

Abstract of “Interactive 3D Drawing for Free-Form Modeling in Scientific Visualization and Art: Tools, Methodologies, and Theoretical Foundations” by Daniel F. Keefe, Ph.D., Brown University, May 2007.

This dissertation investigates free-form modeling tools driven by 3D drawing-style input and their utility for design, illustration, and visualization in science and art. 3D drawing interactions are developed and analyzed with respect to control, artistic expression, and the unique design methodologies they make possible.

A toolset of new interactive algorithms for 3D drawing called Drawing on Air is presented. The additional control provided by Drawing on Air was measured in a quantitative user-study in which it significantly outperformed multiple alternative techniques. Additionally, qualitative evaluations driven by art and scientific illustration lead to successful 3D modeling results of subjects previously too challenging to address effectively with 3D drawing techniques. To better understand this style of 3D computer input, statistical models grounded in theories of human perception and motor control were developed and analyzed with respect to experimentally collected data. Results help to highlight differences between the input techniques tested and lead to the definition of an index of difficulty for controlled, 3D, drawing-style input. Finally, a series of four experiments using 3D drawing tools for design of scientific visualizations is presented. These lead to several tool refinements and a new methodology for collaborative design of virtual reality visualizations called Scientific Sketching.

The major conclusions of the dissertation can be summarized briefly: 1. New tools for free-form modeling based on 3D drawing-style interaction can increase artists’ expressive power. 2. This style of computer interaction is useful for depicting challenging subjects in science and art. 3. We can better understand controlled, continuous 3D computer input through statistical models based on theories of human perception and motor control. 4. Visual experts, such as artists, can make important contributions to visual problems in science, but appropriate tools are required to make these collaborations productive.

Interactive 3D Drawing for Free-Form Modeling in Scientific Visualization and Art:
Tools, Methodologies, and Theoretical Foundations

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

May 2007

© Copyright May 2007 by Daniel F. Keefe

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 _____
_____ Fritz Drury, Reader
(Rhode Island School of Design)

Date _____
_____ John F. Hughes, Reader

Date _____
_____ Andries van Dam, Reader

Approved by the Graduate Council

Date _____
_____ Sheila Bonde
Dean of the Graduate School

Vita

Daniel Keefe was born, a Red Sox fan, on December 27, 1976 in Framingham, MA. Soon thereafter, he moved to North Carolina, where he lived with his family until attending Tufts University in Medford, MA in 1995. He received the bachelor's degree in computer engineering summa cum laude from Tufts University in 1999. In the fall of 1999 he moved to Providence, RI to attend Brown University, where in 2001 he received the master's degree in computer science for work developing the CavePainting system for artistic 3D modeling in virtual reality. He remained at Brown University to pursue a doctorate degree in computer science.

While at Brown, he has been a member of the scientific visualization research lab and the graphics group. He has also worked closely with collaborators at the Rhode Island School of Design, where he has co-taught courses, mentored several students, and continued his studies within the illustration department. Connections to both science and art are important themes in his research. He has published in the areas of scientific visualization, computer graphics, human-computer interaction, and digital art. He has also shown work in several art exhibitions. Selected publications and exhibitions relating to the work presented in this dissertation are cited below.

Selected Publications

- Daniel F. Keefe, Robert C. Zeleznik, and David H. Laidlaw. Drawing on air: Input techniques for controlled 3D line illustration. *IEEE Transactions on Visualization and Computer Graphics*, 2007. In press.
- Daniel F. Keefe, Daniel Acevedo, Jadrian Miles, Fritz Drury, Sharon M. Swartz, and David H. Laidlaw. Scientific sketching: An experimentally refined methodology for designing VR visualizations with artists. 2007. In review.

- 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, 2005.
- 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.
- Michael M. Kirby, Daniel F. Keefe, and David H. Laidlaw. Painting and Visualization. *The Visualization Handbook*, pages 873–891. Elsevier Inc., 2005.
- 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.
- 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.
- Joseph LaViola, Daniel Acevedo, Daniel F. Keefe, and Robert C. Zeleznik. Hands-free multi-scale navigation in virtual environments. In *Proceedings of ACM Symposium on Interactive 3D Graphics 2001*, pages 9–15, 2001.
- Cullen Jackson, Daniel Acevedo, David H. Laidlaw, Fritz Drury, Eileen Vote, and Daniel Keefe. Designer-critiqued comparison of 2D vector visualization methods: A pilot study. ACM SIGGRAPH 2003 Sketches and Applications, July 2003.
- Cynthia B. Rubin and Daniel F. Keefe. Hiding spaces: A cave of elusive immateriality. ACM SIGGRAPH 2002 Conference Abstracts and Applications, July 2002.

Exhibitions

- Rochester Institute of Technology Digital Arts Competition and Exhibition, with Gallery r (Rochester, NY), "Interactive Virtual Reality 3D Drawing: An Art Process" (Invited Exhibit), Video, Rochester, NY, May 11 – 19, 2007.
- Byte-Sized, Digital Miniatures, Viento y Agua Gallery, "La Guitarrista Gitana", Dye-sublimation print of CavePainted virtual sculpture, Long Beach, CA, August 8 – September 4, 2004.
- Byte-Sized, Digital Miniatures, Viento y Agua Gallery, "Wedding Day, Walking Away", Dye-sublimation print of CavePainted virtual sculpture, Long Beach, CA, August 8 – September 4, 2004.
- Boston CyberArts Festival, Bell Gallery and Creative Arts Council at Brown University, "Sailing a Dhow in Tanzania: A Cave Painting", Immersive, body-controlled virtual reality environment with sound, Providence, RI, April 26 and May 3, 2003.
- Boston CyberArts Festival, Bell Gallery and Creative Arts Council at Brown University, "La Guitarrista Gitana: A Cave Painting", Immersive, body-controlled virtual reality environment with sound, Providence, RI, April 26 and May 3, 2003.
- SIGGRAPH 2002 Art Gallery, "The Making of La Guitarrista Gitana: A Cave-Painting", Video, San Antonio, TX, July 21 – 26, 2002.

Acknowledgments

I would like to thank the many people who have contributed their ideas and energy to the work leading to this dissertation.

Members of the Brown scientific visualization lab and graphics group have been extraordinary collaborators in discussion, implementation, and paper writing. I am grateful, in particular, for the contributions of Bob Zeleznik, Daniel Acevedo, Andy Forsberg, Sharon Swartz, David Karelitz, David Eigen, Misha Zaitzeff, Tomer Moscovich, Morgan McGuire, Andy Wenger, Cullen Jackson, Joseph LaViola, Eileen Vote, Jadrian Miles, Loring Holden, Song Zhang, Cagatay Demiralp, and the Brown and RISD students of the course “Virtual Reality Design for Science.”

Likewise, I am grateful to artists and instructors at RISD for their support, feedback, and enthusiastic work with my computer tools. I would especially like to thank Harrison Love, Helen Zhu, Eric Benedict, Cynthia Beth Rubin, and Gisel Florez.

I am particularly grateful to my advisor and friend David Laidlaw, who has been a role model in my development as a scientist, teacher, and mentor. He has also provided an excellent environment for research and convinced me of the great excitement and rewards that can be found in interdisciplinary scientific research. I will always be grateful for his guidance and trust during my time at Brown. I am also thankful for his family’s understanding, support, and generosity, which have on many occasions been truly appreciated by my own family.

My appreciation also goes to the readers of this dissertation, Fritz Drury, John Hughes, and Andries van Dam, for their guidance in refining the work and their support throughout my time at Brown.

Fritz Drury also deserves thanks for his role as an advisor in visual studies. His consistent guidance and countless hours of support in my attempts to connect with RISD and the world of art, design, and illustration were invaluable to this work and

to my development as a teacher, artist, and scholar.

Finally, I would like to thank my friends and family. The “Gapas” (my rock band), the “Dingers” (my softball team), and the “Hackers” (my basketball team) have been there for the highs and also the “Ph.D. Blues” (which the Gapas actually recorded).

To my family, Delia, Frank, Anne, and John, thanks for your love and support. To my wife, Susi, I could not have done this without you. Thank you for your constant encouragement and love, and thank you for showing me the world, quite literally! To “Baby Keefe” ... can’t wait to meet you, thanks for a great deadline for the final push!

Contents

List of Tables	xiv
List of Figures	xv
1 Introduction	1
1.1 Why 3D Drawing?	4
1.2 Thesis Statement	6
1.3 Overview of the Dissertation	7
2 Review of Artistic Free-Form Modeling	8
2.1 Techniques Utilizing Direct 3D Input	10
2.1.1 Early Work in 3D Computer Input	10
2.1.2 Recent Work in 3D Computer Input	13
2.2 Techniques Utilizing Large-Scale 2D Input	15
2.3 Techniques Utilizing Desktop-Scale 2D Input	15
2.4 Summary of Related Work	16
3 Drawing on Air: Techniques for Controlled and Expressive 3D Line Illustration	17
3.1 Related Work	21
3.1.1 Bimanual Approaches to Drawing Lines	21
3.1.2 Freehand 3D Drawing Systems	21
3.1.3 Haptic-Aided Drawing and Modeling	22
3.2 Drawing on Air	22
3.2.1 One-Handed Drag Drawing	24

3.2.2	Two-Handed Tape Drawing	28
3.2.3	Integrated One- and Two-Handed Drawing	32
3.2.4	Brush Model for 3D Geometric Pigment	35
3.3	Illustration Results	36
3.3.1	Illustrations of Bat Flight	37
3.3.2	Artistic Anatomy	40
3.4	Discussion	44
3.4.1	Strategies for Effective Use	44
3.4.2	Controls for Line Weight	45
3.4.3	3D Tape Drawing Without Haptics	46
3.4.4	Nondominant-Hand Offset Mode	47
3.5	Conclusion	47
4	The Drawing Control Experiment	48
4.1	Conditions and Hypotheses	48
4.2	Methodology	49
4.3	Training	50
4.4	Participants	51
4.5	Results	51
4.6	Analysis of Results	53
4.7	Discussion	57
4.7.1	Appropriateness of the Task	57
4.7.2	Importance of Training and Depth Cues	58
4.7.3	User-Guided Drawing and Controlled 3D Input	59
4.8	Summary and Conclusions	62
5	Models for User Performance in 3D Tracing Tasks	63
5.1	Review of Action Laws in HCI Literature	64
5.2	Review of Related Theory in Neuroscience and Motor Control	66
5.2.1	The Power Law	66
5.2.2	Anisotropy in Control of Visually Guided 3D Motion	68
5.3	Local Models for 3D Tracing Performance	70
5.3.1	Structure and Methodology	71

