Abstract of "Artistic Free-From Modeling for Scientific Visualization and 3D Illustration
PhD Thesis Proposal" by Daniel F. Keefe, Ph.D., Brown University, January 2006.

In this thesis proposal / initial thesis draft we present a methodology for creating scientific visualizations with collaboration from artists, illustrators, and designers. Our chief contribution is a set of free-form modeling tools driven by natural, artistic input that are controllable and expressive enough for artists to tackle difficult modeling subjects in science, medicine, and visual art.

Our approach uses virtual reality (VR) and free-form modeling techniques based on the paradigm of sweeping, 3D input. Several similar approaches to artistic, free-form modeling exist in the literature [15, 22, 51, 55], but our work differs in several key ways. First, our Cave-based tools make specific use of the large scale and unique properties of the Cave VR form factor to investigate its use as an unique artistic medium. Second, our two-handed, haptic techniques enable artists to create smooth and controlled marks directly in 3-space. With haptic feedback guiding the artists' movements, our two-handed approach is able to function as a far more direct method of drawing 3D curves than previous two-handed approaches [26] which require users to draw two 2D curves to create a single 3D curve. We also present several interface techniques that leverage the power of haptics to create rich styles of line, including the ability to adjust line weight by pushing against a force as an artist pushes against paper when drawing with charcoal.

These interfaces are motivated by a desire to enable artists to address more sophisticated subjects with free-form modeling techniques. In particular, we are interested in supporting artistic collaboration in scientific visualization.

Our 3D modeling tools are useful in illustrating science in the mode of traditional medical or scientific illustration. But, we also explore artistic collaboration in more exploratory styles of visualization. As part of this effort, we present a data-driven visualization design tool that combines an intuitive artistic modeling interface with a strong link to underlying multi-variate datasets. This tool enables artists to design complex, glyph-based scientific visualizations of fluid flow datasets directly within

virtual reality. Our scientific efforts are illustrated through applications in visualizing arterial blood flow and the complex flight and anatomy of bats.

We present findings from two pilot user studies and a proposal for a refined study of our haptic-based tools that will provide qualitative and quantitative insight into their use for depicting complex subjects.

Finally, we present a vision of the expected contributions from portions of the thesis which remain to be completed, including the study mentioned above and more advanced two-handed, haptics-based tools, along with a plan and timeline for completing this work.

Artistic Free-From Modeling for Scientific Visualization and 3D Illustration
PhD Thesis Proposal

by
Daniel F. Keefe
B. S., Tufts University, 1999
Sc. M., Brown University, 2001

A dissertation submitted in partial fulfillment of the
requirements for the Degree of Doctor of Philosophy
in the Department of Computer Science at Brown University

Providence, Rhode Island
January 2006

This dissertation by Daniel F. Keefe is accepted in its present form by the Department of Computer Science as satisfying the dissertation requirement for the degree of Doctor of Philosophy.

Date ______________ ________________________________

David H. Laidlaw, Director

Recommended to the Graduate Council

Date ______________ ________________________________

Andries van Dam, Reader

Date ______________ ________________________________

John F. Hughes, Reader

Date ______________ ________________________________

Fritz Drury, Reader

Approved by the Graduate Council

Date ______________ ________________________________

Sheila Bonde
Dean of the Graduate School

Contents

List of Tables

List of Figures

Chapter 1

Introduction

"The figure of a laborer–some furrows in a ploughed field–a bit of sand, sea and sky–are serious subjects, so difficult, but at the same time so beautiful, that it is indeed worth while to devote one's life to the task of expressing the poetry hidden in them."

– Vincent Van Gogh in a letter to Theo

In the age of computers, there is no tool more powerful to convey the complex, the serious, and the beautiful around us. Yet, within the world of three-dimensional computer graphics, the level of intuitive control over the computer medium needed to represent the beauty that Van Gogh describes is still relatively absent. While photo-realistic renderings of the natural world continue to become increasingly convincing, the ability for an artist to intuitively create with a computer, and achieve a representation or feeling as rich as what we find in nature, remains elusive.

In this thesis, we examine the artistic modeling process in detail. The fundamental challenge that we address is discovering how best to leverage artistic skill in creating 3D models, while providing a rich and controllable enough modeling environment for artists to tackle difficult modeling subjects.

The subjects that are most interesting to us are natural and complex. We are interested in scientific and medical illustration of 3D phenomena, designing 3D scientific visualizations for display in virtual reality, and artistic representation of human anatomical subjects. These subjects are difficult or impossible to model with popular tools such as Maya or CAD-based systems. They also require a level of sophistication and clarity of representation that is not currently attainable via free-form modeling approaches, but which artists are clearly able to achieve in physical mediums, such as drawing and sculpture. We believe artists can make significant contributions to the fields of science, medicine, visual art, and more if we can provide them with more appropriate means of investigating these types of serious, difficult subjects.

1.1 Free-Form Modeling and Current Limitations

Our approach is based on free-form modeling. This has come to mean many different things in the computer graphics and interaction community, but typically, "free" refers to the style of the 3D form generated in the modeling system. In most free-form modeling the 3D form is less rigid and constrained than in more traditional modeling approaches.

Often, free-form modeling systems utilize a style of input which may be described as freehand or sweeping. A common characteristic of many free-form modelers is the tight connection between the input, often sweeping movements or gestures, and the form that is created. This tight coupling usually yields a more immediate and physical modeling experience for the user than what is found more traditional systems. For some subjects this type of interface proves to be much more expressive and artistically interesting than more traditional modeling tools, where the mapping from user input to form generation is far less direct.

While this approach places a great deal of expressive power directly in the hands of the artist, the major limitation of modeling based on such direct, artistic input is controlling the input enough to create clear and refined forms. The loose, gestural models that often characterize free-form modeling, particularly in systems based on sweeping input paradigms like drawing, are interesting artistically, but due to their loose, sloppy-by-necessity style, they can be too limiting or even misleading for illustration problems in science, medicine, and artistic anatomy.

In this thesis, we present novel approaches to free-form modeling that make it a more controllable modeling medium, thereby allowing artists to address important subjects in scientific and artistic illustration that require artistic tools that are both highly expressive and controllable.

1.2 Scientific Visualization via Artistic Collaboration

As the ability for scientists to collect and compute more exciting and often voluminous data continues to rapidly increase, the ability to explore and present this data through effective scientific visualization becomes increasingly important. [32] However, effective visualizations of challenging data are difficult to create. Relatively new computer mediums, such as immersive virtual reality, offer great potential for visualizing today's most complex multi-variate, time varying, 3D datasets. [67] Yet, these technologies provide a special visualization challenge because they are so different from more traditional, extensively studied visual mediums. This makes it difficult for us to draw upon effective guidelines for good visual depiction in mediums like virtual reality because, for the most part, the guidelines are yet to be determined. As a result, some researchers have turned to cognitive science as a way of discovering truths about human perception and then strived to build more effective visualizations by leveraging their understanding of the human visual system.[31, 69] Others have turned to the study of successful artistic technique in traditional mediums and strived to apply the lessons learned to computer-based visualization. [39, 40, 43]

In the work presented in this thesis, we take a related but different approach to visualizing these challenging new datasets. We build upon the idea of learning from successful artistic technique, but rather than studying existing art and transferring the resulting knowledge to a new medium, we collaborate directly with artists. This collaboration is often difficult to achieve, in large part because few good design tools exist for artists to explore the types of representations we need from within the visualization target mediums.

Our contribution to the field of visualization is the presentation of a methodology and accompanying tools to aid this style of collaboration. Our efforts target experts in art, illustration, and design without extensive knowledge of computers, and certainly without knowledge of programming. Currently, these are prerequisites for most VR-based visualization design. Within these constraints, our work addresses the tension between providing designers with enough control over the design medium to accurately depict the science and providing interfaces that are intuitive, immediate, and expressive. Many of our insights and contributions result from exploring the balance between controlled, scientific representation and intuitive, expressive power.

1.3 Overview of Contributions

The main contributions of this thesis fall within the fields of human computer interaction, graphics, and visualization. This work also represents a significant contribution in the field of art, and has resulted in several publications in artistic venues along with works presented in juried art shows. Our efforts in visualization have contributed to scientific understanding in the fields of our collaborators, and the methods presented here have potential for significant long term impact in improving discovery, understanding, and communication within the sciences.

The contributions listed in this overview are the results of work that is completed in the sense that it has reached at least the threshold of being publishable. At the end of each contribution, I list where the work has been published, or note that it is part of the paper in progress that we plan to submit to SIGGRAPH 2006. In the SIGGRAPH 2006 category, the distinction between this section and the next "Proposed Contributions" section is essentially that this work, in terms of coding and experimentation, is already complete at the time of writing this proposal.

- A haptic interface for creating smooth 3D curves with varying line weight. *SIGGRAPH 2005 award winning poster [38], part of proposed SIGGRAPH 2006 paper submission.*

- A 3D charcoal drawing extension to the haptic interface for drawing smooth swaths of surface that enables the artist to simultaneously control variation in the swath shape and color value of the mark while creating it. *Part of proposed SIGGRAPH 2006 paper submission.*

- A data-driven visualization design system that combines intuitive artistic input and multivariate datasets to enable quick prototyping of glyph-based, time varying, VR visualizations. *IEEE CG&A 2005 [37], Vis 2003 Poster [33]*

- A methodology for collaborating with artists to design VR-based visualizations, along with analysis and discussion of key factors in teaching and enabling collaboration. *IEEE CG&A 2005 [37]*

- Applications in artistic anatomy: results and discussion from a formal, expert critique of styles of representing form with the CavePainting and haptic-based tools. *Unpublished*

- CavePainting: A Cave-based, free-from modeling system with a hand crafted aesthetic. *I3D 2001 [36], SIGGRAPH 2002 sketch [49], YoungVR 2001 [54], International Symposium for Electronic Art (ISEA 2002) [50], Several art shows*

- Application of pop-through buttons to the CavePainting interface. *VR 2002 [73]*

- Discussion of insights about Cave interaction and unusual use of the space of the Cave gathered from extensive user experience with CavePainting and building custom input devices for the Cave space. *Mostly unpublished, some discussion in VR 2004 Workshop [41]*

- Transfer of CavePainting interaction techniques to Cave-based scientific visualization applications. *Mostly unpublished, some discussion in Particle Flurries, IEEE CG&A 2004 [62]*

1.4 Proposed Work and Contributions

- A 2-handed haptic technique for smudging 3D marks in space to suggest smooth surfaces. *Part of the proposed SIGGRAPH 2006 paper submission.*

- A haptic extension to the 3D tape drawing interface for backing up and editing via overdrawing while creating a mark *Part of proposed SIGGRAPH 2006 paper submission.*

- A model for virtual charcoal that wears down with use and insights about its effect on richness and accidental discovery in VR. *Part of proposed SIGGRAPH 2006 paper submission.*

- A user study of artists trained with the tools via a RISD winter session course. Qualitative feedback from structured but open ended modeling tasks and quantitative feedback establishing the relative importance of aspects of the two-handed, haptic interface. *Suitable for publication in CHI, UIST, or similar venue.*

- An application of the illustration tool to visualizing and then presenting new data from Dr. Sharon Swartz' lab describing the material properties of bat bones, in particular, their unusual ability to bend dramatically during flight. *After seeing some initial sketches with our newest interfaces, Dr. Swartz thinks our illustrations may push this project over the threshold needed to submit a paper to Nature.*

1.5 Organization

This proposal is in the style of an initial thesis draft. The "Proposed Contributions" section above highlights some of the anticipated differences between this document and final thesis. Similarly, Section 9.1 within the Conclusions of the document describes proposed additions in more detail along with a plan and timeline for completion of the work. The work presented in Chapters 5, 7, and 8 is currently in progress and initial results along with expected outcomes are presented in these chapters.

The synergy between art and science is an important theme in this thesis: artistic insight aids scientific discovery, while scientific discovery pushes the limits of artistic representation and understanding. This interplay makes its way into the organization of the thesis as the artistic techniques described in one chapter play a key role in the scientific methods and discoveries of the next.

We introduce our work with a review of related techniques in artistic 3D modeling in Chapter 2. Chapter 3 presents CavePainting, our initial investigation into artistic 3D modeling tools using a Cave VR system. This is followed by a presentation of our methodology for artistic collaboration in scientific visualization and our data-driven visualization design tool in Chapter 4. Then, with the motivation for addressing more complicated visual subjects in mind, we turn to haptic interfaces for more controlled free-form 3D illustration in Chapter 5. We present applications of our tools for artistic anatomy in Chapter 6 and for scientific visualization in Chapter 7. Then, we present a user study of our haptics-based tools in Chapter 8, and finally a conclusion in Chapter 9.

Chapter 2

Related Work in Artistic 3D Modeling

Since Ivan Sutherland's Sketchpad system [64], marking the creation of the field, modeling form has been a key challenge in interactive computer graphics. Sutherland's vector graphics-based system was years beyond its time, including features commonly found in today's CAD tools and animation software. As the field has matured, several subareas of modeling have evolved. One of the first to become practical and popular is CAD or computer aided design. This type of modeling is characterized by geometrical exactness. Since CAD models are typically used as blueprints or manufacturing specifications, lengths must be exact and are often input as precise dimensions using a keyboard or mouse in a snap-to-grid mode. While extremely useful for many industrial applications, CAD programs fall short as artistic tools and even as tools for initial concept design for industrial applications in part because of the requirement of exactness and also due to the training required to learn them. Paper and pencil are often preferred to CAD-like systems for initial design sketches and more gestural drawings.

Of course, paper and pencil do not translate so well to 3D, and the notion of finding a 3D extension to the immediacy and gestural qualities of traditional mediums like paper and pencil introduces a second subarea of computer modeling. Where CAD-style modeling fills the role of the T-square and drafting pencil, free-form modeling is often seen filling the role of the artist's charcoal pencil. In fact, even within the area of free-form modeling there is great variation in things like the geometrical constraints that may imposed on the modeling process and the level of connection that the user input has with generating form. But, in general, free-form modeling approaches tend to be highly guided by the user and the resulting 3D forms tend to be more organic, with fewer straight, flat surfaces and right angles, than that produced by CAD and like minded systems.

We have found the following criteria to be helpful as a means of classifying and understanding 3D modeling tools and their application to modeling challenging artistic subjects.

- The dimensionality of the input used to create the form.
- The extent to which the input is freehand or geometrically constrained.
- The format of the output: implicit surface, mesh, points, curves, etc..

Dimensionality of input: One of the clearest distinctions between 3D modeling systems is the dimensionality of the input used to produce the model. Many, in fact the majority, of 3D modeling systems use 2D, not 3D, input to create form. Depending on the modeling problem and the disposition of the artist or researcher designing the tool, 2D input can provide either a great opportunity to simplify a difficult problem, or it can provide indirect and unwieldy interface for one of the most inherently 3D tasks commonly performed with computers.

Freehand vs. geometrically constrained: Geometrically constrained 3D modeling systems, such as CAD systems, are typically very precise and clear in the form that they create. Their shortcomings for our driving applications lie in their inability to easily represent organic subjects and often also in their lack of appropriate interfaces for leveraging the skills of trained artists.

Freehand tools, on the other hand, are usually only good at representing organic subjects, but are almost always too difficult to control to produce refined and clarified depictions of these subjects. The tight coupling between user input and output that usually exists in freehand tools often requires the user to be a trained artist. These tools often do leverage existing artistic skill. However, perhaps because they require users, and sometimes developers, to have a considerable amount of artistic skill, there is a trend in freehand modeling research to target quick sketching interfaces and rough output forms. Extremely compelling systems, like Igarashi et al.'s Teddy[30], exist for turning a few gestural freehand movements of an input device into very interesting 3D models. But models from these types of sketch-based systems are usually blobby forms that lack the specificity required to tackle important subjects in medical illustration or artistic anatomy.

Output format: 3D modeling is certainly a broad topic and even within the sub-area of artistically relevant 3D modeling there is considerable variation in application area and desired output format. In industrial design, the primary application area for SensAble's FreeForm tool[58], models are often sent to rapid prototyping machines for output as true physical shapes cast in plaster. These rapid prototyping 3D printers require watertight, thick triangle meshes as input and this requirement dictates the type of modeling approaches that can be used to create them. Likewise, animation software, such as Maya or 3D Studio Max, are built to import models as triangle meshes or NURBs surfaces, so 3D models intended to be used for this type of animation need to represented as such.

In general the models produced with the techniques described in this thesis are intended to be viewed, and not necessarily animated or printed, so we are less concerned with adhering to a specific output format than typical modeling systems. This gives us a bit more flexibility to explore alternative styles of representation.

In the following sections, we present a thorough review of related artistic modeling approaches with an emphasis on those that utilize direct 3D input in the style of the work we present.

2.1 Techniques Utilizing Direct 3D Input

2.1.1 Early Work

Sutherland introduced the first head-mounted, tracked, stereo display in 1968. [65] Eight years later, Clark combined this hardware with a tracked wand device to make the first virtual reality system for direct manipulation of 3D surfaces. [10] Clark's work introduced the idea that the intuitive control afforded by direct 3D interaction coupled with a 3D display could be a better paradigm for 3D modeling tasks than more indirect methods involving 2D projections of 3D objects or mapping 2D input to a 3D space. He also discussed the implications that this approach has for the process that designers might take. Freeing designers from the requirements of knowing particular numerical values of coordinates describing a surface or orientations of coordinate system axes allows them to concentrate fully on the design task. These concepts resonate with our work and many others who have followed Clark.

Schmandt [56] was the first that we know of to implement a modeling system that could generate form from sweeping movements of a 3D input device. Schmandt's system used an early Polhemus 6D device to create a magic wand that could emit 3D paint in space. He used half-silvered mirrors, conventional video monitors, and shutter glasses to produce a stereo view of the 3D paint. Schmandt's work was one of the first experiments into the interactive capabilities of stereo displays. His results indicated a good natural correspondence between the wand and the 3D paint, but lag and distortion in the tracking field were problems for the interactive feel of his tool.

Galyean and Hughes [22] used direct 3D input with a monoscopic display to produce a compelling voxel-based modeling approach. They even created a "poor man's" 3D force feedback device to assist in controlling the input tool by suspending a Polhemus tracker in a cube with eight elastic cords attached to the corners of the cube. This work was the first of the 3D sculpting systems that use a clay metaphor and one of the first to examine creating artistic, blobby models that were a drastic departure from the more rigid geometric shapes produced with CAD programs. In this system, material could be cut away or added to a block of clay. A great deal of free-form modeling work has been based on this metaphor of working with digital clay, and it remains one of the most successful approaches for obtaining output in the form of triangle meshes or voxel-based models, which is useful for applications in rapid prototyping and animation. Advances in computer and haptic hardware continue to enable a wide variety of extensions to this work, providing for more control of the interface and more sophisticated forms. CavePainting and related approaches presented here build on the model of sweeping out 3D form presented in this work and in that of Schmandt. One of the key differences in CavePainting and Clay-based free-form modeling approaches, like that of Galyean and Hughes is the departure from a physically realizable or solid, surface-based model. In CavePainting, form does not need be watertight or even connected, and the strokes that are produced can be infinitely thin. This is actually a significant hindrance if the goal is to export the model from the system and send it to a rapid prototyping machine or an animation package, but it allows the creation of a different style of form that is interesting for art and illustration.

At roughly the same time as Galyean and Hughes, Sachs et al. presented the 3-Draw system [51], which used direct 3D input with both hands to create an intuitive CAD tool. This work was ground breaking in terms of establishing computers as a viable tool for the initial phases of industrial design. 3-Draw models consisted of lines only, but the lines can easily be thought of as defining surfaces. The tool allowed for both unconstrained and various types of semi-constrained sweeping input to specify complex 3D curves as well as interaction techniques for specifying start and end points for curves and several curve editing operations. This work was important in establishing free-form 3D input as a viable tool for serious design problems and in establishing interaction techniques for enhancing control over free-form input, an area explored in detail in later chapters.

The 3DM system [8] presented by Butterworth et al. built on Clark's conception of a head-mounted display modeling environment with the additional goal of providing the user with an interface that was as easy to learn and use as what was commonly available in 2D drawing programs like MacPaint. In addition to performing geometrically constrained CAD-style modeling operations, 3DM incorporated a successful modeling mode based on sweeping 3D input for creating surfaces by extruding curves. One of the important contributions of 3DM with respect to this thesis is the discussion of user experiences with this extrusion tool. Users were reported to have difficulty aligning 3D objects, keeping two triangles parallel, and doing other geometrically constrained operations in this VR environment, but for the first time they could easily perform highly complex free-form extrusions. Extrusions had already been proven to be extremely useful modeling tools in desktop-based programs. The difference in 3DM was the user's ability to perform the extrusion directly in 3D by dragging and twisting an extrusion curve along a free-form path. In this way, users could naturally control more than one parameter (position and twist) of the extrusion while creating it. CavePainting's extrusion stroke uses this exact concept and our feedback from users builds on the comments reported by Butterworth et al. Artists using CavePainting have found this ability to control several aspects of a curve at once an essential, defining aspect of the tool. In fact, with CavePainting, this seems to be most useful with strokes that are far simpler than extrusions. The simplest strokes, like ribbons, are perfect for controlling both position and twist along a curve. The result is an aesthetic that looks intriguingly hand-made, a rarity a computer graphics.

Deering's Holosketch [14] was the first system to combine a head-tracked, stereo VR environment with a modeling system that was geared towards artistic creation. Holosketch was a strikingly complete modeling and animation package with a fully developed menu of modeling modes and operations. Several of the drawing modes in Holosketch were based on continuous sweeping input, including a toothpaste mode reminiscent of Gaylean and Hughes additive sculpting, a wire-frame lines mode reminiscent of the 3D line drawings produced in 3-Draw, and a mode where clouds of random small triangle particles were left behind the wand as it was swept through the air. Holosketch worked in a fishtank VR [13] setup using a 20 inch CRT. The drastically different scale of this system compared to that of the Cave's 8 foot cube space is one of the important differences between Holosketch and CavePainting. Scale has important implications for the gestures, widgets, and props used in the two approaches, and it also has an impact on the types of gestural motions that artists can use to create form. For example, feedback from CavePainting artists suggests that in the large volume of the Cave our own bodies can be used as a useful tool in controlling free-form input. Just as traditional painters learn to do when working on large canvases, CavePainters often lock their elbow to form a natural fulcrum at the shoulder in order to make more controlled arcing movements of the hand. This type of large scale, gestural interaction that leverages the mechanical properties of our entire arm, if not the whole body, is impossible in smaller, desktop-based systems.

Deering reports going to extreme lengths to calibrate the head and wand tracking in Holosketch, even to the point of correcting for distortions due to the curvature of the CRT, the index of refraction of the glass of the CRT, and changes in interocular distance due to rotation of the viewer's eyes in their sockets. This "calibration fanaticism," as he describes it, leads to tracking accurate to within 0.25cm. This makes his reports of user feedback about controlling the tool of special significance. Even with this precise calibration, Deering reports that users had difficulty accurately aligning 3D objects in space. However, adding the ability to adjust the scale of the model coupled with a 10x reduction mode for user input and armrests for the user is reported to eliminate this accuracy issue, at least for moving objects around. It is unclear from Deering's published work whether any of these additional features could be applied in Holosketch to help control the creation of geometry, like his "toothpaste" forms.

2.1.2 Recent Approaches

In more recent work, two classes of free-form modeling techniques that utilize direct, sweeping 3D input have emerged: those based on large scale movements of a tracked device in the air, as in CavePainting, and those based on haptic feedback, usually utilizing a clay sculpting metaphor as first introduced by Gaylean and Hughes.

Schkolne's Surface Drawing [55] is an interesting example of free-form modeling utilizing a large screen display device. Unlike CavePainting, the display used for Surface Drawing is a flat table top. This display lends itself to the use of physical props that can be placed on the table when not in use. Natural use and selection of appropriate props and metaphors is the main thrust of Schkolne's work. [52, 53] To create form in the system, the artist uses his hand augmented with a bend sensing glove as a device for sweeping out bits of surface in space. These surface fragments are then stitched together to create a solid triangle mesh model. The gestural, full-bodied input characteristic of Surface Drawing is the closest to CavePainting out of all these related approaches. The similarity in the feel of the tools is probably best attributed to the large scale at which artists work in both environments. Surface drawing and CavePainting differ in the particular interaction techniques used, which are tailored in each system to the display environment. From an artistic modeling standpoint, the most interesting variation in the two approaches is in the type of forms they strive to create. Although surface drawing is almost entirely an additive process, it is similar to many free-form tools based on a sculpting metaphor in that it is tailored towards creating solid forms composed of surfaces depicted by a single, or sometimes a small number of triangle meshes. In CavePainting the individual strokes are only visually fused into a model, they are never actually stitched into a single triangle mesh. This makes the artistic process of creating form and thinking about modeling, as well as the results strikingly different.

Several other large scale, open air input systems exist for free-form modeling. The FreeDrawer [70] system runs on a responsive workbench and is a good example of a modern approach to a spline-based modeling system that utilizes sweeping 3D input, but with constraints setup to enforce valid spline based representations for curves. Of particular artistic note is the work of the artist Mäkelä, who has teamed with researchers at Helsinki University of Technology to create a system where form may be generated by tracking fingertips. [46] The fingertip control, achieved through a custom ultrasonic input device, adds the ability to control the thickness of swaths of form as they are swept out in space, a feature similar to that provided by the haptic interface presented in Chapter 5. Mäkelä's work also illustrates a compelling visual effect of combining point and surface-based representations. The BLUI [6] system has been the topic of several sketches at SIGGRAPH in recent years in which physical printouts, both 2D and 3D of free-form objects created with the system have been presented.

In the area of haptics based free-form modeling, SensAble Technologies, the makers of the PHANToM force feedback device, introduced what was probably the first haptics modeling system that allowed artists to feel the 3D shapes they modeled while pushing, pulling, and deforming them. The tool, which has undergone considerable refinements and is still available today is called *FreeForm*™ [58]. The current incarnation of the tool targets product designers working anywhere from the early conceptual stages of product design to final steps where output from the modeling tool can be sent to rapid prototyping machines for production of physical models. The models that skilled artists can create with this system are impressive in their clarity and refined aesthetic, two of the main goals for the work described in the remainder of this thesis. However, this tool and other similar approaches [23], including those based on implicit surface representations [29] are only useful for creating solid models, that is they cannot be used to directly explore the loose, stroke-based aesthetic described in CavePainting.

One haptics-based modeling system that takes a different approach than the rest is the springs and constraints presented by Snibbe et al. [61]. The haptic drawing interfaces presented in Chapter 5 build upon this work. Snibbe et al.'s approach is different in that it proposes dynamic haptic models to help artists control the creative process, but these models are not based on interacting with a static geometry or properly simulating contact forces. They help to control and guide free-form input, but tend to allow artists to remain gestural in their interactions. The new medium described in this work is left relatively unexplored artistically, but it is stroke rather than geometry based and therefore these interaction techniques could fit nicely into a CavePainting-style conception of model building.

Several other systems describe techniques termed 3D painting, sometimes including haptic feedback as part of the interface. [28] These refer to methods of painting color or texture onto a 3D model, not creating the 3D form of the model itself.

Direct 3D input has also been used as a way of modeling by tracing. [27] This concept has also been explored extensively artistically with the Surface Drawing tool. [24]

Of final note is Schroering et al.'s work [57] in which a light pen is used to draw on a tablet that is tracked as it is moved through the air. The drawing/modeling application presented in this work is quite limited in its scope, but we could imagine this input style as having important implications for free-form modeling.

2.2 Techniques Utilizing Large Scale 2D Input

Some of the most interesting work improving the controllability of free-form input for modeling has been that of Buxton's group on two-handed techniques for drawing smooth curves. [3, 9, 25] This work presents a digital version of tape drawing, a technique commonly used in car design for creating large scale drawings made from smooth curves. In an extension of this technique to 3D [26], 3D curves are created by drawing two 2D curves, one of which is projected onto a surface defined by the other. This two step process is described as easier to control than a more direct 3D input of the curve. This may be true in industrial applications, like car design, where it is traditional to think about curves in space in terms of their projections onto planes. However, for artistic purposes and in applications where artists wish to draw many, many curves rather than just a few principal ones a two step curve drawing operation is impractical.

Another technique from the same research group uses a high degree-of-freedom input device to create curves. [66] The approach of using “shape tape” to easily describe complex curves seems promising artistically since it appears to be such an immediate method. However, the same two-step approach is required for creating a 3D curve from this style of input, making a direct application of this technique impractical for the quick, gestural aesthetic found in CavePainting.

