

Force Shading for Shape Perception in Haptic Virtual Environments

by

Hugh Brian Morgenbesser

S.B., Massachusetts Institute of Technology (1994)

Submitted to the Department of Electrical Engineering and
Computer Science

in partial fulfillment of the requirements for the degree of
Master of Engineering in Electrical Engineering and Computer
Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1995

© Massachusetts Institute of Technology 1995. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
August 31, 1995

Certified by
Mandayam A. Srinivasan
Principal Research Scientist
Department of Mechanical Engineering
Thesis Supervisor

Accepted by

Frederick P. M. Mowenthaler

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

Chairman, Department Committee on Graduate Theses

JAN 29 1996

LIBRARIES

ARCHIVES

Force Shading for Shape Perception in Haptic Virtual Environments

by

Hugh Brian Morgenbesser

Submitted to the Department of Electrical Engineering and Computer Science
on August 31, 1995, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

This thesis proposes a novel haptic rendering algorithm and describes two experiments that demonstrate its effectiveness in causing humans to perceive shapes in haptic virtual environments. Haptic rendering algorithms control haptic interfaces, the robotic devices which enable manual interaction between humans and virtual environments. The algorithm proposed in this thesis, called force shading, refers to a controlled variation in the direction of the force vector output by the haptic renderer for the purpose of creating the illusion of a non-flat shape on a nominally flat surface. The two experiments, a matching experiment and a classification experiment, were done on 5 subjects.

In the matching experiment, subjects were asked to change the size of a virtual “matching” bump until it matched a virtual “reference” bump. The reference bump was always a circular cylinder, while the matching bump was either another circular cylinder or one of three force shaded polyhedral approximations. As the number of polygons increased in the polyhedral approximations, the subjects’ performance approached their performance in matching cylinders. Even for the nominally flat surface with force shading, subjects correlated their matching radii with the reference radii. When visual stimuli was added to the polyhedral approximations, performance was even closer to that in matching two cylinders.

In the classification experiment, subjects were asked to classify the shape of various virtual polyhedral approximations to a circular cylindrical bump. Without force shading, the subjects accurately identified the polyhedral nature of the virtual bumps. With force shading, however, the subjects identified the polyhedral approximations as conveying approximately the same feel as a smooth cylinder.

Thesis Supervisor: Mandayam A. Srinivasan
Title: Principal Research Scientist
Department of Mechanical Engineering

Acknowledgments

Thanks to:

Srini, for introducing me to haptics and giving me the opportunity to work on such a great project.

My other colleagues at the Touch Lab, who have made it a wonderful place to work this past year.

My experimental subjects for their patience, and most importantly their data.

Yev, Steve, Len, Pete, Kwong, Paulo, Danielle, and Dan, for their friendship and company this summer as I finished this.

SensAble Devices for providing me with a Phantom, and for their help in keeping it up and running- especially for their immediate, last minute repair.

And my family, for their love.

Contents

1	Background	10
1.1	Virtual Environments	10
1.2	Haptics in Virtual Environments	10
1.2.1	Machine Haptics	12
1.2.2	Human Haptics	14
1.3	Multimodal Perception in VEs	15
2	Haptic Rendering Algorithms	16
2.1	Introduction	16
2.2	Vector-Field Based Rendering Algorithms	17
2.2.1	Simulation of Thick Walls	18
2.2.2	Simulation of Cylinders	18
2.3	Problems with Vector-Field Based Approach	19
2.4	God-Object Idea	20
2.5	Height Maps	21
2.5.1	Computation of Surface Normals	21
3	Force Shading for Shape Perception	23
3.1	Introduction to Force Shading	23
3.2	The Force Shading Algorithm	24
3.3	Application to Polygonal Mesh	25
4	The Bump Display Algorithms	27

4.1	Introduction	27
4.2	The Geometry	29
4.3	The Cylindrical Bump	31
4.4	The Unshaded Polygonal Approximations	31
4.4.1	Single Polygon Without Force Shading	31
4.4.2	Two Polygons Without Force Shading	32
4.4.3	Three Polygons Without Force Shading	33
4.5	The Shaded Polygonal Approximations	34
4.5.1	Single Polygon With Force Shading	35
4.5.2	Two Polygons With Force Shading	35
4.5.3	Three Polygons With Force Shading	35
4.6	Comments	35
5	Experimental Design	36
5.1	Experimental Goals	36
5.2	The Theta-critical Hypothesis	36
5.3	Experimental Paradigms	38
5.4	The Experimental Setup	38
5.5	The Matching Experiments	39
5.5.1	Instructions to Subjects	44
5.5.2	Exit Interviews	45
5.5.3	History Recording	45
5.6	The Classification Experiment	46
5.6.1	Instructions to Subjects	49
6	Results	50
6.1	The Matching Experiment	50
6.1.1	Cylindrical Bump Matching	51
6.1.2	Flat Bump Matching	58
6.1.3	Two Polygon Matching	62
6.1.4	Three Polygon Matching	66

6.2	The Classification Experiment	70
7	Discussion	72
8	Future Work	75

List of Figures

1-1	The interrelationship between the two main areas of research for building high-performance haptic VE systems.	11
1-2	Cross section of two sample virtual objects.	13
2-1	Block Diagram for Input/Output of Haptic Renderer	16
2-2	A piece of a wall with normal vector, \hat{n} , through x_0	18
2-3	A heightmap function.	22
3-1	Method for Applying Force Shading Algorithm to Smooth a Polygonal Mesh	26
4-1	Cross Section of a Cylindrical Bump on a Flat Plane	28
4-2	Cross section of the four different nominal geometries used in the bump display algorithms.	28
4-3	Labels for the Geometry for the various Haptic Simulations	30
4-4	Geometry for a Two-Polygon Approximation to the Cylindrical Bump	33
4-5	Geometry for a Three-Polygon Approximation to the Cylindrical Bump	34
5-1	Profiles of three different constant-curvature bumps with the same width.	37
5-2	The Experimental Setup	38
5-3	The visuals for the matching experiment	40
5-4	The heights of the bumps as a function of radius	41
5-5	The amount of angle subtended by the cylinders divided by the number of polygons used to approximate the cylinder.	43
5-6	The initial values of the reference and matching bumps.	43

5-7	The instructions given to the subjects before the matching experiment.	44
5-8	The visuals for the classification experiment	48
5-9	The instructions given to the subjects before the classification experiment.	49
6-1	Position profile for the cylindrical bump matching task.	54
6-2	Results for Cylinder, No Vision	55
6-3	Results for Cylinder, With Vision	55
6-4	Satisfaction Index Results for Real No Vision	56
6-5	Satisfaction Index Results for Real With Vision	56
6-6	Radius profiles for the cylindrical bump matching task.	57
6-7	Position profile for the flat bump matching task.	59
6-8	Results for Flat Bump Matching, No Vision	60
6-9	Results for Flat Bump Matching, With Vision	60
6-10	Satisfaction Index Results for Flat Bump Matching, No Vision	61
6-11	Satisfaction Index Results for Flat Bump Matching, With Vision	61
6-12	Position profile for the two polygon bump matching task.	63
6-13	Results for Two Polygon Bump Matching, No Vision	64
6-14	Results for Two Polygon Bump Matching, With Vision	64
6-15	Satisfaction Index Results for Two Polygon Bump Matching, No Vision	65
6-16	Satisfaction Index Results for Two Polygon Bump Matching, With Vision	65
6-17	Position profile for the three polygon bump matching task.	67
6-18	Results for Three Polygon Bump Matching, No Vision	68
6-19	Results for Three Polygon Bump Matching, With Vision	68
6-20	Satisfaction Index Results for Three Polygon Bump Matching, No Vision	69
6-21	Satisfaction Index Results for Three Polygon Bump Matching, With Vision	69

List of Tables

- 5.1 Matching Experiment Conditions 39
- 5.2 Classification Experiment Conditions 46

- 6.1 Summary of Matching Results for the Five Subjects for Cylindrical Bump Without Vision 53
- 6.2 Summary of Matching Results for the Five Subjects for Cylindrical Bump With Vision 53
- 6.3 Summary of Matching Results for the Five Subjects for Flat Bump Without Vision 59
- 6.4 Summary of Matching Results for the Five Subjects for Flat Bump With Vision 59
- 6.5 Summary of Matching Results for the Five Subjects for Two Polygon Bump Without Vision 62
- 6.6 Summary of Matching Results for the Five Subjects for Two Polygon Bump With Vision 62
- 6.7 Summary of Matching Results for the Five Subjects for Three Polygon Bump Without Vision 66
- 6.8 Summary of Matching Results for the Five Subjects for Three Polygon Bump With Vision 66
- 6.9 Summary of Classification Results for the Five Subjects, as a percentage of the total number of trials for each case. 71

- 7.1 Summary of Matching Results for the Five Subjects for Three Polygon Bump With Vision 72

Chapter 1

Background

1.1 Virtual Environments

Virtual environments (VEs) are computer generated environments with which human users can have a sense of natural interaction. In order to accomplish this, VE systems consist of interface devices which mimic the sensory and motor interactions that the user normally encounters in the real environment. The three main perceptual modalities which give humans a sense of presence in a VE are vision, audition, and haptics. This thesis is an investigation of issues at the boundary between the design of haptic interfaces and research into fundamentals of human haptic sensory abilities.

1.2 Haptics in Virtual Environments

In VE systems, haptic interfaces are the components interposed between the user and the environment for manual sensing and manipulation. Examples of haptic interfaces include the computer mouse, a joystick, and force-reflecting robotic devices. At present, the fidelity of visual and auditory stimulators is much higher than that of available haptic interface systems. Because by comparison the visual and auditory VE systems are so sophisticated and well understood, some of the most significant improvements to multimodal VE systems consist of advancing the capabilities of haptic interfaces.

The work required to improve the haptic component of VE systems is in two broad categories: machine haptics and human haptics. The current needs in machine haptics are to build better haptic interface mechanisms and their software controllers. The current needs in human haptics are to find out what needs to be displayed so that the user experiences an intended perceptual effect. These two areas, although distinct, are not separate in their development. Figure 1-1 shows the interaction between the various parts of this field. Human haptics aids research in machine haptics by discovering how humans use the devices for their haptic interactions. This fundamental research is done using experimental haptic interfaces, so that human haptics research also relies on advances in machine haptics. The figure also illustrates that machine haptics can be further subdivided into hardware and software design, and human haptics can be subdivided into psychophysical and physiological research. This thesis integrates software design with human psychophysical research by demonstrating how a perceptual limitation can be exploited in order to increase the effectiveness of a haptic interface.

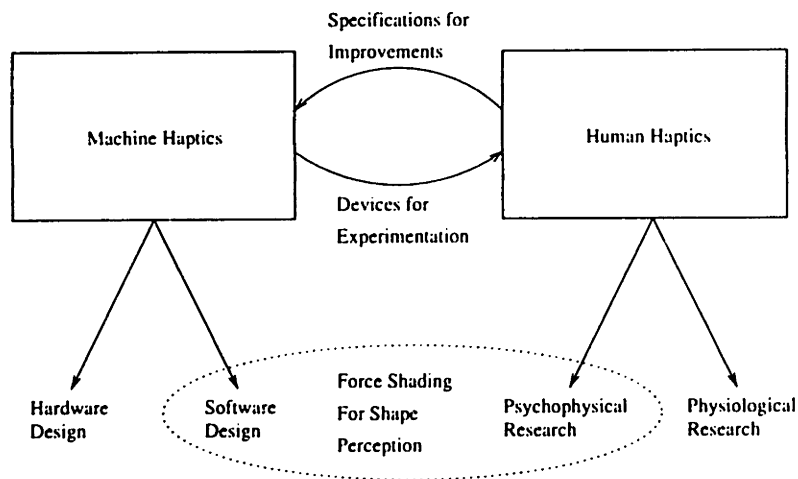


Figure 1-1: The interrelationship between the two main areas of research for building high-performance haptic VE systems.