5.3.2	Connection to the “Tunnel-Width” Term of the Steering Law	71
5.3.3	Model 1: Local Orientation	72
5.3.4	Model 2: Incorporating Local Curvature	74
5.3.5	Model 3: Incorporating Interactions	78
5.3.6	Comparison and Analysis of Local Models	79
5.4	Global Models for 3D Tracing Performance	83
5.5	Suggested Follow-On Investigations	86
5.5.1	Bending the Steering Law: Curved Paths and Paths of Equal Length	86
5.5.2	Considerations in Testing 3D Path Steering	88
5.6	Summary and Conclusions	90
6	Designing VR Visualizations with Artists: Experiments and Methodology	91
6.1	Related Work	95
6.1.1	Artistic Collaboration for Scientific Visualization	95
6.1.2	Software Rapid Prototyping	96
6.1.3	Toolkits for Rapid Visualization Development	97
6.2	Background in Artistic Critique	97
6.3	Experiments in Visualization with Artists	97
6.3.1	Experiment 1: CavePainting Visualizations	98
6.3.2	Experiment 2: Data-Driven Glyph Sketching	101
6.3.3	Experiment 3: Appropriate Artistic Control for Science	103
6.3.4	Experiment 4: Time-Varying, Interactive Prototyping	107
6.4	Scientific Sketching: A Collaborative VR Design Methodology	111
6.4.1	Paper Sketching	112
6.4.2	VR Sketching	113
6.4.3	VR Prototyping: Drawing, Wizardry, and Connections to Data	113
6.4.4	Implementation of the Resulting Visual Specification	114
6.5	Analysis	115
6.5.1	The Value of Collaborative Design	115
6.5.2	Characterization of Design Tools	117
6.6	Discussion	119

6.6.1	Toward Renaissance Visualization Designers	119
6.6.2	Limitations and Future Directions	120
6.7	Summary and Conclusions	120
7	Critique-Based Evaluation of Art Applications	122
7.1	Introduction to Evaluations in Artistic Anatomical Illustration	123
7.2	Hypothesis	124
7.3	Methodology	124
7.4	Results and Individual Critique Analysis	125
7.4.1	Torsos: Descriptive and Dynamic	125
7.4.2	Arms: Descriptive and Dynamic	128
7.4.3	Hands: Descriptive and Dynamic	130
7.4.4	Figures: Gesture and Economy of Line	135
7.4.5	Faces	136
7.5	Conclusions from Anatomical Studies	138
7.5.1	Sculpting vs. Suggesting Form	138
7.5.2	Economy of Line and Occlusion	138
7.5.3	Use of Color Together with CG Lighting	139
7.5.4	Stereo Verses 2D Viewing	139
7.5.5	Comparison of Cave and Fishtank Environments	139
7.6	Additional Artistic Investigations	140
7.6.1	Contrasts in Styles of Depiction: Painting vs. Sculpting	140
7.6.2	Interactive Environments Utilizing the Space of the Cave	142
7.7	Discussion of Cave-Based Artistic Media	144
7.7.1	Movement in the Space of the Cave	144
7.7.2	Large, Body-Scale Interaction	145
7.7.3	Prop-Based Interaction	145
7.7.4	Importance of Artistic Process	147
7.8	Conclusions	149
8	Discussion and Conclusions	150
8.1	Discussion and Future Work	150
8.1.1	Drawing Circles and Other Specific 3D Input Tasks	150

8.1.2	Reuse of Artist-Created Visualization Designs: Visualization by Analogy	152
8.1.3	Additional New Directions	153
8.2	Summary of Primary Contributions and Conclusions	154
8.2.1	Interactive Algorithms for 3D Drawing	154
8.2.2	Statistical Models for 3D Tracing Tasks	155
8.2.3	Artistic Collaboration in Scientific Visualization	156
8.2.4	Computer Tools for 3D Artistic Illustration	157
8.2.5	General Conclusions	158

Bibliography **159**

★ Parts of this dissertation have been previously published with co-authors, Robert C. Zeleznik and David H. Laidlaw [56].

List of Tables

4.1	Results from the Drawing Control Experiment.	52
5.1	Best values for the exponent, β , from the Power Law.	78

List of Figures

1.1	Lines drawn with charcoal.	2
1.2	Charcoal drawing of a runner.	3
1.3	3D structure of bat wings during flight.	4
3.1	3D illustration of a bat posed in flight created with Drawing on Air.	18
3.2	Drawing on Air virtual reality hardware.	20
3.3	Progression of a Drawing on Air one-handed, drag-mode interaction.	25
3.4	Custom elastic controller device for adjustment of line weight while drawing.	28
3.5	Progression of a Drawing on Air tape mode interaction.	30
3.6	Drawing on Air visual feedback.	33
3.7	A 3D bat illustration created with Drawing on Air.	37
3.8	A 3D bat skeleton drawn on top of motion-capture data.	39
3.9	3D bat skeleton and muscles drawn with Drawing on Air.	41
3.10	Bearded man.	42
3.11	A Swahili bride wearing a green veil.	43
3.12	Muscles and tendons in the hand.	43
4.1	The Drawing Control Experiment tracing task.	49
4.2	Comparison of positional error for four input conditions.	53
4.3	Comparison of directional error for four input conditions.	54
4.4	Comparison of drawing time for four input conditions.	54
4.5	One participant's best tracing results.	56
4.6	Requirements for controlled 3D drawing-style input.	59
5.1	Target touching task relating to Fitts' Law.	64

5.2	Steering through a tunnel.	64
5.3	Steering Law derivation.	65
5.4	Figure eight drawing task typical of Power Law studies.	67
5.5	Trends in measured directional error based on local curve orientation.	73
5.6	Trends in correlation of D_K with experimental data for a range of values for β for directional error.	75
5.7	Trends in correlation of D_K with experimental data for a range of values for β for positional error.	76
5.8	Trends in correlation of D_K with experimental data for a range of values for β for drawing time.	76
5.9	Scatter plot of local curvature and measured directional error.	77
5.10	Scatter plot of local curvature and measured positional error.	77
5.11	Scatter plot of local curvature and measured drawing time.	78
5.12	Comparison of local models for directional error.	79
5.13	Comparison of local models for positional error.	80
5.14	Comparison of local models for drawing time.	80
5.15	Local measured directional error plotted against the local index of difficulty, D_3	82
5.16	Local measured positional error plotted against the local index of difficulty, D_3	82
5.17	Local measured drawing times plotted against the local index of difficulty, D_3	83
5.18	Trends in directional error versus global models for difficulty.	84
5.19	Comparison of global models for directional error.	85
5.20	The spiral tunnel steering task tested by Accot and Zhai.	87
5.21	Alternative spiral tunnel steering tasks.	88
6.1	Designs for VR scientific visualizations in traditional artistic media.	92
6.2	Experiment 1: A free-form 3D modeling tool, CavePainting, was used to sketch designs for visualizations directly in VR.	99
6.3	Design for a data-driven flow glyph inspired by natural sea creature forms.	101
6.4	Experiment 2: An animated, data-driven visualization sketch.	102

6.5	Experiment 3: Additional control provided by Drawing on Air aids in depicting science.	106
6.6	Experiment 4: Progression of a visualization design from paper to interactive VR prototype.	110
6.7	Overview of the Scientific Sketching methodology.	111
6.8	Goals and deliverables for each stage of design.	112
7.1	A descriptive depiction of a torso.	125
7.2	A dynamic depiction of a torso.	126
7.3	A descriptive depiction of the scapula, shoulder, and arm bones. . . .	128
7.4	Sketches of arms in action.	129
7.5	Hands pulling in a fishing net.	130
7.6	Hand holding a baseball.	131
7.7	Hands wringing out a towel.	131
7.8	Anatomical drawing of a hand.	132
7.9	Seated woman.	134
7.10	Woman undressing.	134
7.11	A Swahili bride wearing a green veil.	136
7.12	A bearded man.	137
7.13	The CavePainting on the left was inspired by the oil painting on the right.	140
7.14	CavePainting of a lion fish, by Helen Zhu.	141
7.15	CavePainting of a falcon, by Helen Zhu.	141
7.16	Cave imagery shifts from the interactive art space, <i>Hiding Spaces</i> . . .	143
7.17	CavePainting prop table.	146
7.18	A physical bucket prop is used to throw virtual paint onto the walls and floor of the Cave.	147
7.19	A timeline widget on the floor of the Cave controls a visualization of the painting process.	148
7.20	Visualization of the painting process for <i>La Guitarrista Gitana</i>	148
8.1	Proposed system for surgical path planning.	151
8.2	Proposed visualization by analogy, inspired by work on Image Analogies. .	152

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 his brother Theo [88]

Van Gogh’s description of his artistic subjects fascinates me. Serious, difficult, beautiful: these words are clearly written by someone driven to intense study of a formidable problem. In trying to capture these difficult subjects, Van Gogh has picked an awesome challenge for his life’s work.

In scientific visualization and computer graphics, we face a similar challenge. We work with serious, difficult, and beautiful subjects, and we strive to make the most effective visual representations of these subjects possible in order to capture the complexity of natural forms, facilitate scientific discovery, and make sophisticated artistic statements.

In the age of computers, there is no tool more powerful for conveying this complexity; yet artists using computers lack something that painters like Van Gogh enjoyed. There is a richness and expressiveness in physical media, such as painting and drawing, that we rarely find in virtual, computer-based media.

In computer graphics, we frequently try to capture the missing ingredient through simulation. We derive models for the absorption and layering of watercolor paint




Figure 1.1: Lines drawn with a piece of vine charcoal (shown at the bottom). Charcoal is the simplest of tools, but its tight connection to the artist makes possible a rich level of expression.

on paper [18, 66] and we create rendering styles based on established artistic technique [34, 50].

An approach that complements simulation of physical media in computer form is capturing rich and expressive artistic computer input. Figure 1.1 is a simple physical-world example of artistic input. Shown here are a sample of marks sketched with vine charcoal in less than two minutes. Vine charcoal, which is nothing more than a burnt stick, is undoubtedly one of the simplest tools used by humans, yet look at the rich variation in expression it provides! In one sense, this richness comes from the physical interaction of the charcoal with the paper, but in this case an even more important aspect of the expression comes from the rich input of the artist. These marks are created by continuously varying position, pressure, speed, angle relative to the page, and orientation of the charcoal within the fingers while drawing. This high-dimensional input space yields what artists describe as a rich, expressive medium. Interestingly, the medium is useful both for quick initial design sketches, the type an




Figure 1.2: A charcoal drawing of a runner about to jump a hurdle. Variation in the strength and energy of the charcoal marks is used to indicate muscles and tendons in action during this pose.



Figure 1.3: Three stages of the downstroke of a cynopterus brachyotis in flight. Studying flight in bats requires tremendous 3D understanding because the motion of their wings is so complex.

architect or painter might make in preparation for a larger work, and for more refined drawings.

Figure 1.2 is a drawing I made in an illustration course on artistic anatomy. The variation in the strength and energy of the lines achievable via charcoal is impressive. That variation is used here to highlight the muscles and tendons that are in action as a runner prepares to jump over a hurdle. Capturing this level of expression in computer-based tools is important because it enables artists to clearly represent complicated visual subjects like this one.

The computer graphics literature contains some important examples of work that captures both the physicality (through simulation) and rich input characteristic of traditional artistic media. A notable example is the haptic-aided painting tool DAB [6], presented at SIGGRAPH 2001. DAB combines a 3D computer model of a paintbrush with a simulation of mixing and layering oil paint. A Phantom input/output device is used to provide high-degree-of-freedom input and force-feedback output to the user through a pen attached to a robotic armature. In chapter 3, we introduce a technique for controlled 3D drawing that employs similar high degree-of-freedom input/output, but for use in 3D modeling rather than simulation of the 2D painting experience.