2.3 Techniques Utilizing Desktop Scale 2D Input

Many, many techniques for generating free-form 3D shape from 2D input have been developed. Most of them are derived or inspired by Zeleznik et al.’s seminal work Sketch. [72] Sketch made creating 3D form quick and intuitive and with the addition of Skin [45], complex, organic shapes could be modeled quite effortlessly.

Teddy [30] is another important system in this line of research. In Teddy a 3D form is inferred by free-form 2D input. This work opened up the door for a whole line of inquiry on the best techniques for translating 2D free-form input into 3D form. Several works have had great success in this area. [34, 48] However, by definition this work guesses or makes assumptions about the 3D form the artist intends to produce. The forms that result from this work can be complicated geometrically, but often appear blobby or unrefined and inappropriate for many styles of art and illustration.

Other work has taken the approach of sketching multiple 2D views of the lines that form an object [2, 5, 4, 35] with the idea that a more refined 3D shape can be obtained with additional 2D constraints. This direction shows promise, especially when used as an editing tool, and it works with standard computer hardware. However, the approach we take with CavePainting and its derivatives is that artists can control 3D input when presented with the right interface and that this 3D input is a much more direct way to specify, refine, and clarify complicated 3D models.

In summary, there is currently a tradeoff artists must make in working with 3D modeling tools. Tools based on direct, 3D artistic input are typically very difficult to control with enough precision to address real subjects in science and medicine. In contrast, full-featured modeling and animation packages, such as Maya, support extremely precise form, but these systems take years to master and lack intuitive interfaces for the creation of complex, natural forms. Techniques based on large scale, 2D input show promise for industrial concept design, but are less appropriate for illustration of natural phenomena. Pen-based systems often leverage artistic skill, but the 2D input they use lacks directness for 3D modeling tasks. Our work builds on previous approaches based on direct 3D input, and explores alternative interfaces and styles of expression that advance the utility of direct 3D input approaches for modeling complex, natural subjects.

Chapter 3

CavePainting

3.1 Background and Motivation

The CavePainting [36] system is the basis and inspiration for much of the work described in this thesis. The interface and artistic process developed for CavePainting is made possible by the unique properties of the Cave [12] VR display. In particular, the large scale of the Cave combined with the experience of being completed immersed within a VR world, yet retaining the ability to see and interact with physical props, allows Cave artists to work and create models with CavePainting that are qualitatively different than other free-form modeling approaches.

The goal behind CavePainting was to provide an intuitive interface for artists to explore unfamiliar ground. The unfamiliar ground is the Cave and the unique potential it offers as an artistic medium both for creation and display of artwork. To traditionally trained artists, just the experience of walking into the Cave, a sterile, 8 ft. x 8 ft. x 8 ft. cube room made of projection screen walls while being wired up with shutter glasses and tracked props can be both daunting and dazzling. Various artists and visitors have recounted as much in descriptions of their encounters with the Cave [47, 68]. Working artistically in this environment is a drastically different process than investigating scientific datasets or doing architectural walkthroughs which are much more common uses of Cave facilities. In fact, asking an artist to walk into this unfamiliar environment and create meaningful art with it is a very difficult

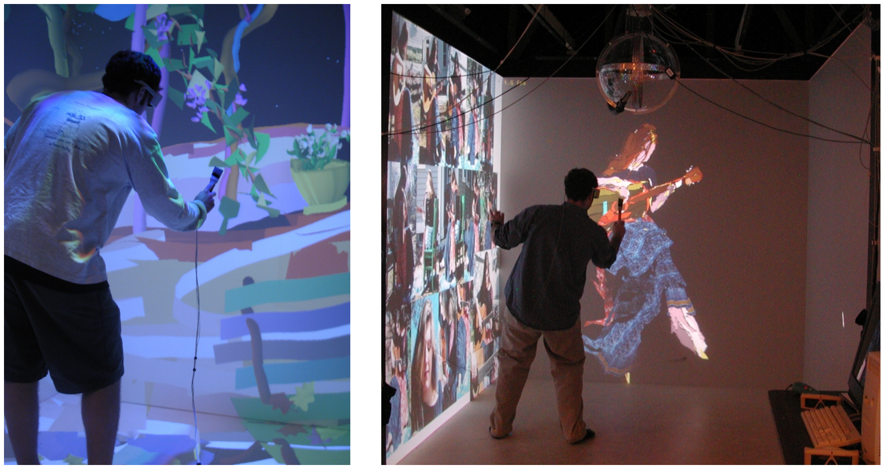

Figure 3.1: Working with CavePainting.

proposition. CavePainting strives to enable a smooth and approachable artistic process within the Cave by providing an intuitive interface derived from the metaphor of traditional oil painting, as illustrated in Figure 3.1.

In fact CavePainting follows the metaphor of oil painting in both its interface and the 3D form generated with it. In traditional painting, when individual brush strokes are clearly visible in the final work, the piece is often described as "loosely painted." Impressionistic paintings, for example, often fit this description [71]. When we examine the brush strokes in these paintings closely, we find subtle and striking variation in color, size, shape, and texture. If we move close enough, our eyes can only see the individual brush strokes. At this level, the amazing variation in the type of stroke used, even within a single painting, is apparent. The layering of strokes on top of each other is also apparent at this close level. As we move away from the painting, we stop seeing individual strokes. Our mind is able to fuse the strokes together and comprehend a complex scene. 3D form created with CavePainting is intended to function in much the same way. 3D brush strokes, that are clearly visible up close, are combined together to suggest 3D shape. Typically, it takes many brushstrokes to completely specify an interesting 3D form. But, as a viewer steps back to view this form, it comes together in a meaningful way, just as the brush strokes that compose a painting fuse into a recognizable scene.

Over the five years since CavePainting was first developed, many artists have used the system both artistically, and in recent years, as a VR illustration tool for

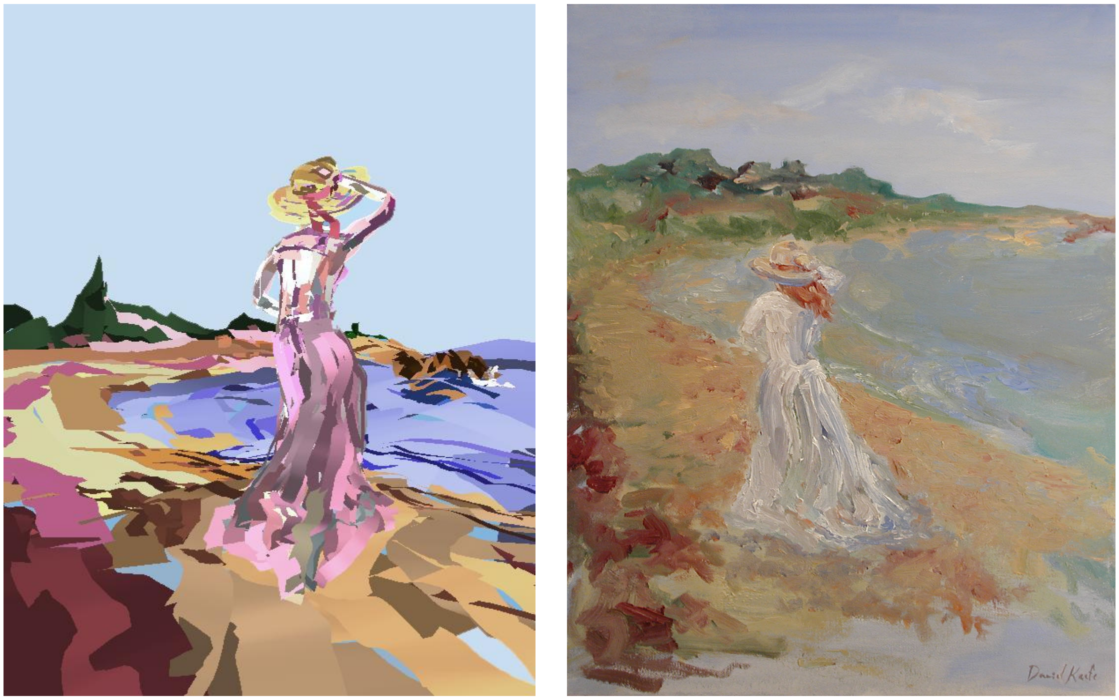

Figure 3.2: The CavePainting on the left was inspired by the oil painting on the right.

describing complicated problems in science and medicine. The insight and feedback from this continuous use has directly guided the research presented in this thesis. In the remainder of this chapter, we describe the specifics of the CavePainting interface and process in more detail.

3.2 Free-Form Modeling with CavePainting

3.2.1 Suggesting Form with 3D Paint Strokes

Unlike traditional 3D modeling approaches, CavePainting models are built from disjoint brushstrokes. These tiny bits of surface combine visually to form the impression of a larger subject, mirroring the function of brush strokes in traditional painting. Thus, the visual goal for the CavePainting modeling system was to explore this new space of 3D forms that can be generated using a painting "build up a form through layering many brushstrokes together" metaphor. The resulting form can exhibit a similar style to that of loosely painting oil paintings as illustrated in Figure 3.2.

This approach to constructing form was heavily influenced by the work of Laidlaw et al. [40, 39] who demonstrated the utility of applying painting metaphors as a way of realizing coherent, complex multi-dimensional 2D visualizations. One of the chief painting techniques utilized in this work was the idea of layering strokes on top of each other. As we began to consider how this layering approach might translate to 3D visualizations, we became interested in the notion of what layering might mean in 3D and in VR. CavePainting was born out of the desire to investigate a layered brush stroke means of 3D representation.

We found that this approach to 3D modeling lends itself to a style that is often described as loosely painted, similar to that commonly found in Impressionist painting. Like Impressionism, CavePaintings suggest rather than precisely define form.

3.2.2 Input Devices and Techniques

Painting with a Tracked Paintbrush Prop

The standard mode of creating form in CavePainting is to make sweeping movements through the air while holding a paintbrush prop. Pressing and holding a button affixed to prop turns the brush on and a trail of 3D form is left behind the brush according to the currently selected stroke type.

We have experimented with several variations of the prop used for this task as well as varying the position of the button on the prop for the custom made devices. Some of the variations tried are shown in Figure 3.3. The prop used does seem to affect performance and comfort, although we have not performed a quantitative study on the subject we feel comfortable making some general statements about the most effective configurations based on our experience and feedback from many, many artists. However, all of our configurations were based on using wired trackers. Some of the lessons learned from our investigations may not translate directly to what we imagine is the ideal situation: a lightweight, wireless brush. Such a configuration may be practical in future VR systems with robust camera-based tracking.

In general, our custom made paintbrush props were more comfortable and easier to control than the commercially available wands we tried. Artists preferred to hold the brushes with their index fingers extended toward the tip of the brush as shown in Figure 3.3a, thus we found the best placement of the button on the brush to be toward the tip of the brush where the index finger would fall naturally. The commercially

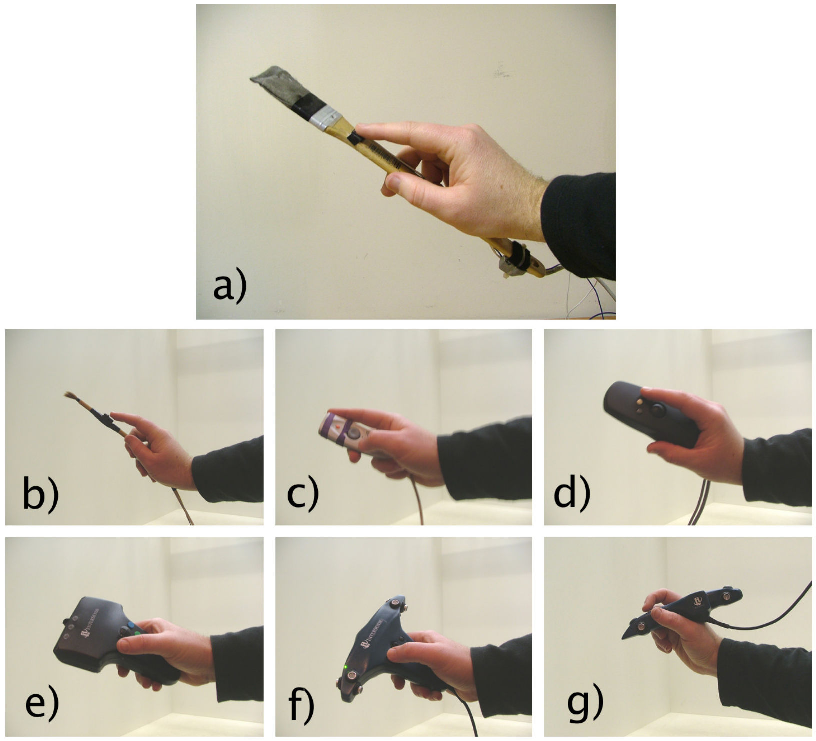

Figure 3.3: Various tracked paintbrush props: a) Our best custom configuration, b) An alternative round brush configuration, c) a modified wireless mouse, d) An Ascension wanda device, e) An InterSense wireless wand, f) An InterSense wand, g) An InterSense stylus.

available wanda device and both the InterSense wand and stylus devices (shown in Figure 3.3d-g) we tried are not designed to be held in this orientation, and that along with the weight of these devices are the main reasons why artists preferred the brush props. Interestingly, the InterSense stylus device (shown in Figure 3.3g) seems to be manufactured specifically for performing pen-based operations, but holding it in a normal pen-like orientation was uncomfortable for the large scale drawing motions of the type found in CavePainting because it severely constrains freedom of movement in the wrist. Artists working on large canvases rarely hold a brush the same way they would a pencil for writing text, so we think this observation has some backing in the evolution of traditional artistic technique.

We were also surprised that artists preferred brush props with a rectangular cross section to those that were round. It is possible that this would change if the brushes were wireless. The rectangular cross section seemed to make the brushes easier to

Figure 3.4: A physical bucket prop is used to throw virtual paint onto the walls and floor of the Cave.

rotate against the tension in the wire attached to the tracker.

The best commercially available solution that we found was to buy a small, lightweight wireless mouse of the shape shown in Figure 3.3c and attach a tracker to it. This device is smaller and flatter than the VR wands we tried, and was nearly as comfortable and usable as the custom brush props.

Weight of the prop itself and also of the cord attached to the tracker is a factor. The position of the tracker wire coming off the prop affects its balance in the hand. Artists were often seen holding the wire coming off the brush prop in their non-dominant hand to ease the tension and perceived weight on the brush while painting. Placing the wire in a position such that it comes off the prop at a point under the center of the artist's palm rather than running it all the way to the end of the brush helps to diminish the negative effects of the wire.

Painting the Walls of the Cave

As an alternative to sweeping brush movements through space, it is also possible to create form in CavePainting by throwing paint from the brush and by pouring it out of a bucket prop as seen in Figure 3.4. In practice, these techniques speak more to the different style of using the space of the Cave in CavePainting than they do to the goal of creating compelling 3D scenes. In these modes, paint appears to splatter onto the physical walls and floor of the Cave. This is a simple idea with a compelling result. This style of interaction mixes virtual reality with the reality of the physical room surrounding the artist or viewer. In most Cave-based programs the projection screen walls are made to disappear. Even though users know they just walked into a 8 foot cube, when the lights go out and the projectors come on, there is no indication of that physical room that we know we are in somewhere in the back of our minds. Artists describe this painting mode as "defining space."

In addition to the sense that the physical space is realized with some virtual form, there may be some perceptual basis for the rather compelling visual effect we observe when we splatter paint on the walls of the Cave. When working in a projection-based VR system we see in stereo, but unlike the real world, our eyes only converge upon one depth, that of the projection plane. Recently, researchers have been working to develop true 3D displays that provide multiple planes of convergence and initial results seem to indicate that viewers sense a heightened awareness of 3D with these more perceptually accurate displays. [18, 42] In the Cave, our eyes converge only on the projection wall surfaces and we have discovered through the use of CavePainting that displaying virtual geometry coincident with those walls appears to improve ones sensation of 3D viewing, even for geometry that floats in the center of the Cave away from the walls.

This phenomenon was compelling enough to investigate in a formal user study [31] and in an art installation [49, 50] developed with CavePainting. Results from the user study did not indicate that displaying geometry on the walls of the Cave has a significant effect on user performance in the task chosen for the study. However, the current hypothesis is that "looking out" tasks, which are often appropriate characterizations of viewing CavePaintings are more susceptible to this effect than "looking in" tasks like the one tested in the study.

In contrast to the results of this user study, our artistic endeavors indicate that a very different sensation is produced when we embrace this sense of virtual and physical coexistence using the Cave walls. This speaks to another motivation for this type of artistic application. Collaboration with artists has great potential for

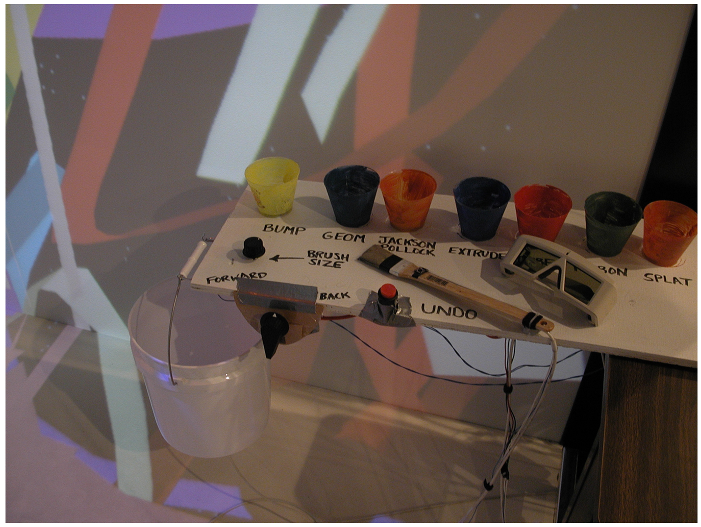

Figure 3.5: This table of props in the Cave mixes physical interaction with the virtual Cave medium.

discovering new modes of virtual representation that can be used in the relatively unexplored medium of VR.

A Physical Prop Table in the Cave

CavePainting uses the space of the Cave in another non-traditional way by placing a table that holds interaction props along one wall of the Cave. There is no question that the interactions performed through this table could be performed more efficiently through virtual popup menus located near the artist or any number of other virtual interaction techniques, but efficiency is not the point of these props. The goal for the painting table, seen in Figure 3.5, is to experiment with a style of movement and physical interaction in the Cave that we have not seen before. In the scientific and architectural programs that dominate the landscape of Cave applications, the typical interaction of the users of the Cave is fairly limited. It generally consists of standing more or less in the center of the Cave, often with a colleague to either side prohibiting movement, and then pointing with a wand and a virtual laser pointer. Often some amount of navigation in the virtual world is done through a flying metaphor, and sometimes the user walks through the space of the Cave to attain a better view of the data being visualized. What is noticeably absent in this picture is interaction that uses the Cave in the style for which we contend it is most useful. In CavePainting, in contrast, artists routinely use almost the full space of the Cave. Just as a painter working on a large canvas paints a few strokes, then steps back to look at the work, CavePainters are in almost constant motion while working, making use of one of the defining characteristics of the Cave, its large walkable size.

To promote this style of walking interaction in the Cave, we place a table inside the Cave along the right wall. The painter dips the brush prop into small cups of "paint" set on the table to switch the type of brushstroke form linked to the brush. This type of physical interaction inside a Cave is rare because the Cave cube is such a sterile place. With all of the walls made of projection screen material, users must be careful not to touch the walls and must take their shoes off when the enter the facility. This makes using props within the Cave a bit tricky, especially since there are no physical bits of furniture upon which to place them when not in use. The painting table is made of a cantilevered board that can be wheeled into the Cave so that it is supported from outside of the Cave and no part of it actually touches the Cave projection screen surfaces.

One experience that illustrates the type of connection that users feel with physical props such as this occurred when a class of ninth graders came to visit our Cave. While demoing CavePainting to this class, we noticed that when the students walked over to the table to dip the brush into one of the paint cups they would stick the brush in and move it back and forth all around the bottom of the cup, as if they were making sure it was completely loaded up with paint before they went back to painting. This sense of presence is something that is hard to quantify, yet it seems to make a real difference, and we believe that factors like feeling comfortable walking throughout the whole space of the Cave have a real effect on performance for tasks like the ones we find in CavePainting.

Pop Through Button Devices and Position Dependent Two Handed Interaction Techniques

Many of the interaction techniques in CavePainting rely on the use of both hands. The brush device described above is held in the dominant hand, and some sort of tracked device is used in the non-dominant hand. In the original implementation of the system a Pinch glove device was worn on the non-dominant hand and used

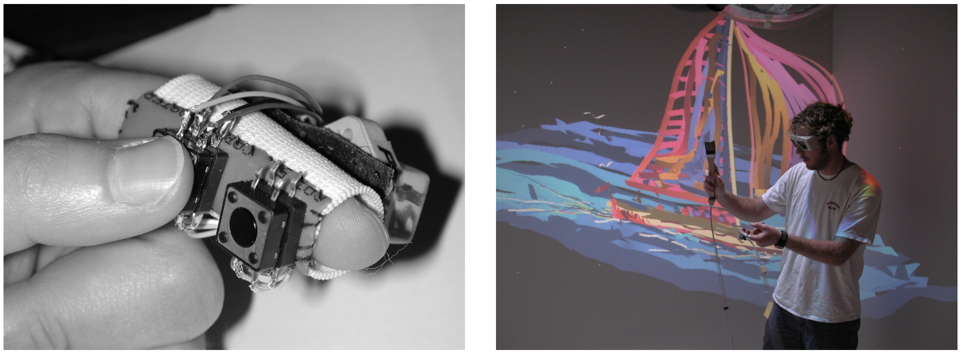

Figure 3.6: Left: The FingerSleeve device with two pop through buttons and a 6 DOF tracker mounted on it. Right: The FingerSleeve is worn on the non-dominant hand while a brush prop is held in the dominant hand.

to activate various functions in the program. After considerable experience with this interface we discovered several inadequacies of the Pinch glove device itself and several areas in our use of it that could be improved. This led us to develop a novel class of input devices that combine pop through buttons with 6 DOF trackers and to redesign the CavePainting interface to take advantage of these devices to produce a cognitively more natural mapping from physical interaction to virtual action.

Pop through buttons have two activation levels corresponding to light and firm pressure. Three characteristics of pop through buttons can be exploited to improve virtual environment interaction:

- Twice as many activation states are available in the same physical surface area (and the corollary that only half as much surface area is needed to achieve the same expressive power) as a traditional button device.
- A bare minimum of additional physical activity is required to activate the additional state, less than that required to activate two different traditional buttons.
- The physical action of popping through one button state to another is arguably cognitively more natural for activating inherently sequential or closely related tasks than pressing separate buttons would be.

The Pinch glove interface used in the original CavePainting system provided users with four distinct contacts which we mapped to four different painting modes: color picking, resizing the virtual brush, translating the world, and toggling scaling mode on and off. These tasks were activated by pinch gestures made between the thumb and one of the fingers. We mapped common tasks to the index, middle, and ring fingers since pinching these fingers is far more comfortable than the pinky. In fact, many users had real difficulty making the pinky pinching gesture, especially if the gloves were too big for their hands. These uncomfortable gestures were one thing we hoped to avoid in our interface redesign. We also discovered some practical limitations of the gloves in that they are slippery and this makes working with other props difficult, and after the consistent, daily use they received in our system the pinch contacts and wiring began to wear out causing many pinch gestures to go unrecognized in the system and frustrate users.

The important conceptual issue that we hoped to avoid in our redesigned interface was the rather arbitrary assignment of a particular function in the program to a finger on the hand. This issue is illustrated when describing the system to a new user. To users of the system, grabbing objects seems very natural, and that is one nice thing about the pinch gloves, but there is no meaningful reason to help users remember that objects should be grabbed with the thumb and middle finger, while colors should be chosen with the thumb and index finger, the world should be rescaled with my thumb and ring finger, etc.. Indeed, these mappings can be learned as a musician learns to play an instrument, but more meaningful mappings could reduce the users' cognitive load and increase the utility of the application.

The redesigned interface makes use of the FingerSleeve device shown in Figure 3.6, which we built specifically for this purpose. This device has four states (from the two multi-level buttons). However, we did not want to map the four original pinch gestures directly onto these states because combining two states into a multi-level gesture does not always make sense cognitively. We ruled out adding additional buttons to the FingerSleeve because we wanted to maintain the simplicity of the device.

To achieve the extra functionality with the FingerSleeve, we make a logical distinction that did not exist in the previous version of CavePainting, we consider button presses to be different depending on the proximity of the FingerSleeve (worn on the non-dominant hand) to the paintbrush (held in the dominant hand). This distinction provides us with the logical equivalent of eight different button presses. When the FingerSleeve and brush are close to each other, the buttons activate modes that control attributes of the brush. Light pressure on the outer button activates a color picker while firm pressure locks in the current color and applies it to the brush. Light pressure on the inner button begins to change the size of the virtual brush, and firm pressure locks in the size change. When the FingerSleeve is not held close to the brush, the buttons affect more global operations. Light pressure on the inner button activates an extensible menuing system that was unavailable in the previous version of CavePainting. Firm pressure on this button selects items from the menu. Light pressure on the outer button is used to reframe the artwork using a traditional navigation by grabbing metaphor. [60] Firm pressure locks in the current artwork frame.

These reframe operations occur quite regularly. They augment the approach artists often use of stepping back to look at their work from a distance, and they are also used to position the artwork properly for the next strokes to be drawn, a style of bi-manual "artwork framing" chronicled by Fitzmaurice et al. [19] The additional feature in our pop through approach is that the artist can quickly return to the previous framing (both translation and orientation) of the artwork. This often makes sense artistically. Again, consider the traditional analogy of a painter working with a large canvas on an easel. There is almost a constant motion of walking close to the canvas to paint something, backing away to observe how the local changes affect the whole, and the returning to continue work. In CavePainting we get some of this same motion (walking through the space of the Cave) and this is often augmented by grabbing onto the artwork, pushing it away, and rotating it a bit to observe the whole work from various angles. The pop through interaction makes it easy to either lock in this new artwork frame if the artists sees a feature that needs to be edited or quickly return to the original frame to continue working as before.

Scaling the artwork is another important feature for artists. While a reframe operation is active the paint brush button on the dominant hand can be activated to initiate a scaling mode where the scale of the painting is adjusted in proportion to the distance between the two hands. This mode is a more constrained version of that described in Mapes' SmartScene [44] techniques, later adopted by Schkolne [55] and others.

Our experience with introducing new artists to the system and then observing them working with the interfaces suggests that the FingerSleeve device is a better fit ergonomically for the wide range of users we see with CavePainting. We have found that the pop through notion takes some time to get used to and when first getting started, users sometimes forget to pop through and lock in operations such as reframing the artwork. We also found it necessary to provide a visual cue for the distinction between the "near" and "far" modes for the two hands discussed above. But, with this addition, users seem to quickly become comfortable with the overloaded functions of the buttons, we think more so than with the more arbitrary mappings used in the pinch glove interface. This comfort level combined with the additional functionality provided by the pop through style interface makes it a more effective interface design for CavePainting, and we suspect many of the considerations and design decisions here can carry over to other virtual reality applications.

3.3 Discussion

3.3.1 Unusual Use of the Space of the Cave

One of the most interesting contributions from our CavePainting work is the discussion of the different use of the space of the Cave that we see with this tool when contrasting it to more traditional, often scientific, uses of the facility. Much of this experience is relayed in the sections above describing interaction with the Cave space, including the walls of the Cave, and with the painting table.

Over several years of observation and feedback from artists with CavePainting and related systems, this unusual use of Cave space continues to be one of the most interesting themes in our work. CavePainting presents an unusually effective mode of operation with the Cave and suggests a paradigm for Cave-based interaction that can extend beyond artistic tools. More compelling and more effective use of a Cave can be achieved with styles of interaction that promote movement and large, body-scale interaction.

Movement within the Cave

In contrast to desktop-based VR and even many head mounted VR form factors, movement inside a Cave display is relatively unencumbered. This is one of the greatest advantages of the space. Navigation by just walking through the space of the

Cave is a feature of basically all VR applications built for the Cave, but very few applications actually encourage this sort of movement. Artistic applications seem to do a better job of this in general than do scientific ones. [49, 50] Yet, scientific applications almost always provide the motivation for establishing and maintaining extremely costly Cave facilities. In CavePainting, we find the movement-heavy style of interaction extremely compelling. So much so, that it seems if an application is not leveraging this characteristic of the medium as a critical means of exploring and becoming proficient with the data being displayed, then it is not leveraging the power of the Cave medium. Perhaps this type of application is best run outside of a Cave, or perhaps a redesigned version of the application would result in a more effective display within the Cave. In order to create Cave applications which take advantage of this principle of movement, application designers might consider the following:

- Is the whole space of the Cave being used? If all the interaction is contained within a wall, then is a PowerWall a better form factor for the application?
- Are there incentives within the tasks of the application to move around spatially? For example, to see a new area of the data?
- Can incentives to move be imposed on the application without interrupting work flow?
- How does physical space within the Cave relate to virtual space? Creating more correspondences between the virtual and the real may lead to richer movement and interaction.