1.2.1 Machine Haptics

Many different types of haptic interfaces have been built in the last decade. They include exoskeletons, tactile stimulators, and robotic devices with a single point of force interaction. Devices with a single point of interaction, such as the Phantom [Massie,1993] and the Sandpaper system [Minsky,1995] enable higher bandwidth through their simplicity. High bandwidth is an important feature in conveying realism with haptic interfaces [Colgate,1993; Srinivasan,1994]. More complicated interfaces, such as a force reflecting exoskeleton for several fingers, have been built. However, they are generally unable to match the fidelity of the simpler systems due to inherent compliances and errors.

The Phantom haptic interface is a three degree of freedom robot manipulator. It is exceptional in terms of its high bandwidth and high spatial and force resolution [Massie,1993]. There is a three degree of freedom gimbal at the end effector which allows the positional input to the VE to have any orientation, allowing the user to move freely. It is equipped with position encoders and motors, but no force sensors in its current design. The algorithms and experiments discussed in this thesis were designed for and implemented on a Phantom, controlled by a Pentium computer. The ideas are applicable to any Phantom-like haptic interface which has position vectors as its input and force vectors as the output that is displayed to the human.

Traditional software control of Phantom-like haptic interfaces is done using a servo loop which continually calculates the output force based on the current input position. This is possible when the contact condition may be accurately determined from the current position. For example, in simulating contact with a wall, the rigidity of the wall may be approximated by a stiff spring. Then the force displayed to the user is set to be proportional to the depth of the current input position inside the virtual wall. If the current position has penetrated the virtual wall, a reasonably accurate haptic simulation may be made by setting the initial point of contact to be the closest point on the surface of the wall. The point of contact determines the contact conditions which determine the direction of the appropriate output force. With more complex

objects there may not be a way to determine the contact conditions accurately from the current position. In Figure 1-2, there are two objects, both of which would not be well displayed by a memoryless vector-field software controller. Whenever the input position penetrates the object and gets into its center, there is no reliable way to determine where the point of contact is. It could be on any of the different sides of the object. Three such possible placements are shown in the figure. The object on the right is a polygonal approximation to the smooth object on the left. For objects composed of polygonal-meshes, a constraint-based God-object method has been developed [Zilles, 1995] for determining an appropriate point of contact based on the current input position and the history of the input positions. This system, although restricted to polygonal-meshes, allows rendering of arbitrary polygonal-mesh solid structures. It has rendered objects composed of over 700 polygons including objects as sophisticated as a miniature model of a space shuttle.




Figure 1-2: Cross section of two sample virtual objects.

Polygonal-meshes are preferable over analytical mathematical formulae for representing models of virtual objects because of their combination of generality and simplicity. The combination of richness and simplicity of the polygonal-mesh class makes it possible for a small set of computer graphics and linear algebra routines to be applicable to virtual objects of all sizes and shapes. A simplified, polygonal-mesh model of objects of every shape and size can be generated by taking samples of the

positions of points on the surface of the object.

1.2.2 Human Haptics

Human haptics research consists of physiological research and perceptual research. Physiological research includes the mechanics and mechanisms of touch. Perceptual research includes the detection, discrimination, and identification abilities of humans in tasks involving haptic stimuli.

Human tactual sensing can be broken down into two parts: tactile sensing and kinesthetic sensing. Tactile sensing refers to the information obtained through the skin from contact with an object, and kinesthetic sensing refers to a person's ability to feel the position and motion of his own arm. Tactile sensing is mediated by the responses of low threshold mechanoreceptors innervating the skin within and around the contact region. Kinesthetic sensing refers to the sense of position and motion of limbs along with the associated forces, conveyed by the sensory receptors in the skin around the joints, joint capsules, tendons, and muscles, together with the neural signals derived from motor commands [Srinivasan, 1994].

Psychohaptics research, (i.e. the psychophysics of haptics), aims to discover the perceptual response of humans to various haptic stimuli and its relationship to their motor performance. Questions in this area are about human performance in detection, discrimination, and identification tasks. Some questions aim at perceptual decomposition, for example, how do the physical parameters stiffness and damping of a virtual wall affect the human user's perception that the object is a wall [Rosenberg,1993]? Other important needed information is in human discrimination abilities, such as the measurement of the Just Noticeable Difference (JND) for various physical parameters. The JND indicates the minimum necessary difference in magnitude between two stimuli for them to appear distinct to a human.

This thesis explores a decomposition of shape perception into two components: position (the location of the hand in 3-D space) and force (the forces acting on the hand). In shape exploration, humans are exposed to both position and force information. In the real world, these two types of information are related, as the force on

the human explorer is in the direction of the surface normal at the point of contact for frictionless objects. In virtual worlds the force direction may be computed independently of the surface shape, enabling experimental investigation of the question: is shape perception based on the perception of the position path that the human's fingers take, or is it based on the shape which best matches the human's perception of the force directions that the hand feels as it travels through the environment?

1.3 Multimodal Perception in VEs

Intersensory psychophysics complicates matters further. When multiple sensory modalities provide the brain with inputs, the brain must sort them out and create a single coherent perception of the environment. In some cases, certain sensory information components may dominate over other sensory components in assessing the overall sensation. Because of this, in virtual environments, perceptual illusions can sometimes be created by intentionally providing conflicting information to multiple senses, causing the human to perceive some events differently than how they physically occur. For example, a visual-haptic illusion was observed at the Touch Lab at MIT's Research Laboratory of Electronics. When subjects were asked to judge the stiffness of a virtual spring, it appeared that the subjects calculated stiffness by taking the ratio of force to visual displacement, which subjects observed visually from a computer monitor, rather than manual displacement, which subjects observed from their kinesthetic senses in their arm. This is an example of visual dominance causing illusory perception.

Additional sensory modes in VE systems can also be used to enhance perception over single mode interactions. For example, proceeding further on the issue of decomposition of shape perception, an additional question is what happens when there is position information from vision in addition to haptic force and position information? In the experiments described here, to investigate this question, visual stimulus was sometimes presented which gave the subjects additional sensory information.

Chapter 2

Haptic Rendering Algorithms

2.1 Introduction

Haptic rendering refers to the control of a haptic interface in order to simulate interactions with a given set of virtual objects. For Phantom-like haptic interfaces, the renderer works as shown in Figure 2-1. The current position of the haptic interface endpoint is the input to the haptic renderer, and from this, along with the information about the virtual objects, and the haptic interface, the haptic renderer outputs the appropriate force values to the user's hand via the haptic interface.

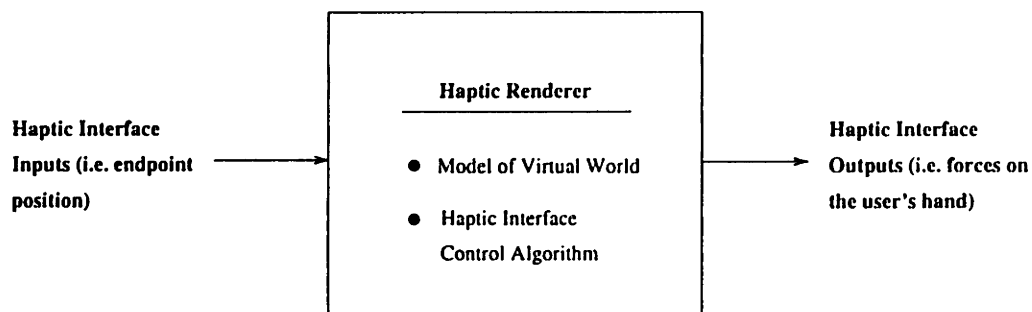


Figure 2-1: Block Diagram for Input/Output of Haptic Renderer

This chapter explains:

- Vector-field based haptic rendering algorithms.

- The constraint based God-Object haptic renderer, which enables rendering of objects composed of large numbers of polygons in a mesh.
- A height map simplification for memoryless vector-field based haptic rendering.

2.2 Vector-Field Based Rendering Algorithms

Vector-field based haptic renderers for Phantom-like haptic interfaces compute the output force of the haptic interface as a function of the input position and the state of the world. Memoryless haptic renderers use a function, \vec{f} , such that the force vector, \vec{F} , is a function of the current input position, \vec{x} , only.

$$\vec{F} = \vec{f}(\vec{x})$$

The memoryless approach is simple, and has been used to provide haptic simulations of various shapes including planar walls, boxes, spheres, cylinders, and more. For these shapes, the renderer models the rigid boundaries as stiff springs. Stiff springs at these boundaries allow a wide range of response forces with only a small range of motion.

In general, vector-field based renderers may use the history of the position as well. The constraint based God-object haptic renderer [Zilles,1995], uses past values of the input position to more accurately determine the initial point of contact between the user and the virtual objects.

The nominal geometry of a vector-field based haptic rendering algorithm refers to the geometry of the region of the input space where there is some force. The nominal geometry gives the shape of the region where there is force, but does not give any information about what the forces are in that region. It does, however, indicate that the force outside of the region is zero. It is often desirable to choose a nominal geometry which is similar to the geometry of the objects which are being represented in the virtual world. Here are two examples of memoryless vector-field based control algorithms which choose a nominal geometry that is the same as the geometry of the

objects being simulated.

2.2.1 Simulation of Thick Walls

A memoryless vector-field based haptic renderer might produce the following algorithm to control a haptic interface if the world contained a single thick rigid wall. The stiffness, k , is chosen to be as high as possible, but is limited by the inertia and spatial resolution of the device as well as the servo rate of the controlling computer. The wall has a unit normal vector of \hat{n} which points away from the wall and the wall goes through the position \vec{x}_0 .




Figure 2-2: A piece of a wall with normal vector, \hat{n} , through x_0 .

$$\vec{F}_{\text{wall}} = \begin{cases} k(\vec{x} - \vec{x}_0)\hat{n}^T\hat{n}, & \text{if } (\vec{x} - \vec{x}_0)\hat{n}^T \geq 0 \\ 0, & \text{otherwise.} \end{cases}$$

2.2.2 Simulation of Cylinders

Here is an example of a simulation of cylinder using this simple vector-field based method. In this case, the cylinder has an axis, \vec{A} , which goes through some point \vec{a}_0 , and has radius R . Again, the surface stiffness, k , is chosen to be the maximum that the system can bear.

The force computation is done by computing the minimum distance from the haptic interface input position to the central axis of the cylinder. If the endpoint is

within a distance R of the central axis, then a force needs to be applied on the user. In order to generate a force which is normal to the surface at the point of contact, the force is generated away from the center of the circle. This computation requires the minimum distance from a point to a line. The point \vec{x} in the reference frame of \vec{a}_0 is $\vec{x}_a = \vec{x} - \vec{a}_0$. The shortest vector from \vec{x}_a to the cylinder's central axis, \vec{l}_a is:

$$\vec{l}_a = \vec{x}_a - (\vec{x}_a \vec{A}^T \vec{A})$$

This vector points radially outward from the central axis of the cylinder. The vector \hat{l} is the unit length version of \vec{l} , and $\|\vec{l}_a\|$ is the magnitude of \vec{l}_a .

$$\hat{l}_a = \frac{\vec{l}_a}{\|\vec{l}_a\|}$$

Then the force is:

$$\vec{F}_{\text{cylinder}} = \begin{cases} k(\|\vec{l}_a\| - R)\hat{l}_a, & \text{if } (\|\vec{l}_a\| - R) \geq 0 \\ 0, & \text{otherwise.} \end{cases}$$

The generation of objects in this manner allows multiple objects to be placed in the same virtual world by adding up the forces from each object in the virtual world.