1.1 Why 3D Drawing?

Many problems in art and science demand a 3D treatment, and our work is driven by these cases. Figure 1.3 shows an example from a scientific problem explored in detail

in this dissertation. In conjunction with Dr. Sharon Swartz of Brown's Department of Ecology and Evolutionary Biology, we have been working to understand the flight of bats. The photographs in Figure 1.3 hint at the 3D complexity of this problem. The mechanics of bat wings are similar in number of degrees of freedom to the human hand. Bats have an amazing ability to reshape their flexible wing membrane during flight. Consequently, bats are among the most highly maneuverable flying animals, and analysis of their flight is far more complex than that of insects, birds, and fixed-wing aircraft. Understanding this problem has potential to unlock the secrets of the evolution of flight in the only flying mammals and also has potential applications to design of unmanned aircraft.

Traditional 2D illustrations of bat anatomy are severely limited in their ability to capture the 3D complexity inherent to this problem. Thus, one of the goals of our work is to create 3D anatomical illustrations of situations, like bats in flight, that demand a 3D treatment. Together with Dr. Swartz, we hypothesize that artistic illustrations of bats posed in flight will improve our ability to understand flight dynamics. In particular, 3D understanding from these illustrations is likely to help formulate hypotheses about which muscles are active at different points during the wing beat cycle.

The use of 3D graphics, especially in virtual reality (VR), is promising for representing this complexity. Many researchers have now quantified the benefits of head-tracked stereo vision in increasing understanding of complex 3D phenomena [82, 105]. Unfortunately, it is difficult for artists to work in VR spaces with the rich expression and visual thinking they enjoy in more traditional media. How can artists best design and create for illustration and visualization in virtual environments?

One possibility is to use desktop-based 3D modeling applications, as are common in industry [1]. As we know from special effects in movies, these tools are quite capable of producing sophisticated 3D depictions. Unfortunately, obtaining this level of detail is extremely time consuming: It can take weeks to develop complex 3D models. Then, additional time is needed for artists to tweak the rendering of the model by "painting" textures onto it and adjusting lighting and background elements to clarify the meaning of the image. Usually, design begins in traditional media, like sketching on paper. Another limitation of traditional 3D modeling packages

is the discontinuity between these initial design phases and the advanced modeling techniques required to generate complex 3D models.

It stands to reason that if we can capture more of the immediacy, expressiveness, and control of media like drawing in a 3D computer graphics context, we may greatly increase the power of artists to investigate complex 3D subjects. In fact, there is a rich history of research into using direct, sweeping 3D input to drive intuitive, artist-accessible 3D modeling. Chapter 2 presents a complete overview of this work.

Our investigations in this area began with a system called CavePainting [54]. In CavePainting, as in other tools based on a 3D drawing paradigm [22, 78, 106], virtual 3D form is created by moving through the air a tracked hand-held stylus. The results of this work had a great impact on the direction of this thesis. In particular, we found that artists can indeed engage with 3D drawing systems to create compelling artistic work [53, 71, 77], but that the loose style characteristic of 3D drawing, while exciting for artistic work, is less appropriate for refined illustration. We confirmed that CavePainting, like other computer interactions based on continuous 3D input, suffers from a lack of precise control. This limits artists' ability to address subjects in refined illustration that go beyond quick design and gesture sketches.

1.2 Thesis Statement

This brings us to the central thesis of this dissertation:

Free-form modeling based on rich, controllable 3D-drawing-style interactions can increase the expressive power of artists in conveying difficult subjects in art and scientific visualization that require illustration-quality precision.

Our research is grounded in the belief that, given a better understanding of the process of 3D drawing and more appropriate algorithms for tying this input to creating 3D form, we can develop techniques that overcome the current limitations in control while simultaneously maintaining the advantages of 3D drawing: immediacy, intuitiveness, and richness. Further, we expect that, given this additional control,

free-form modeling based on 3D drawing-style interaction will be useful in two driving application areas. The first is artistic anatomical illustration. We focus our artistic efforts specifically on illustrations of anatomy because they are difficult and require high precision in the drawing tools. The second area is scientific visualization. The premise behind these investigations is that visual experts in art, illustration, and design have valuable insights that can help solve visual problems arising in science. Making these collaborations possible in today's visualization environments requires developing appropriate visual design tools. This dissertation works to support the claim that modeling based in 3D drawing holds an important key to effective visual design and visual thinking in 3D visualization environments, such as VR.

1.3 Overview of the Dissertation

Chapter 2 presents a review of artistic free-form modeling. Chapter 3 introduces Drawing on Air, a new toolset of interactive algorithms for controlled and expressive 3D drawing. Chapter 4 presents the Drawing Control Experiment, a formal, quantitative comparison of user control with Drawing on Air compared to alternative techniques. Also in this chapter is a discussion of controlled 3D drawing and user-guided drawing techniques, of which Drawing on Air is an example. In Chapter 5, local and global statistical models based on theory in human-computer interaction, psychology, and neuroscience are derived for describing user performance in 3D tracing tasks. The models are evaluated with experimental data, and the results suggest important directions for future research. Chapter 6 presents four experiments related to refining artistic design tools for VR scientific visualization. These lead to a methodology for artist-scientist-technologist collaboration in visualization design called Scientific Sketching. Chapter 7 presents evaluations and conclusions from artistic applications of 3D drawing, including use of CavePainting and Drawing on Air, in creating virtual, interactive art spaces and artistic anatomical illustrations.

Chapter 2

Review of Artistic Free-Form Modeling

Since Ivan Sutherland's Sketchpad system [89], marking the creation of the field, modeling form has been a key challenge in interactive computer graphics. Sutherland's vector-graphics-based system, years ahead of its time, included 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 (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 snap-to-grid mode. While extremely useful in 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 lengthy training they require. 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 well to 3D, and the goal of finding a 3D extension to the immediacy and gestural qualities of traditional media 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 can be seen as filling the role of the artist's charcoal pencil. 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 similar systems.

Several characteristics are useful in classifying and understanding 3D modeling tools and their application to modeling challenging artistic subjects.

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

Each of these is described in turn below to provide a basis for the discussion of related approaches in artistic 3D modeling that follows.

Dimensionality of input: One of the clearest distinctions between 3D modeling systems is the dimensionality of the input used to produce the model. Most 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 an indirect interface to one of the most inherently 3D tasks commonly performed with computers.

Freehand versus geometrically constrained: Geometrically constrained 3D modeling systems, such as CAD systems, are typically precise in the form they create. Their shortcomings in 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 appropriate only for representing organic subjects, but are almost always too difficult to control to produce refined depictions of these subjects. Models created with freehand tools often look loose or blobby and lack the specificity required for important subjects in medical illustration or artistic anatomy.

Output format: Even within the subarea of 3D modeling techniques appropriate for artists, there is considerable variation in the desired output format. In industrial design, the primary application area for SensAble's FreeForm tool [83], models are often sent to rapid prototyping machines for output as true physical shapes cast in plaster. When a physical realization is required, the model must conform to the

specifications of the 3D printer and the properties of the physical world, gravity for example. These constraints, or similar ones that may be imposed by animation or simulation software, may be built into the modeling tools. In general, the techniques presented in this dissertation are intended to create models to be viewed in virtual reality. This allows some flexibility in the visuals created: they are not required to adhere to physical laws, such as gravity, and are not required to be fully enclosed surfaces, thus allowing artists to explore alternative means of representation in virtual reality.

2.1 Techniques Utilizing Direct 3D Input

Early and recent work in techniques utilizing direct 3D input are described in this section. Then, in sections to follow, large-scale and desktop-scale 2D input techniques for producing 3D models are reviewed.

2.1.1 Early Work in 3D Computer Input

Sutherland introduced the first head-mounted, tracked, stereo display in 1968 [90]. 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 [16]. 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 of this approach for the design process. Freeing designers from needing to know particular numerical coordinate values describing a surface or orientations of coordinate system axes allows them to concentrate fully on the design task. These concepts have inspired this work and that of many others who have followed Clark.

Schmandt [79] was the first we know 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. Schmandt's work was one of the first experiments

on 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 detracted significantly from the interactive feel of the tool.

Galyean and Hughes [33] used direct 3D input with a monoscopic display to produce a voxel-based modeling technique. They 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 cube corners. This work was the first of the 3D sculpting systems to 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 from or added to a block of clay. A large body of work based on this metaphor has followed, and virtual clay remains one of the most successful approaches to obtaining output in the form of voxel-based descriptions or enclosed triangle meshes, which are useful for applications in rapid prototyping and animation. Advances in computer and haptic hardware continue to make possible a wide variety of extensions to this work that provide more control of the interface and more sophisticated forms.

At roughly the same time as the work by Galyean and Hughes, Sachs et al. presented the 3-Draw system [73], which used direct 3D input with both hands to create an intuitive CAD modeling system. This work was ground breaking in 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 both unconstrained and various types of semi-constrained sweeping input to specify complex 3D curves and provided 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 presenting interaction techniques for enhancing control over free-form input, one focus of this dissertation.

The 3DM system [13] 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 that commonly available in 2D drawing programs of the time, such as MacPaint. In addition to

performing geometrically constrained CAD-style modeling operations, 3DM incorporated a modeling mode based on sweeping 3D input for creating surfaces by extruding curves. One of the important contributions of 3DM with respect to this dissertation is its 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.

Deering's Holosketch [22] was the first system to combine a head-tracked stereo VR environment with a modeling system 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 Holosketch's drawing modes were based on continuous sweeping input, including a toothpaste mode reminiscent of Gaylean and Hughes' additive sculpting [33], a wire-frame-lines mode reminiscent of the 3D line drawings of 3-Draw [73], and a mode in which clouds of random small triangle particles were left behind the wand as it was swept through the air. Holosketch worked in a fishtank VR [21] setup using a 20-inch CRT and was also explored in alternative VR form factors.

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 0.25 cm and 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.

2.1.2 Recent Work in 3D Computer Input

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 and those based on haptic feedback, usually utilizing the clay sculpting metaphor first introduced by Gaylean and Hughes [33].

Schkolne’s Surface Drawing [78] is an example of free-form modeling utilizing a large-screen display device. Unlike our CavePainting work [54], (described in some detail in chapter 7), 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 [75, 76]. 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 bigger triangle-mesh model.

Several other large-scale, open-air input systems have been created for free-form modeling. The FreeDrawer [106] 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. 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 in which form can be generated by tracking fingertips [64]. 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 3. Mäkelä’s work also illustrates compelling visual effects that are possible when combining point- and surface-based representations. The BLUI system [11] 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. The IMAX SANDDE™ digital system, used in the films June [24], Falling in Love Again [25], and Moon Man [67], is also based on freehand 3D drawing with a tracked wand in a stereoscopic environment.

In the area of haptics-based free-form modeling, SensAble Technologies, 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 FreeFormTM [83]. The current incarnation 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 precision and refined aesthetic, two of the main goals of the work described in the remainder of this dissertation. However, this tool and other similar approaches [35], including those based on implicit surface representations [44], are useful only for creating solid models; that is, they cannot be used to explore directly the loose, stroke-based aesthetic seen in our work.

One haptics-based modeling system that takes a different approach from the rest is the springs-and-constraints 3D drawing system presented by Snibbe et al. [85]. The haptic interfaces presented in chapter 3 build upon this work. Snibbe et al.'s approach is different in that it uses dynamic haptic models to help artists explore various modes of creation, but these models are not based on interacting with a static geometry or properly simulating contact forces. Rather, they help to control and guide free-form input, but tend to allow artists to remain gestural in their interactions. The artistic potential of the new medium described by Snibbe et al. is left relatively unexplored, but because it is stroke- rather than geometry-based it fits nicely into the style of form creation explored in this dissertation.