Large, Body-Scale Interaction

Another defining characteristic of the Cave form factor is its scale. Movements can be made in the Cave that are not possible with many other VR displays. One of the most compelling examples of this from CavePainting is the way that artists use their body as a tool for making more controlled movements. The idea of locking one's elbow in place to create a rigid lever hinged at the shoulder as a way of making a controlled arc or circle is one example. Various other painting gestures that rely on whole arm or body movement on a large scale help to make certain types of strokes in the Cave.

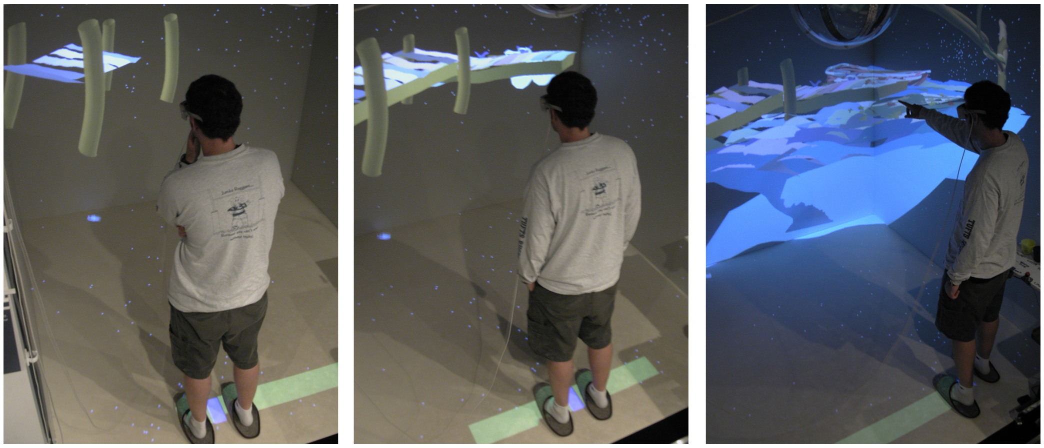

Figure 3.7: A timeline widget on the floor of the Cave controls a visualization of the painting process.

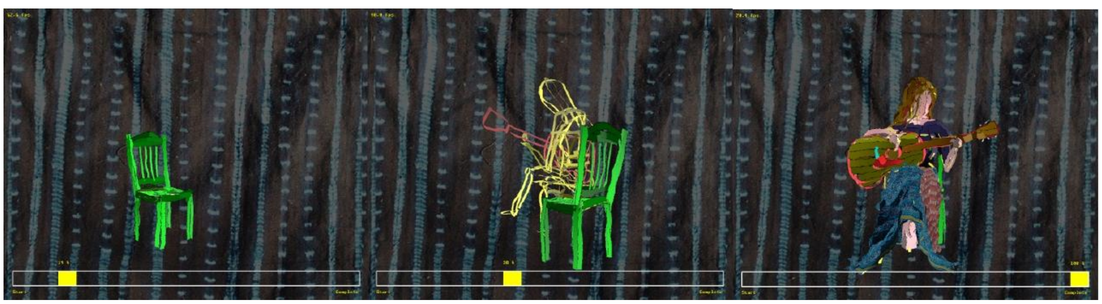

Figure 3.8: Visualization of the painting process for *La Guitarrista Gitana*.

3.3.2 Artistic Importance of the Painting Process

Some of the most interesting feedback that we received from artists was that we should consider the process of working with CavePainting part of the finished artistic result. That is, a CavePainting is an artistic happening. Viewing a CavePainting without some sense of the dynamic movement of the brush and body that created it is only looking at part of the art.

The art world is often extremely interested in the process used to create a work. Unfortunately, this is usually very difficult to determine by looking at a finished work. For example, a skilled painter or art historian can often tell the order in which portions of a painting were created, but it is impossible to completely peel back each brush stroke of a masterpiece to see what lies beneath. There are a few examples of successful artworks in which the artist gives us a rare glimpse into the artistic process. For example, Pablo Picasso's series of lithographs of a bull. [1] In this series, Picasso made 18 prints of 11 different states of his lithograph stone. The first states are detailed realistic representations of a bull. As we look at the next states in order, we see the image progress from an intricate, realistic image of a bull to a abstract line representation. Using this record, printmakers can gain insight into Picasso's process from a technical and intellectual standpoint.

In CavePainting we try to capture the same type of record of the artistic process and then present it in an insightful way. We use a special viewing mode, shown in Figure 3.7. In this mode a timeline widget is displayed on the floor of the Cave. When an observer stands on this widget, his movement along the timeline controls the state of the painting that is displayed. At the far left, we see a blank canvas, as we gradually move to the right, more and more strokes are displayed in the order that they were created by the artists, until we reach the end of the timeline when we see the final result.

This mode of interaction returns to the theme of making use of the space of the Cave and creating a connection between the physical and the virtual. Here, we encourage, in fact require, movement through the space of the Cave as a means of controlling the display, and again, we interact with the physical surface of the Cave in a meaningful way.

This visualization of process provides real insight into the artistic process used in creating this different style of form. An interesting example is shown in Figure 3.8 where we see the development of human figure. Initially a chair is created to define the space for the form. Then, a scaffolding is setup as a way of establishing correct proportion. Finally, the scaffolding is covered up by textured brush strokes that define the form of the final work.

3.4 Results

Our collaborations with artists have pushed this tool in many different artistic directions. Through describing some of the most interesting artistic results in this section, we will also present a contrast a variety of different artistic styles of representation that were developed with the tool.

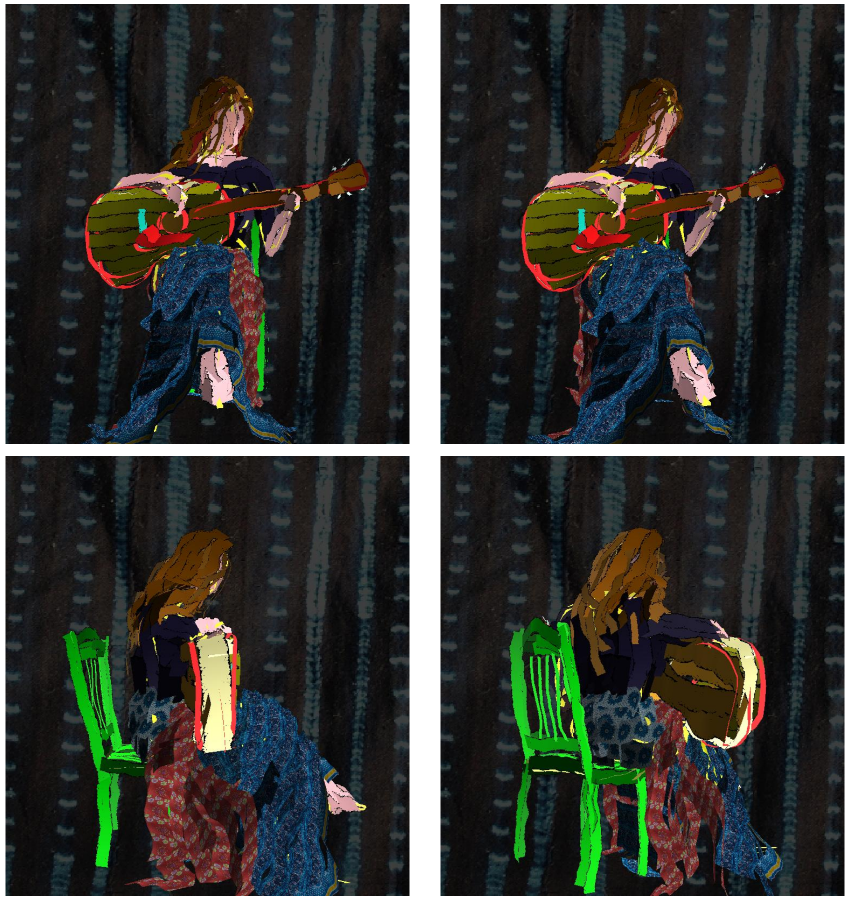

Figure 3.9: Four snapshots taken from around the CavePainting *La Guitarrista Gitana.*

3.4.1 Loose, Impressionistic Forms

These results are of a style that most reflects the original motivation for exploring a medium like CavePaiting: to explore the types of forms that could result from layering of small brushstrokes in 3D. They have a characteristic loose style, where we clearly see individual strokes of paint. This gives them a hand crafted aesthetic that is rare in computer graphics and has been well received in the art community. *La Guitarrista Gitana* shown in Figure 3.9 has been shown in a juried show at SIGGRAPH 2002, and been exhibited as an invited piece in the Boston CyberArts Festival and at the Viento y Agua Gallery in Long Beach CA. *Sailing a Dhow in Tanzania* (Figure 3.10) was also exhibited at the Boston CyberArts Festival.

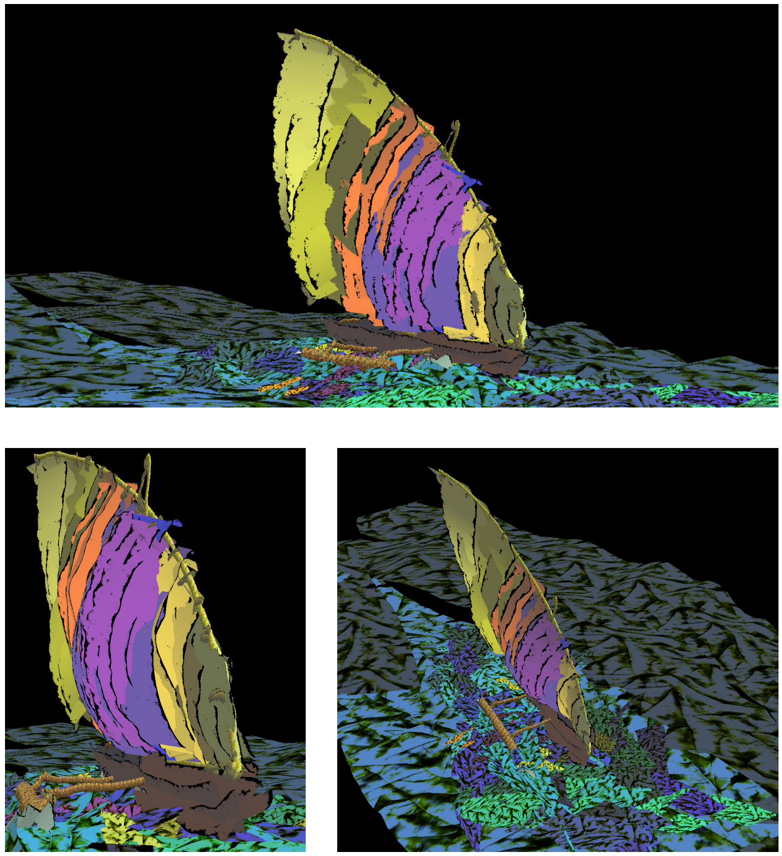

Figure 3.10: Four snapshots taken from around the CavePainting *Sailing a Dhow in Tanzania*.

3.4.2 Solid Free-Form Modeling

Some artists have chosen to approach CavePainting as more of a solid modeling tool. While the results still have a loose aesthetic created with individual brush strokes, they feel more like a solid mesh. Helen Zhu, an illustration student at the Rhode Island School of Design, has worked extensively with this style as part of an independent study project under the direction of Fritz Drury. Zhu worked extensively with a tapered tube shape brush stroke. The lighting effects produced from this primitive contribute to the volumous sense evoked in her CavePaintings, shown in Figures 3.11 and 3.12.

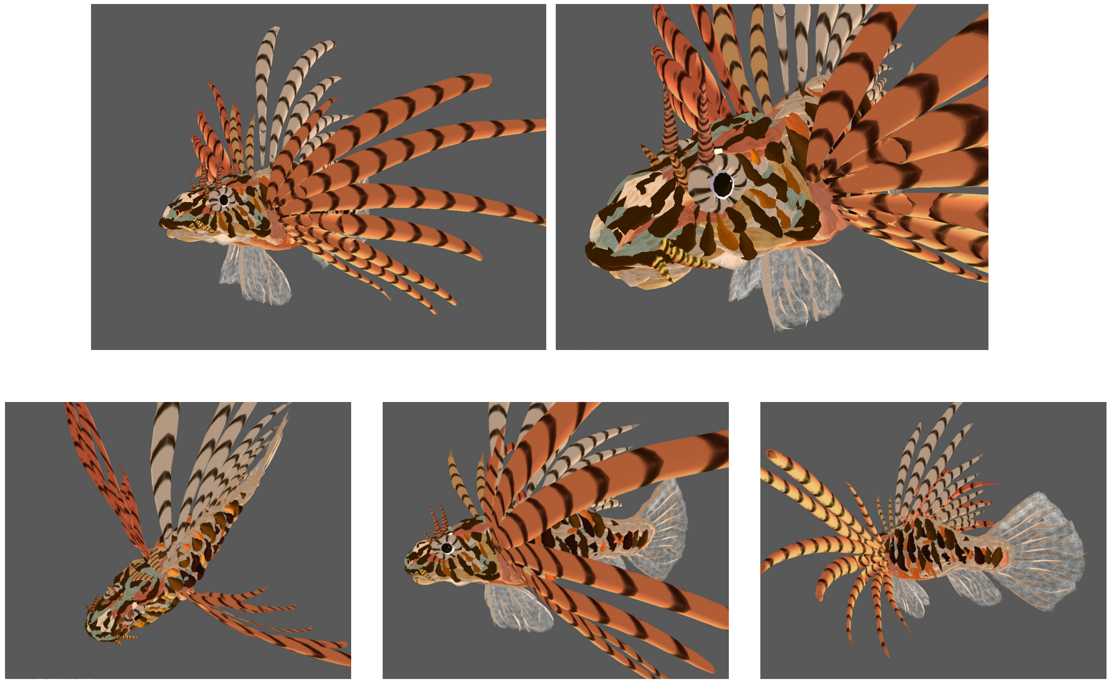

Figure 3.11: CavePainting of a Lion Fish, by Helen Zhu.

3.4.3 Interactive Installations Utilizing the Space of the Cave

Another style of CavePainting investigates more fully the theme of space, movement, and the interplay between physical and virtual reality within the Cave. In *Hiding Spaces* [49, 50], (details in Figure 3.13) CavePainted tree forms twist up from the floor and into the walls of the Cave where they meet carefully constructed imagery textured only the walls and floor of the space. The images on the walls of the Cave are a key component in this interactive installation because they undergo color and composition shifts as the viewer walks within the space of the Cave. The resulting ambiguity in the viewer's perception of space compels him to explore the work. As he ducks under a tree branch to see what lies on the other side, the wall imagery shifts to reveal the image of a doorway, but as he moves to investigate this door, it fades into a dull representation that becomes tangled in the web of foliage surrounding it. In some cases imagery on the walls becomes partially transparent and we notice that the CavePainted tree forms extend beyond the space of the physical Cave into the undefined region beyond its walls. The two versions of the window shown in Figure 3.14 illustrate the types of color shifts used in the wall imagery.

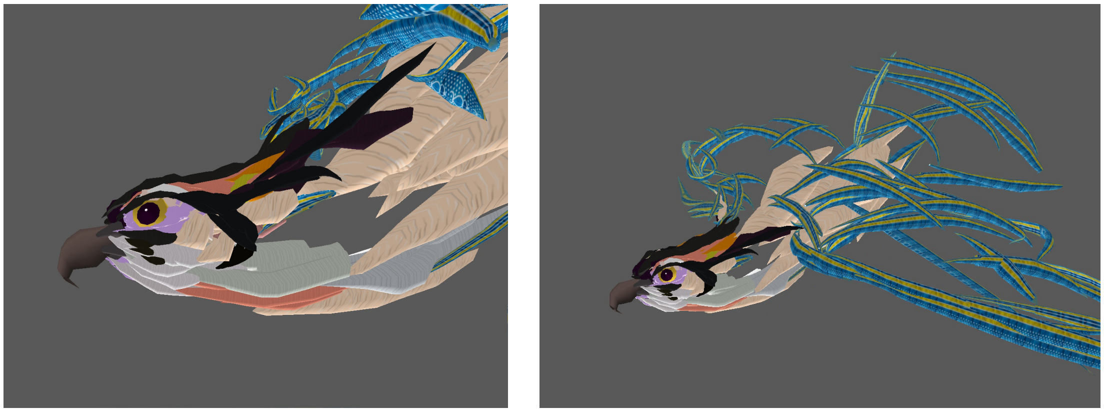

Figure 3.12: CavePainting of a Falcon, by Helen Zhu.

Figure 3.13: In *Hiding Spaces* imagery on the floor and walls of the Cave shifts in response to the user's movement.

Shifts in the imagery are triggered by the viewer's movement within the space. We determine the viewer's location by projecting his head position onto the plane defined by floor of the Cave. We divide this plane into a grid of 2 ft. x 2 ft. squares and lookup the textures to display on the Cave walls based on the square that the viewer currently occupies. When moving from one square to the next the two sets of imagery are blended together to produce a smooth transition.

This work in particular, illustrates to utility of CavePainting as a tool for expression and exploration of artistic concepts. *Hiding Spaces* is a thorough, months long, collaborative artistic investigation into the ambiguity and tension in visual perception that occurs when we display and interact with rich imagery projected directly on the walls of the Cave. Several aspects of this are interesting. First, this idea that something intriguing happens perceptually when we work with the walls of the Cave was born out of exploration done with CavePainting. This speaks to the utility of

Figure 3.14: An example of shifts that occur in the background imagery as the observer walks through the space of the Cave in *Hiding Spaces.*

having good design tools in new mediums like VR. When you sketch, design, and experiment, you often discover something interesting about your medium. Second, once we had discovered this idea, we were able to use CavePainting to thoroughly explore it artistically. We reworked and refined the visuals in this piece again and again to try to produce the most compelling visual result possible. Along the way we learned to direct viewer attention through shifts in the wall imagery and we explored many different artistic techniques for defining and redefining space within the Cave. Third, this work demonstrates an interesting outlet for CavePainted form. Rather than viewing CavePaintings as static models, this work incorporates CavePainted form into an interactive art installation in a style that has been well received in the art world.

In this chapter, we presented the CavePainting system and insights from years of artistic exploration with the tool. We discussed CavePainting's unusual hand-crafted modeling aesthetic and its unusual use of the space of the Cave floor and walls. We present prop-based interaction techniques for use within a Cave. We also present a novel, pop-through buttons device and our experiences with redesigning the CavePainting interface to leverage the expressive power and cognitive benefits of pop-through buttons. Finally, we present results from artistic explorations with the tool and describe how this work has furthered our understanding of Cave-based interaction.

Chapter 4

Artistic Collaboration in Scientific Visualizations for Virtual Reality

4.1 Background and Motivation

In addition to strictly artistic applications of free-form modeling, a second theme which runs in parallel throughout this thesis is enabling artists to contribute to advances in science. We address this challenge through scientific visualization and through providing tools for artists to contribute to this effort. Within the scientific visualization community, a fair amount of research has been conducted recently on visualization techniques inspired by the study of art. [39, 40, 43] Artists are skilled practitioners of conveying meaning through imagery, and these works strive to leverage this skill through study and reproduction of successful artistic technique. The motivation for the work presented in this thesis is similar, but we take a slightly different approach. Rather than learning from existing works of art, we strive to include artists in the process of scientific visualization. In this chapter, we discuss our methodology for this collaborative approach to visualization, and present several research tools that we developed to aid this collaboration.

We present two main contributions in this chapter: 1. A methodology for collaborating with artists in visualization design and discussion of how to foster and teach collaboration using this process. 2. A data-driven software tool for artists, illustrators, and other visual experts without knowledge of computers and programming to

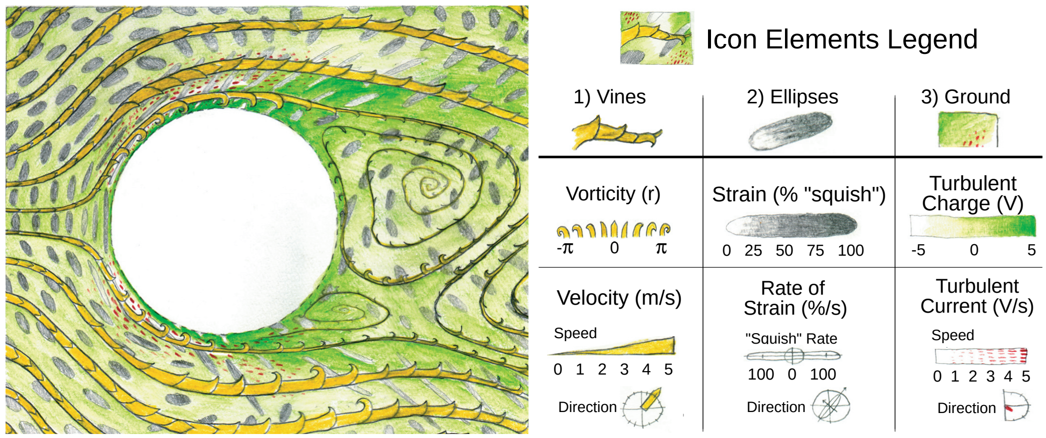

Figure 4.1: An art student's visualization design of 2D steady fluid flow past a cylinder. Courtesy of Deborah Grossberg.

quickly explore the design space of glyph-based multivariate visualizations in VR.

In our experience, artists routinely provide an unique source of visual insight and creativity for tackling difficult visual problems. They are also expertly trained in critiquing and refining visual works, an essential task in the iterative visualization process. Unfortunately, it is currently very difficult to collaborate with artists, illustrators, and designers on VR design problems because there is a lack of appropriate design tools to support them and their role.

The lack of appropriate design tools is particularly evident in visualizations using new technologies, such as virtual reality (VR) or volume rendering. It is difficult for artists to get involved in design in these visual spaces, since, with rare exceptions, one needs to know how to program in order to create within them. Unfortunately, these are also the types of technologies that offer great potential for visualizing many of today's complex datasets. [67] Additionally, they are probably the technologies in which we can most benefit from artistic insight, since guidelines for good visual depiction are far less developed in unconventional visual spaces, such as virtual reality, than in more traditional 2D media.

We begin by describing one of our recent major collaborative efforts, a class on designing virtual reality scientific visualizations that was co-taught with professors and students from Brown's computer science department and from the Rhode Island School of Design (RISD)'s illustration department. The class has been taught

Figure 4.2: Students prepare for a critique of arterial blood flow visualization designs.

twice now, and many of the experiences and conclusions relayed here are the results of preparing and teaching this class. We then present a data driven, artistic visualization design platform which combines CavePainting-style interactions with multi-variate data visualization techniques. Our final contribution in this chapter is a demonstration and presentation of several guidelines for enabling artistic collaboration in VR visualization design derived from our experiences.

4.2 Collaborative Teaching of VR Visualization Technique

Our interdisciplinary visualization class brought artists and computer scientists together to solve visualization problems driven by science. Students worked in teams on visualization and design assignments, following the interdisciplinary Renaissance team model presented by Donna Cox. [11] We began the semester with 2D fluid flow visualization assignments, as in Figure 4.1, and gradually built up to the final projects, which were virtual reality visualizations. In our first version of the class we focused on data describing pulsatile blood flow through a branching coronary artery. In the second version, we worked with a different dataset describing the motion of a flying bat. As we describe our approach and contributions in this chapter, we will present examples from both of these scientific problems.

Although artists rarely work with complex scientific data, they do train to convey information effectively through imagery, given the constraints imposed by their media, employers, or audience. In this abstract sense, normal artistic practice is not such a far cry from typical visualization design tasks. The images in Figure 4.1 show one art student's result from an early visualization design assignment. The students created visualizations and legends that convey eight continuous variables describing a steady, 2D fluid flow in a single picture. This is a very difficult visual problem; in fact, it is still being actively researched in the visualization community. We found that artists were adept at investigating visual problems like this one when we could clearly convey the scientific goals and constraints of the problem.

Artists also excelled in the initial design and conceptualization stages of the scientific research process, often prompting new insights on the part of the scientist team members. Scientists rethought their hypotheses, clarified their experimental goals, and even altered the way they collect data in response to feedback from artists. As Vibeke Sorensen explains in her discussion of the artist's contribution to scientific visualization [63], this role for artists is a departure from the norm: artists are typically thought to be useful only in the last stage of the research process, dissemination. However, this is a limited use of the artist. Our experiences support Sorensen's claims that artists can be involved in many more stages of research, with conceptualization being perhaps the most important with regard to scientific visualization.

Collaboration was sometimes difficult to manage. In early assignments, such as in Figure 4.1, the right tools for the job were colored pencil, oil paint, gouache, watercolors, and Photoshop. In later assignments, the essential tool for the job moved closer and closer to programming. At this point, the art students often had visual insights to offer, but had difficulty conveying them. It was easy for the non-programmers to feel left out of the loop. As Fritz Drury (the RISD illustration professor who co-taught the class) remarked, the programmers are the ones with the ultimate power: they have the final say about what ends up on the screen.

One device that helped us keep artists, computer scientists, and fluid flow researchers on the same page is the critique, a common teaching tool in art classes. All the class work was displayed on a wall, as seen in Figure 4.2, and as a class, we discussed important design lessons in relation to each work. We critiqued the work both from a visual and a scientific standpoint. Visually, we explored color, scale, form, metaphor, and narrative. Scientifically, we learned about the data we were trying

Figure 4.3: Student visualization designs for the artery visualization using traditional artistic media.

to represent and critiqued the work on the basis of how truthfully and completely the science was represented, given the tasks our scientists wished to perform. We have now adopted “crits” into the visualization development process for many of our projects.

We were able to investigate and refine many designs though this critique-based model. Some of the original student designs are reproduced in Figures 4.3 and 4.4. Notice the storyboards in particular. Here, the students are already starting to address one of the main challenges in designing for VR. It is a different medium. It is highly interactive and immersive, these traits are very difficult to portray in traditional artistic mediums. In the design on the bottom right in Figure 4.3 you see a sculpture of a visualization idea for a visualization of blood flow in an artery. We explored mediums like sculpture and video as ways to bridge the gap between design outside of VR and VR-based visualization. All of these approaches certainly have their place in the design pipeline, but in the end we felt that we could accomplish more with a tool that would enable design directly in VR.

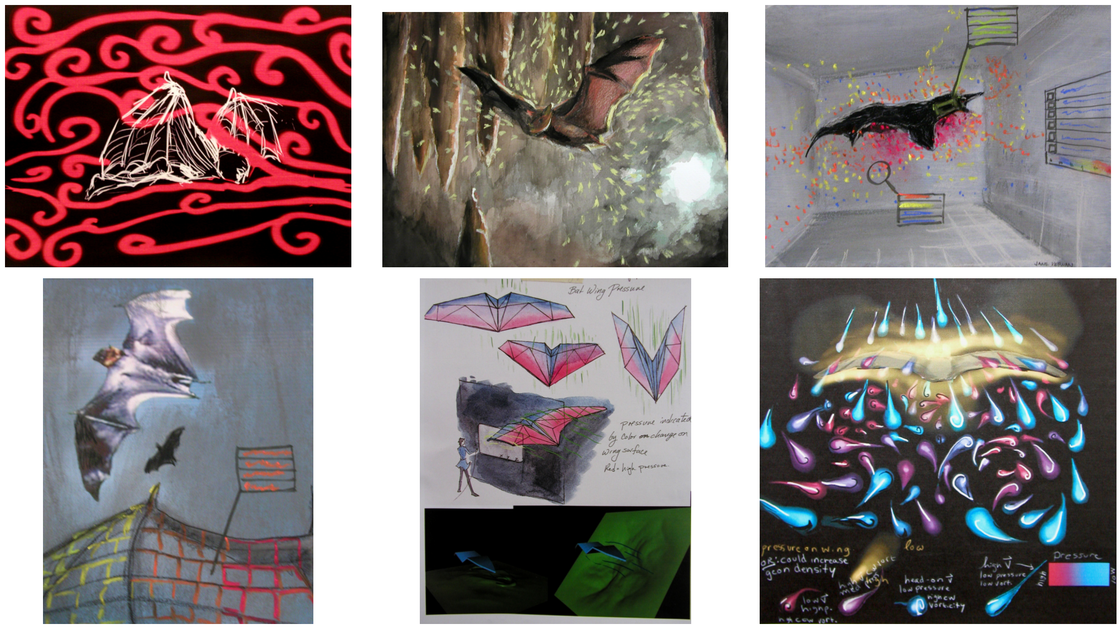

Figure 4.4: Student visualization designs for the bat flight visualization using traditional artistic media.

