the textures felt the same or different but much more difficult to decide which was roughest. This might be caused by the range of stimuli generating two distinct notions of roughness. Frequencies of 10 and 15 were perceived as "bumpy" or "corrugated roughness" whereas frequencies from 20-35 were perceived as "sharp" or "sandpaper roughness".

## CONCLUSIONS AND FUTURE WORK

The results of the study further illustrate the complex nature of the concept of roughness as well as providing some guidelines as to how to present perceptually distinct virtual roughness through force feedback interaction. The addition of audio information to such force feedback textures might increase the range and/or resolution of textures available to the designer through such devices alone. Work currently underway is investigating the effects of multimodally presented textures on the perception of virtual roughness.

## ACKNOWLEDGMENTS

This research is supported under EPSRC project GR/L79212 and EPSRC studentship 98700418. Thanks also go to the SHEFC REVELATION Project, SensAble Technologies and Virtual Presence Ltd.

## REFERENCES

1. Jansson, G., Fanger, J., Konig, H, Billberger, K. (1998), Visually Impaired Person's use of the PHANToM for Information about texture and 3D form of Virtual Objects, *Proceedings of the Third PHANToM Users Group Workshop*, Cambridge, MA: Massachusetts Institute of Technology.
2. Katz, D. (1989) *The World of Touch*, (Translated by Krueger, L.E.), Erlbaum, Hillsdale, NJ. Original work published in 1925.
3. Lederman, S.J. (1974). Tactile roughness of grooved surfaces: the touching process and effects of macro- and microsurface structure, *Perception and Psychophysics*, 16, 2, pp. 385-395.
4. Oakley, I., McGee, M.R., Brewster, S. & Gray, P. (2000) Putting the Feel in 'Look and Feel', *In Proceedings of ACM CHI 2000*, The Hague, ACM Press, Addison-Wesley, pp.415-422.
5. West, A.M. & Cutkosky, M.R. (1997), Detection of Real and Virtual Fine Surface Features with a Haptic Interface and Stylus, *Proceedings of the ASME Dynamic Systems and Control Division*, 61, pp. 159-166.

# Mixed Feelings: Multimodal Perception of Virtual Roughness

Marilyn Rose McGee, Philip Gray, Stephen Brewster

Glasgow Interactive Systems Group  
Department of Computing Science  
University of Glasgow, Glasgow, G12 8QQ, UK

*Tel: (0141) 330 3541*

*Email: {mcgeemr, pdg, stephen}@dcs.gla.ac.uk*

*Web: <http://www.dcs.gla.ac.uk/~mcgeemr>*

## Abstract

*The texture of a real or virtual surface can both increase the sense of realism of an object as well as convey information about an object's identity, type, location, function, and so on. It is important therefore that interface designers understand the range of textural information available to them through current interaction devices in virtual environments. Previous work (e.g. [2]), has examined the perceived roughness of a set of force feedback generated textures (conveyed via a PHANToM device) in order to work towards such an understanding. In doing so, this work has highlighted the possible perceptual limitations involved in reliably and confidently judging the relative roughness of a set of haptic textures. How many textures can we distinguish between for example and how likely is it that we reliably judge any one as rougher than, less rough than or the same as the other? The work presented here empirically investigates the effects of adding auditory textural cues to the existing haptic textures. Does the existence of an additional cue (in the auditory modality) change the answers to our questions above for example? We propose that the addition of auditory stimuli will increase the potential range and resolution of texture roughness percepts available through force feedback interaction.*

## Keywords

Haptic, auditory, force feedback, texture, roughness perception, multisensory, multimodal interaction.

## Introduction

In a virtual world texture information can both increase the sense of realism of an object as well as convey information about what the object is, where it is, and what it is for [2]. Through force feedback interaction in particular

we can provide texture information in virtual environments that we can literally feel through our haptic (touch) modality.