Several other systems describe techniques called 3D painting, sometimes including haptic feedback as part of the interface [41]. 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 [40]. This concept has also been explored extensively artistically with the Surface Drawing tool [36].

Of final note is Schroering et al.'s work [81] 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 this input style could have important implications for free-form modeling in 3D.

2.2 Techniques Utilizing Large-Scale 2D Input

Some of the most influential work for the techniques presented in this dissertation is that of Buxton’s group, which has extensively explored using two-handed computer interaction techniques to draw controlled 2D curves [5, 15, 37]. This work was inspired by tape drawing, a technique commonly used in car design to create large scale drawings on a wall. In an extension to 3D drawing [38], 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 particularly useful in industrial applications, like car design, where it makes sense to think about curves in space in terms of their projections onto planes. For more organic subjects, however, a two-step input technique like this one can sometimes become tedious.

Another technique from the same research group uses a high-degree-of-freedom input device to create curves [39]. This approach, based on the “shape tape” technology, is promising artistically since it is such an immediate input method. However, the same two-step approach is required for creating a 3D curve. Our work investigates more direct alternatives for 3D construction.

2.3 Techniques Utilizing Desktop-Scale 2D Input

Many techniques for generating free-form 3D shapes from 2D input have been developed, most of them derived from or inspired by Zeleznik et al.’s seminal work Sketch [108]. Sketch made creating 3D form quick and intuitive and, with the addition of Skin [65], complex, organic shapes could be quickly modeled.

The Teddy [45] system builds on this work and takes it in a direction more appropriate for organic modeling subjects. In Teddy organic 3D forms are inferred from free-form 2D input. This work spurred a succession of investigations on the best algorithms for translating 2D free-form input into 3D form. Several systems have produced successful results in this area [51, 69].

Other work takes the approach of sketching multiple 2D views of an object [4, 9, 10, 52], 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 might complement the techniques presented in this dissertation.

2.4 Summary of Related Work

In summary, artists working with 3D modeling tools must currently make a tradeoff. Tools based on direct, 3D artistic input are typically very difficult to control with enough precision to address real subjects in science, medicine, and representational art. In contrast, full-featured modeling and animation packages, such as CAD systems or Maya [1], support extremely precise form but take years to master and lack intuitive interfaces for creating complex, natural forms. Techniques based on 2D input often succeed at leveraging artistic skill, but 2D input lacks directness for 3D modeling tasks. The work presented in the following chapters builds on previous approaches based on direct 3D input, and explores strategies that advance the precision and expressiveness of direct 3D computer input for modeling complex, natural subjects.

Chapter 3

Drawing on Air: Techniques for Controlled and Expressive 3D Line Illustration

Three-dimensional modeling approaches based on direct sweeping hand input, as in the direct 3D input tools described in the previous chapter, typically offer artists immediacy, intuitive interfaces, and exciting new artistic directions. The drawback with these tools is that artists cannot control them enough to address challenging subjects, such as those in scientific visualization [55] and even in representational art. While more traditional 3D modelers used in industry (typically driven by tablet, mouse, keyboard, and programming input) can achieve the precision needed to address these subjects, these systems are not accessible to an artist who has not trained with them, and they lack the physicality and directness that artists find so compelling in hand-based 3D interfaces. In this chapter, a toolset of alternative 3D, hand-based drawing interfaces called Drawing on Air is introduced. Drawing on Air maintains the advantages of direct 3D input but improves control and precision to the point that artists feel comfortable addressing challenging 3D subjects. Modeling based on a 3D input paradigm has already proven useful for initial concept design and for artistic gesture sketching. We hope these tools will lead to a new application area that goes beyond quick 3D sketches and moves toward illustration and more controlled drawing of difficult subjects.




Figure 3.1: A view of a 3D line illustration of a bat in flight created with Drawing on Air, a toolset of controllable 3D drawing computer interfaces. 3D input techniques inspired by tape drawing enable artists to create smooth, controlled 3D lines, as we see in the wing bones, with far more precision than is possible with freehand 3D drawing. The inset picture is a zoomed-in view of the wing from a different angle, showing artistic use of line weight to highlight joint locations.

In 2D, one of the most controlled approaches to drawing lines on a surface is tape drawing, a two-handed technique employed by car designers and recently adapted to digital media [5]. While such a deliberate approach to drawing is not always needed for 2D illustration, it is often used in car design because of the unusual constraints of the field. First, tape drawing is used for large scale drawings, typically full-scale, or near full-scale. Second, the curves in these drawings must be exceptionally smooth and controlled,¹ since measurements for blueprints are taken directly from the drawings. Tape drawing techniques overcome many of the difficulties of drawing controlled lines on such a large scale.

Just as drawing full-scale cars in 2D is considerably more challenging than smaller-scale drawing, drawing precisely in 3D, even at small scales, is much more difficult than traditional drawing. We hypothesize that elements of tape drawing, that are helpful in addressing challenges in 2D car design, may also be useful in computer-based, 3D drawing interactions. In this chapter, we introduce and evaluate two 3D drawing interactions inspired by tape drawing. Our first technique is a true 3D variant of tape drawing in which, just as in car design, both hands are used together to draw precisely. For the second technique, first proposed in 2D by Balakrishnan et al. [5], just one hand is used to draw. The one-handed approach proves to be easier to learn and simpler for drawing certain types of shapes in 3D, while the two-handed approach is very precise for expert users and adapts well to many styles of curves.

Both styles of drawing have their advantages and both belong in a complete 3D tool set. In fact, it can be useful to transition between the two even in the middle of drawing a curve. We show how to handle this situation and produce smooth, tangent-preserving, $G1$ -continuous transitions. Recovering gracefully from a mistake is particularly important, since 3D lines are harder to draw than their 2D counterparts. Users often want to back up to redraw portions of the line. Both our interfaces support this style of fluid, immediate, oversketching editing. Finally, Drawing on Air supports creating stylized 3D lines by allowing line parameters (orientation, thickness, and color) to be adjusted while drawing. These parameters serve as a 3D counterpart to line weight in traditional drawing.

One of our scientifically motivated illustration results is shown in Figure 3.1 and

¹Experienced car designers are said to be able to step back from their drawings and evaluate whether individual curves are $C2$ continuous simply through visual inspection.



Figure 3.2: Drawing on Air uses a stereoscopic desktop display. A Phantom haptic device and 6-DOF trackers are used for two-handed input.

the VR drawing environment used to create it is shown in Figure 3.2. To create this model, the artist needs a great deal of control over line shape, line weight (thickness and color variation), and 3D proportion. Drawing on Air enables artists to create 3D drawings like these. Note that the smooth shape of the bones of this bat would be nearly impossible to draw using freehand 3D input.

In the next section, we contrast our techniques with related approaches in bimanual drawing, freehand 3D modeling, and haptic-aided modeling, and then we describe our methods in detail. We present results in artistic anatomy and medical illustration and finally some lessons learned, future directions, and conclusions.

3.1 Related Work

This work builds on several areas of related research in 3D modeling and controllable 3D input that are described in detail in the previous chapter. Here we revisit several of the most relevant connections to related work and highlight differences between these approaches and the work described in this chapter.

3.1.1 Bimanual Approaches to Drawing Lines

Our bimanual approach to drawing lines builds on tape drawing, which was first introduced in digital form by Balakrishnan et al. [5] and later extended to a 3D application [38] that required drawing two 2D curves to construct a single 3D curve. High-degree-of-freedom input devices have also been used to create 3D curves using a similar two-step approach [39]. This approach is practical, and potentially preferable, in some applications in industrial design where parts fit together and curves can be constructed on the basis of constraints imposed by related curves. However, a more direct 3D approach to constructing curves is desirable for depicting organic subjects in an illustration style. Our technique introduces a form of tape drawing based on true 3D input coupled with haptic constraints.

3.1.2 Freehand 3D Drawing Systems

As noted in the previous chapter, there have been many approaches to using direct 3D input for geometric modeling [22, 54, 64, 73, 78]. In all of these completely freehand approaches, refinement of line or surface is difficult to achieve. HoloSketch uses arm and sometimes wrist rests, which are impractical for our approach because the arm and wrist need to move freely to specify orientation as well as position. A 10x input reduction mode can also be used in HoloSketch to change the mapping from input to output space. This reduces the apparent effects of muscular error, but also reduces the size of the curve that can be drawn in a single gesture. Surface Drawing uses a multiple-pass approach in which smoothing and magnet tools can be brushed over the form to edit and refine the resulting triangle mesh. Even with multi-step approaches like this, the form that typically results from 3D freehand modeling systems is characteristically loose, gestural, and sketchy. These are fine qualities for

artistic work — in fact, they offer a hand-crafted aesthetic that is rare and exciting in computer graphics — but they are inappropriate for problems in more refined illustration.

Alternative approaches, such as FreeDrawer [106] and Fiorentino et al.’s stroke segmentation and filtering [28], filter freehand input into smooth spline approximations. Tape-drawing-based approaches like ours act as user-driven filters. We avoid the difficult problem of separating noise from artistic intent, and the resulting errors that often frustrate artists, by having the artist drive the filtering process explicitly. Some additional filtering may help but does not seem necessary. We implemented an anisotropic filter in the style of Fiorentino et al.’s initial processing step, but found it of little utility in our situation because the haptic friction and viscosity forces seem to reduce muscular noise and help users hold their hand still at roughly the same level as the filter.

3.1.3 Haptic-Aided Drawing and Modeling

Our use of haptics is closely related to the springs and constraints for 3D drawing in Snibbe et al. [85] in that both approaches use haptic forces to create drawing guides rather than simulate realistic surface contact forces. While Snibbe et al. focus on exploratory doodling, our focus is on controlled drawing.

The DAB system [6] contains a sophisticated 3D haptic model of a brush that, like traditional painting and drawing, inherently supports adjusting line quality by twisting and pushing the brush against the canvas. Our work achieves similar continuous variation in line weight, but with a 3D “canvas” and a simplified 3D brush model.

3.2 Drawing on Air

Drawing on Air integrates two complementary approaches to drawing 3D curves, one-handed drag drawing and two-handed tape drawing. Both techniques have advantages. One-handed is generally easier to learn to control than two-handed, while two-handed feels more controllable to expert users. One-handed is also more appropriate for circular shapes that would require one’s arms to cross if drawn with the

two-handed approach.

The key to both techniques is giving the user explicit control of the tangent of the curve being drawn. This direction-explicit approach to drawing can be described in terms of two subtasks: 1. Defining the direction (tangent) of the drawing, and 2. Advancing the line along this direction. In the one-handed drawing mode, both these operations are performed with one hand. The artist drags around the brush like a water skier being towed behind a boat, and the drawing is constrained to move along the “tow rope” that describes the tangent of the curve. In the two-handed case, control of the two subtasks is separated; the drawing direction is set by moving the nondominant hand and the line is advanced by moving the dominant hand.

This direction-explicit approach to drawing helps with control at both a low, motor-control level and a higher cognitive level. At a low level, the techniques function as user-guided filters, greatly reducing tracker and muscular noise while biasing results toward important styles of curves. When an unsteady hand causes some jitter in the 3D input, the effects are minimized by the lever arm formed by the tangent. 3D positional error at the end of the lever shows up as a much smaller angular error at the point of the brush.