4.3 Using CavePainting to Design Scientific Visualizations

As we move from 2D visualizations into more complex 3D situations such as virtual reality, collaboration with artists becomes much more difficult to facilitate. The first theme we have derived from our experiences is that visualization design needs to occur within the visualization target medium. This sounds simple, but it has fairly significant ramifications for the visualization media we often use. For example, it is very difficult for anyone, and nearly impossible for an artist who is not a programmer, to create visualizations or simply experiment with design ideas in virtual reality.

A starting approximation for designing within VR is to design with more traditional, often 2D, media as seen in Figures 4.3 and 4.4. The difficulty with this approach is the drastic difference between what we can convey on paper and what we can convey in VR. We use a four-wall Cave VR display environment for much of our research. So much changes when we enter the Cave: scale, interaction, stereo vision, vividness of color, and contrast. Due to the drastic differences in the mediums, it is very difficult to trust or evaluate traditional designs of VR ideas.

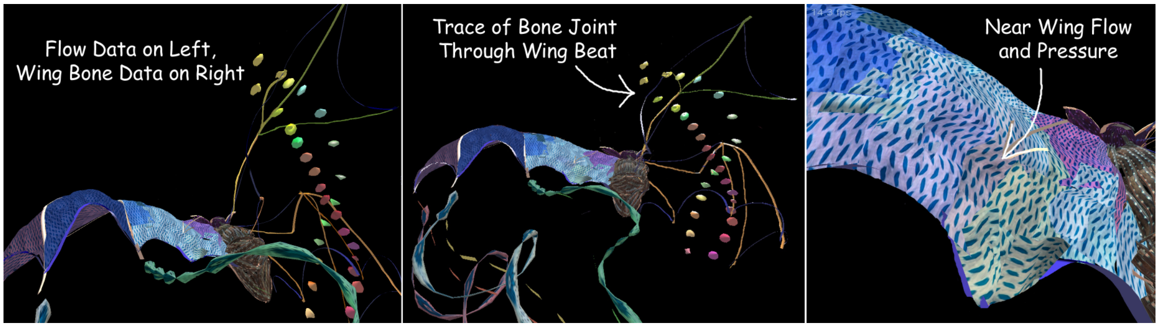

Figure 4.5: A CavePainting visualization design of bat flight data.

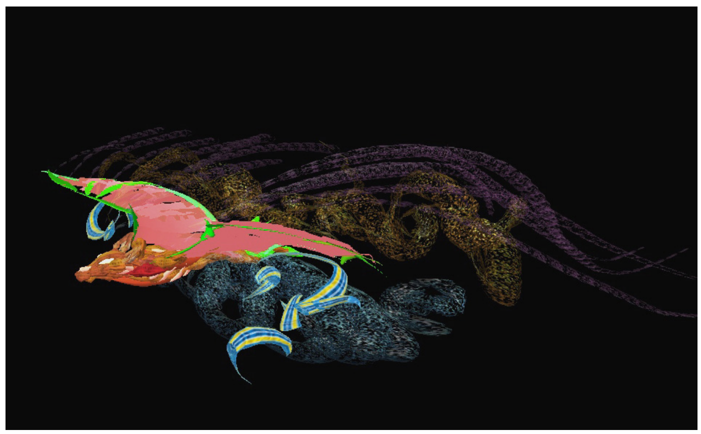

Figure 4.6: A CavePainting visualization design of bat flight data.

Ben Shedd notes a similar dilemma in a comparison between traditional filmmaking and new Imax^{TM}-style filmmaking. [59] In giant screen films, as in VR, the projected images extend beyond our peripheral vision. This significant change has required filmmakers to begin to invent a new visual language and prompted Shedd to call for redefining giant screen filmmaking tools. This is one of the issues Shedd explores in his interdisciplinary class at Princeton.

In VR visualization, we are also defining a new visual language, and we deal with a similar lack of appropriate traditional tools to do the job. When designing traditionally with an eye towards VR, we face the problem that a good 2D design does not necessarily translate into a good 3D, much less VR, design. Further, it is very difficult to evaluate or propose refinements to a design without actually seeing it implemented in the Cave. We lose the power of the critique, which we have found so useful. We need to be able to design and critique within VR.

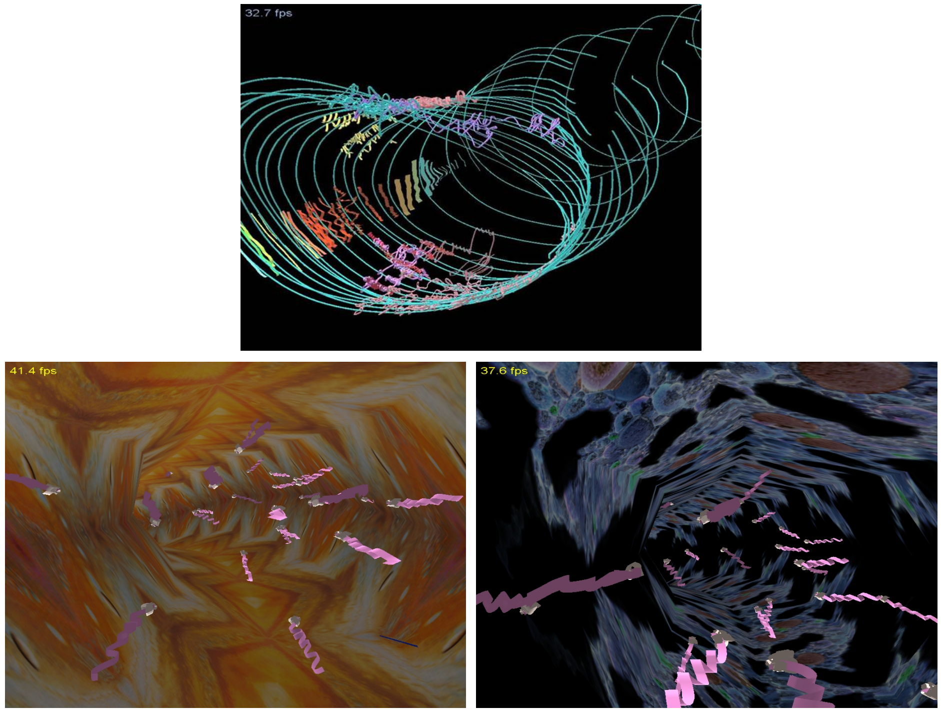

Figure 4.7: CavePainting visualization designs for depicting arterial blood flow.

With this motivation in mind, we began exploring ways to work with artists to design visualizations directly within the Cave. Figures 4.5, 4.6, and 4.7 show snapshots of VR visualization designs created with our system.

The basis for our tools is the CavePainting program [36] described in Chapter 3. Several additions to the original software, including extensive facilities for working with textures were required to enable designs like those described above.

There are several benefits to working directly in the Cave with a tool like CavePainting. The most important is that the design can be easily critiqued and refined with proper attention to the nuances of the target medium. In practice, we have gained valuable insight from these critiques. We have made several alterations to our initial bat visualization designs based on feedback from Dr. Swartz and her collaborators after meeting for critiques in the Cave. During these critiques we have even been able to quickly sketch modifications to designs and discuss them immediately.

Using CavePainting to design visualizations also has the advantage that we can investigate, refine, and converge on a successful visual design at an early stage in the process. With the usual approach of implementing before visual refinement, it might

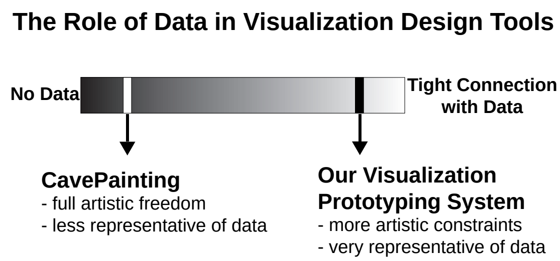

Figure 4.8: Design tools can have a stronger or weaker built-in connection with data. Tools at both ends of the spectrum are useful.

take weeks or months of implementation before we discover our design is flawed from a visual standpoint, and once we notice a problem and brainstorm another design, it could take another few weeks before we are ready to visually critique that one. Thus, particularly in VR, where implementation can be difficult and time consuming, putting visual design decisions in series with implementation can extend the time between iterations on a visual design. Designing directly in VR, on the other hand, lets us converge upon a visually successful design early in the implementation process. We can quickly work through many more iterations of the design because we do not have to wait for them to be implemented before critiquing them in the Cave.

4.4 Data-Driven Artistic Visualization Prototyping

In the CavePainted visualization designs described above, there is no programmed link between the visuals and scientific datasets. Designs such as this lie at one end of the spectrum shown in Figure 4.8. Despite the lack of a low level link with the data, this type of design is extremely useful. The designers have imagined some representative data and sketched it out. The visualization is not far fetched; they have seen previous attempts at bat or artery visualizations and talked with the scientists about their goals. Essentially, they know enough about the structure of the data to paint a typical situation so that we can meet with the scientists and critique the visualization idea in the Cave.

The danger in going too far in the design process without a program-level connection of the visuals to the data is that we may converge upon a design that works

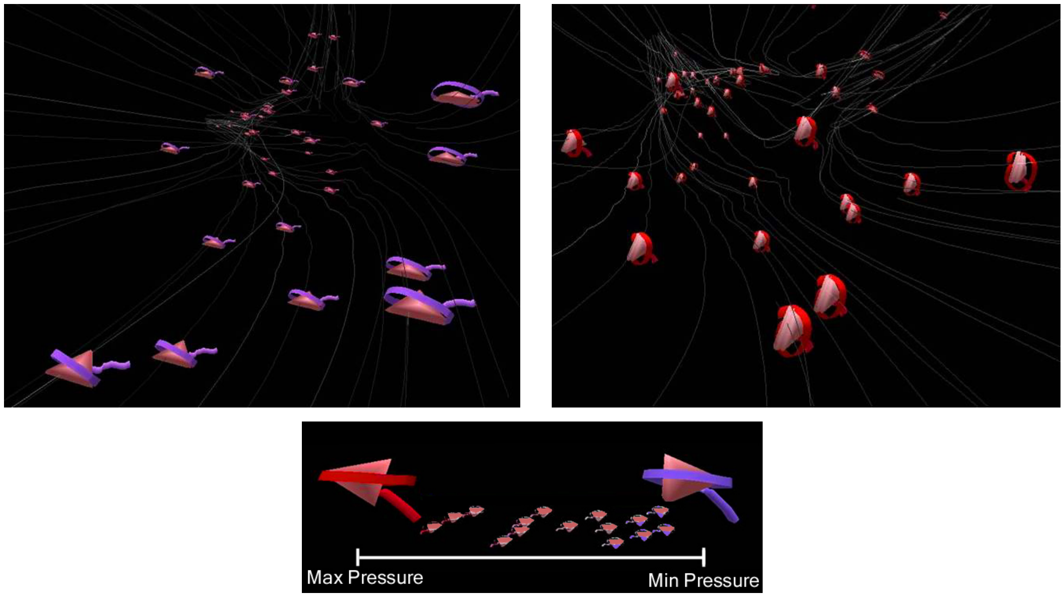

Figure 4.9: Two time steps from a glyph-based, data-driven visualization prototype of arterial blood flow.

well for our perception of the data but not as well with the actual data. In an effort to explore this issue, we built a data-driven design tool that lies at the end of the spectrum where artistic design is closely coupled with scientific datasets.

Our data-driven visualization prototyping system [33] targets glyph or icon-based, time varying, multi-variate, VR visualizations of fluid flow. These visualizations pose particularly challenging design problems. Because the data lies in three dimensions and occluding distant data is big concern in our designs, we cannot rely upon the layering methods described in Kirby et al.'s multivariate 2D visualization work. [39] Rather, we build upon one of the most successful approaches to visualizing multi-variate 3D flow data. We populate the visualization of the dataset with icons that move through the flow and change their appearance in response to the underlying data. We follow the framework setup by Sobel et al. [62] for presenting a synoptic view of the data, where at all points in the animated visualization, a set of icons is visible that represents the global structure of the full flow dataset, and we guarantee that within a reasonable time window an icon passes through each point in the dataset.

The design challenge for this style of visualization is to create a glyph that represents as many variables as possible as clearly as possible. For example, a color map may be applied to the glyph as a way of representing pressure in the flow. When the

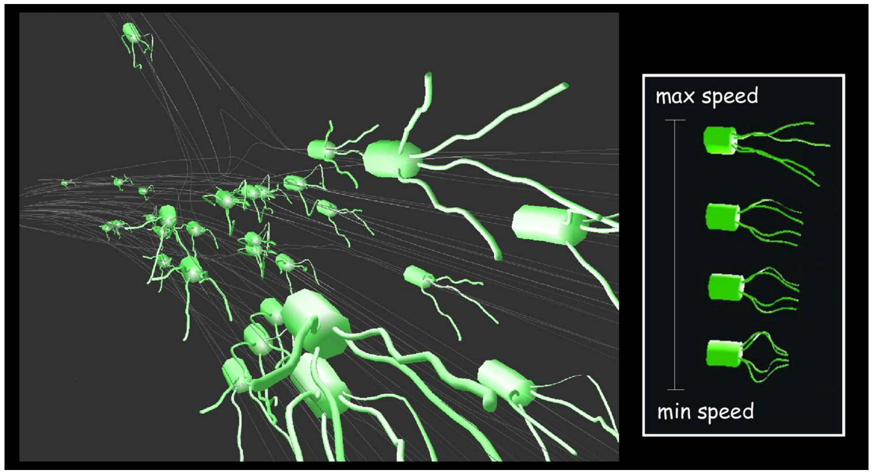

Figure 4.10: A data-driven visualization prototype of arterial blood flow depicted with a squid-like glyph.

pressure is small the glyph may be drawn in blue. As pressure increases, the color shifts toward red in an effort to draw the scientist's attention to the high pressure areas of the flow. With color representing pressure, another aspect of the glyph's appearance, scale for example, must be used to represent the speed of the flow. The icon might be oriented in the flow so that it lines up with the flow direction. Other aspects of its appearance may be linked to shear stress or any number of other flow variables. To create a successful glyph-based multi-variate visualization, care must be taken in the visual design to maximize the readability of the visuals and minimize clutter and distractions.

Our contribution is a design tool that allows artists, illustrators, and other visual experts without knowledge of computers and programming to quickly explore this design space. Again, we build upon CavePainting inspired interfaces. CavePainting interactions are used to draw key frame glyphs. For example, when designing a glyph that changes color in response to pressure changes, an artist might draw one version of the glyph that should appear when pressure is at a minimum and one that should appear when pressure is at a maximum. Copy and paste operators along with stroke editing operations were added to the CavePainting interface to facilitate this style of glyph creation. Once the key frame glyphs have been created, they are added to the appropriate spot on a legend displayed by the system. A legend for a glyph describing pressure change is displayed in the bottom image in Figure 4.9. The two large versions

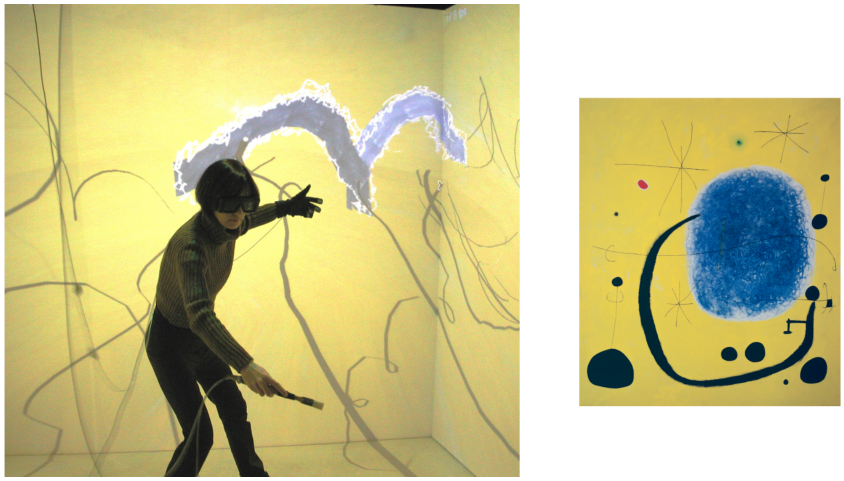

Figure 4.11: The 3D visualization design on the left, also for the bat flight problem, was inspired by the Miró painting “The Gold of the Azure” on the right. “The Gold of the Azure” (C) 2004 Successió Miró / Artists Rights Society (ARS), New York / ADAGP, Paris.

of the icon at either side of the legend are the key frame glyphs that were drawn by the artist and added to the legend. The system automatically determines the difference between the two glyph drawings and computes a morph between them. A preview of this morph is shown in the legend by the smaller icons. Key frame glyphs can be added to multiple legends so that a glyph may respond to more than one variable. When the legends are complete, the artist instructs the system to create a visualization based on the design. Two time steps of this visualization are also displayed in Figure 4.9. Note that this is time varying data and the visualization is animated. Since we are visualizing pulsatile flow in this artery, pressure changes as the heart pumps and that is why we see the glyph responding to pressure changes by changing color at the two time steps that are shown.

Figure 4.10 is a snapshot of a second visualization design for the arterial blood flow dataset. In this design, the squid’s tentacles morph in response to data values. At high speeds, they straighten out and the squid appears quite streamlined. At lower speeds, they flail out to the sides, as the squid assumes a more sluggish posture.

This prototyping tool has been useful in evaluating several different designs for arterial blood flow visualization. Since we are working with time-varying, pulsatile fluid flow, the ability to see the design animated, with glyphs flowing down the artery and changing shape in response to the data, is critical in evaluating the design's success. This would be a difficult display to realize without a program-level link to the flow data.

Despite the success of this approach in achieving these animated visualization designs, we have had difficultly moving beyond these relatively simple cases to the more complex ones required for many of our driving scientific problems.

This experience illustrates the tradeoff that exists in many design systems based on the role they provide for data. Given plenty of preprogrammed connections to data, design tools can produce visualization designs that are so representative of the data that they can be trusted and critiqued as completely accurate visualizations. However, preprogrammed connections to data can be constraining to the artist. For example, in our current implementation of the prototyping tool, the icons must be drawn in a special way in order to establish a solid correspondence for our morphing algorithm. This means that the artist must have this in the front of her mind while working on the design. Creating very complex designs, for example icons that respond to six different variables, can become almost impossible to manage cognitively. Again, these difficult design tasks are the ones our driving scientific problems require and the ones in which we can most benefit from artistic insight. We need to continue to develop intuitive design tools that provide this type of solid connection to the data, but also allow artists to work naturally.

4.5 Enabling Artistic Process for Science in VR

A final theme that has emerged again and again in our collaborations is the need to support continued, evolving work with VR tools. This has been evident in two areas. First, getting started in VR is hard. Often our artists have done several preparatory sketches or studies, like the designs from our classes shown at the beginning of this chapter, before entering the Cave to work. We need to make it easier for them to build on those sketches when they get to the Cave, rather than shutting out the real world and concentrating only on VR. Second, we need to facilitate returning again and again to a design to rework and refine it. The real-world problems with which we anticipate artists will work are sufficiently complex that they will require many design iterations to complete. Tools to facilitate artistic collaboration in visualization need to be accessible to artists in these ways if we want to support artistic involvement in difficult visual problems.

Let us look at CavePainting again as an example of an artistic design tool to see how it can be difficult to get started on a VR design. When the program begins, we walk into the Cave, a dark blank room of projection screen walls. We carry a tracked paintbrush prop and a pair of glasses. Once we put on our stereo glasses, it is too dark to see any paper or other real objects we might have brought in with us. By default, we start with a completely blank canvas and no external inspiration, something designers almost never want to do.

One approach that has been helpful to us is to import our design inspiration, often 2D work such as sketches or paintings, directly into the VR design program. In Figure 4.11, we see a 3D CavePainted design inspired by a Miró painting [21, 16]. One of our designers saw the painting and it prompted an idea for visualizing the bat dataset. We cut out subregions of the Miró and imported them into CavePainting as brushstroke textures. Then our designer was able to work directly with elements of the inspirational imagery to create the 3D design she imagined. This gave her a jump-start on her 3D design and helped her quickly create a coherent design.

The ability to return to a design and refine it again and again is just as important as starting with something in VR. The design task is necessarily an iterative one, with critiques by other designers, implementers, and scientists all playing an important role in refining each iteration. Normally artists refine work in two ways. First, they add additional layers of clarification. In painting, for example, additional layers of paint conceal what lies below. A rough outline of a face can be laid down as a place holder for a much more complex rendering to come later, applied with additional paint layers. Second, they create many studies of an idea, sometimes ending up with a studio full of renderings and re-renderings. At the end of this period the idea is clear enough in the artist's mind that she feels ready to produce a final work.

These approaches are not at all mutually exclusive, however we have difficulty supporting either with our current design tools. In the first case, we can add some additional layers of clarification with the CavePainting system, but this can have the effect of distorting the original form. We are a little closer to supporting the second style of refinement, which amounts to letting an artist quickly reel off many sketches before creating a final work. However, it is unclear how to refer back to several studies while working on a new piece, since each design is usually intended to be viewed in the full space of the Cave. These issues are among the most important to address before working closely with artists on design problems, since they can be very frustrating and limit the amount of real design work that can be accomplished.

In this chapter, we presented a methodology for collaborating with artists in visualization design and discussion of how to foster and teach collaboration using this process. As part of this effort, we described a data-driven software tool for artists, illustrators, and other visual experts without knowledge of computers and programming to quickly explore the design space of glyph-based multivariate visualizations in VR. We demonstrated these results with visualizations of arterial blood flow and bat flight designed by illustration students. We found that artists can fill an important role in the visualization process. Our approach to collaboration is successful, yet we identify several areas of future research that will help in fostering collaboration, including the ability to create more specific and clarified 3D models, the issue we explore in the next chapter.

Chapter 5

Haptic Interfaces for Controlled Free-Form 3D Illustration

5.1 Background and Motivation

In the quest to address more complicated visual subjects with direct, sweeping input-based approaches, the fundamental problem that needs to be solved is providing the artist with enough control over the medium so that he can depict the subject with a sufficient level of clarity. This problem can be broken down into two subareas: 1. Providing sufficient control over defining the shape of the form. 2. Providing a rich enough style of form so that the artist can effectively convey the subject.

The interfaces presented in this chapter are an exploration into using the rich space of haptic-based interaction techniques as a way to explore both aspects of this problem. The two handed approaches to drawing with haptics described here aim to provide artists with more control over the precise shapes of the 3D marks they create. Then, the ability to adjust color value and line width while drawing, as well as the ability to smudge out 3D pigment to suggest surfaces, combine to yield a richer visual style than previously available for this type of modeling, thereby enabling artists to address more complicated and meaningful subjects.

At the time of the writing of this proposal, some of this work is completed and portions are still under development. This discussion represents a specification for tools that we plan to implement as part of a SIGGRAPH 2006 paper submission.

Therefore, we anticipate having much more of a complete description of these techniques along with more meaningful results soon, even by the time of the oral proposal defense.

5.2 Drawing Smooth 3D Curves with Varying Line Weight

The goals for this interface are: 1. to help artists create complex 3D curves that are smooth, i.e. have less jitter and bumpiness than what we often see when inputting a curve by drawing a path directly in the air with a 3D tracker, and 2. to provide artists with a natural way to vary the line weight of a mark as they are creating it. When drawing with pencil artists continually vary the pressure of the pencil against the paper to control the weight of the mark. They use this technique to help convey lighting and form, and also to establish a style and feeling in the drawing. When we move a tracker through the air to input a 3D curve, as is commonly done in free-form modeling tools like CavePainting, there is no paper to push against. The richness of this physical interaction is absent, and with it the ability to intuitively control line weight. Our interface attempts to restore some of this physical interaction for artists working with free-form modeling with the hope that providing more artistic control over the medium will allow artists to explore much more complicated and valuable subjects.

Our system uses a fishtank (desktop-based) virtual reality setup as shown in Figure 5.2 with two Polhemus magnetic trackers, one tracking the artist's head and one tracking his non-dominant hand. The stylus of a SensAble Phantom force feedback device is held in the dominant hand.

The two-handed drawing approach that we use is inspired by tape drawing and recent work on digital implementations of the technique [26]. Traditional (non-digital) tape drawing has been used for years by car designers to make life size drawings of cars, as depicted in Figure 5.1. Both hands and a roll of thin black tape are used to draw. One hand pushes the tape against the wall to advance the curve being drawn, while the other pulls the tape taught in a straight line, essentially specifying the tangent for the next part of the curve to be drawn. Both hands must move together

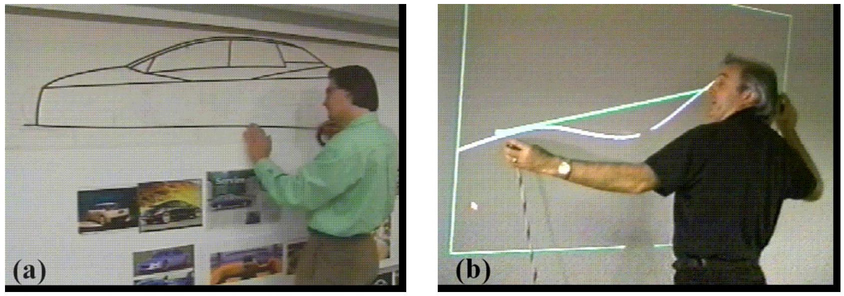

Figure 5.1: Traditional and 2D digital versions of tape drawing.

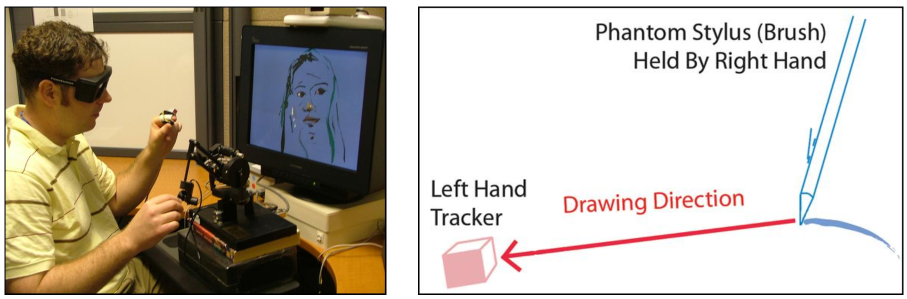

Figure 5.2: The 3D, haptic drawing interface setup.

to draw a curve, which takes some practice to control, but the great advantage of this drawing technique is that it produces very, very smooth curves.

Our interface implements a 3D, digital version of tape drawing. The interface is shown in Figure 5.2. The magnetic tracker on the non-dominant hand is used to specify the direction of drawing. The dominant hand holds the stylus of the Phantom device. When the Phantom's stylus button is pressed and held, drawing begins. From that point on, the drawing direction is constrained (through force feedback) to move toward the non-dominant hand. The force to keep the Phantom drawing along the correct line is modeled as a damped spring force that inhibits movement in all directions except toward the non-dominant hand. In our implementation of 3D tape drawing, the non-dominant hand can move anywhere, so lines can be drawn in any direction in 3-space.

As the artist draws, he can push against the force applied by the Phantom device to control the line weight of the mark. Unlike flat paper that we use in the "real" world, our virtual paper effect is 3D. One way to think of it is as if the Phantom's stylus is stuck inside a very thin tube of paper, as illustrated in Figure 5.3, where the tube is

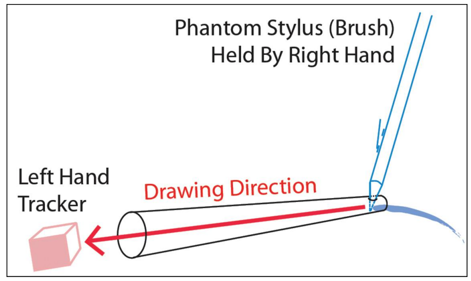

Figure 5.3: Haptic forces can be thought of as constraining the stylus tip to stay stuck inside a thin tube of virtual paper.

