

Evaluation and Optimization of a Multi-Point Tactile Renderer for Virtual Textures

Matthew Philpott
Department of Physics
University of Exeter

Evaluation and Optimization of a Multi-Point Tactile Renderer for Virtual Textures

Submitted by Matthew Philpott, to the University of Exeter as a thesis for the degree of
Doctor of Philosophy in Physics, October 2013.

This thesis is available for Library use on the understanding that it is copyright material
and that no quotation from the thesis may be published without proper acknowledgment.

I certify that all material in this thesis which is not my own work has been identified and
that no material has previously been submitted and approved for the award of a degree
by this or any other University.

(Signature).....

Matthew Philpott

October 21, 2013

Abstract

The EU funded HAPTEX project aimed to create a virtual reality system that allowed a user to explore and manipulate a suspended virtual textile with the thumb and index finger. This was achieved through a combination of a tactile renderer on the fingertips for surface textures and a force feedback system for deformation of the virtual material.

This project focuses on the tactile rendering component of this system, which uses a tactile display developed at the University of Exeter. The 24 pin display is driven by piezoelectric bimorphs. Each of the pins can be driven independently, allowing for a variety of different sensations to be transmitted to the fingertip.

The display is driven by rendering software that uses a spatial spectrum of the intended surface, in combination with the frequency response of touch receptors in the skin, position on the surface, and exploration velocity to produce a signal that is intended to recreate the sensation of exploring the surface texture. The output signal on each of the 24 contactors is a combination of high (320 Hz) and low (40 Hz) frequency sine waves.

In this project, the tactile renderer is initially evaluated based on its ability to recreate the sensations of exploring particular textured surfaces. The users were asked to rank virtual textures in order of similarity to a real target texture. The results of the initial test were disappointingly low, with a $38.1 \pm 3.1\%$ correct identification rate. However, feedback from this initial test was used to make improvements to the rendering strategy. These improvements did not give a significant improvement in identification ($41.3 \pm 1.6\%$).

Finally, the tests were repeated with a target virtual texture instead of the real one used in previous tests. This test yielded a higher identification rate ($64.1 \pm 5.5\%$). This increase in identification suggests that the virtual textures are distinguishable but that they not always accurate recreations of the real textures they are mimicking.

Acknowledgements

When I started this project in 2009, having completed my MPhys in Astrophysics, I had minimal knowledge of haptics as an area of study. For me, the subject had mostly been an interesting curiosity but not something I had ever seriously looked into. However, when my PhD supervisor, Doctor Ian Summers, offered me the opportunity to work in this field, I found it hard to resist. The chance to dive into something almost completely new, connected to such a wide variety of subject areas such as Biology and Psychology, yet still not be completely out of my depth, thanks to my previous programming and scientific knowledge, was something I could not pass up.

Over the course of this project, Ian has been a great help to me. His ability to find solutions and new approaches to problems never ceases to impress me. I am grateful to him for the initial offer of this project and the support he is given me throughout.

Several other people also helped to make this project what it is today. My thanks go to Doctor Dennis Allerkamp, whose PhD work with the HAPTEX project created the original rendering software that was the starting point for this project. His help and documentation when I was first starting was a great help in my understanding of the renderer's code. I also want to thank Alasdair Allan for helping me while I was struggling with compiling the renderer's firmware. Without him, I would probably still be working on that now.

My mentor, Doctor Jenny Patience, has also been very helpful over the project's duration. While I may not have seen her as often as I possibly should have, she's always been ready to listen to what I have to say and help me sort through any issues I may have.

Many thanks go to the Biomedical Physics group and the staff and PhD students of the Exeter Physics department at large, who are always good company, regardless of their own pressures. Our conversations at tea-breaks and lunchtimes, scientific or otherwise, have helped make the day-to-day existence in the PhD lab far more enjoyable than I ever expected.

Outside of the lab, thanks go to my wider community of friends who have put up with me over the course of my undergrad, PhD and beyond. Members of the Science Fiction and Tolkien societies, both old and new, who've listened to my complaints about equipment not working or code not compiling and provided the diversions and discussions I've needed. You've kept me from driving myself mad over any challenges I've encountered.

Many thanks go to Amanda for listening to me vent my concerns and frustrations and encouraging me to stick with it even when I was most doubtful of my abilities. You have

more than made up for any help I gave you during your write up.

And, of course, I wouldn't be here without the love and support of my family. My thanks go to my parents, Clive and Gill, who gave me the freedom and encouragement to follow my desire for understanding and my brother, Adam, for reminding me what I'm aiming for. I also want to give thanks to my grandparents and other extended family members, who have always been supportive and happy to listen to what I've been doing, even if my explanations have left them somewhat perplexed at times.

“When you are studying any matter, or considering any philosophy, ask yourself only: What are the facts, and what is the truth that the facts bear out. Never let yourself be diverted, either by what you wish to believe, or what you think could have beneficent social effects if it were believed; but look only and solely at what are the facts.”

(Bertrand Russell, BBC Interview on “Face to Face”, 1959)

Contents

1	Introduction	18
2	The Sense of Touch	21
2.1	Psychophysics	21
2.1.1	Just-Noticeable Difference and Weber's Law	22
2.1.2	Psychometric Measures	22
2.1.3	Multi-Dimensional Scaling	23
2.1.4	Perception of Objects	23
2.2	Biology of Touch	25
2.2.1	Human Nervous System	26
2.2.2	Mechanoreceptive Afferents	29
2.2.3	Perception of Fine Textures	31
2.2.4	Frequency Response	31
2.3	Haptic Perceptual Spaces	32
2.4	Summary	38
3	Tactile Rendering Device	40
3.1	Existing Devices	40
3.1.1	Electrostatic Display	41
3.1.2	Piezoelectric Devices	42
3.1.3	Rheological Fluid Display	43
3.1.4	Electroactive Polymer Displays	43
3.1.5	Ultrasound Display	44
3.1.6	Surface Acoustic Waves	44
3.1.7	Shape Memory Alloy Displays	45

3.1.8	Electrocutaneous Display	45
3.1.9	Pneumatic Display	47
3.1.10	Electrovibration Display	47
3.1.11	Shear Force Display	48
3.2	The Tactile Rendering Setup	49
3.2.1	Principles of Operation	49
3.2.2	Tactile Display	53
3.2.3	Driving Electronics	56
3.2.4	Rendering Software	61
3.2.5	Texture File Creation	70
3.3	The HAPTEX Project	78
3.3.1	Overview of System Development	80
3.3.2	Physical Simulation	82
3.3.3	Force Feedback	83
3.3.4	Tactile Renderer	83
3.3.5	Evaluation	85
3.4	Improvements to Renderer	89
3.4.1	Direction of Motion	89
3.4.2	Textile Frequency Range	90
3.4.3	Modifications to Amplitude Calculation	90
3.4.4	Filter Modifications	91
3.4.5	Amplitude Scaling Factors	92
3.4.6	Buffering Calculated Amplitudes	94
3.4.7	Contact Surface Area	95
3.4.8	Dynamic Range and Spectral Balance	95
3.4.9	Summary of Modifications	97
3.5	Summary	98
4	Characterisation of Renderer	99
4.1	Characterisation Strategy	100
4.2	Characterisation of Experimental Renderer	101
4.3	Comparison with HAPTEX Renderer	110

4.4	Characterisation of Scaled Texture Files	111
4.5	Summary	115
5	Public Demonstration	116
5.1	Demonstration Environment	117
5.2	Demonstration Setup	118
5.3	Feedback	121
5.4	Summary	127
6	Virtual Texture Evaluation	128
6.1	Experimental Design	129
6.2	Experiment 1	131
6.3	Experiment 2	141
6.4	Experiment 3	148
6.5	Perceptual Dimensions	153
6.6	Experienced Users	157
6.7	Summary	157
7	Conclusions	161
7.1	Summary of Thesis	161
7.2	Overview of Project Conclusions	163
7.3	Discussion and Future Work	164
7.4	Potential Applications	166
	Appendices	168
	References	168
A	Texture Database	177
B	Spectra of Texture Files	189
C	Chance Result Derivation	190
D	Ethics Application	193
E	Additional Discriminability Plots	196

List of Tables

2.1	Comparison of different mechanoreceptor types based on their receptive field sizes and adaptation rates	30
2.2	The four mechanoreceptor populations, their sensitivity ranges, and their associated functions	33
2.3	List of the car seat material samples and their main characteristics	34
2.4	Co-occurrence matrix showing the number of participants who grouped specific pairs of stimuli together	35
2.5	Correlation coefficients between each adjective scale and each dimension of the MDS texture space	37
2.6	A comparison of the sensory perceptual spaces suggested in the literature	38
3.1	The characteristic values that are output by the KES-F	73
5.1	Results of the survey conducted to gather feedback about the tactile renderer installation	125
5.2	Results of the survey conducted to gather feedback about the presented facsimiles	126
6.1	The recorded details of the eight subjects who participated in experiment 1	132
6.2	The sample set used in experiment 1	132
6.3	Cumulative scores for matching real and virtual textiles in experiment 1	135
6.4	The number of subjects that chose the virtual textiles as their first choice match to the target textile	139
6.5	The sample set used in experiment 2	141
6.6	The recorded details of the eight subjects who participated in experiment 2	141
6.7	Cumulative scores from experiment 2 for matching real and virtual textiles	143
6.8	Dissimilarity matrix generated from the results of the second experiment, measuring dissimilarity between each real textile and the 16 virtual textile	145

6.9	The recorded details of the eight subjects who participated in experiment 3	149
6.10	Cumulative scores from experiment 3 for matching virtual textiles	149
6.11	Dissimilarity matrix generated from the results of the third experiment . . .	151
6.12	The one- and two-dimensional solutions to the Multi-Dimensional Scaling .	154
6.13	Comparison of the correct identifications between the subjects in the experiments and more experienced users	157
A.1	Abbreviations of the fibre contents of the textiles	177
A.2	Database of the available textures	178
B.1	The spectral lines of the texture files	189
C.1	The number of microstates available for each number of subjects that selected the most popular choice, and the total number of microstates	192

List of Figures

2.1	Movement patterns associated with each of the exploratory procedures . . .	24
2.2	The maximum-likelihood estimate integration of two hypothetical situations	25
2.3	Signal propagation in neurons	27
2.4	A schematic of a triggered action potential	28
2.5	The mechanoreceptors in the glabrous skin	29
2.6	Micrograph of a Pacinian corpuscle	30
2.7	Locations of Pacinian corpuscles within the right index distal phalange . . .	31
2.8	Threshold measurements for different populations of mechanoreceptors in the thenar eminence of the hand	32
2.9	Four dimensional MDS space presented as two two-dimensional planes . . .	36
3.1	An example of an electrostatic actuator, including structure and function .	41
3.2	Schematic diagrams of the piezoelectric display mechanism showing a cross- section side view and a top view	42
3.3	An example of an electrorheological actuator	43
3.4	A schematic of an electroactive polymer actuator	44
3.5	An example of an ultrasound display, consisting of an annular array of 91 ultrasound transducers packed in an hexagonal arrangement	44
3.6	A schematic of a surface acoustic wave tactile display	45
3.7	A side view of an example shape memory alloy actuator	46
3.8	A schematic of an electrocutaneous contactor	46
3.9	Schematic of the pneumatic circuit associated with an pneumatic actuator .	47
3.10	Schematic illustration of the electrode-skin interface	48
3.11	The ShiverPad generating a net rightward force	49
3.12	As the exploratory points move across the surface, they are displaced in a direction normal to the surface	49

3.13	A representation of the tactile workspace and how the spatial spectrum is mapped onto temporal frequency according to velocity of exploration	50
3.14	The spectrum of temporal frequencies is reduced to components at 40 Hz and 320 Hz, and the output is specified by their amplitudes, A_{40} and A_{320} . . .	52
3.15	Schematic drawing of the Piezoelectric effect	54
3.16	Schematic of a parallel configuration Piezoelectric bimorph	54
3.17	Measured frequency response of a piezoelectric drive element with and without the mechanical load presented by the skin	55
3.18	The tactile display used in this project and the arrangement of the static and dynamic contactors	56
3.19	The driving electronics for the tactile display	57
3.20	A schematic of the operation of the drive electronics	58
3.21	The board that hosts the USB controller, data bus and variable voltage supply	58
3.22	The circuit galvanically isolating the USB controller from the rest of the system	59
3.23	Amplification stages of the variable voltage supply	60
3.24	The board hosting the DAC, address logic, and the amplifiers	60
3.25	Passive, first order low-pass RC filter	61
3.26	The two amplitude stages of each channel	62
3.27	An example of an output signal, as generated by the driving electronics . .	62
3.28	The tactile rendering framework, showing the four threads within the software	63
3.29	The Kalman filter uses a predictor-corrector feedback loop	64
3.30	An example of the graphical display presented to the user	66
3.31	The schematic operation of the software's tactile rendering thread	66
3.32	Average detection thresholds as a function of frequency of the tactile receptors in the skin for Pacinian and non-Pacinian channels used in this project	68
3.33	The filter functions H_{40} and H_{320}	69
3.34	The HAPTEx selection process	71
3.35	The KES-F surface tester	74
3.36	Schematic of the KES-F roughness tester contactor	75
3.37	An example of a Hamming window	76
3.38	Derivation of a set of spatial spectra from a Kawabata line profile	76

3.39	The surface profiles recorded by the KES-F roughness tester while travelling across the right-side surface of textile 03, as well as the first 20 lines of the mean of the spectra	77
3.40	The surface profiles recorded by the KES-F roughness tester while travelling across the right-side surface of textile 20, as well as the first 20 lines of the mean of the spectra	79
3.41	The target virtual scenario for the HAPTEX project	80
3.42	Schematic of the concept of the proposed HAPTEX device	81
3.43	The different threads within the HAPTEX system	81
3.44	The development levels of the HAPTEX system	82
3.45	The GRAB system from PERCRO	83
3.46	A CAD model of the layout of the HandExos device	84
3.47	The tactile array used in the HAPTEX system	85
3.48	Finger positions and manipulations used to test the properties	86
3.49	Comparisons between mean ratings from subject 1 and subject 2 for the properties of tensile stiffness, surface roughness, surface friction, and bending stiffness	87
3.50	The relation between the average subjective ratings of the virtual textiles for both subjects and the corresponding physical values for tensile stiffness, surface roughness, surface friction, and bending stiffness	87
3.51	The relation between the average subjective ratings of the real textiles for both subjects and the corresponding physical values for tensile stiffness, surface roughness, surface friction, and bending stiffness	88
3.52	The relation between the average subjective ratings of the virtual textiles for both subjects and the corresponding average subjective ratings of the real textiles for both subjects for tensile stiffness, surface roughness, surface friction, and bending stiffness	88
3.53	The original band-pass filters that were used to highlight the fundamental frequency components for tactile rendering, compared to the new filter functions	91
3.54	The sum in both channels of the maximum amplitudes that were calculated for a 10 cm s^{-1} exploration of the various surface textures	93
3.55	Without the buffer for the amplitudes, the output wave is updated when the new amplitudes are received by the firmware from the rendering software	94
3.56	The Exeter tactile display before and after the addition of caps to the contactors	95

3.57	Comparison between the sum of mean characteristic amplitudes of sample set 1, and the subjective best match to the real textiles, as established by the author, and the fitted scaling power law	96
4.1	The characterisation plots for the tactile renderer, based on the texture data used in the first experiment, showing the mean and CoV values of A_{40} and A_{320} during exploration of texture surfaces at 10 cm s^{-1}	102
4.2	Comparison between the root-sum-square of warp and weft geometrical roughness from the KES-F and the sum of the mean 40 Hz and 320 Hz amplitudes calculated for the characterisation plots	103
4.3	The characterisation plots for the tactile renderer, based on the texture data used in the first experiment, showing the mean and CoV values of A_{40} and A_{320} during exploration of texture surfaces at 3 cm s^{-1}	104
4.4	The characterisation plots for the tactile renderer, based on the texture data used in the first experiment, showing the mean and CoV values of A_{40} and A_{320} during exploration of texture surfaces at 30 cm s^{-1}	105
4.5	The mean amplitude values for the two channels as a function of exploration speed for 1 cm s^{-1} to 30 cm s^{-1}	106
4.6	The CoV of the amplitude values for the two channels as a function of exploration speed for 1 cm s^{-1} to 30 cm s^{-1}	107
4.7	A comparison of the anisotropy of the mean intensity experienced by a user travelling in each of the four directions of motion across a surface, between textures 37_R, 39_R, 03_R, and 54_R	109
4.8	Characterisation plots based on the tactile renderer and texture data that existed at the beginning of the project, showing the mean and CoV values of A_{40} and A_{320} during exploration of texture surfaces at 10 cm s^{-1}	112
4.9	A comparison of the mean intensity experienced by a user travelling in each of the four directions of motion, for texture 54_R, between the experiments renderer and the original HAPTEX renderer	113
4.10	A comparison of the sum of the mean amplitudes between the scaled and unscaled texture files, as well as the scaling that was used	113
4.11	The characterisation plots for the tactile renderer, based on the texture data used in the second and third experiments, showing the mean and CoV values of A_{40} and A_{320} during exploration of texture surfaces at 10 cm s^{-1}	114
5.1	The setup being used as part of “The Big Bang South West Fair”	116
5.2	Two other strategies presented at “Touching the Untouchable”, allowing for physical and virtual explorations of a “Lewis Chessmen” piece	118
5.3	A replica of the queen “Lewis Chessmen” piece	119

5.4	Diagram of the ghost touch setup	119
5.5	Use of the ghost touch setup to explore the surface of the “Lewis Chessmen” piece	120
5.6	Diagram of the haptic pen setup	120
5.7	Use of the haptic pen setup to explore the surface of the “Lewis Chessmen” piece	121
5.8	The textile that was the subject of the recreation for “Touching the Un- touchable” project: the original, archaeological, textile along with a modern recreation	122
5.9	Schematic of the installation used at the museum	122
5.10	The installed tactile renderer, featuring the Exeter tactile display and a visual overlay to navigating the virtual environment	123
5.11	Two facsimiles of the Falkirk tartan	124
5.12	A group of the facsimiles presented at the “Touching the Untouchable” event	125
6.1	The tactile rendering setup used for the experiments	130
6.2	How the tactile display is held during use	131
6.3	Characterisation plots for the tactile renderer for the sample set of 16 tex- tures used in the first experiment, showing the mean and CoV values of A_{40} and A_{320} during exploration of texture surfaces at 10 cm s^{-1}	133
6.4	An example of one of the virtual workspaces presented during experiment 1, showing four virtual textiles that are to be ranked against the target textile	134
6.5	An example of the presentation of a real textile as a target	134
6.6	A breakdown of the results of experiment 1, showing the percentage of times each of the 8 subjects assigned the correct textile with a particular rank of similarity to the target	135
6.7	A breakdown of the results of experiment 1, showing the percentage of times each of the 16 correct virtual textiles was assigned a particular rank of similarity to the target	136
6.8	Mean scores from 5 subjects for matching real and virtual textiles, as a function of the number of textiles in the stimulus set	137
6.9	From experiment 1, breakdown of the percentages of how the “most pop- ular” choice scored across all the test items, compared to what would be expected from chance	138
6.10	A comparison of the distribution of mode choices between cases where the mode is the correct choice and cases where it is not the correct choice	139

6.11	The error patterns obtained from the first choices of the subjects made during experiment 1	140
6.12	Characterisation plots for the tactile renderer for the sample set of 16 textures used in the second experiment, showing the mean and CoV values of A_{40} and A_{320} during exploration of texture surfaces at 10 cm s^{-1}	142
6.13	A breakdown of the results of experiment 2, showing the percentage of times each of the 8 subjects assigned the correct textile with a particular rank of similarity to the target	143
6.14	A breakdown of the results of experiment 2, showing the percentage of times each of the 16 correct virtual textiles was assigned a particular rank of similarity to the target	144
6.15	From experiment 2, breakdown of the percentages of how the “most popular” choice scored across all the test items, compared to what would be expected from chance	146
6.16	Graph showing the relationship between the correlation coefficient between virtual textures and the Euclidean distance in the mean characterisation space	147
6.17	Graph showing the relationship between the correlation coefficient between virtual textures and the Euclidean distance in the CoV characterisation space	147
6.18	An example of one of the virtual workspaces presented during experiment 3, showing four virtual textiles that are to be ranked against the target textile	148
6.19	Comparison of the results from the 3 experiments	149
6.20	Graph showing the relationship between the correlation coefficient between virtual textures from experiment 3 and the Euclidean distance in the mean characterisation space	150
6.21	Graph showing the relationship between the correlation coefficient between virtual textures from experiment 3 and the Euclidean distance in the CoV characterisation space	152
6.22	A plot comparing the dissimilarity scores for identifying target textiles, experiment 2 against experiment 3, along with a line of equality	152
6.23	Scree plots comparing the stresses of various multi-dimensional solutions	153
6.24	The two-dimensional MDS solution	154
6.25	Schematic of the mean characterisation space being compared to the calculated MDS space	155
6.26	Comparison of the mean characterisation space and the two-dimensional MDS solution, rotated through 193.1°	158
D.1	The first information sheet given to potential subjects	194

D.2	The second information sheet given to potential subjects	195
D.3	The consent form to be signed by individuals who agree to participate in the study	195
E.1	Graph showing the relationship between the correlation coefficient between virtual textures and the Euclidean distance between mean 40 Hz amplitudes	196
E.2	Graph showing the relationship between the correlation coefficient between virtual textures and the Euclidean distance between mean 320 Hz amplitudes	197
E.3	Graph showing the relationship between the correlation coefficient between virtual textures and the Euclidean distance between CoV 40 Hz amplitudes	197
E.4	Graph showing the relationship between the correlation coefficient between virtual textures and the Euclidean distance between CoV 320 Hz amplitudes	198

List of Acronyms

AHRC	Arts & Humanities Research Council
CNS	central nervous system
CoV	coefficient of variation
DAC	digital-analogue converter
EP	exploratory procedure
EPSRC	Engineering and Physical Sciences Research Council
FA	fast adapting
FAST	Fabric Assurance by Simple Testing
FFT	Fast Fourier Transform
HAPTEX	HAPtic sensing of virtual TEXtiles
IC	integrated circuit
IDT	interdigital transducer
jnd	just-noticeable difference
KES-F	Kawabata Evaluation System for Fabrics
MDS	Multi-Dimensional Scaling
MLE	maximum-likelihood estimate
NP	non-Pacinian
P	Pacinian
PNS	peripheral nervous system
PZT	Lead Zirconate Titanate
RAM	random access memory
root-sum-square	root of the sum of the squares
SA	slowly adapting

SAW	surface acoustic wave
SMA	shape memory alloy
USB	Universal Serial Bus
VR	virtual reality
warp	vertical
weft	horizontal

Chapter 1

Introduction

The creation of virtual environments has become a commonplace task, with many devices existing for the recreation and transmission of sight and sound sensations. This virtual reality (VR) allows for a person to experience an environment where exposure to the real environment is not possible, either because the environment is fictitious, in the case of computer games, or would place the person, or others, in jeopardy, such as in medical or flight training.

In most situations, the transmission of sight and sound is enough for the VR as these senses are most commonly employed by humans in the transfer of information. However, the inclusion of haptic information, in addition to adding general depth to the virtual environment, may be a requirement for some VR applications to be truly useful. For example, in medical training, a trainee surgeon would desire haptic feedback from the different virtual tissues being manipulated during a practice surgery session.

Unlike sight and hearing, the receptors of which are concentrated within a relatively small area and can be excited remotely with light and sound waves, the skin has a large area, being between 1.5 to 2 square metres, with receptors throughout at various densities and currently requires direct contact to the skin to stimulate. The result is that haptic devices typically focus on a small area of the skin to receive the haptic sensations, such as the fingertip and, even on a limited area, require comparatively sophisticated hardware to generate appropriate haptic stimuli.

Most haptic applications involve conveying the sensations through proxy tools that exist both in the virtual and real spaces, serving as a haptic bridge between the two environments. These tools can range from the more general purpose (thimble or pen-like probe) to the specific (scalpel handle or aircraft yoke), as appropriate. The main drawback to this approach, regarding haptic sensation, is that the user is not touching the virtual object directly. This limits the possible realism of the VR system.

The tactile renderer used in this project was created as part of the EU funded HAPTEX project, the purpose of which was to develop a VR system for the virtual recreation of textiles, visually and haptically. The scenario that this system was expected to address was where a virtual textile is suspended from a fixed support. It was expected that the user should be able to manipulate the virtual textile with the thumb and index finger,

feeling appropriate tactile stimulation at the fingertips. This scenario combines tactile feedback for the fingertips and force-feedback for the position of the fingers in space.

The tactile component of this VR system, which this project is based on, was developed from a combination of tactile stimulator hardware, developed at the Biomedical Physics group at the University of Exeter, and haptic rendering software to compute the tactile signals, developed at the Welfenlab at the Leibniz Universität Hannover. The development of this tactile rendering system is described in detail in [1].

Time constraints did not allow for a proper evaluation of the efficacy of the tactile rendering before the conclusion of the HAPTEX project. The aim of this project, therefore, is to examine the ability of the renderer to recreate the textured surfaces associated with textiles and make improvements where necessary, based on feedback from users. This evaluation was conducted in the Biomedical Physics group of the University of Exeter and was funded by the Engineering and Physical Sciences Research Council (EPSRC), with additional funding from the Arts & Humanities Research Council (AHRC).

Chapter 1 introduces the reader to the work presented here, beginning with an introduction haptic feedback and to how it relates to wider VR situations. It details the context for the creation of the tactile renderer used in this project, and provides an overview of the contents of each of the chapters within this work.

Chapter 2 begins with a discussion of the field of psychophysics, which described the relationship between an external stimulus with the subjective perception of that stimulus, an overview of Multi-Dimensional Scaling, which is a technique commonly used in psychophysics for establishing perceptual spaces, and a description of how haptic perception of an object is achieved. This is followed by an overview of the biology of the human touch system, including the properties of the various touch receptors that exist within the skin. The chapter then concludes with an example from the literature of how Multi-Dimensional Scaling was used to establish the perceptual dimensions of the sense of touch. It also provides an overview of the touch perceptual spaces established within the literature from the exploration of real surface textures.

Chapter 3 starts with an overview of various tactile devices that exist within the literature, and their capabilities, which provides context for the choice of tactile device used within this project. Following this is a detailed explanation of the tactile display used in this project, as well as a description of the tactile rendering system, both hardware and software, used to recreate the surfaces of textiles virtually. As the renderer was developed as part of the HAPTEX project, an overview of the project is given here. This includes details of the development of the integrated system and the evaluation that was done towards the end of the project. The chapter then concludes with a summary of the improvements and modifications made to the system over the course of this project.

Chapter 4 documents work done to characterise the output from the tactile renderer, something that was not done during the HAPTEX project due to a lack of time. The start of the chapter describes how the output of the renderer may be characterised, and how that characterisation may be validated against measurements of the surface properties made by the KES-F. This characterisation is then performed on the tactile renderer as

it existed at various stages of the project. Specifically, the version that existed at the end of the HAPTEX project/beginning of this project, the version that existed for the initial experiment of this project, and the version that existed for the second and third experiments of this project. There is also some discussion on the effects of the speed of exploration and anisotropy of the surface texture on the characterisation values.

Chapter 5 documents a public demonstration of the renderer that was done at the National Museum of Scotland as part of a project by the Archaeology department of Exeter University to explore the reception of touch based interaction into a museum environment. The tactile renderer was one of the virtual interfaces that was presented, along with various other physical and virtual recreations of different artefacts. Feedback from the visitors was used to gauge the efficacy of the renderer in a realistic setting, as well as suggest any changes that needed to be made to the renderer.

Chapter 6 describes the experiments undertaken in this project to demonstrate the efficacy of this rendering system in creating virtual textured surfaces. The first experiment evaluates the efficacy of the renderer in recreating surface textures by asking test subjects to rate virtual textiles in similarity to a real textiles. The second experiment repeats the first experiment's protocol but with modifications made to the rendering strategy based on the feedback from both experienced users and the test subjects. This second experiment also uses a different sample set to reflect this change in strategy. The third experiment replaces the target real textile with a virtual one to test the discriminability of the virtual textures from each other. The results from this final experiment were used to construct a perceptual space for the virtual textures. Finally, the experiments were repeated by individuals who were previously familiar with the device, so as to investigate the effect of experience on texture discrimination.

Chapter 7 serves to summarise the work described in this document, and the various conclusions that can be drawn from it. It also suggests what possibilities exist for future work within this field, as well as what potential applications the renderer might have in real world settings.

Chapter 2

The Sense of Touch

In order for the previously developed rendering strategy to be understood, and thus evaluated and improved, it is necessary to understand how the sense of touch works and how it leads to a perception of the world. The purpose of this chapter is therefore to provide the reader with an overview of the area of human haptic sense. From this, it can be understood in later chapters how the tactile sense may be fooled by the renderer into perceiving the sensations associated with exploring textured surfaces.

Initially, this chapter introduces the concepts of psychophysics, which establishes a theoretical relationship between a real-world stimulus and perceived sensation, Multi-Dimensional Scaling, an analysis technique which may be used to examine the relationship between stimulus and perception, and how the haptic sense can be used to perceive objects and obtain estimates of particular properties, which can be improved via integration with other modalities. This is followed by an explanation of the biology of the touch system that exists within humans, including the number, type, and properties of the various touch receptors within the skin, including the frequency response associated with the different receptor types. This chapter then concludes with a brief overview from the literature of the perceptual dimensions that have previously been identified for the sense of touch, based on the exploration of real textures, including an example of how Multi-Dimensional Scaling can be used to establish these perceptual dimensions.

2.1 Psychophysics

Psychophysics, a term coined in 1860 by Gustav Theodor Fechner in “Elemente der Psychophysik” (Elements of Psychophysics), is the field of study that relates physical stimulus and perceived sensation. Fechner observed growing knowledge of both the material world and of the mind but noted the lack of a solid foundation relating the two, even though the two are connected by a single reality. Hence, psychophysics was developed to determine the relationship between physical and psychological worlds [2]. Psychophysical applications include uses in digital signal processing, where it has informed the development of lossy compression to allow for the compression of video and audio signals with very little loss of perceived signal quality.

2.1.1 Just-Noticeable Difference and Weber’s Law

A just-noticeable difference (jnd) is defined as the minimum detectable difference in some attribute (e.g. intensity) between two stimuli. The jnd in the intensity ΔR for a given stimulus of intensity R is given approximately by the equation known as Weber’s Law, where the constant c_w depends on the kind of the stimulus (equation 2.1).

$$\Delta R = c_w \cdot R \tag{2.1}$$

The law is named for Ernst Heinrich Weber who, in 1834, studied jnds and found the relationship which holds true for weight, visual and tone discrimination, although Weber himself never stated the law in this mathematical form. Instead, it was Fechner in 1860 who put it in this form [3].

2.1.2 Psychometric Measures

In order to determine the efficacy of the tactile rendering strategy, it is necessary to compare the difference between the perception of the real texture, and the perception of the virtual texture. Although the subject could be asked to directly compare the sensations generated by the virtual and real textures, it is more useful to provide the subject with a set of virtual sensations. From this set of virtual sensations, the subject is asked to identify the appropriate virtual texture for a given real texture. This creates a theoretical probability that an arbitrary subject could select the correct virtual texture from the available options through random chance alone.

Reference [4] provides a list of six standard discrimination methods that can be used to establish the psychometric distances (i.e., subjective differences between stimuli), which are variants on the user being presented with a sensation and attempting to choose the corresponding sensation from a set of other sensations. One example of these methods is the Same-Different method, where the subject is presented with a pair of samples and must decide if the received pair is different, or the same. Alternatively, there is the Duo-Trio method, where the subject is presented with a set of 3 samples (one of the six combinations of the two products being investigated). The first of the set is considered the ‘control’ and the subject must choose one in the remaining set to match the control.

For the experiments used in this project, subjects were presented with four virtual textured surfaces. They were also given a fifth surface (real or virtual, depending upon experiment) and asked to rate the four virtual textures in order of similarity to the given texture. If the subjects were unable to distinguish between two or more textures, then they were asked to give their best guess. One of the virtual textures is intended to match the “target” texture so, if the subjects were simply guessing at random, we would expect a successful identification rate of 25%.

2.1.3 Multi-Dimensional Scaling

There are many situations where one dimension is insufficient to relate a stimulus and the perceived sensations. Additionally, in many of these situations, the dimensions of the stimulus are not known. For these situations, Multi-Dimensional Scaling (MDS) may be used to find the perceptual distance between stimulus points by embedding them in a multidimensional space, from the given similarity of pairs of objects/stimuli. The distances between the points may relate, to a greater or lesser extent, to the measured similarity of the various object pairs, depending on the number of dimensions used for the spaces. The error in the generation of the spaces, referred to as the stress, reduces with increasing dimensions with zero stress at dimensionality $m \geq n - 1$, where n is the number of objects, which allows for the distances between the points to exactly match the measured similarity. There are three standard measures for the goodness-of-fit of the MDS solution; Stress which relates to the perceptual distances [5]; S-stress which relates to the square distances [6]; and $1 - R^2$, which relates to the proportion of the variance not accounted for by the MDS solution [7].

The specific details of calculations relating to MDS are not given here. Instead, the interested reader is directed to [8] for more information on how the positions of the points is computed. MDS is quite versatile and maybe applied to a wide variety of situations, including the perception of colour [9], the evaluation of textures (section 2.3), and clinical psychology [10], among others.

2.1.4 Perception of Objects

The haptic perception of an object's properties can be achieved through a variety of manipulations of the object, with particular patterns of exploration, referred to as exploratory procedures (EPs), being observed when people are asked to investigate particular properties. [11] describes the relationship that is observed between the various EPs and the particular properties of the object being investigated, defining a total of eight EPs:

- “Lateral Motion” EP, exploring surface texture by sideways movement between the skin and the object's surface.
- “Pressure” EP, exploring hardness by applying torque of normal forces to a part of the object, while another part is stabilized or an opposing force is applied.
- “Static Contact” EP, exploring the temperature by one hand resting, passively, on the object's surface.
- “Unsupported Holding” EP, exploring the weight of the object by lifting the object away from the supporting surface and held in the hand.
- “Enclosure” EP, exploring the global shape and/or volume of the object by the hand maintaining simultaneous contact with as much of the envelope of the object as possible, often moulding the hand to more precisely follow the object's contours.

- “Contour Following” EP, exploring exact shape and/or volume by moving the hand along and maintaining contact with a contour of the object, stopping or shifting direction when a contour segment ends.
- “Function Testing” EP, exploring the function of the object by applying movements that perform the function of the object in question.
- “Part Motion Test” EP, exploring the relative motions of different parts of the object by applying force to one part while stabilising or applying counter force to the rest of the object.

These EPs are illustrated graphically in figure 2.1. Six of these EPs can be applied to all objects while the two remaining EPs (function test and part motion test) only apply in certain situations, depending on the nature of the object being explored. Function testing is only applicable to objects that perform a function, such as placing a hand or finger into a container or pinching the ends of a pincer together. Similarly, the part motion test is only relevant if the object in question has one or more moving parts to explore.

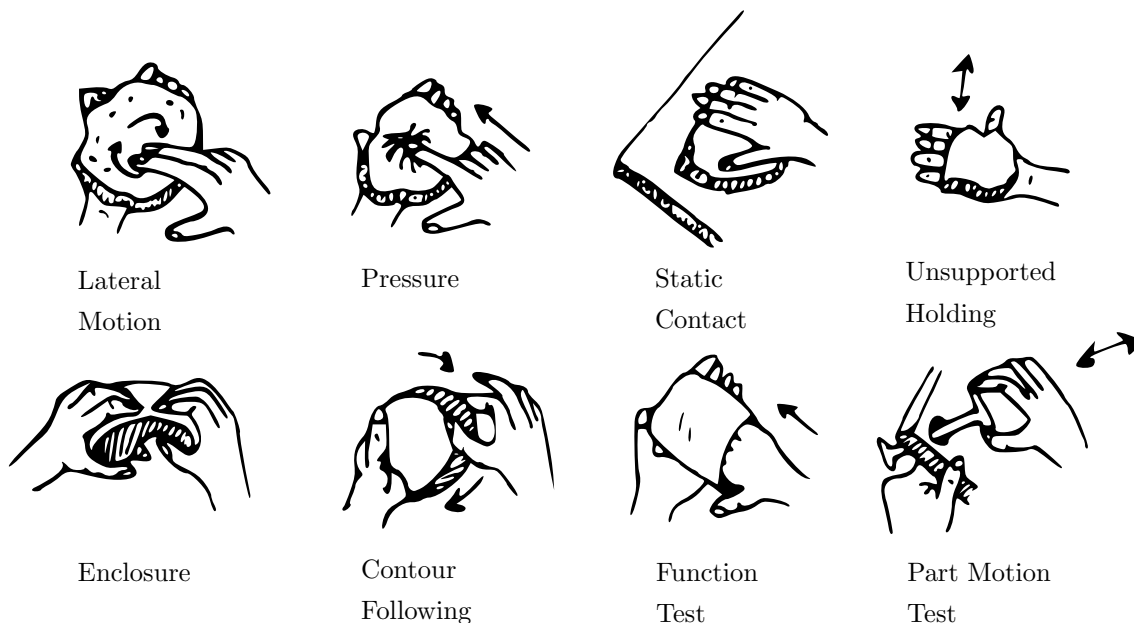


Figure 2.1: Movement patterns associated with each of the exploratory procedures (recreated from [11]).

While estimates may be made on an objects properties based solely on particular EP, the integration between different sensory modalities produces a more accurate assessment of the object than the independent sensory modalities. It is proposed in [12] that the integrated system of the visual and haptic senses combines the sensory information in a way that is similar to a maximum-likelihood estimate (MLE) model. In this model, the estimation of a property \hat{S}_i in a modality i is given by

$$\hat{S}_i = f_i(S) \quad (2.2)$$

where S is the property being estimated and f_i is the operation of the sensory system that is estimating the property. Each estimate has a noise associated with it, which is assumed

to be independent and Gaussian with a variance of σ_i^2 . Under the MLE rule, the sensor estimates are combined by addition, weighted by their normalized reciprocal variances. Thus the MLE of the property is given by

$$\hat{S} = \sum_i w_i \hat{S}_i \text{ with } w_i = \frac{1/\sigma_i^2}{\sum_i 1/\sigma_i^2} \quad (2.3)$$

The task presented in [12] related to the integration of cues from the visual and haptic modalities. Therefore, using the MLE rule, the variance of the estimate S_{VH} , which is a combination of visual and haptic estimates \hat{S}_V and \hat{S}_H , can therefore be defined by

$$\sigma_{VH}^2 = \frac{\sigma_V^2 \sigma_H^2}{\sigma_V^2 + \sigma_H^2} \quad (2.4)$$

The estimate that results from combining the visual and haptic modalities, as illustrated in figure 2.2, has a combined variance that is lower than either the visual-alone or haptic-alone estimations.

The results of the supporting experiments described in [12], where subjects were asked to visually and haptically estimate the height of stimuli, appears to validate this approach, with combined visual and haptic estimations similar to those predicted by the MLE integrator. Although verification of this combination of sensory modalities is done between the haptic and visual senses, the author of [12] suggests that this could also work between different combinations of sensory modalities, or even different cues within the same modality.




Figure 2.2: The maximum-likelihood estimate integration of two hypothetical situations (recreated from [12]). (Left panel) The visual and haptic variances are equal ($\sigma_H^2/\sigma_V^2 = 1$). In this case, the mean of the combined estimation is equal to the mean of the visual-alone and haptic-alone estimations and the variance of the combined estimate is reduced to half of the individual measurements. (Right panel) The haptic variance is four times greater than the visual variance ($\sigma_H^2/\sigma_V^2 = 4$). In this case, the combined estimation is shifted towards the visual estimation with a variance that is 0.8 times the visual-alone variance.

2.2 Biology of Touch

The human haptic sense can be divided into two categories; the kinaesthetic sense and the tactile sense. The kinaesthetic sense uses signals from the muscles and skin receptors to provide information on the static and dynamic posture of the body, in addition to

muscular effort. Touch sensation that is mediated entirely through cutaneous stimulation is referred to as the tactile sense. Information about events and objects conveyed from both the cutaneous and kinaesthetic senses is collectively referred to as haptic perception. [13]

The focus of this project is to exploit the tactile sense to generate the sensation of exploring a surface. To generate these sensations, an understanding of the biology of human touch perception is needed, as detailed in this section. Details of the nervous system described in this section are mainly based on reference [14].

2.2.1 Human Nervous System

Neurons, are electrically excitable cells, which form the nervous system. Each cell consists of a main cell body (soma), dendrites, and an axon (figure 2.3). A typical neuron receives its signals through its dendrites, which collect the information and forwards it to the soma. From there, the information is transmitted by the axon which, at its end, branches into several terminals. Neurons are able to transmit information to each other through a synapse, which is a point of contact between two neurons.

The neurons are not themselves connected. Rather, there is a 20 nm gap (the synaptic cleft) between the two. To transmit the signals, the gap is crossed by a group of chemicals called neurotransmitters. These chemical messengers are bound within synaptic vesicles on the pre-synaptic side of the synapse and are released across the synaptic cleft. Upon reaching the post-synaptic side of the synapse, they bind with receptors on the membrane of the receiving neuron to illicit an effect.

The nervous system allows for information to be collected, transported throughout the body, and responded to. Sensory neurons transmit information detected from external stimuli to the central nervous system (CNS), where the information is analysed and interpreted by interneurons. In humans, the CNS consists of the brain and spinal cord, the interneurons of which have the greatest complexity of connections. While the spinal cord can act independently in sensory processing, such as in the cases of reflex actions or basic locomotion, it is also used to transport sensory information to the brain where it can be analysed by the appropriate cortices, allowing for more complex behaviour. The integration of the sensory signals can result in the production of other sensory signals that may lead to excitation of appropriate efferent, or motor, neurons and thus excitement of the muscles they are connected to. Together, the sensory and motor neurons form the peripheral nervous system (PNS).

Any given neuron can receive information from potentially thousands of other neurons, via the synapses. Each synapse influences the membrane potential of the neuron, either exciting (increasing) or inhibiting (decreasing) it. The effect of a synapse on the membrane potential depends on both the location, and time since the last signal of the synapse, with the effect decaying over time. Neurons will sum the effects of multiple synapses, either through spatial summation where the effects of different synapses in different locations are combined, or temporal summation, which combines the affects of several signals from




Figure 2.3: Signal propagation in neurons. (Image source: U.S. National Institute on Aging)

a single synapse. The latter only occurs if the time between two signals is short enough that the first signal has not yet decayed before the second one is received.

Membrane potential is the potential of the neuron relative to the outside of the cell. A neuron in its resting state has a potential of -70 mV . As the neurotransmitters for the synapses influence the membrane, this potential changes. If the membrane reaches the threshold potential of -50 mV to -55 mV , an action potential is triggered (figure 2.4). This action potential of $+35\text{ mV}$ begins in the soma and propagates down the axon. At the terminals of the axon, the potential triggers the release of neurotransmitters to continue propagation of the signal. After this spike, the membrane potential returns to the resting potential.




Figure 2.4: A schematic of a triggered action potential. (Image source: Wikimedia Commons)

High intensity stimuli are encoded, not by the amplitude of the action potentials, but rather by their frequency, denoted by impulses per second. Thus, the upper limit on the frequency is given by the time taken for a neuron to recover to its resting potential following an action potential.

In order to allow for perception of the external environment, various transducers exist throughout the body that convert physical stimuli into a neural signal which can be interpreted by the CNS. These transducers emit neurotransmitters when they are excited by the relevant stimulus, generating the signals in neighbouring neurons via synapses which can then propagate through the nervous system. In [14], these sensory transducers are grouped into five categories, based on the nature of the stimuli they transduce:

- “Mechanoreceptors”, which sense the physical deformation that is caused by mechanical activity (pressure, touch, sound, etc.).
- “Chemoreceptors”, which respond to chemical stimuli, either as general receptors, detecting information about total solute concentration such as the osmolarity of blood, or as specific receptors, detecting particular molecules such as oxygen or glucose.
- “Electromagnetic Receptors”, which detect various forms of electromagnetic radiation (such as visible light or body heat).
- “Thermoreceptors”, which react to different temperatures (either hot or cold).
- “Nociceptors”, also known as pain receptors, react to potentially harmful stimuli (such as extreme pressure or temperature).

While the sensory cells typically contain a particular type of receptor for a particular stimulus, various kinds of receptors can also serve as nociceptors in cases where the relevant stimulus is too extreme. For example, a thermoreceptor which detects heat can trigger a pain response if the stimulus is too hot, and therefore would be potentially damaging.

2.2.2 Mechanoreceptive Afferents

The human tactile sense is comprised of a set of four types of mechanoreceptors that are embedded in the glabrous skin, as illustrated in figure 2.5. They can be divided into various different groups based on their adaptation speeds or their response fields, with table 2.1 showing how the mechanoreceptor types are distributed among the various groups. Those which adapt quickly are denoted as fast adapting (FA) and those which adapt slowly are denoted as slowly adapting (SA). Additionally, those with comparatively small receptive fields with distinct borders are part of group I while those with large receptive fields without distinct borders are part of group II. The specific properties of the various mechanoreceptors, including density and receptive field area, are summarized below, from [15] and [16].




Figure 2.5: The mechanoreceptors in the glabrous skin (from [17]).

Table 2.1: Comparison of different mechanoreceptor types based on their receptive field sizes and adaptation rates (adapted from [18]).

		RECEPTIVE FIELDS	
		Small, sharp borders	Large, obscure borders
ADAPTATION	Fast (No static response)	FA I Meissner Corpuscles	FA II Pacinian Corpuscles
	Slow (Static response present)	SA I Merkel Neurite Complex	SA II Ruffini Corpuscles

For the afferent nerves with small fields, the FA units are associated with the Meissner corpuscles and the SA units are associated with the Merkel cell neurite complex. These units have receptor around 10 mm^2 . The densities of the units vary throughout the hand with the highest concentrations in the finger tip at 140 cm^{-2} (Meissner) and 70 cm^{-2} (Merkel), which decrease down to 25 cm^{-2} (Meissner) and 8 cm^{-2} (Merkel) in the palm of the hand.