An important observation is that the force has the form of a product of a magnitude and a direction. These two values can be independently controlled by the haptic renderer.

2.3 Problems with Vector-Field Based Approach

The main problem with the memoryless vector-field based approach is that a decision must be made about the magnitude and direction of the force for every point in space. There is not always a clear way of determining this, as shown in Figure 1-2. In the figure there is no way to reliably decide which direction it would be appropriate for the force to point.

Even in cases when the shape is symmetric and it seems that there is an obvious

choice of what the forces should be, there can be other problems. Consider the example of a thin wall which can be felt from either side. This thin wall is very different from the thick wall discussed earlier. Now, depending on what side of the wall the user feels it from, the force needs to be in totally opposite directions. The amount that a user may penetrate a wall just to get back enough force to perceive that it is an immovable object might exceed the thickness of the wall. Clearly a system that does not make use of the history of the input position will not be able to determine appropriate forces in this case.

Another problem is illustrated in the simulation of a cube. For this simulation, it is desirable to determine which of the six faces of the cube was used to enter the cube, so that the haptic interface can push the user's hand in the correct direction. One possible method is to assume that the user entered through the face of the cube which is closest to the current input position. If the input position changes enough to change which face is closest, then there is an abrupt change of force direction, which effectively throws the user out of another side of the cube.

2.4 God-Object Idea

An extension of these basic algorithms which solves these basic problems of determining which face or wall is the active wall is called a constraint based God-object haptic rendering algorithm [Zilles,1995]. The constraint based algorithm keeps track of the position of an imaginary object, called the God-object. This object tries to follow the input position around, but is unable to penetrate walls like the actual input position can. When the input position penetrates a wall, the God-object is forced against the wall that it cannot enter. By keeping track of where the God-object is, it is easy to tell which wall the user is making contact with.

In addition to this, there are other rules which cause the God-object to update its position while the input position changes. This extension has enabled renderings of polygonal mesh structures composed of over 700 polygons as intricate as a miniature model of a space shuttle.

The constraint based God-object method is a solution for objects which are composed of polygons. But, when an object has curved surfaces, they may require a lot of polygons in order to create an illusion of the curvature. An important question is how small must these polygons be to create an illusion of curvature instead of a sequence of flat surfaces for humans? The force shading algorithm, explained in the next chapter, shows a way to decrease the number of polygons necessary to create the perception of curvature.

2.5 Height Maps

This section shows a memoryless haptic rendering algorithm for an analytical height map surface. In these surfaces, there must be a one to one correspondence between points on some region of a plane and the height map. This analysis assumes that the height map is a scalar function of two variables which is continuous and piecewise differentiable. A simple example of a height map is shown in Figure 2-3. It is not a requirement that the values always be positive.

Here is a method for rendering such a heightmap in the form of $z = f(x, y)$. The height map approximation eliminates the possibly difficult task of determining the closest point on the surface to the input position for the determination of the direction and magnitude of the force. Instead, the point of contact when the input position is at position (x_1, y_1, z_1) is assumed to be $(x_1, y_1, f(x_1, y_1))$. The height map approximation also makes it easy to compute the amount of indentation to be $z_1 - f(x_1, y_1)$.

2.5.1 Computation of Surface Normals

The output force is the product of a magnitude and a direction. The height map approximation made it easy to determine the point of contact. Here is how to calculate the direction of the surface normal at the point.

Consider first a two dimensional case. If a curve $z = f(x)$ is the nominal surface, then the slope everywhere is $z' = f'(x)$. Then the slope of the normal would be the negative reciprocal, $\frac{-1}{z'}$. But because the surface normal should always point away

from the surface, the z -component of the force is always positive. This means:

$$\vec{N}_{2D}(x) = \frac{-z'}{\sqrt{1+z'^2}}\hat{x} + \frac{1}{\sqrt{1+z'^2}}\hat{z}$$

This is the direction of the surface normal for two dimensions. For the height map $U = f(x, y)$ in three dimensions, the surface normal vector is:

$$\vec{N}_{3D}(x, y) = \frac{-\frac{\delta U(x, y)}{\delta x}\hat{x} - \frac{\delta U(x, y)}{\delta y}\hat{y} + \hat{z}}{\sqrt{\left(-\frac{\delta U(x, y)}{\delta x}\right)^2 + \left(-\frac{\delta U(x, y)}{\delta y}\right)^2 + 1}}$$

So for an input position of (x_1, y_1, z_1) , the force is:

$$Force = \begin{cases} K(U(x_1, y_1) - z_1)\vec{N}(x_1, y_1), & \text{if } |(U(x_1, y_1) - z_1)| \geq 0 \\ 0, & \text{otherwise.} \end{cases}$$




Figure 2-3: A heightmap function.

Chapter 3

Force Shading for Shape Perception

3.1 Introduction to Force Shading

Force shading refers to a controlled variation in the direction of the force vector output by the haptic renderer for the purpose of creating the illusion of a non-flat shape on a nominally flat surface. Specifically, it may be used to create the illusion of a smoothly curved shape which is actually nominally composed of a mesh of polygons. In the real world, it is difficult to create illusions like this, but VEs permit the creation of many illusions because many of the sensory inputs to the human can be independently controlled. This allows the VE designer to alter the relationship between modalities or submodalities thus facilitating the creation of illusions.

In exploration and perception of shapes in virtual environments, humans are exposed to a variety of different types of sensory inputs. The sensory inputs in shape perception are positional information, obtained from both the visual and haptic systems, and force information, obtained only from the haptic system.

Force shading is derived from the idea that during haptic exploration of a shape, in the presence of a non-normal relationship between positional and force information, the human may perceive that the shape felt was the shape that is consistent with the force information. This illusion may be strengthened even further by reinforcement

with a visual stimulus that is also consistent with the force information, so that only the positional information derived from haptic senses would be inconsistent.

3.2 The Force Shading Algorithm

The force shading algorithm presented here is for mapping any height map onto a flat polygon. It can be extended to non-height map objects if those objects may be broken up into polygons. The height map requirements are that it be a one to one function with a portion of the plane, and that the function be continuous and piecewise-differentiable. This discussion is limited to frictionless, rigid objects.

In force shading, if the haptic interface input position is at (x_1, y_1, z_1) , then the force is chosen to be in the direction of the surface normal at $(x_1, y_1, U(x_1, y_1))$, where $U(x, y)$ is the height map defined on the xy plane. The direction of this surface normal is computed as described in Section 2.5.

The force magnitude, however, is computed according to a spring law with the plane as the object surface, not the height map. The nominal geometry of this simulation, therefore, is the plane, rather than the actual shape of the height map.

So, for a height map, $U(x, y)$, defined on the xy plane, a force shading renderer would compute the following force for an input position at (x_1, y_1, z_1) :

$$F = \begin{cases} K(-z_1)\vec{N}(x_1, y_1), & \text{if } |(z_1)| \leq 0 \\ 0, & \text{otherwise.} \end{cases}$$

$$\vec{N}(x, y) = \frac{-\frac{\delta U(x, y)}{\delta x} \hat{x} - \frac{\delta U(x, y)}{\delta y} \hat{y} + \hat{z}}{\sqrt{\left(-\frac{\delta U(x, y)}{\delta x}\right)^2 + \left(-\frac{\delta U(x, y)}{\delta y}\right)^2 + 1}}$$

This is a computationally easy way to add height maps to polygons in a mesh for haptic rendering, because it doesn't add much complexity over the job that the renderer would do even if it were just the polygon mesh. This algorithm uses the same exact method for computation of the force magnitude as if the shape were indeed a flat surface, so that the only complication is that the force direction needs to be

computed according to the procedure described in Sections 2.5 and 3.2.

3.3 Application to Polygonal Mesh

One focus of this thesis is to investigate human perception of shapes rendered using these algorithms. A generic polygonal mesh renderer which does force shading has not been implemented, but, here is how these ideas could be incorporated into a more generic polygonal mesh renderer. The constraint based God-object method would be used [Zilles,1995] to choose which polygon is being contacted and where the contact point is. When the force is computed, an additional step must be taken to determine the direction of the force. If force shading is intended to be used to display a texture on the polygon, then the algorithm would look it up from a texture file. If it were being used for smoothing, then some interpolation scheme must be chosen to smooth out the normal vectors.

This is a proposed interpolation scheme: Each polygon has associated with it a normal vector, \vec{N}_i . Consider an individual triangle in the polygon mesh, with three adjacent neighboring triangles, as in Figure 3-1. The normal vector for the triangle in the figure is \vec{N}_0 , and the adjacent polygons have normal vectors \vec{N}_1 , \vec{N}_2 , and \vec{N}_3 .

$$\vec{N}_{01} = \frac{\vec{N}_0 + \vec{N}_1}{2}, \vec{N}_{02} = \frac{\vec{N}_0 + \vec{N}_2}{2}, \vec{N}_{03} = \frac{\vec{N}_0 + \vec{N}_3}{2}.$$

Then for a input position which causes the God-object to be located at \vec{x} with distances x_1 , x_2 , and x_3 from the edges of the polygon the haptic renderer should choose the force direction to be:

$$\vec{N}(\vec{x}) = \frac{x_1(\vec{N}_{02} + \vec{N}_{03}) + x_2(\vec{N}_{01} + \vec{N}_{03}) + x_3(\vec{N}_{01} + \vec{N}_{02})}{2(x_1 + x_2 + x_3)}$$

Other interpolation schemes may also be worth considering, as this is merely an example. This method only relies on some arithmetic based on position of \vec{x} and the normals of the adjacent polygons, so it is not very computationally intensive. A potential problem with this is that it only takes into account the three adjacent

polygons. Another system might instead compute a normal vector for each vertex by averaging the normals of all faces that include that vertex. Then, the the values for the normals for each of the three vertices for each triangle would be interpolated in a similar manner for the force values for the given polygon. More sophisticated interpolation methods might use a three dimensional bandlimited interpolator.




Figure 3-1: Method for Applying Force Shading Algorithm to Smooth a Polygonal Mesh

Chapter 4

The Bump Display Algorithms

4.1 Introduction

This describes seven different vector-field based haptic display algorithms for representing frictionless cylindrical bumps protruding from a flat plane, as shown in Figure 4-1. The figure shows a visual representation of the virtual object which these haptic algorithms are meant to give the perception of. The seven different methods are:

1. Cylindrical bump.
2. Single polygon approximation without force shading.
3. Single polygon approximation with force shading.
4. Two polygon approximation without force shading.
5. Two polygon approximation with force shading.
6. Three polygon approximation without force shading.
7. Three polygon approximation with force shading.

All seven of these methods are intended to convey the perception of the environment shown in the figure, but they do not all have the same nominal geometry.

The nominal geometry of a vector-field based haptic display algorithm is the virtual boundary composed of the locus of all positions of the input such that the magnitude of the reflected force changes from zero to a nonzero value. The four different nominal geometries used for these bumps are shown side by side in Figure 4-2.