At a higher, cognitive level, three factors aid control. 1. The visual guidelines displaying the direction of drawing help to measure space and plan drawing. 2. Backup and redraw features help artists “explore” a curve, redrawing sections of it as they go. 3. Artists can work deliberately without introducing additional jitter, advancing the line only when they see it is going in the right direction.

In our implementation, drawing takes place at a fishtank (desktop-based) VR set up as shown in Figure 3.2, with two Polhemus magnetic trackers, one tracking the artist’s head and one tracking his nondominant hand. The tracked device worn on the nondominant hand also has a button on it which is used primarily to clutch and reframe the virtual artwork, as is done frequently while working to examine the model and position it appropriately for the next curve to be drawn.

The stylus of a SensAble Phantom force-feedback device is held in the dominant hand, and small friction and viscous force effects are applied to the stylus throughout the interaction to give the user some slight resistance as the pen is moved through the air. In this form factor, an offset exists between the physical working space of the

hands and their virtual representations on the screen. Alternative hardware designs that allow collocation and maintain a wide range of motion for both hands might be possible. We expect such a design would further enhance control.

In the sections below, we describe the details and implementation of the two drawing modes. Then, in the third section, we describe how to transition between the two modes while drawing.

3.2.1 One-Handed Drag Drawing

In the one-handed drag drawing technique, a virtual brush from which the curve is drawn is towed behind the physical stylus manipulated by the user. The “tow rope” used can be thought of as a rope of length l , in that when the stylus is a distance l away from the brush, the rope pulls tight and the brush is dragged directly toward the stylus. When the stylus moves toward the brush the rope goes slack, and the stylus is free to move anywhere within a radius of l of the brush without doing any towing. The position of the brush at each new frame $b(t)$ can be updated from the latest reading returned from the Phantom for the stylus position $s(t)$ as follows. Let \vec{d} be the current drawing direction

$$\vec{d} = s(t) - b(t-1) \tag{3.1}$$

then, when the brush is in a drawing state, $b(t)$ is computed as

$$b(t) = \begin{cases} b(t-1) & \text{if } |\vec{d}| < l \\ s(t) - l\vec{d} & \text{if } |\vec{d}| \geq l \end{cases} \tag{3.2}$$

In two cases, this metaphor becomes slightly more complex. First, when the artist first begins to draw, it is annoying to start with the tow rope slack because quite a bit of movement is required just to begin to draw. To address this, we start with a very small tow rope and gradually lengthen it while the curve is being drawn. The second case arises when we introduce the ability to back up and redraw portions of the curve. We discuss both of these in more detail in the next sections.




Figure 3.3: The progression of a Drawing on Air one-handed, drag-mode interaction. When drawing first starts (from position a to b), the drag line grows to its maximum length l . From position c to d, the user has backed up slightly and then made a sharp change in direction before continuing to draw until position e. He then backs up to within a distance $l/2$ of the end of the drawn curve and begins to erase a portion of the curve (position f to g). The haptic constraint imposed during the erasing motion guides the user toward a tangent-preserving transition when he begins to draw again (position h).

Dynamic Tow-Rope Lengthening

The length l of the tow rope changes dynamically so that drawing starts almost immediately when the brush button is pressed, and there is no tow rope ($l = 0$) when the brush is turned off. As drawing begins, the user moves the stylus a minimum distance l_{min} (0.5 cm in our implementation) away from the brush before any drawing occurs. This distance should be just far enough that the user can establish an initial drawing direction, but not so far that he becomes frustrated because he is trying to draw but instead is only lengthening the tow rope. Then, the tow rope gradually grows to its maximum length l_{max} (4.5 cm) as the curve is drawn according to the following relationship, where a is the arc length of the curve drawn so far:

$$l = \begin{cases} \max(l_{min}, a) & \text{if } a < l_{max} \\ l_{max} & \text{if } a \geq l_{max} \end{cases} \quad (3.3)$$

The lengthening of the tow rope is represented in the first two illustrations in Figure 3.3. The blue pen is the stylus s that the artist holds. At position a, the virtual brush, $b(t)$ in Equation (3.2), is at the end of the black mark where it meets the green tow rope. Here, the tow rope represented by the green line is growing longer as the curve is just starting to be drawn. By the time the stylus reaches position b in Figure 3.3, the tow rope is at its maximum length, where it will stay as the rest of the curve is drawn.

Haptic-Aided Curve Redrawing

To signal the beginning of a backup and redraw operation, the stylus is moved backward to be within a distance $l/2$ of the brush. A better metaphor for the tow rope here is a stiff rod of length $l/2$ (the red portion of the tow rope in Figure 3.3) attached at one end to the brush and at the other end to a rope of length $l/2$ (the green portion of the tow rope in Figure 3.3) that is tied to the stylus. Notice the two circles in the third illustration of Figure 3.3. The outer circle is at a distance l from the brush and marks the region outside of which any stylus movement will drag the brush along and add to the drawn curve. The inner circle marks the backup region with radius $l/2$. Stylus movement between these two regions does not cause the brush to move, and this allows artists to create sharp discontinuities in the curve, as seen in Figure 3.3

between positions c and d. When the stylus is moved to the edge of the inner circle, the redrawing mode is engaged.

While backing up to erase the curve, haptic forces steer the stylus to a position from which forward drawing results in a smooth transition in the tangent of the curve. The rod in our metaphor needs to be swept around so that it always points in the direction of the tangent of the last sample of the curve. This is achieved through a haptic polyline constraint, which we render to the Phantom with SensAble’s OpenHaptics toolkit. The polyline sweeps out the arc formed by each sample of the curve offset by $l/2$ times the tangent vector at that sample, similar to the dashed blue line between positions f and g in Figure 3.3. Thus, if the drawn curve C is defined by a set of samples $c_0..c_n$, the haptic constraint polyline P is defined by $p_0..p_{n+1}$ such that

$$p_i = \begin{cases} c_i + \frac{l}{2} \vec{d}_i & \text{for } i = 0, \dots, n \\ c_n + l \vec{d}_n & \text{for } i = n + 1 \end{cases} \quad (3.4)$$

where \vec{d}_i is the direction of drawing (tangent) at sample i along C . The additional line segment added for the $(n+1)$ th sample lets the user move easily out of the backup region and begin forward drawing while preserving tangent consistency.

We back up and erase the portion of the curve that we pass up to the sample c_{backup} . Thus, we can rewrite Equation (3.2) to a more complete form that includes the case for curve redrawing:

$$b(t) = \begin{cases} c_{backup} & \text{if } |\vec{d}| \leq \frac{l}{2} \\ b(t-1) & \text{if } \frac{l}{2} < |\vec{d}| < l \\ s(t) - l \vec{d} & \text{if } |\vec{d}| \geq l. \end{cases} \quad (3.5)$$

In informal testing of several users with a non-haptic version of this technique, some users had trouble engaging the mode initially and then frequently moved out of the backup region and accidentally started drawing backwards. In addition to making possible a tangent preserving transition, the haptics serve to keep the user’s pen on track to execute this operation quickly, avoiding the miscues we found without the haptics.

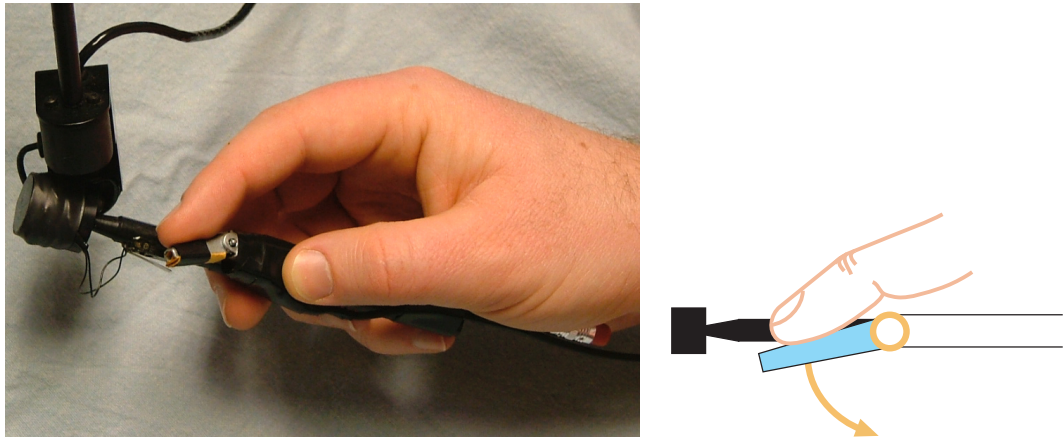


Figure 3.4: A custom elastic controller is mounted on a second pen attached to the side of the Phantom stylus. For 3D drawing, comfort and range of wrist motion are improved by holding the stylus with the fingertips as an artist holds a piece of charcoal. In this position, the index finger is properly positioned to apply pressure to deflect the spring-loaded hinge as shown in the diagram on the right.

Varying Line Weight

To begin drawing and then adjust the line weight of the curve, the user pushes with his fingertip on a custom elastic controller, shown in Figure 3.4, made from a spring-loaded hinge fastened to the Phantom stylus. As more force is applied and the hinge deflects, the width of the mark is expanded and the color is adjusted to create a heavier 3D line. Releasing the spring device makes a thinner line.

Colors are interpolated from a gradient selected by the user. Artists often import their own color palettes and adjust the gradients to increase contrast with the background color as pressure increases.

As with traditional artistic tools and other pressure-based interfaces [70], visual feedback is important for control. The width of the curve geometry and a pressure meter drawn to the right of the brush model provide continuous visual indications of the current line weight.

3.2.2 Two-Handed Tape Drawing

Drawing a curve with the two-handed tape-drawing interface requires coordinated movement of both hands, as depicted in Figure 3.5. Throughout the interaction, the

tape-drawing tangent or drawing direction \vec{d} is updated on the basis of the last sample of the curve, c_n , and the latest tracker reading for the hand, $h(t)$:

$$\vec{d} = \begin{cases} h(t) - s(t) & \text{if } n = 0 \\ h(t) - c_n & \text{if } n > 0. \end{cases} \quad (3.6)$$

In the initial case, the stylus location is used instead of the last curve sample.

The brush is advanced along the drawing direction by movement of the stylus:

$$b(t) = \text{projection of } s(t) \text{ onto the line segment } (h(t), b(t-1)). \quad (3.7)$$

Straight lines can easily be drawn by holding the nondominant hand in place and moving the stylus directly along the tangent line. To draw a curve, the nondominant hand is moved while drawing to change the tangent dynamically as the dominant hand advances along the tangent, as we see in Figure 3.5 from position a to c. The artist can stop his dominant hand at any point and make a drastic change in the curve tangent before proceeding to create jagged or bumpy lines.

Force feedback in the form of a dynamic line constraint is used to constrain the stylus tip to remain on the line segment connecting the two hands. This helps the user concentrate on specifying the drawing direction and advancing deliberately along this tangent rather than thinking too much about the 3D position of the dominant hand.

It is unclear whether users have a consistent preference for the role of each hand in tape drawing. Traditionally, 2D tape drawers draw from left to right regardless of handedness. In this 3D interaction, artists seem to be most comfortable drawing toward the nondominant hand, so that the dominant hand can play the key role in adjusting line weight via haptic interaction, as described below.