Receptors with large fields are the Pacinian corpuscles (figure 2.6), which are FA units, and the Ruffini endings, which are SA units. Their receptive fields are around 100 mm^2 . Unlike the small field receptors, the distribution of large field units is fairly uniform throughout the hand at 0.22 mm^{-2} (Pacinian) and 0.09 mm^{-2} (Ruffini) respectively. The receptor distribution is not completely uniform, however, as shown in [19], where the locations of Pacinian receptors were identified from an MRI scan of subjects' fingertips. Figure 2.7 shows that the distribution of Pacinian receptors is highest towards the sides of the fingerpad and a lower density towards the centre.




Figure 2.6: Micrograph of a Pacinian corpuscle (Field of view is $500 \mu\text{m}$ by $300 \mu\text{m}$) (from [20]).

Since highest density touch receptors within the skin is on the order of 1 mm^{-2} , it can be suggested that the detection threshold of the tactile sense should also be on that order. Experiments performed in [21] seem to verify that is the case, with spatial acuity in the

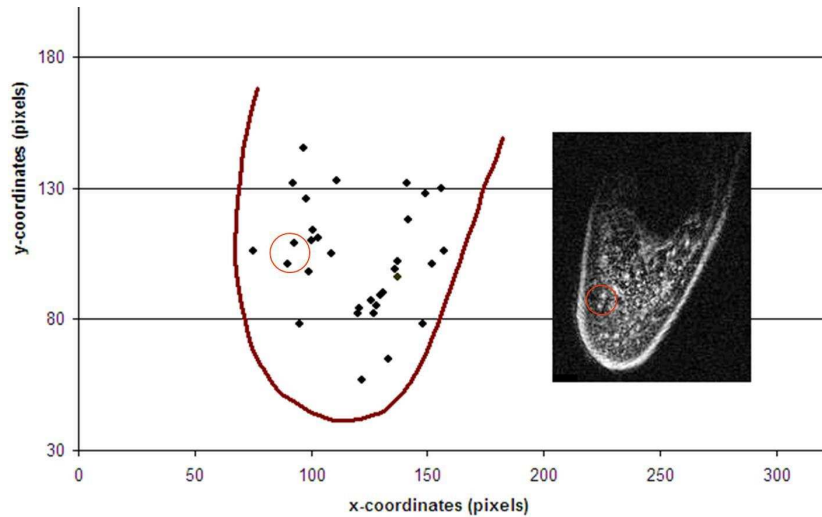


Figure 2.7: Locations of Pacinian corpuscles within the right index distal phalange (from [19]). The outline of the finger is indicated by the U-shaped line. The inset shows a single slice of the 3D MRI data.

fingertip found to be of around 1 mm.

2.2.3 Perception of Fine Textures

The limit on the perception of roughness cannot solely be based on the resolution of the touch receptors however, because humans are able to perceive fine textures that have structure sizes below the spatial resolution of the SA I receptors. The experiments described in [22, 23, 24, 25] indicate that the Pacinian corpuscles are involved in the perception of fine textures, in terms of temporal aspects (i.e., vibration).

The authors of [26] investigated the importance of fingerprints in the perception of fine textures by comparing two biomimetic tactile sensors with geometrical and mechanical properties consistent with the human fingertip, where one of the sensors was patterned by parallel ridges to mimic fingerprints. When the fingerprinted sensor is used on a randomly textured surface, the resulting spectrum of vibrations is dominated by a particular frequency. This frequency is determined by the ratio of scanning speed to the interridge distance and is within the frequency range consistent with the Pacinian receptors. This frequency is not present on the equivalent, smooth-surfaced probe.

2.2.4 Frequency Response

For each of the mechanoreceptors, there exists a detection threshold, which represents the minimum detectable amplitude of a stimulus by that receptor from a signal of particular frequency. The literature includes many psychophysical studies where test subjects are asked to respond to closely controlled stimuli, typically displacement sine waves of known amplitude and frequency, that were at or near the threshold levels (e.g., [27], [28], [29], [30], [31]). Figure 2.8 shows that the experimental results can be interpreted in terms of perception via four different channels, one for each mechanoreceptor, whose sensitivity

varies with frequency. Experiments also suggest that detection threshold is also dependent on the density of mechanoreceptors, such as in [19], where the centre of the fingertip, which had a lower density of Pacinian receptors, had a higher detection threshold than towards the side of the fingerpad, where the receptor density is lower.




Figure 2.8: Threshold measurements for different populations of mechanoreceptors in the thenar eminence of the hand. The points marked by the square used a stimulator with a large contactor and those marked by a circle used a stimulator with a small contactor (recreated from [30]).

Each of the mechanoreceptors is sensitive to a different frequency range and serves a different function in touch perception, which is summarized in a table 2.2. The Meissner corpuscles are sensitive to temporal changes in skin deformation, ranging in frequency from ~ 5 Hz to ~ 40 Hz. This detection of low frequency vibrations contributes to precision grasping and manipulation of objects, which the other mechanoreceptors also contribute to. The Merkel cell neurite complexes are sensitive to frequencies of less than ~ 5 Hz, allowing for the perception of form and textures. The Pacinian corpuscles are sensitive to higher frequencies, ranging from ~ 40 Hz to ~ 400 Hz, allowing for fine texture perception. Ruffini endings are not subject to temporal changes in skin deformation and are instead sensitive to sustained downward pressure and lateral skin deformation. This allows for the detection of the direction and force of a moving object.

2.3 Haptic Perceptual Spaces

In order to demonstrate how psychophysics may be applied to quantifying perceived sensations, the following details the experiments used in [42] which identified some perceptual dimensions used in the evaluation of real fabric textures. The identification of the percep-

Table 2.2: The four mechanoreceptor populations, their sensitivity ranges, and their associated functions (recreated from [32]).

Population	Maximum Feature Sensitivity	Primary Function
SA I	Sustained pressure; maximally sensitive to very low frequencies ($< \sim 5$ Hz) [34]; spatial deformation [36]	Very-low frequency vibration detection [33] Coarse texture perception [35] Pattern/form detection [37] Stable precision grasp and manipulation [38]
FA I	Temporal changes in skin deformation (~ 5 Hz to ~ 40 Hz) [34]; spatial deformation [36]	Low-frequency vibration detection [33] Stable precision grasp and manipulation [38]
FA II	Temporal changes in skin deformation (~ 40 Hz to ~ 400 Hz) [34]	High-frequency vibration detection [33] Fine texture perception [24] Stable precision grasp and manipulation [38]
SA II	Sustained downward pressure, lateral skin stretch [39]; low dynamic sensitivity [34]	Direction of object motion and force due to skin stretch [40] Stable precision grasp and manipulation [38] Finger position [41]

tual dimensions, as presented in this case, may be broken down into three components; the identification of possible dimension labels, the generation of the perceptual space, and identifying the dimensions from the list of possible labels. For this study, the fabrics consisted of a set of samples of 24 different car seat materials (listed in table 2.3), selected for their differing tactile properties.

During the first part of the study, forty subjects were asked think about various fabrics they encounter in everyday life (in general, not limited to the 24 listed in table 2.3) and the tactile impressions produced by touching these fabrics. They were then asked to produce, from memory, a list of adjectives that describe the tactile impressions of fabrics, as the adjectives came to mind. The average subject produced a list of ten descriptors with a combined total of 27 descriptors being suggested by at least 10% of the subjects. Analysis of these descriptors, by comparing mean frequency of production to mean order of production, indicates that more common descriptors are evoked earlier in the subject's mind.

In order to establish the perceptual dimensions of tactile sensations, for the second part of the study twenty subjects participated in a free-sorting task where they were asked to sort the 24 textures of table 2.3 into groups based on perceived similarity. This was followed by a characterisation task where the the participants were asked to produces labels for each of the resulting groups. In order to establish the perceptual dimensions, a co-occurrence matrix was generated based on the number of times (maximum of twenty) that two textures were put into the same group (table 2.4). Using this matrix, the stimuli may be positioned in a perceptual space based on their perceived dissimilarity (figure 2.9). As well as multi-dimensional scaling, correspondence analysis was also used with

Table 2.3: List of the car seat material samples and their main characteristics (from [42]).

Sample	Main characteristics
t01	Plastic alveolar fabric
t02	Plastic fabric with wavy relief
t03	Velvet fabric with circular inscription
t04	Velvet fabric
t05	Velvet fabric with ribbed relief
t06	Velvet fabric with leaf-like inscriptions
t07	Plastic alveolar fabric
t08	Woven fabric with light squared inscriptions
t09	Velvet fabric
t10	Velvet fabric with fuzzy inscriptions
t11	Suedette
t12	Velvet fabric with squared inscriptions
t13	Jeans-like fabric
t14	Woven fabric
t15	Velvet fabric with squared inscriptions
t16	Velvet fabric with squared inscriptions
t17	Velvet fabric with fine ridges
t18	Velvet fabric with oblique ridges
t19	Woven fabric with fine stripes
t20	Woven fabric with fine stripes
t21	Velvet fabric with large ribs
t22	Velvet-like fabric with large ribs
t23	Woven chequered fabric
t24	Woven chequered fabric

a contingency table, which has descriptors in rows and stimuli in columns and whose numbers represent the mean weighted frequency at which each descriptor was employed for a stimulus, to generate an additional perceptual space, which is not included here.

For the final part of this study, subjects were asked to rate the textures on a set of 10 provided unipolar, 9-point adjective scales, selected based on the results of the previous experiments (soft, harsh, pleasant, thin, thick, hard, mellow, relief, smooth and rough). The mean scores of the adjective scales could then be computed for each stimulus. Linear regression analyses were used to suggest the degree of relation between an adjective scale and a dimension of the proposed perceptual space (table 2.5). In addition to a strong positive correlation between some adjectives and dimensions, some adjectives produce a strong negative correlation. This can help to identify possible meaningful dimensions, particularly if the negatively correlated adjective can be considered to be an opposite of a positively correlated one (Thin vs Thick, for example), as this suggests a continuous perceptual dimension between the two adjectives. The adjective scales have been included in the MDS space (figure 2.9) such that the relative length of each scale reflects its coefficient of regression (R^2), with the sign of the correlation coefficients used to determine the orientation of the scales within the space. For clarity, as in [42], the scales extend from the origin to a point defined by $2R^2$.

The perceptual space suggested at the conclusion of this study is a three-dimensional space, with dimensions of Soft/Harsh, Thin/Thick, and Relief. The author of [42] does

Table 2.4: Co-occurrence matrix showing the number of participants (maximum 20) who grouped specific pairs of stimuli together (from [42]).

	t01	t02	t03	t04	t05	t06	t07	t08	t09	t10	t11	t12	t13	t14	t15	t16	t17	t18	t19	t20	t21	t22	t23	t24	
t01	20																								
t02	13	20																							
t03	0	1	20																						
t04	1	0	1	20																					
t05	7	6	2	1	20																				
t06	1	0	0	7	1	20																			
t07	17	15	0	0	7	0	20																		
t08	0	0	4	0	6	0	1	20																	
t09	1	0	0	12	1	12	0	0	20																
t10	1	0	0	11	1	13	0	0	17	20															
t11	0	0	0	10	1	11	0	0	15	12	20														
t12	1	0	2	5	1	13	0	0	9	10	7	20													
t13	0	1	4	1	2	0	0	8	1	0	2	0	20												
t14	1	3	2	0	1	0	2	6	0	0	0	0	6	20											
t15	1	0	2	4	1	8	0	0	4	5	4	13	0	0	20										
t16	1	0	2	4	1	8	0	0	4	4	4	13	0	0	18	20									
t17	0	0	0	17	0	7	0	0	13	12	12	5	2	0	3	3	20								
t18	1	1	1	1	4	0	2	7	1	0	1	0	12	12	0	0	1	20							
t19	1	2	2	1	1	1	0	6	1	1	0	1	9	12	1	1	0	9	20						
t20	0	2	1	0	0	1	0	4	0	1	1	1	10	14	1	0	1	9	15	20					
t21	0	0	9	0	3	0	0	4	0	0	0	0	2	2	0	1	0	2	1	1	20				
t22	0	0	9	0	3	0	0	4	0	0	0	0	2	2	0	1	0	3	1	1	20	20			
t23	6	3	1	0	5	0	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	20	
t24	6	3	1	0	5	0	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	20	20

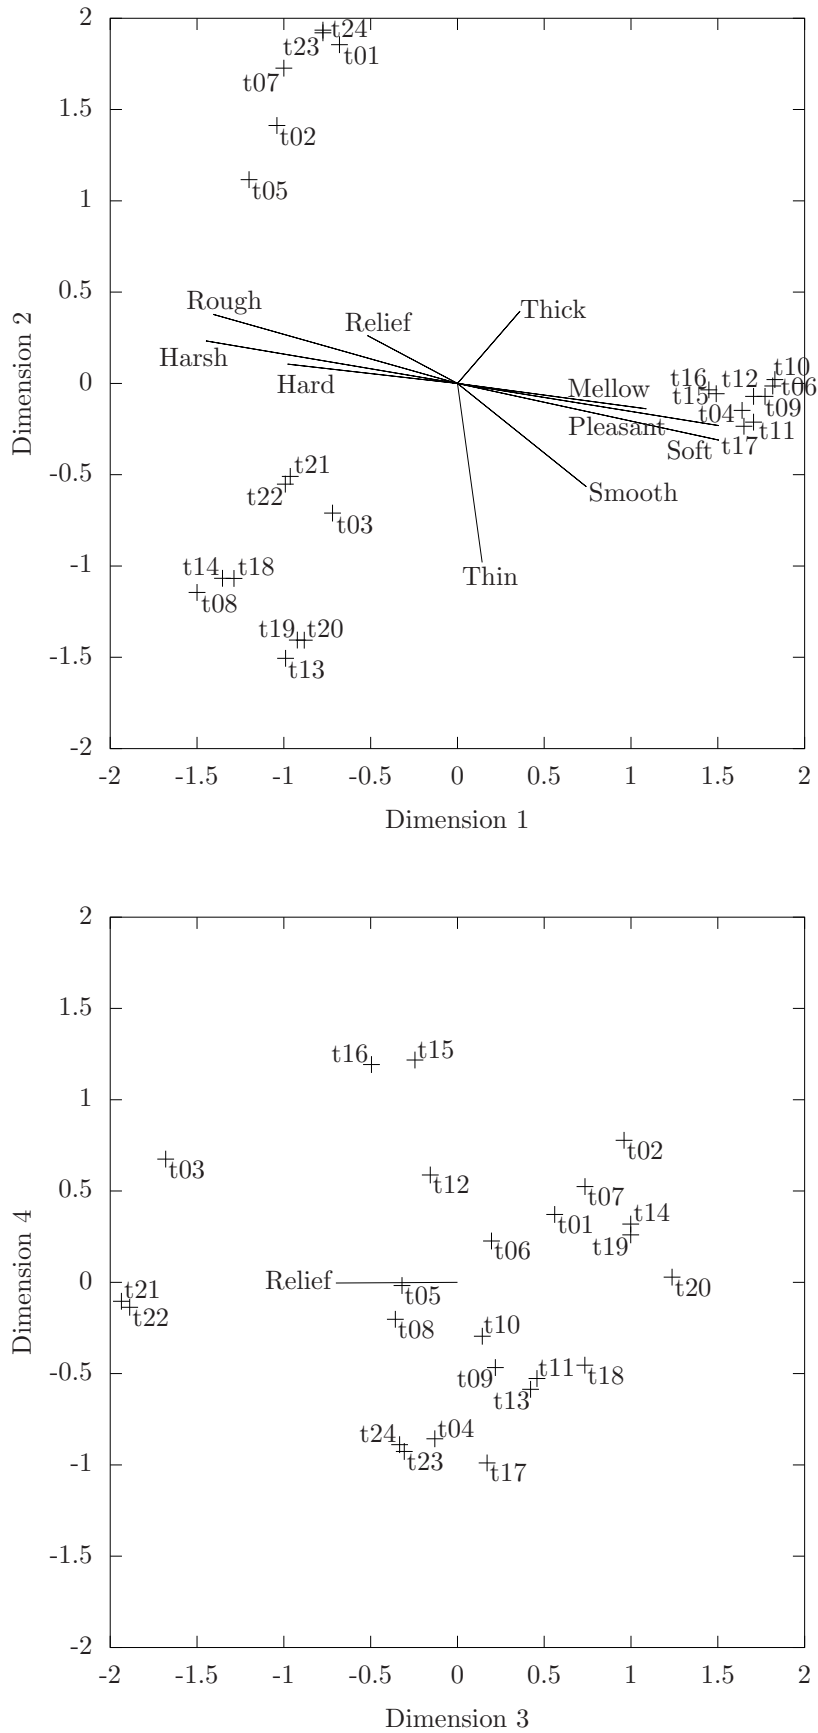


Figure 2.9: Four dimensional MDS space presented as two two-dimensional planes, with the first and second dimensions in the upper panel and the third and fourth dimensions in the lower panel. Adjective scales are regressed in the space, the endpoints defined by $2R^2$ (recreated from [42]).

Table 2.5: Correlation coefficients between each adjective scale and each dimension of the MDS texture space (from [42]).

	Soft	Harsh	Smooth	Rough	Thin	Thick	Hard	Mellow	Relief	Pleasant
Dimension 1	0.86	-0.86	0.60	-0.85	0.24	0.41	-0.71	0.73	-0.52	0.86
Dimension 2	-0.39	0.35	-0.53	0.44	-0.70	0.45	0.24	-0.26	0.37	-0.34
Dimension 3	-0.03	-0.03	0.37	-0.05	0.17	-0.12	-0.14	0.03	-0.60	0.08
Dimension 4	0.12	-0.13	-0.27	-0.10	-0.17	0.20	-0.37	0.36	0.29	0.17

however note that they did not identify roughness as a perceptual dimension, as would be expected from the existing literature, nor did the adjective ‘rough’ align with the adjective ‘smooth’ within the texture space. This may be a result of differences in language as the experiment was conducted and the results recorded in French but the publication is written in English. The author of [42] therefore suggests that roughness may be accounted for, to some extent, by the soft/harsh dimension.

This is one example of the creation of tactile perceptual spaces. Table 2.6 contains a collection of perceptual spaces from the literature and the labels associated with the various directions. Although there are some dimensions that appear a number of different times (rough/smooth, hard/soft), many possible dimensions are only accepted tentatively, suggesting that no definitive perceptual space for the tactile sense yet exists.

Table 2.6: A comparison of the sensory perceptual spaces suggested in the literature. Labels that are tentative, either because the researchers did not give a label, or because the provided label was explicitly stated as tentative, are denoted by question marks (adapted from [43]).

Study	Stimuli	Dimensions
[43], Experiment 1	33 sensory words	Rough/Smooth Wet/Dry Hot/Cold
[43], Experiment 2	5 textured stimuli	Rough/Smooth Wet/Dry? Texture?
[44]	124 diverse materials	Soft/Hard? Rough/Smooth? Featural Regularity?
[45]	26 fabrics	Soft/Harsh Thin/Thick Supple/Stiff
[42]	24 car seat materials	Soft/Harsh Thin/Thick Relief
[46]	17 diverse materials	Rough/Smooth Soft/Hard Sticky/Slippery?
[47]	17 diverse materials	Rough/Smooth Hard/Soft Elasticity?

2.4 Summary

Psychophysics is the field of study that relates the subjective perception of a particular stimulus to the objective properties of that stimulus. One of the established relationships within the psychophysics field is Weber’s Law, which describes, for a given stimulus, the minimum change that must occur before the difference is perceivable. This is referred to as the just-noticeable difference (jnd) and it suggests that, within the tactile sense, there is a possibility for similar surface textures to be mistaken for each other. Multi-Dimensional

Scaling (MDS) is commonly used in psychophysical studies to construct perceptual spaces based on the perceived dissimilarity between stimuli, as a one-dimensional solution may not be enough to fully account for the perceptual distances between stimuli. When exploring the haptic properties of an object, humans have been observed to employ a number of specific exploratory procedures (EPs) in order to obtain estimates of the properties, with the precise nature of the EP depending on the specific property being investigated. While the resulting haptic estimates are subject to a degree of uncertainty then can be combined with estimates for the same properties obtained from other modalities to produce an estimate with less uncertainty than any of the modalities individually.

Physically, human senses can be attributed to the various receptors that exist throughout the body. Different receptor populations react to different stimuli, including mechanical deformation (referred to as ‘mechanoreceptors’) which form the basis of the human tactile sense and so are of interest here. The various types of mechanoreceptor have different properties, including varying densities throughout the skin, difference in response areas and difference in adaptation speeds. Each of the different receptor types has a frequency dependent detection threshold, representing the minimum amplitude that the type of receptor can detect from a stimulus of a particular frequency.

Throughout the existing literature, efforts have been made to determine the relevant properties that are used to distinguish between different textures. The literature suggests that the tactile sense is best addressed with a three-dimensional solution, but the labels associated with those dimensions are, at present, ill-defined. Properties that have been proposed as the dimensions of texture perception include, among others, roughness, softness and thickness.

Chapter 3

Tactile Rendering Device

Although there are many devices that convey sound and visual information to us in our daily lives, devices that convey tactile information are not as prevalent. However, tactile devices potentially have a variety of uses, including adding more information into a virtual environment for the purpose of virtual training [48] or to provide additional feedback about the environment to blind/visually impaired individuals [49]. In the case of this project, the tactile rendering system was created as the tactile component for a VR system developed by the HAPTEX project. The purpose of this VR system was to create virtual recreations of textiles that the users of the system could explore haptically, either through the sensation of the surface texture via a tactile display, or through manipulation of the virtual textile via a force-feedback device.

The beginning of this chapter compares some existing tactile displays within the literature that use a variety of methods to generate and convey tactile sensations to the user. This is followed by a more detailed description of the setup that was used for this project specifically, including detailed descriptions of the principles behind the operation of the display, details of the Exeter tactile display itself, which is a 24 channel piezoelectric display driven by sine waves of 40 Hz and 320 Hz, as well as the supporting hardware and software. As the tactile renderer was developed as part of the HAPTEX project, section 3.3 gives an overview of the project, including the different stages of development and the evaluation that was done during the project. This is then concluded by a review of the various changes that have been made to the setup in order to improve the rendering strategy from the version that existed at the end of the HAPTEX project. An explanation is given as to why those changes were made, including some changes that were made based on analysis of the results from the first experiment (section 6.2).

3.1 Existing Devices

There exists a variety of tactile displays, each of which uses one of a number of possible methods to provide a tactile stimulus to the user. A common method is to use an array of contactors that stimulate the receptors in the skin, although some displays have only a single contact point. Different patterns of contactor movement lead to different tactile

sensations in the stimulated skin.

The displays can vary in a number of properties, including density of contactors, frequency bandwidth, expandability in number contactors, and compactness of completed display. All of these must be considered in relation to the intended use when selecting a display method. For example, the contactors of a particular display maybe compact enough to increase the density, and therefore the resolution, of the display. However, if the display is intended to be moved with the user as they explore a textured surface then the increased contactor numbers might reduce the compactness of the display to a point where it may no longer be feasible to move the display with the exploring fingertip.

The following is a selection of the various tactile display types that exist, based on a variety of principles. This is not an exhaustive list as that is beyond the scope of this work.

3.1.1 Electrostatic Display

In [50], [51], and [52], a layered dielectric polymer is used to drive the actuators in the tactile display. By applying an electric field across the dielectric, as shown in figure 3.1, the polymer contracts between the electrodes due to electrostatic attraction and thus, by varying the field strength, the height of the stimulator can be varied correspondingly, i.e. the stronger the field, the lower that stimulator height. Contactors of this nature have a compact display surface (< 1 mm between contactors) but have the disadvantage of requiring automation to effectively construct.




Figure 3.1: An example of an electrostatic actuator, including structure and function (replicated from [50]).

The necessity of automation comes from the requirement of the dielectric film being kept to a small thickness (a few microns) to reduce the voltages required to adjust the height of the contactor. However, this limits the change in height that a single layer can achieve. This is addressed through the use of a large number of alternating dielectric film and electrode layers, as shown in figure 3.1, allowing for greater absolute displacement of the contactor at the same voltage, thus a much smaller voltage than would be required to

make the same displacement from a single layer of dielectric. In [50], the dielectric layers are $10\ \mu\text{m}$ thick, requiring the use of 1000 layers to produce appropriate displacements.

3.1.2 Piezoelectric Devices

Actuators based on piezoelectric materials, which deform when a potential difference is applied across them, are used in a variety of tactile display designs. [53] details a prototype display that uses piezoelectric actuators to create movement of 50 stimulator pins at a $2\ \text{mm}$ pitch in a direction normal to the surface of the skin with amplitudes that can be adjusted from $5\ \mu\text{m}$ to $57\ \mu\text{m}$, but only at a frequency of 250 Hz. The same piezoelectric actuators are used in [54] in a 10 pin array, at a $3\ \text{mm}$ spacing and a maximum amplitude of $22\ \mu\text{m}$. Similarly, [55] has 48 contact pins at a resolution of about $8\ \text{mm}^2$ per contactor and allows for amplitudes up to $0.7\ \text{mm}$ but the bandwidth is limited to 20 Hz.

The display used in the present project is a piezoelectric display and is discussed in detail in section 3.2.2. It is a modified version of the display detailed in [56], which has 100 contact pins covering an area of $1\ \text{cm}^2$ and a bandwidth from 20 Hz to 400 Hz. Figure 3.2 shows how the piezoelectric bimorphs were arranged. The cross-section view shows 2 of the 20 stacks of 5 bimorphs. Each bimorph is connected to an L-shaped wire, the free end of which acts as the contactor. The top view shows the arrangement of the bimorph stacks to give 5 tiers, each with a ring of 20 bimorphs. The geometry of the wire links is varied between tiers to address all 100 contactor positions.




Figure 3.2: Schematic diagrams of the piezoelectric display mechanism (not to scale) showing a cross-section side view (left) and a top view (right) (replicated from [56]).

In an alternative method to movement of the actuators in the normal direction, [57] and [58] detail piezoelectric displays which use bimorphs that move laterally to the skin's surface. [57] has an array of 100 contactors at a resolution of $1\ \text{mm}^2$ per contactor and [58] has an array of 60 contactors at a resolution of about $2\ \text{mm}^2$ per contactor.

Lateral motion of contactors is also utilised in [59] but in a manner different to that of [57] or [58], where the lateral motion is generated by the piezoelectric actuators directly. Instead, the actuators stretch or contract in a direction normal to the skin's surface and this motion is converted to the lateral deformation of the skin by a 'crown' of contactors, which connect neighbouring actuators together. The amount of lateral displacement of

the contactor is therefore determined by the relative heights of neighbouring actuators, which are themselves dependent on the relative applied voltages.

3.1.3 Rheological Fluid Display

Rheological fluids transform from a liquid to a pseudo-solid state in the presence of an electric [60] or magnetic [61] field. [62] and [63] details the use of magnetorheological and electrorheological fluids in tactile displays. In these cases the resistance of the display to deformation is controlled, rather than the displacement.

Figure 3.3 shows an example of an electrorheological (ER) contactor. When no external voltage is applied, the ER fluid remains in a liquid state and, when the users finger is pressed onto, or dragged across, the flexible membrane, the user will feel little opposing force. When the external voltage is applied, however, the ER fluid in the region between the conductive rubber membrane and the Aluminium element will experience an electric field and so will stiffen into a pseudo-solid state. Now, when the users finger is pressed into, or dragged across, the membrane, they will experience greater resistance to their motion in the area above the Aluminium element but little difference outside of this area.




Figure 3.3: An example of an electrorheological actuator (recreated from [63]).

3.1.4 Electroactive Polymer Displays

The actuators used in [64] are based on electroactive polymers gels, which make a bending motion by internal ionic motion when an electric field is applied between the surfaces. The display constructed had an array of 55 actuators, separated by 1 mm in width and 1.2 mm in length, with a bandwidth of up to 100 Hz. Figure 3.4 shows a schematic of an electroactive polymer actuator that is used in [64]. The electroactive polymer in this actuator is a composite of PFS (PerFluoroSulfonic acid) and thin Platinum surface layers.

A limitation of this display technology is that it requires moisture within the polymer gel for the actuators to respond. Attempts to address this limitation have been made in [65], where a protective coating was used to maintain the hydration within the polymer and extend the operational lifespan, from a few minutes in dry conditions, to about four months.




Figure 3.4: A schematic of an electroactive polymer actuator (recreated from [64]).

3.1.5 Ultrasound Display

As opposed to the direct contact method of producing tactile sensations on the skin, [66] illustrates an alternative non-contact method of producing the sensation using airborne ultrasound (figure 3.5). The display uses ultrasound transducers to produce stress fields in three-dimensional space using the phenomenon of acoustic radiation pressure. The prototype that is described in [66] has a wide frequency range (up to 1 kHz) but is limited by its relatively low spatial resolution (20 mm).




Figure 3.5: An example of an ultrasound display, consisting of an annular array of 91 ultrasound transducers packed in an hexagonal arrangement (from [66]).

3.1.6 Surface Acoustic Waves

The displays described in [67] and [68] generate tactile sensations in the skin using surface acoustic waves (SAWs), which propagate through the surfaces of the material. These vibrations may be perceived as roughness, friction, or shear sensations by the mechanoreceptors in the fingertip. These devices have the advantage of being thin with a simple structure and a wide operating frequency range, however they cannot produce spatio-variant signals.

Figure 3.6 shows the SAW display used in that study. The sensations are generated by an interdigital transducer (IDT) on a piezoelectric substrate. When an alternating voltage is applied to the IDT, a SAW is excited and propagates along the substrate. If the user were

to touch the wave directly, it would be absorbed by the viscoelasticity of the fingertip. This problem is avoided by using steel balls as an intermediate step between the substrate and the user's finger, allowing the SAW to be transmitted faithfully to the fingertip.




Figure 3.6: A schematic of a surface acoustic wave tactile display (recreated from [67]).

3.1.7 Shape Memory Alloy Displays

References [69], [70], and [71] use shape memory alloys (SMAs) - materials whose shape is dependent on the temperature - in a few different methods to transfer the movement to the pins. However, SMAs have a slow response due to the need for cooling between heating cycles. [69] uses a display with 24 contact pins, at 2 mm element spacing. The alloy is used to drive levers that raise and lower the contact pins (figure 3.7) resulting in a relatively large maximum amplitude (3 mm) but a bandwidth of only 10 Hz. A 64 element display covering an area of 8 cm² is used in [70]. The contact pins are pushed up by springs and varied in height by the alloy, which pull the pins downward as they contract. The result is a display with a low bandwidth of between 1 Hz to 3 Hz. [71], which has a wider bandwidth of up to 30 Hz, features a display with a line of 10 contactor pins at 2 mm spacing. The high bandwidth is achieved with the use of a recirculating bath of water to aid in the cooling of the SMAs. The way the contactor is constructed is such that pin is raised as the SMA is shortened, as opposed to [70] where the shortening lowers the pin.

3.1.8 Electrocutaneous Display

Electrocutaneous displays, such as those in [72] and [73], stimulate the touch receptors of the skin with electrical impulses. Each contactor consists of an electrode 'pin', surrounded by an insulating air gap. An additional metal plate serves as the return path for the current. When a connection is made across the air gap by skin, an electrical impulse travels across the skin and a touch-like sensation is felt by the user. Figure 3.8 shows a cross-section of an electrocutaneous contactor.




Figure 3.7: A side view of an example shape memory alloy actuator (recreated from [69]).




Figure 3.8: A schematic of an electrocutaneous contactor (recreated from [73]).

The electrodes are compact, allowing for a small inter-electrode spacing (center-to-center distance of 2.54 mm and 2.45 mm, respectively). This compactness allows arrays to be made with a range in the number of contactors, in a wide variety of configurations. However, the sensation is unpleasant for some users, because of variation in skin resistance, the current of the display needs to be carefully controlled to avoid pain or injury.

3.1.9 Pneumatic Display

The display described in [74] is an array of 16 pneumatic contactors covering 15 mm wide by 15 mm long that can generate vibrations of 20 Hz to 300 Hz. The display is also designed to impart lateral motion using pneumatic actuators mounted to each lateral face of the array to replicate friction/drag forces. Similarly, [75] uses pneumatic actuators in a 5 by 5 array covering a combined area of 25 mm². However the display only has a bandwidth of up to 5 Hz.

Figure 3.9 shows the pneumatic circuit that drives this form of actuator. The circuit can be divided into two sections that address different aspects of the tactile signal that is output by the pin. The effects of static contact and/or low frequency textures during dynamic contacts (up to 20 Hz) is controlled by the pressure module, which consists of the filling valve, the venting valve, and the pressure sensor. Higher frequency vibrations (up to 300 Hz) are controlled by the texture generation module (texture valve).

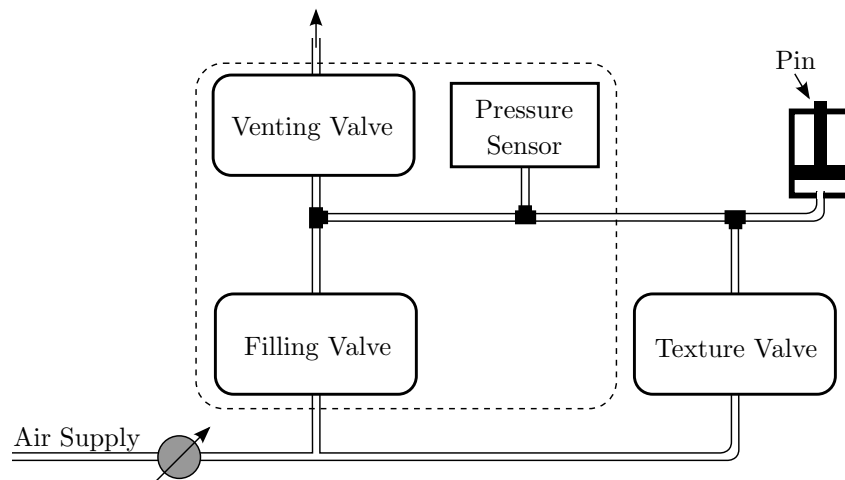


Figure 3.9: Schematic of the pneumatic circuit associated with an pneumatic actuator (recreated from [74]).

3.1.10 Electro vibration Display

Another method for generating touch sensations is through the use of electrovibrations, as explored in [76] and [77]. Here, the display features a smooth metal surface with a low current alternating potential which is experienced as a tactile sensation as the fingertip is drawn across the surface. The current involved in this setup is small (a few microamperes), so the sensory nerves associated with touch perception are not stimulated by the electrical signals directly. Instead, the tactile sensation is generated by the electrostatic attraction

that exists between the metal surface and the conductive layers of the skin. Relative motion of the skin to the surface is required because the shallower, non-conductive layers of the skin contain no mechanoreceptors and so do not detect the compression due to the electrostatic forces. Instead, the tangential motion induces a shear force to deform the skin, which is detected by the mechanoreceptors.

Figure 3.10 illustrates the electrode-skin interface of an electrovibration device. The system can be considered as a parallel-plate capacitor, with the electrode and the conductive skin tissues each serving as an electrode, and the non-conductive skin and coating layers each serving as a dielectric. Thus the electrostatic force that the user experiences will be dependent on various factors, including the thicknesses and relative permittivities of the insulators (skin and coating) and the total potential between the two electrodes.




Figure 3.10: Schematic illustration of the electrode-skin interface (recreated from [77]).

Some prototype devices that use this electrovibration effect are discussed in [78] and [79]. Displays that use this effect can be high resolution with a large number of contactors (180 contactors with 2.54 mm center-to-center spacing in [79]) with the electrovibration effect being potentially detectable over a wide range (5 Hz to 1000 Hz, [77]).

3.1.11 Shear Force Display

In electrovibration displays, electrostatic forces between the electrode and the skin induce shear forces in the skin. There are, however, other methods for inducing shear forces in the skin, such as in [80] and [81]. With these devices, as illustrated in figure 3.11, a block is moved under the fingertip to impart shear forces to the fingertip. In [80] the block is moved under the fingertip with a constant normal force, producing a net zero friction force. In [81] however, as illustrated in figure 3.11, the design allows for a non-zero net force through the use of a squeeze film of air, which can reduce the amount of shear force imparted to the fingertip, and so the perceived friction, during the motion of the shear block that is counter to the desired sensation.

Whereas most tactile devices attempt to recreate the perceived roughness of the surface, the shear force devices described here attempt instead to replicate the friction of the




Figure 3.11: The ShiverPad generating a net rightward force (recreated from [81]). (Left) The pad is moved right creating a rightward impulse on the finger. (Right) The squeeze film is turned on and the low friction pad is moved left to prepare for the next rightward impulse.

surface. They have can a range of bandwidths (~ 5 Hz in [80] versus 20 Hz to 100 Hz in [81]) but are typically limited to a single contactor on the fingertip.

3.2 The Tactile Rendering Setup

The following details the tactile renderer that is used in this project. The tactile sensation of the virtual textiles is conveyed to the user through a 24 pin piezoelectric tactile display which is driven by electrical signals generated by driving electronics. The amplitudes of the electrical signals that the electronics send to the display is determined by calculations in the tactile rendering software, using spatial frequency spectra of the real textile in combination with the current position and velocity of the display.

3.2.1 Principles of Operation

The fundamental principle behind the renderer that was developed during the HAPtic sensing of virtual TEXTiles (HAPTEX) project was to recreate the normal displacement of a set of exploratory points as they move across the profile of a surface, as shown in figure 3.12. This can be envisaged as representing the displacement of the skin on an exploring fingertip. Therefore, if we can recreate these same displacements on a finger that is making corresponding exploratory movements in a virtual environment, the sensations experienced by the fingertip would be those of the original surface profile.




Figure 3.12: As the exploratory points move across the surface, they are displaced in a direction normal to the surface.

The physical profile of the surface that might be measured by, for example, optical means, is not the surface profile used in this case. Instead, the surface profiles used by the HAPTEX renderer originate from the Kawabata Evaluation System for Fabrics (KES-F),

which uses a specially designed probe and filtering system to extract the surface profiles (see section 3.2.5). The probe of the KES-F approximates to an artificial finger which is used to extract aspects of the surface profile that are associated with the subjective “feel” of the surface from the “true” surface profile. Different profiles are recorded and used for the four primary directions on the textile surface - left, right, up and down - so as to maintain the anisotropy of the surface of real textiles. Exploratory movements that are made within the two-dimensional virtual environment are recorded as separate components of movement along the horizontal (left/right) and vertical (up/down) axes. Each of the components is separately processed by the rendering software, using the direction of the motion in that component to identify the recorded surface spectrum appropriate to the movement (up or down in the vertical (warp) direction, left or right in the horizontal (weft) direction). The results of the software’s processing of the separate components are then combined by a root of the sum of the squares (root-sum-square) to define the output of the renderer.

The tactile display features a set of 24 contactors that can each be controlled independently by channels within the tactile rendering software. The software envisages each channel as an exploratory point moving across the Kawabata profiles with appropriate position and velocity. The Kawabata profiles associated with a textile surface are stored as spatial spectra (i.e, decomposed into different spatial frequencies), which vary point-by-point, with a 1 mm-by-1 mm resolution over the surface (figure 3.13). The spatial spectra that are used to calculate the output from the display are defined by the position of the exploratory point on the virtual surface. The separate components of the points’ velocity are used to convert the spatial frequencies of the relevant spectra into temporal frequencies, using equation 3.1.




Figure 3.13: (Left panel) A representation of the tactile workspace, the virtual surface is specified on a 1 mm \times 1 mm matrix (grey circles). The virtual fingertip is represented by 6 \times 4 points at a 2 mm spacing, corresponding to the contactors on the tactile array (black circles). (Right panel) The spatial spectrum is mapped onto the temporal frequency according to velocity of exploration in direction i .

$$f_i = v_i \xi_i \quad (3.1)$$

where f_i is the temporal frequency, v_i is the component of velocity, and ξ_i is the spatial frequency (the inverse of wavelength) in direction i , where i is either x or y .

It would be possible to stimulate the fingertip with any mix of the temporal frequencies defined by equation 3.1 and, ideally, the tactile display would present broadband stimuli in the tactile frequency range of 20 Hz to 500 Hz, containing all the frequency components that would be present in the “real” situation, in order to achieve a “true” representation of a virtual tactile surface. While the display and the driving electronics used in the HAPTEX renderer are capable of generating signals over this wide range of frequencies, the resolution of tactile frequency is poor – the jnd for frequency is around 30% for single-channel stimulation [82] and is expected to be worse than this for the spatially varying signals on an array. This suggests the possibility of using a low-resolution representation of the tactile spectrum based on only a few, widely spaced, frequency components, since such a reduction of spectral resolution is expected to have only a limited effect on the perceived tactile sensation.

The results of [83], using a single-channel system, suggest that the spectrum can be reduced to only two sine wave components, providing these are chosen to separately target Pacinian (P) and non-Pacinian (NP) receptor populations. This two-component strategy [84] was adopted for the HAPTEX system (and for the present study), by using tactile signals which are a mixture of sine waves at frequencies of 320 Hz, targeting the P receptors, and 40 Hz, targeting the NP receptors. Anecdotal reports suggest that the user is unaware of the discrete nature of the tactile spectrum and instead reports that the 320 Hz and 40 Hz signals combine to give a realistic overall sensation of surface roughness, with the 40 Hz component suggesting the gross topology of the surface and the 320 Hz component suggesting a “hairy” or “tickly” aspect of the virtual textile. The amplitudes of the component sine waves, A_{320} and A_{40} for the 320 Hz and 40 Hz respectively, may be calculated by applying bandpass filters that are centred on the appropriate frequencies to the full temporal spectrum of the surface, as illustrated in figure 3.14.

These two-component tactile signals are, in a sense, analogous to the three-component video signals which are tailored to match the three channels of colour vision in the eye [84]. This is not a perfect analogy, however, as while the combination of visual signals can produce a near perfect perceptual match to most colours, the combination of the 40 Hz and 320 Hz tactile signals produces only an approximation to the sensation of intermediate frequencies because the periodicity to which touch receptors are sensitive [85] is not accurately represented.

This mixing of vibrotactile cues can also be exploited for other effects, such as in [86] and [87], where separate actuators at different locations were used to produce the perceived sensation of a travelling wave. In [86], a series of actuators are attached to the arm of the subject and, when two of the actuators are activated, the subject perceives an illusionary sensation at a location that is between the two actuators. The position of the illusionary sensation is dependent of the relative intensities of the two actuators, with the perceived location being shifted towards the actuator with the greater intensity, and so, by continuously adjusting the relative intensities of the actuators, the illusionary sensation is perceived to travel along the length of the subject’s arm. Similarly, in [87] a set of

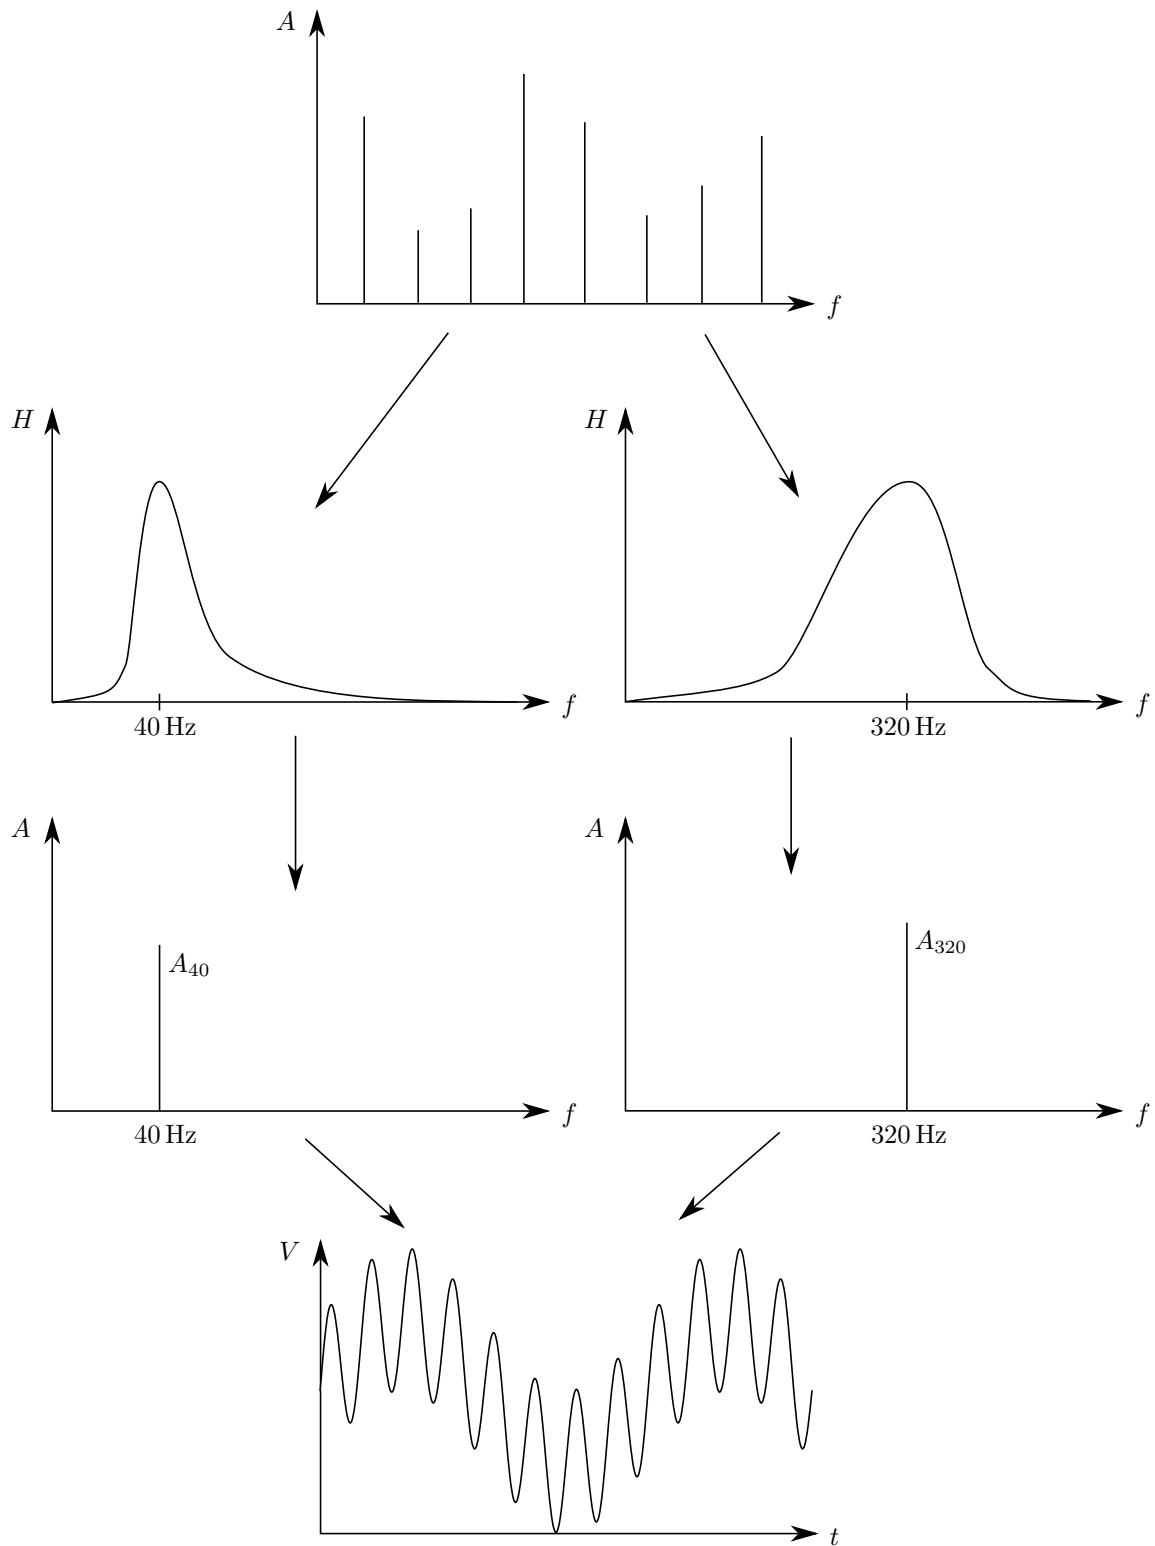


Figure 3.14: The spectrum of temporal frequencies is reduced to components at 40 Hz and 320 Hz, and the output is specified by their amplitudes, A_{40} and A_{320} .

actuators are used, but they are instead connected to an independent device, rather than the subject directly, and the relative activation times of the actuators adjusted, as opposed to the intensity. In this case, the location of the illusionary sensation is shifted towards the actuator with the later activation time. A series of activations in rapid succession, with differing relative activation times, give the illusion of a wave that is travelling across the surface. The user can perceive this travelling wave by touching the surface of the device at a location between the two actuators.