Figure 4-1: Cross Section of a Cylindrical Bump on a Flat Plane




Figure 4-2: Cross section of the four different nominal geometries used in the bump display algorithms.

These bumps, which form the basis for the two experiments described in this thesis, were chosen because they compare an analytical form for a smooth shape with various polygonal approximations. In practice, complex shapes are approximated with polygons in order to simulate them in a virtual environment. In addition to the cylindrical geometry and the three polygonal approximations, there are three force shaded polygonal approximations. These force shaded bumps have the same nominal geometry as their unshaded counterparts, but have a different distribution of force

directions which is intended to convey the perception that the shapes are smooth (i.e. surface slopes are continuous), like the analytical solution.

As more polygons are used to approximate a given object, the nominal geometry better approximates the actual geometry that is being simulated (the cylindrical bump), but not without expense—larger numbers of polygons require more memory and more computation, and can result in significant time delays in rendering. In a VE, it is therefore desirable to use the minimum number of polygons that is satisfactory to the human user.

4.2 The Geometry

The following conventions will be used throughout the explanations of the various bump display methods.

The coordinate system has the x -axis increasing to the right, the y -axis increasing into the page, and the z -axis increasing upward. For simplicity, the cylinder axes are all aligned parallel to the y axis and on the yz plane, and they protrude from the xz plane. Figure 4-3 shows a labelled diagram of a cross section of a cylinder partially exposed through a planar floor. The unit normal pointing away from the plane is \hat{z} and the plane goes through the origin. The width of the region where the bump is exposed on the plane is w . The radius of the bump is R , making the curvature of the bump $1/R$. The axis of the cylinder, on the yz plane, is located at a depth of d beneath the plane. The axis depth, d , is equal to $\sqrt{R^2 - (\frac{w}{2})^2}$. The input position of the haptic interface is denoted as \vec{x} , which is equal to $\vec{x} = x\hat{x} + y\hat{y} + z\hat{z}$.




Figure 4-3: Labels for the Geometry for the various Haptic Simulations

4.3 The Cylindrical Bump

The display algorithm for the cylindrical bump is as follows. Compute the axis depth, d , from R and w . Then the output force is the sum of two components:

$$F_{\text{total}} = F_{\text{cylinder}} + F_{\text{wall}}$$

These are the same F_{cylinder} and F_{wall} which were listed in Chapter 2.

4.4 The Unshaded Polygonal Approximations

These three bump display algorithms are based on the height map approach given in Section 2.5. All of these forces are a product of a force magnitude and a force direction. The force magnitude is given by the stiffness constant, representing the stiffness of the object's surface, multiplied by the indentation, and the direction is given by the surface normal.

$$H(x, y) = \text{Height Map}$$

$$M(x, y, z) = \text{Force Magnitude}$$

$$\vec{D}(x, y) = \text{Force Direction}$$

$$\vec{F}_{\text{total}} = (\text{Force Magnitude}) \times (\text{Force Direction})$$

$$\vec{F}_{\text{total}} = M\vec{D}$$

4.4.1 Single Polygon Without Force Shading

The nominal geometry of this single polygon approximation to the cylindrical bump does not have enough polygons to add any shape to the plane over which the bump is superimposed. This approximation, without any shading, is exactly the same as a flat wall, and is clearly inadequate for representing the intended cylindrical bump. Following are expressions for the height map, force magnitude, and force direction for

this simulation.

$$H_{1p}(x, y) = 0,$$

$$\vec{D}_{1p}(x, y) = \hat{z}$$

$$M_{1p}(x, y, z) = \begin{cases} K(H_{1p}(x, y) - z), & \text{if } H_{1p}(x, y) - z \geq 0 \\ 0, & \text{otherwise.} \end{cases}$$

4.4.2 Two Polygons Without Force Shading

Here, the indentation is calculated from a new nominal geometry. The two polygon approximation is made by sampling the cylinder at three points, one at the peak and two at the sides. The height of the bump is exactly the same as the circular cylinder bump since one of the samples is at the peak. The height is equal to $R - \sqrt{R^2 - (\frac{w}{2})^2}$. The height of the bump as a function of x (shown in Figure 4-4) is:

$$H_{2p}(x, y) = \begin{cases} 0, & \text{if } |x| > \frac{w}{2} \\ h(1 - \frac{2|x|}{w}), & \text{if } |x| \leq \frac{w}{2} \end{cases}$$

And the surface normal, as a function of x is,

$$\vec{D}_{2p}(x, y) = \begin{cases} \frac{(\frac{2h}{w} \text{sgn}(x)\hat{x} + \hat{z})}{\sqrt{1 + \frac{4h^2}{w^2}}}, & \text{if } |x| \leq \frac{w}{2} \\ \hat{z}, & \text{if } |x| > \frac{w}{2} \end{cases}$$

The magnitude of the force is proportional to the indentation into the height map.

$$M_{2p}(x, y) = \begin{cases} K(H_{2p}(x, y) - z), & \text{if } H_{2p}(x, y) - z \geq 0 \\ 0, & \text{otherwise} \end{cases}$$




Figure 4-5: Geometry for a Three-Polygon Approximation to the Cylindrical Bump

$$D_{3p}(x, y) = \begin{cases} \hat{z}, & \text{if } |x| \leq \frac{w'}{2} \\ \frac{1}{\sqrt{1 + \frac{4h'^2}{(w-w')^2}}} \left(\frac{2h'}{w-w'} \text{sgn}(x) \hat{x} + \hat{z} \right), & \text{if } \frac{w'}{2} < |x| \leq \frac{w}{2} \\ \hat{z}, & \text{if } |x| > \frac{w}{2} \end{cases}$$

The corresponding magnitude of the force is

$$M_{3p}(x, y, z) = \begin{cases} K(H_{3p}(x, y) - z), & \text{if } H_{3p}(x, y) - z > 0 \\ 0, & \text{otherwise.} \end{cases}$$

4.5 The Shaded Polygonal Approximations

Here the forces are computed for the polygonal approximations with force shading. The nominal surface and magnitudes of the forces are computed exactly the same as the unshaded case. But now, the force direction is computed to match the surface normal from the circular cylinder.

$$D_{shaded}(x, y) = \begin{cases} \hat{z}, & \text{if } |x| > \frac{w}{2} \\ \frac{x}{R} \hat{x} + \frac{\sqrt{R^2 - x^2}}{R} \hat{z}, & \text{if } |x| \leq \frac{w}{2} \end{cases}$$

4.5.1 Single Polygon With Force Shading

$$F_{1p\text{-shaded}}(x, y, z) = M_{1p}(x, y, z)\vec{D}_{\text{shaded}}(x, y)$$

4.5.2 Two Polygons With Force Shading

$$F_{2p\text{-shaded}}(x, y, z) = M_{2p}(x, y, z)\vec{D}_{\text{shaded}}(x, y)$$

4.5.3 Three Polygons With Force Shading

$$F_{3p\text{-shaded}}(x, y, z) = M_{3p}(x, y, z)\vec{D}_{\text{shaded}}(x, y)$$

4.6 Comments

In the section above it may appear that the polygonal approximations are more complicated to compute than the cylindrical bump, whose display they are supposed to simplify. This is because the cylindrical bump has a simple analytic form that is regular and smooth. Had the initial form not been analytic, or had the shape been more complicated, or even unknown, then polygonal approximations are the best choice. The force shading algorithms described above are applicable for a generic rendering system where typical objects are represented in terms of polygons intended to simulate surfaces either exactly (e.g. a cube) or approximately (e.g. a sphere).

Chapter 5

Experimental Design

5.1 Experimental Goals

These experiments were designed to find out about how humans explore convex shapes and curvatures, especially to discover how well the force shading algorithm can create the illusion of smooth shapes. Some general questions which these experiments were designed to answer are:

- How much haptic shape perception arises from position sensing versus force sensing?
- How does vision affect this perception? For example, in the absence of vision does the brain use the haptic position sensors more than in the presence of vision?
- How effective is force shading at displaying convex shapes? How does this vary with the radius of the shape?

5.2 The Theta-critical Hypothesis

Intuitively, it seems that force shading will work best for height maps where the height map has negligible height, and will perform most poorly for very large height maps. Certainly when the height map includes height variations that are much larger than

the human JND for position sensing, it is likely that force shading will be inadequate because of the lack of sensation of really moving the hand up and down.

Figure 5-1 shows three different cylinders protruding from a plane. In each, there is an angle labelled θ which represents the angle subtended by the exposed portion of the cylinder. The theta-critical hypothesis suggests that there is a maximum angle that can be subtended by the exposed portion of a cylinder which can be convincingly simulated by a force shaded nominally flat surface. This maximum angle is called $\theta_{critical}$.




Figure 5-1: Profiles of three different constant-curvature bumps with the same width.

The centers of the cylinders are drawn to show θ , the angle subtended by the exposed portion of the cylinder. The rightmost bump, then, has a θ of 180° , whereas the leftmost bump has a θ of 70° .

Here we propose the hypothesis that there is some $\theta_{critical}$ which is the threshold below which there is no need to break up a surface into more shaded polygons, since the human is unable to distinguish between the shape of a force shaded flat nominal surface and a smoothly curved nominal surface. More psychophysical experiments are needed to determine $\theta_{critical}$ and its dependence on various display parameters.

When large surfaces are approximated by breaking them up into small sections, each of which is approximated by a force shaded polygon, this angle, $\theta_{critical}$ can be used to determine how big the individual polygons need to be. Each polygon needs to be chosen so that if it were to be approximated by a constant curvature region

then the total angle through the polygon must be less than $\theta_{critical}$.

5.3 Experimental Paradigms

Two experiments were designed to help gain insight into these questions. They involve using two different psychophysical experimental paradigms, a *matching paradigm* and a *classification paradigm*. In the matching paradigm, the subjects were given control over one “matching” stimulus and asked to modify it until it felt most like the “reference” stimulus. When they thought that they had matched the two stimuli as well as they could, they were asked to give a satisfaction index to indicate how alike the two were. In the classification paradigm, the subjects were given a stimulus and told to decide what type it was from among a list of choices. They were forced to answer from among the choices.

5.4 The Experimental Setup




Figure 5-2: The Experimental Setup

The experimental setup. The Phantom haptic interface (on the right) was controlled by a Pentium computer. The monitor provides subjects with the visual stimuli.

Both experiments were done using the Phantom haptic interface which was con-

trolled by a Pentium based computer running DOS. Figure 5-2 shows a photograph of the setup. The subjects put their hands through an opaque cloth to grasp the stylus of the Phantom in order to prevent them from looking at their hands during the experiments. The computer programs for the control of the experiments were all written in C++ and compiled using the Borland C++ 4.0 Compiler.

5.5 The Matching Experiments

The matching experiment was designed to discover how suitable the force shaded polygonal approximations to the real bump are at substituting the real bump. In this experiment, subjects were asked to change the size of a “matching” bump to make it feel as much like a “reference” bump as was possible. The subjects were given control over the radius value of the “matching” bump. Several parameters were varied so that their effects could be studied. These parameters are: the presence or absence of a visual stimulus, the algorithm used to generate the “matching” bump, the size of the “reference” bump, and the initial size of the “matching” bump.

Table 5.1 enumerates the different experimental conditions. A total of 2400 trials were done, using 5 different subjects. The duration of each trial was about a minute and a half, so that there were 60 total hours of experiments distributed over the 5 subjects.