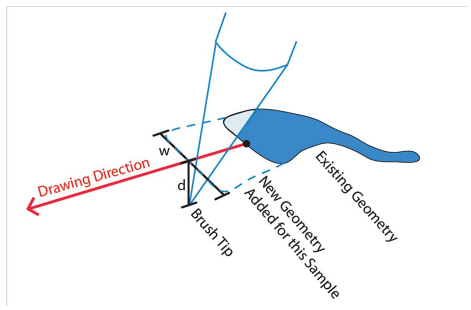

Figure 5.4: Line weight is adjusted in proportion to the distance that the artist forces the tip of the phantom's stylus to move off line formed by the drawing direction.

oriented such that its axis is aligned with the current drawing direction specified by the position of the non-dominant hand. Note that this direction is constantly changing as the artist moves his non-dominant hand. Forces are applied to the Phantom to keep the stylus stuck inside this thin tube, but allow it to slide easily down the length of the tube. When the artist pushes against the side walls of the tube the line weight of the mark is adjusted in proportion to the amount of force he applies, which is captured by measuring the distance, d shown in Figure 5.4, that he pushes the brush tip off the line established by the drawing direction.

Figures 5.5 and 5.6 are snapshots of a 3D illustration of a bat in flight that was created completely with this interface. In Figure 5.5 notice the smooth curving characteristic of the bones. These smooth curves would be impossible to draw with such control with a freehand interface either in the Cave or Fishtank environments. Within the highlighted region in Figure 5.6 we can see an artistic use of the varying line weight feature. Notice how the bones increase in thickness at the ends to visually indicate joints in the skeleton. Increasing the level of control and ability to depict

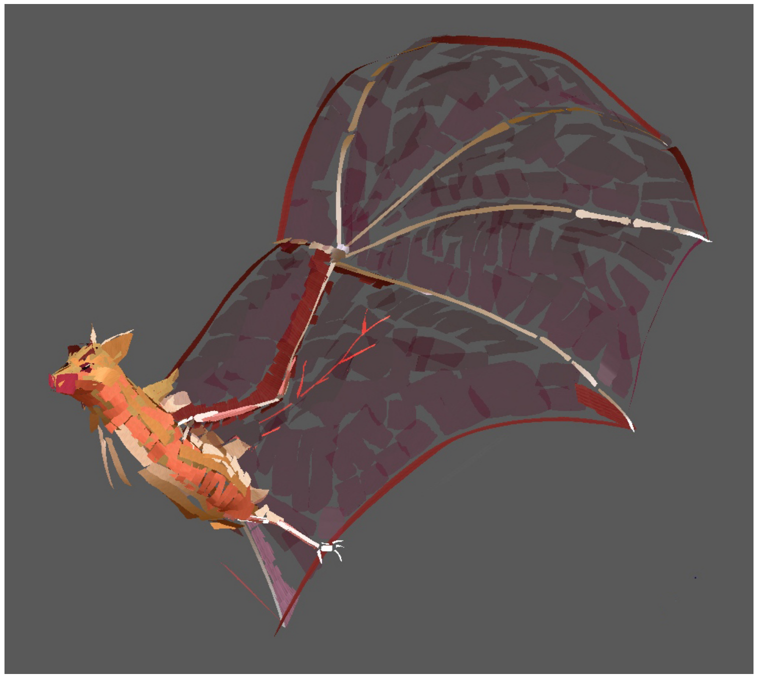

Figure 5.5: Illustration of a bat wing during flight.

specific form like this is our goal as we try to enable artists to address more challenging subjects with these tools.

The small working space of our Phantom device is one of the main limitations of our implementation. Our version of tape drawing is on a scale that is at least ten times smaller than the scale used in car design. Interestingly, even at this small scale, the technique produces very smooth curves, but two handed movements do feel constrained in the small working space.

Tape drawing is the real world technique that provides the motivation for this work, but tape drawing is really a 2D technique. We've discovered at least one real world related approach that is truly 3D, and we expect that many others of this form exist. It turns out that model airplane builders fasten parts on their models by pulling out some sort of stringy adhesive and then ironing it down onto the airplanes. This is a two hand task and often the paths the string of adhesive must follow along the plane are fairly complicated 3D shapes. Our ability to perform real "physical" tasks like this supports our notion that with practice, artists are able to master haptically aided two-handed drawing techniques like this one.

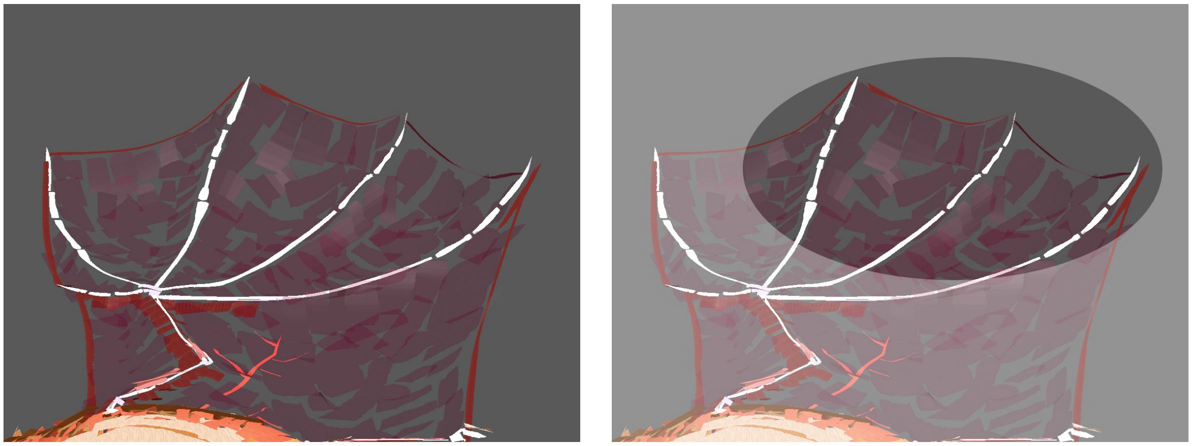

Figure 5.6: Detail of smooth bone shapes and effect of line weight.

5.3 Extension for Charcoal Style Drawing

Drawing with charcoal is an important artistic medium and an essential part of many artists' process. Interestingly, this holds true even when the artists intend to produce a painting, a sculpture, a car, or some other work where the final result contains absolutely no charcoal. Charcoal is so important because it has an amazing ability to be both quick and expressive. As such, it is used, almost exclusively, in design schools to teach drawing and illustration, and often used in early conceptual design sketches.

Charcoal is so expressive because it is a very rich medium. By simply varying pressure of the charcoal against the paper and adjusting the angle at which the charcoal is held against the paper, an amazing variety in the style of line can be produced, some of which are shown in Figure 5.7.

With the two handed, haptic drawing interface described up to this point, we have introduced the ability to vary "line weight" while drawing a mark, but as the examples in Figure 5.7 illustrate, with a medium as rich as charcoal on paper, line weight is not so simple, it can have many visual realizations. In an attempt to explore some of these and capture more expressive power in the tradition of charcoal drawing, we developed the following extension to the basic drawing interface.

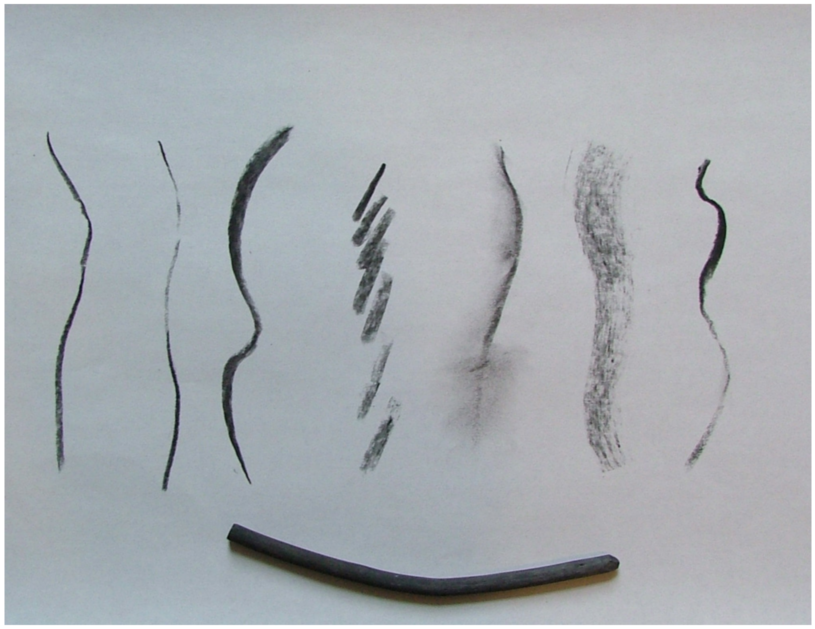

Figure 5.7: A piece of vine charcoal (at the bottom) and example styles of line that are easy to produce with real charcoal.

5.3.1 Drawing Interaction

We model the piece of charcoal that the artist holds as a cylinder. This is controlled via the phantom stylus; there is a one to one correspondence between movement (position and orientation) of the phantom and the charcoal cylinder. In contrast to the previous approach, in this technique we also work with the notion of a virtual piece of paper that is continuously repositioned and oriented in space.

The paper plane is repositioned each frame as the charcoal and the left hand move around based on the following constraints: 1. the tip of the charcoal stylus must lie within the plane. 2. the hand tracker must lie within the plane. and 3. the orientation of the plane is linked to the orientation of the hand tracker. In the images in Figure 5.8, the blue cylinder on the right is the piece of charcoal. The drawing plane is indicated by the grid drawn in space. In the image on the right the dark gray form is the mark swept out by the drawing interaction, notice the complexity of the form achieved with this simple interaction.

Drawing is constrained to use the 3D tape drawing metaphor described above, so the drawing direction is always constrained to be toward the left hand. Instead of

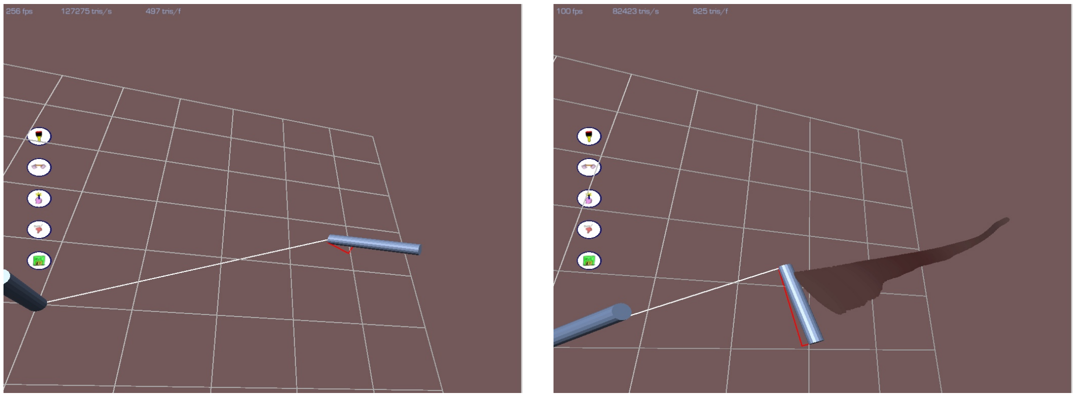

Figure 5.8: Charcoal drawing interface.

using the orientation of the stylus to determine the normal for the geometry produced, the normal is specified by the drawing plane. This allows the artist to hold the stylus at any angle desired while contacting the paper plane. As the stylus is pushed against the paper, we compute the intersection of the cylinder with a plane just slightly offset above the drawing plane. The polygonal intersection computed specifies the shape of the geometry produced as the stylus moves. The stylus can be held on its side relative to the plane to produce a thick swath of pigment or held on its end to produce a thinner mark.

One of the key features that this interface provides is the ability to separate control of color value and line thickness, two aspects of what artists traditionally refer to as “line weight”. The thickness of the line is specified by the orientation of the stylus as it contacts the drawing plane. The color on the other hand is varied based on the pressure applied by the artist against the haptic force produced by the Phantom. This mirrors the characteristics of real world charcoal drawing wherein the charcoal can be laid on its side to produce a thick line, but within this orientation of the tool the artist can drag the charcoal lightly across the paper, or push heavily against the paper to produce a much darker mark. As in the physical interaction, the orientation and pressure in our interface can be adjusted independently and continuously as a mark is being drawn, yielding a very rich space of possible forms for the mark.

5.3.2 Wearing Down the Charcoal

As physical charcoal begins to be worn down, interesting shapes are produced on the tip. This constant wearing down of the physical tool can be a problem. When the charcoal gets too small, artists need to switch to a new piece. Additionally, some shapes that develop are better than others for particular types of lines. Wearing down the tip into a nice sharp point for example helps it to be useful for drawing very thin lines. Wearing it down into a wedge shape makes it useful for lines that vary between thin and thick widths.

Artists can describe this physical property of charcoal as either a benefit or a nuisance, but whatever one's opinion, there is no argument about it being characterized as a rich aspect of the interaction. Strange shapes that develop can lead to new styles of line almost by accident. This type of richness and accidental discovery is something that artists miss when they move to a virtual medium. We believe that finding ways to include richness like this in virtual mediums is a worthy research direction. Thus, despite our sense that we may find that charcoal that never wears out is "better" than real charcoal, we decided to investigate virtual charcoal that wears down with use with the hope that the added richness that it might add to the interaction would inspire new thought, new styles of lines, and experimentation.

5.4 Redraw-Style Editing with a Pulling Metaphor

In order to back up and redraw sections of a mark as it is being drawn, we implement a haptic enabled pulling interaction whereby an artist can pull against the force feedback with enough force to "pop" out of the tape drawing mode and into an editing mode. Once in editing mode, the artist tugs backwards along the mark that was just drawn to wipe away the virtual pigment. A faint trace of the original line is left displayed as a guide when the artist resumes normal drawing and the guideline is removed when the mark is complete.

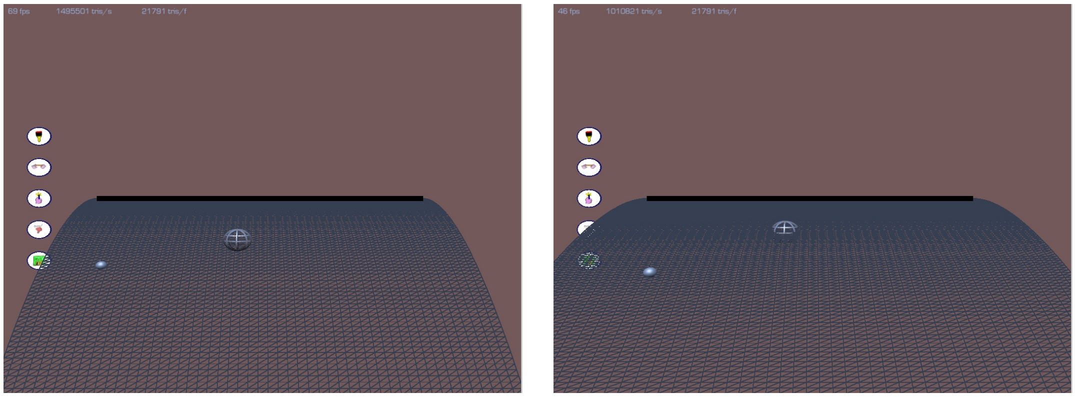

Figure 5.9: Smudging surface prototype implementation with bending controlled by the position of the non-dominant hand.

5.5 Haptic Smudging of Point-Based Pigment

This is an important technique because it adds significantly to the visual aesthetic that can be produced with these tools. When these models are viewed in VR the 3D shape is readily apparent. Just a small line can give all the indication that is needed to illustrate the pose of an arm or a head. Likewise, a small indication of a surface is all that is needed to suggest the sense that a larger surface exists. Establishing this suggestion is the primary purpose for smudging. We want to be able to take a 3D line and smear it around just a bit, sometimes with just a hint of a curve to the surface added. In this way we allow the viewers mind the fill in the details of a larger surface by making just a few marks that serve to suggest a larger form. Smudging is a perfect technique to use for this purpose because it creates a visual ramp that leads the eye from the solid form of a line down to empty space that the eye can easily interpret as a continuation of a suggested surface.

The following description represents our current conception of how the smudging algorithm will work. We have some initial prototypes of this idea, but the full implementation is not yet complete.

Smudging begins by selecting a ribbon-style mark in the model. Think of the ribbon as a small segment of a larger ruled surface S specified by the base curve $b(t)$, which is the path the brush took through space to create it, and rulings $r(t)$, specified by the orientation of the flat brush as it was swept through space. When smudging begins, the only visible portion of this ruled surface is the section that was defined and drawn when the mark was created. The goal for the smudging is to extend this visible portion outward along the rulings to specify a wider mark. In addition, we want to bend the rulings slightly as we smudge to suggest a curving surface.

When in smudging mode haptic forces are used to constrain the brush to lie within S or S', as defined below. To create a smudged out region, the artist holds down the brush button, moves over the original ribbon form to pick up some pigment and then slides the brush over this surface to smear out the pigment. While this is happening the non-dominant hand is controlling the bending of the rulings. The non-dominant hand is held in space somewhere close to the extended rulings. We find the closest ruling to the location of the hand and compute a coefficient α for the quadratic $y = \alpha x^2$ which describes the bending of the ruling. We calculate α such that the ruling bends to intersect the location of the hand. Two states of such a surface are shown in Figure 5.9, created with our prototype implementation. A similar bend is applied to neighboring rulings to create the smooth surface S'. As pigment is applied to portions of S' the newly defined portions of the surface become increasingly fixed in place. Spots on the surface become more and more fixed each time the smudging brush moves over them. As they become fixed, they respond less and less to the bending controlled by the non-dominant hand.

5.6 Discussion and Future Work

When using the system one obvious difference between the haptic feedback produced and the sensation artists are used to with real charcoal is the fact that all the forces are passed through the tip of the stylus. In addition, force feedback is given through position only, there is no rotational component to the force. While our goal was not to make a realistic force model of the interaction of charcoal with paper, the fact that many of these charcoal-based techniques rely so heavily on rotational components and a sense of feeling along the length of the charcoal tool, makes future experimentation with more complete force models and devices an interesting avenue for future work. This could take several forms. One would be a richer contact model, in the style of UNC's DAB system. Another approach or way to augment this approach could be the use of a six degree of freedom haptic device such as SensAble's Phantom 6DOF.

With this device artists would be able to feel the sensation of the full piece of charcoal as it touches the paper, not just the tip. This would certainly mimic the physical world more closely.

In this chapter, we presented a series of tools that enable artists to make more clarified and expressive models within a framework of a direct, 3D approach to inputting form. We use a two-handed drawing technique derived from tape drawing, an approach commonly used in automotive design, and we make it practical to extend this real world technique to a 3D, computer-based one by including haptic feedback to guide the artist. Haptic feedback is also used as a means of enabling more expressive line qualities. We presented several interfaces to achieve expressive qualities of line motivated by traditional charcoal drawing.

Chapter 6

Applications in Artistic Anatomy

6.1 Background and Motivation

The contribution of this chapter is to present the insight that comes from intensive use of our modeling tools and critiques with an artistic expert and to present modeling results of subjects in artistic anatomy. The works presented here were created with both the CavePainting system and the haptic-based tools described in the previous chapter, and they represent difficult modeling subjects that require a high level of expressiveness from a modeling tool. The aesthetic along with the abilities and shortcomings of the tools to convey complex forms like these is the main subject of this chapter. We are interested in learning how best to work with the aesthetic our tools provide and what we can learn about depiction in virtual reality through study and expert visual analysis of artwork generated with our modeling techniques.

Chapter 8 is related in that it presents an evaluation of our work, but here the evaluation is from the perspective of an outside artistic expert who critiques the success of the medium as a means of representing anatomical forms. These insights come from an independent study project I organized with Fritz Drury of RISD's Illustration department. For this study, I wanted to ignore the urge to program as a computer scientist, and instead use these tools in a strictly artistic sense to address driving artistic problems. The goal was to critique the work with Drury's help and arrive at new artistic insights through focusing on a specific subject and trying to advance our ability to render it over the course of several months.

Key components of our methodology are: 1. We focused on depicting human anatomy, which is a challenging artistic subject that demands a level of clarity that is difficult to achieve with computer-based tools. 2. We met regularly throughout the project to evaluate the success of the work I was producing and revise the direction of inquiry based on the outcomes of these critiques. 3. Our approach of assigning "descriptive" drawing assignments together with "dynamic" drawing assignments, based on the format of Drury's advanced Artistic Anatomy drawing course, helped us to work with a purpose in mind and provided a structure for our critiques. 4. A final critique at the end of the study served as a means of wrapping individual insights into more general principles of depiction for the medium and thoughts on useful future directions for our artistic tools.

Our typical structure during this investigation was to meet biweekly in the Cave or the Fishtank VR room to critique new work. Often the assignments I worked on took the form of a pair of studies, as seen in the example of two torsos later in this chapter. The first work is intended to be "descriptive." In our case, this usually meant drawing the bones, tendons, and/or muscles that were responsible for placing a figure in a particular pose. The second work is intended to be a "dynamic" representation of the body in that pose. This line of study follows the format Drury uses regularly in his Artistic Anatomy drawing class at RISD. The descriptive drawing typically informs the dynamic one. In a successful dynamic drawing we can find evidence in the marks drawn of the anatomical structure that is typically represented in detail in the companion descriptive drawing. This structure helped to guide our critiques. We often evaluated the success of a work based on how well it represented a body in motion, with appropriate active muscle groups represented clearly. An interesting parallel can be drawn here between this approach and that of our applications in scientific visualization. In both situations, the work is created with a purpose in mind, and in both situations the success of the work can be evaluated with respect to this purpose. Through the evaluation of the artistic pieces in this chapter, we arrived at insights about appropriate use of the tool and about the nature of the tool itself.

6.2 Formalization of Insights Obtained Through Expert Critique

Several major themes arise from critiques of these works. We state them briefly here, and then describe them in more detail by referencing specific examples in the works and critiques recounted below.

1. Sculpting vs. suggesting form: If many, tightly packed brush strokes are used to create a form, these tools can produce results that look like familiar computer graphics meshes of objects. In this sense, the use of the tool is like sculpting, and the resulting form is a solid object. For some applications these sculpted forms feel more pleasing and comfortable, they are certainly more familiar to us. However, an interesting new aesthetic can be produced when the form is simply suggested, often by working with brush strokes that indicate a surface but leave gaps and holes in the surface for the viewer's eye to fill in. When used correctly, this aesthetic can be exciting artistically.

2. Economy of line and occlusion: If we embrace the idea of suggesting form with few lines as described above, then it turns out we can accurately specify some highly complex shapes with minimal use of line, especially when the drawings are viewed in stereo. But, if we go too far down this path, we end up with occlusion problems. While our minds are happy to imagine a surface between two curves in space, this illusion can be broken when we can see through that surface to objects that lie behind it. If we can overcome this occlusion problem in our software, for example through view-dependent rendering, we may have a much more powerful new aesthetic with which to work when we consciously work with minimal use of line.

3. Use of color together with lighting: In traditional drawing or painting, color often functions to indicate lighting on an object which helps to define the object's form. In computer graphics this type of shading is done as a routine part of the graphics pipeline. The question with respect to these tools is when to rely on CG shading and when to rely on adjusting color by hand. There does seem to be a role for working with color here, but successful use of color probably implies a much more limited role than we imagined before this investigation.

4. Stereo verses 2D viewing: The faintest hint of the tip of an elbow can be all we

need to suggest the pose of the arm when drawn in stereo. On the other hand, when many of these works are taken out of stereo and projected onto a page, they look remarkably flat. So much so, that often the essence of the piece is really lost. When we critique these 2D projections, we find that many fundamental drawing principles for indicating depth are ignored in the rendering which gives rise to several ideas for incorporating art-based rendering techniques into the drawing of these forms.

5. Differences between the Cave and Fishtank: The fishtank drawing interface with haptics tended to lend itself to working with fewer marks per model and trying to make each mark the right one. The Cave tended to lend itself to more gestural drawings with a sense of feeling and movement. The fishtank was better for exactness, but proportion and incorrect proportion were easier to identify in the Cave. Initial drawing attempts at the fishtank tended to come out flatter than intended. Often guide lines had to be drawn and checked from multiple viewpoints as a strategy for establishing correct working proportions.

6.2.1 Critiques organized by artwork

Descriptive Torso, Figure 6.1: The strongest aspect of this work is the pelvis, which is defined fairly well. The pelvis is one of the most complicated shapes to draw. Drury tells his students that if they want to become better drawers, a good drill is to draw as many pelvises as possible. The weak aspect of this work is the use of the blue forms. These were intended to function as shading information. I applied them as a painter would use color to suggest areas of the form that are in shadow. The conflict in this situation is that they create an ambiguity. Are these marks representing form or lighting? They seem to do a poor job of indicating lighting. One of our first major observations from this work follows out of this and the next critique. The role for color here is tricky. The ribs are flat forms compared to the pelvis, Drury suggested using more of a tube form or several ribbons close together to indicate more of a volume in the ribs.

Dynamic Torso, Figure 6.2: This is a loose, gestural work, it suggests a sense of twisting movement. This is a bit less of the body than what Drury normally assigns for this assignment, and that fact motivates the decision to investigate other parts of the body next. There is something about torsos that looks beautiful and magical.

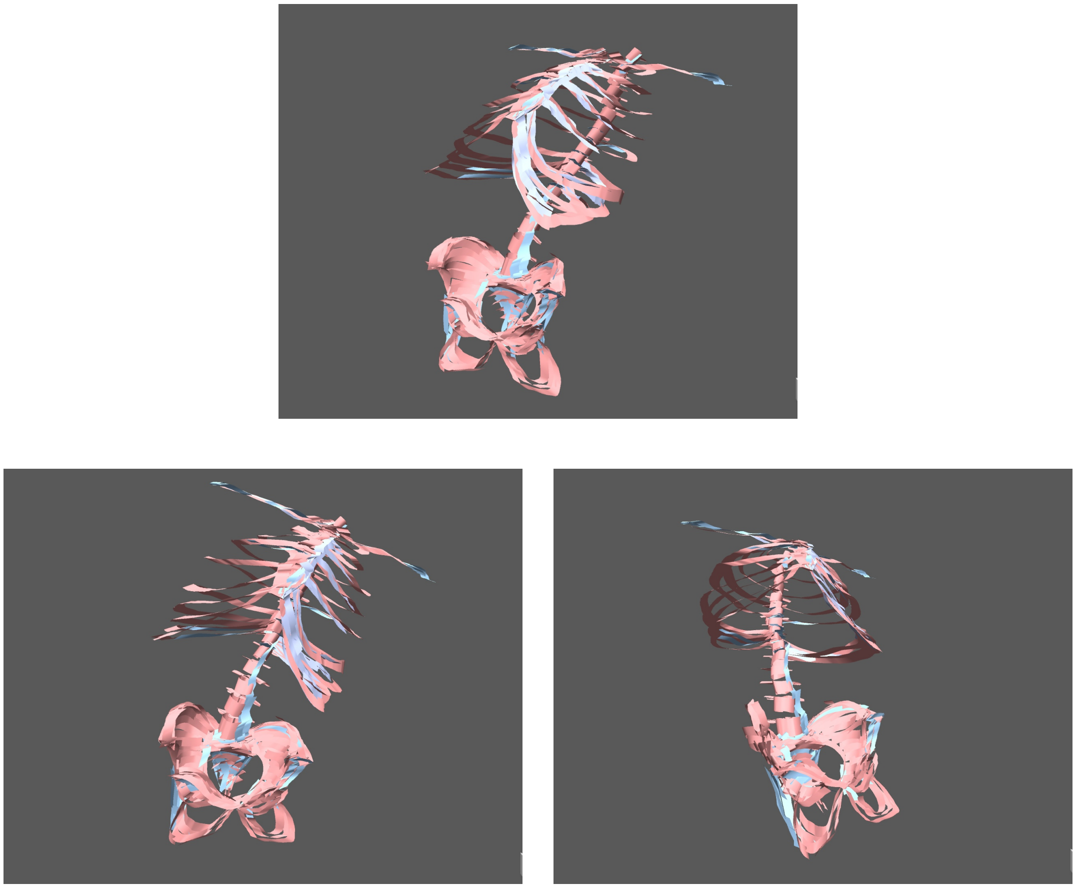

Figure 6.1: A descriptive depiction of a torso.

The challenge is to get the same sense of gesture and movement with an arm, or with a whole body. One of the most elegant and provocative parts of this piece is the slight indication of the head direction achieved with just a few strokes. When seen in stereo this part of the form works exceptionally well. It is slight, but without it the pose would not be complete. This speaks to the potential for depicting forms with minimal use of line that exists in stereo viewing, an area worth exploring in our visualization techniques where visual space is always at a premium.