At any given instant, the rendering software calculates the amplitudes A_{40} and A_{320} for each of the 24 exploratory points. These amplitudes are updated every 25 ms and change with the position and speed of the virtual fingertip as it is moved over the surface, with signals from the resolved weft and warp calculations being combined into a single output for each point.

Each of the display's contactors is driven by a drive signal which is a mixture of 40 Hz and 320 Hz sine waves, the amplitudes of which are defined by A_{40} and A_{320} respectively. Each contactor therefore delivers a tactile stimulus, the intensity and spatial content of which reproduce the intensity and spatial content of the displacement of the corresponding exploratory point as it moves over the Kawabata profile.

3.2.2 Tactile Display

For this project, the display used is developed from [56] which was developed by the University of Exeter Biomedical Physics group as a psychophysical research tool for the passive presentation of stimuli to the subject (i.e., no movement made by the subject). The stimuli are generated using piezoelectric bimorphs to convert electrical signals to mechanical movement of the stimulator pins.

The piezoelectric effect occurs when a piezoelectric material is mechanically deformed, causing a potential difference to be induced when the centres of the positive and negative ions are displaced from each other (figure 3.15). The inverse effect is used here, where a potential difference is applied across the material in the piezoelectric bimorphs, resulting in mechanical deformation of the material, to move the contactors.

In this application, the piezoelectric effect is exploited using a piezoelectric bimorph, which features two layers of a piezoelectric material that are wired up in such a way that an applied voltage causes one layer to expand while the other layer contracts, causing the bimorph to bend in a direction perpendicular to the applied voltage. With one end of the bimorph held into a fixed position, the magnitude of this effect is proportional to the square of the unsupported (free) length of the bimorph, i.e. a longer free length of bimorph gives a larger amplitude of displacement. The tactile display employed here uses commercially available bimorphs that have a thickness of 0.6 mm, 2.1 mm width, a free length of 32 mm, a maximum driving voltage of 150 V, producing a maximum displacement amplitudes of 100 μm . As the blocked force (the maximum force an actuator can generate if blocked by an infinitely rigid restraint) of the bimorph is dependent upon the width of the bimorph, two bimorphs were used per actuator, connected together, in order to convey the sensations required.




Figure 3.15: Schematic drawing of the Piezoelectric effect. The undeformed material (left) has no potential difference while the deformed material (right) has a difference induced by compression by force F .

The bimorphs are Stripe™ Actuators which have been purchased from APC International, Ltd. and use Lead Zirconate Titanate (PZT), an inorganic compound chemical formula $(\text{Pb}[\text{Zr}_x\text{Ti}_{1-x}]\text{O}_3 \ 0 \leq x \leq 1)$, as the piezoelectric material. PZT is a ceramic material that is a white solid that is insoluble in all solvents. In addition to being piezoelectric, PZT is also pyroelectric and ferroelectric with a large dielectric constant, ranging from 300 to 3850 depending on the orientation and doping of the material. These bimorphs are covered by a varnish layer to electrically insulate the surface, as well as protect them from environmental hazards, such as dust and humidity.

The individual bimorphs used in this case are in what is known as the parallel configuration, where the polarisation of the layers are aligned. Electrical connections are made with the two outer faces of the layers and are connected so that they each receive the same voltage while a third, central electrode receives a different voltage (figure 3.16). The result is that both layers receive the same voltage, but in opposite senses (the electric field is aligned with the polarisation in one layer and is anti-aligned in the other).




Figure 3.16: Schematic of a parallel configuration Piezoelectric bimorph.

The alternate configuration, known as serial configuration, is where the polarisation of the two layers are anti-aligned and only the electrical connections are made with the two outer faces (no third, central connection). However, in this case, each outer layer connection receives a different voltage, resulting in the same electric field throughout the entire bimorph. Again, the result is that one layer's polarisation is aligned with the field and the other's is anti-aligned. This configuration has the advantage over the parallel

configuration of simpler wiring (two connections in serial vs three in parallel) but at the cost of twice the voltage being required to produce the same displacement.

Figure 3.17 shows the frequency response obtained from a piezoelectric bimorph actuator of the same material and configuration as those used in this display, as measured by a miniature accelerometer, both with and without the load cause by the skin, which adds a significant stiffness (100 N m^{-1}) and resistance (0.1 N s m^{-1}) to the system, although the additional mass (10^{-5} kg) can be considered negligible. The effect of the load is to increase the resonant frequency of the bimorph but the narrowness of the resonance is reduced. As the tactile stimuli consists only of a superposition of 40 Hz and 320 Hz sine waves, the bimorphs were chosen because their principal resonance lies between these two operating frequencies, as illustrated by the vertical lines on figure 3.17.

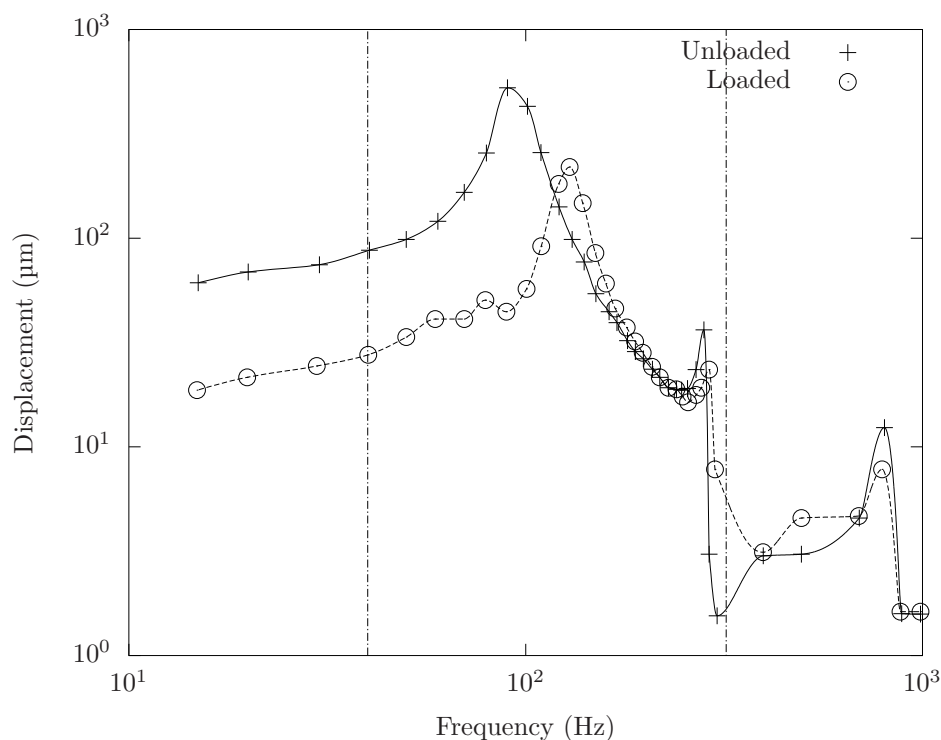


Figure 3.17: Measured frequency response of a piezoelectric drive element with and without the mechanical load presented by the skin (recreated from [84]). Vertical lines indicate the 40 Hz and 320 Hz operating frequencies.

The display detailed in [56] is a 100 pin square array covering a total area of 1 cm^2 (figure 3.2). Each contactor therefore covers an area of 1 mm^2 , which is approximately the spatial resolution of the fingertip (see section 2.2.2 or [21]). The investigation in [88] however demonstrates that it is difficult to distinguish between stimuli presented at 1 mm and 2 mm resolutions. Therefore the display used in the project is a 24 contactor variant at 2 mm^2 resolution in a 6-by-4 array (the Exeter tactile display, figure 3.18), resulting in a display which covers the same display area but with reduced manufacturing and computational cost of the 100 contactor version. An additional advantage of the 24-contactor device is the reduced mechanism compared to the 100-contactor device, which allows the display to be moved with the finger during the active surface explorations.

The compactness required of the display array means that the bimorphs cannot directly




Figure 3.18: The tactile display used in this project (Exeter tactile display, left panel) and the arrangement of the static and dynamic contactors (right panel, empty and filled circles respectively).

interact with the skin without interfering with each other. Therefore, the stimuli are conveyed from the piezoelectric bimorphs to the surface of the skin through contactor pins that were made of 0.6 mm brass wire. This wire was glued into a rectangle of a thin polystyrene sheet that provided a large area for the glue joint, insulated the wire from the bimorphs, and served to join the free ends of the two bimorphs into a single actuator. The wire was bent through a 90° angle beyond the end of the actuator and the end of the wire was passed through a base plate that served to keep the ends of the wire, which are the display's contactor, in the correct geometric location.

The use of the wire allows the contactors to be packed much more tightly than the bimorphs alone would allow, and for the actuators to be arranged in ways other than that implied by the display array. The limiting condition is that the contactor wire must reach from the end of the actuator to the base plate. In this case the actuators for a row of four contactors is divided into two column of two actuators, with the outer most contacts being driven by the lower row of actuators.

In addition to the active twenty-four display pins, there are an additional intermediate thirty-five static pins that were incorporated into the design in an attempt to make the surface of the array feel smoother and to provide a stiffer surface. The latter provides gross mechanical forces to the fingertip, i.e. something the users are able to rest their fingertip on so as to not distort the base levels of the contactors/actuators but still allow for the full range of sensation from the contactors.

3.2.3 Driving Electronics

Each of the 24 actuators in the tactile display is independently driven by a signal which is a mixture of 40 Hz and 320 Hz sine waves, the amplitudes of which are specified by the rendering software based on the description of the surface texture and the exploratory movements of the user (position and velocity) (as explained in section 3.2.1). The driving

signals are generated by a set of driving electronics (figure 3.19), based on the data received from the rendering software, running on the connected PC. Figure 3.20 shows a block schematic of the drive electronics, consisting of a Universal Serial Bus (USB) controller to receive and process the tactile information, a data bus to distribute the output signals to the 24 channels, each of which has a digital-analogue converter (DAC) and a drive amplifier.




Figure 3.19: The driving electronics for the tactile display.

The electronics are controlled via the USB controller, which is connected to the host computer. USB is an extremely popular interface for connecting peripheral hardware to computers, to the point that nearly all modern personal computers have at least one USB connector. Therefore, by connecting the device via USB, the setup can be used on the broadest range of devices.

The USB controller used for this setup is a CY7C64713 integrated circuit (IC) produced by the Cypress Semiconductor Corporation, which includes a USB transceiver and an Intel 8051 microprocessor with 16 kilobytes of on chip random access memory (RAM), and is responsible for receiving, processing and forwarding the data to the DACs. The USB controller is hosted on the board show in figure 3.21, along with the data bus and the variable voltage supply.

The data that is received from the rendering software consists of the amplitudes that have been calculated for each component of the output signal, for each channel, totalling forty-eight 8-bit numbers: 24 channels with 2 component amplitudes per channel. New amplitudes are calculated and transmitted to the USB controller so that the amplitudes of the 24 output signals can be updated every 25 ms.

The amplitude data is processed by the firmware; software which runs on the USB controller. This receives the amplitudes from the rendering software, converting them to the drive signals that run the display. Each of these is a superposition of sine waves $f(\omega_1 t)$ at 40 Hz and $f(\omega_2 t)$ 320 Hz. The drive signal g_n for channel n at time t is defined from

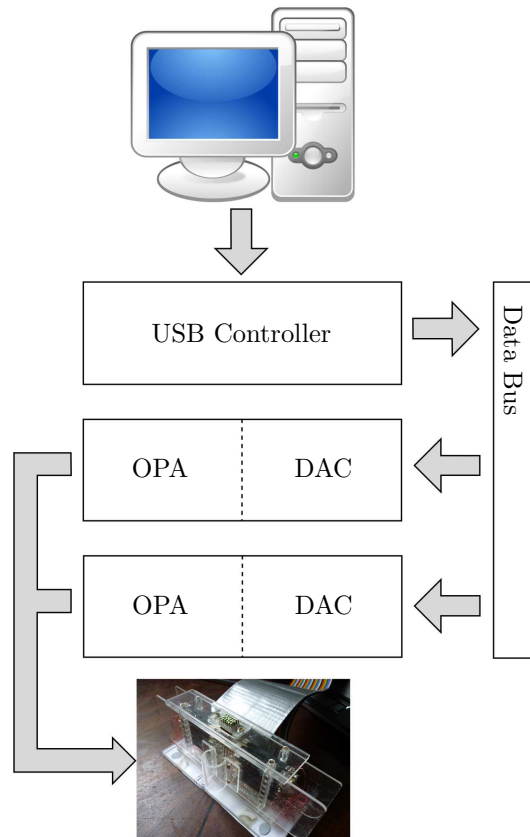


Figure 3.20: A schematic of the operation of the drive electronics (recreated from [1] with the computer graphic from Wikimedia Commons).




Figure 3.21: The board that hosts the USB controller, data bus and variable voltage supply (recreated from [1]).

amplitudes $A_{n,1}$ and $A_{n,2}$ according to equation 3.2. (For this setup, the periodic functions f are sine waves, but the shape of the wave is defined by the rendering software and, as such, can be defined to be one of any number of possible functions.)

$$g_n(t) = A_{n,1}(t)f(\omega_1 t) + A_{n,2}(t)f(\omega_2 t) \quad (3.2)$$

When the values of the signals at time t have been calculated for the channels, they are transmitted to the DACs via a bus system. The data bus, which can be seen in the middle of figure 3.21, receives the new value (specified by 8 data lines) of the signal associated with a particular channel and writes the new signal value into the DAC buffer appropriate to the channel being updated (specified by 8 address lines). When the buffers for all of the channels have been updated, the signal is given by the USB controller for the new signal values to be used, ensuring that the DACs update their outputs simultaneously.

The address of each of the remaining boards is specified by the address logic circuit, which can be seen in lower left of figure 3.24, which consists of a 6 bit dual line-in package switch. This switch configures the address range of the board and a comparator IC compares the address of the board to the address of the buffer being written into.

To avoid problems with noise from the rest of the system, and to amplify the signal from the USB controller, the data bus galvanically isolates the USB controller from the rest of the system using an optical coupler. This is achieved using the circuit depicted in figure 3.22 which is part of an 6N137 IC. An inverter is inserted between the output of the USB controller and the input of the optical coupler because the optical coupler inverts the signal, and the output current of the USB controller is insufficient to drive the LEDs inside the optical coupler.




Figure 3.22: The circuit galvanically isolating the USB controller from the rest of the system.

Also on the USB controller board is a variable voltage supply (left side of figure 3.21), which can be used to adjust the global gain of the system by setting the reference voltages for the DACs ($-V_{ref}$ and $+V_{ref}$), which can be configured by the USB controller. A separate DAC exists for setting reference voltages, independent of the DACs used to convert the drive signals, which sets the reference voltage depending on the digital input to the DAC.

In this case the DAC circuit is a AD558 IC. This output is inverted twice by the circuit depicted in figure 3.23, which is implemented by a L272 IC.

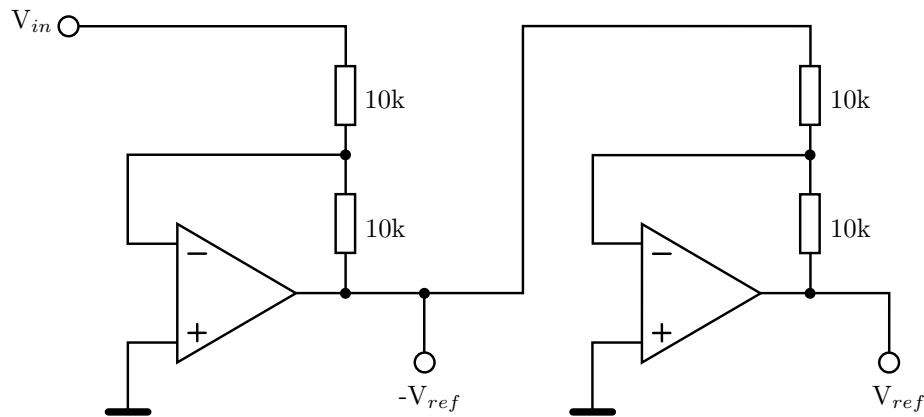


Figure 3.23: Amplification stages of the variable voltage supply.

In addition to the board hosting the USB controller, there are a series of boards that receive new digital amplitudes for the output signal. This signal is then converted to an analogue voltages and amplified into signals that can drive the tactile display. There are a total of six of these boards within the driving electronics, one of which is shown in figure 3.24, with each one serving four of the channels of the tactile display.




Figure 3.24: The board hosting the DAC, address logic, and the amplifiers (recreated from [1]).

Each board has a DAC that receives the data from the bus and stores it in the buffer until the load signal is given to use the new value as the output. To buffer the digital signals and convert them into an analogue output, the MAX5100 IC is used. This provides four independent channels (hence why each board serves four channels). The DACs can be seen in the middle left of figure 3.24.

The output values of the DAC are discrete, rather than continuous. Therefore, a passive, first order, low-pass RC filter, as depicted in figure 3.25, is used to reduce the imaging effect cause by this quantisation.




Figure 3.25: Passive, first order low-pass RC filter.

The cut-off frequency of the filter, f_c , is defined by equation 3.3 with resistor and capacitor values in this case of $33\text{ k}\Omega$ and 4.7 nF respectively, resulting in a cut off frequency of approximately 1 kHz , which is above the maximum frequency content and below the sampling rate of the signals.

$$f_c = \frac{1}{2\pi RC} \quad (3.3)$$

In order to produce strong tactile sensations, the piezoelectric bimorphs require relatively high voltages (up to 100 V). To increase the output level to what is required by the bimorph actuators, the signal goes through the drive amplifiers which use a bridge design running from $\pm 24\text{ V}$ rails. The amplifiers can be seen on the right side of the board in figure 3.24.

The amplifiers used in this case are the OPA551 IC, with the amplifying circuits shown in figure 3.26. The first of the two amplifiers is a non-inverting amplifier (left side of figure 3.26). The exact amplification factor of this amplifier can be adjusted with a potentiometer. The output from the first amplifier is inverted, with respect to the ground voltage, by a second, inverting amplifier which has a fixed amplification factor of -1 (right side of figure 3.26). The result of this amplification is an output signal with a maximum output range of about 90 V , peak-to-peak.

The resulting output of the drive electronics is a superposition of the two sine waves at the required frequencies with the amplitudes that were calculated by the rendering software. An example of such an output is given in figure 3.27.

3.2.4 Rendering Software

The amplitudes that the firmware processes into the digital driving signals are determined by rendering software running on the host computer. To calculate them, the software uses a combination of information about the textile's surface, in the form of a spectrum of the surface profile, as well as band-pass filters for the different frequency components, and the

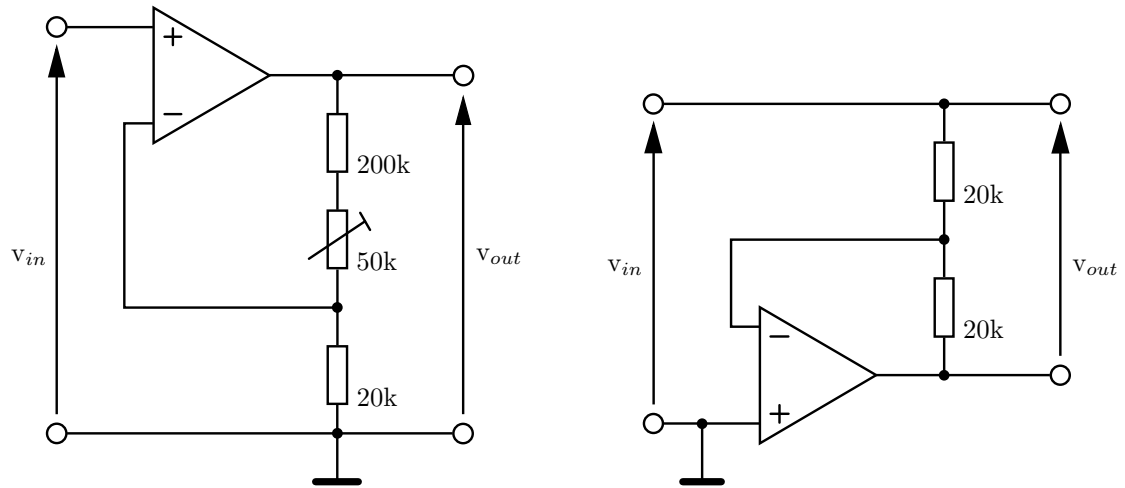


Figure 3.26: The two amplitude stages of each channel. A non-inverting amplifier with an adjustable gain factor (left panel) and an inverting amplifier with gain factor -1 (right panel).




Figure 3.27: An example of an output signal, as generated by the driving electronics.

velocity and position of the display, as taken from an associated input device. This section contains a detailed description of the rendering software; for a more general overview of the principles behind the software, see section 3.2.1. The software used in this project was written in the C++ language and run under a Ubuntu distribution of Linux.

The information about the surface of the textiles is stored as a library of texture files, each one representing either the front or back surface of a single textile. Every file contains a set of five matrices, four of which define the surface profile through a series of spatial spectra derived from Fourier transformations of the recorded surface profile along a particular direction (up/down/left/right). Each column of the matrix represents the Fourier transformation of a particular 4.37 mm section along the direction of the profile. A Hamming window, shifted by 1 mm increments along the profile between transformations, is used to remove the discontinuities between the beginning and end of the sampled data that would be associated with a rectangular window.

The original measurements from the KES-F system cover a total of 20 mm along the textile surface therefore, with each transformation covering 4.37 mm of the profile and a 1 mm shift between each transformation, each matrix therefore defines the surface profile for a total of 16 mm of textile along a particular direction. These 16 mm of textile are then repeated over the full length of the textile. Each spectrum contains the amplitude values for 50 spectral lines, with the spatial frequencies that these lines represent being stored in the fifth matrix in the file. A full description of the surface profile recording and the creation of the texture files is given in 3.2.5.

The rendering software was developed within a framework consisting of four threads that are executed in parallel, as shown in figure 3.28. The network and position tracking threads are responsible for acquiring the input data which is used by the graphics and tactile rendering threads to produce their respective outputs.

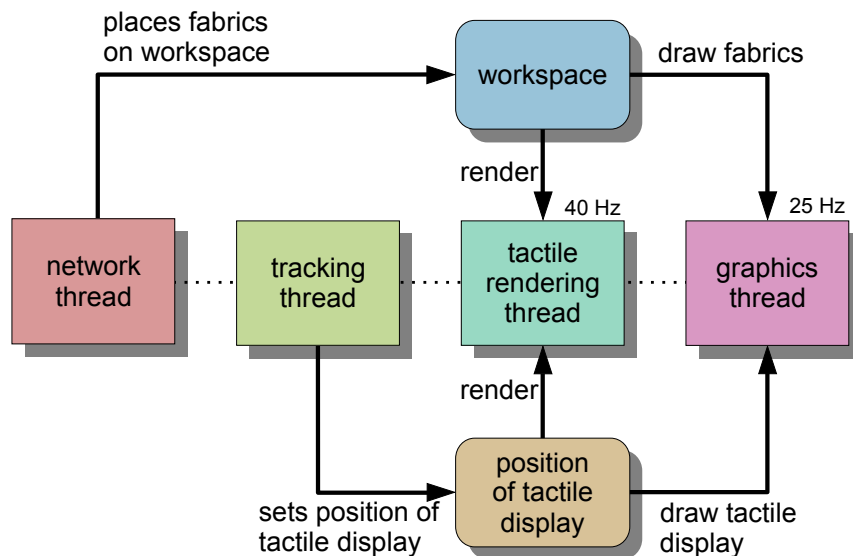


Figure 3.28: The tactile rendering framework, showing the four threads within the software (recreated from [1]).

User requests, such as changes to the gain factors or the inclusion of new textile in the workspace, are received by the network thread. This thread listens on the network port

for these requests and then places textiles into the virtual workspace accordingly. This thread allows the software to be controlled by a remote computer, as well as the local one.

The tracking thread continuously receives positional data from the device that is tracking the tactile display's position, which is fed into the tactile rendering algorithm. In this project, the tracking device was a commercially available graphics tablet, which was connected to the display by the stylus being held by a cylinder of material to the outside of the display, as seen in figure 3.18, with the combined stylus and display being moved together across the tablet's surface. In addition to the graphics tablet, a standard computer mouse has also been used to give position information for the two dimensional application of the renderer, with the main difference between these two devices being that the mouse provides the renderer with relative position of the display, whereas the tablet provides the display with absolute position. The renderer has also been connected to a force feedback device for use of the renderer in a three dimensional workspace, i.e. the integrated HAPTEX system.

As well as the position of the display, the rendering algorithm requires the velocity of the display, which the tracking device does not, itself, provide. Instead, a Kalman filter is used to compute the velocities that correspond to the measured positions, as well as reduce the noise on the measurements. What follows is a brief description of the Kalman filter. For more detail, the interested reader is invited to read [89].

The Kalman filter is based on the idea of estimating the state $x \in \mathbb{R}^n$ of a discrete-time controlled process, as described by equation 3.4, by predicting the state of the next time step, then correcting the predicted state using using the corresponding measurement of the process, as depicted in figure 3.29.

$$x_k = Ax_{k-1} + Bu_k + w_{k-1} \quad (3.4)$$

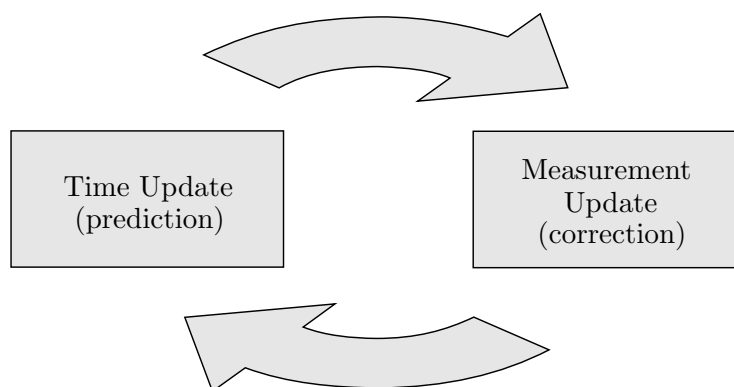


Figure 3.29: The Kalman filter uses a predictor-corrector feedback loop (recreated from [1]).

For position tracking, the state, x , is chosen to be

$$x = \begin{pmatrix} p \\ v \\ a \end{pmatrix} \quad (3.5)$$

where p , v and a denote the position, velocity and the acceleration of the state respectively. The model used assumes acceleration to be constant and thus matrix $A \in \mathbb{R}^{n \times n}$, which relates the state at time step $k - 1$ to the state at time step k , is chosen to be

$$A = \begin{pmatrix} 1 & \Delta t & \frac{1}{2}\Delta t^2 \\ 0 & 1 & \Delta t \\ 0 & 0 & 1 \end{pmatrix} \quad (3.6)$$

where Δt is the time between two time steps. The term denoted by Bu_k describes the driving force, which does not exist in the model used for position tracking. The process noise is represented by the random variable w_{k-1} . The noise is assumed to be normally distributed, white noise with a mean of zero and a covariance matrix Q , defined as

$$Q = \begin{pmatrix} \frac{1}{4}\Delta t^4 & \frac{1}{2}\Delta t^3 & \frac{1}{2}\Delta t^2 \\ \frac{1}{2}\Delta t^3 & \Delta t^2 & \Delta t \\ \frac{1}{2}\Delta t^2 & \Delta t & 1 \end{pmatrix} \cdot \text{var}(a) \quad (3.7)$$

The variance of the acceleration is estimated as 90 ms^{-2} . With this model, the influence of the user on the position of the display is modelled as the noise of the acceleration.

A measurement $z_k \in \mathbb{R}^m$ is related to the state x_k by

$$z_k = Hx_k + v_k \quad (3.8)$$

where H is the measurement model. As only the position is measured in this case, H has the form

$$H = \begin{pmatrix} 1 & 0 & 0 \end{pmatrix} \quad (3.9)$$

The random variable v_k represents the measured noise which, like the process noise, is assumed to be white, normally distributed and with a mean of zero. The covariance matrix on this variable, R , is a 1×1 matrix that is estimated to be

$$R = \begin{pmatrix} 0.001 \text{ m} \end{pmatrix} \quad (3.10)$$

In order to allow the user to navigate their way within the virtual environment, a graphical display is presented which features square areas indicating the locations of the different tactile surfaces. Additionally, there is a pointer consisting of 24 green squares, arranged in a 6-by-4 matrix, reflecting the arrangement of the tactile display. This indicates the location of the tactile display within the virtual space and shares the same motion in the virtual workspace as the real tactile display across the surface of the graphics tablet. The graphics thread in the rendering software is responsible for drawing this display and updating it every 40 ms. Figure 3.30 illustrates an example virtual environment that the user may encounter. In this case the renderer is displaying four different virtual textiles.

The virtual workspace may have overlapping textiles but, in that case, the textile that would be experienced by the user in a given area is the textile that has most recently been called.




Figure 3.30: An example of the graphical display presented to the user.

The rendering algorithm, which calculates new amplitudes from the drive signals of the tactile display based on the positional data from the tracking thread and the configuration of the virtual workspace, is called by the rendering thread every 25 ms while the software remains running. Figure 3.31 is a block diagram showing the operation of the rendering thread, beginning with the initialisation of the thread, which includes the starting of the rendering firmware and the setting of the initial gain values that are used in the calculations. After initialisation of the thread, the calculation of the new amplitudes begins, using a combination of the surface data of the relevant textile and the display's current position and velocity.

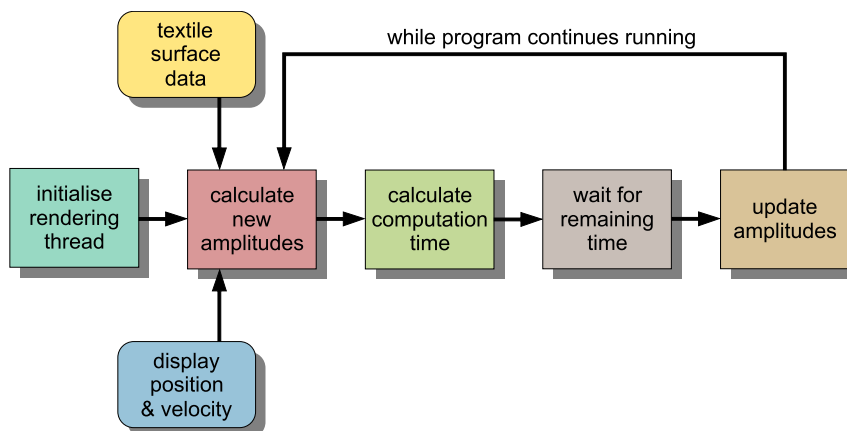


Figure 3.31: The schematic operation of the software's tactile rendering thread.

Once the calculations of the new amplitudes are complete, they are sent to the firmware for conversion into display driving signals after the remaining portion of the 25 ms section has

elapsed. This ensures that the amplitudes are updated with a frequency of 40 Hz, timing the amplitude updates so that the update occurs at the end of the 40 Hz component of the output wave, i.e. one full 40 Hz cycle and eight 320 Hz cycles are completed before the amplitudes are updated. If the calculations of the new amplitudes are not completed within the 25 ms time frame, the rendering thread does not update the amplitudes and instead completes the calculations and waits for the completion of the next 25 ms update window before updating. A while loop then returns the rendering thread to the beginning of the amplitude calculation stage for the computation of the amplitudes at the next time-step, which continues for as long as the rendering software remains running.

The amplitudes that are output for the different frequency components of the tactile display drive signal, A_{40} and A_{320} , are calculated by the rendering thread using equation 3.11, where A_i is the output amplitude associated with frequency i , either 40 Hz or 320 Hz. The calculations are done on the amplitudes of the frequency spectrum of a textile, a , that has been calculated by a Fourier transform of the surface profile for that textile that was recorded by the KES-F. As well as the spectral amplitudes, two band-pass filter functions, H_{40} and H_{320} have been created to apply weighting factors to the spectral amplitudes and thus highlight the amplitudes associated with the frequency components, i.e. those associated with 40 Hz and 320 Hz respectively.

$$A_i = \sqrt{\sum_f H_i^2(f) \cdot a^2(f)} \quad (3.11)$$

The set of band-pass filter functions are originally derived from the frequency responses associated with the mechanoreceptors that exist within the skin of the have, as described in section 2.2.4. As is mentioned in section 3.2.1, the two frequencies that have been chosen so as to target two groups of the receptors in the skin, defined as the Pacinian (P) channel and the non-Pacinian (NP) channel. Figure 3.32 shows the frequency dependent, minimum detection thresholds associated with each of these two groups. The threshold for the P channel, T_P , shows that this channel is most sensitive at ~ 300 Hz, and the threshold for the NP channel, T_{NP} , similarly shows that this channel is most sensitive at ~ 50 Hz.

To convert these threshold curves into the band-pass filters used in the tactile renderer, we begin by considering the sensation experienced by a user from an input stimulus of amplitude A , at frequency f , in the two mechanoreceptor channels. The sensation experienced from the stimulus in a particular channel will have the magnitude of the original stimulus weighted by a frequency-dependent factor inherent to the channel. This weighting factor is assumed to be related to the minimum detection threshold curves by $\frac{1}{T_{NP}(f)}$ and $\frac{1}{T_P(f)}$ in the NP and P channels respectively. The sensation experienced in the user in a channel, S_i , where i is the receptor channel (NP or P), is therefore related to the stimulus by

$$S_i = \frac{A}{T_i(f)} \quad (3.12)$$

In the renderer, an equivalent stimulus is split into two components by using the set of




Figure 3.32: Average detection thresholds as a function of frequency of the tactile receptors in the skin for Pacinian (T_P) and non-Pacinian (T_{NP}) channels used in this project.

band-pass filters that we are deriving, H , which give the frequency-dependent weighting factor of the various frequency channels. The result of applying a filter to the original stimulus signal is a single amplitude (one per filter), a_j , that will serve as the amplitude of that channel, where j denotes the frequency of the channel (40 Hz or 320 Hz), in the renderer's output signal. This amplitude is given by

$$a_j = AH_j(f) \quad (3.13)$$

The intention is that the sensation experienced by the user from the renderer is equivalent to that experienced from the original stimulus directly. We can therefore equate the sensation from the input stimulus directly into the skin, S_i as described in equation 3.12, with the sensation experienced from the renderer. Unlike the original stimulus, which was a signal of single frequency, f , and amplitude, A , the stimulus caused by the renderer is a combination of two signals, one of frequency 40 Hz with an associated amplitude, a_{40} , and one of frequency 320 Hz with an associated amplitude, a_{320} . As the stimulus from the renderer is transmitted into the skin, it too undergoes the damping that is associated with the minimum detection thresholds, as described in equation 3.12. It can therefore be said that the stimulus from the renderer that is equivalent to the stimulus in equation 3.12 can be defined as

$$S_i = \frac{a_{40}}{T_i(40)} + \frac{a_{320}}{T_i(320)} \quad (3.14)$$

Equation 3.14 may then be combined with equations 3.12 and 3.13 to give the set of

equations

$$\begin{aligned}\frac{A}{T_{NP}(f)} &= \frac{AH_{40}(f)}{T_{NP}(40)} + \frac{AH_{320}(f)}{T_{NP}(320)} \\ \frac{A}{T_P(f)} &= \frac{AH_{40}(f)}{T_P(40)} + \frac{AH_{320}(f)}{T_P(320)}\end{aligned}\quad (3.15)$$

which may be further simplified into the matrix equation

$$\begin{pmatrix} \frac{1}{T_{NP}(40)} & \frac{1}{T_{NP}(320)} \\ \frac{1}{T_P(40)} & \frac{1}{T_P(320)} \end{pmatrix} \begin{pmatrix} H_{40}(f) \\ H_{320}(f) \end{pmatrix} = \begin{pmatrix} \frac{1}{T_{NP}(f)} \\ \frac{1}{T_P(f)} \end{pmatrix}\quad (3.16)$$

As T_{NP} and T_P are known, this equation can easily be solved for the values of the band-pass filters H_{40} and H_{320} at a range of frequencies. The tactile renderer used in this project's experiments had filter functions that ranged from 1 Hz to 1 kHz.

The basic filter functions, obtained from the solution to equation 3.16, were then modified as detailed in section 3.4.4, resulting in the filter functions there were implemented within the rendering strategy (equation 3.11) used the experiments detailed in this work (figure 3.33).




Figure 3.33: The filter functions H_{40} and H_{320} .

The amplitude calculations also consider a number of additional factors, such as whether the display is travelling in the positive or negative direction along the x and y axes, utilising the appropriate surface profiles. The recorded surface profiles include information for 16 mm of textile in each of the four principle directions and so the position p_n determines which of the set available spectra to include in the calculation. When the display has travelled across the 16 available spectra of a direction, it cycles back to the first spectrum (i.e. $n = 17 \rightarrow n = 1$).

Additionally the velocity of exploration of the surface is included in the calculation from the texture files, where the frequencies of the spectra are stored as spatial frequency. Thus, the frequencies f are determined by multiplying the spatial frequencies by the velocity component in that direction (equation 3.1). Thus the final amplitude calculation for channel i is defined as

$$A_i = \sqrt{\sum_{n=x;y} \sum_{f=1 \text{ Hz}}^{1 \text{ kHz}} a_i^2(f, p_n) \cdot H_i^2(f)} \quad (3.17)$$

where the first summation combines the signals derived from movements in the x and y directions.

An additional consideration on nature the output drive signal is the spectral balance between the two frequency components, i.e. the ratio of the amplitudes A_{40} and A_{320} . Figure 3.32 shows that, even at their most sensitive, the mechanoreceptors associated with the NP channel have a higher minimum detection threshold than the P channel. It can therefore be said that, in order to produce the equivalent level of experienced sensation, the signal into the NP channel must be stronger than that into the P channel.

This spectral balance is handled within the rendering software by a gain factor for each channel. The amplitudes calculated from the rendering strategy are multiplied by the relevant gain factor (value of between 0 and 1) before being transmitted to the driving electronics. The specific details of the spectral balances used in the particular experiments are available in chapter 6, with an overview of all the balances used in section 3.4.8.

3.2.5 Texture File Creation

Each surface texture is stored as a texture file that defines that particular surface within the tactile rendering framework. In order to generate these files, it is necessary to convert the surface features of the physical textile into a digital format. A number of standardized testers exist, allowing textile researchers to evaluate textiles based on a variety of properties and obtain characteristic values for them. These testers can be used to extract the surface profile from the textiles and convert them into the required digital information. Two of these standard systems are the Kawabata Evaluation System for Fabrics (KES-F) [90], used to obtain the texture files for the present work, and Fabric Assurance by Simple Testing (FAST) [91].

A total of 54 textiles were used in the present study (table A), which were the same as those used in the HAPTEX project. During the selection process, three criteria were identified as relevant considerations for textile selection. These criteria, as shown in figure 3.34, were fibre type (natural, synthetic, etc.), structure (woven, knitted, etc.), and dimension (thickness, density, etc.). A total of three sets of textiles were selected over the course of the HAPTEX project.

The initial set of textiles (textiles 1 - 32) was selected to give the greatest variety in terms of the above criteria, with each group represented in the selection, as described

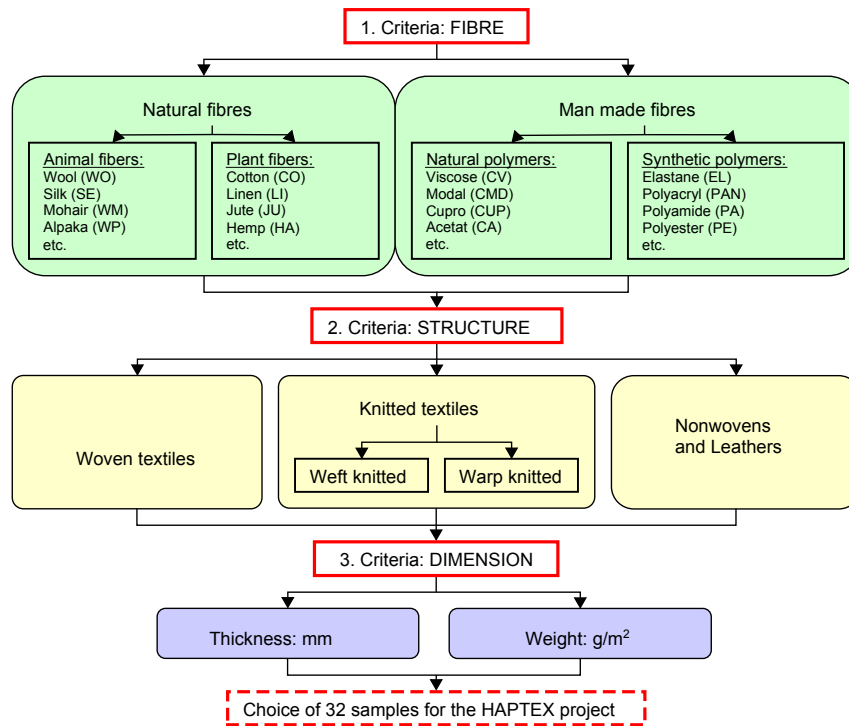


Figure 3.34: The HAPTEX selection process (recreated from [92])

in [92]. The first criterion, fibre, was chosen as different fibre properties influence the parameters of the textile. Samples 1 - 10, 11 - 21, and 22 - 32 are made of natural, blended and synthetic fibres respectively. The second criterion, structure, also influences the textile properties, such as knitted textiles which are more elastic than woven textiles. Textiles with the same fibre composition but different structures were selected so as to investigate the effect of the structure on the measured hand properties, totalling 23 woven textiles, 7 knitted textiles, and 1 nonwoven and leather. The final criterion, dimension, is based on the physical aspects of the textiles, specifically thickness and weight. Textiles of the same fibre composition and structure may still differ among these parameters, caused by differences in yarn number, yarn and loop densities or finishing treatments, and thus give different haptic parameters. By selecting a broad range of textile samples, the group represents a wide range in terms of both the physical parameters and the expected haptic characteristics of textiles, and so provides a greater opportunity for the comparison between the real and virtual surfaces.

The selection of samples was expanded by a second textile set (33 - 42) which aimed to increase the size of the selection of samples, focusing particularly on samples of men's woven suit fabric or samples that have elastic fibres/elasthane in their construction [93]. This set contained a total of 10 samples, consisting of six thin relatively and four relatively thick samples. The samples also vary by fibre composition, structure, and densities. The third set (43 - 54) expands the selection, and thus the range of properties being tested, even further and consisted of fabrics from the automotive industry, specifically automotive upholstery or automotive interior trim [94]. These samples were made from synthetic fibres and consisted of foam laminated and non-laminated warp knitted textiles. The laminated fabrics have a cover fabric which was laminated with polyether or polyester foam of different densities, which typically makes them thick and compressible, and the non-

laminated fabrics have been treated and finished in different processes (dyeing, sueding, raising, cropping, etc.).