Vision	Bump Type	Subject	Reference Radius	Initial Matching Radius	Repeated Trials
<ul style="list-style-type: none"> • Yes • No 	<ul style="list-style-type: none"> • Cylindrical • One-Poly • Two-Poly • Three-Poly 	<ul style="list-style-type: none"> • 1 • 2 • 3 • 4 • 5 	<ul style="list-style-type: none"> • 600 mils • 800 mils • 1000 mils • 1200 mils • 1400 mils 	<ul style="list-style-type: none"> • 1400 mils • Same • 600 mils 	<ul style="list-style-type: none"> • 4

Table 5.1: Matching Experiment Conditions

Presence of Visual Stimulus. The force shading illusion is strengthened by coupling it with a visual display of the virtual world and the haptic interface endpoint. In order to strengthen the illusion, the visual stimulus was modified to always show the endpoint as though it had not penetrated the virtual objects. Furthermore, the two bumps were always shown to be the same size. If the visual bump changed size during the matching task it would be too easy for the subjects to rely on the visuals and to ignore haptics. By comparing the subject's performance with and without the visual information, it was possible to see how much more effective the illusion was when vision was added.



Figure 5-3: The visuals for the matching experiment

The Bump Display Algorithms. Four different bump display algorithms were used. In all cases, the left bump, which was the reference bump, was implemented using the algorithm for the circular cylinder given in Chapter 4. The four different displays for the right bump were: the circular cylinder bump, the flat (one polygon) bump with shading, the two polygon bump with shading, and the three polygon bump with shading. By comparing the subject's abilities to perform the matching

task for the different bumps, it was possible to determine how adequate the polygonal approximations with shading were at relating curvature information. The matching task with the real bump provided an important control to estimate the resolution of the subjects independent of the effects of force shading. The bump display algorithms are all explained in Chapter 4.

Reference Bump Sizes. The bumps were always 1000 mils (one inch) wide. Five different reference radii were used: 600, 800, 1000, 1200, and 1400 mils. The heights of the different bumps as a function of radii are shown in Figure 5-4. Also, the amount of subtended angle for each of the polygons in the various approximations to the cylinders is shown in Figure 5-5.

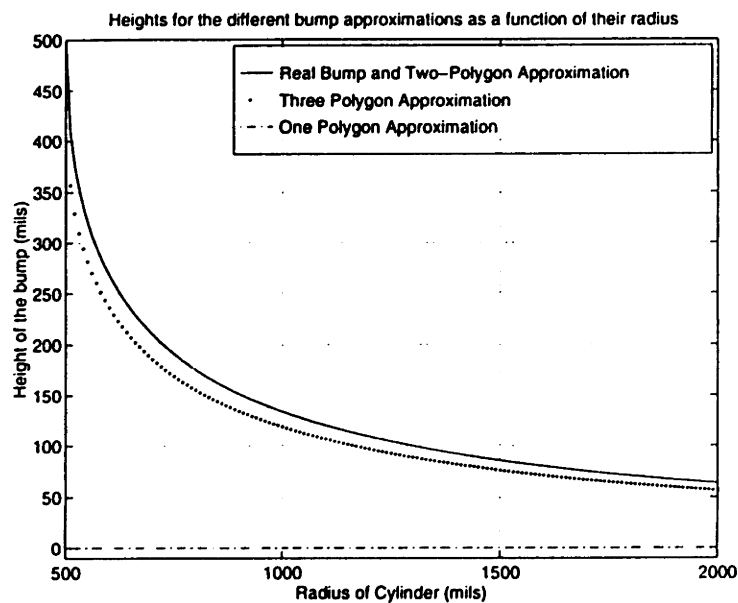


Figure 5-4: The heights of the bumps as a function of radius

Initial Values for Matching Bump Sizes. In order to avoid biasing the results, it was important to choose various initial values for the matching bump. If the subjects did not explore the full range of radii for the matching bump, then the data might appear to have properties that would actually just be artifacts of the initial values. For this reason, three different initial values were chosen for each reference bump.

One value was fixed at a large radius, one value fixed at a small radius, and one value at the same value as the reference bump. The large radius was the largest radius used as a reference bump in the experiment, 1400 mils. The small radius was the smallest used as a reference in the experiment, 600 mils. Figure 5-6 shows the initial layout of the matching values. The small initial radius points are marked with an \circ , the large initial radius points are marked with a $+$, and the initial radius points equal to the reference radius are marked with an \times . On the final data plots, the same convention is used so that it can easily be seen whether the final matching radii were heavily biased by the initial radii. The pairs of reference and matching initial values were presented to the subjects in a random order.

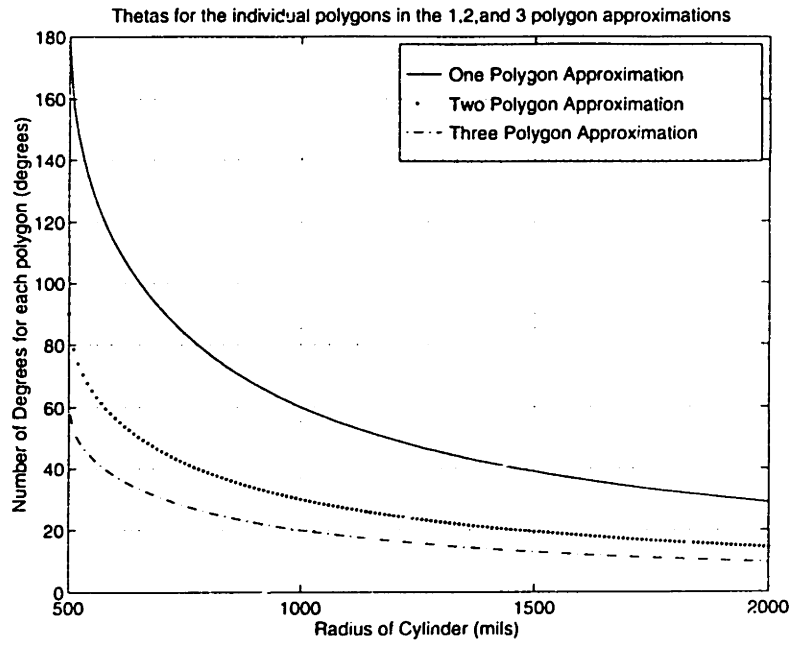


Figure 5-5: The amount of angle subtended by the cylinders divided by the number of polygons used to approximate the cylinder.

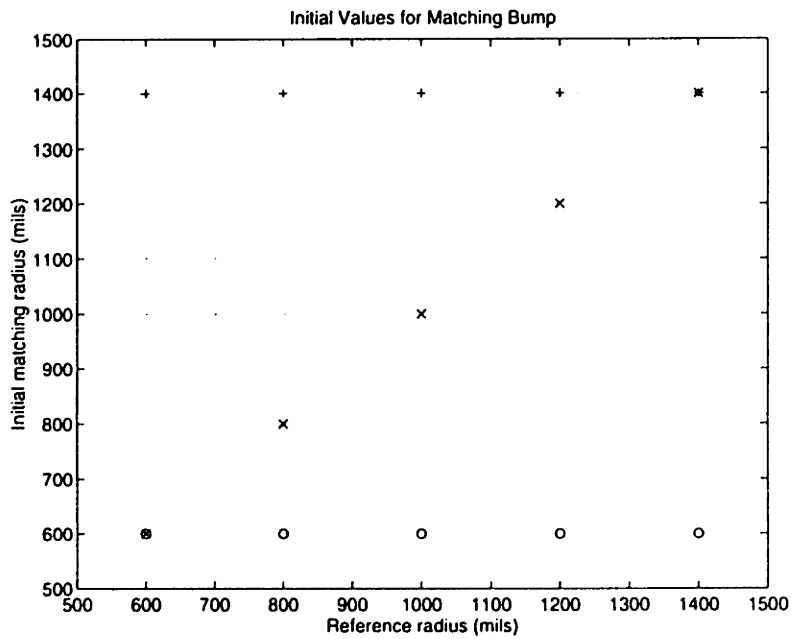


Figure 5-6: The initial values of the reference and matching bumps.

5.5.1 Instructions to Subjects

The subjects were presented with the instructions shown in Figure 5-7

Instructions for Shape Matching Experiment

You will be presented with a series of pairs of bumps protruding from a flat plane. For each pair, your goal is to make the two bumps feel as alike as possible. The left bump is the reference bump- you cannot change its size. The right bump is the matching bump, which you can make bigger or smaller. To make the bump smaller, hit the 'j' key on the keyboard. To make the bump bigger, hit the 'k' key on the keyboard. Continue making the bump bigger or smaller until you think that you have made the two bumps as alike as possible. When you are sure that you have done the best that you can, hit the 'l' key to enter in your Satisfaction Index and to then move on. An explanation of the Satisfaction Index is given below.

Explanation of Satisfaction Index

The purpose of the Satisfaction Index is to relate how closely you feel that the two bumps matched. You are asked first to try your best to make the two bumps as alike as possible, but then indicate how close you think they are.

A value of 10 should be used to indicate examples where you do not think there is any difference at all between the two bumps. A value of 1 should be used when you wish to indicate that the two bumps felt completely different, even after you did your best at matching the bumps. Finally, the values in between 1 and 10 should be used to indicate the cases in between these, where the higher values should be used for cases where the bumps feel very similar, and lower values should be used for cases where the bumps feel only slightly similar.

Figure 5-7: The instructions given to the subjects before the matching experiment.

5.5.2 Exit Interviews

Upon finishing this set of experiments, the subjects were asked to fill out a questionnaire. The questions were aimed at finding out what the subjects consciously thought about the shapes of the bumps, and whether or not they had any particular strategies for matching the bumps. The following were the questions on it:

1. Describe the strategy you used to match the way that the bumps felt. Try to describe the way in which you felt the bumps. (For example, did you stroke both bumps in one pass left to right, or did you go over the left bump a few times and then the right bump a few times?)
2. What was the most difficult part of matching the way that the two different bumps felt?
3. Draw the shapes of the bumps that you can remember feeling in some part of the experiment.
4. Describe the strategy you used to give values for satisfaction index. What caused you to give low ratings? What caused you to give high ratings?
5. Do you have any other general comments or impressions that you would like to relate to the experimenter? Feel free to use the space below.

5.5.3 History Recording

Insight is often gained about subject's responses and about their methods and strategies by recording their process. To facilitate this, in addition to recording the subjects final values for the matched radii and for the satisfaction indices, the following information was recorded.

Position Profile. As the subjects explored the virtual environment, their positions were recorded into a data file. This data can be used to recreate the paths the subjects

Bump Type	Subject	Reference Radius	Repeated Trials
<ul style="list-style-type: none"> • Cylindrical • One-Poly • One-Poly • Two-Poly • Two-Poly • Three-Poly • Three-Poly 	<ul style="list-style-type: none"> • 1 • 2 • 3 • 4 • 5 	<ul style="list-style-type: none"> • 600 mils • 1000 mils • 1400 mils 	<ul style="list-style-type: none"> • 24

Table 5.2: Classification Experiment Conditions

took in the virtual environment. It can also be used to calculate the forces which were actually exerted on the subjects.

Radius Matching Profile. As the subjects made changes to the radius of the “matching” bumps, these changes were recorded. This data can be used to find out how wide a range of radius values were considered by the subject before they decided on their final answer.