Figure 3.5: The progression of a Drawing on Air tape-mode interaction for a left-handed user. The drawing direction is determined by the position of the hand and the endpoint of the curve being drawn. To draw a curved path, both hands must move together (position a through c). As the user backs up to redraw a portion of the curve (d), a virtual offset (shown as a magenta vector) is applied to the hand position so that a tangent-preserving transition is made when forward drawing resumes (e).

Varying Line Weight

The haptic line constraint provides a control for varying line weight that mimics physical media. Just as a brush or a piece of charcoal is pushed against the paper to make a dark, thick line, users push against this line constraint to change the weight of the mark. The pressure from this interaction, p_{tape} , is combined with the pressure from the elastic finger controller, p_{finger} , to produce a total value for the line weight of the mark:

$$w = p_{finger}/p_{maxfinger} + p_{tape}/p_{maxtape}. \quad (3.8)$$

This value is used to adjust the color and width of the mark being drawn.

As the user pushes against the haptic constraint, the stylus physically moves slightly away from the line constraint. In fact, the distance that it moves off the line serves as our measure for p_{tape} , but its virtual position is constrained in software to remain precisely on the tangent line so that a smooth curve is drawn.

Haptic-Aided Curve Redrawing

As with drag drawing, we extend the basic tape drawing interaction to support backing up and redrawing the mark. Balakrishnan et al.’s 2D tape drawing [5] included a similar procedure for lifting up tape. We extend this to 3D and remove the need for a button press to enter redrawing mode.

This feature requires the use of four haptic states: Brush Off, Drawing Forward, Backing Up, and Hands Too Close. When the brush is off, no haptic forces (other than the constant friction and viscosity) are rendered, allowing both hands to move freely. The Drawing Forward state is also straightforward and entails rendering the normal line segment constraint going from the last sample of the curve to the hand position, as described above.

To enter the Backing Up state, the user pushes backwards against a small force until he “pops” into the new state. Initially, this feels as though the stylus is trapped by the bounds of the normal forward-drawing constraint so that it cannot move back. Once a sufficient force is applied, this constraint is lifted. (See Komerska and Ware for discussion of similar haptic pop-through effects [60].) Now the stylus slides effortlessly along a haptic polyline constraint defined by all the previous samples of the curve and connected in the forward direction to the position of the hand. By following the

haptic guide, the user can slide easily backward to erase a portion of the curve, or start moving in a forward direction to resume drawing.

As the curve is erased by moving backward, a virtual offset is applied to the location of the nondominant hand in order to setup a tangent preserving transition when forward drawing resumes. Notice the positioning of the hand at point c in Figure 3.5. At position d, the brush has backed up but the hand remains in the same place. An offset (the magenta line) is applied to the virtual position of the hand, so that when the brush starts moving forward again (from position d to e) a smooth, tangent-preserving transition is made between the old part of the curve and the newly drawn portion.

The final state, Hands Too Close, is entered if the hand and the brush positions are closer than 2.25 cm to each other. At such close distances, the tracker readings can cross quickly, drastically changing the direction of the curve tangent and causing the haptic simulation to become unstable. If the hands reach this state, we render a haptic point constraint to hold the brush at its current position and indicate visually to the user to move his hands apart.

Visual Feedback

As in traditional tape drawing, sighting and measuring space with the tangent guideline both in preparation for drawing and as an interactive preview while drawing is extremely important. As seen in Figure 3.6, feedback is rendered with an orange line connecting the center of the brush model to the position of the nondominant hand; the orange line is surrounded by a black rectangle indicating the surface upon which a ribbon form will be drawn and the maximum width of that ribbon. A pressure meter drawn with yellow and red bars to the side of the brush indicates the current line weight and a yellow cross-bar at the tip of the brush also changes length in response to pressure.

3.2.3 Integrated One- and Two-Handed Drawing

Drawing on Air begins by default in the drag-drawing mode. The user transitions to tape-drawing mode by pressing and holding the button in the nondominant hand; to return to drag mode, the button is released. Through each of these transitions,




Figure 3.6: Visual feedback while drawing a blue ribbon form with tape mode. The sphere on the left is the location of the nondominant hand. A yellow cylinder facing out of the screen marks the location and orientation of the brush.

virtual offsets are applied to the position of the hand and brush in order to maintain a smooth transition in the drawing direction. The calculation for line weight is also adjusted to maintain a constant value across the transition.

To begin a line with tape rather than drag mode, the user holds down the button on the nondominant hand before starting to draw. Recall that when not drawing, this button is usually used to clutch and reframe the artwork. To disambiguate these two operations, we make a logical distinction based on the positioning of the two hands when the button is pressed: if the hands are close together, the button press activates tape mode; if they are far apart, it activates the reframing operation.

Drag-to-Tape Transition

Upon the transition from drag to tape mode, the mapping from the stylus to the virtual brush must be adjusted. In drag drawing, the stylus tows the brush behind it, but in tape mode the stylus and the brush are collocated. To make the transformation, an offset from the raw stylus input values to a virtual location is maintained. The offset is set to zero at the beginning of each line, and for each drag-to-tape transition, $\overrightarrow{(b - s)}$ is added to the offset.

This alone does not guarantee a smooth transition in the tangent of the curve,

since the tape-mode drawing direction is also determined by the hand location. Thus, a second offset is applied to the hand, also initialized to zero. When the transition occurs, the “goal” hand position is the closest point to the hand along the line defined by the last sample on the curve and the tangent previously defined in drag mode. The hand offset adjusts the raw hand input so that it matches the goal hand position.

Tape-to-Drag Transition

To transition from tape drawing to drag drawing, the stylus position needs to jump forward along the drawing direction so that it is again pulling the brush through space. A new ideal position for the stylus is

$$s_{new} = c_n + l\vec{d} \quad (3.9)$$

where c_n is the last sample on the curve, l is the length of the drag rope, and \vec{d} is the drawing direction established by tape drawing. The stylus offset described above is adjusted to make the current raw input match the value of s_{new} . The hand offset is reset to zero on this transition to avoid accumulating a large offset if multiple transitions are made while drawing the same line. Accumulating a large offset is not a problem for the stylus, since the offsets applied there are always small.

In tape mode, two pressure terms contribute to the line weight calculation, pressure from the elastic finger controller and pressure from pushing against the haptic line constraint. On the transition from tape to drag mode, the line constraint term (p_{tape} in Equation (3.8)) goes to zero. Thus, the mapping from finger pressure to line weight needs to be adjusted so that: 1. The total line weight stays constant through the transition, and 2. The line weight returns smoothly to zero as the finger controller is released. To accomplish this, the gain of the device is adjusted by changing the $p_{maxfinger}$ term from Equation (3.8):

$$p_{maxfinger} = p_{finger}/w. \quad (3.10)$$

$p_{maxfinger}$ is reinitialized to 1.0 at the beginning of each line.

Reverse Tape Drawing

During the transition from drag drawing to tape drawing, the drawing direction established by the drag technique may point away from rather than toward the hand.

This creates a problem in maintaining a consistent drawing direction before and after the transition to tape mode. When we encounter this situation, we switch to a technique we call reverse tape drawing where all calculations based on \vec{d} are performed with $-\vec{d}$. Rather than drawing toward the hand, users draw directly away from it. In practice, this technique is far harder to control than normal tape drawing, but it is useful for drawing small sections of a curve in this situation.

3.2.4 Brush Model for 3D Geometric Pigment

A variety of geometries could be generated by the user’s input, which includes several continuously varying parameters: position, orientation, and pressure along a controlled 3D path through space. We have found two simple geometric forms, ribbons and tubes, useful in creating a variety of artistic and scientific line illustrations. Ribbons are useful for depicting 3D form because they look like tiny patches of evenly lit surfaces that, when seen in stereo, are composed by the human visual system into a larger coherent 3D surface. Tubes are unlike ribbons in that they do not merely *suggest* a larger form, but instead, evoke the sensation of *being* the form. Thus, while a few appropriately placed ribbons can effectively suggest the skin moving over the cheekbone on a face, tubes are inappropriate for suggesting this form, but are better than ribbons for depicting thin tendons or muscles that can be specified completely with one stroke.

Ribbons require the user to specify an orientation as the curve is drawn. Care must be taken in designing the mapping from user input to the ribbon surface normal so as not to make the user move his wrist into uncomfortable positions while drawing in order to maintain a correct normal. For ribbons drawn roughly within a plane parallel to the film plane, using as the normal the component of the brush stylus handle that is perpendicular to the drawing direction works well:

$$\vec{n}_{default} = \vec{h} - \vec{h}(\vec{h} \cdot \vec{d}) \quad (3.11)$$

where \vec{h} points in the direction of the brush handle. However, for more difficult-to-draw curves that move in and out of the screen, \vec{h} and \vec{d} become roughly parallel and the normal becomes unstable or gradually spins as the curve progresses through a turn. The user can avoid this situation by carefully adjusting the handle of the brush

while drawing, but doing this can become uncomfortable and annoying.

A solution attempts to do what the user typically expects to happen in these unstable situations: to maintain something very close to the previous normal throughout the period of instability. The following pseudocode describes the algorithm:

$$\begin{aligned}
 &\text{if } |\vec{h} \cdot \vec{d}| < 0.7 \text{ then} && (3.12) \\
 &\quad \vec{n}_{new} = \vec{n}_{default} \\
 &\quad \vec{n}_{lastgood} = \vec{n}_{default} \\
 &\text{else if } |\vec{h} \cdot \vec{d}| < 0.8 \text{ then} \\
 &\quad a = (|\vec{h} \cdot \vec{d}| - 0.7)/0.1 \\
 &\quad \vec{n}_{new} = \text{linearInterpolate}(\vec{n}_{default}, \vec{n}_{lastgood}, a) \\
 &\text{else} \\
 &\quad \vec{n}_{new} = \vec{n}_{lastgood}.
 \end{aligned}$$

We default to returning the normal as the component of the brush’s handle not pointing in the drawing direction. If the handle and drawing direction are close to parallel, then we return the last good value for the normal, and in a small range of values between these two cases we linearly interpolate between the two potential values for the normal to achieve a smooth transition between the cases.

3.3 Illustration Results

More than twenty artists have used Drawing on Air for sessions ranging from an hour to regular use over the course of a year. Feedback from these artists has been critical in refining these tools. We have encouraged the the few artists we have been able to support in long-term, repeated use to work on scientific illustration problems, in particular, 3D illustrations of bat flight, since it is a real-world illustration problem that requires a 3D treatment. In this section we report these results in detail. Also touched upon are the more artistically motivated 3D drawings that I have created. These are described in additional detail in chapter 7, which also presents findings from artistic critiques of the work. Each of these models took between two to five hours to create. They are designed specifically to be viewed in stereo and, while they




Figure 3.7: An artist's initial attempt at a 3D illustration of a bat posed in flight.

lose a great deal of their impact and 3D character when printed on paper, they still demonstrate in many instances significant control of form and line quality.

3.3.1 Illustrations of Bat Flight

The illustrations presented here were made by two different artists as part of an on-going collaboration with an evolutionary biologist studying bat flight. Traditionally, almost all anatomical illustrations, and even preserved specimens of bats, have been posed with the wing membrane and skeleton completely flattened, as typical of a bird's wing or a fixed-wing aircraft. However, recent research has demonstrated that bat flight is far more complex than that of birds, in large part because the flexible wing membrane and bones undergo tremendous 3D deformations during flight [91].