For this work, the Kawabata Evaluation System for Fabrics (KES-F) system was used to provide the information needed to create the texture files. It was developed by Sueo Kawabata in 1980 and features a set of standardised testers for measuring the various properties of textiles. The standard output from this system is set of values, one for each of the properties, that may be considered as characteristic of that textile being examined. The measured properties are placed into six groups based on their testing method:

- Tensile Properties, where the tester applies tensile deformation (“strip biaxial deformation”) along the length of the sample until the tensile force reaches a particular value (500 gf cm^{-1}). After the tensile force reaches this value, the recovery process is also measured. Strain rate is kept constant at $4.00 \times 10^{-3} \text{ s}^{-1}$. The characteristic values are calculated from the relationship between tensile force and tensile strain.
- Bending Properties, where the tester bends the sample textile between curvatures -2.5 cm^{-1} and 2.5 cm^{-1} at a constant rate ($0.50 \text{ cm}^{-1} \text{ s}^{-1}$). The sample is held vertically to eliminate the effect of gravity. The bending characteristic values are derived from the relationship between the curvature and the bending moment per unit length of specimen.
- Surface Properties, where the surface of the sample textile is moved under the probe of the surface tester with a speed of 0.1 cm s^{-1} , while the probe is pressed into the surface with a specified force (10 gf for roughness and 50 gf for friction). The sample is held under tension of 20 gf cm^{-1} as it is moved. Profiles of thickness and coefficient of friction are recorded as a function of position, which are used to calculate the characteristic values.
- Shearing Properties, where the tester applies a shear force to the sample while the sample is held under tension (constant at 10 gf cm^{-1}), which is applied in a direction orthogonal to the shearing force. The relationship between angle of deformation and the applied shear force is then used to define the characteristic values. The shear deformation overlaps with the tensile deformation.
- Compressional Properties, where the sample is compressed between two circular plates, each with an area of 2 cm^2 , with a typical compression velocity of $20 \text{ } \mu\text{m s}^{-1}$ (women’s thin fabrics have a reduce compression velocity of $6.66 \text{ } \mu\text{m s}^{-1}$). The recovery process is then measured with the same velocity. The relationship between thickness and pressure is used to derive the characteristic values.
- Weight and Thickness, where the thickness is measured at the same time as the compressional property and recorded when the sample is under a pressure of 0.5 gf cm^{-2} , and the weight is recorded as the weight per unit area of the sample.

Table 3.1 summarises all of the characteristic values that the KES-F system calculates in the various groups, together with their respective physical units. The system is described

in more detail, including the origins of the tests, full details on the testing parameters, and how the characteristic values are defined, in [90]. The full sample set of 54 textiles was analysed during the HAPTEX project, by the Smartwear Lab in Tampere, Finland, using the KES-F and the values resulting from the testing of each of the parameters are recorded in [92] (samples 1 to 32), [93] (samples 33 to 42), and [94] (samples 43 to 54). As these values are considered characteristic of the textiles, they were later used to define the behaviour of the equivalent virtual textiles when they were simulated as part of the HAPTEX project.

Table 3.1: The characteristic values that are output by the KES-F (recreated from [90]).

Parameter		Definition	Unit
Tensile	LT	Linearity	-
	WT	Tensile energy	gf cm cm ⁻²
	RT	Resilience	%
	EM	Max extension	%
Bending	B	Bending rigidity	gf cm ² cm ⁻¹
	2HB	Hysteresis	gf cm cm ⁻¹
Shearing	G	Shear stiffness	gf cm ⁻¹ °
	2HG	Hysteresis at 0.5°	gf cm ⁻¹
	2HG5	Hysteresis at 5°	gf cm ⁻¹
Compression	LC	Linearity	-
	WC	Compressional energy	gf cm cm ⁻²
	RC	Resilience	%
Surface	MIU	Coefficient of friction	-
	MMD	Mean deviation of MIU	-
	SMD	Geometrical roughness	µm
Weight & Thickness	W	Weight per unit area	mg cm ⁻²
	T	Thickness at 0.5 gf cm ⁻²	mm

The unusual unit of gram-force is used by the KES-F system. This unit is defined as the magnitude of the force exerted on one gram of mass by standard gravity (9.806 65 m s⁻²). Thus, a force of 1 gf is equivalent to 0.009 806 65 N.

Although these systems are designed to only produce characteristic values of the texture properties in question, the tester was, in this case, used to obtain full measurements of the surface profile. The relevant tester is the KES-F surface tester (figure 3.35), which features a sensor with a contactor of 0.5 mm steel piano wire bent into a square shape, as schematically shown in figure 3.36. During the testing, the contactor is pressed into the texture with a contact force of 10 gf. The textile is moved under the contactor, supported by a smooth steel plate, under tension of 20 gf cm⁻¹ along the measurement direction, at a speed of 1 mm s⁻¹ for 20 s. The height of the contactor, in micrometres, is recorded by the sensor at a sampling frequency of 938 Hz, resulting in a recording of 18,756 samples (20 mm) along the surface of the textile in that direction. The tester then reverses the direction of the motion of the textile sample to record the height profile associated with travel across the surface in the opposite direction for the same distance (20 mm). The process was repeated so that recordings were made of the sample textiles in both the weft and warp directions across the surface, with some textiles having both their right and back sides recorded. These height profiles were also measured and recorded by the Smartwear

Lab in Tampere, Finland, and were later used to create the texture files which define the virtual surfaces.




Figure 3.35: The KES-F surface tester (from [92]).

All signals from the surface tester are passed through a high-pass filter such that signals of wavelength smaller than 1 mm may pass through it. The transfer function of this filter is described in equation 3.18, where $\omega_n = 2\pi\text{rad s}^{-1}$, $\xi = 0.6$ and ω is the angular frequency.

$$G(j\omega) = \frac{(j\omega)^2}{(j\omega)^2 + 2(j\omega)\xi\omega_n + \omega_n^2} \quad (3.18)$$

Each of the textiles is stored in a separate texture file (.FAB), containing a total of five matrices. Four of the five matrices are Fourier transforms of the surface profile for that textile from the KES-F roughness tester, with one matrix for each direction and the fifth matrix containing the list of frequencies used in the transformation. Each of the direction matrices contains sixteen spectra that each describe the profile of the surface at different locations along the surface. In order to obtain the spectra at these different points, a Hamming window, w_i , which is defined by equation 3.19 and shown in figure 3.37 for range $i = 0, \dots, N - 1$ where N is the length of the window (4096 samples, which corresponds to a length of 4.37 mm), was used at different offsets to remove signal outside the area of interest of each spectrum. The signal resulting from the combination of the original signal and the Hamming windows then undergoes a Fast Fourier Transform (FFT) to extract the various spatial frequency components of the surface at that point along the surface profile (see figure 3.38). The window is then shifted by 938 samples (~ 1 mm) along the axis for the calculation of the next spectrum. By using a window of this shape frequency artefacts associated with a rectangular window, due to discontinuity between the beginning and end of the sampled data, are avoided.




Figure 3.36: Schematic of the KES-F roughness tester contactor (recreated from [1]).

$$w_i = 0.54 - 0.46 \cos\left(\frac{2\pi i}{N-1}\right) \quad (3.19)$$

Finally, before being stored as the texture file, the full 2048-line spectra that are produced by the FFT of the surface profiles are reduced to a lower-resolution spectrum of 50 lines for use in the renderer. The spectral lines are chosen on a logarithmic scale, with many lines stored individually at lower spatial frequencies, where much of the detail of the surface profile is expressed, and lines of higher spatial frequencies being binned together, by a root-sum-square of the individual spectral lines. This allows any surface detail that may exist at the higher frequencies to be included in the texture files without needing to store the full spectrum and so save on both storage capacity and processing time. Additionally, some lines in the original spectra are always ignored, as they consistently contain spectral artefacts. More detail on the choices associated with the storage of the spectra is given in section 3.4.2.

The result of this processing of the recorded surface profiles is a surface definition file containing five matrices, four of which are the spectra of the surface profile at different points along the surface along the four different, principle directions, with the fifth matrix defining the frequencies associated with the individual spectral lines.

Figures 3.39 and 3.40 show two examples of surfaces that were recorded by the KES-F roughness surface tester, as well as the first 20 lines of the mean of the 16 spectra that result from the processing of these recorded surface profiles, along each of the four directions. These plots illustrate how the surface can differ based on the direction of exploration, as well as the corresponding effect that has on the spectrum. The two textiles presented here are 03 and 20, respectively.

Textile 03, Cord, is strongly anisotropic, with a set of ridges present along the weft direction. These ridges result in a very rough sensation while the surface is explored along the




Figure 3.37: An example of a Hamming window.

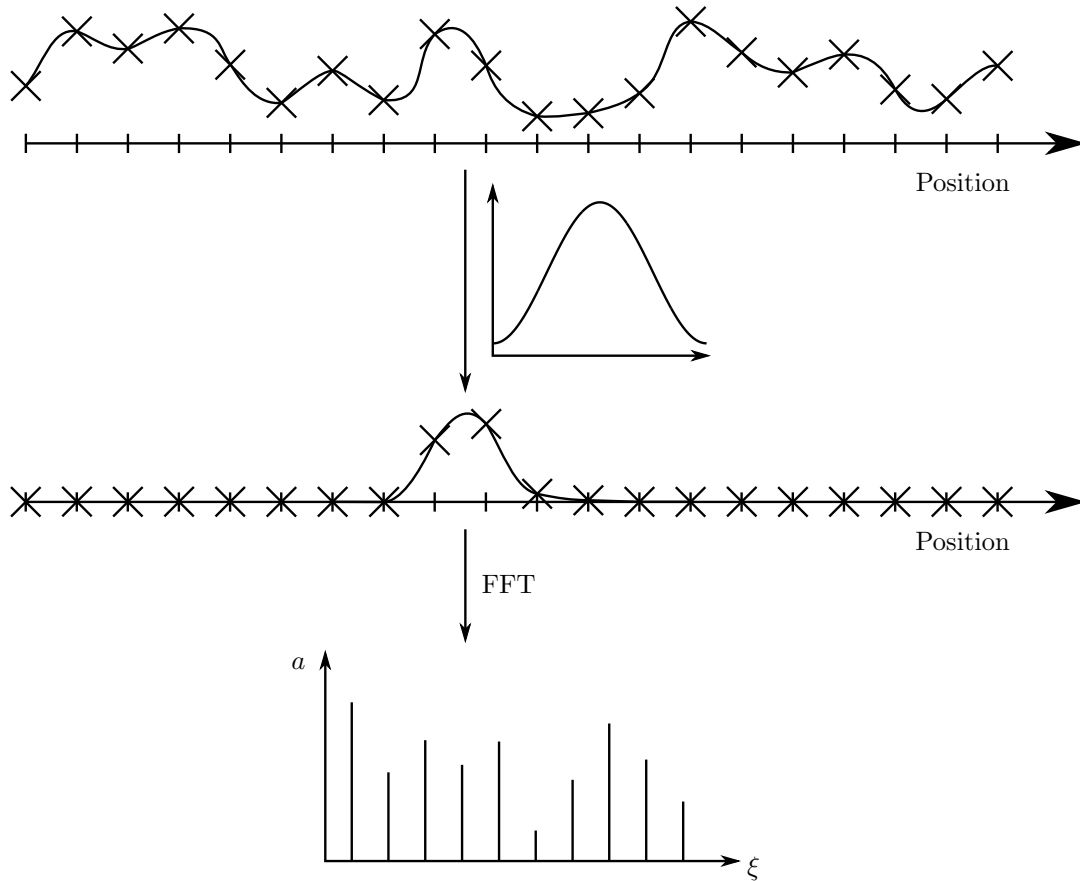


Figure 3.38: Derivation of a set of spatial spectra from a Kawabata line profile.

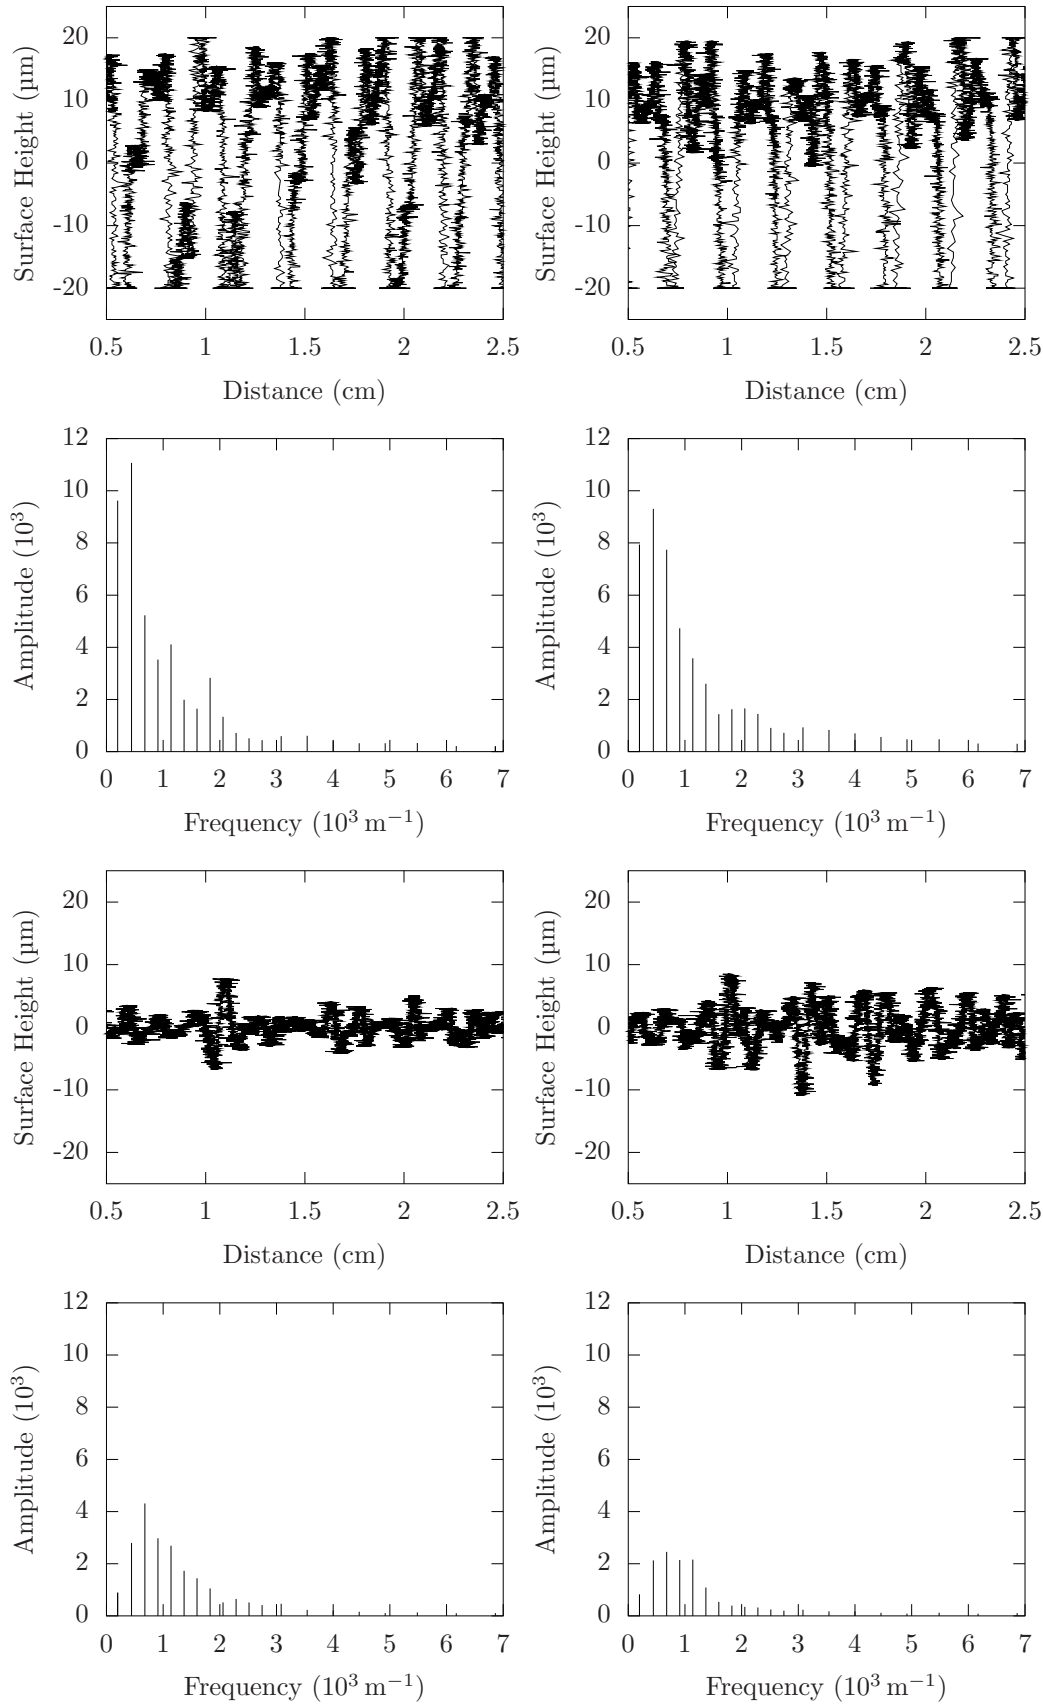


Figure 3.39: The surface profiles recorded by the KES-F roughness tester while travelling along the right (top-left), left (top-right), up (bottom-left), and down (bottom-right) directions across the right-side surface of textile 03 (above), as well as the first 20 lines of the mean of the spectra calculated from those surface profiles (below).

weft direction (right/left) but a very smooth sensation along the warp direction (up/down) and this is also reflected in the associated spectra, with larger amplitudes associated with the various spectral components in the weft direction than in the warp direction.

Textile 20, Warp knitted velour, is also anisotropic but, unlike textile 03 where the sensation difference exists between exploring along the weft and warp axes but is similar irrespective of direction of movement along those axes, the pile of the surface produces a sensation is different depending on the direction of motion along an axis. In this case, exploring along the positive warp direction (up) results in a much rougher sensation than exploring in the negative warp direction (down). This is also reflected in the spectra corresponding to those surface profile, with greater amplitudes recorded along the forward direction than the backwards direction.

3.3 The HAPTEX Project

The tactile renderer that is used within this project was developed as part of the EU-funded HAPTEX project (IST-6549) as part of a Future and Emerging Technologies initiative. The aim of HAPTEX was to create a virtual reality system that is capable of visually and haptically recreating textiles. Specifically, the scenario that was to be recreated was a two-digit interaction with a square of textile hanging from a stand (figure 3.41).

The HAPTEX consortium comprised of members from five different institutions, each specialising in a specific area within the wider project. Those institutions were:

- University of Exeter Biomedical Physics group, which provided experience in conducting psychophysics studies and expertise in the tactile field (both in hardware and in rendering strategies).
- WelfenLab at the University of Hanover, responsible for creating the software framework that brings the components of the system together.
- MIRALab at the University of Geneva, who specialise in computer graphics and complex computer animation.
- SmartWearLab at Tampere University of Technology brought with them knowledge of fabrics and the KES-F system for objective assessment of hand properties.
- PERCRO Laboratory at the Scuola Superiore SantAnna (Pisa) have extensive experience in developing haptic interfaces.

This section describes the virtual reality system that was developed over the course of the HAPTEX project, which is summarized in [96], beginning with an overview of the intended aims and development sequence of the system, followed by a description of the implemented system and concluding with evaluation of the system that was done as part of the project.

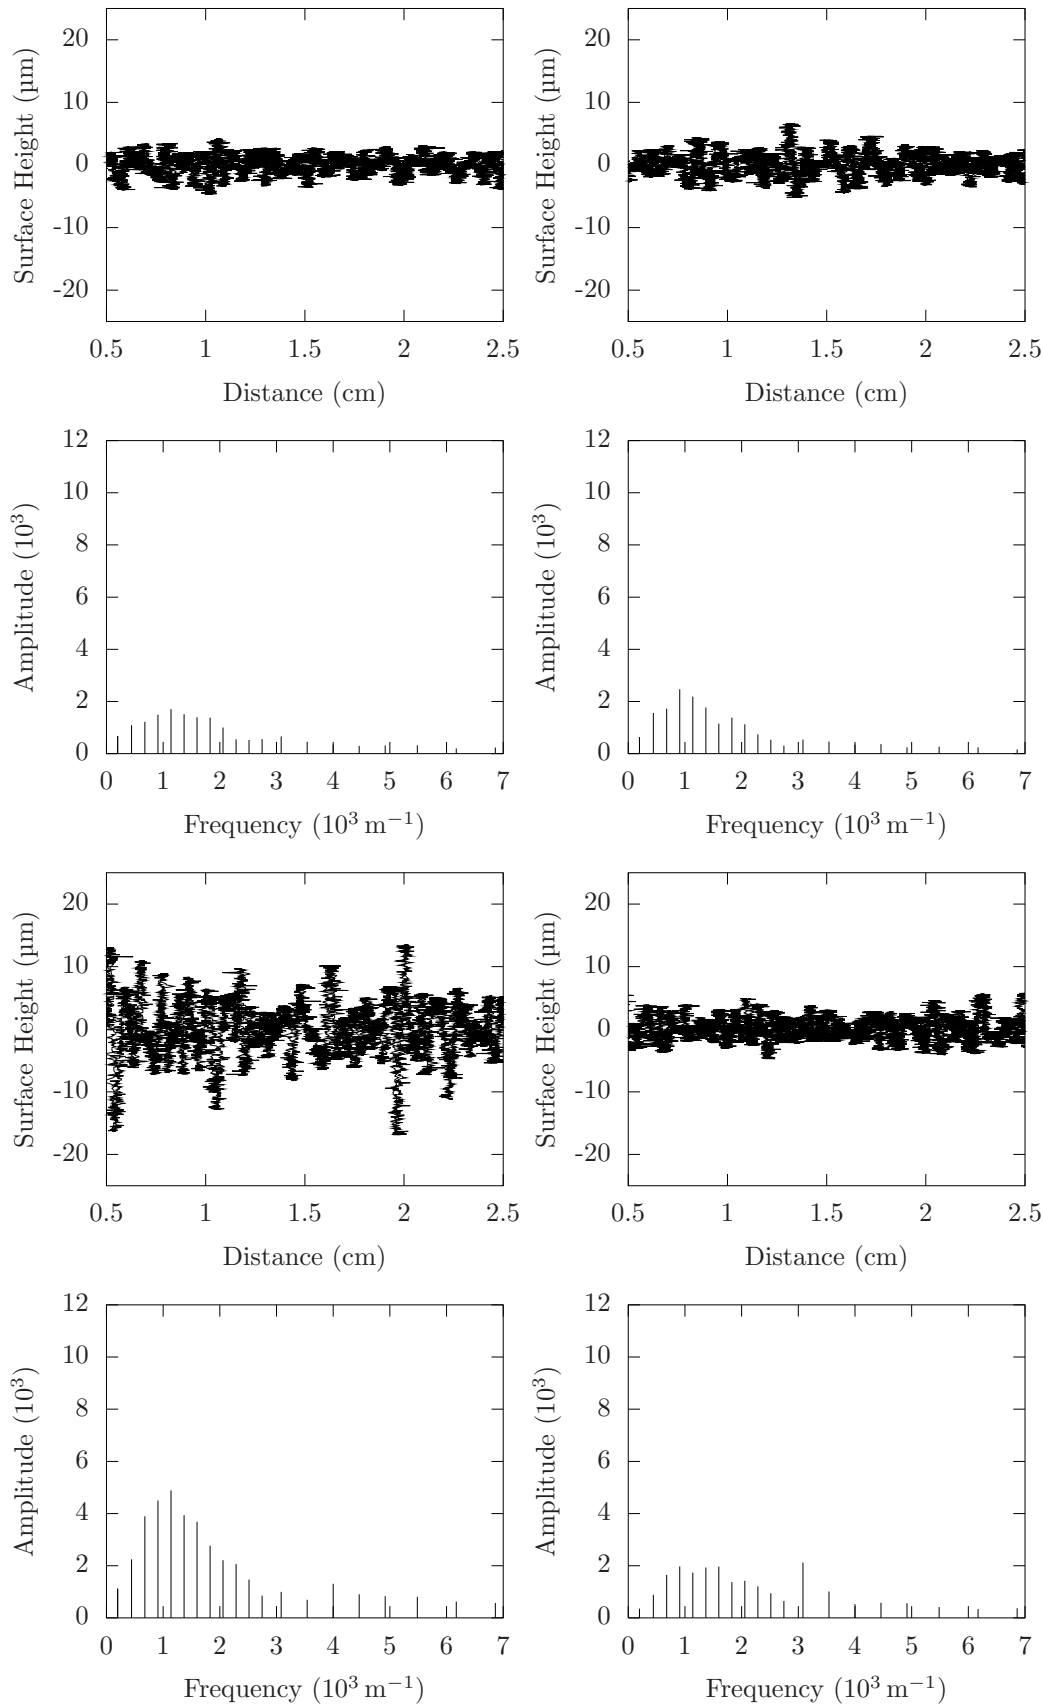


Figure 3.40: The surface profiles recorded by the KES-F roughness tester while travelling along the right (top-left), left (top-right), up (bottom-left), and down (bottom-right) directions across the right-side surface of textile 20 (above), as well as the first 20 lines of the mean of the spectra calculated from those surface profiles (below).




Figure 3.41: The target virtual scenario for the HAPTEX project (originally from [95]).

3.3.1 Overview of System Development

In order to produce a system capable of recreating the haptic sensations of manipulating textiles, the core focus of the project was the creation of a display that integrated both a tactile array for small, distributed forces (the textile surface as it is stroked), and a force-feedback device for the gross mechanical properties (stretching/shearing/etc. of the textile).

The concept of the device from the initial proposal (figure 3.42) shows a wrist-mounted device with two articulated force-feedback devices attached. The forces are conveyed to the thumb and index finger of the user via a pair of thimbles which are additionally equipped with a tactile display on the inner surface.

Each of the subsystems (tactile rendering, force-feedback rendering, and physical simulation) are run concurrently within the final system, with figure 3.43 showing the structure of the integrated system software. Before being integrated into the final system, each of the subsystems was implemented and tested separately, through a series of development levels.

Figure 3.44 shows the levels of development the HAPTEX system underwent during the project. Development level 1 consisted of a two-digit force-feedback system allowing users to manipulate the virtual textile. Development level 2 featured a one-digit tactile display to explore the surface of the virtual textile. Instead of the intended development level 3, which was not implemented, a related level designated as 4a was undertaken which connected one-digit of tactile display to a force feedback device. This integration was continued into level 4b, which expanded the system from one to two digits. The final development level, level 5, replaces the force-feedback system used in previous portions of development with a purpose built hand exoskeleton, thus bringing the device in line with the initial proposal. All development levels include visual feedback for navigation of the




Figure 3.42: Schematic of the concept of the proposed HAPTEX device (originally from [97]).

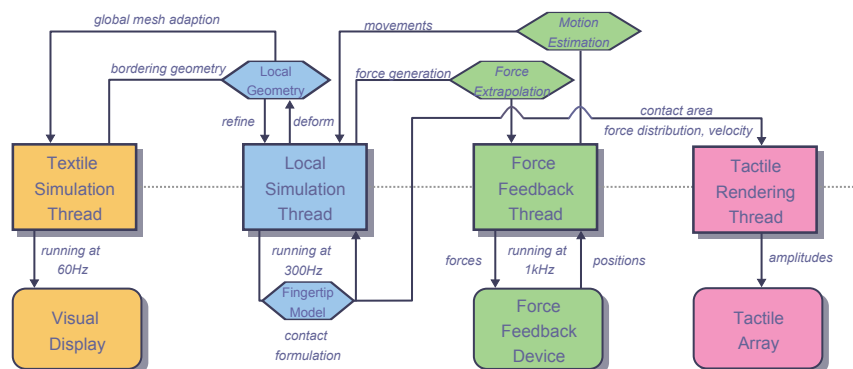


Figure 3.43: The different threads within the HAPTEX system (recreated from [96]).

virtual textile.

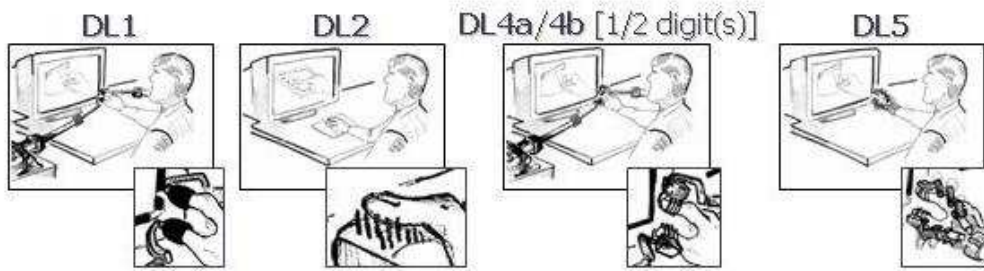


Figure 3.44: The development levels of the HAPTEX system (originally from [98]).

3.3.2 Physical Simulation

When a textile is selected by the user, a square of side length 20 cm is simulated hanging from a stand. This virtual textile may be squeezed, stretched, rubbed or lifted by the user using their thumb and/or index finger using the force feedback device. The software described in [99], [100] and [101] is used for the integrated system. Although the software is able to simulate the entire textile accurately, it does so at the cost of a relatively low refresh rate, running at 60 Hz in the integrated system. While this is sufficient for updates in the visual display, it is not enough for accurate haptic interaction. The physical simulation is therefore divided between two threads; a low frequency, high accuracy textile wide simulation and a high frequency, less accurate simulation of the “local geometry” (the textile surrounding the point of contact between the user’s virtual finger and the simulated textile).

The simulated textile consists of a finite set of discrete particles of defined mass and position. These particles are kept together by internal forces that act on adjacent masses which reflect the tensile behaviour of the real textile. External forces, such as gravity, and damping forces may also be included in the model. The simulation used within HAPTEX includes stretch/tensile, shear, bending, and damping forces, the magnitudes of these properties being found from the Kawabata Evaluation System for Fabrics (see section 3.2.5).

In order to avoid errors in the haptic interactions, the haptic rendering needs to be able to respond to the user’s movements with time-scales on the order of milliseconds. This is not possible for the simulation of the whole textile so a smaller, quicker simulation of the local geometry is used as an intermediate step between the full simulation and the haptic rendering. The local simulation receives the fingertip dimensions as well as their current and predicted positions to check for collisions with the virtual textile. In the event of a collision, the deformation of the textile is calculated based on the fingertip model being used. Forces that are generated during contact are sent to the force-feedback thread.

3.3.3 Force Feedback

The force feedback hardware for the HAPTEX project was developed by the PERCRO laboratory and is talked about in detail in [102]. For most of the project, hardware was a modified version of their existing GRAB system [103] (figure 3.45). This device uses two identical and independent manipulators which are fixed to the desktop modified to include new end-effectors which integrate a tactile display that rests on the user's thumb or index fingertip. Each manipulator can generate forces of up to 20 N, with errors in force in a range of about ± 0.1 N, within a workspace of volume 400 mm wide by 300 mm high by 200 mm deep. The workspaces of both manipulators overlap.

The tactile arrays were attached to the ends of the GRAB arms via gimbals which were designed to counterbalance the attached array. Each gimbal also featured position encoders, to give information about the orientation of the array, and force transducers, to measure the force being delivered to the finger tip.

The force feedback thread of the software is responsible for computation and generation of the forces occurring at the contact points between the fingertip models and simulated textile. The force feedback devices, together with the new gimbals, provide the software with the information required for its rendering calculations (position and orientation of the user's fingers, etc.).




Figure 3.45: The GRAB system from PERCRO (originally from [104]).

The second iteration of the force feedback hardware was intended to be a compact, wearable device that integrated force feedback and tactile elements into a hand exoskeleton, hence the device is known as HandExos. Figure 3.46 shows a CAD model of how the device would have worked. According to [104] a working prototype was demonstrated with a compatible tactile array but time pressures prevented the two being integrated together.

3.3.4 Tactile Renderer

Before the tactile renderer and force feedback devices were integrated together within the HAPTEX project, a stand-alone tactile renderer that can be used to explore a two dimensional workspace was created. The chosen display was a 6 by 4 array, with a 2 mm spacing




Figure 3.46: A CAD model of the layout of the HandExos device (originally from [105]).

which uses piezoelectric actuators to drive the contactor pins (Exeter tactile display). Details of this system and its operation are given in section 3.2, although a brief introduction follows in the next paragraph. For the HAPTEX project, two configurations of this system were developed. The first, which is used in this project, uses a graphics tablet to provide position information about the user's explorations to the rendering software. The second configuration uses a mouse to provide the information but is otherwise the similar.

The task of the tactile rendering component is to recreate the sensations to the user of exploring the surfaces of the textiles. This is done by the rendering software sending signals to the tactile display that make displacements on the users fingers as they explore the virtual textile that correspond to displacements that would occur if the user were exploring the equivalent real textile in a similar fashion. The displacements are calculated by the rendering algorithm using a combination of bandpass filters, the position and velocity of the user on the textile's surface, and the surface profiles of the textile, which were recorded by the KES-F. The displacements that are calculated are used as the amplitudes of the tactile display's driving signal, causing the actuators to move with a corresponding amplitude. The tactile rendering algorithm is the same in both the stand-alone two-dimensional and integrated three-dimensional cases.

The configuration of the display had to be modified from the version that was used in the tactile testing, which is designed to work on a two-dimensional surface, into a version that could be mounted onto the force feedback device. The modified displays had a number of design constraints. The display needed to fit into the space available in the GRAB gimbal as well as minimising the weight and bulk of the display. A further constraint is, so that the thumb and finger could be brought together to grasp the virtual textile, very little of the display's mechanism can be below the fingertip.

The proposed display solution, the HAPTEX tactile display, is shown in figure 3.47. In this design, the display's structure is reduced leaving only the finger-plate and the frame for the actuators. The piezoelectric actuators are held above and around the finger with the contactor pins being wrapped around and through the finger-plate to provide the

sensations to the fingertip. The finger-plate serves as both a replacement for the GRAB force-plate, which measure and transmits the fingertip contact forces, and a substitute for the static pins used in the two dimensional display. As the actuator pins are quite vulnerable, a safety shield is fitted to the underside of the display in the full system. The total thickness of the device below the finger is 5 mm, which allows for the required grasping action.




Figure 3.47: The tactile array used in the HAPTEX system (HAPTEX tactile display, originally from [105]).

3.3.5 Evaluation

The evaluation that was done as part of the HAPTEX project is available in [98], which records the evaluation of the separate force feedback and tactile rendering elements, and [105], which focuses on the evaluation of the integrated system. The latter is also the focus of this section.

The evaluation procedure used a series of defined manipulations of the textile so as to provide information about the property being investigated. The properties being assessed were:

- tensile stiffness
- surface roughness
- surface friction
- bending stiffness/weight/drapeability

Manipulations were also defined for two additional properties that were not included in the assessment:

- shear stiffness - sticking of the GRAB mechanism during horizontal movements meant that it was not possible to reliably perform the shear manipulation
- compressional stiffness - this property was not implemented in the software models

The manipulations (figure 3.48) were based on manipulations used for evaluation fabric hand properties, but also accommodated the limitations of the HAPTEX system [106]. The manipulations were performed on a sample set of 5 virtual textiles, which were chosen from the HAPTEX library prior to the evaluation based on having both real and virtual versions available and having the fullest range and best spread of the KES-F values. The textiles were manipulated and the properties were rated on a scale of 1 to 5, with reference samples provided to define the ends of the ranges. Each sample was presented twice in a sequence of 10 items. These ratings were compared to those obtain from equivalent manipulations of the corresponding real textiles. If the HAPTEX system is effective in the creation of virtual textiles, then the rank order for a given property of a set of virtual textiles should be the same order as the property of set of equivalent real textiles. The rank order should also be matched by the order of the KES-F measurements of that property, as these values were used to create the software models. Full details of the experimental procedures are available in [105].




Figure 3.48: Finger positions and manipulations used to test the properties (recreated from [105]).

Data for evaluating the HAPTEX system was obtained in Hanover from two users who had some familiarity with the system, the first (S1) was from the University of Exeter and the second (S2) was from the University of Hanover. As mentioned above, each textile

was assessed by each user twice. Figure 3.49 shows how the mean ratings from each the two subjects compare to each other, along with the correlation coefficient, R .




Figure 3.49: Comparisons between mean ratings from subject 1 (S1) and subject 2 (S2) for the properties of (left to right) tensile stiffness ($R = 0.91$), surface roughness ($R = 0.93$), surface friction ($R = 0.76$), and bending stiffness ($R = 0.61$) (recreated from [105]).

Between the mean subjective ratings of subjects 1 and 2, we see a good correspondence for the assessment of the tensile stiffness and surface roughness but a less strong correspondence for the assessment of the surface friction and the bending stiffness.

Figure 3.50 shows how the subjective ratings of the textiles (averaged over the two subjects), compare to the corresponding physical values for tensile stiffness (measured at a force of 10 N), surface roughness, surface friction, and bending/drapeability that were obtained from the KES-F, and so were incorporated into the software models. The physical values are represented on a logarithmic scale. As the visual assessment of bending stiffness is effectively an assessment of drapeability, the corresponding physical value was taken to be the ratio of bending stiffness to weight. The correlation coefficient between the subjective rankings and the measurements of physical properties, R , is also given.




Figure 3.50: The relation between the average subjective ratings of the virtual textiles for both subjects (S1 & S2) and the corresponding physical values for (left to right) tensile stiffness ($R = 0.85$), surface roughness ($R = 0.97$), surface friction ($R = 0.90$), and bending stiffness ($R = 0.48$) (recreated from [105]).

Tensile stiffness, surface roughness and surface friction can all be seen to have a good correspondence between the subjective ratings of virtual textiles and the physical parameters used in the software models, suggesting that the HAPTEX system is delivering appropriate cues to the users on the nature of the textile. It is suggested in [105] that the low correlations in the assessment of bending, both in figures 3.49 and 3.50, is the result of the difficulty in making the assessment.

Corresponding data based on the manipulations of the real textiles was obtained in Tampere from two subjects, one from MIRALab who had experience in the fashion industry

(*S3*), and an expert in the textile industry, recruited by SmartWearLab (*S4*). Figure 3.51 shows the average subjective ratings, based on the real textiles and averaged over subjects 3 and 4 compared to the corresponding physical values obtained from the textiles by the KES-F for tensile stiffness, surface roughness, surface friction and bending/drapeability, with the latter being defined as it was previously. The correlation coefficient between the subjective rankings and the measurements of physical properties, R , is also given.




Figure 3.51: The relation between the average subjective ratings of the real textiles for both subjects (*S3* & *S4*) and the corresponding physical values for (left to right) tensile stiffness ($R = 0.85$), surface roughness ($R = 0.86$), surface friction ($R = 0.70$), and bending stiffness ($R = 0.74$) (recreated from [105]).

Each real textile was assessed under two conditions: “with vision” and “without vision”. The data presented here is from the “with vision” condition for the tensile and bending manipulations and from the “without vision” for the roughness and surface manipulations (which were chosen to be comparable with the wire-frame view used when assessing the virtual textiles). There is good correspondence between the ratings of the real textiles and the measured physical properties in all four cases

Figure 3.52 shows how the average ratings of subjects 1 and 2 (*S1* & *S2*), who were rating the virtual textiles, compare to the average ratings of subjects 3 and 4 (*S3* & *S4*), who were rating the equivalent real textiles, as well as the correlation coefficient between the two, R . These plots show good correspondence in the rating of the virtual and real textiles for tensile stiffness, surface roughness and surface friction but no correspondence for bending stiffness which may be attributed, as described previously, in the difficulty in making the assessment in bending stiffness for virtual textiles.




Figure 3.52: The relation between the average subjective ratings of the virtual textiles for both subjects (*S1* & *S2*) and the corresponding average subjective ratings of the real textiles for both subjects (*S3* & *S4*) for (left to right) tensile stiffness ($R = 0.96$), surface roughness ($R = 0.90$), surface friction ($R = 0.93$), and bending stiffness ($R = 0.00$) (recreated from [105]).

Overall, it can be said that the system that was developed for the HAPTEX project is

effective in the creation of virtual textiles. The evaluation that was done towards the project's end shows a generally good correspondence in the rank orders of the virtual textiles between different users of the virtual system, as well a good correspondence between the average rank order across the users to both the rank order of the KES-F measurements, that were used in the creation of the software models, and the average rank order from the equivalent test performed on the corresponding real textiles.

3.4 Improvements to Renderer

Since receiving the original renderer following the conclusion of the HAPTEX project, a number of improvements have been made to the renderer. In addition to some minor changes within the software, some notable alterations have also been made to what surface texture information is used, and how the amplitudes for the output wave are calculated and implemented in the fundamental rendering strategy. This section details each of the notable changes, as well as the rationale for each change.

Unless it is otherwise stated in the respective section, these changes were made before the first experiment (described in section 6.2) was performed. Essentially, this section describes the transition that occurred from the original HAPTEX renderer to the renderer used in the various experiments.

3.4.1 Direction of Motion

The intention of the HAPTEX project is to recreate the surface of textiles based on recordings of the surface profile that were made by the KES-F system. A recording of the surface was made covering 30 mm in one direction of either the weft or warp axis, followed by another recording in the opposite direction along the same axis by moving the textile under the sensor and recording the height profile as a function of position. The recording of both directions was made as one continuous recording so, to eliminate effects such as the acceleration of the textile under the sensor as well as the change in the direction of motion of the textile, the total usable section of recorded surface covered 20 mm in each direction.

The initial form of the tactile renderer (that was received from the HAPTEX project) used only the recordings made from the first direction of motion in the weft and warp directions to represent the surface profiles in the texture files (total of three matrices; two containing spatial spectra and one defining the frequencies of the spectral lines). The render then used these spatial spectra to represent the surface along that direction, irrespective of the direction of travel along that axis.

However, as is illustrated in sections 4.2 and 4.3, some surface textures are not direction independent in terms of the sensation experienced by the user in the way this would imply. Thus, in order to more accurately recreate the surfaces of the textiles, one of the earliest modifications that made was to alter the texture files to instead include the surface information for both directions on the weft and warp axes, now totalling five matrices per

texture file (four containing the spatial spectra and one defining the frequencies of the spectral lines).

The renderer itself had to also be modified to account for this change. As well as now being able to load and identify the new matrices, the renderer also had to distinguish between the positive and negative directions of motion along the weft or warp axes so as to select the appropriate matrices for use in the amplitude calculation. This was done based on the sign of the velocity of the tactile display, as output by the combination of the graphics tablet and Kalman filter. During the conversion from single direction to both directions, the parts of the surface profiles that had previously been included in the renderer were defined as positive/forward motion along the surface and the new section of surface profile were defined as negative/backward motion along the surface.

3.4.2 Textile Frequency Range

The texture files were also changed in the frequency range that they cover. Initially, the frequency range covered by the texture files was linear, i.e, the spacing between the spectral lines was a consistent, fixed value. Specifically, the frequency range defined by the first 50 spectral lines recorded from the full FFT of the surface profile. This linear spacing has been replaced by a new scale based on a logarithmic approach. The result is still a 50 line spectrum but one that is instead based on the first 500 lines of the full FFT spectrum of a profile. Table B.1 compares the frequencies that define the spectral lines of the two approaches.

The logarithmic scale has a high level of detail at lower spatial frequencies, almost matching that of the linear approach. However, for higher frequencies, the logarithmic approach is to bin together a number of the original spectral lines into a single value, using a root-sum-square formula, with higher frequency sections having a larger number of the lines in the full spectrum binned together.

As the majority of the textile surface detail exists at lower spatial frequencies, as illustrated in figures 3.39 and 3.40, the use of a logarithmic scale to define the spectra avoids unnecessary detail at the high frequencies and allows the upper limit of the file to be expanded from $11\,448\text{ m}^{-1}$ up to $109\,559\text{ m}^{-1}$. This ensures that any materials that have fine surface details, which would be encoded at a higher spatial frequency, have these features included in the texture files without any significant increase in storage requirements (texture files are approximately the same size) or computational time (50 lines per spectrum are still being used in the calculation of the output amplitudes).

As the spatial frequencies that define the spectral lines are also stored within the texture files, the tactile renderer itself did not need to be modified to account for these changes, unlike the changes described in section 3.4.1.

3.4.3 Modifications to Amplitude Calculation

The basis of the rendering strategy is to reduce the spectrum of the surface texture down to its fundamental frequencies by applying band-pass filter functions. In the original

rendering strategy, the output amplitude, A , was calculated by an equation of the form in equation 3.20, which sums the products of the spectral amplitudes, $a(f)$, and filter values, $H(f)$, over all of the available temporal frequencies, f .

$$A = \sum_f a(f)H(f) \quad (3.20)$$

However, when combining components of different frequencies, the power of the combination is the sum of the powers associated with the individual amplitudes and power is proportional to the amplitude squared, rather than simply amplitude. It was therefore decided that the rendering software should instead calculate the output amplitudes using a root-sum-square approach to combining the different components, as described in equation 3.11.

3.4.4 Filter Modifications

The band-pass filter functions that were provided at the start of the project were simply the solutions to the set of simultaneous equations 3.16 over the frequency range of 0 Hz to 600 Hz. These filters are depicted in figure 3.53, where they are also directly compared to the new filter functions that were used during the experiment (figure 3.33).




Figure 3.53: The original band-pass filters that were used to highlight the fundamental frequency components for tactile rendering (recreated from [1]), compared to the new filter functions.

The minimum detention thresholds shown in figures 2.8 and 3.32 are displayed in terms of the displacement level relative to $1\ \mu\text{m}$, x , rather than the absolute displacement, d .

To calculate the filter functions, it is necessary to convert the displacement levels back to absolute displacement. The conversion of the original filters was found to be based on the formula that is typically associated with intensity or power (equation 3.21). As the amplitudes of the waves are the subjects of the calculations, and the detection thresholds refer the minimum detectable amplitudes, the version of equation 3.21 that is associated with amplitude (equation 3.22) would instead be more appropriate in this case.

$$x = 10 \log\left(\frac{d}{d_{ref}}\right) \quad (3.21)$$

$$x = 20 \log\left(\frac{d}{d_{ref}}\right) \quad (3.22)$$

As well as the conversion equation, the revised filters have some additional changes compared to the original filters, the results of which are illustrated in figure 3.53. For the low frequency filter, H_{40} , the high frequency lobe that originally existed was removed by setting the values in that range to zero, along with the negative values. For the high frequency filter, H_{320} , the negative values were set to zero, and its end point, which was extended from 600 Hz to 1 kHz, was also set to zero. A smooth variation of the filter function from 500 Hz to the zero value at 1 kHz was computed in OCTAVE. So that the two filters are of the same size, the end of the low frequency filter was also extended up to 1 kHz, with all new values in the extended range set to zero.