The holes in the form are of course one of the most distinguishing characteristics of this model. Drury agrees that this technique works well, it is interesting (not a problem) to have these areas left blank.

The proportion in the model is exceptional. The use of color is the weakest aspect here. Drury's suggestion is to rely much more heavily upon the built-in computer

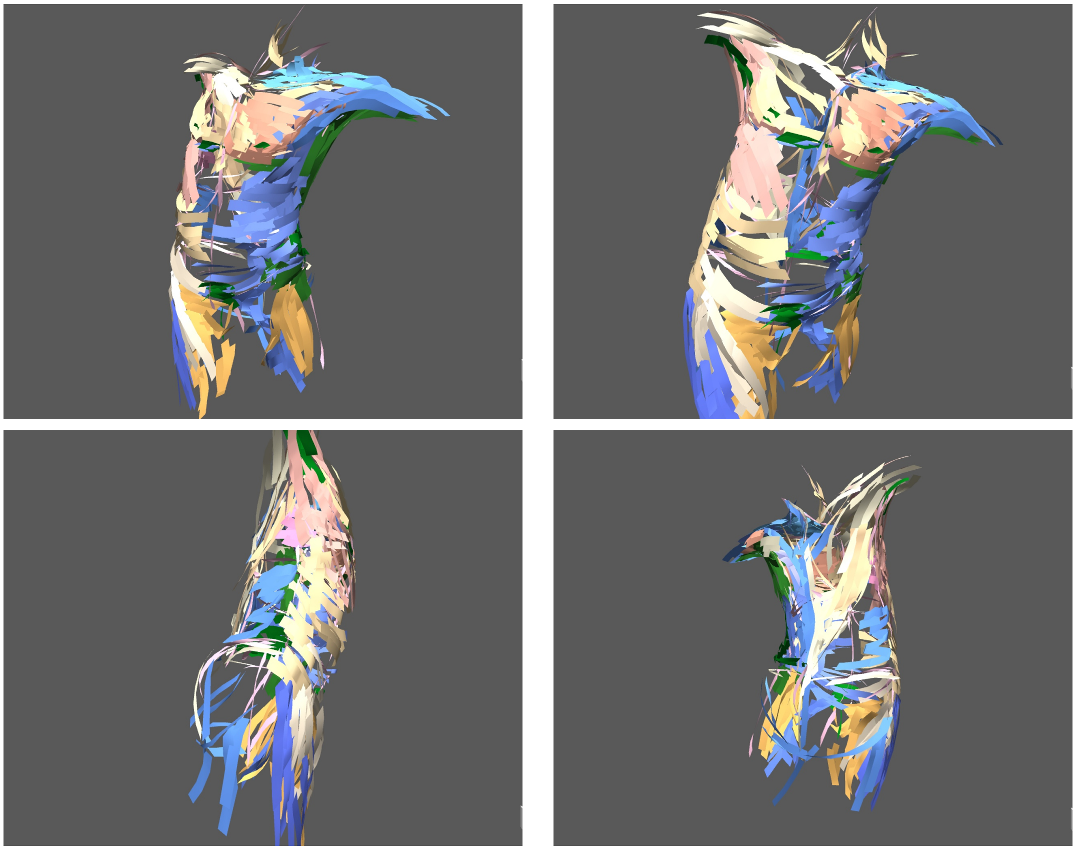

Figure 6.2: A dynamic depiction of a torso.

graphics lighting rather than specifying it with color changes in the form. There may be a role for color suggesting lighting here, but probably a bit more limited role than the way it is used here. While this was the goal while working on the form, Drury notes that some of the marks are not tied closely enough to anatomical structure. Particularly when working with this style where we can see through portions of the form, the suggestion is to look for features that aid in depiction, like the muscles on the side of the torso that pull the eye around the figure.

Sketches of Arms in Action, Figure 6.3: These sketches represent a definite struggle. In each of them, I had trouble achieving a dynamic pose and then determining how much of the form to depict. Following our evaluation of the success of the approach of leaving some holes in the representation of surfaces, I tried to stay away from depicting every inch of the form, but my sense was that I did not always pick the best part of the form to draw. Drury's assessment was that it was not as much

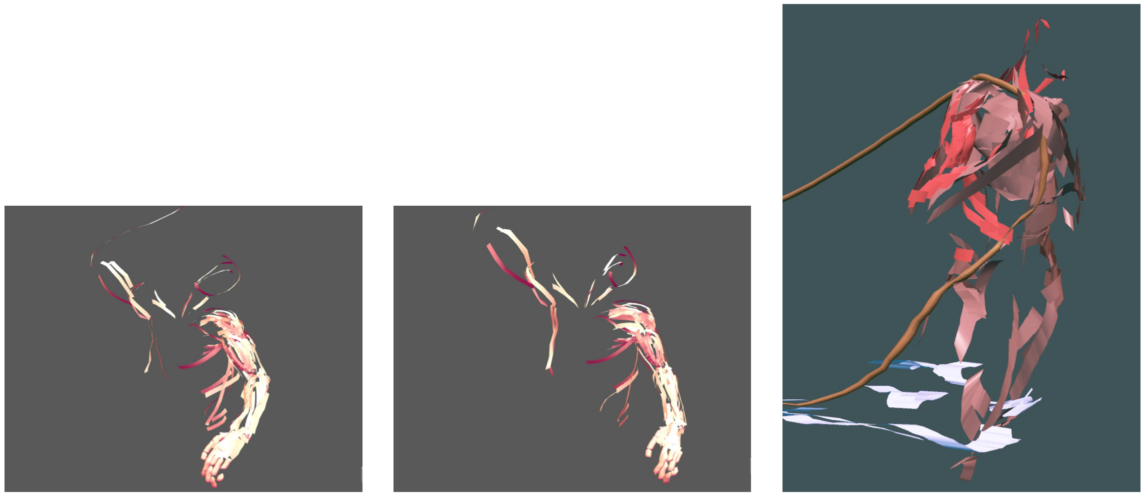

Figure 6.3: Sketches of arms in action.

a question of picking the right lines as not getting the right pose and then clarifying the important parts of the pose.

Again, there is a nice parallel to our visualization practices here. In the image on the right, for example, I was trying to depict a fisherman dragging a net. There is a purpose to this drawing. This fisherman is doing hard work. How do we convey that information as clearly as possible in the imagery we produce? The first step is to get the pose right. My fisherman is far too upright. He is not struggling against the weight of the net and the rope. Instead, he looks like he is just relaxing with a rope hanging off his shoulder. The second step is to highlight the muscles in action. We should see the tension in this man's back and in the muscles of his arms and torso that struggle to haul his catch. We are after trying to convey this level of detail and information in our work. That is why this effort and these critiques are useful in driving our research. When we can convey this level of information in our work, then we have achieved a truly expressive tool, the kind of tool that can also be useful in describing challenging scientific or medical subjects.

Descriptive Shoulder and Arm, Figure 6.4: The scapula is the main accomplishment in this model. I zoomed way in on this and defined it with many, many smaller marks. When we step back and look at it from a normal scale it has the feel of a solid form. This is a completely different style than the works above, and represents an exploration into sculpting out form. It works well in that you get a real sense of

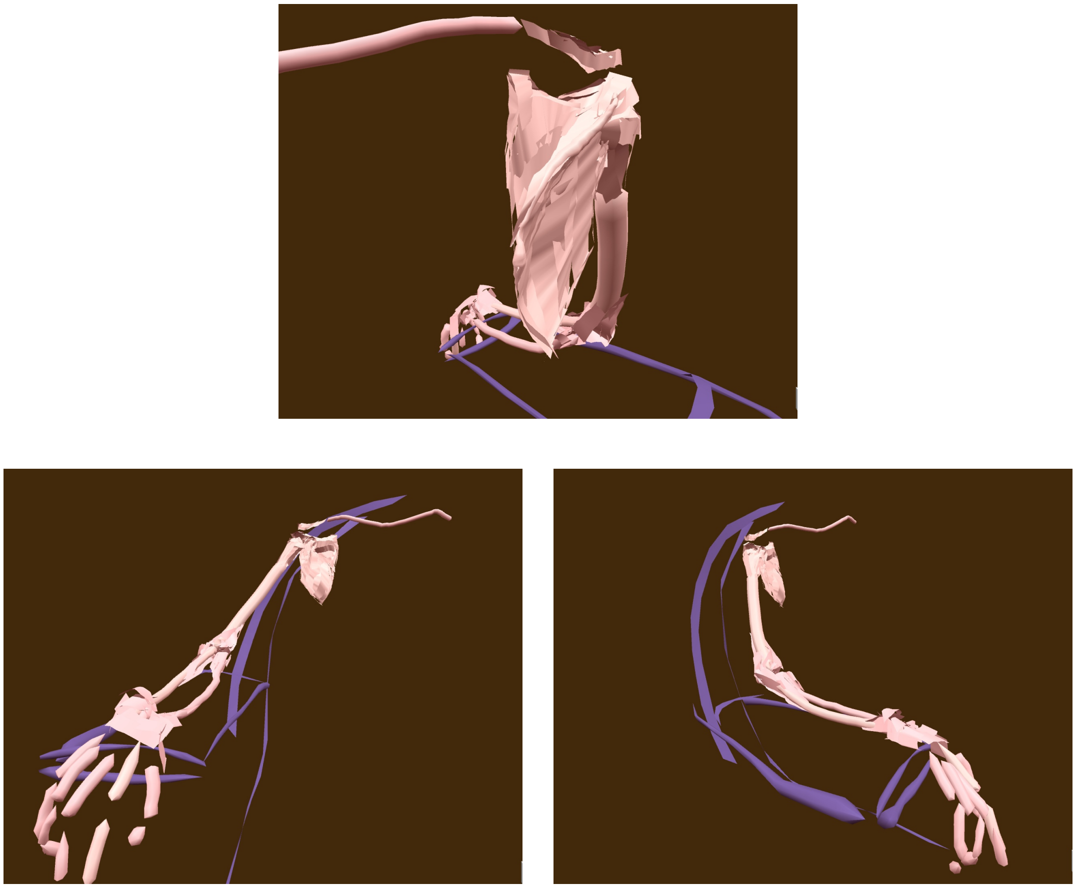

Figure 6.4: A descriptive depiction of the scapula, shoulder, and arm bones.

a solid form, much more so than in the descriptive torso model above, for example. Tube-shaped forms were used to make the rough shape of the bones, and then ribbons were used to specify the change in form at the ends of the bones. The style breaks down somewhat in areas where the ribbons overlap significantly and the form looks sloppy with jagged edges that distract the eye. The proportion is not as good in this drawing and the form of the clavicle is drastically oversimplified.

Hands Pulling in a Fishing Net, Figure 6.5: This work is a second exploration into a more sculpted aesthetic. In some sense, we look at this aesthetic and think it looks better than the other drawings because it is such a solid model. But, this style is difficult to create with CavePainting. It does not lend itself to working with surfaces like this and in some sense, this is a style that we are used to in computer graphics. The more interesting aesthetic is the more gestural, hand crafted one we

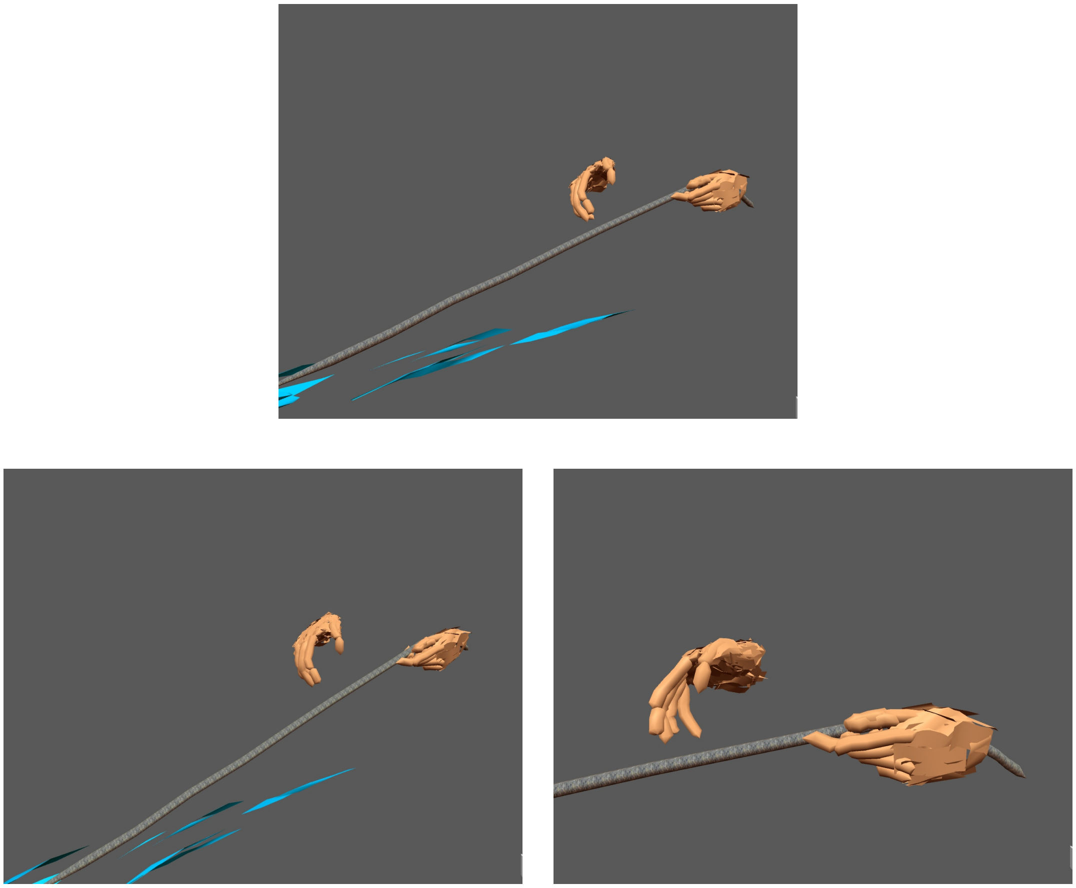

Figure 6.5: Hands pulling in a fishing net.

see when we move a bit more away from the sculpture inspired look. In terms of this work in particular, Drury commented on the posing of the hands. The pose of the hand that is about to grab the rope is fairly good. The one that is holding the rope is not as good. It looks like the rope is just lightly resting in the hand, rather than the hand grabbing onto it and pulling with force. One interesting point is that drawing out the skeleton like structure of the fingers was easy with the tube stroke. For these anatomical structures, it might be interesting to think of some sort of skeletal system that underlies the model that you could quickly sketch out. Then you might even think of bending and modifying that structure and having it deform the marks around it.

Hand Holding a Baseball, Figure 6.6: This work was a final attempt at this more sculptural style. The entire form is specified with ribbon-like mini surfaces. The

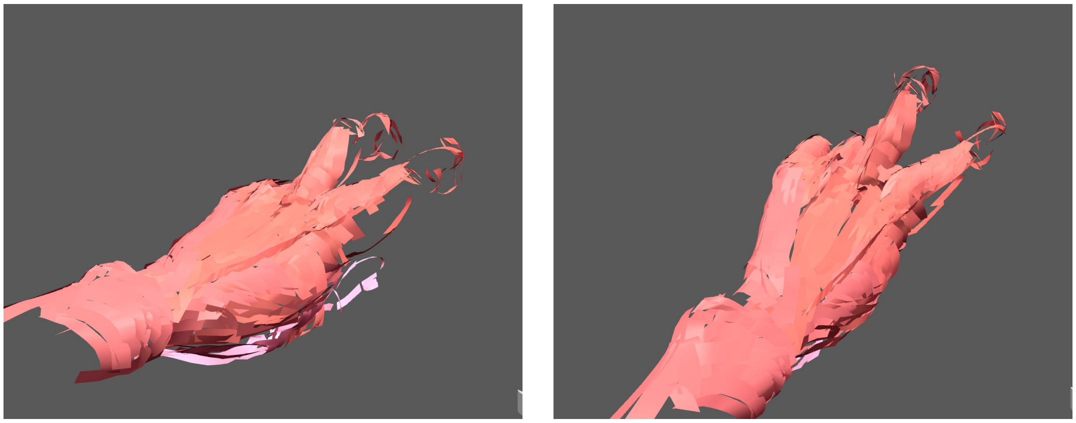

Figure 6.6: Hand holding a baseball.

knuckles are successful in this work, but again we end up with a strange style because we are trying to create solid surface from incredibly tightly packed smaller surfaces.

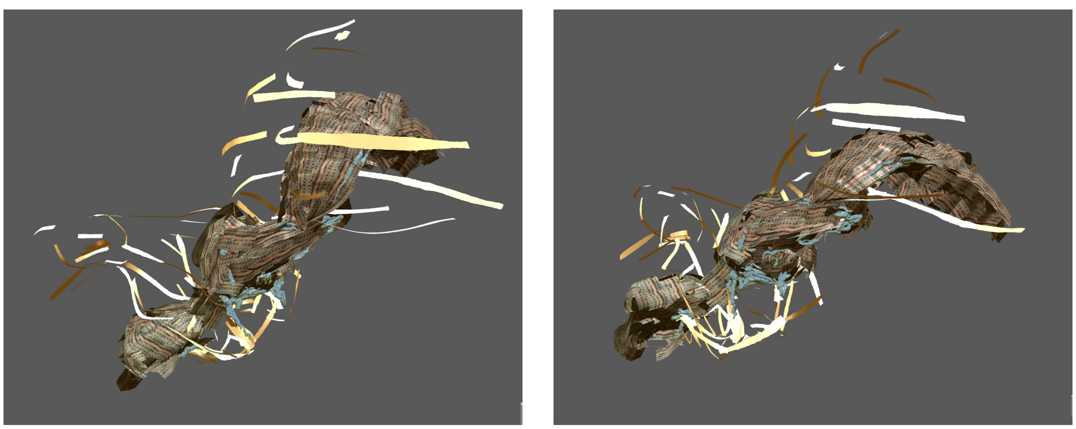

Figure 6.7: Hands squeezing a towel.

Hands Squeezing a Towel, Figure 6.7: This model represents an interesting idea which is difficult to see without the aid of stereo viewing. One of the themes that came out of our investigation to this point, is that seeing through a form suggested with a more minimal use of line is not necessarily such a bad thing. Especially when seen in stereo. A dot can suffice for specifying an elbow for example, because with stereo we can easily tell the exact location of that elbow. In this model, I tried to specify as little of the hands as possible, but to completely specify the towel on the inside of the hands. You can see through the hands to the towel that they are deforming. The idea is interesting, but as Drury noted, the implementation did not work as well.

One reason may be because the representation for the towel is so clumsy and jagged that it completely dominates visually.

This idea may hold some promise for future investigations. For example, if we import another smooth and refined model, such as one of our visualizations of brain matter, and then use this style to convey the context of the skull and other annotations around it, this style may work extremely well. This work was created with the two-handed haptic-aided interface at the fishtank. Notice the straight smooth curves highlighting the tendons in the hands that are a departure in style from the type of marks we see in the Cave-based work before this.

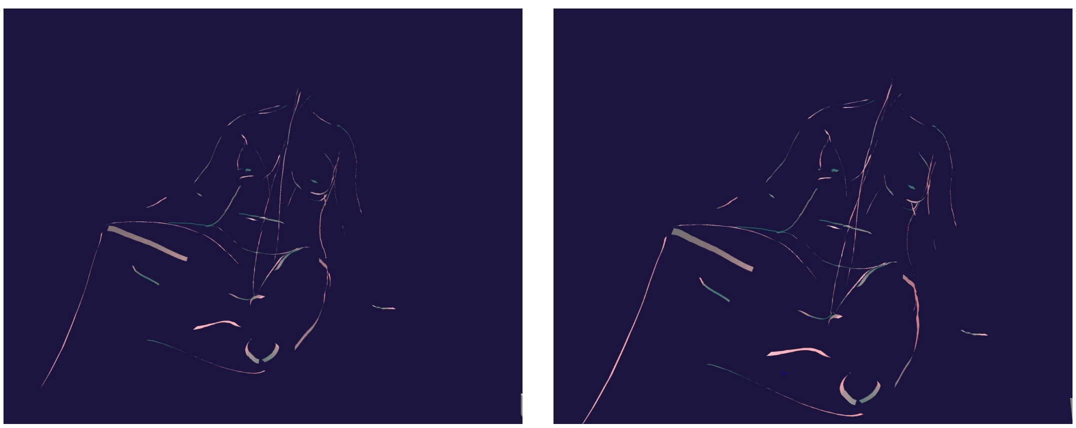

Figure 6.8: Seated woman.

Seated Woman, Figure 6.8: Drury notes that this model is the most successful example of minimal use of line. It looks good from many viewpoints, but few lines are used to depict the model. Drury points out that I picked the “correct” lines to draw, but that they weren’t always the same. For example, on each leg, I found a representative feature to draw. On one it is the outline of a quad muscle, on the other, the kneecap. This style works extremely well. The indication of the folds in the tummy area are exceptional. Drury notes that the lines indicating the shoulder are also beautiful. The one problem area is the breasts. These actually pose a difficult problem for this style of representation because they protrude from the form. If we put a contour line along the center of the form, then this looks strange from some angles because visually we cannot understand the straight line running across the smooth form of the breast. On the other hand, if we leave this type of indication out of the picture, as is the case now, then from many views, we loose the 3D form of them. In this work, we seem to have found a successful style of depiction, but we have also discovered a limit of it. The obvious extension to the tool inspired by this example is view dependent rendering. For ninety percent of the model, we can pick a representative line that does a good job of specifying the form from almost all angles, but in order to have a representation that truly works from all angles for that last ten percent of the form, perhaps we need to be able to draw the view differently depending on the orientation of the viewer.

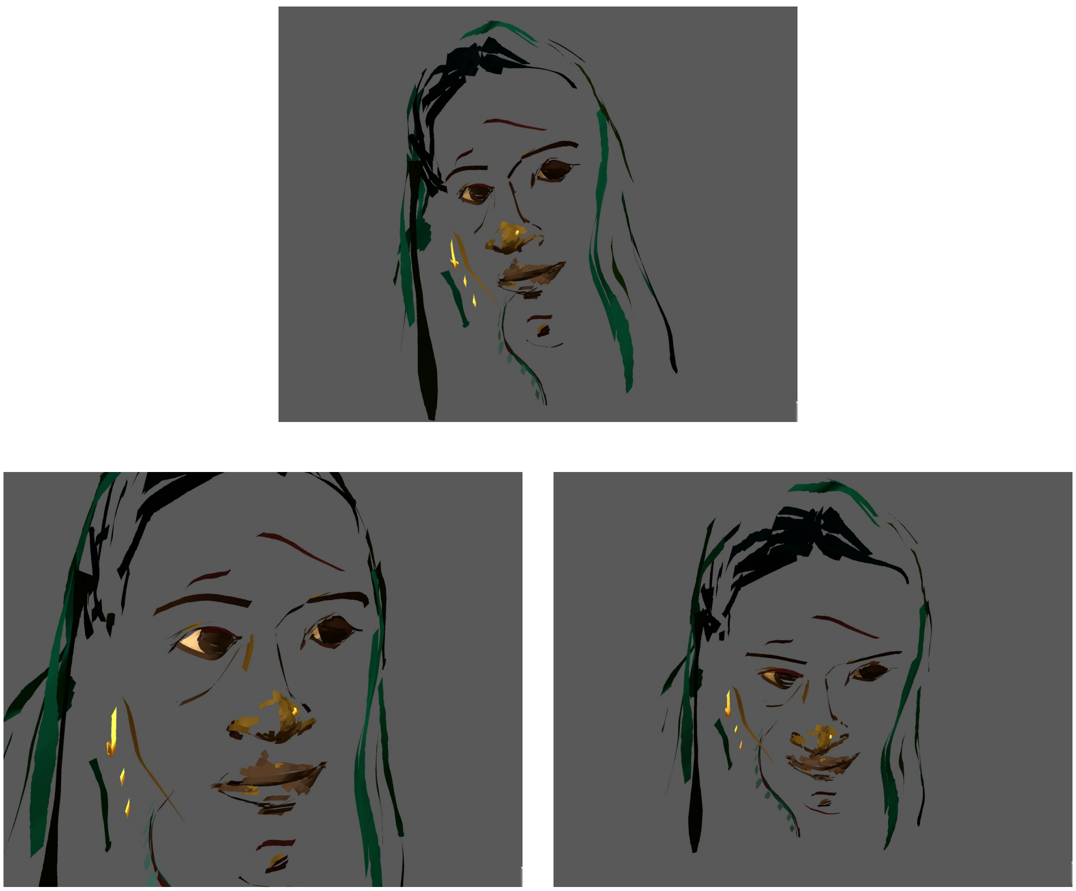

Figure 6.9: A Swahili bride wearing a green veil.

Swahili Bride, Figure 6.9: This work is a good example of a face that has good 3D proportion, so it looks good from most angles. Most of the lines are in the ”right” place spatially, but because there is fairly minimal use of line, we find that the lines I picked to draw are not necessarily the best ones to draw from all angles. Good nose. Fairly good eyes. The simplified color palette works well for this model.

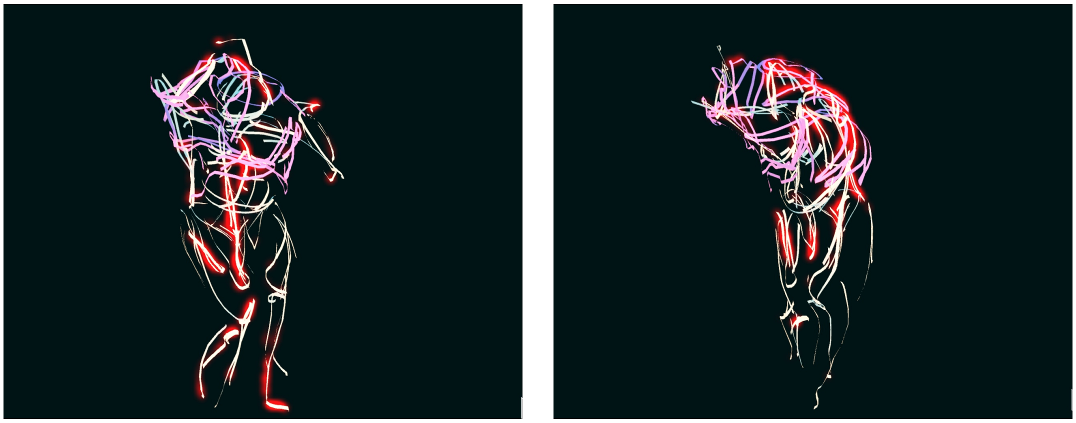

Figure 6.10: Woman undressing.

Woman Undressing, Figure 6.10: This model illustrates one of the advantages of working in the Cave. This is an extremely gestural drawing that would be much harder to do in the fishtank system because it is so constrained. There is also a glowing effect in these images achieved in real time via tone-mapping. This look is inspired by Picasso's light pen drawings in the air. Drury was intrigued by the glowing rendering effect, but the challenge is being able to control it while creating the form. This is the type of 2D effect which may help our projected images appear less flat and boring when printed on paper.

Bearded Man, Figure 6.11: This model has an extremely well defined nose. While working on this model with the haptics-based interfaces, I discovered that tracing along an existing path is much easier with this interface because the two-handed tape drawing technique includes a guideline. The artist can draw slowly, but still get a smooth curve. Drury liked the 3D placement of the brow of the head. When looking from above we get a sense that it jets out of the form. The right eye if off in its bilateral symmetry. Again the sparseness of representation worked well. The indication of the hair for example works well.

Anatomy of a Hand, Figure 6.12: Several interesting ideas are at play in this work. Again, there is an attempt at economy of line. Again, it works well from some angles, not as well from others. There is a nice use of transparency used to define the muscles. We still see a lack of clarity in things like the insertion points for tendons and muscles. Specifying this level of clarity will require a fine level of control at a smaller scale. The translation from stereo viewing to mono viewing is so drastic with

Figure 6.11: A bearded man.

the model. To capture the interest of artists and make our 3D illustrations useful as 2D projections, we need to follow more of the basic guidelines of suggesting depth with drawing. For example, color of the marks should be adjusted with the respect to the background color to help clarify depth relationships.

In this chapter, we analyzed the success of our tool and presented findings on means of successful artistic representation in VR, insights gained through the methodology of critique with an expert visual artist. We reported on insights in the areas of the balance between sculpting and suggesting form, economical use of line in depiction, the use of color along with CG lighting, stereo verses 2D viewing, and differences between cave and fishtank VR environments. These insights are described in large part through the a presentation of artwork and summaries of their critique.