5.6 The Classification Experiment

The classification experiment was designed to explicitly compare the various shaded bumps with their unshaded counterparts. The matching experiment did not address the question of whether the shaded bumps were any better than the unshaded approximations. In this experiment, subjects were asked to feel a single bump and then classify it as one of four which they saw on the screen. The four pictures on the screen (shown in Figure 5-8) depicted: a flat plane, a two-polygon bump, a three-polygon bump, and a cylindrical bump. This experiment had two different parameters to vary: the bump display algorithm, and the size of the bump. The experimental conditions are enumerated in Table 5.2. A total of 2520 trials were done, using 5 subjects. These trials took only about 10 seconds each, so that the whole experiment took about 7 hours to complete.

The Bump Display Algorithms. All seven of the bump display algorithms listed in Chapter 4 were used. Only one bump was presented at a time in this experiment.

The Bump Sizes. The bumps were all 1000 mils (one inch) wide. Three different values for the radius of the bumps were chosen (600, 1000, and 1400 mils). This was to find out whether there were differences in the effectiveness of the force shading algorithm for the different radii of the bumps.




Figure 5-8: The visuals for the classification experiment

5.6.1 Instructions to Subjects

The subjects were presented with the instructions shown in Figure 5-9.

Instructions for Shape Classification Experiment

You will be presented with a series of single bumps protruding from a flat plane. For each bump, feel the bump and decide which of the four shapes that you see on the screen is most like the bump that you just felt.

Do not focus on the sizes of the bumps on the screen, but instead focus on the distinguishing characteristics of the different shapes.

When you decide which of the four bumps on the screen is the one which you have felt, type in the number of the bump ('1', '2', '3', or '4') and then press return for the next bump.

Figure 5-9: The instructions given to the subjects before the classification experiment.

Chapter 6

Results

The experiments explained in Chapter 5 were all performed on five subjects. These subjects were all right-handed MIT undergraduates studying engineering. They were paid for their time. The five subjects are called TP, IU, AL, IG, and KN.

6.1 The Matching Experiment

The matching experiment used four different haptic rendering methods to display the “matching” bump. Here we discuss the results of each of these different parts of the experiment in successive sections. They are:

1. Cylindrical bump matching.
2. Flat bump matching.
3. Two-polygon matching.
4. Three-polygon matching.

For each of these four different experiments, the results are listed separately for the trials with the visual stimulus present and for those without the visual stimulus.

6.1.1 Cylindrical Bump Matching

Tables 6.1 and 6.2 summarize the results of the cylindrical bump matching experiment. Table 6.1 lists the results for the trials without the visual stimulus, and Table 6.2 lists the results for the trials with the visual stimulus. The data is summarized for each subject individually, and for the five subjects combined. The following data is presented in this table:

- **Slope.** The slope given in the table is that of the best-fit line in the plots of the “matching” radius values chosen by the subjects versus the “reference” radius values. This may be clearly seen in Figures 6-2 and 6-3. As in Section 5.5, the data points on these plots are marked with o’s, x’s, and +’s in order to indicate whether the initial value of the matching radius was 600 mils, equal to the radius of the reference bump, or 1400 mils, respectively.
- **Mean SD (%).** In this column, the average standard deviation of the subjects responses is given as a percentage of the mean matching radius. This was computed by first finding the standard deviation of all of the “matching” radii for each individual “reference” radius, and then converting that to a percentage of the mean value of the “matching” radius. The average of these five values, which are listed individually on the scatter plots of the matching and reference radius values as “varper”, is the Mean SD in percent.
- **Mean SI.** The mean Satisfaction Index (SI) is computed by taking the average of all of the satisfaction index responses for all trials with the given initial conditions. This average value does not give information about variations in their SI over different radii for a particular experimental condition, but instead gives an overall measurement of the subjects’ SI for that type of bump display. The values are also plotted in Figures 6-4 and 6-5 as a function of the reference radius value. These plots give an indication of the variation in responses over repeated trials and across different values of the “reference” radius.

- **Mean Number of Steps.** The mean number of steps is computed using the recorded data on how the subject arrived at the matching radius. This number indicates the average number of changes that the subject made to the matching bump prior to deciding on the final value that was the best match. Some examples of paths that subjects took as they explored different values of the radius of matching bumps are given in Figure 6-6. Each time the subject changed the radius size, an event occurs and there is another point in the trace. When the subject decided that the best match had been attained the trace ended.

In addition to the tables and figures mentioned above, a typical position trace is given in Figure 6-1. This position trace shows the path that subject AL took in one part of one trial when both the reference and matching bumps had a radius of 1000 mils. During this trial, the visual stimulus was not present.

Subject	Slope	Mean SD (%)	Mean SI	Mean # of Steps
TP	0.85	13.0	7.7	15
IU	1.04	13.3	8.8	37
AL	0.99	6.7	8.1	22
IG	1.03	15.3	8.8	21
KN	0.92	7.9	7.2	64
Overall	0.97	12.3	8.1	32

Table 6.1: Summary of Matching Results for the Five Subjects for Cylindrical Bump Without Vision

Subject	Slope	Mean SD (%)	Mean SI	Mean # of Steps
TP	1.01	13.2	7.8	16
IU	0.88	14.0	8.6	40
AL	0.99	9.1	8.2	19
IG	0.95	14.7	8.6	17
KN	0.88	8.1	8.3	51
Overall	0.94	12.9	8.3	29

Table 6.2: Summary of Matching Results for the Five Subjects for Cylindrical Bump With Vision

If the subjects were ideal, then for this circular cylinder matching experiment, they would have attained a slope of 1. Some subjects achieved slopes very close to 1 in both experiments, with and without vision. The SI values, around 8 on a scale of 1 to 10, are indicative of the highest values that these subjects were willing to give, since the two bumps were implemented in exactly the same way and exact matches were always possible. The Mean SD of two subjects (AL and KN) stood out as being much lower than that of the other three subjects. These two subjects were the most consistent from trial to trial in all the experiments.




Figure 6-1: Position profile for the cylindrical bump matching task.

This trace shows the position profile for a short amount of time while subject AL felt two cylindrical bumps with radii of 1000 mils. The solid line shows the nominal geometry of the virtual objects. The dots show the location of the endpoint of the Phantom at different moments in time. The endpoint is inside of the nominal surface because in order for the user to feel any force according to the spring law, there must be some finite indentation. This trace shows a relatively constant indentation which suggests that the force magnitude, which is proportional to the indentation, was roughly constant.

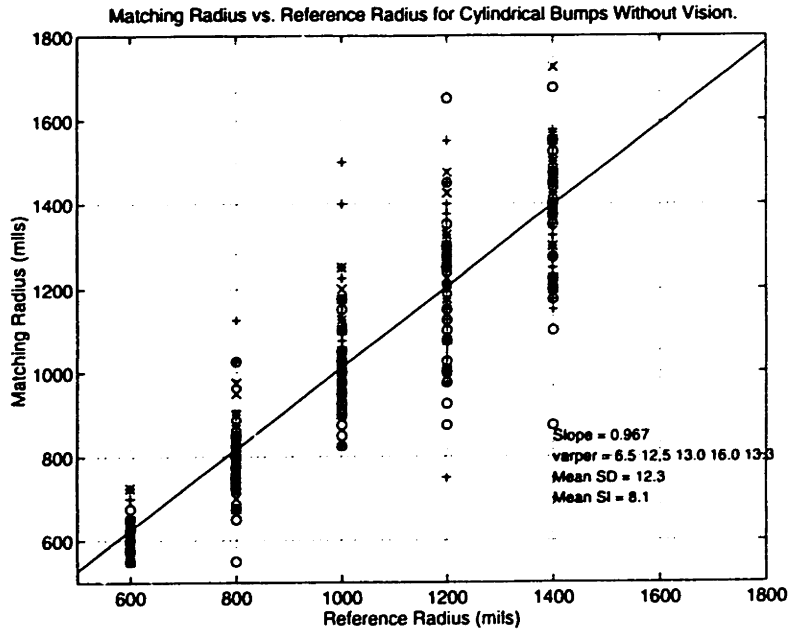


Figure 6-2: Results for Cylinder, No Vision

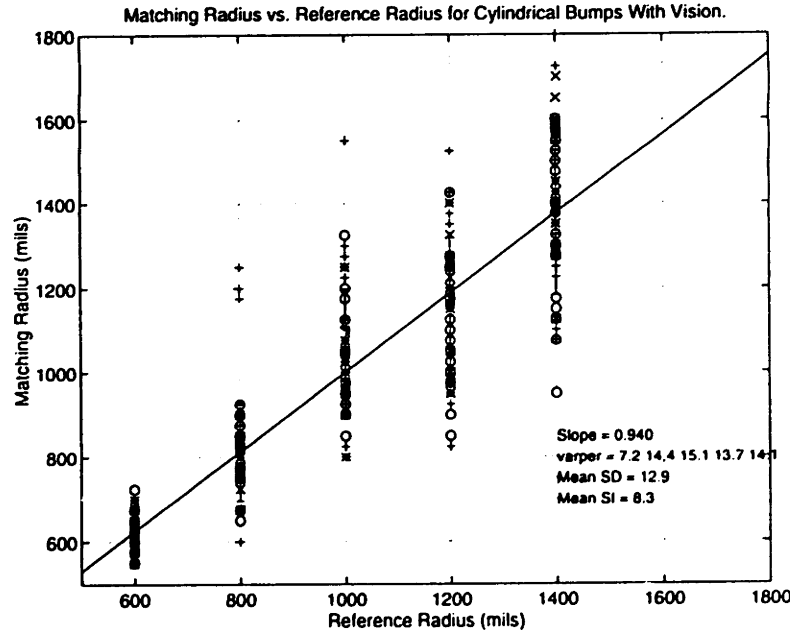


Figure 6-3: Results for Cylinder, With Vision




Figure 6-4: Satisfaction Index Results for Real No Vision




Figure 6-5: Satisfaction Index Results for Real With Vision




Figure 6-6: Radius profiles for the cylindrical bump matching task.

These three examples of the paths taken by three different subjects in matching a cylinder of 1000 mils radius show three generally distinct behavioral patterns. KN tended to change radius values a lot, and go above and below his final answer several times before finally stopping. AL made it very accurately to the final values, and would generally go past his final value and then go back. Another type of behavior is shown in the plots at the top of the figure. IG tended to bring the matching bump toward the final value, but there was a lot of variability in his final value. Perhaps this subject's sensitivity to fine differences was not as good as AL's, or perhaps he wasn't working hard enough, like KN. AL seemed to have the most natural ability in that his variability was very low and his accuracy was very high, and he also used a relatively few number of steps to get to his final value.

6.1.2 Flat Bump Matching

Tables 6.3 and 6.4 show the data for the matching experiments where the “matching” bump was implemented using the force shaded single polygon approach. The format for the data is the same as explained in the previous section. In these trials, the “matching” bump had a nominally flat surface, and the subjects did not have any control over the nominal surface, but instead could only change the force directions. Figure 6-7 shows a typical position trace for part of one trial, which indeed shows that the subject did move along the nominally flat surface even when force shading was present. Based on force directions alone, with no haptic positional cues, the subjects were able to correlate larger “matching” radius values for larger “reference” radius values. This can be seen by the positive slopes of the best-fit lines drawn in Figures 6-8 and 6-9. The subjects’ responses showed a high degree of variability. The subjects who did not exhibit much variability tended to choose very low values for the “matching” radius independent of the value of the “reference” radius. The SI values are rather low, especially for the sessions without vision. The SI plots (shown in Figures 6-10 and 6-11) show that the subjects were especially dissatisfied with their matches when the “reference” radius was low.