The result of these modifications is the two band-pass filters given in figures 3.33 and 3.53, which are the filters that were subsequently used in this project. By making these alterations to the base filters, it is ensured that high frequency information is not lost to the P channel, while the NP is appropriately restricted to low frequency information and does not include information from higher frequencies in the amplitude calculations.

3.4.5 Amplitude Scaling Factors

In order to scale the amplitudes that are calculated by the rendering software to an appropriate level where the firmware can use them in the waveform, two scaling factors exist within the system to control and restrict the amplitude levels. Each factor serves a different purpose within the renderer and, as such, is derived from different aspects of the system.

The first scaling factor operates entirely within the rendering software, and is designed to normalise the value of the calculated amplitudes so that they may be sent to the firmware. To achieve this, each amplitude is divided by a pre-determined maximum value so that the resulting value is between 0 and 1, with values greater than 1 being capped at 1.

These first factors, in the original version of the renderer, had a value of 50000 in both the NP and P channels. Before the experiments began, a new value for these factors was established by calculating the amplitudes that might be generated by the exploration of the textile surface, from a variety of textiles at a typical exploration speed and in a

variety of directions. This establishes the range of amplitude values that these textiles occupy. Specifically of interest here is the maximum value that a textile might generate during typical exploration of the surface. Figure 3.54 shows the sum of the maximum amplitudes obtained in both channels for the various textiles. From this, the new gain factors were then chosen to be 30427 for both the NP and P channels, shown in the figure, which accounted for the maximum possible amplitudes for most textiles (70%). Although this doesn't cover all of the textiles, it is important to consider both that the values used for figure 3.54 are the theoretical maximum for any given texture, thus the amplitudes typically calculated by the rendering software will be lower, and that the amplitudes associated with the individual, separate channels will also be lower.




Figure 3.54: The sum in both channels of the maximum amplitudes that were calculated for a 10 cm s^{-1} exploration of the various surface textures, as well as the line representing the new scaling factor.

These factors, as well as placing an upper limit of the output of the amplitude calculations, also establish what the resolution of the amplitudes that may be output by the renderer is. This is most obvious with respect to the output from the DAC, which is not continuous from minimum to maximum value. Instead, as the input is defined by an 8-bit digital number, the analogue output is correspondingly divided into 256 discrete steps between the minimum value of zero and the maximum value as defined by the reference voltage. The use of the two fundamental frequencies means that each channel has 128 of the steps available to define its respective output. Thus, by selecting a lower maximum amplitude value, a greater resolution in the output amplitudes is available to the renderer ($30427/128 = 237.7$ per step, rather than $50000/128 = 390.6$ per step).

The second scaling factors are a set of gain factors that are used in the amplitude calculations and firmware to allow for dynamic control of the intensity of the signal by the user

by scaling the amplitudes of the output signal. The gain factors (one for the 40 Hz NP channel and one for the 320 Hz P channel) may be altered separately, controlling not only the strength of each channel, but also allowing for different balances between the channels; the latter being useful for accounting for factors such as the difference in sensitivity between the two mechanoreceptor channels, as described at the end of section 3.2.4.

The factors present in the original HAPTEx renderer were 0.7 for the NP channel and 0.5 for the P channel but later subjective comparison of virtual textiles with their real equivalents performed by experienced users, including the author, suggested that these gain factors in the renderer should instead have values of 1.0 for the NP channel and 0.7 for the P channel. It was these new gain factors that were used in the initial experiment, although they were changed again for further experiments, as described in section 3.4.8.

3.4.6 Buffering Calculated Amplitudes

In the original version of the renderer, after the amplitudes are calculated they are immediately sent to the firmware to be implemented directly in the output waveform. There is, however, no control over when the waveform is updated with the new amplitudes and, as the firmware is continuously running to output the waveform, this can potentially create problems with the waveform that is controlling the tactile display. Figure 3.55 shows one such problem, where the amplitudes have been update at the incorrect point in the cycle. Cases such as this occur on the occasions where the calculation of the amplitudes has exceeded the typical 25 ms time frame. These errors risk creating anomalous high frequency signals in the touch receptors of the skin, thus detracting from the two frequency rendering strategy.




Figure 3.55: Without the buffer for the amplitudes, the output wave is updated when the new amplitudes are received by the firmware from the rendering software.

To correct this problem, a buffer was added to the rendering software to time the updating of the wave amplitudes in the firmware to the wave zero-point, every 25 ms. In the event of the calculation time exceeding the 25 ms window, the amplitudes are held in the buffer until the next update opportunity, i.e. the completion of the next 25 ms cycle, before the update is made and the amplitude calculations are continued as normal.

3.4.7 Contactor Surface Area

When the tactile renderer was presented to the public (chapter 5), it was generally well received. However, a consistent negative component of the feedback was the “prickly” sensation of the Exeter tactile display. Some even suggested that this broke the illusion that it was the surface of a textile they were exploring. In order to address this, the area of each of the 24 contactors was expanded through the use of caps that were fitted over the ends of each of the individual contactors, increasing the area of each contactor to $\sim 1 \text{ mm}^2$, up from the $\sim 0.28 \text{ mm}^2$ of the contactor end. Figure 3.56 shows the tactile display with half of the caps in place.




Figure 3.56: The Exeter tactile display before (left) and after (right) the addition of caps to the contactors.

The caps made the sensation from each individual contactor less pronounced within the rendering setup, and gave the sensation of a more continuous surface texture. Experienced users of the renderer were asked for feedback about the new sensation given by the display across a variety of textures. They suggested that this sensation was an improvement over previous sensations from the display without the caps. This improvement is, however, limited to textiles of medium or high roughness, where the output from the display serves to mask the sensation caused by the presence of the individual contactors on the skin’s surface. For smooth textiles, where the vibratory output of the display is very weak, the sensation of the multiple, individual contactors is still apparent.

3.4.8 Dynamic Range and Spectral Balance

Anecdotal evidence from the experiment described in section 6.2 suggests that the texture files used in the experiment covered too wide a dynamic range, with high intensity textiles being too strong and low intensity textiles being too weak. To address this, new versions of the texture files were created intended to occupy a narrower dynamic range.

To do this, the textiles would have to be scaled according to a power law, such as in equation 3.23, with a value of the exponent b of < 1 . The values of the output y would then be increased for values of x less than the constant a , approximately the same for values of x around the value of the constant a , and decreased for values of x greater than the constant a .

$$y = ax^b \quad (3.23)$$

To decide the values that would be assigned to the constants in the power law, the author performed a subjective analysis on various textiles, focusing on the sample set used in the first experiment as they, by design, covered a broad range of textures. The gain factors of the output signal were altered until the author felt the subjective sensation of the virtual textile matched the equivalent real textile. Figure 3.57 shows the sum of the mean characterisation values of the textiles in sample set 1 compared to the equivalent sum of the amplitudes the author deemed to be the subjective best-match between the real and virtual textiles. These scaled values were then fitted based on the assumption that the relationship between the original and scaled values were consistent with the power law, and that the value of the exponent was < 1 .




Figure 3.57: Comparison between the sum of mean characteristic amplitudes of sample set 1, and the subjective best match to the real textiles, as established by the author, and the fitted scaling power law.

The resulting values from the fitting suggest that the textiles should be scaled according to equation 3.24 (illustrated by the line in figure 3.57), where A_o is the original intensity amplitude, as determined by the sum of the mean values of A_{40} and A_{320} , and to give the scaled amplitude, A_s . The value of A_s/A_o is then used as a scaling factor on the matrices

within the texture files.

$$A_s = 0.65A_o^{0.38} \quad (3.24)$$

Because of the anisotropy associated with certain textiles, the scaling was done on each matrix separately, with scaling factors calculated from the sum of the mean amplitudes generated by travel in that direction across the surface of the textile. As an alternative, it would also have been possible to compress the dynamic range between textiles but retain the spatial variability within textiles by scaling all of the spectral matrices within the texture file with a single factor, regardless of the textiles anisotropy.

After the alteration of the dynamic range of the textiles, the spectral balance (the ratio of the P and NP channels) was also considered. Although a ratio of 0.7:1.0 (P:NP) was used for experiment 1 (sections 3.4.5 and 6.2), further investigation by the subjective matching of virtual textiles to their real equivalents suggests that a ratio of 0.83:1.0 provides a more effective recreation of the textile surfaces and this ratio is used in later experiments, though the “best-match” ratio of the two channels was observed to vary between textiles.

3.4.9 Summary of Modifications

Although the basic operation of the renderer and display follows the same principles as those that were developed as part of the HAPTEX project (as described in section 3.2.1), various aspects of the specifics of the renderer’s operation have been altered during the course of this project, partly to correct errors and omissions in the original software, and partly as a result of feedback from system trials. Particular changes that were made are:

- the renderer and texture files now distinguish between the direction of motion along the weft and warp axes, in case of surface anisotropy
- the range of spatial frequencies recorded within the texture files has been increased to preserve surface detail with a high spatial frequency
- the amplitudes that are output by the various channels of the renderer are calculated based on a root-sum-square algorithm, rather than a linear one, so as to better combine the sensations from the individual components
- the characteristics of the band-pass filters for the NP and P channels have been changed to ensure the filters isolate the correct information for the calculations of their respective amplitudes
- the amplitudes A_{40} and A_{320} in each of the 24 channels have been normalised to allow for a greater resolution in the output signal
- the implementation of the calculated amplitudes is now delayed until an appropriate time, to avoid errors in the wave that is output by the renderer
- the contactor surface on the display has been improved to reduced the sensation of a series of distinct contactors on the fingertip

- the dynamic range of the texture files has been compressed and the output signal rebalanced, based on feedback from both the first experiment and experienced users, so as to more accurately match the sensation of the virtual textiles to their real equivalents

3.5 Summary

There exist a number of tactile displays that use a variety of different methods to convey tactile information to the user. The best choice of tactile display and rendering strategy depends strongly on the requirements of the situation, particularly what information is needed, and under what circumstances it is to be conveyed.

Here, piezoelectric bimorph actuators are used in the tactile display as they can be easily fit into a movable device for the required active exploration. The actuators are driven by a mixture of sine waves at frequencies of 40 Hz and 320 Hz, which is generated by a set of driving electronics. The display is only required to operate at these frequencies, i.e. a broadband response is not required.

The amplitudes of the two components of the output signal are defined by the tactile rendering software that is running on the host computer, which uses the rendering algorithm, together with information about the surface of the textile, as defined by the spatial spectra stored in the appropriate texture file, band-pass filters centred around the 40 Hz and 320 Hz frequencies, as well as the display's position and velocity in the virtual workspace, to calculate the output amplitudes. These amplitudes are then sent to the drive electronics to be used in the output signals before new amplitudes are calculated. The amplitudes of the output signal are updated every 25 ms.

The display, driving electronics, and principles of operation were all developed as part of the EU-funded HAPTEX project, which intended to virtually recreate textiles both haptically and visually. The tactile display used in this project (Exeter tactile display) was part of the second stage of the development of the HAPTEX system. The completed HAPTEX system featured versions of the tactile display that were modified to minimise the amount of display mechanic beneath the fingertip (HAPTEX tactile display), which were mounted onto force-feedback devices. These displays could then be grasped by the thumb and index finger and then used to manipulate than explore the virtual textile. Some evaluation was done towards the end of the project, the results of which suggested that the system was effective in the virtual recreation of textiles.

Over the course of this project, a number of improvements have been made to the system based on the analysis of experimental results and feedback obtained from experienced users or public demonstrations. These changes are intended to make the sensation of a virtual textile experienced by the user even closer to the sensation from the equivalent real textile. Factors that have been modified include the fundamental algorithm for the rendering strategy, the surface of the tactile display, the spectral balance between the two frequency components, and the dynamic range occupied by the virtual textiles.

Chapter 4

Characterisation of Renderer

The intention of this chapter is to demonstrate that the time-varying patterns presented to the user during active exploration of the virtual environment can be considered to be indicative of a particular surface texture. To do this, a strategy is used that establishes, for each textile, a set of characterisation values that suggests how the virtual textiles may be related to each other. Time constraints towards the end of the HAPTEX project meant that the developed tactile renderer received very little characterisation at the time and so this chapter serves to correct that omission and documents work done to characterise the renderer used in the present study.

This chapter begins with a discussion of how characterisation of the tactile renderer can be achieved and how that characterisation may be validated against pre-existing, KES-F measurements. This is followed by the characterisation of the renderer that was used in the first experiment, as well as discussions on both the effect of exploration speed on the characterisation, and on the anisotropy of the surface texture that can be observed from different directions of travel across the surface.

Once characterisation of the initial experimental renderer is achieved, it is compared to the renderer that existed at the beginning of the current project/end of the HAPTEX project. The feedback obtained from the first experiment was used to alter the renderer further and the characterisation of this incarnation of the tactile renderer concludes the chapter.

It was desired that the stimuli used for the experiments should cover a wide range of available sensations. Therefore, the characterisation strategy presented here was used to guide the selection process of the sample sets that were used in the experiments detailed in chapter 6. These characterisations were also compared to the multi-dimensional space that was created from the results of experiment 3 in order to possibly identify which parameters are used by the user during texture discrimination. The characterisation strategy has previously been presented in [107].

4.1 Characterisation Strategy

As the surface texture is presented to the user as a set of time-varying amplitudes (A_{40} and A_{320}), there are a number of properties that can be calculated from those amplitudes to characterise the output signal for a given textile. In this case, the chosen properties are the mean of the amplitudes that are output by the renderer, μ , to give a measure of the intensity of the signal component, and the coefficient of variation (CoV) of the amplitudes, c_v , to give a measure of the variability of the signal component. The CoV is defined as the ratio of the standard deviation, σ , to the mean (equation 4.1). Both of these sets of properties are considered during characterisation as we might expect a signal of a particular mean level but with high variation to feel different to one of the same mean level but with low variation. These values can be presented as a set of two-dimensional spaces that illustrate how the different signals that are output by the renderer from different textures may relate to each other.

$$c_v = \frac{\sigma}{\mu} \quad (4.1)$$

In order to produce a representative value for each of the chosen properties, the amplitude values that are output by the renderer for the two components of the tactile signal are calculated using the rendering strategy, as described in section 3.2.4, for a particular exploration speed. The strategy is implemented in an OCTAVE script so that a large number of textiles can be characterised quick succession. The calculations also iterate across all available locations on the virtual surface, as well as a large number of directions of motion across the surface. This ensures that amplitude sets are representative of a wide variety of exploratory movements that may be made by the user across the surface. The results of these calculations can then serve as the sample population from which the mean and CoV values are calculated for the texture in question. For the characterisations, a typical surface exploration speed was considered to be $\sim 10 \text{ cm s}^{-1}$, supported by [108].

In order to assess the validity of the values used in these two-dimensional plots, the calculated mean amplitude values of the components may be compared to values that are part of the standard output by the KES-F (see section 3.2.5 for an overview of the KES-F system, or reference [90] for a full description). The relevant value in this case is the geometrical roughness (SMD), which is part of the KES-F surface tester, which was the tester used to record the surface profiles, and may be considered analogous to the mean amplitude values calculated during characterisation.

The SMD value is the mean deviation of the surface thickness in micrometres and is calculated from the surface that is scanned by the KES-F surface tester, after it has been passed through the high pass filter described by equation 3.18, by the equation

$$\text{SMD} = \frac{1}{X} \int_0^X |T(x) - \bar{T}| dx \quad (4.2)$$

where x is the position of the contactor on the fabric, $X = 2 \text{ cm}$, $T(x)$ is the thickness

of the fabric at position x and \bar{T} is the mean value of T . The mean deviation of surface thickness is analogous to the mean amplitude values calculated for the 2-dimensional plots.

The KES-F produces one SMD value each for warp and weft directions. As both the chosen characterisation and KES-F calculations produce a set of two values, but that are separated by different aspects, they must be combined in some fashion to be compared. The KES-F values, as they are separated by direction of motion, are combined by the root-sum-square, as they are in the rendering strategy that is implemented in the tactile renderer. Unfortunately, there is no clear way to combine the two values from the characterisation calculations, which represent different frequency components. Possible considerations for the combination include whether or not one of the channels should be scaled with respect to the other, as well as the precise nature of the combination (sum or root-sum-square). However, because we are dealing with the objective case of comparing to the KES-F measurements, and not the subjective case of how the signal will feel to the user, it was decided that the comparison should be done without scaling of the amplitudes, using the sum of the two amplitudes as an estimate of the level of the signal.

4.2 Characterisation of Experimental Renderer

The initial characterisation work was done on the renderer and texture files that would later be used in the first experiment (section 6.2). The characterisation plots in figure 4.1 show the mean and CoV values resulting from the characterisation calculations for the 74 possible textures, for 72 different directions of motion at 10 cm s^{-1} , with no scaling of either the 40 Hz or 320 Hz signal.

The mean and CoV of the values that are output by the tactile renderer for the various textiles cover a full order of magnitude of each of the characterisation axes. This range of values suggests that the output from the renderer is dependent on the textile being represented and that the render is capable of delivering a wide range of sensations to the user. These characterisation values were later used as the basis for the selection of the experimental sample set, intending that the sensations presented to the users should cover as much of the available range as possible.

When the sum of the mean output A_{40} and A_{320} values is compared to the root-sum-square of the weft and warp KES-F SMD values (figure 4.2), we see that there is a strong positive correlation between the two ($R^2 = 0.95$). This shows that the two values are directly comparable, confirming that the amplitudes A_{40} and A_{320} relate to the surface profiles in the expected way.

Figures 4.3 and 4.4 show the effects that the speed of exploration has on the characterisation values (mean and CoV, showing the values associated with speeds of 3 cm s^{-1} and 30 cm s^{-1} respectively). These show that the speed of exploration can, potentially, have a strong effect on the sensations experienced by the users during surface exploration. Figures 4.5 and 4.6 show, more directly, how the mean and CoV of the amplitudes of the channels are effected by the speed of surface exploration.




Figure 4.1: The characterisation plots for the tactile renderer, based on the texture data used in the first experiment, showing the mean (top panel) and CoV (bottom panel) values of A_{40} and A_{320} during exploration of texture surfaces at 10 cm s^{-1} , for 74 different textures.




Figure 4.2: Comparison between the root-sum-square of warp and weft geometrical roughness from the KES-F and the sum of the mean 40 Hz and 320 Hz amplitudes calculated for the characterisation plots.

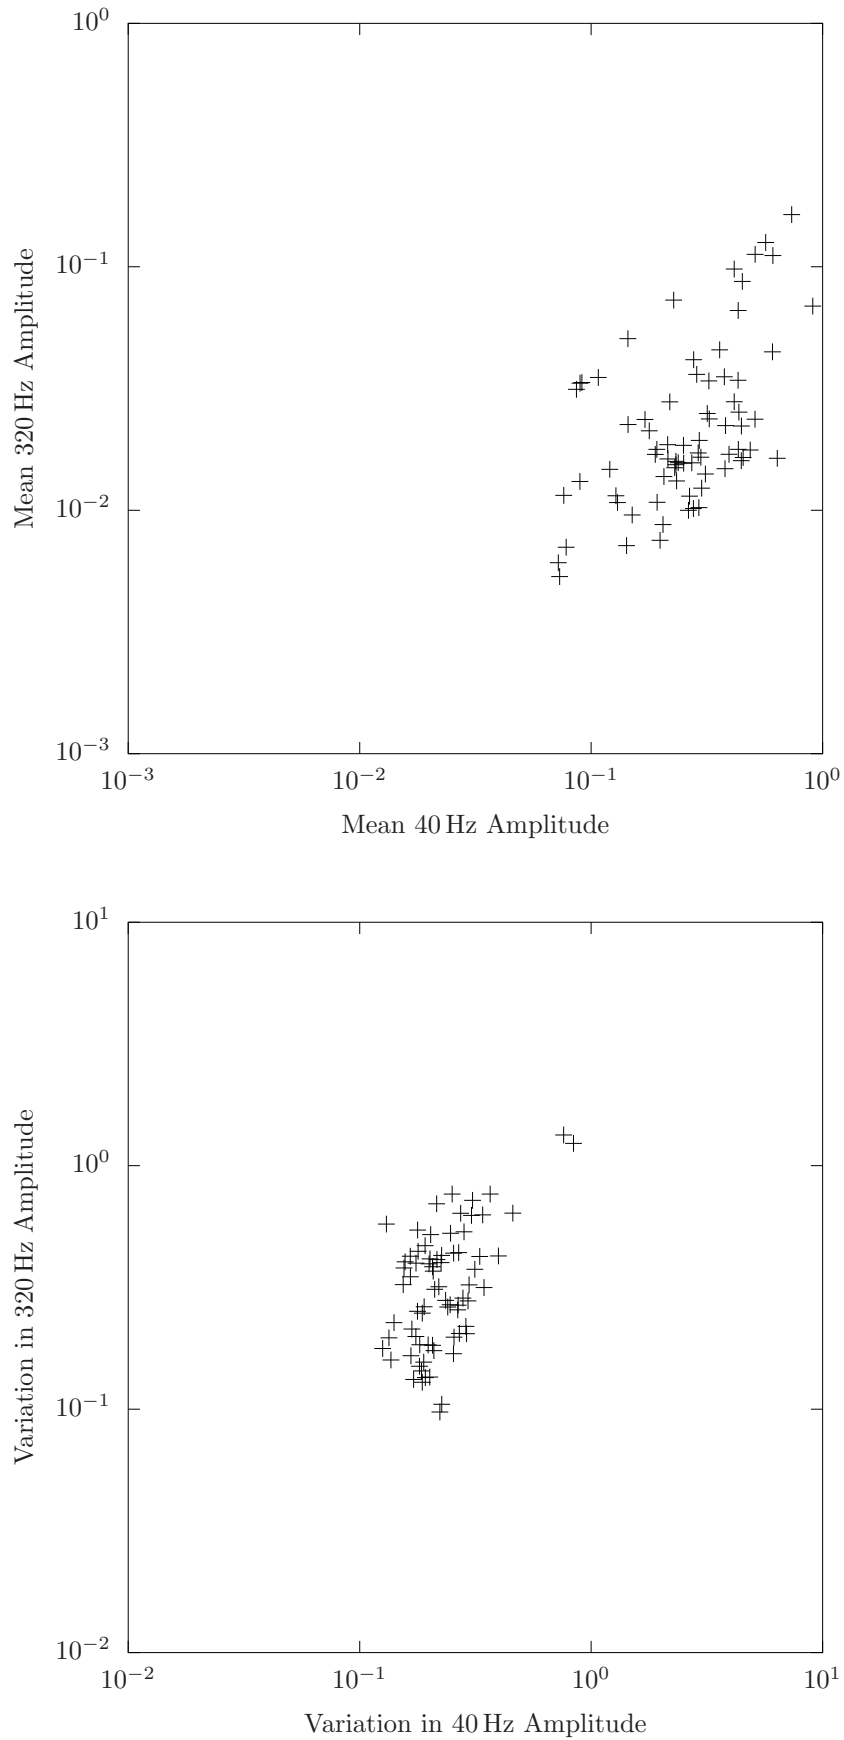


Figure 4.3: The characterisation plots for the tactile renderer, based on the texture data used in the first experiment, showing the mean (top panel) and CoV (bottom panel) values of A_{40} and A_{320} during exploration of texture surfaces at 3 cm s^{-1} , for 74 different textures.




Figure 4.4: The characterisation plots for the tactile renderer, based on the texture data used in the first experiment, showing the mean (top panel) and CoV (bottom panel) values of A_{40} and A_{320} during exploration of texture surfaces at 30 cm s^{-1} , for 74 different textures.




Figure 4.5: The mean amplitude values for the two channels as a function of exploration speed for 1 cm s^{-1} to 30 cm s^{-1} . 26_R, 29_R, and 37_B are highlighted in red in both channels.




Figure 4.6: The CoV of the amplitude values for the two channels as a function of exploration speed for 1 cm s⁻¹ to 30 cm s⁻¹.

Typically, in the low frequency NP channel, the maximum amplitude is experienced by the user at exploration speeds of less than 15 cm s^{-1} while the maximum amplitude is experienced in the P channel at speeds generally greater than 15 cm s^{-1} , although there are some exceptions, such as textiles 26_R (Organza), 29_R (Woven leisure), and 37_B (Woven outdoor leisurewear fabric) (which are, for clarity, highlighted as red in figure 4.5), which have maximum amplitudes in the P channel at speeds less than 10 cm s^{-1} . This is, however, matched by the maximum amplitude in the NP channel also being at a low exploration speed compared to the other textures characterised ($\leq 1 \text{ cm s}^{-1}$).

This pattern of maximum amplitudes as a function of speed is what would be expected, where features on the surface of the texture that have a particular spatial frequency would be experienced by the user with a temporal frequency closer to 40 Hz at slower exploration speeds and closer to 320 Hz at higher exploration speeds, with the speeds associated with the maximum amplitudes in each channel dependent of the spatial frequency of the features (higher frequency features require a lower exploration speeds to reach maximum amplitude). This suggests that the main surface features of the three highlighted textiles (26_R, 29_R, and 37_B) have spatial frequencies that are higher than the frequencies of the main surface features of the majority of the other characterised textiles.

The effect of exploration speed on the CoV measurements is a lot more varied, in both channels, as shown in figure 4.6. Some of the characterised textures increase in CoV with an increased exploration speed while others decrease in CoV for the same increase in speed. In the NP channel, many textures appear to go through an initial decrease in CoV followed by an increase once the minimum value is reached. In the P channel, the opposite trend seems to exist, with many textures increasing in CoV until a maximum is reached, which is followed by a decrease. However, as there are many exceptions to both of these trends, it is difficult to suggest that these are in any way indicative of a general trend of CoV as a function of exploration speed in either channel.

As the rendering algorithm has surface profiles from along both directions of each of the weft and warp axes available for its calculations, two of which (one per axis) are selected to be included in the calculations, the sensation that the user experiences is dependent on the direction of travel across the virtual surface. The distinction between directions is not significant for most textures, as the textures are similar in many, if not all, of the directions and so the user would experience similar sensations irrespective of the direction of travel. In some textures however there are differences between directions, resulting in the user experiencing a comparatively rough sensation in one direction of travel but a comparatively smooth sensation in the another direction.

The values that are calculated to characterise the mean and CoV of the surface texture along a direction of travel, either Up (U), Down (D), Right (R), or Left (L), and so quantifying this anisotropy, are the same as those calculated to find the mean and CoV from exploring the entire surface of the texture. In this case however, the calculation of the mean and CoV values is limited to only cases where the direction of the exploration vector is mostly pointed in the direction of interest (i.e. less than 45° from the direction of interest).

Figure 4.7 shows the root-sum-square of the characteristic mean amplitudes of the two channels along the U, D, R, and L directions of travel for various sample textures. There are four textures chosen to illustrate the effect of this anisotropy. Two of these textiles that have little anisotropy and vary only in overall amplitude; 37_R (upper left plot), which has a relatively low overall intensity in all directions and 39_R (upper right plot), which has a relatively high overall intensity in all directions. In contrast, two textile do have the discussed anisotropy; 03_R, (lower left plot), where the intensity is higher when travelling in the weft direction (R/L) than when travelling in the warp direction (U/D) and 54_R (lower right plot), where the intensity in the warp direction is determined by the direction of motion along that axis, with the intensity in the U direction being much lower than the intensity of the D direction.

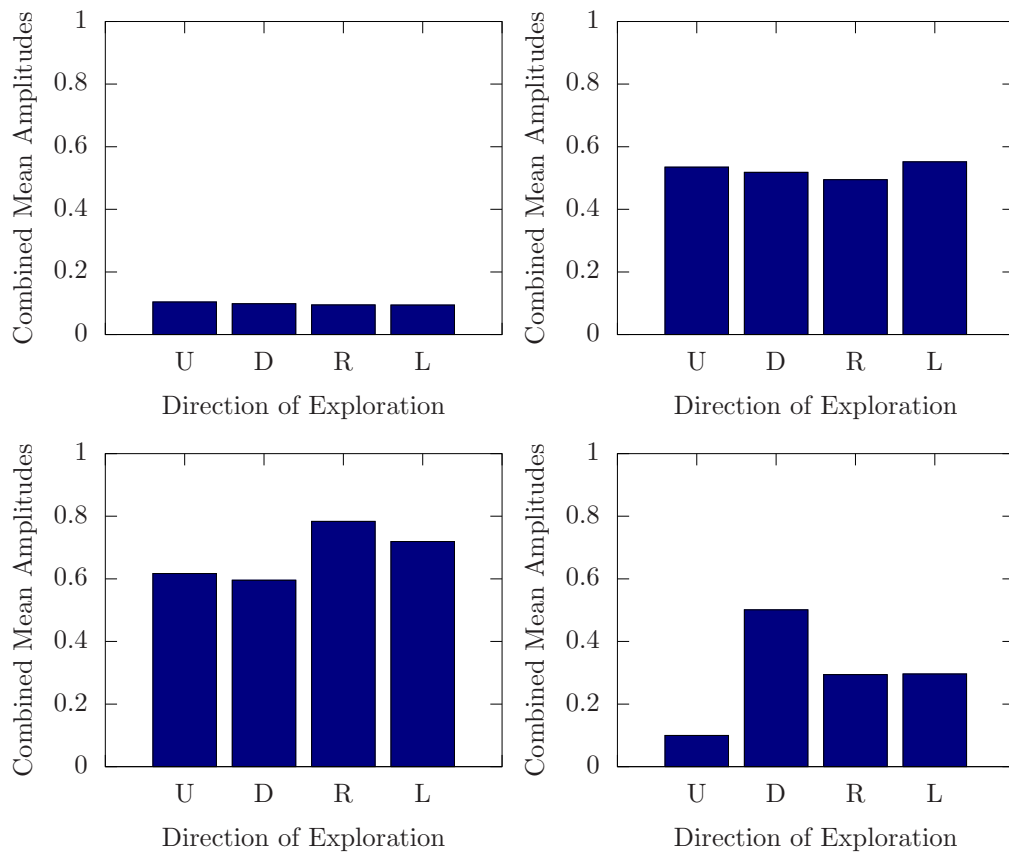


Figure 4.7: A comparison of the anisotropy of the mean intensity experienced by a user travelling in each of the four directions of motion across a surface, between textures 37_R (upper left), 39_R (upper right), 03_R (lower left), and 54_R (lower right).

These plots compare favourably to their real equivalent textures when one compares the surface feature and roughness sensation of the surfaces. Textures 37_R and 39_R are fairly smooth and rough, respectively, but they otherwise have a uniform surface and the sensation is similar irrespective of direction of exploration. The surface of 03_R features a very strong pattern of vertical ridges which makes the surface much rougher along the weft direction than in the warp. Visually, 54_R has no pattern that would suggest an anisotropy, but exploration with the fingertip reveals a definite pile to the surface. The user experiences a smooth sensation in one direction of motion (U) but the reverse of the direction (D) feels comparatively rougher. This surface anisotropy would reasonably serve

as an additional characterisation criteria, as two textures that might otherwise give similar characterisation values but have a different anisotropy would present different sensations to the user during the exploration of the virtual surface and so should be considered while investigating the recorded results from the various experiments.

However, attempts to reliably quantify the anisotropy of a textile surface have proven unsuccessful. An often suggested characterisation measure is based on the range of the directional mean amplitudes for a textile, divided by the overall mean level of the textile. Other proposed anisotropy characterisations share the similar idea of using a measure of spread of directional mean measurements as a means of quantifying the anisotropy. However this does not seem to be a reliable as a system for quantifying the anisotropy.

For example, if we compare the values of the proposed anisotropy measure (the range divided by the mean) for textiles 48_R and 49_R, which are 1.31 and 0.26 respectively, we would expect 48_R to feel much more anisotropic than 49_R when, in fact, the opposite is true. Both textiles are anisotropic but the effect is much less pronounced on 48_R than 49_R. A similar case exists between 43_R and 44_R, which have anisotropy values of 0.61 and 0.22 respectively. Again, this suggests that 43_R should be the more anisotropic textile of the two and, again, this is not the case.

Based on this measure, the textile 42_R has a value of 0.24 which would suggest that the surface is mildly anisotropic when it is, in fact, not anisotropic to the touch. The profile of the surface appears different in the weft than in the warp direction but these differences are fine enough that the user would not experience them. If the KES-F system recorded these differences as it scanned the surface profile however, this may explain the false anisotropy measurement. Similar effects may have occurred on the previously mentioned textile, with the KES-F recording surface features differently to how a user might experience the surface.

This observation adds weight to the suggestion that is based on the results of chapter 6 that, while the sensations experienced from the display are “textile-like”, the KES-F measurements are not themselves an appropriate source for the surface descriptions.

4.3 Comparison with HAPTEX Renderer

The characterisation work discussed so far was for the renderer that had already undergone a number of changes from the renderer that existed at the beginning of the project/end of the HAPTEX project. These changes are described in section 3.4, with the exceptions of the changes to the dynamic range and spectral balance (described in section 3.4.8), which were not applied as they were made based on the results of experiment one.

Figure 4.8 shows the characterisation plots generated using the renderer and texture files that existed at the beginning of the project, before any of the changes were made to the renderer. When compared to figure 4.1, the effects of these changes to the renderer on the output signal can be seen. Both versions cover similar dynamic ranges in both mean value and variation, with changes to the characterisation values associated with the individual

textures. There are some differences between the two sets of characterisation though. The dynamic ranges associated the mean amplitudes, while similar, are generally increased in magnitude in the experimental renderer over those from the original HAPTEX renderer. Particularly the dynamic range of the 40 Hz mean amplitude appears to be narrower in the experimental renderer when compared to the original renderer.

The most fundamental difference between the tactile renderer that existed at the start of the project and the one used during the experiments was the separation of the different directions of the motion along weft or warp axes (section 3.4.1). Figure 4.9 compares how the two renderer versions present 54_R to the user. To allow for a more direct comparison, the characterisation values for the different renderer versions were both calculated using the texture file that was part of the characterisation of the experimental renderer (section 4.2). To allow the HAPTEX renderer to be characterised with this texture file, it was necessary to adapt the file into a state comparable to that of the texture files that were used by the HAPTEX renderer (one matrix per degree of freedom). In both cases, the surface profile data representing a positive direction of motion was used.

As the renderer from the beginning of the project used information from only one direction of motion along each axis to produce the virtual textile, the user experienced the same sensation in both directions. Alternatively, the renderer used for the majority of the project, including the experiments, uses information from both directions along the axis, which allows for differences in the surface profile. As a result, the sensation from the display from the renderer used in the project, particularly in cases with this kind of anisotropy, is closer to the real texture than it was from the renderer that existed at the end of the HAPTEX project.

4.4 Characterisation of Scaled Texture Files

Analysis of the results from experiment 1 (section 6.2) suggested that the dynamic range of the renderer was too large, leading the subjects to incorrectly identifying the various textures. To address this, the output of the renderer was indirectly scaled. This was done by scaling the texture files that were used to create the virtual surfaces, according to the formula specified in section 3.4.8.

Figure 4.10 directly shows the effect of the scaling on the sum of the mean amplitude values. For comparison, a line that follows equation 3.24, which defined the scaling, is also shown. As expected, textures with relatively low overall intensity have been increased in intensity and so the sum of the mean amplitudes is increased, while those with relatively high intensity have been reduced in intensity, reducing the sum of the mean amplitude.

Figure 4.11 shows the characterisation plots that are generated from these scaled texture files. As expected, the figure suggests that the scaled texture files have a reduced dynamic range of mean sensation levels compared to the previous, unscaled texture files (figure 4.1), but with comparable ranges in the variation of the sensation. These characterisation values were used in the selection of the sample set that was used in experiments two (section 6.3) and three (section 6.4).

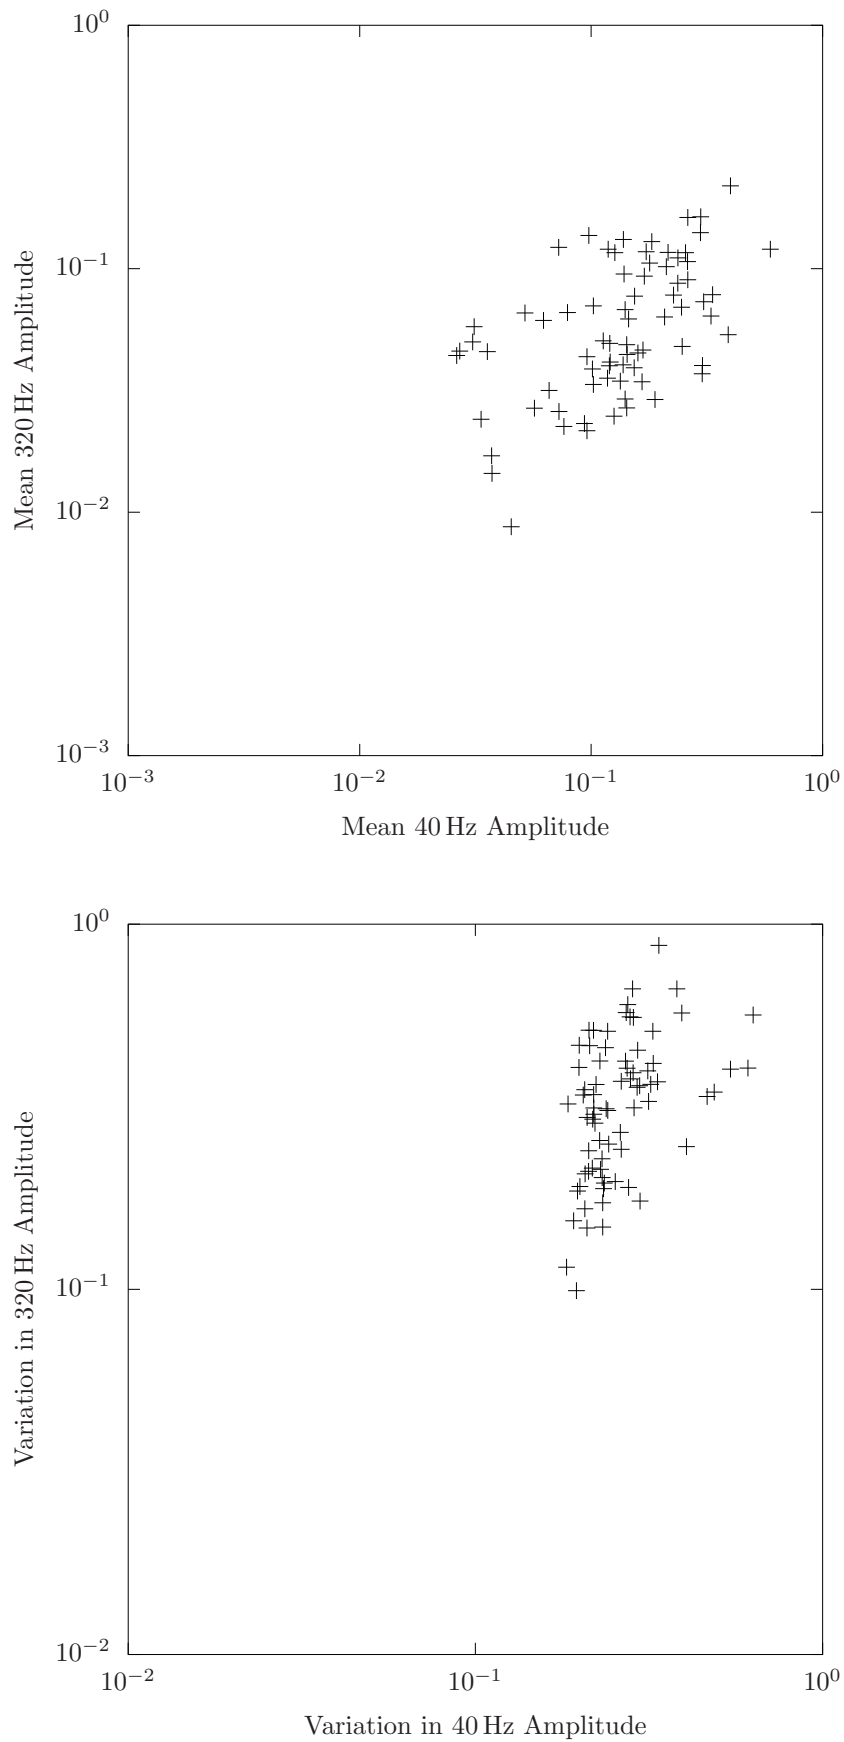


Figure 4.8: Characterisation plots based on the tactile renderer and texture data that existed at the beginning of the project, showing the mean (top panel) and CoV (bottom panel) values of A_{40} and A_{320} during exploration of texture surfaces at 10 cm s^{-1} , for 74 different textures.




Figure 4.9: A comparison of the mean intensity experienced by a user travelling in each of the four directions of motion, for texture 54_R, between the experiments renderer and the original HAPTEX renderer.




Figure 4.10: A comparison of the sum of the mean amplitudes between the scaled and unscaled texture files (crosses), as well as the scaling that was used (line).




Figure 4.11: The characterisation plots for the tactile renderer, based on the texture data used in the second and third experiments, showing the mean (top panel) and CoV (bottom panel) values of A_{40} and A_{320} during exploration of texture surfaces at 10 cm s^{-1} , for 74 different textures.

4.5 Summary

There was only limited characterisation done of the tactile renderer at the end of the HAPTEX project because of time restrictions. This chapter has corrected that omission by establishing a method for characterising the output of the tactile renderer. This characterisation is done in terms of the mean and CoV of the amplitudes that are output by the two channels of the renderer during active surface exploration. These characterisation values are shown in the form of a series of two-dimensional spaces which suggest how the sensations experienced by the user from the virtual textures may be related to each other. The results of this characterisation were then compared to measurements recorded by the KES-F to validate the characterisation method. The results of these characterisations are used to help select a sample set of textures for the experiments. It is intended that these sensations cover as full a range as possible, with the characteristic variables later being considered as possible dimensions for the multi-dimensional, perceptual space that is occupied by the virtual textiles.

Typically, characterisation is based on a typical exploration speed of 10 cm s^{-1} , although it is possible to characterise the renderer for other speeds, which will affect the characterisation values that are found for any particular texture. During characterisation, it is also important to consider the anisotropy of the surface texture, as the user would experience difference sensations from textures with different anisotropy, even if the textures have the same characterisation values otherwise.

Three different forms of the renderer are characterised within this chapter. The initial characterisation was the done for the render that was used for the first experiment. This version was used for validation and to explore the effects of speed and anisotropy on the characterisation. The second to be characterised was the renderer that existed at the end of the HAPTEX project, which underwent a few key changes before becoming the renderer used in the first experiment. This version is compared to the renderer used in the first experiment to explore the effect of the alterations on the characterisation values. The third renderer to be characterised was the renderer that was used during the second and third experiments. This version underwent some changes from the renderer used for the first experiment as a result of feedback from that experiment. This characterisation can also be used to observe the effect of these changes on the output signal.

Chapter 5

Public Demonstration

Demonstrating that the rendering system works within a laboratory setup is one challenge. Another, very different, challenge is demonstrating its effectiveness in the real world, or a public space. To that end, over the course of the project, the renderer has been presented to the public on a number of occasions in an attempt to gauge reaction to the device. Most of the public presentations simply involved the renderer displaying a variety of the stock textures that were part of the HAPTEX project, in an abstract setting (such as university open days or science fairs, as in figure 5.1). The public presentation described in this chapter, however, was where the renderer was used to provide a tactile recreation of an archaeological textile.



Figure 5.1: The setup being used as part of “The Big Bang South West Fair”.

This chapter begins by describing the wider environment in which the demonstration took place, including the project that it was a part of and some of the other demonstrations involved in the project. The chapter then goes on to describe which artefact was being virtually recreated and how the renderer was used in this situation. The final part of this chapter details the feedback from the visitors about the display and how it compares to

feedback given for physical recreations of the same artefact.

5.1 Demonstration Environment

In January of 2011, the tactile render was taken to the National Museum of Scotland as part of the project “Touching the Untouchable”, a project by the Department of Archaeology of the University of Exeter, funded by a network grant between the EPSRC and the AHRC, as part of the Science and Heritage programme of the AHRC. The aim of this project was to explore the use of various techniques in order to allow visitors to handle the museum’s artefacts, which would otherwise be inaccessible, through the use of recreations generated from a variety of techniques. In addition to some physical recreations of the artefacts, a number of the chosen display techniques used virtual recreations, through touch-based computer interfaces, to allow the visitor to explore the surfaces of the artefacts. The suggested advantages of this are that the virtual recreations allow for multiple objects within a portable computer systems (as opposed to transporting large numbers of physical recreations), and that physical recreations may be fragile and hence be damaged by public handling. This project’s principal investigator was Dr Linda Hurcombe of the Department of Archaeology, University of Exeter, with co-investigators of Dr Mark Paterson of the Department of Geography, University of Exeter, and Prof Stephen Brewster of the Department of Computer Science, University of Glasgow.

In addition to the tactile renderer, other possible solutions were presented at the same event, for a number of different museum artefacts. Figure 5.2 illustrates two other suggested solutions. Unlike the renderer, which is recreating the surface of a textile artefact, these solutions are adding touch-based sensations to a three-dimensional model. In this case, the model used is one of the “Lewis Chessmen”. These demonstrations were developed as a collaborative effort between Dr Mark Wright, then of the Department of Informatics at Edinburgh University, and Dr Mariza Dima of the Edinburgh College of Art and Edinburgh University. Of the two solutions depicted in the figure, one is a physical solution, and the other is a virtual solution.

The “Lewis Chessmen” are a set of 78 chess pieces from the 12th century that are made from Walrus Ivory and Whales’ teeth. They were discovered in Uig bay on the Isle of Lewis in 1831 and constitute one of the few surviving medieval chess sets. The ownership and exhibition of the set is divided between the British Museum in London, which has 67 of the original pieces, and the National Museum of Scotland in Edinburgh, which has 11 pieces. The piece being recreated by these demonstrations is the queen piece, a replica of which is shown in figure 5.3, which is one of the pieces on display in Edinburgh.