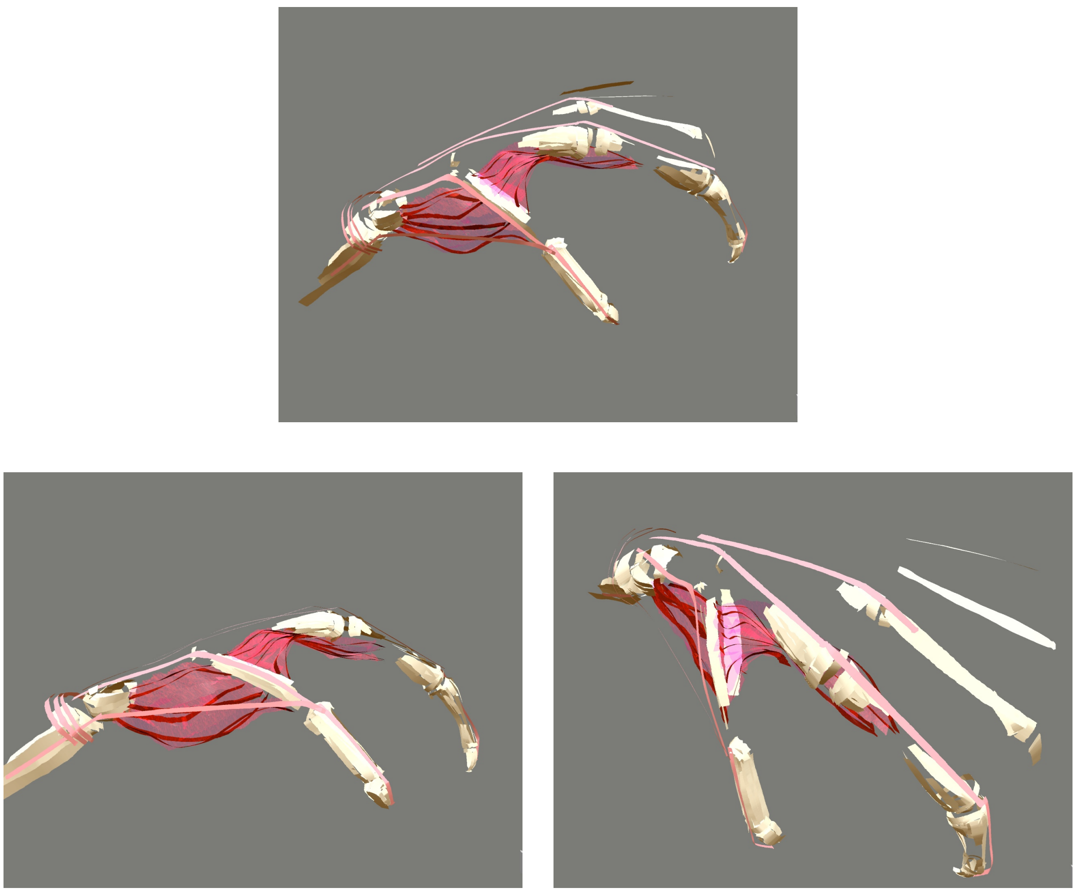

Figure 6.12: Anatomical drawing of a hand.

Chapter 7

Applications in Scientific Visualization and Medical Illustration

In contrast to the artistic explorations in the previous chapter, the works presented here target specific visual problems in scientific representation. These visualizations and illustrations were all created with our haptic-based tools, which provide the control and expressive power necessary in order to tackle these complex subjects.

7.1 Proposed Illustration of Bat Wing Bone Deformation During Flight

Our group has been heavily involved in an ongoing collaboration with Dr. Sharon Swartz of the Evolutionary Biology Department at Brown. We have been helping Dr. Swartz visualize experimental and computational data describing bat flight with VR-based visualizations like the one in Figure 7.1. Recently, Dr. Swartz has collected a new dataset that quantifies for the first time the amazing variation in the mechanical properties of bat wing bones. This incredible variation is an anatomical feature that is unique to bats among all mammals. The rigidity of human bones, for example, is for the most part identical in all of the bones of the skeleton, with the only exception being the bones in the ear. Bats account for something like eighty percent of the

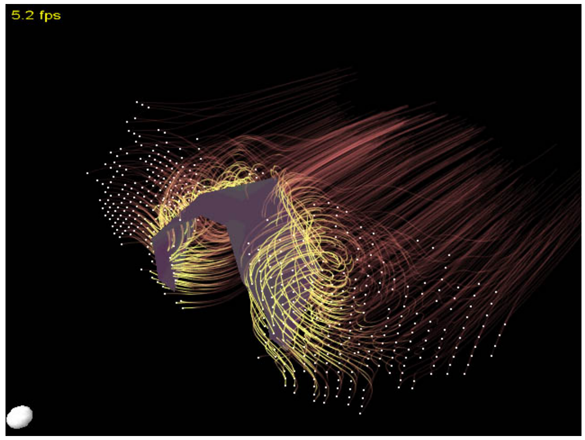

Figure 7.1: Current bat flight visualization methods describe air flow, but not the changing wing geometry of the bat.

variation in bone rigidity in all mammals. Understanding the function of this unique feature of their anatomy may provide valuable insight into how bats fly.

The specific data collected by Swartz's lab describe the bendability of each of the wing bones for a particular species of bat. A challenge in understanding this data is making the connection between this set of numbers and the changing shape of the bat wing as it flies. Bat wing motion is highly complex, far more so than what we typically find in birds. Bird wing motion is often described in texts with figures that plot a projection onto front, side, and top views of the joints in the wing as they move through the wing beat cycle. For bat flight, this 2D visualization method is too difficult to interpret because the motions of so many of the joints in the wing are highly three dimensional. In order to understand this new data within the context of the complex bat wing motion, we need 3D illustrations that incorporate both the unusual pose of the wing at important points within the wing beat and the mechanical properties of the bones in that pose.

7.1.1 Illustration for Data Exploration

We propose to use the haptic-based 3D illustration tools presented in this thesis as means of helping Dr. Swartz explore her new data. For this particular species of bat we have images of the bat in flight, but no 3D positional data. This makes it difficult and time consuming to generate a completely data driven picture of both the wing geometry and bone material properties, yet as described earlier, the bat flight problem is so incredibly three dimensional that we need such a picture simply to begin to understand the data.

Our proposed solution to this data exploration problem is to work with artists to help reconstruct a useful 3D illustration from the available photographs and bone data. We expect this will involve importing the bone properties data into our drawing system and displaying it in some visual manner for the illustrators. For example, given a basic initial skeleton sketch that the artist has drawn the data could be imported to display a series of height bars next to each bone that illustrate the degree of bending that it exhibits at a particular point in flight. Once the data is depicted visually in this way, the illustrators can work to modify their sketch to present an illustration of the data.

The obvious shortcoming of this approach is the potential that the illustrators will misrepresent the data. Our sense is that this is far outweighed by the potential for the scientist to quickly see a 3D representation of her data. By quick here, we anticipate turn around times of days, not the months that can be involved when serious programming is required in order to produce a visualization. We also feel secure in our approach given the background of the illustrators with which we will work. As we do in the field of graphics and visualization, these artists work for years and years to train themselves to represent forms and ideas with clarity and understand any biases they may introduce into a picture through the method of representation they choose. As long as we all understand these biases, these tools present scientists with a wonderful opportunity to get a glimpse of their highly complex 3D data quickly.

7.1.2 Illustration for Data Presentation

The second stage of this problem arises once we have gathered some scientific insight from the data and want to present that insight to peers. We can think of this as a classic problem in medical illustration, except in the case of problems like our bat, 2D illustrations typically fall short. We propose to use these 3D drawing tools to make 3D illustrations, iterate on them collaboratively with Dr. Swartz, and arrive at computer-based 3D illustrations of her findings that are useful for presentation to her peers.

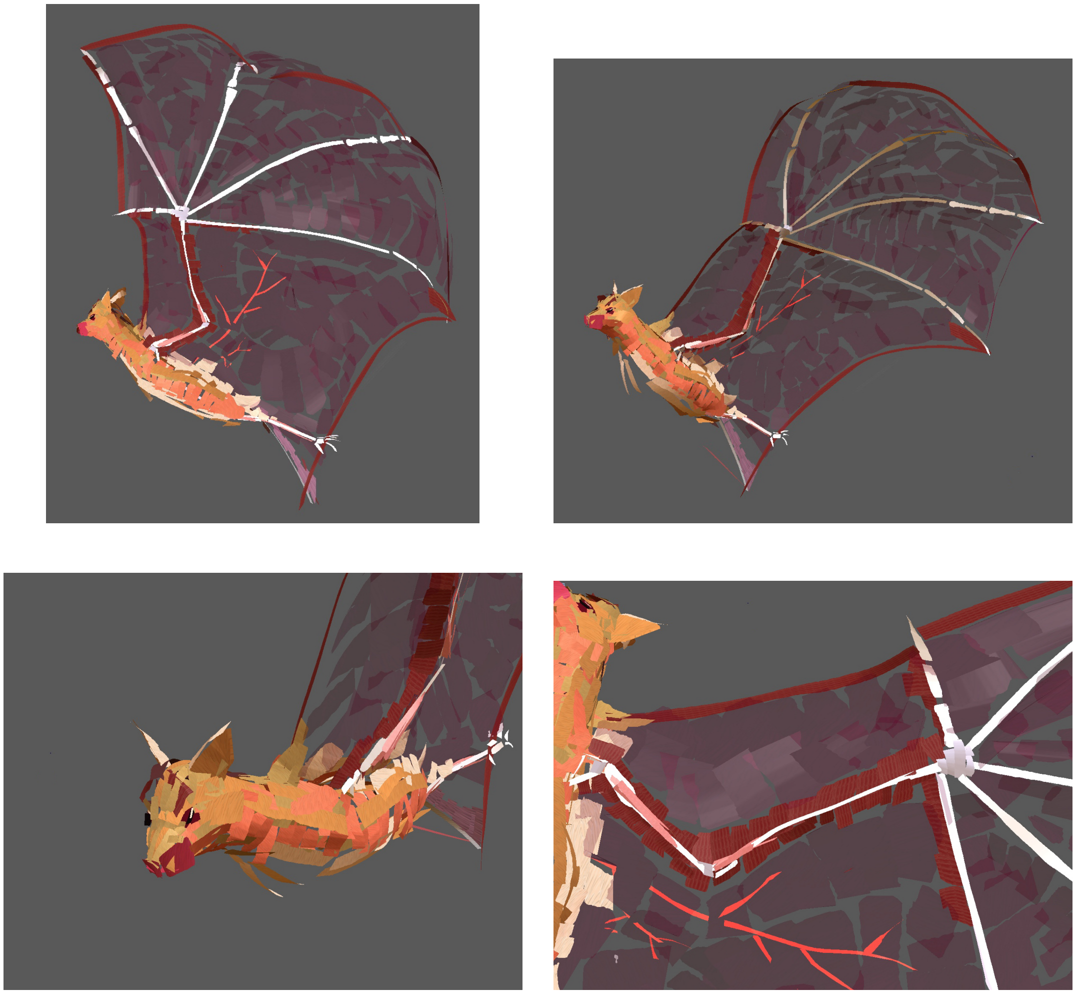

Figure 7.2: Initial bat bone illustration.

This is a very natural role for these tools. Dr. Swartz was very encouraged to think of this approach as a potential solution. In some of our initial meetings we showed her images from an initial illustration of a bat done with our most recent tools, as seen in Figure 7.2. In this model, for the first time, our tools are precise enough to represent very smoothly curving bones in the wing and joints where the arm bones meet the shoulder and in the hand. The artist has even reached the point where he is beginning to explore representing veins in the wing membrane, tendons, and bones in the foot. Dr. Swartz was impressed with this new level of clarity and is excited about the future of this project. She thinks that with the added illustration power provided by these tools, this could push the work over the threshold needed to

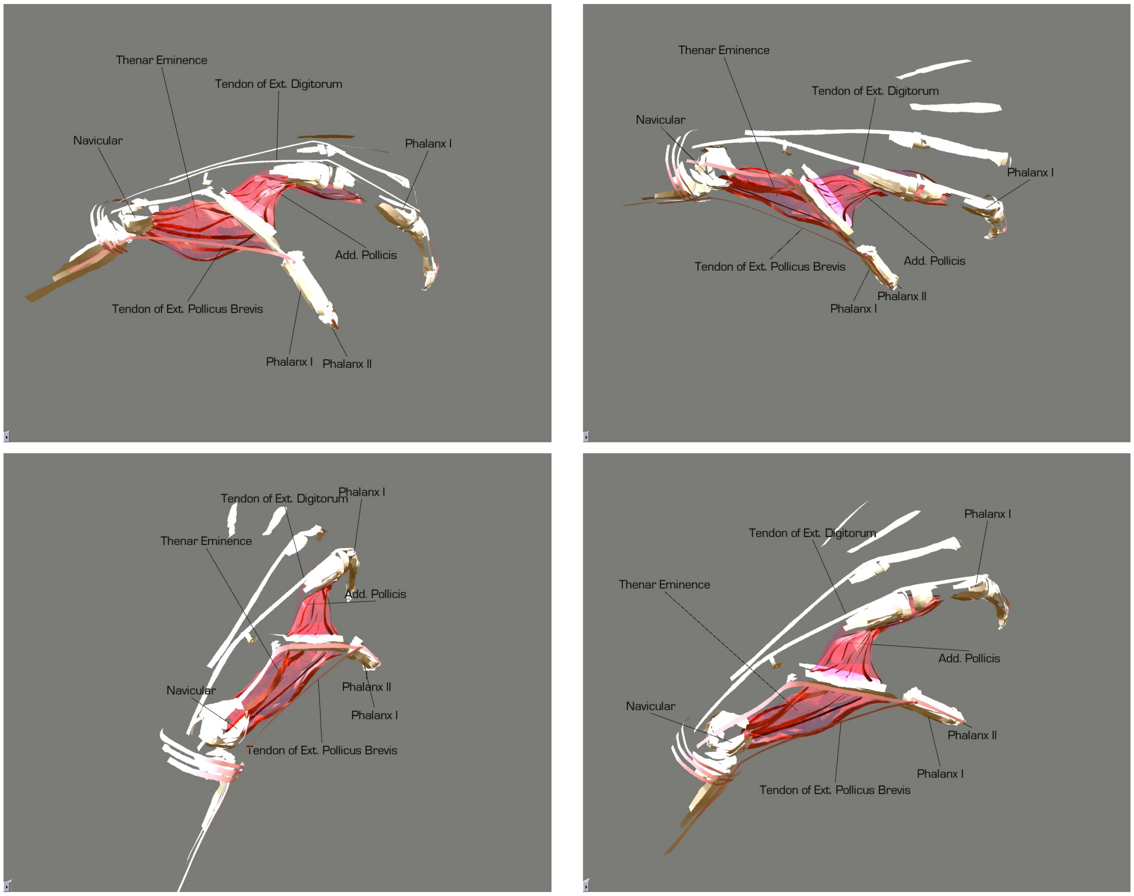

Figure 7.3: An annotated illustration of hand anatomy.

submit it for publication in *Nature*.

7.2 Anatomical Illustration of a Hand

The illustration of bones, tendons, and muscle in the hand in Figure 7.3 provides us with a sense of the type of 3D medical illustration that may be possible with this tool. The ability to create 3D annotations was added to the software to enable the creation of this model. We are exploring potential collaborations with faculty in the Brown medical school who have become interested in custom medical illustration like this for use as a teaching tool that can help students understand human anatomy.

In this chapter, we present scientific visualization results created by artists using our approach to free-form modeling. Because of the scientific subject matter, these works demand a level of sophistication and clarity in representation that was unavailable in previous free-form modeling techniques. Now that artists can work with an intuitive but controlled and expressive approach to 3D illustration, they are able to play a much more important role in addressing scientific problems with visualizations based in virtual reality.

Chapter 8

Analysis of Control in Artistic Free-Form Input

8.1 Background and Motivation

The results and discussion of the previous chapters motivate the need for control when working with direct 3D input strategies. In general, our approach is to leverage the immediacy and intuitiveness of direct, freehand 3D input and develop novel approaches to help artists overcome limitations sometimes imposed by the unconstrained nature of this style of input. This chapter presents our investigations into evaluating the success of our techniques. We break this evaluation into two main components. The first and more subjective component is validating that the style of input and modeling we present is useful for artists. We hope to establish some sense of how easy it is for skilled artists to learn and use our techniques. We would like to know the sort of subjects for which they anticipate the tool being most useful, and we would like to gain some sense of how useful it is for subjects in medicine and science, two of our target application areas. The second component of this investigation deals with trying to better understand the theory of the interfaces. To do this, we attempt to determine the relative importance of aspects of the interface. For example, is the two handed nature of the drawing more or less important for control of the technique than the friction and viscosity forces exerted by the haptic device.

At the time of this proposal, we have completed some pilot studies that begin to

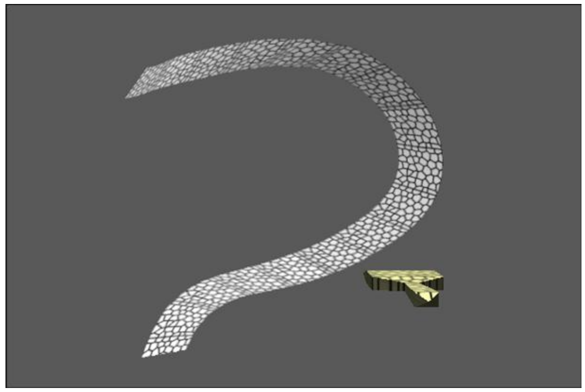

Figure 8.1: Example display for the tracing task.

address these questions. We have developed a collaborative study RISD winter session class to be taught in conjunction with Fritz Drury of RISD's Illustration department in January-February 2006 to train four or five Illustration students to use our modeling tools, and we have developed an experimental methodology that includes collecting feedback from artists trained in the winter session class and quantitative data collected from studies of artists drawn from this pool of trained users.

In the rest of this chapter, we will explain the methodology and insights gained from our pilot studies, discuss how these have shaped and refined our new proposed methodology, and then describe that methodology and our hypotheses in detail.

8.1.1 Pilot Studies

One of the first interesting insights we had after porting CavePainting to a Fishtank VR setup is that the scale of the drawing interaction plays a major role in control. To investigate this issue more closely, we set out to do a comparison of control in a Cave environment verses a Fishtank environment.

Pilot 1: Cave vs. Powerwall vs. Fishtank

For this pilot study we created two experimental tasks for artists to perform. The first is tracing, and the second is replicating. We wanted to measure artists performance at the tasks performed with the CavePainting software running in different VR form factors. Our hypothesis was that artists would be better at these tasks when performed in the large scale environments, like the Cave and the Powerwall as

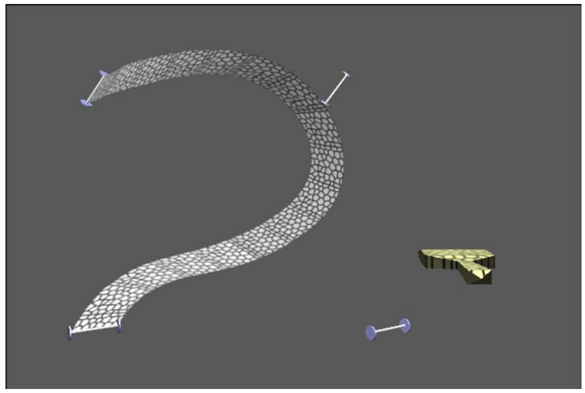

Figure 8.2: Example display for the replicate task.

compared to the Fishtank. We expected a slight preference for a Cave configuration over a Powerwall with the idea that users would be able to perceive 3D form more easily and naturally in the fully immersive Cave where they are free to walk almost completely around the virtual forms.

In the tracing task, the subject is presented with an example 3D swath like the one in Figure 8.1. For these experiments, the 3D swath was created with a CavePainting ribbon stroke, a triangle strip depicting the ruled surface generated from sweeping and twisting a flat brush through space. The subject observes the prompt ribbon for five seconds before he is allowed to start tracing it. He is free to move about the space of the environment while observing and while tracing. After the initial 5 second forced observation period, the subject has 30 seconds to begin tracing the ribbon. The subject pushes and holds a button on the brush device to start tracing and gets one try to trace the prompt as accurately as possible keeping in mind that the goal is to capture both the position and the twist of the ribbon.

The replicate task is very similar. The first 5 second observation period is the same. But after the initial period, the prompt ribbon moves away from the subject by a foot or two, depending on the form factor. It remains in view at this distance, and in its place in the original position are two dumbbell shaped markers, shown in Figure 8.2, marking the start and end points for the original shape. The subject's task is to replicate the original ribbon as accurately as possible between the start end points.

These two tasks approximate what we would really like to test. Ideally, we want to know: Does the artist have enough control with the interface to draw exactly what he is thinking in his head. Unfortunately, since we don't know what is in his head, it is impossible to create an error metric for this task. Instead we try to approximate it with tracing and replicating. In these cases, we know what is in the artist's head because we put it there by making them reproduce a prompt mark.

In the tracing task, the artist has a very good sense of the form of the prompt mark because it lives in exactly the same space as the drawn mark. So, we're fairly confident that the right shape is in the artist's head. As he draws he gets constant feedback on how closely the two marks match, making this a fairly ideal drawing situation. However, this task is somewhat unrealistic because in normal use, artists do not trace existing marks, they create them from scratch.

In the replicate task, we experiment with a second task that tries to get closer to the realistic situation. Here, we are less sure that the artist has the exact form of the prompt mark in mind when the drawing begins, but the drawing is done in more of a normal context. You can think of the replicate task as looking at a still life setup in the distance and then trying to reproduce it locally – a far more typical use of the tool than tracing.

In both of these tasks we wanted to use prompt marks that are representative examples of the types of marks commonly created with the tool. To achieve this we captured marks from existing CavePaintings made in each of the VR form factors along with characteristic curves from laser scanned real sculpture. This last category of marks is like the holy grail for our input techniques. With them we hoped to test how close artists could come with our interfaces to reproducing the S-curve of the spine or characteristic shape of the clavicle in celebrated works of sculpture.

To facilitate comparisons between the environments, the prompt marks were rescaled to fit comfortably within the typical working volume of each environment. For example, marks drawn originally in the fishtank were scaled up to larger and more completely fill the environment of the Cave.

We ran across several issues with this design which eventually doomed the results. Many of these informed future designs, and are therefore useful to discuss here. The first was our inability to control for variation in the tracking devices used in the different setups. In the fishtank, the magnetic tracker we used produces extremely noisy data when close to the CRT display. But, the display is small enough that if you want to work at a reasonable size you need to be very close to it. This was just an impractical tracking situation, so we tried working with an offset applied to the tracking system. This allowed us to move the brush at the fishtank in a slightly larger area at a distance of a foot or two from the CRT. This improved tracker accuracy, but at the cost of adding the confounding variable of the offset. Even with the offset applied, accuracy in the tracking data fell off at a much quicker rate than with the Cave long-range tracking system as the brush was moved away from the tracking transmitter.

At the fishtank setup, noise in the tracking data also seemed to be more of a problem because everything is drawn at a smaller scale. If jitter and noise in the tracking data is at a scale that is one percent of the size of the shape being drawn, then it is probably lost in the shape. But, if it is ten percent, it can be very apparent in the final shape and dominate its visual appearance.

In another attempt to control for some of this variation, we employed a mechanical tracking device at the fishtank. For this we used a phantom haptic device with the force feedback feature turned off. The phantom reports extremely accurate tracking data, but it also requires an offset to be applied between the hand and the virtual location for the brush and its working volume is extremely small. There are also some strange inertial sensations that arise from using a tracker attached to a mechanical arm. Even with this highly accurate tracking situation, the fishtank setup proved to be difficult to control. Our sense is that a big part of this difficulty is a fundamental difference between working with hand, wrist, and arm mechanics in such a small scale and that this is one situation where there is really an advantage to working in large scale environments like a Cave. But, without being able to control for these types of tracking issues, this is a very difficult hypothesis to prove.

We also imagined this study as a situation where subjects that were trained artists, but probably not trained in VR would be able to come in and spend an hour running through the study. We found that this was unrealistic. Subjects were certainly able to learn the interface within several minutes, but they were not comfortable in the sense that you could imagine them creating a work of art with the system. This helped us to understand that we are far more interested in how the interfaces work for artists once they get over the novelty of VR and are trying to do real work with the systems.

The two tasks we picked are good candidates for study because it is easy to compute many different error measures from the difference between the prompt mark and the drawn mark. We computed errors in beginning directly on the start point and finishing on the end point. We also found the area between the centerlines of both marks and divided this by arc length as a measure of error along the curve, and used a similar strategy with the edges of the curve to get a measure of error in the twisting of the mark along the curve. We found that while an almost countless number of error metrics could be constructed here it was very difficult to come up with metrics that have a reasonable physical and artistic meaning. Often, what we really want to know as artists is: Is this line good enough? Good enough depends almost entirely on the context around the line and often has as much to do with the emotion or character in the line as it does with the accuracy of the line relative to other marks. For example, if the rest of the drawing is done with very smooth lines, then a jaggy line that very closely approximates the centerline of the prompt is probably worse than a smooth line that may move off the prompt mark somewhere in the middle of the mark. In this study, the tasks are performed with no surrounding context, so it is impossible to make these sort of context-based evaluations of error, and it is difficult to say whether any of the techniques are "good enough." Without the context we have no notion of acceptable error, even at the level of asking the artist himself.

Another issue we discovered through observation of our pilot testers was that subjects were approaching the tasks in ways that were unnatural artistic uses of the tool. There was far too much attention paid to the individual mark. Subjects often practiced tracing a line several times before doing a real tracing. Because they were tracing rather than drawing, they often moved the brush so slowly that extra jitter and shakiness in their hand produced a jagged mark that would not have occurred if they had just drawn it without trying to reproduce the example.

Pilot 2: Bimanual Haptic Drawing at the Fishtank

In a second pilot study we decided to avoid the issue of controlling for variation in tracking systems by using the phantom device in all conditions. In this study we were interested in accuracy of a 3D mark traced in three different experimental conditions: 1. using the 2-handed tape drawing interface described in Chapter 5, 2. using the same interface but without any force feedback, and 3. using the phantom stylus for completely freehand drawing, without the tape drawing interface and without force feedback.

In this pilot, we used the tracing task described above, but with a tube shaped form rather than a ribbon, so the goal was to match the position of the prompt mark only, there was no twist component. Our subjects were the developers of the software for the project (they are also artists as it turns out), so they were skilled at drawing and also highly trained with the two-handed tape drawing technique. We were unsatisfied with reports from the error metrics described in Pilot 1 above, so we added a self-rating component to this pilot as a way of ranking one's own performance.

The study was arranged in blocks of trials with ten trials per block. There were ten different prompt marks in the study and each one was shown once per block in a randomized order. At the end of each block the input condition was changed. Subjects completed 12 blocks of trials in total, which gave us 4 blocks for each input condition, yielding 4 attempts at tracing each of the prompt marks for each input condition.

After this part of the study was complete, the subjects evaluated their own performance using the following tournament-based strategy. First, the subject picked the best of the four tracings done for each mark in each condition. This was done by displaying two marks side by side on the screen overlaid on top of the prompt mark and forcing the subject to pick the better of the two. This forced choice process proceeded in a tournament style, with the two winners in the first comparison going up against each other in a second round, until we arrived at the best tracing of each mark, for a given input condition. At this point we have 30 "winning" marks, one for each input condition (3) x one for each prompt mark (10). Then we continued with the tournament style evaluation, comparing the best marks done with each of the input conditions against each other until there was one winner for each prompt.

We found immediately that that third condition (completely freehand, no tape drawing) was so much worse than the others that it was almost not worth testing. It was never picked by the subjects as the best tracing of a given prompt, even when the original prompt was creating using the freehand drawing condition. The other interesting result was that the tape drawing without haptics performed almost as well as the tape drawing with haptics. Subjects selected the winning mark from each condition with roughly the same probability and remarked that in many cases, the forced choice for the best mark was difficult to make because they both seemed to be of equal quality.

Our experience with this pilot prompted several insights. For many of the marks we are interested in drawing, haptics may not be necessary for getting accurate positioning of the mark. But, our sense is that it becomes more and more helpful when the curves make sharp turns and when twist in addition to position is being recorded. The completely freehand condition was so much worse than the others that it made us wonder why we would even test it. It does offer a useful baseline for comparison, but it also takes up valuable time during the experiment. Because this format of study did not force subjects to make use of the line weight control feature of the haptic-based interface we were not using it up to its full potential. In some sense, this points to the idea that we were not testing as interesting a question as we might be able to.