Previous work investigating the perception of real surface textures has shown that an overall understanding of the physical properties of textures remains somewhat elusive [e.g. 2,3,4]. Virtual textures are therefore not necessarily straightforward to produce. Despite the complex nature of textures it has been accepted that roughness (along with perhaps hardness) is one of the primary properties of a surface used to identify and classify an object. We have chosen therefore to focus our research on the dimension of roughness of virtual surfaces.

Force feedback devices convey texture specifically by actuating kinesthetic forces on the users' finger, hand, or body. This type of interaction relies on forces created through kinesthetic movement or displacement of the device and user limbs or joints while much of the texture perception we are used to comes through tactile stimulation of the mechanoreceptors on or just below the surface of our skin [3]. High fidelity force feedback devices (such as the PHANToM) are becoming increasingly realistic interaction tools in a variety of applications where the texture of a virtual surface may be of great importance. The exact quantity and quality of textural information available through such devices must therefore be explored.

Previous unimodal studies of the perceived roughness of a set of force feedback generated textures have shown some possible limitations in reliable roughness discrimination (for full details see [5]). It was found for example that participants did not necessarily judge identical textures as the same roughness. Nor did they necessarily judge adjacent textures in a set as reliably different in terms of roughness. The current experiment examines the effects of multimodality (adding auditory cues) on the perceived

roughness judgments of an equivalent set of force feedback textures.

This multimodal approach offers a cost-effective solution to overcoming the possible perceptual limitations of the currently available devices and texture models. Such a solution exploits the human ability to combine and integrate information from multiple sensory modalities into a fused and meaningful and whole percept. We hypothesise that presenting combined haptic and audio percepts of roughness will increase the reliability and confidence with which people can make comparative roughness judgements of force feedback textures.

## Overview of Experiment

### The Design

A within subjects (N=18) design was used with two independent variables - Modality of judgment, and Frequency of texture. The dependent measure was the Relative perceived roughness rating. The effect of texture frequency on perceived roughness ratings was evaluated as well as the effect of the modality of the judgments on those perceived roughness ratings. Computing Science students with no prior experience of the PHANToM participated in the experiment. All participants experienced all texture comparisons in all conditions. The order in which the modality conditions were experienced and the texture comparisons were presented within each condition were counterbalanced.

### The Force Feedback Device

The PHANToM 1.0 force feedback device by SensAble Technologies (Fig. 1) was used to generate the virtual textures. Optical sensors detect changes in the device's configuration and mechanical actuators apply forces back to the user. Users interact with the device by holding a pen-like stylus attached to a passive gimbal on the device.



Figure 1: The PHANToM 3D force feedback device from SensAble Technologies.

By scraping this stylus/probe back and forth across the textured area the appropriate forces or sounds can be

calculated from the positional information of the tip of the probe in combination with the stored algorithmic models of the textured surface.

### Haptic Textures

Haptic textures were generated as sinusoidal waves or gratings on a rectangular patch on the back wall of the workspace. Figure 2 shows a diagrammatic view of the profile of a texture and the forces generated as a result of this profile. The resulting texture profiles depended therefore on the amplitude and frequency of the sinusoidal waves. The textures had fixed amplitude of 0.5mm and variable frequency (cycles per 30mm). Higher frequencies were more tightly packed waves and lower frequencies were more loosely packed waves. The result of these textures was a bump felt at the peak of each wave.




Figure 2: (a) diagrammatic view of the profile of the texture; (b) indication of forces resulting from amplitude and frequency of haptic texture wave.

The frequencies that were used in the experiment varied from 10 – 35 cycles. These boundaries were selected due to the observations from our previous work [5]. Participants commented that textures of 5 cycles felt more like individual bumps than elements and that those of 40 and 45 were more of a smooth vibration when compared with the other more 'corrugated' or 'jagged' textures. The perceived roughness scores also confirmed this.

### Multimodal (Haptic-Auditory) Textures

Multimodal textures were generated from the same sinusoidal waves on a rectangular patch on the back wall of the workspace. The resulting profile still depended on the amplitude and frequency of the waves as in the unimodal haptic case. The result of dragging the PHANToM pen across these textures was a single MIDI note generated from and heard at or near the peak of each wave as well as the haptic forces as indicated above.