Because 3D understanding is critical in this problem, 3D presentations of bats posed in flight are important tools for the biological researcher. Figures 3.1, 3.7, 3.8, and 3.9 show initial results working toward the goal of an animated 3D anatomical illustration of a bat, including bones, muscles, and tendons with clear insertion and attachment points. The illustration in Figure 3.1, our initial proof of concept, highlights several features of *Drawing on Air*. First, the smooth curves of the wing bones are clearly indicated. Since these bones actually bend during flight, their shape is important but would be impossible to convey accurately with a freehand approach to drawing. Note also how the artist has adjusted the line thickness in the bones (see inset detail in Figure 3.1) to indicate the joints.

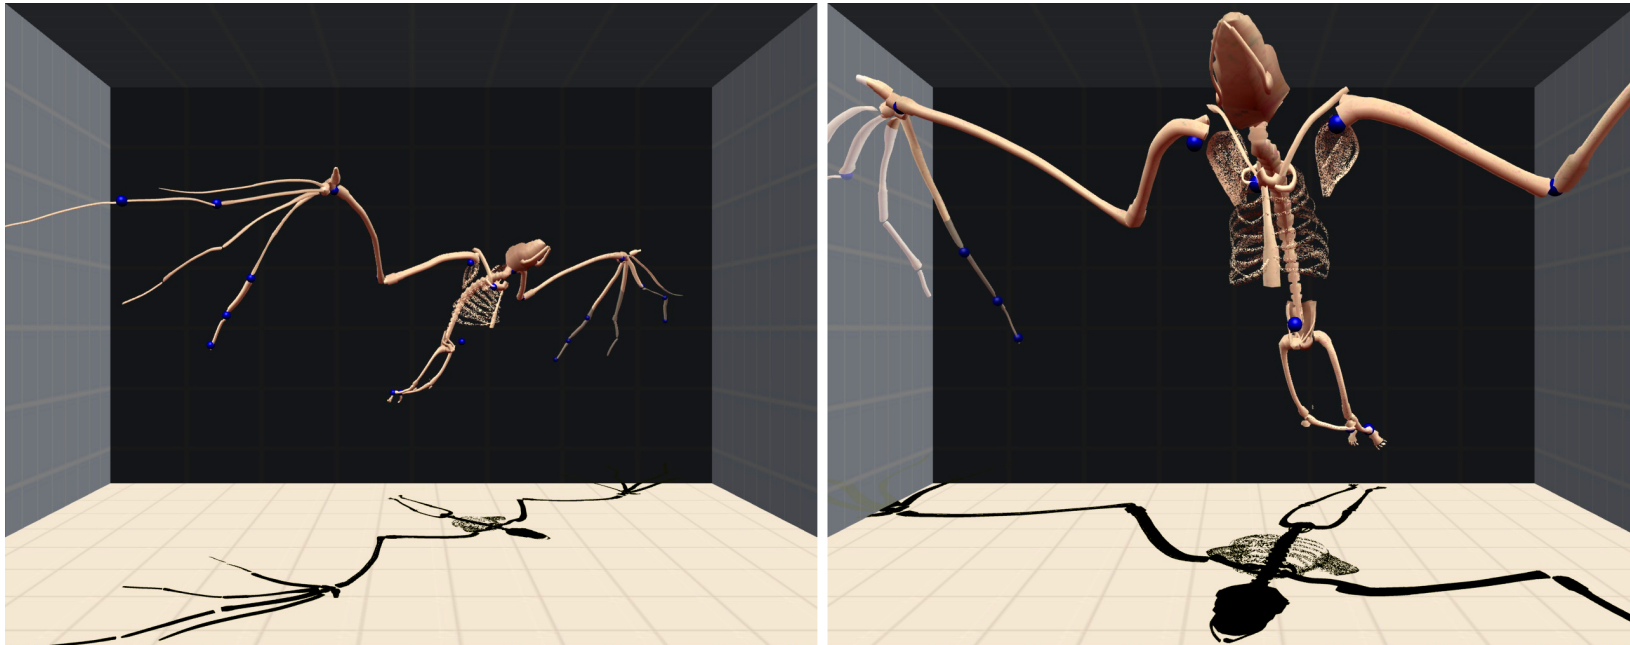


Figure 3.8: Two views from an illustration of a bat skeleton posed in flight. The blue spheres are markers from motion capture data of a bat flying in a wind tunnel.

Figures 3.7 through 3.9 were drawn by a second artist. Figure 3.7 is a view of the first bat 3D drawing created by this artist. The model shown in Figure 3.8 is a second, considerably refined attempt that was drawn on top of 3D markers imported into the system from data collected by flying bats in a wind tunnel. Twelve markers were placed on important joints in the wing and tracked by cameras. A frame of the resulting motion data was imported into the drawing tool and the markers were displayed as blue spheres. Then, the artist drew within the reference frame given by the markers to create an illustration that is faithful to the scientific data, yet stylized to clarify the role of the skeletal system in flight.

Ultimately, our scientist collaborator would like to see 3D anatomical illustrations complete with bones, muscles, and tendons. The model shown in Figure 3.9 is a result from a first attempt at including muscles in the 3D illustration. Our investigations are still too preliminary to say that illustrations like this have helped direct the underlying scientific study, but we are moving steadily in this direction. In the future, one concept that we believe will be of great utility in investigations like these is artistic augmentation of 3D displays of experimental data. The drawing on top of simple motion-capture markers in Figure 3.8 is an example. This approach leaves sufficient room for artistic visual clarification, but also grounds the visual display in real experimental data, making this both an exciting visual problem for artistic investigation and a valuable scientific result.

3.3.2 Artistic Anatomy

A brief treatment of some results in artistic anatomy is also warranted here. Additional discussion of these results, which I created and then critiqued in virtual reality with the help of RISD illustration professor, Fritz Drury, is presented in chapter 7. Three results from this investigation are pictured in Figures 3.10, 3.11, and 3.12. Looking at these figures shows that interesting stylistic variation is possible with Drawing on Air. While the first bat illustration (Figure 3.1) and the bearded man are quite sculpted, the Swahili bride is created with minimal use of line. One theme emerging from critiques of this work was the effectiveness of this minimal style. When seen in head-tracked stereo, a line drawing like this contains enough depth cues to exert a tremendous 3D presence. The artistic effect is as compelling, if not more so,

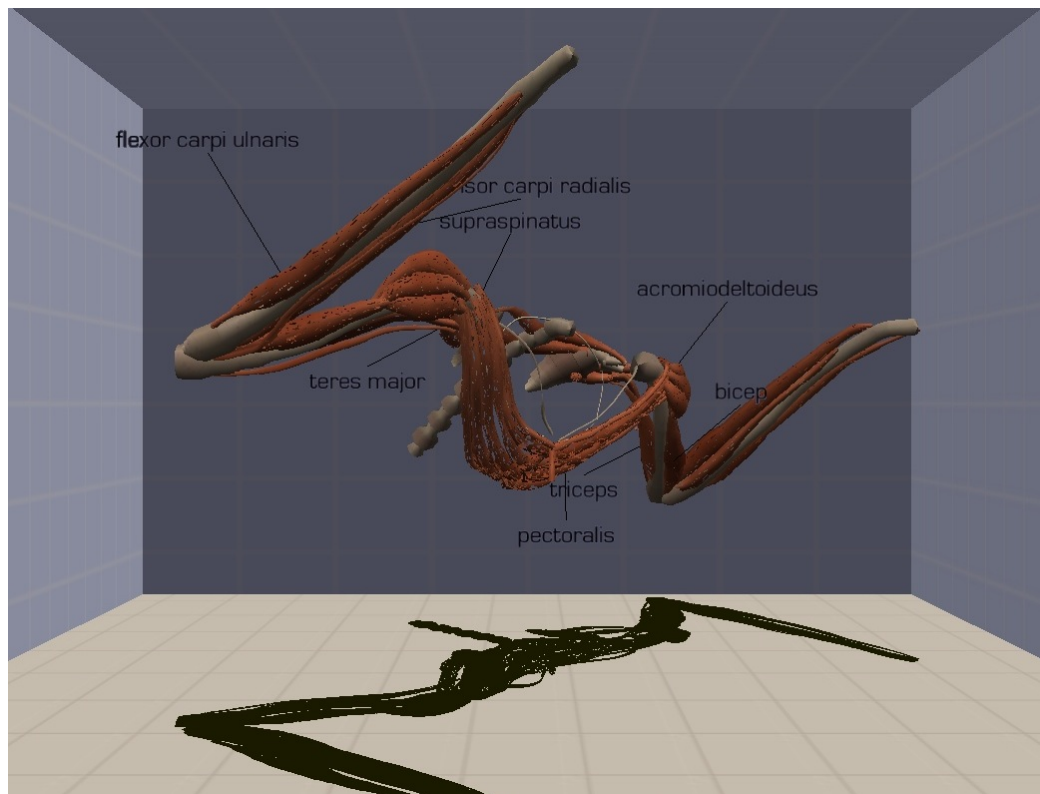


Figure 3.9: 3D illustration of bones and muscles in a bat's torso that are of importance in controlling flight.



Figure 3.10: Bearded man.



Figure 3.11: A Swahili bride wearing a green veil.

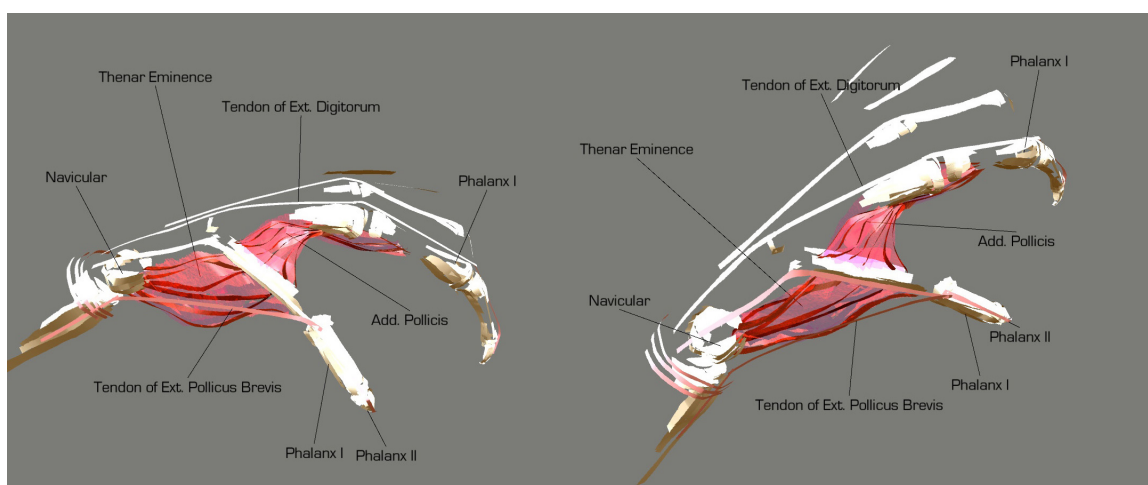


Figure 3.12: Muscles and tendons in the hand.

as that of a more traditional, full-surface representation of the face.