8.2 Proposed Methodology

Based on these insights provided by these pilot investigations, we propose the following approach to continuing this line of inquiry.

One of the most important findings from our pilot work is the need to give artists time to learn the interfaces we present. With this constraint, the most useful approach for a study is to have artists work regularly with the tool over the course of at least several weeks. In our new design we plan to take this approach with the primary focus being a structured setup for gathering qualitative feedback from the artists. A secondary focus of the study will be gathering quantitative measures that help us to better understand the relative importance of different aspects of the interface.

We think this approach is the best format for helping us to answer the following important questions:

1. Can the newer minimalistic style enabled by the line weight control and charcoal blending be useful artistically? The main advantages of this approach over say sculpting are: You can suggest a form with just a few marks rather than sculpting out the whole thing, and you specify the look and the form in the same motion. However, this style is unconventional it would be nice to get some confirmation that artists think it is an interesting one.

2. How easy is the technique for artists to learn? In fact, we know it will take some time to learn from our pilot work, but we would like to be more confident in our statements that trained artists can pick it up and be able to back these statements up with data from artists that were not involved in the development of the tools as most of our pilot users were.

3. We assume the charcoal technique is much harder to use than the more direct drawing technique because it takes more coordination and requires more attention to the interface while drawing. Does the charcoal addition add enough stylistically that it is still useful?

4. Subtleties of the technique: How much does the friction and viscosity help with control of form? How important are the line constraint haptics? These are more directed questions. We would like to get a sense of the relative importance of different aspects of the technique in order to understand more of the theory of it and what areas are most promising for future investigations.

Since so many of these questions (particularly questions 1, 2, and 3) depend on artists using the tool for serious work and also depend on them having enough time to learn the tool, we think the most beneficial way to investigate these questions is to arrange for artists to make continued, real use of the tool as opposed to recruiting them for an afternoon-long study. To this end, we have arranged to teach a winter session class in conjunction with RISD in which several art students will work with our system over a six week period. The goal for the class will be to familiarize them with the tool. We will do this through a series of assignments. Some will be directed, like working from still life subjects. Others will be open ended, affording the artists the chance to apply their own creativity to the tool and see where that

Position	Line Weight
1. Freehand with viscosity and friction	twist
2. 3D tape drawing no line constraint	twist
3. 3D tape drawing with line constraint, viscosity and friction	twist
4. 3D tape drawing with line constraint, viscosity and friction	pressure
5. 3D tape drawing with line constraint, viscosity and friction	twist + pressure
6. 3D tape drawing with line constraint, no viscosity and friction	twist + pressure

Table 8.1: Experimental matrix for haptic interface study.

leads. We hope to answer many of the questions posed above with feedback from this pool of artists taken in the form of interviews and questionnaires. Feedback from the directed assignments will help us to evaluate how well the interface works for very specific tasks. Feedback from the more open ended assignments will help us form an impression of the overall expressive power of the tool and gauge its potential artistic contribution.

To address the fourth area of inquiry outlined above, the subtleties of the technique, we will run a more quantitative study with subjects drawn from the pool of artists trained as part of the class. The tasks in the study will consist of some combination of the tracing task outlined in the pilot studies above and perhaps a variant of it that incorporates more context. The prompt marks and context may be taken from works created by the students in the course. For these tasks or perhaps a similar one to be determined soon, we plan to test several experimental conditions. They are outlined in the matrix in Table 8.1, where the column labeled Position specifies the style of the interface used to establish the position of the curve that is drawn and the column labeled Line Weight specifies the interface for adjusting the line weight of the mark.

Note that in this setup, all the interfaces control both position and line weight in some way. Since we have determined that this is a useful aesthetic, we are interested in how well artists can control both of these parameters at once. The ability to control position alone, as tested in pilot study number 2 is far less interesting.

Another change from the pilots is the omission of the completely freehand condition. Here, it is replaced with a freehand with viscosity and friction forces. While the completely freehand condition is so bad that nobody would ever use it, if we add a viscosity and friction effect via the haptic device, the freehand condition becomes easier to control and may provide a useful baseline in this study. Twist in the table above refers to the artists' ability to spin a brush with a flat tip to create a line that is thinner or thicker. In the physical world, line weight is a combination of this twisting and pressure applied against the paper. The final three conditions examine the artists ability to control various combinations of these line weight actions in our 3D setting.

8.3 Hypotheses

We expect that all of the 3D tape drawing-based interfaces will out perform the freehand one due to the added control and smooth quality of the resulting lines with tape drawing. For marks that curve only slightly, we expect the 3D tape drawing without a haptic line constraint to do nearly as well as the haptic version for specifying the position of the line. However, we expect the versions that incorporate haptics to out perform the non-haptic versions when looking at accuracy for both the position of the mark and its line weight. The viscosity and friction effect seems to aid control in both the freehand and tape drawing cases. We expect the best overall performance with condition number five. Condition six is the same as five without the added viscosity and friction effects, we expect subjects to be less accurate with this condition, but are not sure how much less. It will be interesting to see how much these additional force effects aid control.

8.4 Analysis

If we can show that the tape drawing-based interfaces outperform the freehand with viscosity and friction condition then we make a very interesting contribution. It is not at all intuitive that a two-handed drawing technique would be easier to control than a one handed in this context. Despite the fact that the benefits of tape drawing on a large scale are well understood in the field of car design, demonstrating a working, computer-based, small scale, 3D version of this technique that outperforms the state of the art in 1-handed direct 3D input would be a significant discovery.

The comparison of condition three with four and five, that is the 3D tape drawing technique with and without a haptic line constraint also promises to yield some interesting results. Bill Buxton, the first to introduce tape drawing to the HCI community, said upon trying our haptic 3D version of the technique that it was the first time he has seen a tape drawing variant actually work in 3D. The haptics clearly contribute something to the drawing experience, but in our second pilot study we noted that for marks that are only slightly curved the haptics may not be as critical as we thought, at least for controlling position. The comparison of performance in these three conditions may help us to establish the relative importance of the haptic feedback.

8.5 Future Work

In this new design, we have realigned our focus to address the most interesting questions that we think we can accurately test at this point. In doing so, we have avoided investigating the issues of scale and total immersion. Recall the initial insights that drove our first pilot study about the difference in control between interactions in the Cave and the smaller scale fishtank VR setup. This is still a very interesting question to explore, perhaps as camera-based tracking becomes more stable and widely available we can control for the tracking differences that plagued us and revisit this line of inquiry.

If haptics on small scale is better than large scale no haptics, but large scale no haptics is better than small scale no haptics, then do we get even better with large scale plus some haptics? Large scale haptics probably would not be as sophisticated as a Phantom, but maybe even light haptics combined with large scale interaction would be an improvement over our best small scale results. From our own informal evaluation, tape drawing on a large scale in the Cave without haptics seems to be more difficult to control than our smaller scale fishtank-based implementations. We think the reason is that it is more difficult to keep the trailing hand close to the path of the mark that has already been drawn, so when that trailing hand moves to advance the drawing, its projection onto the mark may not be where the artist expects, causing the drawn portion of the mark to advance more quickly than intended.

In summary, we present two pilot studies investigating the issue of controlled free-form modeling. We describe insights gained from these initial investigations and a plan for a new methodology that addresses more interesting and refined questions than our pilot work. Our new direction will yield both quantitative and qualitative insights into the ease of use and level of control afforded by our haptic-based tools.

Chapter 9

Conclusions and Future Work

Figure 9.1: Study of a free water jet issuing from a square hole into a pool. Leonardo da Vinci, c. 1508-1509

The combination of artist and scientist working together is perhaps most celebrated in the life and work of Leonardo da Vinci. Even today as students struggle to understand the complexities of fluid dynamics (a relevant scientific topic in this thesis), we turn to da Vinci's intricate studies and visualizations of water for insight and understanding. The work reproduced in Figure 9.1 probably represents the world's first use of scientific visualization to study turbulent flow. da Vinci's drawing and his words describe intricate details of the flow. He writes, "the water has eddying motions, one part of which is due to the principal current, the other to the random and reverse motion." [17] da Vinci's work is a timeless reminder of successful collaboration between art and science.

As science has advanced in the years since da Vinci, we have developed new tools for collecting scientific data and new formats for displaying scientific visualizations. In many ways, the tools have changed, but the power of the artist and scientist collaboration remains the same. [11, 20]

The contributions of this thesis are driven by the goal of enabling more effective collaborations between artists and scientists in today's 3D, virtual reality based visualization mediums. The fundamental challenge we address is discovering how best to leverage artistic skill in creating 3D models, while providing a rich and controllable enough environment for artists to tackle difficult modeling subjects. For artists to make significant contributions with today's scientific tools, they need to be able to design and create directly within mediums like VR. They need to be able to work with natural interaction styles that leverage their existing artistic skill and technique. Finally, the interaction and rendering styles that they employ need to be sufficiently controllable and expressive for them to be able to depict extremely complex subjects. Our work advances the state of the art in all of these areas.

In the style of the computer science toolsmith described by Fred Brooks [7], our applications in science and visual art drive the development of our tools. A chief motivator in our work is the ability to depict clarified and refined 3D form. This has obvious benefits for depicting scientific phenomena. We present scientific visualizations of bat flight that would be impossible to create without the additional control provided by our tools. Our driving artistic applications are chosen carefully to push this issue of control as well. For example, subjects in artistic anatomy demand a level of sophistication and clarity of representation that is not needed for more abstract or simplified subjects.

Our approach is based upon direct, sweeping 3D input achieved through six degree of freedom input devices coupled with virtual reality display technologies. As part of our initial work, we present the CavePainting system, which is an exploration into artistic interfaces motivated directly by the Cave VR form factor. In contrast to related free-form modeling approaches in virtual reality [14, 55], CavePainting interactions are tailored toward exploring the unique possibilities of the Cave environment. CavePainting is also unique among free-form modeling tools in the aesthetic that it fosters. Rather than completely specifying surfaces of 3D forms, the form is suggested through many brush strokes which visually combine to reveal the sense of a more solid object. This style is rooted in traditional painting, like that of the Impressionists, but our 3D version of this style of brushwork appears quite different and provides a novel, hand-crafted aesthetic for computer graphics.

As part of the effort to bring this style of artistic tool to bear on scientific problems, we present a methodology for collaborating with artists on visualization design problems directly within VR. A key contribution of this work is a data-driven, artistic visualization design tool that combines the natural interface of CavePainting with a strong connection to multi-variate data. This provides artists and designers with a platform for intuitive design of glyph-based, multi-variate visualizations of fluid flow datasets in VR. Many of our contributions in this area are structured around teaching artists, scientists, and computer scientists to collaborate on visualization problems in the style of the renaissance teams described by Donna Cox. [11] We present methods, results, and insights gathered from these experiences with teaching collaborative visualization design to students from Brown and the Rhode Island School of Design.

One of the most significant contributions of this thesis is our presentation of two-handed, haptic drawing interfaces for additional control in modeling via sweeping, free-form input. The need for these techniques is motivated by the level of clarity of form required for artists to work with subjects in science and medicine. Our approaches are easier to learn and use than CAD or Maya style, full featured, 3D modeling systems, and provide more direct, artistic control over the resulting form than 3D modelers based on intuitive 2D input like Teddy. [30] In addition, by working with a true 3D implementation of a two-handed drawing technique, rather than a two-step 3D curve drawing process as in related approaches [26], we make working quickly and creating many 3D curves as part of model practical. The addition of haptic constraints makes this true, 3D interaction work for artists. The haptics also allow us to investigate more expressive visual qualities. We present techniques that allow artists to adjust the line weight of a 3D mark while drawing it, an extension motivated by charcoal drawing that allows separate simultaneous control of line thickness and color value, and a two handed approach to smudging charcoal-like pigment in 3D to suggest the form of a surface.

Finally, we present a user study of our haptic-based tools that yields qualitative and quantitative insight into the ability for artists to control them for use in serious artistic exploration. The study is organized in conjunction with a RISD winter session class that serves as a framework for familiarizing artists with VR and our tools. Subjects in the study are drawn from students in the class. The study explores the relative importance of different aspects of the haptic interface to try to better understand the theory behind it, and also strives to establish some notion of how quickly artists can learn to work with these tools, giving us a measure of how intuitive our two-handed, haptic drawing techniques are.

9.1 Proposed Additions

The contributions outlined above reflect our sense of the contributions of the completed thesis, and are presented in the format of a draft of the thesis. As described in the outline of contributions in the introduction and throughout this thesis, several of these contributions are still works in progress.

There are three major pieces of the thesis that still need to be completed. The first is finishing the development of the two-handed, haptic-based interactions for working with virtual charcoal, including the editing feature and smearing of pigment as described in Chapter 5. At this point, we are confident we have these interactions fairly well thought out and we have some preliminary implementations of each of them. We are proceeding to work on them with the SIGGRAPH 2006 paper deadline in mind. Thus, we anticipate wrapping up this item into a solid contribution by late January.

The second area of work still in progress is the user study of the haptics-based tools. This is also in full swing at the moment. We have just begun teaching a RISD winter session course to train artists with our tools in preparation for this user study. The course is structured as a six week long collaborative study project under the guidance of Fritz Drury. As the course progresses and the artists become more familiar with working in VR and with our tools, we will be able to draw from this pool of trained artists to evaluate our tools with quantitative and qualitative assessments. The timeline for this class is to finish by February 10, so we anticipate having some

Work	Expected Completion Date
Editing via pulling with haptic-based system	Jan. 25 (SIGGRAPH)
2-handed smearing of pigment with haptic-based system	Jan. 25 (SIGGRAPH)
Refinement of charcoal drawing feature and examples of use	Jan. 25 (SIGGRAPH)
User study with RISD artists	Feb. 10 (end of winter session)
Analysis of data from user study	March 10
Import of data for bat visualization	March 10
Initial data-less designs for bat visualizations	March 10
Final bat visualizations	April 10

Table 9.1: Remaining significant portions of the thesis.

Thesis Proposal Presentation	Feb. 6, 2:30-4:00pm, in Lubrano
The Last Week Before SIGGRAPH 06	July 24 - 28

Table 9.2: Other significant dates.

meaningful results by that point or soon after.

The third area of work still in progress is the visualization of bat flight data described in Chapter 7. In fact, some of the students in the RISD winter session class have expressed interest in working with this subject matter during the class. In that case, work will continue on this subject very quickly. Dr. Swartz is very excited about this project as am I. We can count on support from her in this investigation and I expect that we can make a contribution of some insightful visualizations and illustrations in return.

Table 9.1 lists a bit more detailed description of remaining significant tasks with expected completion dates. Note that many of these are bounded by hard deadlines: The SIGGRAPH paper submission deadline, and the end of the RISD winter session class. Thus, there is a significant motivation to complete this work within the timeframe setup here.

Table 9.2 lists some other important dates. My proposal presentation is scheduled for Feb. 6. I would like to schedule my final defense during the summer at some point before everyone leaves for SIGGRAPH, which would mean late July.

In summary, the proposed additions to this thesis are setup to fall in place within a reasonable timeframe. With a potential SIGGRAPH paper, insight from a user study of trained artists, and the application of our latest techniques to a difficult visualization problem, these remaining contributions can add significantly to the impact of this body of work.

Bibliography

[1] *Picasso the Printmaker: Graphics from the Marina Picasso Collection*. Dallas Museum of Art, 1983.

[2] John Alex. *Hybrid Sketching: A New Middle Ground Between 2- and 3-D*. PhD thesis, Massachusetts Institute of Technology, 2005.

[3] Ravin Balakrishnan, George Fitzmaurice, Gordon Kurtenbach, and William Buxton. Digital tape drawing. In *ACM symposium on User interface software and technology*, pages 161–169, 1999.

[4] David Bourguignon. *Interactive Animation and Modeling by Drawing – Pedagogical Applications in Medicine*. PhD thesis, Institut National Polytechnique de Grenoble, 2003.

[5] David Bourguignon, Marie-Paule Cani, and George Drettakis. Drawing for illustration and annotation in 3d. In *EUROGRAPHICS Computer Graphics Forum*, pages C114–C122, 2001.

[6] Bill Brody and Chris Hartman. Painting space with BLUI. In *SIGGRAPH '00 Conference Abstracts and Applications*, page 242, 2000.

[7] Frederick P. Brooks Jr. The computer scientist as toolsmith - studies in interactive computer graphics. In *Proceedings of Information Processing 1977*, pages 625–634, 1977.

[8] Jeff Butterworth, Andrew Davidson, Stephen Hench, and Marc. T. Olano. 3dm: a three dimensional modeler using a head-mounted display. In *symposium on Interactive 3D graphics*, pages 135–138, 1992.

[9] William Buxton, George Fitzmaurice, Ravin Balakrishnan, and Gordon Kurtenbach. Large displays in automotive design. *IEEE Computer Graphics and Applications*, 20(4):68–75, 2000.

[10] James H. Clark. Designing surfaces in 3-d. *Communications of the ACM*, 19(8):454–460, 1976.

[11] Donna Cox. Using the supercomputer to visualize higher dimensions: An artist's contribution to scientific visualization. *Leonardo*, 21(3):233–242, 1988.

[12] Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the cave. In *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 135–142, 1993.

[13] Michael Deering. High resolution virtual reality. In *SIGGRAPH '92 Conference on Computer graphics and interactive techniques*, pages 195–202, 1992.

[14] Michael F. Deering. Holosketch: a virtual reality sketching/animation tool. *ACM Transactions of Computer-Human Interaction*, 2(3):220–238, 1995.

[15] Michael F. Deering. The holosketch vr sketching system. *Communications of the ACM*, 39(5):54–61, 1996.

[16] Jacques Dupin. *Miró*. Abrams, New York, N.Y., 1993.

[17] Mohamed Gad el Hak. *FLow Control: Passive, Active, and Reactive Flow Management*. Cambridge University Press, 2000.

[18] Gregg E. Favalora. Volumetric 3d displays and application infrastructure. *IEEE Computer*, 38(8):37–44, 2005.

[19] Balakrishnan-R. Kurtenbach G. & Buxton W. Fitzmaurice, G. An exploration into supporting artwork orientation in the user interface. In *Proceedings of the 1999 ACM Conference on Human Factors in Computing Systems*, pages 167–174, 1999.

[20] Felice Frankel. *Envisioning Science: The Design and Craft of the Science Image*. MIT Press, 2002.

[21] Fundació Joan Miró. Joan Miró, the gold of the azure, 1967, 2004.

[22] Tinsley A. Galyean and John F. Hughes. Sculpting: An interactive volumetric modeling technique. In *Proceedings of SIGGRAPH 91: Computer Graphics Proceedings, Annual Conference Series*, pages 267–274, 1991.

[23] Arthur D. Gregory, Stephen A. Ehmann, and Ming C. Lin. intouch: Interactive multiresolution modeling and 3d painting with a haptic interface. In *IEEE Virtual Reality Conference*, page 45, 2000.

[24] Jen Grey. Human-computer interaction in life drawing, a fine artists perspective. In *In Proceedings of the Sixth International Conference on Information Visualisation*, 2002.

[25] Tovi Grossman, Ravin Balakrishnan, Gordon Kurtenbach, George Fitzmaurice, Azam Khan, and Bill Buxton. Interaction techniques for 3d modeling on large displays. In *Symposium on Interactive 3D graphics*, pages 17–23, 2001.

[26] Tovi Grossman, Ravin Balakrishnan, Gordon Kurtenbach, George Fitzmaurice, Azam Khan, and Bill Buxton. Creating principal 3d curves with digital tape drawing. In *SIGCHI conference on Human factors in computing systems*, pages 121–128, 2002.

[27] Song Han and Gérard Medioni. 3dsketch: modeling by digitizing with a smart 3d pen. In *ACM international conference on Multimedia*, pages 41–49, 1997.

[28] Pat Hanrahan and Paul E. Haeberli. Direct wysiwyg painting and texturing on 3d shapes. In *Proceedings of SIGGRAPH 90: Computer Graphics Proceedings, Annual Conference Series*, pages 215–223, 1990.

[29] Jing Hua and Hong Qin. Haptics-based dynamic implicit solid modeling. *IEEE Transactions on Visualization and Computer Graphics*, 10(5):574–586, 2004.

[30] Takeo Igarashi, Satoshi Matsuoka, and Hidehiko Tanaka. Teddy: a sketching interface for 3d freeform design. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 409–416, 1999.

[31] Cullen Jackson, David Karelitz, Sean A. Cannella, and David H. Laidlaw. The great potato search: The effects of visual context on users feature search and recognition abilities in an IVR scene. In *Poster Proceedings of IEEE Visualization*, 2002.

[32] Chris Johnson, Robert Moorhead, Tamara Munzner, Hanspeter Pfister, Penny Rheingans, and Terry S. Yoo. Nih-nsf visualization research challenges report. Draft October 2005, October 2005.

[33] David B. Karelitz, Daniel F. Keefe, and David H. Laidlaw. Using CavePainting to create scientific visualizations. IEEE Visualization 2003 Poster Compendium, October 2003.

[34] Olga Karpenko, John F. Hughes, and Ramesh Raskar. Free-form sketching with variational implicit surfaces. *Computer Graphics Forum*, 21(3), 2002.

[35] Olga Karpenko, John F. Hughes, and Ramesh Raskar. Epipolar methods for multi-view sketching. In *Eurographics Workshop on Sketch-Based Interfaces*, 2004.

[36] Daniel F. Keefe, Daniel Acevedo Feliz, Tomer Moscovich, David H. Laidlaw, and Joseph J. LaViola Jr. CavePainting: A fully immersive 3D artistic medium and interactive experience. In *Proceedings of ACM Symposium on Interactive 3D Graphics 2001*, pages 85–93, 2001.

[37] Daniel F. Keefe, David B. Karelitz, Eileen L. Vote, and David H. Laidlaw. Artistic collaboration in designing VR visualizations. *IEEE Computer Graphics and Applications*, 25(2):18–23, March/April 2005.

[38] Daniel F. Keefe and David H. Laidlaw. A haptic interface for creating smooth 3d curves with varying line weight. SIGGRAPH 2005 Poster, July 2005.

[39] Michael Kirby, H. Marmanis, and David H. Laidlaw. Visualizing multivalued data from 2D incompressible flows using concepts from painting. In *Proceedings of IEEE Visualization 1999*, pages 333–340, 1999.

[40] David H. Laidlaw, Eric T. Ahrens, davidkremers, Matthew J. Avalos, Carol Readhead, and Russell E. Jacobs. Visualizing diffusion tensor images of the mouse spinal cord. In *Proceedings of IEEE Visualization 1998*, pages 127–134, 1998.

[41] Joseph J. LaViola, Daniel F. Keefe, Robert C. Zeleznik, and Daniel Acevedo Feliz. Case studies in building custom input devices for virtual environment interaction. VR 2004 Workshop: Beyond Glove and Wand Based Interaction, March 2004.

[42] LightSpace Technologies. *LightSpace DepthCube Technology White Paper*.

[43] Eric B. Lum and Kwan-Liu Ma. Hardware-accelerated parallel non-photorealistic volume rendering. In *Proceedings of the International Symposium on Nonphotorealistic Animation and Rendering*, 2002.

[44] D. P. Mapes and J. M. Moshell. A two-handed interface for object manipulation in virtual environments. *Presence*, 4(4):403–416, 1995.

[45] Lee Markosian, Jonathan M. Cohen, Thomas Crulli, and John Hughes. Skin: a constructive approach to modeling free-form shapes. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 393–400, 1999.

[46] Wille Mkel. Working 3d meshes and particles with finger tips. towards an immersive artists' interface. In *IEEE Virtual Reality Workshop*, pages 77–80, 2005.

[47] Stephen Ornes. High tech spelunking: Science and art cohabitate in a 21st century cave. 2005.

[48] Shigeru Owada, Frank Nielsen, Makoto Okabe, and Takeo Igarashi. Volumetric illustration: designing 3d models with internal textures. *ACM Transactions on Graphics*, 23(3):322–328, 2004.

[49] Cynthia B. Rubin and Daniel F. Keefe. Hiding spaces: A cave of elusive immateriality. ACM SIGGRAPH 2002 Conference Abstracts and Applications, July 2002.

[50] Cynthia B. Rubin and Daniel F. Keefe. Hiding spaces: A cave of illusive immateriality. In *The 11th International Symposium on Electronic Art Poster Proceedings*, 2002.

[51] Emanuel Sachs, Andrew Roberts, and David Stoops. 3-draw: A tool for designing 3d shapes. *IEEE Computer Graphics and Applications*, 11(6):18–26, 1991.

[52] Steven Schkolne. *3D Interfaces for Spatial Construction*. PhD thesis, California Institute of Technology, 2003.

[53] Steven Schkolne, Hiroshi Ishii, and Peter Schroder. Immersive design of dna molecules with a tangible interface. In *Proceedings of IEEE Visualization*, pages 227–234, 2004.

[54] Steven Schkolne and Daniel F. Keefe. Visual space after virtual reality. In *Proceedings of Young VR 2001*, Korea, December 2001.

[55] Steven Schkolne, Michael Pruett, and Peter Schröder. Surface drawing: creating organic 3d shapes with the hand and tangible tools. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 261–268, 2001.

[56] Christopher Schmandt. Spatial input/display correspondence in a stereoscopic computer graphic work station. In *SIGGRAPH '83: Proceedings of the 10th annual conference on Computer graphics and interactive techniques*, pages 253–261, 1983.

[57] Mark Schroering, Cindy Grimm, and Robert Pless. A new input device for 3d sketching. In *Vision Interface*, pages 311–318, 2003.

[58] Inc. SensAble Technologies. *Freeform: Feel the Difference*.

[59] Ben Shedd. Exploding the frame: Seeking a new cinematic language. *Society of Motion Picture and Television Engineers 135th Conference*, 1994.

[60] William Sherman and Alan Craig. *Understanding Virtual Reality: Interface, Application, and Design*. Morgan Kaufmann, 2003.

[61] Scott Snibbe, Sean Anderson, and Bill Verplank. Springs and constraints for 3d drawing. In *Proceedings of the Third Phantom Users Group*, 1998. M.I.T. Artificial Intelligence Laboratory Technical Report AITR-1643.

[62] Jason S. Sobel, Andrew S. Foresberg, David H. Laidlaw, Robert C. Zeleznik, Daniel F. Keefe, Igor Pivkin, George E. Karniadakis, Peter Richardson, and Sharon Swartz. Particle flurries: Synoptic 3d pulsatile flow visualization. *IEEE Computer Graphics and Applications*, 24(2):76–85, March/April 2004.

[63] Vibeke Sorensen. The contribution of the artist to scientific visualization, 1989.

[64] Ivan Sutherland. *Sketchpad: A Man Machine Graphical Communication System*. PhD thesis, Massachusetts Institute of Technology, 1963.

[65] Ivan Sutherland. A head-mounted three dimensional display. In *Proceedings of the AFIPS Fall Joint Computer Conference*, volume 33, pages 757–764, 1968.

[66] Karan Singh Tovi Grossman, Ravin Balakrishnan. Pointing and manipulation: An interface for creating and manipulating curves using a high degree-of-freedom curve input device. In *SIGCHI conference on Human factors in computing systems*, pages 185 – 192, 2003.

[67] Andries van Dam, Andrew Forsberg, David Laidlaw, Joseph LaViola, and Rosemary Michelle Simpson. Immersive virtual reality for scientific visualization: A progress report. *IEEE Computer Graphics and Applications*, 20(6):26–52, 2000.

[68] Kathy Walker. Virtual da vinci. *Shift Magazine*, 10(2), 2002.

[69] Colin Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann Publishers, 2000.

[70] Gerold Wesche and Hans-Peter Seidel. Freedrawer: a free-form sketching system on the responsive workbench. In *ACM symposium on Virtual reality software and technology*, pages 167–174, 2001.

[71] Barbara Ehrlich White. *Impressionists side by side: their relationships, rivalries, and artistic exchanges.* A.A. Knopf, New York, 1996.

[72] Robert C. Zeleznik, Kenneth P. Herndon, and John F. Hughes. Sketch: an interface for sketching 3d scenes. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 163–170, 1996.

[73] Robert C. Zeleznik, Joseph J. LaViola, Daniel Acevedo, and Daniel F. Keefe. Pop through buttons for virtual environment navigation and interaction. In *Proceedings of Virtual Reality 2002*, pages 127–134, 2002.