Subject	Slope	Mean SD (%)	Mean SI	Mean # of Steps
TP	0.65	11.6	5.6	27
IU	0.72	11.7	7.3	35
AL	0.03	1.6	3.4	23
IG	0.24	10.2	4.7	28
KN	0.51	6.9	5.2	72
Overall	0.43	14.4	5.2	37

Table 6.3: Summary of Matching Results for the Five Subjects for Flat Bump Without Vision

Subject	Slope	Mean SD (%)	Mean SI	Mean # of Steps
TP	0.73	15.4	7.4	23
IU	0.70	19.5	8.5	39
AL	0.34	9.7	6.7	20
IG	0.54	18.2	7.8	17
KN	0.53	8.2	6.8	51
Overall	0.57	20.0	7.5	30

Table 6.4: Summary of Matching Results for the Five Subjects for Flat Bump With Vision




Figure 6-7: Position profile for the flat bump matching task.




Figure 6-8: Results for Flat Bump Matching, No Vision

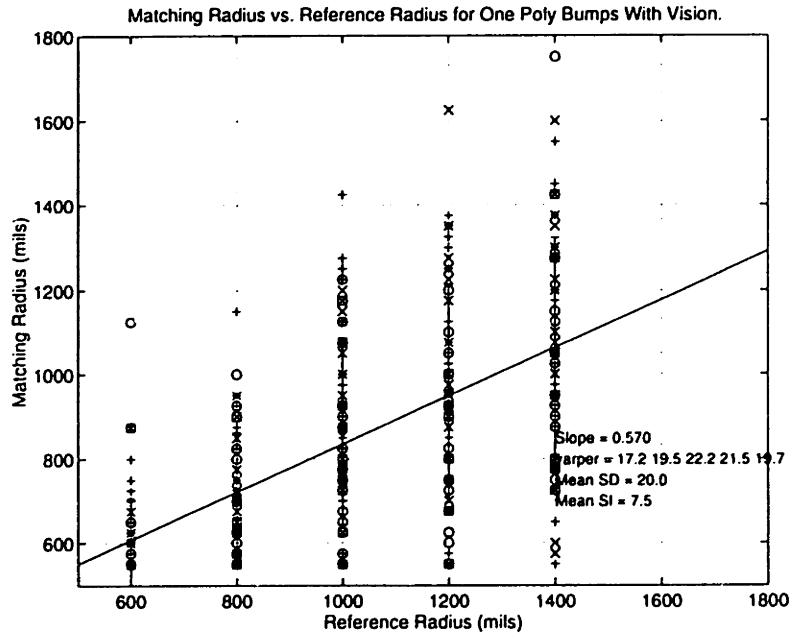


Figure 6-9: Results for Flat Bump Matching, With Vision

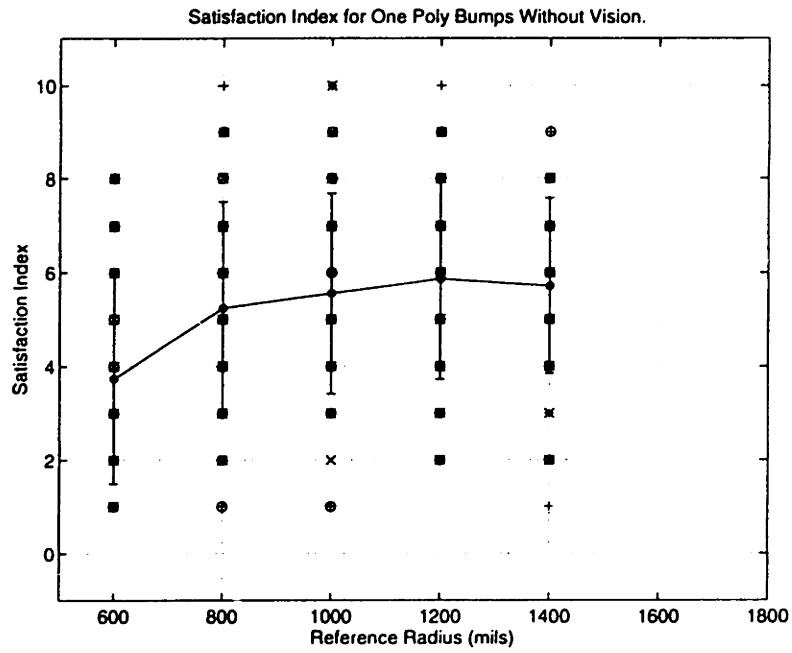


Figure 6-10: Satisfaction Index Results for Flat Bump Matching, No Vision

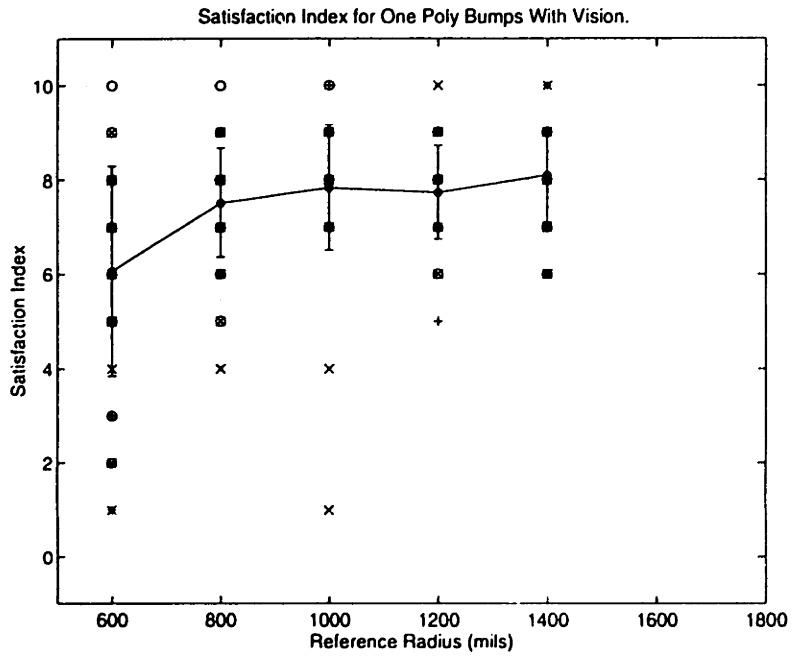


Figure 6-11: Satisfaction Index Results for Flat Bump Matching, With Vision

6.1.3 Two Polygon Matching

Tables 6.5 and 6.5 show the summarized results for the matching experiments where the “matching” bump was implemented using the force shaded two polygon approach. The overall slopes were less than 1, indicating that they chose larger “matching” bumps than the exact two polygon approximations to the reference cylinder. The SI values were higher for the sets with vision. Figure 6-12 shows a position trace, which verifies that in these trials, the subjects did actually trace a path close to the nominal geometry of the bumps.

Subject	Slope	Mean SD (%)	Mean SI	Mean # of Steps
TP	0.77	10.3	7.1	19
IU	0.79	10.4	8.1	35
AL	0.81	6.5	7.6	28
IG	0.41	15.7	7.0	20
KN	0.60	5.4	7.1	67
Overall	0.67	11.7	7.4	34

Table 6.5: Summary of Matching Results for the Five Subjects for Two Polygon Bump Without Vision

Subject	Slope	Mean SD (%)	Mean SI	Mean # of Steps
TP	0.84	14.3	7.4	16
IU	0.72	15.3	8.6	36
AL	0.80	14.3	7.7	25
IG	0.57	11.1	8.4	19
KN	0.64	7.2	8.4	59
Overall	0.71	14.7	8.1	31

Table 6.6: Summary of Matching Results for the Five Subjects for Two Polygon Bump With Vision




Figure 6-12: Position profile for the two polygon bump matching task.




Figure 6-13: Results for Two Polygon Bump Matching, No Vision

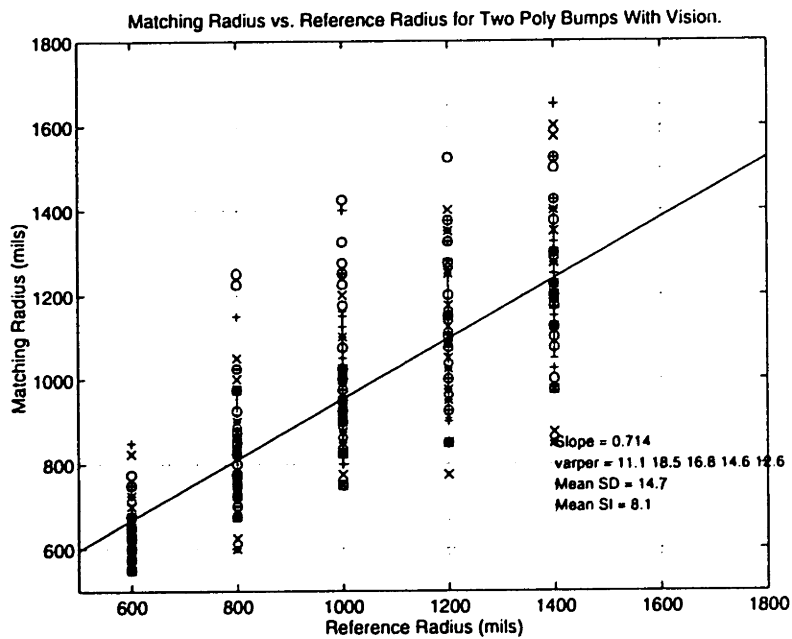


Figure 6-14: Results for Two Polygon Bump Matching, With Vision

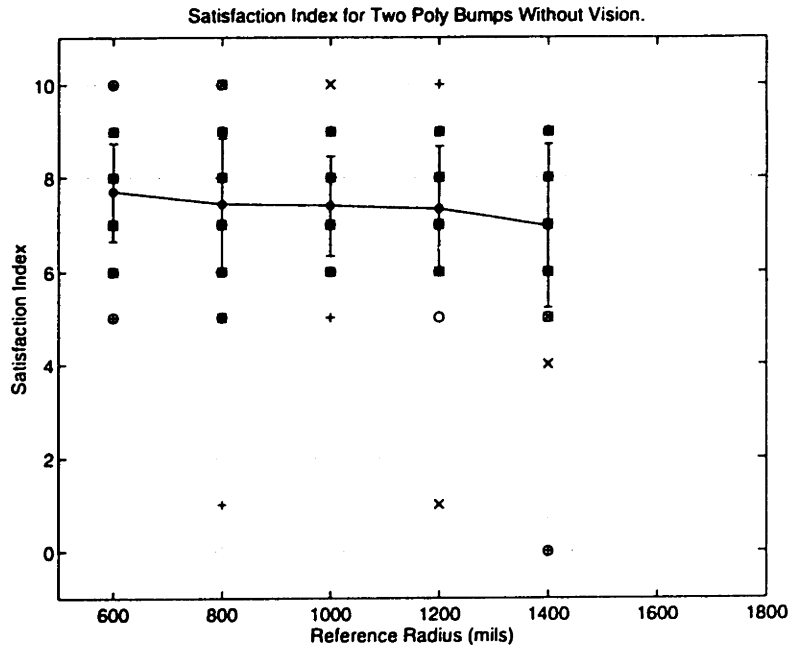


Figure 6-15: Satisfaction Index Results for Two Polygon Bump Matching. No Vision




Figure 6-16: Satisfaction Index Results for Two Polygon Bump Matching. With Vision

6.1.4 Three Polygon Matching

Tables 6.7 and 6.7 show the summarized results for the matching experiments where the “matching” bump was implemented using the force shaded three polygon approach. The slopes are both less than 1, indicating that they chose larger “matching” bumps than the exact three polygon approximations to the reference cylinder. The SI values were higher for the sets with vision. Figure 6-17 shows a position trace, which verifies that in these trials, the subjects did actually trace a path close to the nominal geometry of the bumps.