Ghost touch, which is the physical solution on the right hand side of figure 5.2, uses a physical recreation of the real artefact that is mirrored along the vertical axis and placed in front of the artefact. Figure 5.4 shows a diagram of the ghost touch setup. The visitor can easily reach and explore the recreation with touch but it is visually obscured view by a cover. The result is that the visitor will have the haptic sense of exploring the recreation but be looking at the artefact, where they see the reflection of their hand. The reflection



Figure 5.2: Two other strategies presented at ‘Touching the Untouchable’, allowing for physical (ghost touch, right) and virtual (haptic pen, front) explorations of a “Lewis Chessmen” piece.

of the visitor’s hand in the display case both aids in navigation and adds to the illusion of exploring the surface of the artefact. To aid in the illusion, both the recreation of the artefact and the inside of the cover are darkened so the visitor sees only the reflection of their hand in the exhibit case and not a reflection of the recreation. Figure 5.5 shows the ghost touch setup in use at the National Museum of Scotland.

The haptic pen, the virtual solution shown at the front of the figure 5.2, uses a haptic stylus to allow visitors to explore the surface of the artefact. Figure 5.6 shows a diagram of the haptic pen setup. The stylus is connected to a force feedback device which is programmed to follow the contours of the artefact’s surface. The visitors can then use the stylus to explore the virtual surface of the artefact as much as they wish. Figure 5.7 shows the haptic pen setup in use at the National Museum of Scotland.

5.2 Demonstration Setup

In this project, the tactile renderer was used to display a recreation of the surface of an archaeological textile, the “Falkirk Tartan”, which is one of the earliest examples of Scottish tartan in existence, being from the 3rd century (figure 5.8). For this, the Exeter tactile display was placed in front and to the right of the textile in its display case (figure 5.9). The virtual textile used for this installation was not one of the stock set created for the HAPTEX project. A custom texture was instead created through the modification of one of the stock fabric file that was found to give similar sensations to those created by the tartan. Feedback from experienced users gave clues about the modifications needed to convert the stock texture into one closer to the sample. The virtual recreation of the archaeological tartan, which also included a representation of fingertip within the virtual



Figure 5.3: A replica of the queen “Lewis Chessmen” piece.

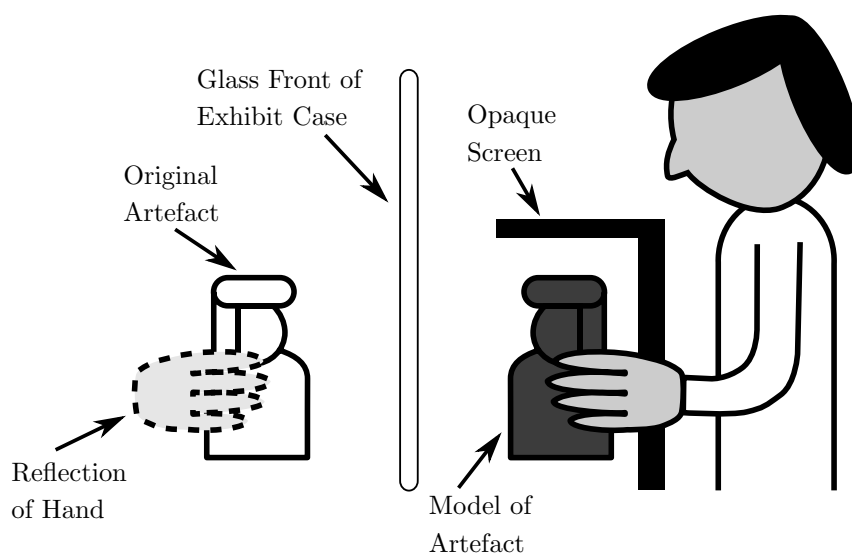


Figure 5.4: Diagram of the ghost touch setup. The model of the artefact is mirrored along the vertical axis. The user explores the model of the artefact with their hands, seeing the reflection of their hands in the exhibit case.



Figure 5.5: Use of the ghost touch setup to explore the surface of the “Lewis Chessmen” piece.

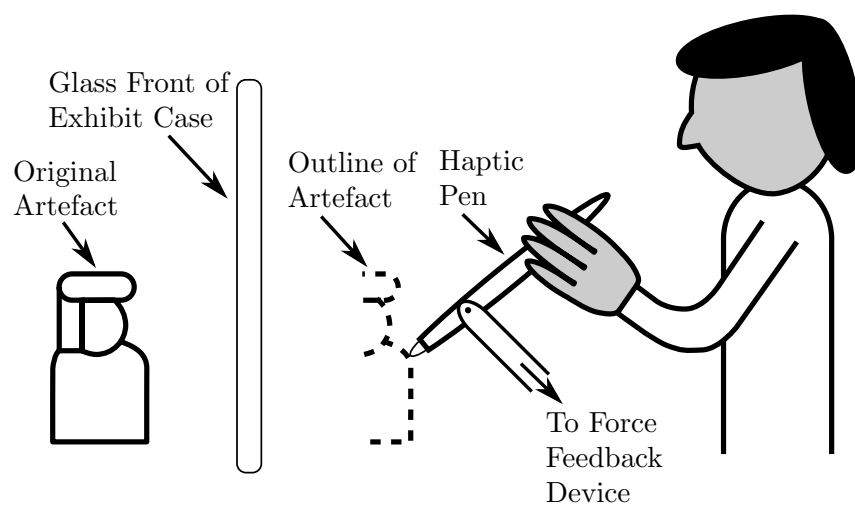


Figure 5.6: Diagram of the haptic pen setup. Users can use the haptic pen to explore the contours of the surface of the artefact.



Figure 5.7: Use of the haptic pen setup to explore the surface of the “Lewis Chessmen” piece.

environment, was also viable in the same position in space as the archaeological textile, as an ‘overlay’.

Figure 5.10 shows the installation in its ready-to-use state. While the installation was available to the museum visitors, they were invited to explore the surface as they wished to do so. By moving the display around, with their fingertip rested on the display, the visitor experiences the sensations equivalent to those associated with exploring the surface of the artefact using their fingertip.

For contrast, two other recreations of the tartan were presented at the event, both of which are shown in figure 5.11. These recreations are facsimiles of the artefact, created from different techniques and they were shown grouped with a number of similar recreations that were based on other artefacts (figure 5.12). The first of these facsimiles is a small sample of a modern recreation of the artefact textile, and the second shows a 3D print made from a laser scan of the original artefact. As with the tactile renderer, visitors were invited to explore these recreations to experience the sensations similar to those that would be experienced from exploring the artefact.

5.3 Feedback

Over the course of two days, a wide range of visitors to the museum took advantage of the ‘Touching the Untouchable’ installations, including around 70 who tried the tactile renderer. To gauge the efficacy of the displays, the visitors were asked to give their feedback on what they thought of them, both in person and through a questionnaire.

Although many visitors seemed to show some scepticism regarding the use of the renderer, this is not unexpected given the museum environment where touch interaction is usually




Figure 5.8: The textile that was the subject of the recreation for “Touching the Untouchable” project: the original, archaeological, textile (right) along with a modern recreation (left).

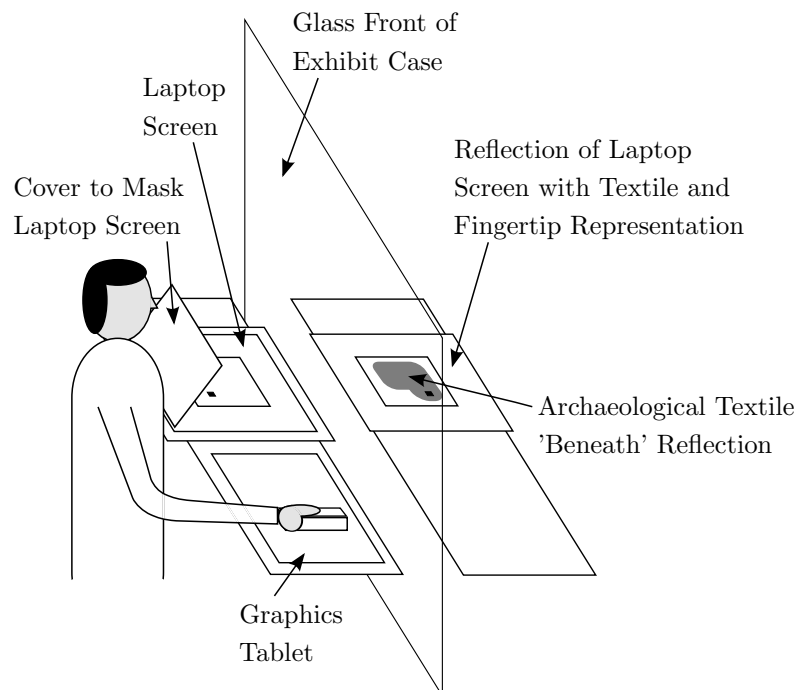


Figure 5.9: Schematic of the installation used at the museum. The user has their right hand on the Exeter tactile display and sees the reflection of the laptop display, showing the textile image and the fingertip representation. The laptop display is inverted to compensate for the inversion caused by the reflection in the glass front of the exhibit case.




Figure 5.10: The installed tactile renderer, featuring the Exeter tactile display and a visual overlay to navigating the virtual environment.

strongly discouraged. Once this initial scepticism was overcome and some initial familiarization given on the use of the display, the visitors (including children) typically responded positively to the display. There were no examples of anyone willing to use the device but who was unable to do so. The success of the device in this situation perhaps indicates the “naturalness” of the experience that is presented.

After the visitors had used the renderer, they were invited to give feedback on their experiences through a survey (written by Dr Linda Hurcombe). By the end of the event, 15 groups of various sizes participated in the survey, totalling 24 adults and 6 children. Table 5.1 shows how many groups chose a particular strength of agreement in response to the presented statements. There was also the option at the end of the survey for the participant to provide extra comments/feedback.

Most groups had some level of agreement with the presented statements, including suggestions that the installation gave an impression of the feel of the object which leads to a better understanding of the artefact. As the response to the final statement of the survey suggests, visitors were generally happy with the display and the virtual texture as presented in this installation. Many suggested, both verbally and through comments, that the system would help engage people (particularly children) into the museum experience.

However, one consistent piece of negative feedback regarded the sensation felt by the user from the display. Many visitors verbally noted that the sensation felt prickly under their fingertips, with one visitor commenting on the survey that

“It was slightly odd sensation and took a little imagination to believe I was touching the fabric.”

The addition of caps to the contactors, as described in section 3.4.7, addressed this feedback by defusing the sensation from each of the contactors across a wider area of the



Figure 5.11: Two facsimiles of the Falkirk tartan. The upper panel shows a modern recreation of the textile (woven by Dr Linda Hurcombe). The lower panel shows a 3D print made from a laser scan of the original artefact (scanned and ordered from the National Conservation Centre, Liverpool).



Figure 5.12: A group of the facsimiles presented at the “Touching the Untouchable” event, with the recreations of the “Falkirk Tartan” on the left of the image.

Table 5.1: Results of the survey conducted to gather feedback about the tactile renderer installation (percentages from a total of 15 surveys, representing 24 adults and 6 children in total).

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
It is important to engage more senses than the visual as part of the museum experience.	00.0	00.0	13.3	40.0	46.7
The installation/replica gave a sense of how the ancient object would feel.	06.7	06.7	06.7	60.0	20.0
The installation/replica was straightforward to use.	00.0	06.7	06.7	46.7	40.0
Because of the installation/replica there was a better understanding of the ancient objects.	00.0	13.3	26.7	40.0	20.0
Overall, the installation/replica enhanced the museum experience.	00.0	00.0	13.3	46.7	40.0

finger tip's surface. The result was that the surface of the display felt smoother to the touch, both during active surface exploration, where the movement of the pins felt less sharp, and outside of surface exploration, where the display's surface felt more like a continuous surface.

The feedback obtained for the renderer can be compared to feedback from the same survey that was obtained from the facsimiles presented (table 5.2, which was received from 24 groups of people, totalling 34 adults and 28 children). While this feedback is for recreations of a wider range of artefacts, it does include the two recreations of the "Falkirk Tartan".

Table 5.2: Results of the survey conducted to gather feedback about the presented facsimiles (percentages from a total of 24 surveys, representing 34 adults and 28 children in total).

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
It is important to engage more senses than the visual as part of the museum experience.	00.0	00.0	00.0	41.7	58.3
The installation/replica gave a sense of how the ancient object would feel.	00.0	00.0	08.3	50.0	41.7
The installation/replica was straightforward to use.	00.0	00.0	08.3	45.8	45.8
Because of the installation/replica there was a better understanding of the ancient objects.	00.0	04.2	16.7	41.7	37.5
Overall, the installation/replica enhanced the museum experience.	00.0	00.0	00.0	25.0	75.0

These results suggest that people prefer the physical recreations to the digital one, with comments on the feedback often comparing the two, such as

"The replicas were more accessible and informative than the digital interfaces."

However, as has been suggested previously, there are a number of potential drawbacks to the use of a large number of physical recreations (such as transport, maintenance, etc.), and users are able to use the tactile renderer quite naturally. Therefore, there may be some merit in sacrificing some small fidelity of the presented recreation in favour of the virtual display, which is able to recreate the surfaces of a large number of artefacts and is relatively easier to transport. This fidelity sacrifice may also be reduced by future improvements to the renderer, such as the addition of the previously mentioned contactor caps, and a more appropriate/accurate recreation of the surface of the artefact textile.

5.4 Summary

Although the renderer has been shown to be successful in a laboratory environment, as explored in chapter 6, in order for the device to be useful for general purposes, then the reaction of the public must be addressed. To that end, the tactile renderer has been presented to the public on number of occasions, including one presentation at the National Museum of Scotland where it was part of a project investigating possible ways to bring touch-based interaction into a museum setting.

The feedback that was received from the interaction of visitors with the device over the course of the two day installation was generally positive. It was a common suggestion that it would help people engage more with the museum environment, although people generally preferred to explore physical facsimiles of the artefacts, such as 3D prints. Some avenues for improvement of the renderer were also identified from the feedback, the most notable of which was the addition of caps to the moving pins to reduce the unevenness of the contactor surface on the fingertip.

Chapter 6

Virtual Texture Evaluation

Time limitations at the end of the project resulted in very little evaluation of the tactile renderer being done during the HAPTEX project (the evaluation that was done is described in section 3.3.5). This chapter remedies that omission by evaluating the renderer that was developed for this study, which is a modified version of the renderer that existed at the end of the HAPTEX project.

In order to demonstrate if the chosen rendering strategy is effective in generating virtual versions of specific real textured surfaces, it is necessary to establish the subjective similarity between the sensations generated by the active exploration of the surfaces of a number of different textiles. To do this, the first experiment compared virtual surface textures to their real equivalents by having the subjects rank a group of virtual textiles in order of similarity to a target, real textile.

Based on the feedback that was received from this experiment, changes were made to both the dynamic range and the spectral balance of the renderer, as documented in section 3.4.8, and the second experiment that was performed followed the same protocol as the first experiment in order to assess the effect that these changes had on the virtual surface textures and their similarity to their real equivalents.

The third experiment involved only the virtual textiles, replacing the real target textile with a virtual one, so as to investigate the discriminability between virtual surface textures. The render and experiment were otherwise unchanged from the second experiment. If the virtual textures are more correctly identified in this experiment than in the previous one, it suggests that the virtual textiles are not appropriate recreations of the real textiles.

From this third experiment, measures of dissimilarity between the virtual textures were calculated and used to create multi-dimensional solutions for the textures which could be then be mapped onto various possible perceptual dimensions, such as the characterisation spaces that were defined in chapter 4, in order to determine what characteristics the subjects use to discriminate between the various textiles.

Finally, the three experiments were repeated with two experienced users as the subjects, so as to illustrate the effect of familiarity and repeated use on the users ability to discriminate between the various surface textures.

These experiments were conducted after receiving approval from the University of Exeter’s ethics committee on the September 16, 2011. Details of the ethics application, as well as copies of the information sheets and consent form that was given to potential subjects can be found in appendix D.

6.1 Experimental Design

A number of possible avenues were considered for the evaluation of the similarity between the real and virtual textiles. The first possibility was to perform a set of inter-textile comparisons on the virtual textiles, in order to determine the relevant perceptual dimensions, and how the various textiles are positioned within the space defined by these dimensions. This procedure could then be repeated for the set of real textiles and the two resulting spaces compared. The limited evaluation carried out during the HAPTEX project, as described in section 3.3.5, was a cut-down version of this procedure, where it was constrained to only one dimension. A limitation of this procedure is that the two sets of textiles (real and virtual) are not directly compared, making it difficult to establish true similarity.

Two possible methods for a direct comparison between the real and virtual textiles are either direct estimates for the similarity between a virtual surface texture and its real counterpart, which may be problematic in some scenarios where the two textures are obviously not equivalent (such as the two textiles featuring a different anisotropy), or a forced-choice matching task, which is the method preferred in this study.

All of the experiments described in this study followed a similar design, using the setup and procedures that are described below, but with differences in the set of sample textiles being used, the subjects who participated in each of the experiments, and the nature of the “target” textile that the subjects are compared against. Approval of the ethics committee of the University of Exeter to perform these experiments, based on application “2010/270 - Evaluating a Setup for Tactile Rendering”, was received on September 16, 2011.

Volunteers who were interested in participating in the experiments were given the information sheets, shown in appendix D, which detailed the nature of the experiment generally (figure D.1) and of the tactile display specifically (figure D.2). The volunteers who were happy to participate in the experiment, after reading the information sheets provided, were asked to sign the consent form (figure D.3) so as to demonstrate that fact, after which the experiment began.

The setup that was used for all of the experiments detailed in this study is shown in figure 6.1. The subjects were seated in front of the setup where they could comfortably reach both the Exeter tactile display/graphics tablet and, if necessary, the real textiles. The computer monitor displayed a representation of the virtual workspace that was being presented to the subject, which the subjects could use to navigate the workspace with the aid of a representation of the tactile display (represented as an array of green squares) whose location on the screen reflects the display’s location within the virtual space. The setup is controlled by the author using a laptop which is running the tactile rendering software (not pictured).




Figure 6.1: The tactile rendering setup used for the experiments.

After the subjects were seated, they were taken through an initial familiarisation with the tactile display. During this familiarisation, the subjects were presented with a series of virtual environments, each of which featured a variety of virtual textiles. The subjects were instructed on how best to hold the display during use, as in figure 6.2, and were invited to freely explore the surface textures of the presented textiles as they wished. During their explorations, subjects were advised to use a combination of circular motions, so as to evaluate the overall feel of the textile, and linear movements (up, down, left, right), to evaluate the directional anisotropy of the textiles. It was also advised that subjects use a similar exploration strategy during the experiment. This familiarisation was intended to serve the dual purposes of both allowing the subjects to become comfortable in the use of the display outside of the experiment, as well as to introduce them to the kinds of sensations that they would experience during the experiment.

After this familiarisation, the subjects were presented with a series of test items, each one consisting of a set of five textiles; the four virtual textiles to be ranked, which were arranged in an approximate square with labels numbered 1 to 4, and the “target” textile (which may be real or virtual, depending on the experiment). Throughout the experiments, all of the textiles, either real or virtual, were presented as squares with a size of 10 cm-by-10 cm.

From the full available texture set, as listed in table A.2, each of the experiments used a sample set of sixteen textures which were selected, based on the characterisation of the output of the tactile renderer that is described in chapter 4, with the intention of covering the dynamic range of the available textures. For each test item in the experiments, the “correct” virtual textile (which matches the “target” textile) is presented with three “incorrect” virtual textiles. These “incorrect” textiles were selected from the set of fifteen available “incorrect” textiles. All “incorrect” textiles were presented with the “correct” textile at some point, making a total of five virtual environments available for each of the sixteen “correct” textures, each with a different set of “incorrect” virtual textiles, totalling



Figure 6.2: How the tactile display is held during use.

eighty test items for each experiment.

For each test item, subjects were free to explore the virtual environment as they wished and were asked to rank the virtual textiles in order of similarity to the “target” textile. The subjects were asked to make a choice even if they claimed to be unable to distinguish between the textiles in level of similarity (forced choice, 25% chance score). Test items were presented to the subjects in blocks of ten to allow rest periods during the experiment, as long periods of exposure to the display has a numbing effect on the fingertip.

During normal operation of the renderer, any virtual textile that is displayed is shown with a visual image that reflects the textile that is being recreated. As this would give a visual indicator of which virtual textile is the correct choice for the subject to make, over the course of the experiment, all of the virtual textiles in any given virtual environment were made to share the same image as that of the “target” textile. Additionally, ear-muffs were used to minimise any acoustic clues from the display during the subjects’ exploration of the surface.

In addition to their ranking orders of each test items, the age, handedness, and gender of the subjects was recorded (tables 6.1, 6.6, and 6.9). Irrespective of their own handedness, during the experiment the subjects were asked to use only their right hand to explore the surfaces of both the real and virtual textiles.

6.2 Experiment 1

This experiment, also described in [109], is a first attempt to gauge the effectiveness of the tactile rendering strategy using the experimental design described in section 6.1. Over the course of the experiment, a set of eight subjects (details listed in table 6.1), were presented with each of the test items in turn to give their rankings of the similarity between the

virtual textiles and the target textile. The test items were generated from the sample set listed in table 6.2, whose locations in the characterisation space given in figure 6.3.

Table 6.1: The recorded details of the eight subjects who participated in experiment 1.

Subject #	Gender	Age	Handedness
1	M	37	R
2	F	30	R
3	F	25	R
4	M	27	R
5	M	22	L
6	F	24	R
7	M	26	R
8	M	25	R

Table 6.2: The sample set used in experiment 1, with the textile number, the face used (Right or Back), and the textile’s description.

Number in Set	Texture Identifier	Textile Description
01	45_B	Warp knit for car seats
02	03_R	Cord
03	36_R	Men’s woven overcoat fabric (twill)
04	30_R	Tulle
05	26_R	Organza
06	31_R	Warp knitted tricot-satin
07	40_B	Warp knitted jersey-based fabric
08	49_B	Laminated warp knit for car seats
09	54_R	Laminated warp knit for car seats
10	09_R	Wild silk (dupion)
11	35_B	Men’s woven overcoat fabric
12	44_R	Warp knit for car seats
13	11_R	Flannel
14	43_B	Warp knit for car seats
15	37_R	Woven outdoor leisurewear fabric
16	27_R	Fleece

For this experiment, the test items consisted of a virtual workspace consisting of four virtual textiles (figure 6.4), which were ranked in order of similarity to the target textile which, in this case, was a real textile with a an available surface area that was the same as the virtual textiles (figure 6.5). The spectral balance of P:NP that was used for the experiment was 0.7:1.0.

Table 6.3 shows the mean values of the textile identification scores, cumulative over the subjects’ rankings. From these results, it can be seen that the textures were correctly identified in $38.1 \pm 3.1\%$ of the time, which is significantly greater than the chance result of 25% (t-value of 4.23, $p < 0.0025$). These results can be further broken down by scores associated with an individual subject (figure 6.6) or texture (figure 6.7). Both of the figures show the percentage of tests where the “correct” textile was given a particular rank of similarity to the target (“1st” shows the percentage of tests where the correct textile was ranked as first in similarity, “2nd” for ranked second in similarity, etc.).




Figure 6.3: Characterisation plots for the tactile renderer for the sample set of 16 textures used in the first experiment, showing the mean (top panel) and CoV (bottom panel) values of A_{40} and A_{320} during exploration of texture surfaces at 10 cm s^{-1} .




Figure 6.4: An example of one of the virtual workspaces presented during experiment 1, showing four virtual textiles that are to be ranked against the target textile.




Figure 6.5: An example of the presentation of a real textile as a target.

Table 6.3: Cumulative scores (% correct) for matching real and virtual textiles in experiment 1: Key: 1 = real textile matches 1st-ranked virtual textile; 2 = real textile matches 1st- or 2nd-ranked virtual textiles; 3 = real textile matches 1st-, 2nd- or 3rd-ranked virtual textiles; 4 = real textile matches 1st-, 2nd-, 3rd- or 4th-ranked virtual textiles.

Similarity Ranking	Mean Score	Standard Error
1	38.1	3.1
2	65.3	3.7
3	90.0	1.7
4	100.0	0.0




Figure 6.6: A breakdown of the results of experiment 1, showing the percentage of times each of the 8 subjects assigned the correct textile with a particular rank of similarity to the target.

Figure 6.6 shows that there is a lot of variation among the individual subjects but that there is also some consistency between the subjects. Each of the subjects correctly identified the textile in 25% or higher of cases, with the subjects selecting the correct textile as either their first or second choice in over 50% of cases. Overall, there are no examples in this experiment of a subject performing worse than would be expected from chance.

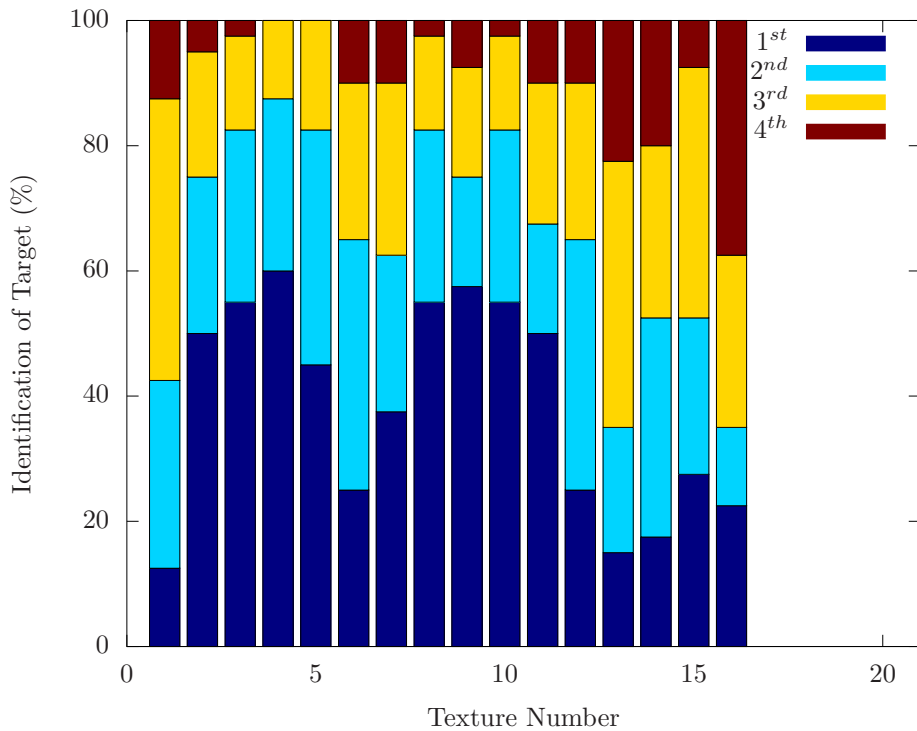


Figure 6.7: A breakdown of the results of experiment 1, showing the percentage of times each of the 16 correct virtual textiles was assigned a particular rank of similarity to the target.

There is much less consistency between the identification of the different textiles, however. As figure 6.7 shows, a number of textiles were correctly identified over 50% of the time, while some others were identified less than 20% of the time. This may suggest that certain virtual textiles do not accurately recreate the sensations of exploring the surface of the equivalent real textile, leading to confusion among the subjects as to the correct textile.

Some of the textures within the sample set, specifically textures 02, 09, and 12, have strongly anisotropic surfaces. We would expect the anisotropy to be a clue for identification of the correct textile. However, as figure 6.7 shows, these textiles are not identified with any more success than the non-anisotropic ones, with texture 12 having a particularly low first-ranked identification rate when compared to the other textures ($< 30\%$). This corroborates the observations regarding the use of anisotropy as a characterisation measurement, as described in chapter 4, in that the textile files used in these experiments are not accurately reproducing the degree of anisotropy present in the real textiles.

The results in table 6.3 are consistent with evaluation work that was done on the renderer previously as part of the HAPTEX project, described in [98]. In that work, the subjects were presented with a target real textile and two virtual textiles, one of which was the virtual recreation of the target, and they were asked to identify which of the two virtual

textiles was closest to the real textile. These results showed that as the number of items in the sample set (test items) increases, starting from two clearly distinguishable textiles and adding more “intermediate” textiles to the sample set, the percentage of correct identifications decreases (figure 6.8). As we have a large number of sample in this experiment (sixteen), the low identification score is, therefore, not entirely unexpected.




Figure 6.8: Mean scores from 5 subjects for matching real and virtual textiles, as a function of the number of textiles in the stimulus set (recreated from [98]).

Analysis of the error patterns within the results, indicates that the errors are often systematic, rather than random errors. When we consider the number of times the “most popular” choice is made by the subjects, i.e. the mode choice, and compare it with what we would expect were the subjects simply selecting textiles randomly, as derived in appendix C, it can be seen that higher mode values appear more often, and the lower mode values appear less often, than we would expect simply from chance, as illustrated in figure 6.9. The mean mode value for the experiment is 4.54 ± 0.13 , with the expected chance mean value of 3.54. The t-value associated with these results is 7.58, which gives a p value of less than 0.0001, suggesting this difference is statistically significant.

Figure 6.10 compares the distribution of mode choices between the cases where the mode choice corresponds to the correct textile (42 of the tests), and the cases where it does not (38 of the tests). There appears to be very little difference between these two distributions. The case of the most popular choice being the correct choice has a higher proportion of modes with a value of 4 and a lower proportion of modes with a value of 5 when compared to the case of the most popular choice not being the correct choice. Aside from this, the two distributions are otherwise equivalent. This is reflected in the similar mean mode score of the two distributions, being 4.50 ± 0.13 for cases where the mode is the correct choice and 4.58 ± 0.13 for cases where it is not. This suggests that the subjects are just as




Figure 6.9: From experiment 1, breakdown of the percentages of how the “most popular” (mode) choice scored across all the test items, compared to what would be expected from chance.

likely to settle on a textile as their collective first choice irrespective of whether it is the correct choice or not, further suggesting that the errors in identification are systematic rather than random.

From these results, in combination with the subjects comments, it was concluded that the textures presented to the test subjects are discriminable and “textile-like”, but the corresponding real and virtual textiles are not always well matched, and sometimes the best match between the real and virtual textile is not the “correct” choice. To explore this possibility, table 6.4 shows which textiles the subjects chose as their first choice matches (columns) for the various real textiles (rows). As there are 8 subjects, and 5 environments for each target textile, each row of the matrix has a total of 40.

When the position of the textile that is chosen as the best match to the target is compared to the position of the target textile within the characterisation space, there emerges a pattern of rougher real textiles tending to be mistaken for less rough virtual textiles and smoother real textiles tending to be mistaken for less smooth virtual textiles. This is illustrated in figure 6.11 which, for clarity, only shows confusions where 4 or more subjects chose the textile as their first choice match to the target.

Based on these results, the dynamic range of the renderer was reduced, as described in section 3.4.8, as well as further investigation into the optimum spectral balance between the P and NP channels (0.83:1.0, P:NP). The aim of these changes is to improve the fidelity of the virtual textures as matches to the real ones on which they are based. The experiment was repeated to directly investigate the effect of these changes on the identification of the




Figure 6.10: A comparison of the distribution of mode choices between cases where the mode is the correct choice and cases where it is not the correct choice.

Table 6.4: The number of subjects that chose the virtual textiles (column) as their first choice match to the target textile (row). The main diagonal is the sum of all 5 available environments.

		Virtual Textile															
		01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16
Real Textile	01	05	06	04	04	01	07	04	06	00	03	00	00	00	00	00	00
	02	02	20	03	03	02	03	02	02	01	02	00	00	00	00	00	00
	03	00	02	22	05	00	03	03	01	03	01	00	00	00	00	00	00
	04	02	02	02	24	00	02	01	03	02	02	00	00	00	00	00	00
	05	00	01	01	00	18	01	01	01	03	04	05	01	01	00	02	01
	06	01	00	00	00	03	10	01	02	03	05	05	03	01	02	03	01
	07	00	02	01	00	02	05	15	07	01	02	03	02	00	00	00	00
	08	00	01	02	02	02	03	01	22	01	03	01	02	00	00	00	00
	09	00	00	00	02	03	02	01	01	23	03	01	01	01	01	01	00
	10	00	01	00	00	04	01	03	01	00	22	03	05	00	00	00	00
	11	02	02	01	01	02	02	02	03	01	01	20	02	01	00	00	00
	12	00	01	01	00	00	06	03	06	07	04	01	10	00	00	01	00
	13	00	03	00	04	05	03	02	05	03	02	05	02	06	00	00	00
	14	01	01	00	01	03	01	00	04	02	02	04	04	05	07	02	03
	15	00	00	00	03	02	01	00	02	03	03	04	05	03	01	11	02
	16	00	00	00	02	02	01	01	05	03	04	05	03	02	02	01	09




Figure 6.11: The error patterns obtained from the first choices of the subjects made during experiment 1. The arrows lead from the stimulus (the correct textile) to the subject's response (the chosen textile). For clarity, only textiles chosen more than three times are shown.

virtual textiles.

6.3 Experiment 2

The second experiment followed the same protocol as the first experiment, as described in section 6.1, but with a new sample set to reflect the fact that the changes made to the dynamic range and spectral balance of the renderer (section 3.4.8) result in the surface textures now generating different output amplitudes to those from when the first sample set was selected. Table 6.5 shows the textures that are in the second sample set and figure 6.12 shows where they are located in the characterisation space. Unlike the previous experiment, there were no anisotropic surface textures featured in this sample set.

Table 6.5: The sample set used in experiment 2, with the texture number, the face used (Right or Back), and the textile’s description.

Number in Set	Texture Identifier	Textile Description
01	37_B	Woven outdoor leisurewear fabric
02	24_R	Satin
03	32_R	Warp knitted tricot-satin
04	27_R	Fleece
05	23_R	Crepe
06	05_R	Crepe
07	11_R	Flannel
08	35_B	Men’s woven overcoat fabric
09	09_R	Wild silk (dupion)
10	46_R	Warp knit for car seats
11	16_R	Lurex knit
12	45_R	Warp knit for car seats
13	30_R	Tulle
14	51_B	Laminated arp knit for car seats
15	10_R	Jute
16	40_R	Warp knitted jersey-based fabric

A new set of eight subjects participated in this experiment (details listed in table 6.6). As with experiment 1, each subject was presented with each test items in turn and asked to rank the four textiles in the virtual environment (figure 6.4) in order of similarity to the target textile, which was a real textile (figure 6.5).

Table 6.6: The recorded details of the eight subjects who participated in experiment 2.

Subject #	Gender	Age	Handedness
1	M	27	R
2	M	27	R
3	F	24	R
4	F	24	R
5	M	33	R
6	M	33	R
7	M	21	R
8	F	26	R

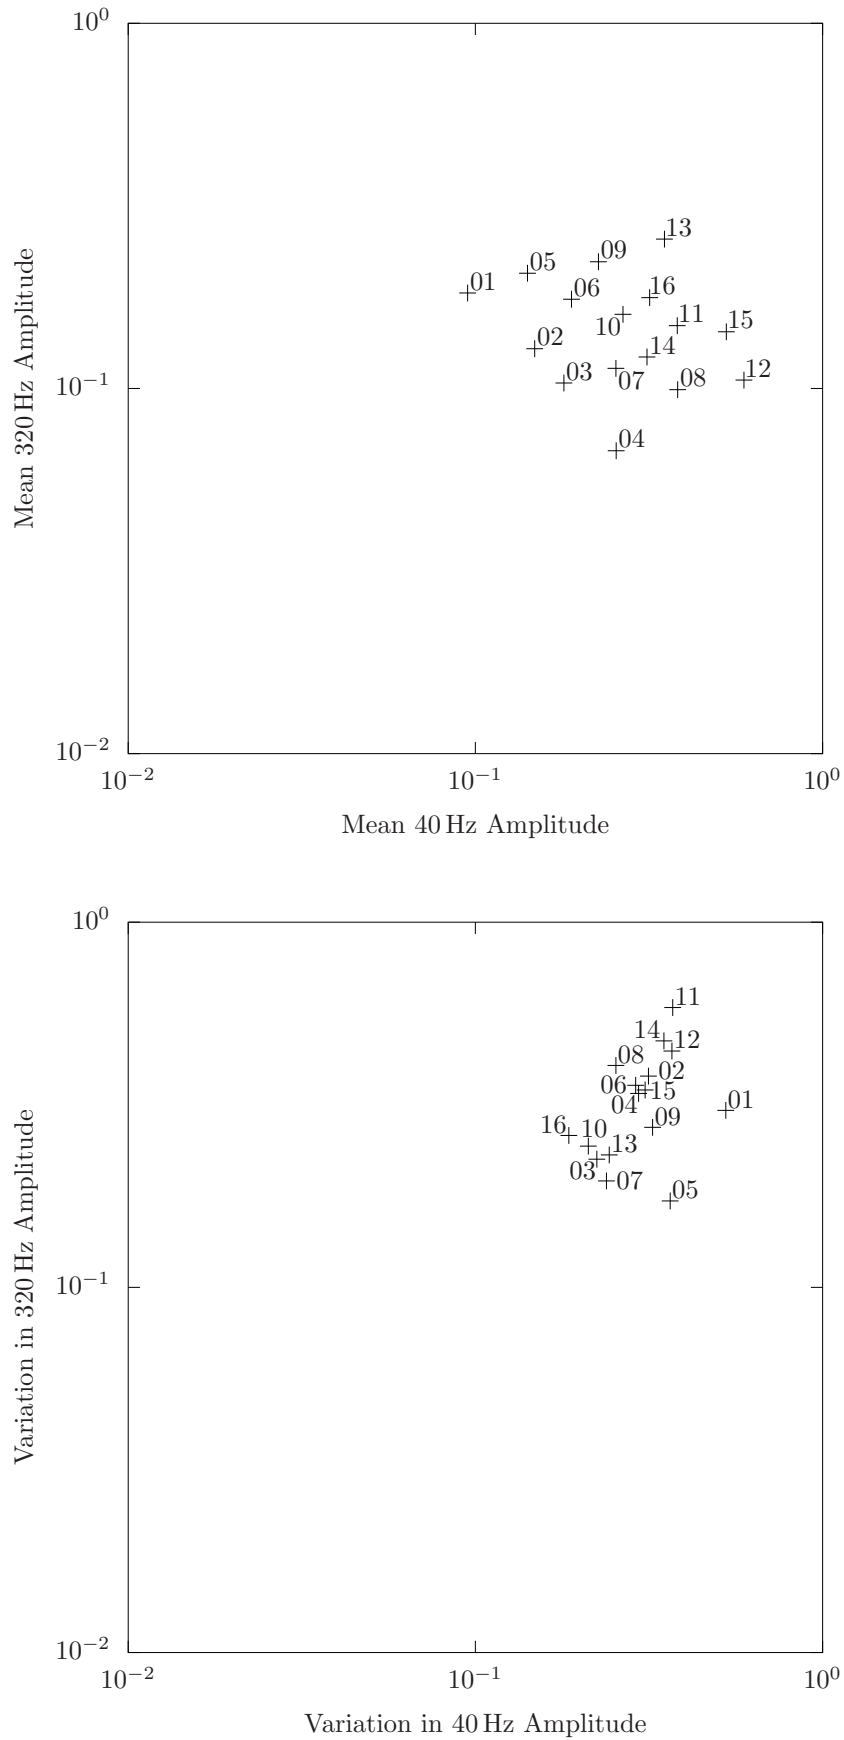


Figure 6.12: Characterisation plots for the tactile renderer for the sample set of 16 textures used in the second experiment, showing the mean (top panel) and CoV (bottom panel) values of A_{40} and A_{320} during exploration of texture surfaces at 10 cm s^{-1} .

Results from the experiment (table 6.7) show results not very different from those of experiment 1. The virtual textiles remain an inaccurate representation of the real textiles, it seems. Feedback from the subjects during the experiment noted that most sets of virtual textures were discriminable but a number were difficult to discriminate between. This was an unexpected consequence of compressing the dynamic range of stimuli - the overall range became more realistic but discrimination on the basis of overall intensity became more difficult. The results of this experiment are broken down by subject (figure 6.13) and by texture (figure 6.14).

Table 6.7: Cumulative scores from experiment 2 (% correct) for matching real and virtual textiles: Key: 1 = real textile matches 1st-ranked virtual textile; 2 = real textile matches 1st- or 2nd-ranked virtual textiles; 3 = real textile matches 1st-, 2nd- or 3rd-ranked virtual textiles; 4 = real textile matches 1st-, 2nd-, 3rd- or 4th-ranked virtual textiles.

Similarity Ranking	Mean Score	Standard Error
1	41.3	1.6
2	65.0	3.0
3	86.1	1.9
4	100.0	0.0




Figure 6.13: A breakdown of the results of experiment 2, showing the percentage of times each of the 8 subjects assigned the correct textile with a particular rank of similarity to the target.

As before, the individual subjects are fairly consistent with each other, all correctly identifying the target in more than 32% of cases and more than 50% of cases with the target being given as the first or second choice. However this experiment shows a lot of variation among the identification of particular textiles. Some textiles, such as number 15, score very highly (77.5%) while others, such as number 12, score quite poorly (22.5%).




Figure 6.14: A breakdown of the results of experiment 2, showing the percentage of times each of the 16 correct virtual textiles was assigned a particular rank of similarity to the target.

If the distributions of the scores of the mode choices is compared to the chance results (figure 6.15), we see a similar result to that observed for the previous experiment in that there are more higher value mode scores and fewer lower value mode scores than we would expect from chance. These results do differ from the previous results in that this time there are mode scores of 2 and 8, while there were none previously. Overall, the results of this experiment give a mean mode value of 4.33 ± 0.14 , which is not very different from the previous result of 4.54 ± 0.13 .

Dissimilarity measures were calculated for each combination of real and virtual textile (table 6.8) based on equation 6.1, where S is the dissimilarity score, i is the ranking position, n_i is the number of subjects that ranked this virtual textile as the i^{th} in similarity, and x_i is the scaling factor for the i^{th} ranking (0 for 1st, 1 for 2nd, 2 for 3rd, and 3 for 4th). To ensure that the real textiles are compared to all virtual textiles, each of the target virtual textiles is used five times during the experiment. Because of this, for the case when the virtual texture is the same as the target real texture, an average dissimilarity score is used across all five occurrences.

$$S = \sum_{i=1}^4 x_i n_i \quad (6.1)$$

The dissimilarity measures used here provide similar information to the subject's first choice matches that were used previously (table 6.4). While the latter is used to provide a simple overview of the subject's responses, the former is used to provide better quan-

Table 6.8: Dissimilarity matrix generated from the results of the second experiment, measuring dissimilarity between each real textile and the 16 virtual textile.

		Virtual Texture															
		01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16
Real Textures	01	06.8	03.0	08.0	07.0	13.0	11.0	11.0	23.0	10.0	12.0	19.0	24.0	19.0	11.0	22.0	13.0
	02	09.0	04.0	13.0	10.0	11.0	11.0	08.0	17.0	18.0	10.0	16.0	18.0	17.0	18.0	20.0	24.0
	03	20.0	07.0	07.6	15.0	12.0	09.0	05.0	12.0	08.0	20.0	18.0	21.0	21.0	08.0	15.0	11.0
	04	09.0	11.0	11.0	07.8	17.0	05.0	10.0	15.0	15.0	13.0	21.0	20.0	15.0	12.0	17.0	10.0
	05	11.0	09.0	09.0	15.0	08.6	17.0	04.0	11.0	13.0	05.0	13.0	23.0	19.0	13.0	24.0	11.0
	06	17.0	15.0	15.0	08.0	19.0	08.8	10.0	12.0	12.0	05.0	08.0	21.0	18.0	04.0	20.0	12.0
	07	16.0	14.0	07.0	11.0	09.0	11.0	08.4	18.0	10.0	15.0	19.0	16.0	13.0	08.0	18.0	13.0
	08	10.0	14.0	15.0	10.0	09.0	10.0	11.0	10.2	13.0	17.0	10.0	18.0	18.0	10.0	14.0	10.0
	09	15.0	13.0	15.0	15.0	11.0	08.0	15.0	04.0	11.2	11.0	07.0	21.0	17.0	08.0	14.0	10.0
	10	09.0	11.0	12.0	07.0	16.0	09.0	10.0	10.0	09.0	09.8	18.0	18.0	14.0	14.0	20.0	14.0
	11	17.0	16.0	17.0	11.0	11.0	12.0	15.0	12.0	12.0	08.0	08.2	18.0	16.0	08.0	13.0	13.0
	12	15.0	15.0	16.0	14.0	17.0	04.0	12.0	09.0	10.0	12.0	07.0	14.8	13.0	05.0	11.0	06.0
	13	20.0	17.0	17.0	18.0	16.0	15.0	08.0	14.0	08.0	07.0	13.0	10.0	08.8	14.0	09.0	10.0
	14	17.0	22.0	19.0	13.0	13.0	14.0	04.0	06.0	14.0	10.0	09.0	14.0	10.0	11.2	12.0	07.0
	15	20.0	17.0	24.0	18.0	21.0	10.0	23.0	10.0	09.0	14.0	09.0	07.0	14.0	16.0	02.4	16.0
	16	20.0	18.0	17.0	15.0	18.0	13.0	04.0	13.0	08.0	12.0	08.0	18.0	13.0	04.0	13.0	09.2

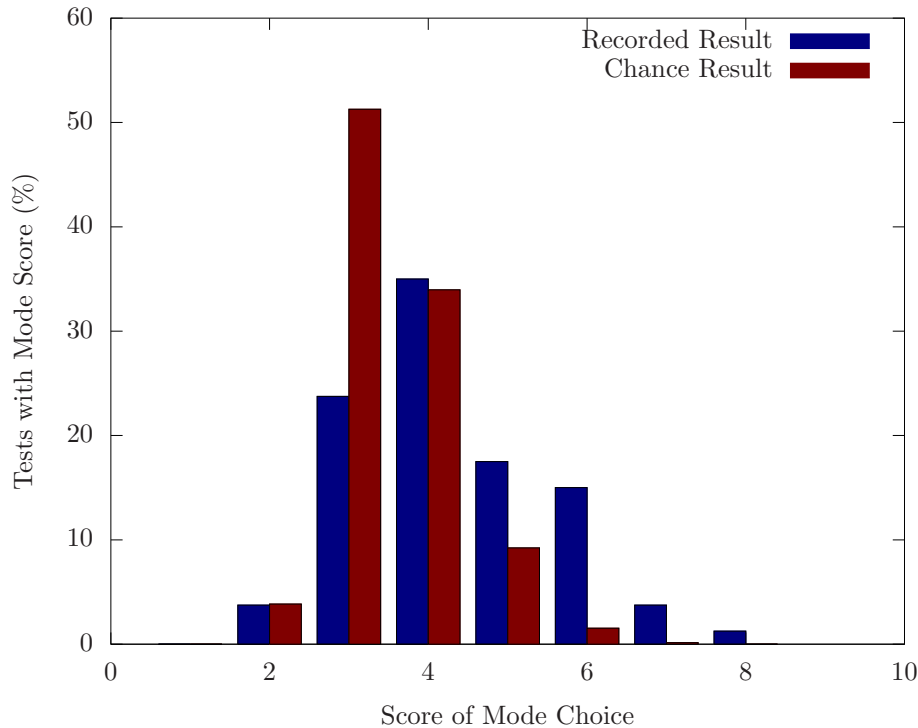


Figure 6.15: From experiment 2, breakdown of the percentages of how the “most popular” (mode) choice scored across all the test items, compared to what would be expected from chance.

titative estimates of the level of dissimilarity between the real and virtual textiles. The correlation coefficients that are calculated between the columns in table 6.8 are indicative of the similarity between the various virtual surface textures, based on their mutual perceived dissimilarity to the each of the real surface textures. Figures 6.16 and 6.17 show these correlation coefficients plotted against the Euclidean distances in the characterisation plots of mean amplitudes and CoV of the amplitudes respectively (figure 6.12). (The distances used in these plots are scaled according to the spectral balance used during the experiment.)