Figure 3: diagrammatic view of the profile of the multimodal texture for the congruent condition.

## Modality of judgement

**Haptic (H):** haptic textures are compared against other haptic texture.

**Multimodal Congruent (MMC):** a haptic texture is compared against a multimodal texture where the haptic frequency and auditory frequency are identical.

**Multimodal Incongruent (MMI):** a haptic texture is compared against a multimodal texture where the auditory frequency is 120% of the haptic frequency.

## Relative perceived roughness ratings

Participants could rate the textures as the same, the one on the left as rougher, or the one on the right as rougher. This allowed us to evaluate (a) whether the participant perceived the two textures as the same or as different in terms of roughness, and (b) the number of times each different texture was rated as the same as, rougher than, or less rough than the other texture.

## Procedure

Participants were instructed to drag the probe of the device over each of the indicated textured surfaces and make a judgment on the roughness of the pair of textures. Participants compared each texture to itself and to each of the others twice (in a random order).

## Hypotheses

### Effects of Frequency on Perceived Roughness

Does increasing frequency lead to increasing perceived roughness?

Does it do so in all conditions regardless of modality of judgment?

### Identical Haptic Stimuli

What is the likelihood that identical textures are judged as the same roughness?

Does this likelihood depend on the frequency of the texture?

Does the modality of the judgment have an effect on the number of times haptically identical textures are judged as the same?

### Different Haptic Stimuli

What is the likelihood that different textures are judged as different in terms of perceived roughness?

Does it depend how far apart the frequencies to be compared are?

Are higher frequency textures always judged as the roughest of a pair?

Does the modality of the judgement affect the likelihood that different textures are judged as different?

## Results

### Effects of Frequency on Perceived Roughness

Increasing frequency leads to increasing perceived roughness in all conditions.

### Identical Haptic Stimuli

What is the likelihood that identical textures are judged as the same roughness?

Does this likelihood depend on the frequency of the texture?

Does the modality of the judgment have an effect on the number of times haptically identical textures are judged as the same?

### Different Haptic Stimuli

What is the likelihood that different textures are judged as different in terms of perceived roughness?

Does it depend how far apart the frequencies to be compared are?

Are higher frequency textures always judged as the roughest of a pair?

Does the modality of the judgement affect the likelihood that different textures are judged as different?

To be presented in full at conference.

Full paper will also contain the results and discussion as well as conclusions and future work.

## Acknowledgments

This research is supported under EPSRC project GR/M44866 and EPSRC studentship 98700418. Thanks also go to the SHEFC REVELATION Project, SensAble Technologies and Virtual Presence Ltd.

## References

1. Hollins, M., Faldowski, R., Rao, S., Young, F. (1993), Perceptual Dimensions of Tactile Surface Texture: A Multidimensional Scaling Analysis, *Perception and Psychophysics*, 54, 6, pp. 697-705.
2. Jansson, G., Fanger, J., Konig, H, Billberger, K. (1998), Visually Impaired Person's use of the PHANToM for Information about texture and 3D form of Virtual Objects, *Proceedings of the Third PHANToM Users Group Workshop*, Cambridge, MA: Massachusetts Institute of Technology.
3. Katz, D. (1989) *The World of Touch*, (Translated by Krueger, L.E.), Erlbaum, Hillsdale, NJ, Original work published in 1925.
4. Lederman, S.J. (1974). Tactile roughness of grooved surfaces: the touching process and effects of macro- and microsurface structure, *Perception and Psychophysics*, 16, 2, pp. 385-395.
5. McGee, M.R., Gray, P.D., and Brewster, S. (2001), Feeling Rough: Multimodal Perception of Virtual Roughness, *Proceedings of 1<sup>st</sup> Eurohaptics Conference*, July, Birmingham, UK.