Subject	Slope	Mean SD (%)	Mean SI	Mean # of Steps
TP	0.68	16.5	7.2	24
IU	0.86	14.0	7.7	40
AL	0.83	8.4	7.8	30
IG	0.60	12.7	7.0	19
KN	0.67	8.7	7.1	58
Overall	0.73	13.6	7.7	34

Table 6.7: Summary of Matching Results for the Five Subjects for Three Polygon Bump Without Vision

Subject	Slope	Mean SD (%)	Mean SI	Mean # of Steps
TP	0.78	13.1	7.3	16
IU	0.91	12.3	8.4	32
AL	0.85	9.1	7.8	25
IG	0.57	10.2	8.4	18
KN	0.62	5.6	8.4	51
Overall	0.75	11.7	8.1	28

Table 6.8: Summary of Matching Results for the Five Subjects for Three Polygon Bump With Vision




Figure 6-17: Position profile for the three polygon bump matching task.




Figure 6-18: Results for Three Polygon Bump Matching, No Vision

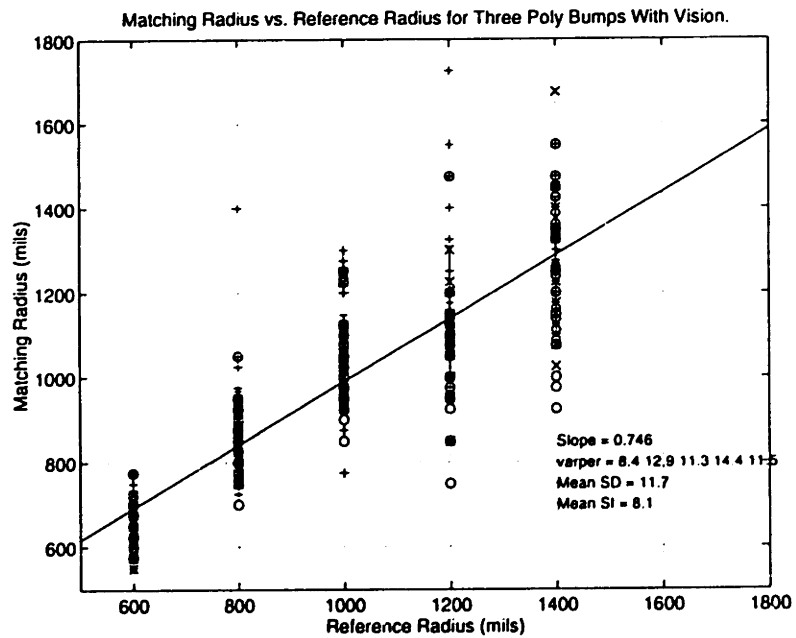


Figure 6-19: Results for Three Polygon Bump Matching, With Vision




Figure 6-20: Satisfaction Index Results for Three Polygon Bump Matching, No Vision




Figure 6-21: Satisfaction Index Results for Three Polygon Bump Matching, With Vision

6.2 The Classification Experiment

The classification experiment provided data on what shape the subjects thought they felt as they explored any one of the seven different types of display algorithms. The results are in Table 6.9. For each type of bump there were three possible radii. The numbers in each column labelled 1 through 4 represent the percentage of responses where the subjects characterized the bump as the given shape. The pictures associated with the various choices are shown in Figure 5-8. Choice 1 was a flat surface, choice 2 was a two polygon bump, choice 3 was a three polygon bump, and choice 4 was a circular cylinder bump. For each row in the table the answer with the greatest percentage of responses is in boldface. For the unshaded bumps, subjects were fairly good at correctly assessing the shape of the bump, although there were a significant number of errors, especially with the two polygon unshaded bump. With the addition of shading to the polyhedral bumps, however, the rounded bump was chosen more often than any other bump.

BumpType	Bump Radius	1	2	3	4
Flat Unshaded	$R = 600\text{mils}$	99	1	0	0
	$R = 1000\text{mils}$	96	1	1	3
	$R = 1400\text{mils}$	100	0	0	0
Two-Poly Unshaded	$R = 600\text{mils}$	0	79	3	19
	$R = 1000\text{mils}$	1	53	3	43
	$R = 1400\text{mils}$	8	47	1	44
Three-Poly Unshaded	$R = 600\text{mils}$	0	3	92	6
	$R = 1000\text{mils}$	0	3	86	10
	$R = 1400\text{mils}$	0	7	87	6
Real	$R = 600\text{mils}$	0	7	19	75
	$R = 1000\text{mils}$	0	9	3	88
	$R = 1400\text{mils}$	0	6	0	94
Flat Shaded	$R = 600\text{mils}$	2	3	38	57
	$R = 1000\text{mils}$	3	11	29	57
	$R = 1400\text{mils}$	2	14	36	48
Two-Poly Shaded	$R = 600\text{mils}$	0	35	16	48
	$R = 1000\text{mils}$	0	12	2	87
	$R = 1400\text{mils}$	1	6	8	85
Three-Poly Shaded	$R = 600\text{mils}$	1	3	39	57
	$R = 1000\text{mils}$	1	3	15	81
	$R = 1400\text{mils}$	0	7	8	86

Table 6.9: Summary of Classification Results for the Five Subjects, as a percentage of the total number of trials for each case.

Chapter 7

Discussion

These experiments demonstrated several important points about the perceptual effects of the addition of force shading to polyhedral approximations.

Bump Type	Slope	Mean SD (%)	Mean SI	Mean # of Steps
Flat, No Vision	0.43	14.4	5.2	37
Flat, With Vision	0.57	20.0	7.5	30
Two Poly, No Vision	0.67	11.7	7.4	34
Two Poly, With Vision	0.71	14.7	8.1	31
Three Poly, No Vision	0.73	13.6	7.7	34
Three Poly, With Vision	0.75	11.7	8.1	28
Cylindrical, No Vision	0.97	12.3	8.1	32
Cylindrical, With Vision	0.94	12.9	8.3	29

Table 7.1: Summary of Matching Results for the Five Subjects for Three Polygon Bump With Vision

Force Shading enhanced the subjects' perception that the object was a smooth cylindrical bump. The classification experiment directly compared force shaded polygonal approximations of the circular cylinder with unshaded polygonal approximations. The five subjects classified the appropriate unshaded polygonal approximations as the actual geometry more often than any other geometry. When shading was added to the polygonal approximations, all the subjects classified the bumps as smooth and cylindrical in character.

As the number of polygons used to approximate the bump increased, human performance approaches that of the real bump. Table 7.1 shows the overall results for each of the eight different tables presented in Chapter 6. The slopes and SI values approach that of the circular cylinder bump as the number of polygons increases.

Polygonal approximations need higher curvatures to be perceptually equivalent to the circular cylinder bump. In all of the matching experiments with force shaded polyhedral approximations, the slopes of matching radii vs. reference radii are less than 1. In general, the lower the slope is, the more the subjects lowered the radii before finding an acceptable match. This increased the size of the matching bump. Since the bump was convex, the polyhedral approximations underestimate the size of the bump in terms of height or cross sectional area. In the flat bump case, this was taken to an extreme. The subjects compensated for this by increasing the effective size of the bump, which in the flat bump case resulted in increased variation in the force direction. As the number of polygons in the approximation increased, the slopes became closer to 1, which was approximately the slope achieved by the subjects for the cylindrical matching experiment.

In the matching experiment, with vision present the subjects gave higher satisfaction index values than when vision was not present. The average satisfaction index values (taken over all subjects) and compared across experiment types, in all four types of matching tasks, was higher when vision was present. It was most apparent for the flat bump matching task when the satisfaction index was 7.5 with vision and 5.2 without (Table 7.1). In these trials, the discordance between the haptic positional and force information was at its greatest. Without vision, the subjects seemed to be a little confused by what they felt. When vision was added, there was an additional sensory cue, so they were more able to perceive the bump as having some shape.

Subjects give higher satisfaction index for higher number of polygons in approximation. As the number of polygons increased, so did the subjects' SI values.

Force shading caused the subjects to perceive a shape by feeling curvature via force direction changes, not position changes. The position plots show that during explorations of force-shaded bumps, the subjects trace out shapes similar to the nominal geometry. In the classification experiments, the subjects identified these shapes as smooth and curved. They did this, even though their hands actually traced out polyhedral shapes consistent with the nominal geometries.

Chapter 8

Future Work

The work reported in this thesis addressed some important issues in haptic rendering and haptic display, and found that force shading can be used to create the perception that a shape is smooth when it is actually nominally composed of polygons. There is still a lot more that is not known about this phenomenon.

Additional studies to investigate the following are recommended:

- Test the theta-critical hypothesis proposed in Chapter 5, and estimate the value of theta-critical.
- Test human perception for force shaded concave shapes. What are the interesting similarities and differences?
- Test human perception for force shaded textures.
- Incorporate the force shading algorithm into a generic haptic rendering engine.

Bibliography

- [1] J. E. Colgate, P. E. Grafing, M. C. Stanley, and G. Schenkel. Implementation of stiff virtual walls in force-reflecting interfaces. In *VRAIS '93*, pages 202–208, Seattle, WA, January 1993.
- [2] F. P. Brooks Jr., M. Ouh-Young, J. J. Batter, and P. J. Kilpatrick. Project GROPE - haptic displays for scientific visualization. *Computer Graphics*, 24(4):177–185, August 1990.
- [3] T. H. Massie. *Design of a Three Degree of Freedom Force-Reflecting Haptic Interface*. SB thesis, Massachusetts Institute of Technology, May 1993.
- [4] M. Minsky, M. Ouh-Young, O. Steele, F. P. Brooks Jr., and M. Behensky. Feeling and seeing: Issues in force display. *Computer Graphics*, 24(2):235–243, March 1990.
- [5] M. D. R. Minsky. *Computational Haptics: The Sandpaper System for Synthesizing Texture for a Force-Feedback Display*. PhD thesis, Massachusetts Institute of Technology, June 1995.
- [6] L. B. Rosenberg and B. D. Adelstein. Perceptual decomposition of virtual haptic surfaces. In *Proceedings IEEE 1993 Symposium on Research Frontiers in Virtual Reality*, pages 46–53, October 1993.
- [7] K. Salisbury, D. Brock, T. Massie, N. Swarup, and C. Zilles. Haptic rendering: Programming touch interaction with virtual objects. In *Proceedings of the 1995 ACM Symposium on Interactive 3D Graphics*, April 1995.

- [8] SensAble Devices, Incorporated, Cambridge, Massachusetts. *Phantom Documentation and Example Code*, 1994.
- [9] M. A. Srinivasan. Haptic interfaces. In N. I. Durlach and A. S. Mavor, editors, *Virtual Reality: Scientific and Technological Challenges*, chapter 4, pages 161–187. National Academy Press, 1994.
- [10] H. Z. Tan, M. A. Srinivasan, B. Eberman, and B. Cheng. Human factors for the design of force-reflecting haptic interfaces. In *Proceedings of the 1994 ASME Winter Annual Meeting*, 1994.
- [11] C. B. Zilles and J. K. Salisbury. A constraint-based god-object method for haptic display. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, August 1995.