Both plots have been fitted to a linear trend line and suggest that as the separation between the virtual textiles increases in the characterisation space, their similarity (correlation) decreases. This decrease is strongest in the case of the mean space, where R^2 value of the fit is 0.37, suggesting a weak relationship between the distance and the correlation. For the CoV space however, the R^2 of the fit is only 0.06, indicating that there is almost no relationship between the two parameters. These plots of correlation versus distance may be further broken down to show how the correlation varies with the individual parameters, either mean or CoV in the 40 Hz or 320 Hz amplitudes (appendix E).

Overall, it can be concluded that the alterations made to the renderer based on the feedback from experiment 1, as described in section 3.4.8, do not significantly improve the fidelity of the virtual textiles as recreations of the real textiles. However, when the similarity between the virtual surface textures is examined, it can be seen that there is some relationship between the measured similarity between two virtual stimuli and the distance in the mean amplitude characterisation space between them. This suggests that




Figure 6.16: Graph showing the relationship between the correlation coefficient between virtual textures and the Euclidean distance in the mean characterisation space. The fitted line has the equation $y = -2.46x + 0.43$, with $R^2 = 0.37$.




Figure 6.17: Graph showing the relationship between the correlation coefficient between virtual textures and the Euclidean distance in the CoV characterisation space. The fitted line has the equation $y = -1.10x + 0.21$, with an $R^2 = 0.06$.

the mean amplitude characterisation space has some relevance to the perceptual space that the virtual stimuli exist within. The experiment was repeated a third time, with a virtual textile replacing the real textile as the target, the intention of this investigation being to determine the discriminability of the virtual textiles between each other.

6.4 Experiment 3

For the third experiment, the real target textile that the subjects were matching against was replaced with the equivalent virtual textile, which was added to the virtual workspace (figure 6.18). The experimental protocol and sample set were otherwise unchanged from experiment 2 (see section 6.1 for details). The spectral balance of the output signal was also unchanged (0.83:1.0 P:NP). A third set of subjects participated in this experiment (details listed in table 6.6), each being presented with each of the test items in turn. The intention of this version of the experiment was to investigate the discriminability of the virtual textures. Table 6.10 gives the results from this experiment, showing the cumulative rankings of the target textiles within the virtual environment.

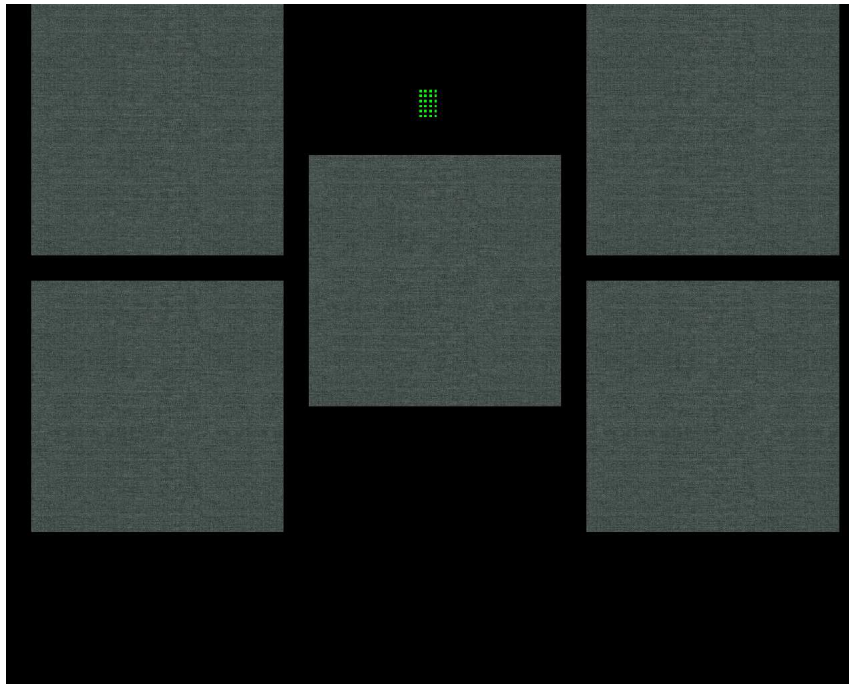


Figure 6.18: An example of one of the virtual workspaces presented during experiment 3, showing four virtual textiles that are to be ranked against the target textile (centre).

In comparing these results to those of the previous experiments (figure 6.19), while the results of the first two do not differ significantly, the results from the third experiment show a much greater accuracy in the identification of matching virtual textures.

The success rate for identification of virtual textures against each other is $64.1 \pm 5.5\%$, providing evidence that the virtual textiles can be discriminated between, but the score is lower than might be expected. This may be due to the virtual textiles approaching the limit of discriminability and so, while some textures can be easily identified as not being similar to the target and eliminated, some are more difficult to discriminate.

Table 6.9: The recorded details of the eight subjects who participated in experiment 3.

Subject #	Gender	Age	Handedness
1	F	24	R
2	M	25	R
3	F	23	R
4	F	22	L
5	M	33	R
6	F	29	R
7	M	23	R
8	M	33	R

Table 6.10: Cumulative scores from experiment 3 (% correct) for matching virtual textiles: Key: 1 = real textile matches 1st-ranked virtual textile; 2 = real textile matches 1st- or 2nd-ranked virtual textiles; 3 = real textile matches 1st-, 2nd- or 3rd-ranked virtual textiles; 4 = real textile matches 1st-, 2nd-, 3rd- or 4th-ranked virtual textiles.

Similarity Ranking	Mean Score	Standard Error
1	64.1	5.5
2	86.7	2.2
3	97.2	0.7
4	100.0	0.0




Figure 6.19: Comparison of the results from the 3 experiments.

Table 6.11 shows the dissimilarity scores for the third experiment, with similarity being between the various virtual surface textures, and the target texture that the subjects are comparing them to. As with table 6.8, multiple occurrences of the same comparison are averaged.

When the similarity between the various virtual textures is measured, using the correlation coefficient between the columns of table 6.11, and this similarity measure is compared to the Euclidean distance in the characterisation spaces between the mean (figure 6.20) and CoV (figure 6.21) values, as was done for experiment two, we again see evidence of a relationship between the similarity and the distance in the mean characterisation space (textures become less similar the further apart they are), but there is no strong evidence of such a relationship in the CoV characterisation space. Based on this, the mean amplitude characterisation space will be used in the development of a perceptual space for the virtual stimuli (section 6.5).




Figure 6.20: Graph showing the relationship between the correlation coefficient between virtual textures from experiment 3 and the Euclidean distance in the mean characterisation space. The fitted line has the equation $y = -2.71x + 0.45$, with $R^2 = 0.43$.

The leading diagonals of tables 6.8 and 6.11 show the dissimilarity between the target textile (a real textile in experiment 2 and a virtual textile in experiment 3) and the corresponding virtual textile. As mentioned above, each of these values are an average of 5 dissimilarity scores from the 5 different virtual environments that exist for each virtual textile. Figure 6.22 shows how the calculated dissimilarity scores compare for each individual virtual environment between experiment 2 and experiment 3, totalling 80 data points (one for each test item in the experiment). For clarity, a line representing equality of score between experiments 2 and 3 is given.

The dissimilarity score in experiment 3 is frequently lower than in experiment 2 (81.25%

Table 6.11: Dissimilarity matrix generated from the results of the third experiment.

		Virtual Texture															
		01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16
Target Texture	01	03.0	10.0	08.0	16.0	08.0	11.0	14.0	21.0	07.0	19.0	21.0	23.0	20.0	12.0	24.0	11.0
	02	12.0	04.6	07.0	11.0	13.0	08.0	05.0	17.0	16.0	09.0	21.0	23.0	18.0	13.0	21.0	23.0
	03	19.0	12.0	03.0	11.0	13.0	08.0	07.0	15.0	10.0	20.0	19.0	19.0	22.0	13.0	21.0	16.0
	04	13.0	14.0	10.0	04.2	17.0	15.0	06.0	10.0	11.0	11.0	23.0	22.0	23.0	20.0	13.0	11.0
	05	04.0	13.0	12.0	19.0	05.8	09.0	13.0	11.0	13.0	06.0	19.0	19.0	21.0	19.0	21.0	12.0
	06	13.0	13.0	13.0	22.0	09.0	06.2	07.0	11.0	06.0	10.0	13.0	21.0	22.0	15.0	21.0	13.0
	07	11.0	13.0	16.0	09.0	20.0	09.0	04.8	07.0	12.0	05.0	13.0	20.0	17.0	24.0	22.0	18.0
	08	14.0	19.0	19.0	08.0	15.0	16.0	10.0	06.0	18.0	14.0	16.0	22.0	10.0	13.0	07.0	09.0
	09	20.0	20.0	13.0	12.0	07.0	11.0	09.0	19.0	03.6	14.0	14.0	24.0	17.0	10.0	22.0	10.0
	10	12.0	19.0	24.0	16.0	14.0	10.0	09.0	16.0	12.0	05.6	11.0	16.0	18.0	06.0	22.0	07.0
	11	21.0	23.0	20.0	22.0	19.0	20.0	17.0	03.0	13.0	11.0	04.4	20.0	11.0	09.0	05.0	04.0
	12	17.0	20.0	23.0	19.0	23.0	11.0	14.0	09.0	14.0	22.0	14.0	01.0	09.0	16.0	08.0	16.0
	13	22.0	12.0	10.0	20.0	20.0	13.0	02.0	06.0	13.0	13.0	08.0	19.0	07.8	17.0	15.0	11.0
	14	23.0	22.0	15.0	15.0	14.0	13.0	17.0	14.0	15.0	07.0	12.0	24.0	10.0	04.2	11.0	07.0
	15	17.0	21.0	24.0	18.0	20.0	10.0	23.0	11.0	10.0	16.0	10.0	05.0	12.0	21.0	01.6	14.0
	16	18.0	19.0	19.0	14.0	18.0	15.0	10.0	08.0	08.0	09.0	13.0	23.0	13.0	09.0	17.0	05.4




Figure 6.21: Graph showing the relationship between the correlation coefficient between virtual textures from experiment 3 and the Euclidean distance in the CoV characterisation space. The fitted line has the equation $y = -1.02x + 0.17$, with $R^2 = 0.05$.




Figure 6.22: A plot comparing the dissimilarity scores for identifying target textiles, experiment 2 against experiment 3, along with a line of equality.

of cases). This corroborates the suggestion that virtual textiles are easier to match to themselves rather than to their real equivalents. This, together with the increase in correct identification in Experiment 3, suggests that the generated virtual textures do not reasonably reflect the surface textures of the real textiles that they are attempting to replicate.

6.5 Perceptual Dimensions

With experiment 3 being a comparison entirely within the virtual context, the dissimilarity scores may be used to create a multi-dimensional space (see section 2.1.3). The ALSCAL algorithm of the SPSS software package was used to create a set of spaces, assuming an ordinal level of measurement, for 1 to 5 dimensions, under the condition that S-stress is minimised in each case. Scree plots (figure 6.23) show the decreasing stress in 3 different measures for an increasing number of dimensions. Typically, the location of a sharp break between the steep, high stress values and more gently sloping, low stress values can be used to identify the appropriate number of dimensions [110].




Figure 6.23: Scree plots comparing the stresses of various multi-dimensional solutions.

Across the calculated solutions, the most noticeable drop in stress is going from a one- to a two-dimensional solution (table 6.12). While it is not the sharp “elbow” suggested by [110], it does greatly reduce the amount of stress from the fitted solution. The validity of the 2D solution is further suggested by the range occupied by the second dimension (figure 6.24). The range is comparable to that of the first dimension suggesting that, if the first dimension is representative of a set of texture properties, the second dimension is also.

On the hypothesis that the multi-dimensional solutions give an accurate representation of perceptual spaces, the question of what properties the various dimensions relate to remains. Initially, the two-dimensional spaces created previously for output characterisation (see chapter 4) were investigated as possible perceptual dimensions, before the search was broadened to include other properties. Because the generation of the space is based solely on the inter-point distances, the space may be rotated freely about the origin and still preserve the inter-point distances. This case is simplest in the two-dimensional case where a point of the solution with co-ordinates (x, y) , after a rotation of angle θ , has the final co-ordinates (x', y') as calculated using equation 6.2. It is possible to assess the appropriateness of a property as a match for the dimension in question by maximising over

Table 6.12: The one- and two-dimensional solutions to the Multi-Dimensional Scaling.

Texture Number	1D	2D	
	Dim. 1	Dim. 1	Dim. 2
01	-1.2431	1.4965	-0.2804
02	-1.2344	1.3560	0.7278
03	-1.1607	1.2946	0.7311
04	-0.9348	0.5745	1.3507
05	-1.0000	1.1263	-0.8426
06	-0.5669	0.8182	-0.5436
07	-0.7210	0.6865	0.8597
08	0.6385	-0.8197	0.8230
09	-0.4916	0.4458	-0.9703
10	0.0055	-0.0231	-1.2192
11	1.2294	-1.4078	-0.6904
12	1.4033	-1.2362	1.3668
13	1.2820	-1.4943	0.3860
14	0.6984	-0.5441	-1.3527
15	1.5018	-1.6420	0.6336
16	0.5937	-0.6312	-0.9795




Figure 6.24: The two-dimensional MDS solution.

different rotation angles the correlation between the property and the proposed dimension, r_i , where i is the dimension being fitted to.

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \quad (6.2)$$

The initial suggestion for the dimensions of the two-dimensional perceptual space were the mean value of the 40 Hz and 320 Hz amplitudes, as established in the two-dimensional characterisation plots in chapter 4. To look for the correspondence between the mean characterisation space and the MDS space, we compare the two as shown schematically in figure 6.25.

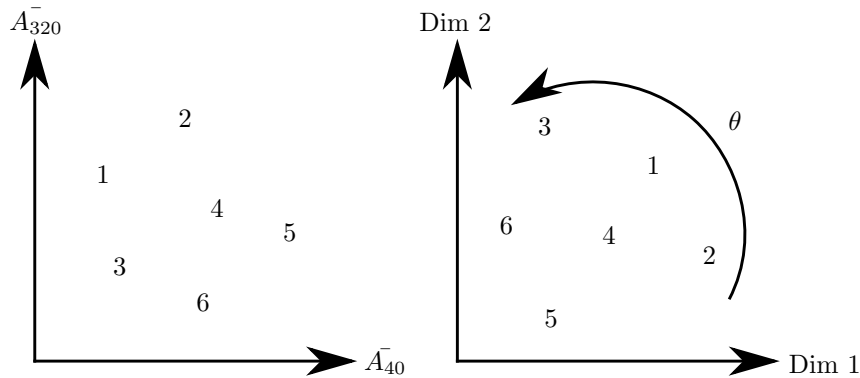


Figure 6.25: Schematic of the mean characterisation space being compared to the calculated MDS space.

The MDS is rotated through various angles, θ , and for each angular position, the mean 40 Hz characterisation values of the textures are correlated with the co-ordinates of the textures on the first dimension, to establish θ_1 , the angle of maximum correlation coefficient, r_1 . Separately, the mean 320 Hz characterisation values of the textures are correlated with the co-ordinates of the textures on the second dimension to establish θ_2 , the angle of maximum correlation coefficient, r_2 . Alternatively, the sum of both correlation coefficients can be maximised with respect to angle of rotation to give the optimum fit of both dimensions into the mean characterisation space, while maintaining the 90° separation between MDS dimensions one and two.

It is useful to calculate the correlation coefficients of the MDS dimensions with the proposed properties independently of each other, as well as together, so as to determine the magnitude of the decrease in correlation coefficient that occurs from maximising the sum of the correlation coefficients instead of maximising the independent coefficients. For example, a chosen dimension might correlate very highly with a characteristic property when fitted by itself, but fitting both dimensions might decrease the correlation coefficient of that dimension (from the difference in the angle of maximum correlation) to a point where an alternate/related property might prove a better fit to the dimension, such as one with a lower correlation coefficient when fitted independently, but whose angle of maximum correlation is closer to that of the angle associated with maximising the sum of the correlations. In addition to the rotation of the solution, the solution may be mirrored about the x- and y- axes arbitrarily, with the inter-point distances still being preserved. It is

therefore appropriate to maximise the sum of the magnitudes of the correlations, $|r_1|$ and $|r_2|$, to find the best fit to the solution.

The first MDS dimension maps optimally onto the mean 40 Hz amplitude at $\theta_1 = 200.2^\circ$ ($r_1 = 0.94$) and the second MDS dimension maps optimally onto the mean 320 Hz amplitude at $\theta_2 = 186.7^\circ$ ($r_2 = 0.56$). When both dimensions are mapped together so that the fitted properties share the same value of θ , a rotation of $\theta = 193.1^\circ$ results in correlation coefficients for the two dimensions of $r_1 = 0.93$ and $r_2 = 0.55$.

The mean 40 Hz amplitude is a good fit for the first dimension but the mean 320 Hz amplitude is not such a good fit to the second dimension. However, the fact that it fits optimally in the direction of the second dimension suggests that it is a good candidate for identification with this dimension. In fact, the first dimension fits either the mean 40 Hz amplitude or, alternatively, the sum of the two mean amplitudes ($r_1 = 0.94$, $\theta_1 = 187.8^\circ$). It can therefore be tentatively identified as a measure of roughness.

While investigating other properties for the second dimension, several possibilities have been considered including different combinations of the mean amplitudes, such as the ratio ($r_2 = 0.82$, $\theta_2 = 318.4^\circ$) and difference ($r_2 = 0.87$, $\theta_2 = 302.3^\circ$), measures of surface anisotropy such as the standard deviation between the mean amplitudes along the four primary directions ($r_2 = 0.74$, $\theta_2 = 306.6^\circ$) or the associated CoV ($r_2 = 0.69$, $\theta_2 = 314.4^\circ$). However, characterisation of the surface anisotropy has been demonstrated to be unreliable, as shown in chapter 4, and so determining that a MDS dimension represents the anisotropy of the surface texture was unlikely.

Another possibility for the second dimension is the characterisation length, which may be determined by a Fourier transform of the amplitudes that are output by the renderer at various positions along the surface of the texture, with a separate length calculated along each direction. From these spectra, the characteristic length along a direction, l , can be calculated from equation 6.3, where a_i is the amplitude of the spectrum at spatial frequency ξ_i .

$$l = \frac{\sum_i a_i}{\sum_i a_i \xi_i} \quad (6.3)$$

The mean feature length and the ratio between mean weft and mean warp feature lengths, as well as some other combinations, have been measured as fits for the second dimension, with correlations of $r_2 = 0.57$ at $\theta_2 = 55.7^\circ$ and $r_2 = 0.44$ at $\theta_2 = 167.6^\circ$ respectively.

Overall, no property has been identified that may truly be considered a good fit to the second dimension of the 2D solution. The best fit is the difference between the mean 40 Hz and 320 Hz values ($r_2 = 0.87$). However, the angle of the fit ($\theta_2 = 302.3^\circ$) suggests that this is not the full picture. If both dimensions are fitted together, the maximum magnitudes of the correlations are at $\theta = 152.6^\circ$ ($r_1 = +0.81$, $r_2 = -0.76$). In fact, the starting point of fitting the two-dimensional MDS solution to the mean characterisation space is almost as good ($r_1 = 0.93$, $r_2 = 0.55$ at 193.1°) and has the merit of simplicity. Figure 6.26 compares the mean characterisation space to the two-dimensional MDS solution, rotated by 193.1° . It is also possible that the second dimension is, in fact, a collection of a number

of different properties used in texture discrimination reduced into a single dimension and thus a higher order MDS solution may be required to fully represent the perceptual space being explored.

6.6 Experienced Users

The subjects of the three experiments were relatively naïve in the use of the device. In order to explore the effect that experience with the device has on the identification of the textiles, two experienced users, both of whom had been involved in the development of the tactile renderer and so had previous experience with using the renderer and knowledge of the underlying rendering strategy, followed the defined experimental protocols to act as a comparison. The results of these experienced user experiments are given in table 6.13, where we see the percentage of correct identification for each of the experienced users as well as the percentage of correct identification obtained from the naïve subjects in each of the three experiments.

Table 6.13: Comparison of the correct identifications (%) between the subjects in the experiments and more experienced users.

	Experiment Subjects	Experienced User 1	Experienced User 2
Experiment 1	38.1±3.1	50.0	40.0
Experiment 2	41.3±1.6	52.5	47.5
Experiment 3	64.1±5.5	86.3	95.0

These results suggest that some familiarity with the device may be helpful in identifying or discriminating between the various textiles that are presented. However, like the naïve subjects, identification between the virtual textiles yielded a greater success than matching real textiles to virtual ones. This comparison of experienced users, as well as a comparison between the results of experiments 2 and 3 for the naïve subjects, were presented in [111].

6.7 Summary

The rendering strategy that is used in the present study may be defined by three distinctive features:

1. The 24 individually specified channels to produce spatio-temporal patterns of stimulation over the fingertip.
2. The two-component stimuli (at 40 Hz and 320 Hz) to provide spectral information.
3. The software model of the surface derived from measurements on the corresponding real textile, i.e., Kawabata surface profiles.

In order to demonstrate the effectiveness of this tactile renderer as a system for generating virtual textured surfaces, a group of subjects were tasked with ranking a set of virtual textiles in similarity to “target” textiles.

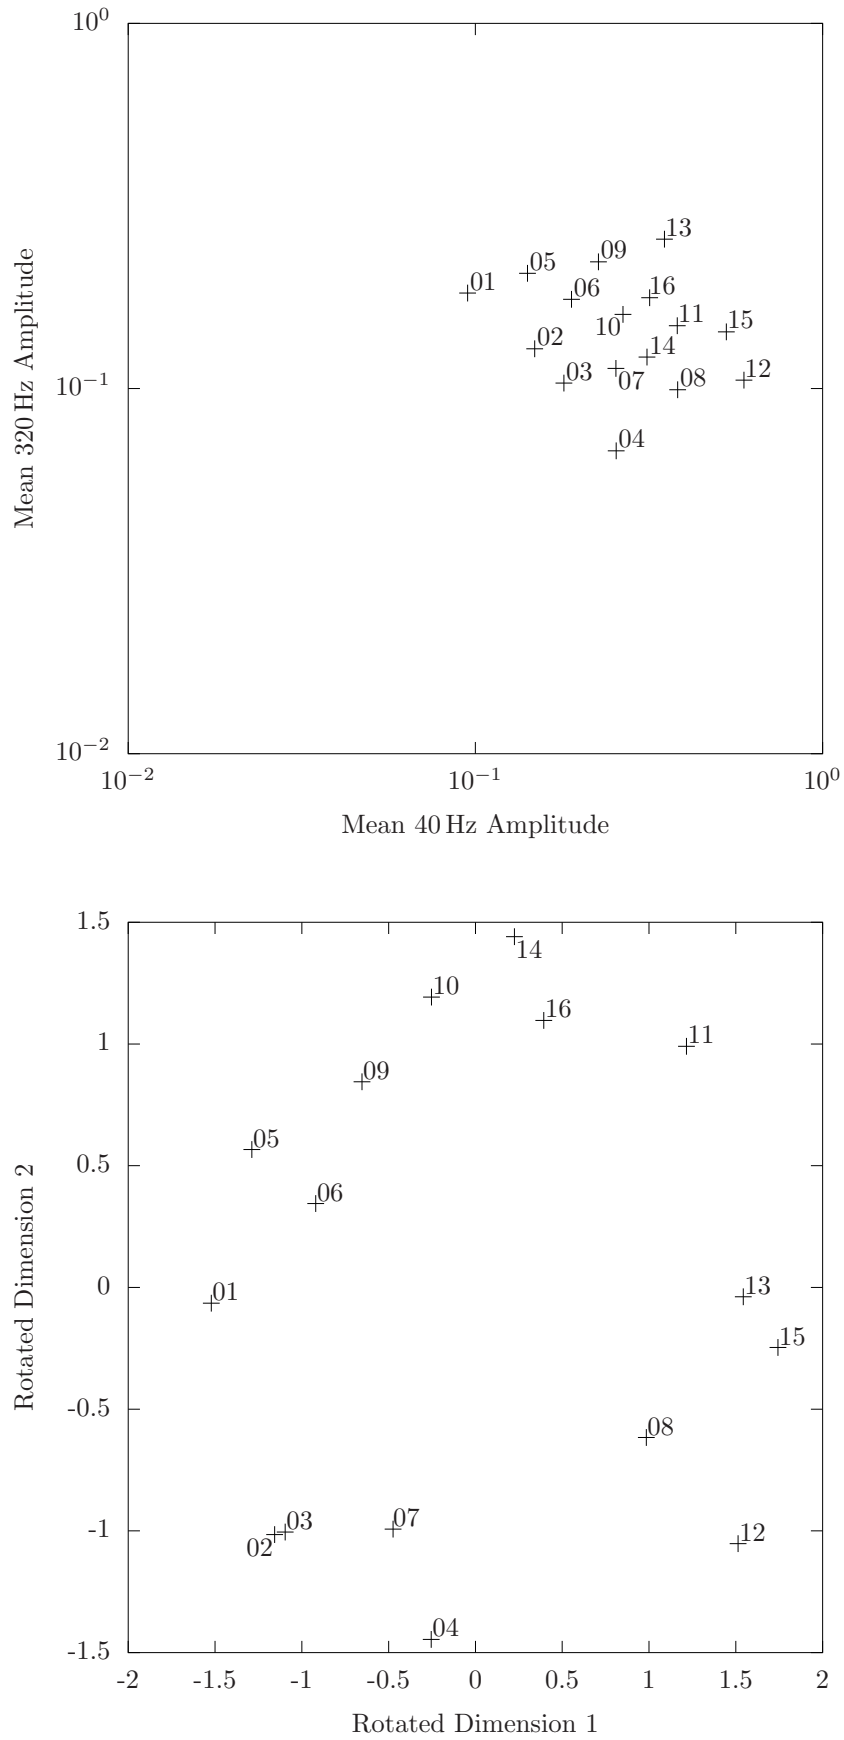


Figure 6.26: Comparison of the mean characterisation space (upper panel) and the two-dimensional MDS solution, rotated through 193.1° (lower panel).

In the initial experiment, the subjects were able to correctly match the virtual textile to the target, which was the textile's real equivalent, in $38.1 \pm 3.1\%$ of cases (table 6.3). This is disappointingly low, given the apparent success of the renderer in producing "textile-like" sensations. Further analysis of the results from this experiment suggests that the "wrong" alternative is sometimes the best match to the target surface texture (figure 6.9), and that the dynamic range of stimuli is too large (figure 6.11).

For experiment 2, the dynamic range of stimuli was corrected with the intention of improving the successful identification rate but the experiment, which repeated the initial experiment with a new sample set, did not show a statistically significant improvement ($41.3 \pm 1.6\%$, table 6.7). However, analysis of the most popular choices of this experiment again suggested showed that the subjects were not simply selecting textiles randomly, but were instead choosing similar textiles as the best match in similarity to the target real textile (figure 6.15).

While experiments 1 and 2 suggests subjects have difficulties in the matching task, when the correlation coefficients that are calculated between the sets of dissimilarity scores associated with the virtual textures from the second experiment are compared to the distances between the virtual textures within the mean amplitude characterisation space, it emerges that the virtual surface textures become less similar to each other at greater distances in the characterisation space (figures 6.20 and 6.21). This, together with the analysis of the most popular choices, suggests that the virtual textures may be discriminable between themselves, but that they are not good representations of the surfaces of the real textiles that they are recreating. This is further corroborated by the third experiment, which follows the same protocol at the previous experiments with the real target textile replaced with a virtual one. Here, the success rate of identification is much higher ($64.1 \pm 5.5\%$, table 6.3).

In combination, the results and subjects' comments from all three experiments suggest that:

- The spatial aspects of the stimulation are adequate, i.e., there would not seem to be an obvious advantage to increasing the density of contactors on the fingertip (decreasing the density seems likely to degrade the perception of corners and edges on the fingertip).
- The spectral aspects of the system are also adequate, although it is difficult to say whether use of additional stimulation frequencies or a continuous spectrum would be an improvement (such strategies would probably require a more elaborate stimulator mechanism).

Hence, of the renderer's three distinctive features listed above, the spatial and spectral features (features 1 and 2) may be considered successful, but the strategy for obtaining software models of the textile surfaces (feature 3) requires more work. Kawabata surface profiles do not appear to provide an adequate basis for the virtual surfaces. While it is possible that Kawabata friction profiles may be useful in this context, it seems more likely that a different strategy altogether is required.

As third experiment compares between virtual textures only, the resulting dissimilarity scores were used to embed the virtual textures into a multi-dimensional space (figure 6.24). The dimensions of the two-dimensional solution were compared with a number of properties in order to identify dimensions of the perceptual space the textures occupy. It was established that the first dimension was some measure of surface roughness; either mean 40 Hz amplitude ($r_1 = 0.94$, $\theta_1 = 200.2^\circ$) or the sum of the mean amplitudes ($r_1 = 0.94$, $\theta_1 = 187.8^\circ$).

The property that is responsible for the second dimension has not been identified, although a number of possibilities have been suggested, with the best fitting being the mean 320 Hz amplitudes ($r_2 = 0.56$, $\theta_2 = 186.7^\circ$) and the difference between the mean 40 Hz and 320 Hz values ($r_2 = 0.87$, $\theta_2 = 302.3^\circ$). It is suggested that, while not the best fitting solution, the comparison between the 2D MDS solution and the mean characterisation space is the best choice at this time, as it fits almost as well as the best fitting solution and has the additional merit of simplicity (figure 6.26). It also remains a possibility that second dimension in fact represents many properties collapsed to a single value and a higher order MDS solution is required to fully represent the perceptual space.

When the identification rates of the naïve experimental subjects were compared to those of some users who were experienced in the use of the tactile renderer, the experienced users show a higher fraction of successful identification of the textiles but still show the same pattern as the naïve users of being able to distinguish between virtual textiles much more easily than to match the virtual textiles to their real equivalents (table 6.13).

Chapter 7

Conclusions

This work seeks to evaluate the effectiveness of a tactile rendering algorithm that was originally developed as part of the EU funded HAPtic sensing of virtual TEXTiles (HAPTEX) project to visually and haptically recreate virtual textiles. The HAPTEX project combined force-feedback devices and tactile displays on the fingertip of the thumb and index fingers of a user, allowing for the exploration and manipulation of the virtual textile. This project uses the tactile component of that system, which allows for the exploration of a variety of surface textures through the use of a tactile display.

7.1 Summary of Thesis

To begin with, Chapter 2 briefly explores the field of psychophysics, which relates the stimulus from the real world to the resulting perceived sensation. This is followed by a description of the human touch system, including the types and properties of touch receptors that exist within the body. The chapter then concludes with a collection of the perceptual spaces that have been established from the exploration of real textures, as well an example of how a perceptual space might be established from measurements of similarity.

The focus of Chapter 3 is the rendering system used in this project, beginning with the display itself. So that the choice of display may be understood, a selection of tactile displays from the literature is provided, illustrating the range of approaches that can be taken in creating a tactile sensation. This is followed by a description of the renderer used in this project, including the specifics on the Exeter tactile display and the driving electronics. With the hardware description complete, the chapter moves on to an explanation of the rendering software used and how the renderer computes the amplitudes that the user experiences through the display. The amplitudes are generated from a series of texture files that the renderer uses to define the surfaces. A detailed description of the creation of these files is also given in the chapter, including how the textures are based upon real surface textures. As this tactile rendering system was developed as part of the HAPTEX project, which intended to visually and haptically recreate textiles virtually, a description of the project is also included, with descriptions of the system development and the evaluation

of the integrated system that was performed toward the end of that project. The chapter ends with a collection of the modifications that were made to the renderer over the course of the project. While the underlying rationale of the renderer remains the same as that which was developed in the HAPTEX project, the specifics of the implementation have been significantly changed.

Very little characterisation of the renderer's output was done towards the end of the HAPTEX project. This omission is corrected in chapter 4, which begins with a discussion of how and by which parameters the virtual textiles may be characterised, and how this characterisation may be verified by comparison against measurements of the real textiles, as recorded by the Kawabata Evaluation System for Fabrics (KES-F). Initial characterisation was done for the renderer version that was used for the initial experiment into the effectiveness of the renderer. This initial characterisation suggests that different textiles result in different sensations being experienced by the user, covering an order of magnitude in both the mean and coefficient of variation (CoV) characterisation spaces. This analysis was also expanded to show the effects of speed of exploration and surface anisotropy on the characterisation values. The renderer that was initially characterised had already undergone a number of alterations since the conclusion of the HAPTEX project and so, for comparison, the HAPTEX version of the tactile renderer was also characterised the same way. The HAPTEX renderer covers similar ranges to the experimental renderer but, fundamentally, it lacks the distinction between different directions of travel across the textile surface which, if the textile is anisotropic, will result in a significant difference between the sensation experienced from the real and virtual textiles. Finally, additional alterations were made to the renderer based on the analysis of the results of the first experiment, as well as feedback from the experimental subjects and other users. This new version of the renderer underwent characterisation so as to examine the effects of the alterations on the output signal. As expected from the dynamic range compression, the range of values in the mean characterisation space was reduced, while the values associated with the CoV remain relatively unaffected.

A public demonstration of the renderer was documented in chapter 5, including the context as to why it was presented within the museum and some of the other installations that were also present for that event. This was followed by a description of the renderer as it was used in the installation, as well as other facsimiles of the artefact that was being recreated by the renderer, which provide context for both evaluating the effectiveness of the recreation and gauging the public's reaction to the renderer installation. Feedback from the use of the display was that most visitors saw it having a positive benefit within the museum environment, but a common complaint relating to the display was that the sensation generated felt "prickly" to the user. To make the sensation more pleasant, caps were added to the moving pins which served to reduce the unevenness of the contactor surface. It can also be said from the feedback that the visitors generally preferred the more direct, physical recreations. However, there are constraints/situations where this virtual system would be preferable over more direct recreations (such as in transporting or flexibility).

Chapter 6 details the experiments performed in this project to evaluate the effectiveness

of the renderer at recreating surface textures. In the initial experiment, subjects were presented with a series of virtual environments containing four virtual textiles. The subjects were tasked with ranking the textiles in order of similarity to a provided real textile. During this experiment, the subjects correctly identified the textile $38.1 \pm 3.1\%$ of the time. After modifications were made to the renderer based on feedback during the initial experiment, the experiment was repeated. This yielded a successful identification rate of $41.3 \pm 1.6\%$, which is not significantly different from the first result. Further analysis suggested that the subjects were able to distinguish between the textures to some degree but that the virtual textures may not be good recreations of their real equivalents. A further experiment, which saw the real target textile replaced with a virtual one yielded a successful identification rate of $64.1 \pm 5.5\%$. Similarity scores were calculated based on the results of this experiment which were used to embed the virtual textures in a series of multi-dimensional spaces. The two-dimensional space was found to be the most practical solution, with one of the dimensions being identified as the roughness of the texture. The source of the second dimension was unable to be identified with confidence.

7.2 Overview of Project Conclusions

The combination of hardware and software that was used in this project forms the basis of a versatile and practical system for the delivery of virtual textures to the fingertip via multipoint stimulation. The device is easy to use and is readily accepted by new users, as is evidenced by the success in the public demonstrations of the renderer (chapter 5). The experience of using the renderer as a physical device is not dissimilar to that from a typical computer mouse, a device a large number of people are very familiar with, and so the experience of using the tactile display in this situation feels “natural” to the users.

The characterisation work (chapter 4) shows that, with the appropriate texture data file, the renderer is capable of producing a wide range of “textile-like” sensations, which are discriminable between each other, as the third experiment suggests (section 6.4, 64%). However, the match between virtual textiles and their real equivalents is not good in the general case ($\sim 40\%$ on average across experiments 1 and 2, sections 6.2 and 6.3), although some virtual textiles are close to their real equivalents ($> 50\%$). Analysis of the “most popular” choices show that subjects tend to make similar choices in virtual environments, suggesting that some virtual textiles are close matches to real textiles, but that they are the “wrong” real textiles.

The fact that one of the experienced users achieved a 50% identification rate in the one-from-four task is indicative of the performance of the system at this time. It suggests that two of the four possible virtual textiles could be discounted from consideration but the choice between the remaining two was simply guesswork. While this is acceptable for a number of applications, especially if combined with other sensory cues (visual, auditory, etc.), it suggests there is still room for improvement.

The fact that the virtual textures are discriminable amongst each other but have variable success in representing the real textiles raises questions over the effectiveness of the KES-F

setup as a means of recording the surface profile. It is possible that the probe that is used deforms the surface, either through compression normal to the surface or through stretching the textile along the direction of scanning. Therefore, a probe which is a more direct analogue to the finger may be a better choice for extracting the surface profile. In such an “artificial finger”, mechanical detectors would register signals corresponding to those detected by the touch receptors in the skin. Work has been done to investigate this possibility, such as in [112] and [113], but more work is required to obtain the extensive data sets for necessary for creating a wide range of virtual textiles.

Investigations into the perceptual dimensions associated with the virtual textures suggest that at least one of the dimensions is a measure of textile roughness, which is a common perceptual dimension within the literature for the haptic exploration of real textiles. The configuration of the perceptual space proposed by this work is a two dimensional space, the axes of which represent the characteristic mean values in the 40 Hz and 320 Hz channels. While this solution has the merit of simplicity and the mean 40 Hz values fit the established perceptual space quite well ($r_1 = 0.93$), the mean 320 Hz values do not fit as well ($r_2 = 0.55$). It is possible that a higher order solution is required, and the presented second dimension may be a reduction of several perceptual dimensions on to a single one.

7.3 Discussion and Future Work

While it can be said that the renderer that currently exists is capable of generating a wide variety of “textile-like” sensations that are able to recreate the surfaces of a range of textiles, albeit with more success on some textiles than others, a number of avenues of investigation still exist that might lead to improvements in the tactile rendering.

A possible initial avenue of further investigation would be to continue to investigate the source of the second perceptual dimension that is discussed in section 6.5, or identifying whether this dimension is in fact a collection of several other properties collapsed into one. However, the amount of confusion between the various virtual textures limits the accuracy of the perceptual space generated from the experimental results. Future investigations may instead be better focused on developing and improving the recreation of specific target textures (see section 7.2) since doing so may also help to increase the discriminability of the virtual textures and allow for the creation of a more accurate perceptual space, as well as provide additional clues to the criteria significant in texture discrimination.

The primary goal of any subsequent research should be to improve the source for the texture surface data. It is suspected that the difficulties in both the discrimination between virtual textiles, as well as the matching of virtual textiles to the equivalent real textiles, is related to the loss of high frequency surface features during the recording by the KES-F. Possibly an alternate sensor, either as part of the KES-F surface tester or for some alternate testing method, and one that is a closer analogue to the human finger, would record a surface with more high frequency information, and thus a surface closer to that perceived by a human exploring the surface. Thus, developing an understanding the mechanics of tactile perception in the fingertip is an important consideration when designing a new

probe, such as the role of fingerprints in tactile perception (see section 2.2.3). Subsequent repeats of the experiments using texture files based on this alternate data source would be expected to yield higher discrimination results than those detailed here if the alternate source is an improvement on the exist surface data.

It is also possible that other features of the textiles will need to be introduced to make the surface textures an accurate recreation of the real textiles. In addition to the roughness measures that are output from the KES-F surface tester, there is also a measure of the friction of the surface that is part of the standard KES-F system that is also recorded by the KES-F surface tester, in a similar way to the roughness in the HAPTEx project. This testing produces measures of friction as a function of position across the textiles surface equivalent to the roughness measurements. These measurements can potentially be combined in a variety of ways (such as root of the sum of the squares (root-sum-square)) with the roughness profiles that are currently being used for the output signal. Whether or not this leads to an improvement in recreation of the surface texture will need to be the subject of further investigation. It should also be noted that, if the roughness contactor is altered to be a closer analogue to the human finger, the friction measurements may also need to be retaken using a similar sensor in order to maintain the equivalence of the two measurements.

Alternatively, as was explored in section 3.1.11, tactile displays that recreate friction sensations often use of a shear block to impart lateral deformation to the fingertip. The Exeter tactile display could be modified to include such a shear block in its design and impart lateral deformation the to users fingertip. Work has, however, been done previously on a similar pin display-shear block combination (see [114]) and the combination did not make a significant difference in the subjects ability to detect and localise embedded objects over the pin display alone.

In section 3.4.2, the conversion of texture files from a linear to a logarithmic scale of spatial frequencies was discussed and it was noted that the change did not significantly increase storage requirements nor computational time. While storage requirements are not, in of themselves, a significant limitation, the amount of processing time required by the renderer is dependent on the number of stored spectral lines (more spectral line requiring longer to process) and processing time is limited by the 40 Hz output refresh rate of the renderer. As most of the surface detail exists at lower frequencies, the conversion to a logarithmic scale preserves the resolution at lower spatial frequencies and increases the range of spatial frequencies that are stored, thus allowing for the inclusion of any high frequency detail, while maintaining the number of spectral lines and thus the processing time. Future researchers should be aware of this limitation of processing time and plan any new texture files, or rendering strategies (which also effect the requited processing time), accordingly. It may also be useful for there to be a future investigation with the intention of formally assessing how much of the spatial spectrum is required for the renderer to produce an effective recreation of the textile surface.

When a new source of the surface texture data is established, whether through the addition of new KES-F measurements, the replacement of the KES-F system completely, or by some other means, it would be useful to re-assess the scaling factors used for the generation of the

output signal, both in the case of the divisors that normalise the signal amplitudes and the case of the spectral balance of the output, as the new source data will result in amplitudes that may, potentially, be very different in value to those that are currently being used. For example, if the new amplitudes are significantly increased, the normalisation factor may remove the variation in high intensity textures. Therefore, for optimum textile recreation, it is important to consider these values if the source data is changed.

One question that has not directly been addressed in this project relates to the discriminability of real surface textures. Repeating the experimental protocol using exclusively real textures should demonstrate the limit of discriminability between real textures and so set a theoretical maximum for the discriminability of the virtual textures against which the results of this project may be compared. It has been assumed in the analysis that real textiles are discriminable with an accuracy of $\sim 100\%$ but this may not be the case. These results of discriminability, like those of Experiment 3, can also be used to create multi-dimensional spaces for the real textures. These can then be compared to the spaces generated for the virtual textures and may help to guide further development of the system.

Future investigators should be aware that the theoretical maximums for the discrimination and identification of the virtual textiles are unlikely to be reached, irrespective of the improvements to the system that are implemented. With the current hardware, certain aspects of some textiles are particularly difficult to replicate, such as very smooth or compressible textiles. In the former case, the tactile sensation is dominated not by the virtual surface texture but by the nature of surface of the tactile display itself (discrete contactors, rather than a continuous surface). In the latter case, our system is limited to a two-dimensional surface texture while mechanical properties, such as compressibility, are better suited to being expressed by a force-feedback device, such as that implemented during the HAPTEX project. While these features are important to achieve truly accurate virtual recreations of textiles, it can be suggested that the amount of development/research time associated with addressing these shortcomings would not be reflected in the amount of improvement in discriminability.

7.4 Potential Applications

With its relatively compact design and Universal Serial Bus (USB) interface, it is possible to run the tactile renderer on a wide range of devices and in a variety of configurations. This potentially allows the system to be installed en masse, either within display environments, such as was demonstrated in the museum, or in private environments, such as the home. In the latter case, online shopping is becoming increasingly common for a wide range of purchases. In some cases however, such as when shopping for clothing or upholstery, where the tactile sensation of the product is a relevant criteria in the purchasing decision, the ability of the customer to make an informed choice about their purchase is limited. This system could correct that by the company providing the customer with a sample texture file, which represents the surface texture of the product, and thus allowing them to decide on their purchase without the previously described limitation.

One area with this display would not necessarily be appropriate however would be within the area of tablet computers. Within this area, the addition of more hardware would diminish the compact and singular nature of the devices, which is often one of the main selling points for the device. Tactile feedback could be integrated into such devices, but the feedback would have to be with the visual display itself, and still allow the user to freely explore the surface of the display. Such devices could follow the same basic principles of the work described here, such as reducing the tactile stimulus into two component sine waves of appropriate amplitudes, but would require significantly different hardware to implement.

As the HAPTEX project demonstrates, this tactile renderer can also be expanded into a three dimensional workspace through the addition of a force feedback device. This could expand upon the previously described application so as to include the mechanical properties of the product that is being considered for purchase (such as having a virtual outfit on a virtual tailor's dummy). While this system was created with the intention of virtually recreating textiles, this system could also be used more generally for simulating deformable objects with realistic surface texture, such as the clothes, hair and skin of a virtual person, which has applications within areas such as video gaming or training simulations.

References

- [1] Dennis Allerkamp. *Generation of Stimuli Supporting Tactile Perception of Textiles in a VR System*. PhD thesis, Gottfried Wilhelm Leibniz Universität Hannover, June 2009.
- [2] Gustav Theodor Fechner. *Elements of Psychophysics*, volume 1. Holt, Rinehart and Winston, Inc, 1966.
- [3] Ernst Heinrich Weber, Helen Elizabeth Ross, and David J. Murray. *E.H. Weber on the Tactile Senses*. Erlbaum (UK) Taylor & Francis, 2nd edition, 1996.
- [4] Jian Bi. *Sensory Discrimination Tests and Measurements: Statistical Principles, Procedures and Tables*. Wiley-Blackwell, 2006.
- [5] J. B. Kruskal. Multidimensional Scaling by Optimizing Goodness of Fit to a Non-metric Hypothesis. *Psychometrika*, 29(1):1–27, 1964.
- [6] Yoshio Takane, Forrest W. Young, and Jan De Leeuw. Nonmetric Individual Differences Multidimensional Scaling: an Alternating Least Squares Method with Optimal Scaling Features. *Psychometrika*, 42(1):7–67, 1977.
- [7] Joseph F. Hair, Rolph E. Anderson, Ronald L. Tatham, and William C. Black. *Multivariate Data Analysis*. Prentice Hall, fifth edition, 1998.
- [8] Ingwer Borg and Patrick J. F. Groenen. *Modern Multidimensional Scaling: Theory and Applications*. Springer, 2nd edition, 2005.
- [9] Gosta Ekman. Dimensions of Color Vision. *The Journal of Psychology*, 38(2):467–474, October 1954.
- [10] G. Konstantakopoulos, K. Tchanturia, S. A. Surguladze, and A. S. David. Insight in eating disorders: clinical and cognitive correlates. *Psychological medicine*, 41(9): 1951–1961, September 2011.
- [11] Susan J. Lederman and Roberta L. Klatzky. Hand movements: a window into haptic object recognition. *Cognitive Psychology*, 19(3):342–368, July 1987.
- [12] Marc O. Ernst and Martin S. Banks. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870):429–433, 2002.
- [13] J. M. Loomis and Susan J. Lederman. Tactual Perception. In *Handbook of Perception and Human Performance*, volume 2, chapter 31, pages 1–41. Wiley, New York, 1986.

-
- [14] Neil A. Campbell and Jane B. Reece. *Biology*. Pearson, Benjamin Cummings, seventh edition, 2005.
- [15] Åke B. Vallbo and Roland S. Johansson. Properties of cutaneous mechanoreceptors in the human hand related to touch sensation. *Human neurobiology*, 3(1):3–14, January 1984.
- [16] Zdenek Halata and Klaus I. Baumann. Anatomy of Receptors. In *Human Haptic Perception: Basis and Applications*, chapter 6, pages 85–92. Birkhäuser, 2008.
- [17] Dale Purves, David Fitzpatrick, Lawrence C. Katz, Anthony-Samuel LaMantia, James O. McNamara, S. Mark Williams, and George J. Augustine, editors. *Neuroscience*. Sinauer Associates, Inc., 2nd edition, 2001.
- [18] Roland S. Johansson and Åke B. Vallbo. Tactile Sensory Coding in the Glabrous Skin of the Human Hand. *Trends in Neurosciences*, 6:27–32, 1983.
- [19] Natalie L. Cornes, R. Sulley, Alan Christopher Brady, and Ian R. Summers. Measurements on Pacinian corpuscles in the fingertip. In *Proceedings of Enactive '06*, pages 117–118, Montpellier, France, 2006.
- [20] Lana Pitts-Yushchenko, Joanne Dale, Peter Winlove, and Ian R. Summers. Dynamic response of pacinian corpuscles. In *EuroMov Inaugural Conference*, pages 93–94, Montpellier, 2013.
- [21] Kenneth O. Johnson, Takashi Yoshioka, and Franciso Vega-Bermudez. Tactile Functions of Mechanoreceptive Afferents Innervating the Hand. *Journal of Clinical Neurophysiology*, 17(6):539–558, November 2000.
- [22] Mark Hollins, Sliman J. Bensmaïa, and S Washburn. Vibrotactile adaptation impairs discrimination of fine, but not coarse, textures. *Somatosensory & Motor Research*, 18(4):253–262, January 2001.
- [23] Sliman J. Bensmaïa, Mark Hollins, and Jeffrey Yau. Vibrotactile intensity and frequency information in the pacinian system: a psychophysical model. *Perception And Psychophysics*, 67(5):828–841, 2005.
- [24] Sliman J. Bensmaïa and Mark Hollins. Pacinian representations of fine surface texture. *Perception And Psychophysics*, 67(5):842–854, 2005.
- [25] Sliman J. Bensmaïa and Mark Hollins. The vibrations of texture. *Somatosensory & Motor Research*, 20(1):33–43, January 2003.
- [26] J. Scheibert, S. Leurent, A. Prevost, and G. Debrégeas. The role of fingerprints in the coding of tactile information probed with a biomimetic sensor. *Science (New York, N.Y.)*, 323(5920):1503–1506, March 2009.
- [27] Ronald T. Verrillo. Effect of Contactor Area on the Vibrotactile Threshold. *The Journal of the Acoustical Society of America*, 35(12):1962–1966, 1963.
-

-
- [28] Ronald T. Verrillo, Anthony J. Fraioli, and Robert L. Smith. Sensation Magnitude of Vibrotactile Stimuli. *Perception & Psychophysics*, 6(6):366–372, November 1969.
- [29] Stanley J. Bolanowski, George A. Gescheider, Ronald T. Verrillo, and C. M. Checkosky. Four channels mediate the mechanical aspects of touch. *The Journal of the Acoustical Society of America*, 84(5):1680–1694, November 1988.
- [30] George A. Gescheider, Stanley J. Bolanowski, and Ronald T. Verrillo. Some characteristics of tactile channels. *Behavioural Brain Research*, 148(1-2):35–40, January 2004.
- [31] Ki-Uk Kyung, Minseung Ahn, Dong-Soo Kwon, and Mandayam A Srinivasan. Perceptual and biomechanical frequency response of human skin: implication for design of tactile displays. In *World Haptics Conference*, pages 96–101. IEEE, 2005.
- [32] Jeremy M. Wolfe, Keith R. Kluender, Dennis M. Levi, Linda M. Bartoshuk, Rachel S. Herz, Roberta L. Klatzky, Susan J. Lederman, and Daniel M. Merfeld. *Sensation & Perception*. Sinauer Associates, Sunderland, MA, 2nd edition, 2008.
- [33] J. Löfvenberg and Roland S. Johansson. Regional Differences and Interindividual Variability in Sensitivity to Vibration in the Glabrous Skin of the Human Hand. *Brain Research*, 301(1):65–72, 1984.
- [34] Roland S. Johansson, U. Landström, and R. Lundström. Responses of Mechanoreceptive Afferent Units in the Glabrous Skin of the Human Hand to Sinusoidal Skin Displacements. *Brain Research*, 244(1):17–25, July 1982.
- [35] David T. Blake, Steven S. Hsiao, and Kenneth O. Johnson. Neural Coding Mechanisms in Tactile Pattern Recognition: The Relative Contributions of Slowly and Rapidly Adapting Mechanoreceptors to Perceived Roughness. *Journal of Neuroscience*, 17(19):7480–7489, 1997.
- [36] Kenneth O. Johnson and Graham D. Lamb. Neural mechanisms of spatial tactile discrimination: Neural patterns evoked by braille-like dot patterns in the monkey. *Journal of Physiology*, 310:117–144, 1981.
- [37] Kenneth O. Johnson and J. R. Phillips. Tactile spatial resolution. I. Two-point discrimination, gap detection, grating resolution, and letter recognition. *Journal of Neurophysiology*, 46(6):1177–1192, December 1981.
- [38] G Westling and Roland S. Johansson. Responses in glabrous skin mechanoreceptors during precision grip in humans. *Experimental Brain Research*, 66(1):128–140, January 1987.
- [39] M. Knibestöl and Åke B. Vallbo. Single Unit Analysis of Mechanoreceptor Activity from the Human Glabrous Skin. *Acta Physiologica Scandinavica*, 80(2):178–195, October 1970.
- [40] H. Olausson, J. Wessberg, and N. Kakuda. Tactile directional sensibility: peripheral neural mechanisms in man. *Brain Research*, 866(1-2):178–187, June 2000.
-

-
- [41] Benoni B. Edin and Niclas Johansson. Skin strain patterns provide kinaesthetic information to the human central nervous system. *Journal of Physiology*, 487(1): 243–251, 1995.
- [42] Delphine Picard, Catherine Dacremont, Dominique Valentin, and Agnès Giboreau. Perceptual dimensions of tactile textures. *Acta Psychologica*, 114(2):165–184, October 2003.
- [43] Steve Guest, Jean Marc Dessirier, Anahit Mehrabyan, Francis McGlone, Greg Es-sick, George A. Gescheider, Anne Fontana, Rui Xiong, Rochelle Ackerley, and Kevin Blot. The development and validation of sensory and emotional scales of touch perception. *Attention, Perception & Psychophysics*, 73:531–550, February 2011.
- [44] Wouter M. Bergmann Tiest and Astrid M. L. Kappers. Analysis of haptic perception of materials by multidimensional scaling and physical measurements of roughness and compressibility. *Acta Psychologica*, 121(1):1–20, January 2006.
- [45] Ivanne Soufflet, Maurice Calonnier, and Catherine Dacremont. A comparison between industrial experts’ and novices’ haptic perceptual organization: a tool to identify descriptors of the handle of fabrics. *Food Quality and Preference*, 15(7-8): 689–699, 2004.
- [46] Mark Hollins, Sliman J. Bensmaïa, Kristie Karlof, and Forrest Young. Individual differences in perceptual space for tactile textures: evidence from multidimensional scaling. *Perception & Psychophysics*, 62(8):1534–1544, November 2000.
- [47] Mark Hollins, Richard Faldowski, Suman Rao, and Forrest Young. Perceptual dimensions of tactile surface texture: a multidimensional scaling analysis. *Perception & Psychophysics*, 54(6):697–705, December 1993.
- [48] Takafumi Terada, Masato Ogata, Takaaki Kikukawa, Shin Hongo, Manabu Nagasaka, Kentaro Takanami, Kagenori Kajihara, and Masaru Fujino. Virtual Human Body using Haptic Devices for Endoscopic Surgery Training Simulator. In *Mechatronics, 2007 IEEE International Conference on*, pages 1–5, New York, NY, May 2007. IEEE.
- [49] Nikolaos Kaklanis, Dimitrios Tzovaras, and Konstantinos Moustakas. Haptic Navigation in the World Wide Web. In *Universal Access in Human-Computer Interaction: Applications and Services, Pt III*, volume 5616 of *Lecture Notes in Computer Science*, pages 707–715, Heidelberger Platz 3, D-14197 Berlin, Germany, 2009. Springer-Verlag, Berlin.
- [50] M. Jungmann and Helmut F. Schlaak. Miniaturised Electrostatic Tactile Display with High Structural Compliance. In *Proceedings of Eurohaptics 2002*, pages 12–17, 2002.
- [51] M. Matysek, P. Lotz, and Helmut F. Schlaak. Braille Display with Dielectric Polymer Actuator. In *Proceedings of Actuator 2006*, pages 14–16, Bremen, Germany, June 2006.

-
- [52] Helmut F. Schlaak, M. Jungmann, M. Matysek, and P. Lotz. Novel Multilayer Electrostatic Solid State Actuators with Elastic Dielectric. In *Proceedings of the SPIE*, volume 5759, pages 121–133. SPIE, May 2005.
- [53] Yasushi Ikei, Kazufumi Wakamatsu, and Shuichi Fukuda. Vibratory Tactile Display of Image-Based Textures. *IEEE Computer Graphics and Applications*, 17(6):53–61, 1997.
- [54] Yasushi Ikei and Masashi Shiratori. TextureExplorer: A Tactile and Force Display for Virtual Textures. In *Proceedings of the 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 327–334. IEEE Comput. Soc, 2002.
- [55] Thorsten Maucher, Karlheinz Meier, and Johannes Schemmel. An interactive tactile graphics display. In *Proceedings of the Sixth International Symposium on Signal Processing and its Applications*, volume 1, pages 190–193, Kuala Lumpur, Malaysia, 2001. IEEE.
- [56] Ian R. Summers and Craig M. Chanter. A broadband tactile array on the fingertip. *The Journal of the Acoustical Society of America*, 112(5):2118–2126, November 2002.
- [57] Jérôme Pasquero and Vincent Hayward. STReSS: A Practical Tactile Display System with One Millimeter Spatial Resolution and 700 Hz Refresh Rate. In *Proceedings of Eurohaptics 2003*, pages 94–110, 2003.
- [58] Qi Wang and Vincent Hayward. Biomechanically Optimized Distributed Tactile Transducer Based on Lateral Skin Deformation. *The International Journal of Robotics Research*, 29(4):323–335, August 2009.
- [59] Vincent Hayward and Juan Manuel Cruz-Hernandez. Tactile display device using distributed lateral skin stretch. In *Proceedings of the haptic interfaces for virtual environment and teleoperator systems symposium*, pages 1309–1314, 2000.
- [60] Harry Block, J. P. Kelly, A. Qin, and T. Watson. Materials and mechanisms in electrorheology. *Langmuir*, 6(1):6–14, January 1990.
- [61] W. I. Kordonsky. Elements and Devices Based on Magnetorheological Effect. *Journal of Intelligent Material Systems and Structures*, 4(1):65–69, January 1993.
- [62] Enzo Pasquale Scilingo, Nicola Sgambelluri, Danilo De Rossi, and Antonio Bicchi. Haptic Displays Based on Magnetorheological Fluids: Design, Realization and Psychophysical Validation. In *Proceedings of the 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2003*, pages 10–15. IEEE Comput. Soc, 2003.
- [63] P. M. Taylor, A. Hosseini-Sianaki, and C. J. Varley. An electrorheological fluid-based tactile array for virtual environments. In *Proceedings of 1996 IEEE International Conference on Robotics and Automation*, volume 1, pages 18–23, Minneapolis, Minnesota, 1996. IEEE.
-

-
- [64] Masashi Konyo, Satoshi Tadokoro, and Toshii Takamori. Artificial tactile feel display using soft gel actuators. In *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, volume 4, pages 3416–3421, San Francisco, 2000. IEEE.
- [65] Y. Bar-Cohen, S. Leary, A. Yavrouian, K. Oguro, S. Tadokoro, J. Harrison, J. Smith, and J. Su. Challenges to the application of IPMC as actuators of planetary mechanisms. In *Proceedings of SPIE'S 7th Annual International Symposium on Smart Structures and Materials*, pages 1–5, Newport, CA, 2000.
- [66] Takayuki Iwamoto, Mari Tatezono, and Hiroyuki Shinoda. Non-Contact Method for Producing Tactile Sensation Using Airborne Ultrasound. In *Proceedings of Eurohaptics 2008*, pages 504–513, 2008.
- [67] M. Takasaki, T. Nara, S. Tachi, and T. Higuchi. A tactile display using surface acoustic wave. In *Proceedings 9th IEEE International Workshop on Robot and Human Interactive Communication*, pages 364–367. IEEE, 2000.
- [68] T. Nara, M. Takasaki, T. Maeda, T. Higuchi, S. Ando, and S. Tachi. Surface acoustic wave tactile display. *IEEE Computer Graphics and Applications*, 21(6):56–63, 2001.
- [69] Dimitrios A. Kontarinis, Jae S. Son, William J. Peine, and Robert D. Howe. A tactile shape sensing and display system for teleoperated manipulation. In *Proceedings of 1995 IEEE International Conference on Robotics and Automation*, volume 1, pages 641–646. IEEE, 1995.
- [70] P. M. Taylor, A. Moser, and A. Creed. The design and control of a tactile display based on shape memory alloys. In *Proceedings of the 1997 International Conference on Robotics and Automation*, volume 2, pages 1318–1323. IEEE, 1997.
- [71] Parris S. Wellman, William J. Peine, Gregg E. Favalora, and Robert D. Howe. Mechanical Design and Control of a High-Bandwidth Shape Memory Alloy Tactile Display. In *International Symposium on Experimental Robotics*, pages 56–66, Barcelona, Spain, 1997. Springer-Verlag.
- [72] Kurt A. Kaczmarek, Mitchell E. Tyler, and Paul Bach-y Rita. Pattern identification on a fingertip-scanned electrotactile display. In *Proceedings of the 19th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 4, pages 1694–1696, Chicago, IL, 1997. IEEE.
- [73] Nels P. Ostrom, Kurt A. Kaczmarek, and David J. Beebe. A microfabricated electrocutaneous tactile display. In *Proceedings of the First Joint BMES/EMBS Conference.*, volume 2, page 838, Atlanta, GA, 1999. IEEE.
- [74] Darwin G. Caldwell, N. Tsagarakis, and C. Giesler. An Integrated Tactile/Shear Feedback Array for Stimulation of Finger Mechanoreceptor. In *Proceedings 1999 IEEE International Conference on Robotics and Automation*, volume 1, pages 287–292, Detroit, MI, 1999. IEEE.
-

-
- [75] G. Moy, C. Wagner, and R.S. Fearing. A compliant tactile display for teletaction. In *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, volume 4, pages 3409–3415, San Francisco, CA, 2000. IEEE.
- [76] Edward Mallinckrodt, A. L. Hughes, and William Sleator. Perception by the Skin of Electrically Induced Vibrations. *Science*, 118:277–278, 1953.
- [77] Kurt A. Kaczmarek, Krishnakant Nammi, Abhishek K. Agarwal, Mitchell E. Tyler, Steven J. Haase, and David J. Beebe. Polarity Effect in Electro-vibration for Tactile Display. *IEEE transactions on bio-medical engineering*, 53(10):2047–2054, October 2006.
- [78] Hui Tang and David J. Beebe. A microfabricated electrostatic haptic display for persons with visual impairments. *IEEE Transactions on Rehabilitation Engineering*, 6(3):241–248, 1998.
- [79] Robert M. Strong and Donald Troxel. An Electrotactile Display. *IEEE Transactions on Man Machine Systems*, 11(1):72–79, March 1970.
- [80] W. R. Provancher and N. D. Sylvester. Fingerpad skin stretch increases the perception of virtual friction. *IEEE Transactions on Haptics*, 2(4):212–223, 2009.
- [81] Erik C. Chubb, J. Edward Colgate, and Michael A. Peshkin. ShiverPad: A device capable of controlling shear force on a bare finger. In *World Haptics 2009 - Third Joint EuroHaptics conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 18–23, Salt Lake City, UT, 2009. IEEE.
- [82] Martin Rothenberg, Ronald T. Verrillo, Stephen A. Zahorian, Michael L. Brachman, and Stanley J. Bolanowski. Vibrotactile frequency for encoding a speech parameter. *The Journal of the Acoustical Society of America*, 62(4):1003–1012, October 1977.
- [83] Lynne E. Bernstein, Silvio P. Eberhardt, and Marilyn E. Demorest. Single-channel vibrotactile supplements to visual perception of intonation and stress. *The Journal of the Acoustical Society of America*, 85(1):397–405, January 1989.
- [84] Ian R. Summers, Alan Christopher Brady, M.K. Syed, and Craig M. Chanter. Design of Array Stimulators for Synthetic Tactile Sensations. In *First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 586–587, Pisa, Italy, 2005. IEEE.
- [85] C. E. Sherrick. A scale for rate of tactual vibration. *The Journal of the Acoustical Society of America*, 78(1 Pt 1):78–83, July 1985.
- [86] Ahmad Barghout, Jongeun Cha, Abdulmotaleb El Saddik, Julius Kammerl, and Eckehard Steinbach. Spatial resolution of vibrotactile perception on the human forearm when exploiting funneling illusion. In *2009 IEEE International Workshop on Haptic Audio visual Environments and Games*, pages 19–23. IEEE, November 2009.
-

-
- [87] Sang-Youn Kim and Jeong Cheol Kim. Vibrotactile rendering for a traveling vibrotactile wave based on a haptic processor. *IEEE Transactions on Haptics*, 5(1):14–20, 2012.
- [88] Ian R. Summers, Craig M. Chanter, Anna L. Southall, and Alan Christopher Brady. Results from a Tactile Array on the Fingertip. In *Proceedings of Eurohaptics 2001*, pages 26–28, 2001.
- [89] Greg Welch and Gary Bishop. *An Introduction to the Kalman Filter*. http://www.cs.unc.edu/~welch/media/pdf/kalman_intro.pdf, 1997.
- [90] S. Kawabata. The Standardisation and Analysis of Hand Evaluation. Technical report, The Textile Machinery Society of Japan, Osaka Science and Technology Centre Bld., 8-4 Utsubo-1-chome, Nishi-ku, Osaka 550, Japan, July 1980.
- [91] Pier Giorgio Minazio. FAST Fabric Assurance by Simple Testing. *International Journal of Clothing Science and Technology*, 7(2/3):43–48, 1995.
- [92] HAPTEX Consortium. *First set of measurements*. HAPTEX Deliverable D3.1. <http://haptex.miralab.unige.ch/>, 2005.
- [93] HAPTEX Consortium. *Second set of fabric measurements*. HAPTEX Deliverable D3.2. <http://haptex.miralab.unige.ch/>, 2006.
- [94] HAPTEX Consortium. *Database of properties for various kind of fabrics*. HAPTEX Deliverable D3.2 Part 2. <http://haptex.miralab.unige.ch/>, 2007.
- [95] HAPTEX Consortium. *Specifications of the Whole Haptic Interface*. HAPTEX Deliverable D4.1. <http://haptex.miralab.unige.ch/>, 2005.
- [96] Nadia Magnenat-Thalmann, P. Volino, U. Bonanni, Ian R. Summers, M. Bergamasco, F. Salsedo, and Franz-Erich Wolter. From Physics-based Simulation to the Touching of Textiles: The HAPTEX Project. *International Journal of Virtual Reality*, 6(3):35–44, 2007.
- [97] HAPTEX Consortium. *Whole Haptic Interface: Report on the Development Activities*. HAPTEX Deliverable D4.3. <http://haptex.miralab.unige.ch/>, 2008.
- [98] HAPTEX Consortium. *Pre-final Demonstrator*. HAPTEX Deliverable D5.1. <http://haptex.miralab.unige.ch/>, 2007.
- [99] Pascal Volino, Pierre Davy, Ugo Bonanni, Guido Böttcher, Dennis Allerkamp, and Franz-Erich Wolter. From measured physical parameters to the haptic feeling of fabric. In *Proceeding of the HAPTEX '05 Workshop on Haptic and Tactile Perception of Deformable Objects*, pages 17–29, Hanover, 2005.
- [100] Pascal Volino, Pierre Davy, Ugo Bonanni, Christiane Luble, Nadia Magnenat-Thalmann, Mailis Mäkinen, and Harriet Meinander. From measured physical parameters to the haptic feeling of fabric. *The Visual Computer*, 23(2):133–142, July 2006.
-

-
- [101] Pascal Volino and Nadia Magnenat-Thalmann. Accurate Anisotropic Bending Stiffness on Particle Grids. In *2007 International Conference on Cyberworlds (CW'07)*, pages 300–307. IEEE Comput. Soc, October 2007.
- [102] HAPTEX Consortium. *Separate haptic and tactile interfaces*. HAPTEX Deliverable D4.2. <http://haptex.miralab.unige.ch/>, 2006.
- [103] C.A. Avizzano, S. Marcheschi, M. Angerilli, M. Fontana, M. Bergamasco, T. Gutierrez, and M. Mannegeis. A multi-finger haptic interface for visually impaired people. In *Proceedings of the 2003 IEEE International Workshop on Robot and Human Interactive Communication*, pages 165–170. IEEE, 2003.
- [104] Alan Christopher Brady. *Tactile Arrays for Virtual Textures*. PhD thesis, University of Exeter, 2010.
- [105] HAPTEX Consortium. *Final Demonstrator and final integration report*. HAPTEX Deliverable D5.2. <http://haptex.miralab.unige.ch/>, 2007.
- [106] Christiane Luble, Minna Varheenmaa, Nadia Magnenat-Thalmann, and Harriet Meinander. Subjective Fabric Evaluation. In *2007 International Conference on Cyberworlds*, pages 285–291. IEEE, October 2007.
- [107] Matthew Philpott, Ian R. Summers, and Dennis Allerkamp. Characteristics of a Tactile Rendering Algorithm. In *Proceedings of Joint VR Conference of euroVR and EGVE*, pages 77–79, Nottingham, UK, 2011.
- [108] Allan M. Smith, Geneviève Gosselin, and Bryan Houde. Deployment of fingertip forces in tactile exploration. *Experimental Brain Research*, 147(2):209–218, November 2002.
- [109] Matthew Philpott and Ian R. Summers. Evaluating a Multipoint Tactile Renderer for Virtual Textured Surfaces. In *Proceedings of Eurohaptics 2012*, pages 121–126, Tampere, Finland, 2012. Springer.
- [110] Raymond B. Cattell. *The Scientific Use of Factor Analysis in Behavioural and Life Sciences*. Plenum Press, 1978.
- [111] Matthew Philpott and Ian R. Summers. Active exploration of virtual surfaces. In *EuroMov Inaugural Conference*, pages 91–92, Montpellier, France, 2013. VTT.
- [112] Y. Mukaibo, H. Shirado, Masashi Konyo, and T. Maeno. Development of a Texture Sensor Emulating the Tissue Structure and Perceptual Mechanism of Human Fingers. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pages 2565–2570. IEEE, 2005.
- [113] P. Dario and G. Buttazzo. An Anthropomorphic Robot Finger for Investigating Artificial Tactile Perception. *The International Journal of Robotics Research*, 6(3): 25–48, September 1987.
- [114] Calum Roke, Adam Spiers, Tony Pipe, and Chris Melhuish. The effects of laterotactile information on lump localization through a teletaction system. In *World Haptics Conference*, pages 365–370. IEEE, 2013.
-

Appendix A





Texture Database

This appendix features a full list of available textures, including the textile number, the available faces (right or back), a description, the content and the structure of the textile (table A.2). Table A.1 gives the abbreviations used for the content section of the database. In the rendering software, particular textures are identified by a combination off the textile number and required face (for example, 03_R, 28_B, etc.). The selection of these textiles, and how they are converted into texture file that may be used by the renderer, is described in section sub:FileCreation, with more detail of the selection of the particular textiles available in [92], [93], and [94].

Table A.1: Abbreviations of the fibre contents of the textiles.






AF	Other fibers	CV	Viscose	PAN	Polyacryl
CA	Acetat	EL	Elastane	PES	Polyester
CLY	Lyocell	JU	Jute	PU	Polyurethane
CMD	Modal	LI	Linen	SE	Silk
CO	Cotton	PA	Polyamide	WO	Wool

Table A.2: Database of the available textures (replicated from [1]).

#	Image	Available Faces	Description	Content %	Structure	Weight g m^{-2}	Thickness mm
01		R	Denim	100 CO	Twill	380	1.60
02		R	Shirt Cotton	100 CO	Combined twill	120	0.61
03		R	Cord	100 CO	Velveteen	330	1.76
04		R	Linin	100 LI	Plain weave	250	1.09


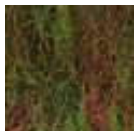
Continued on next page

Continued from previous page

#	Image	Available Faces	Description	Content %	Structure	Weight g m^{-2}	Thickness mm
05		R	Gabardine	100 WO	Twill	175	0.55
06		R	Crepe	100 WO	Plain weave	145	0.93
07		R	Silk	100 SE	Plain weave	15	0.10
08		R	Natural silk (bourette)	100 SE	Plain weave	150	0.80
09		R	Wild silk (dupion)	100 SE	Plain weave	80	0.44






Continued on next page

Continued from previous page

#	Image	Available Faces	Description	Content %	Structure	Weight g m^{-2}	Thickness mm
10		R	Jute	100 JU	Plain weave	300	1.44
11		R	Flannel	80/20 WO/PES	Twill	290	1.53
12		R	Denim	62/35/3 PES/CO/EL	Twill	275	1.13
13		R	Plaid	35/35/30 PES/AF/WO	Twill	270	1.14
14		R	Tweed	66/14/10/10 AF/WO/PES/CMD	Combined twill	270	3.90






Continued on next page

Continued from previous page

#	Image	Available Faces	Description	Content %	Structure	Weight g m^{-2}	Thickness mm
15		R	Velvet	92/8 CO/CMD	Velvet	300	1.88
16		R	Lurex knit	70/30 PES/PA	Held stitch knit	215	2.94
17		R	Crepe-jersey	85/15 PES/EL	Rib knit	135	0.73
18		R	Motorcyclist, coated	72/18 PA/PU	Plain weave	90	0.39
19		R	Woven easy care	65/35 PES/CO	Twill	180	0.57






Continued on next page

Continued from previous page

#	Image	Available Faces	Description	Content %	Structure	Weight g m^{-2}	Thickness mm
20		R	Warp knitted velour	90/10 PA/EL	Warp knit velour	235	1.56
21		R	Weft knitted plain	98/2 CLY/EL	Single jersey	172	1.21
22		R	Taffeta	100 CA	Plain weave	125	0.33
23		R	Crepe	100 PES	Plain weave	85	0.25
24		R	Satin	100 PES	Satin	125	0.30





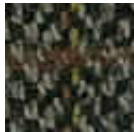
Continued on next page

Continued from previous page

#	Image	Available Faces	Description	Content %	Structure	Weight g m^{-2}	Thickness mm
25		R	Felt	100 PES	Non-woven	155	1.25
26		R	Organza	100 PES	Plain weave	25	0.16
27		R	Fleece	100 PES	Weft knit	250	3.99
28		B/R	Woven jacquard upholstery	100 PES	Woven	600	2.38
29		R	Woven leisure	100 PES	Plain weave	90	0.20






Continued on next page

Continued from previous page

#	Image	Available Faces	Description	Content %	Structure	Weight g m^{-2}	Thickness mm
30		R	Tulle	100 PA	Warp knitted tulle	10	0.30
31		R	Warp knitted tricot-satin	100 PA	Warp knitted tricot-satin	100	0.40
32		R	Leather	100 Leather	N/A	815	1.68
33		B/R	Men's woven suit fabric (plain)	60/38/3 WO/PES/EL	Plain weave	195	0.57
34		B/R	Men's woven suit fabric (herringbone)	100 WO	Broken twill	232	0.83


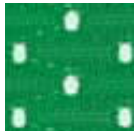



Continued on next page

Continued from previous page

#	Image	Available Faces	Description	Content %	Structure	Weight g m^{-2}	Thickness mm
35		B/R	Men's woven overcoat fabric	80/20 WO/PA	Modified plain weave	324	2.64
36		B/R	Men's woven overcoat fabric (twill)	59/25/11/5 CO/PAN/WO/PES	Twill	460	3.23
37		B/R	Woven outdoor leisurewear fabric	100 PES	Plain weave	98	0.26
38		B/R	Weft knitted jersey fabric	48/48/4 CO/CMD/EL	Single jersey	208	1.09
39		B/R	Weft knitted terry fabric	55/45 CV/PES	Weft knit terry	288	1.09






Continued on next page

Continued from previous page

#	Image	Available Faces	Description	Content %	Structure	Weight g m^{-2}	Thickness mm
40		B/R	Warp knitted jersey-based fabric	100 PES	Warp knit jersey-based	154	0.51
41		B/R	Warp knitted mesh fabric	100 PES	Warp knit mesh	128	0.51
42		B/R	Brushed warp knitted fabric	100 PES	Brushed warp knit	215	0.98
43		B/R	Warp knit for car seats	89/11 PES/EL	Pile warp knit	278	1.22
44		B/R	Warp knit for car seats	100 PES	Flattened pile warp knit	168	1.10






Continued on next page

Continued from previous page

#	Image	Available Faces	Description	Content %	Structure	Weight g m^{-2}	Thickness mm
45		B/R	Warp knit for car seats	100 PES	Raised pile warp knit	147	1.60
46		B/R	Warp knit for car seats	100 PA	Raised pile warp knit	98	1.55
47		B/R	Laminated warp knit for car seats	Face fabric: 100 PES, Foam: E-35 PF (35 kg m^{-3} polyether), Backing: 100 PES	Laminated brushed warp knit	386	4.09
48		B/R	Laminated warp knit for car seats	Face fabric: 89/11 PES/EL, Foam: E-35 PF (35 kg m^{-3} polyether), Backing: 100 PA	Laminated pile warp knit	438	4.19
49		B/R	Laminated warp knit for car seats	Face fabric: 100 PES, Foam: E-35 PF (35 kg m^{-3} polyether), Backing: 100 PES	Laminated flattened pile warp knit	335	4.11

Continued on next page

Continued from previous page

#	Image	Available Faces	Description	Content %	Structure	Weight g m^{-2}	Thickness mm
50		B/R	Laminated warp knit for car seats	Face fabric: 100 PES, Foam: L30PS (30 kg m^{-3} polyester), Backing: 100 PES	Laminated raised pile warp knit	332	3.54
51		B/R	Laminated warp knit for car seats	Face fabric: 100 PA, Foam: E-25 TX (25 kg m^{-3} polyether)	Laminated raised pile warp knit	170	4.03
52		None	Backing	100 PES	Warp knit	53	0.42
53		B/R	Backing	100 PA	Interlock	121	0.81
54		R	Laminated warp knit for car seats	Face fabric: 89/11 PES/EL, Foam: L30PS (30 kg m^{-3} polyester)	Laminated pile warp knit	287	3.35

Appendix B

Spectra of Texture Files

This table compares the spatial frequencies of the spectral lines associated with the linear (lin) and logarithmic (log) approaches to the textile files. Both of these set of spectral lines originate from the full 2048-line spectrum that results from the Fast Fourier Transform of the textiles surface profile. In the linear case, the first 50 lines are taken directly as the spectra and are stored in the textile file. In the logarithmic case, the first 500 lines are taken and compressed down to 50 lines by binning a number of the lines together using a root of the sum of the squares (root-sum-square) calculation. At low spatial frequencies these bins are small, containing one or two lines, and, as spacial frequency increases, the number of lines included in the bins increases.

Table B.1: The spectral lines of the texture files, numbered 1 to 50. The linear (lin) and logarithmic (log) values are spatial frequencies, all with units of m^{-1} .

Line #	Lin	Log	Line #	Lin	Log	Line #	Lin	Log
01	229	198	18	4121	5484	35	8013	28592
02	458	443	19	4350	6172	36	8242	31222
03	687	677	20	4579	6860	37	8471	34082
04	916	909	21	4808	7548	38	8700	37280
05	1145	1139	22	5037	8344	39	8929	40827
06	1374	1369	23	5266	9261	40	9158	44713
07	1603	1599	24	5495	10178	41	9387	48948
08	1832	1828	25	5724	11204	42	9616	53522
09	2061	2057	26	5953	12350	43	9845	58553
10	2290	2287	27	6182	13606	44	10074	64044
11	2519	2516	28	6411	14981	45	10303	69992
12	2747	2745	29	6640	16465	46	10532	76508
13	2976	3082	30	6869	18070	47	10761	83712
14	3205	3542	31	7098	19783	48	10990	91604
15	3434	4000	32	7327	21726	49	11219	100183
16	3663	4459	33	7556	23898	50	11448	109559
17	3892	4917	34	7784	26190	-	-	-

Appendix C

Chance Result Derivation

The probability distribution of the number of subjects that chose the “most popular” (mode) choice, that is recorded from the experiments, may be compared to the probability distribution of the case where the subjects were simply picking textiles randomly from the virtual environments, i.e. the case where a number of subjects might randomly select the same textile from the virtual workspace. For the comparison it is therefore necessary to calculate what such a random probability distribution would look like.

There are a total of 7 possible values of the mode choice, ranging from 2 to 8, and so the number of microstates associated with each mode choice value must be calculated. The total number of available microstates is $4^8 = 65536$, which covers all possible distributions of the 8 subjects within the 4 virtual textiles. Table C.1 summarises the distribution of the microstates among the possible mode scores.

The simplest case is where all 8 of the subjects choose the same textile. As the order of choice does not matter, all subjects choosing the same textile is a single microstate. As there are 4 textiles available, there is therefore a total of 4 microstates available for this case.

In case of 7 of the subjects choosing the same textile, in addition to the 4 textiles that could be selected as the mode choice, the question of which 7 of the 8 total subjects made the most popular textile their choice must be addressed. The number of microstates must therefore include the number of possible combinations that exist for 7 subjects taken from a total pool of 8. Finally, the question remains of the choice of the final subject, of which there are 3 possibilities. The total number of available microstates for this case is therefore

$$4 * \binom{8}{7} * 3 = 96. \tag{C.1}$$

The case of 6 subjects selecting the mode choice is similar to the case of 7 subjects, with two main differences. The first is that the number of microstates is related to the number of combinations of 6 from 8, rather than 7. The second difference is that two remaining subjects could have any distribution within the remaining 3 textiles, which has 3^2 possible combinations. The number of microstates available for this case is

$$4 * \binom{8}{6} * 3^2 = 1008. \quad (\text{C.2})$$

The case of 5 is also similar to 7, again with the number of possible combinations being 5 out of a pool of 8 and the remaining subjects being distributed amount the remaining textiles. The number of microstates available in this case is therefore

$$4 * \binom{8}{5} * 3^3 = 6048. \quad (\text{C.3})$$

At first glance, the case of a mode score of 4 would be similar. However, there is the possibility that the remaining subjects would also collect on the same textile, resulting in the subjects being evenly split among two textiles. For each of the 4 textiles, there are 3 cases where the other subjects would group to one textile. Half of these cases will have the subjects will be divided between the same two textiles as the other, for example in the the case of textiles 1 and 2 being “chosen” versus the case of 2 and 1 being “chosen”. Therefore is is necessary to remove half of these cases from the microstates consideration, making the total

$$4 * \binom{8}{4} * \left(3^4 - \frac{3}{2}\right) = 22260. \quad (\text{C.4})$$

Similarly, for cases of the mode score of 3, the cases where the remaining subjects cluster onto one textile to give 4 or 5 scores need to be removed, as well as have of the case where the other subjects cluster to give a 3 score, as previously with the 4 score. For the 5 case, where the subjects are all collected onto one textile, there are 3 possibilities, as there are three other textiles which the 5 choice may occupy. For the 4 case, there are still 3 possibilities but multiplied by the number of combinations available from selecting 4 subjects from a pool of 5, and further multiplied by 2 due to the position of the final subject in the remaining two textiles. The case of the 3 score is similar to the 4 but, as with the 6 mode, the combinations are 3 from a pool of 5 and the distribution of the remaining textiles is increase to 2^2 . The resulting total number of microstates for this case is then

$$4 * \binom{8}{3} * \left(3^5 - 3 - \left(2 * 3 * \binom{5}{4}\right) - \frac{\binom{5}{3} * 2^2 * 3}{2}\right) = 33600 \quad (\text{C.5})$$

For the case of the mode score being 2, there is only one situation where that is possible i.e. all of the textiles contain two subjects from the available 8. The number of available microstates can therefore be calculated as the number of combinations of 2 available from the pool of 8, multiplied by the number of combinations of 2 available from the remaining pool of 6, further multiplied by the number of combinations of 2 from the remaining pool of 4, resulting in a total of

$$\binom{8}{2} * \binom{6}{2} * \binom{4}{2} = 2520 \quad (\text{C.6})$$

Table C.1: The number of microstates available for each number of subjects that selected the most popular choice, and the total number of microstates.

Score of Mode Choice	Number of Microstates
8	4
7	96
6	1008
5	6048
4	22260
3	33600
2	2520
Total:	65536

Appendix D

Ethics Application

The experiments detailed in chapter 6 could only be performed following the approval of the ethics committee at the University of Exeter. The application to the committee, “2010/270 - Evaluating a Setup for Tactile Rendering”, included a description of the research project which read,

“This project uses similar protocols and techniques to those used in previous projects in the School of Physics (Dr Syed Mustafa Kamal and Dr Sarah Carr). For this project, computer generated virtual textiles are applied to the fingertip by means of a tactile display, consisting of an array of small contactors. These contactors provide vibratory stimulation at a comfortable level of sensation. There is no reason to believe that such touch stimulation poses any risk to the test subject.

During the experiments, the subjects will be asked to use the tactile display to explore a number of real and/or virtual fabric surfaces. The aim of the experiments is to test the effectiveness of the tactile renderer using psychophysical techniques to determine the discriminability of the virtual surface textures.

Test sessions will be approximately 30 minutes in duration, preceded by a short demonstration / training session. During testing, subjects wear earphones that provide a comfortable level of background noise to mask any acoustic cues. There is no reason to believe that such acoustic stimulation poses any risk to the test subject.”

In addition to the description, the application included copies of the information sheets that would be given to a potential subject, detailing what the subject could expect from the project (figure D.1), details of the tactile display being used (figure D.2), as well as a consent form that would be signed, should they agree to participate (figure D.3). These information sheets and consent form were originally printed at A4 size.

The ethics approval to perform the experiments detailed in this work was received on September 16, 2011.

2

What will happen to the results of the research study?

The results from this study will be published in the form of papers submitted to scientific journals and also presentations at relevant scientific conferences. You will not be identified in any report or publication.

What if I have any questions?

If you have any question at any time, please feel free to contact Matthew Philpott on **01392 724142** or Email: **mp236@ex.ac.uk**

Can I change my mind?

You can stop being part in the study at any time without giving a reason and there will be no problem with this.

Matthew Philpott

PhD Researcher

This project has been reviewed and approved by the Local Research Ethics Committee

School of Physics
University of Exeter
Stocker Road
EX4 4QL

Matthew Philpott
Tel: 01392 724142
mp236@ex.ac.uk

1

Information Sheet**Invitation**

We would like to invite you to participate in our research study into evaluating our setup for virtually recreating the feel of textiles. Before you decide you need to understand why the research is being done and what it would involve for you. Please take the time that you need to read the following information and talk to others about the study if you wish. Please ask us to clarify if there is anything that you want more information about or is not clear to you. Please take your time in deciding whether or not you wish to participate in this study.

The purpose of the study

The creation of virtual surfaces could add another layer to virtual reality environments by adding touch feedback to existing visual and audio feedback. The aim of this study is to evaluate the setup we have for creating virtual textiles by comparing the virtual surfaces to their real equivalents. This experiment will be carried out in Physics Building, Exeter University.

Why have I been invited to take part?

The people who will take part in the study will be in good health between ages 18 and 35. We believe you are appropriate to participate in this experiment.

What will happen to me if I take part?

This study uses a tactile display to provide the skin with touch stimulation (**please see information sheet 2**). One of your fingertips will be resting lightly on the tactile display which will vibrate in a non-painful and non-invasive way as you explore the textile surfaces. The same fingertip will also be used to explore the surface of a real textile. The duration of the experiment will be 30 minutes. During the experiment you will be asked to identify which of the presented virtual surfaces is the closest match to the presented real textile.

If you decide to be a participant in the study there are some forms you will ask to fill in before we start the experiment. These forms are registration form (**to record your details**) and the consent form (**you agree to be part in this study**).

School of Physics
University of Exeter
Stocker Road
EX4 4QL


Matthew Philpott
Tel: 01392 724142
mp236@ex.ac.uk

Figure D.1: The first information sheet given to potential subjects, detailing the purpose of the experiment and what subject could expect if they agree to participate.

Information Sheet 2

The Tactile Display

The tactile display consists of an array of contactors, arranged in a 6×4 array (see the image below). During the experiment the pins will vibrate in a way that will non-invasively stimulate the skin’s touch receptors. The sensation should be comfortable.



The tactile display



The 24 pins

This project has been reviewed and approved by the Local Research Ethics Committee.

School of Physics
University of Exeter
Stocker Road
Ex4 4QL

Matthew Philpott
Tel: 01392 724142
mp236@ex.ac.uk

Figure D.2: The second information sheet given to potential subjects, detailing the tactile display used in the experiment.

Consent Form for the participant

I have read all the information sheets concerning this study and understand what it is about, all my questions have been answered to my satisfaction (please initial each box).

- 1- I confirm I have read and understood the information sheets. I have had the opportunity to consider the information, ask questions and have had these answered satisfactorily.
- 2- I understand that my participation is voluntary and that I am free to withdraw at any time without giving any reason.
- 3- I understand the results and the data that are relevant to me may be looked at by individuals from the research team. I give permission for these individuals to have access to my records.
- 4- I will provide information about my age and left or right handedness.
- 5- I agree to take part in this study.

Signed Participant _____ Date _____

Signed Researcher _____ Date _____

This project has been reviewed and approved by the Local research Ethics Committee

School of Physics
University of Exeter
Stocker Road
Ex4 4QL

Matthew Philpott
Tel: 01392 724142
mp236@ex.ac.uk

Figure D.3: The consent form to be signed by individuals who agree to participate in the study.

Appendix E

Additional Discriminability Plots

This appendix contains the additional plots that detail relationship between the correlation coefficients from the columns in the confusion matrix generated from the results of experiment 2, and so measuring the similarity between that virtual textiles, and the distances along various proposed characterisation axes, in this case the mean and coefficient of variation (CoV) values of the 40 Hz and 320 Hz channels. The distance between any two textiles is the Euclidean distance between the textures on the axis which, given that these plots represent only one dimension, is simply the modulus of the difference between the two characteristic values.




Figure E.1: Graph showing the relationship between the correlation coefficient between virtual textures and the Euclidean distance between mean 40 Hz amplitudes. The fitted line has the equation $y = -2.22x + 0.35$, with an $R^2 = 0.34$.




Figure E.2: Graph showing the relationship between the correlation coefficient between virtual textures and the Euclidean distance between mean 320 Hz amplitudes. The fitted line has the equation $y = -1.43x + 0.10$, with an $R^2 = 0.01$.




Figure E.3: Graph showing the relationship between the correlation coefficient between virtual textures and the Euclidean distance between CoV 40 Hz amplitudes. The fitted line has the equation $y = -0.91x + 0.11$, with an $R^2 = 0.02$.




Figure E.4: Graph showing the relationship between the correlation coefficient between virtual textures and the Euclidean distance between CoV 320 Hz amplitudes. The fitted line has the equation $y = -1.09x + 0.16$, with an $R^2 = 0.05$.