The use of ribbons as the drawing primitive is important in making this style work because they suggest a small portion of a larger surface. Figure 3.12 is an experiment in using this minimal style for medical illustration. The end points of the bones are drawn out in detail, but the anatomically less interesting flat regions in the middle of the bones are merely suggested. In many ways, this focus on detail in important regions mimics the way an illustrator would work with 2D physical media. Notice the control in the lines of the tendons running over the knuckles in this example, drawn with the tape mode.

3.4 Discussion

Drawing in 3D and drawing with both hands are new ways for most artists to work. In this section, we discuss what we learned about effective strategies for drawing with correct proportion and picking the right lines to draw to make compelling 3D illustrations. We also discuss the design for our elastic finger controller and some nuances of working with tape drawing in 3D.

3.4.1 Strategies for Effective Use

Frequent reframe operations are an important part of the artistic process with Drawing on Air. Repositioning and rotating the form increases 3D understanding of the shape. It is also important in positioning the model appropriately to draw the next mark. This has special importance for tape mode, where users have a clear preference for orienting the artwork so that lines can be drawn toward the nondominant hand. Reframing and scaling is also necessary to deal with the limited range of the Phantom device.

Artists also find it useful to create guidelines or scaffolding for refining 3D proportion before drawing a final version. Our application supports drawing layers that can be turned on and off. Often at least one layer is used for rough guidelines and working out 3D proportion.

Picking the right lines to draw is far more difficult in 3D line illustration than in 2D because multiple viewpoints must be considered. Silhouettes are often used in 2D

to bound and define a form, but they break down in 3D when the form is intended to be viewed from multiple directions. Rather than using multiple 2D silhouettes, 3D illustrations are much more compelling when they are composed of lines that cross many planes of the form, often following along some important feature. In figure drawing, for example, the serratus and oblique muscles of the side of the torso are a good choice for this type of characteristic curve because they naturally spiral around the form from almost every viewpoint. When the edge of one of these muscles is traced out with a ribbon, the orientation of the ribbon at each point in space helps clarify the 3D shape to the viewer and lead the eye around the form. This style of depiction mirrors the more scientifically motivated styles developed by Interrante [46]. In both, pattern and 3D “line” are shown to be effective visual elements in describing complex volumetric virtual forms.

One feature we plan to explore in the future is view-dependent rendering of these 3D line illustrations. The minimalist style of 3D drawing, as opposed to more sculptural approaches, lends itself to creating models with clear transparent regions. In some situations, such as the faces in Figures 3.10 and 3.11, looking from the perspective of the rear of the model is distracting because we see the features of the face inside-out, ruining the impression of the back of the head. Thus, for some subjects, faces in particular, view-dependent display of marks, including the ability to hide marks from certain views, might help create clearer illustrations.

3.4.2 Controls for Line Weight

Pushing against the haptic line constraint in the tape-drawing mode to adjust line weight mimics the approach used in traditional media of pushing the drawing implement into the paper to thicken and darken a line. The difficulty with moving this approach directly into 3D is knowing where to simulate the surface of the paper, since in the general case, it is impossible to predict where the artist will want to draw next. Tape drawing avoids this prediction problem by separating the controls for setting the drawing direction and advancing along it.

We did explore a one-handed technique that also separates these controls. The position of the brush is used to advance the line, while its orientation sets the drawing direction. This allows us to push against the linear drawing direction constraint as

we do with tape drawing to adjust line weight. The drawback of this approach is that it requires extreme and unnatural bending of the wrist in order to move in and out of a plane and create complex marks.

The elastic hinge device for line-weight adjustment is the result of several design iterations beginning with the mouse-based control in HoloSketch [22]. We also explored a touchpad-based variant and an isometric pressure sensor used in either hand. In informal testing, the isometric controller was preferred to the isotonic mouse and touchpad, and the final elastic design was preferred to the isometric one. Zhai notes greater ease of learning with elastic versus isometric controllers in 6-DOF manipulations [110]. We note a similar preference for the elastic over isometric controller for untrained users, but have yet to explore how this may vary with additional experience.

Both methods for controlling line weight (the elastic hinge and pushing against the tape drawing haptic constraint) have their advantages. The finger-based method is easier to hold at a constant value while drawing intricate curves, while the haptic-based method seems easier to control for simple curves, especially those lying roughly within a plane.

3.4.3 3D Tape Drawing Without Haptics

We tried a non-haptic digital implementation of 3D tape drawing, but without the haptic constraint, the trailing hand can stray from the curve. This makes our tape-mode control for line weight impossible to realize, and even with that feature disabled, users are still frustrated by their lack of control. As in 2D digital implementations [5], we advance the drawing along the tangent line by projecting the position of the trailing hand onto the tangent. Drifting of the trailing hand slightly off the curve is not significant enough in 2D to pose a problem, but in 3D, keeping the trailing hand close to the curve is much more difficult. When it drifts too far, its projection onto the tangent drawing guide can be in an unexpected place. As the 3D nature of the curves becomes more complex, this drifting increases until an unexpected projection makes the drawing seem to jump forward to the user.

3.4.4 Nondominant-Hand Offset Mode

One of the limitations of tape drawing is the difficulty of drawing complete circles and other shapes that require the hands to cross. We explored a mode in which the position of the nondominant hand is offset horizontally by six inches toward the dominant hand for all calculations. In this mode, the hands can easily cross virtually without crossing physically, allowing the user to draw full circles with tape drawing. While in principle this solves the circle limitation, in practice so much concentration is required of the user to work with such a large hand offset that the technique is infeasible.

3.5 Conclusion

Drawing on Air enables artists to work with direct, hand-based 3D input to create controlled 3D models in an illustration style. It provides simultaneous control of position, orientation, and line weight of a 3D mark through two modes of interaction, each appropriate for important classes of 3D curves. Mechanisms for transitioning from one-handed to two-handed drawing preserve the fluidity of the interaction and the smooth quality of the curves. Haptic-aided curve redrawing, which preserves smoothness, enables artists to explore subtle line variations with precision. Drawing on Air, like other VR tools, leverages the benefits of working directly in space, but also provides the rich, controllable interaction necessary for refined 3D illustration.

Our illustration results demonstrate that artists can effectively address challenging visual subjects in both visual art and science using Drawing on Air. We attribute this to the increased control afforded by Drawing on Air, coupled with the ability to adjust line weight. Drawing on Air is an important first step toward making refined 3D illustration as accessible as drawing on paper. We believe advances toward this goal will make possible both important artistic results and more effective scientific and medical illustrations.

Chapter 4

The Drawing Control Experiment

In this chapter we present the drawing control experiment, a formal evaluation of artistic control in 3D drawing techniques. The drag-mode and tape-mode drawing techniques that comprise Drawing on Air are compared to a baseline of two freehand drawing techniques to understand differences between the two Drawing on Air modes and establish the benefits of working with Drawing on Air over standard approaches. We asked users who know how to draw with physical materials to participate in this study. Indeed, a subsidiary goal of this study was simply to see whether these users, generally inexperienced with computers, would be able to learn to use Drawing on Air in what was typically their first exposure to VR.

4.1 Conditions and Hypotheses

The study tested four conditions corresponding to the four input techniques for drawing 3D lines. The first, “drag”, is the drag mode technique of Drawing on Air. The second, “tape”, is the 3D tape drawing mode of Drawing on Air. For the purpose of the study, transitioning from one mode to the other was disabled along with backing up to redraw a line and adjusting line weight. The third condition, called “sand”, is a freehand drawing technique. There are no constraints on movement of the stylus or the resulting line, but the friction and viscosity forces that are part of Drawing on Air are applied to the Phantom. Users describe the effect as being like moving the brush through a bucket of loose sand. The final technique, “free”, is also freehand,

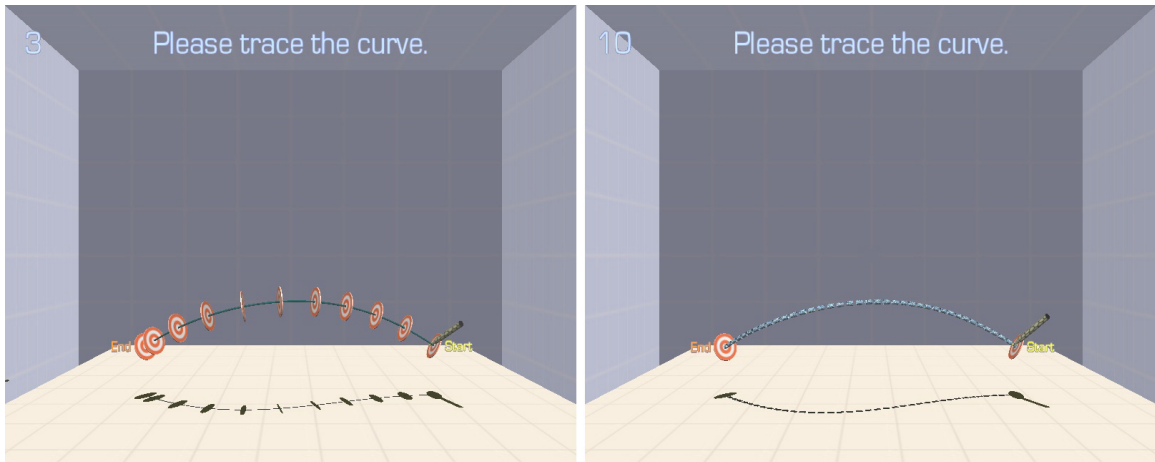


Figure 4.1: Participants’ view of the tracing task during training (left) and post-training (right).

but without any haptic forces. All techniques used the Phantom device for input.

Our hypotheses entering the experiment were: 1. That drag and tape would considerably outperform sand and free, and 2. That sand would outperform free by a significant margin, but less than the difference between the Drawing on Air techniques and the freehand ones.

4.2 Methodology

Users performed repeated tracing tasks under each of the four conditions. Each participant used each of the four 3D drawing techniques, thus the study was a within-subjects design. A Latin square was used to randomize the ordering of the drawing techniques across participants. Measures of positional and directional accuracy and drawing time were computed for each tracing trial.

Tracing was performed directly on top of a 3D curve displayed in VR. In each trial, the participant was asked to trace one of five *prompt* curves designed to be characteristic of 3D anatomical illustrations. The same prompts were shown repeatedly in blocks of five throughout the experiment. Within each block of five trials, each one of the prompt curves appeared once in a random order.

Care was taken to place the curves appropriately within the working volume of the Phantom to avoid accidentally reaching the limits of the Phantom’s armature. The

curves were also oriented to minimize drawing from left to right. For right-handed users, drawing in this direction is difficult enough with the tape technique that artists typically reposition their drawing rather than draw with their arms crossed. Users were required to cross their hands slightly to complete some of the curves, but were never required to draw an entire curve in such an orientation. The orientation and position of the prompts and the direction of drawing was held constant for all drawing techniques.

4.3 Training

Participants were trained in two stages. The first stage was a scripted introduction to VR and to each of the four techniques. Participants were shown how to hold the pen, as in Figure 3.4, and practiced drawing a straight line and several curved lines with each technique. They also practiced tracing some of the lines that they had drawn. Participants were also instructed about the keys to drawing controlled lines with each of the techniques. For both freehand techniques, the key described was finding the right balance for the drawing speed. Drawing too fast lacks precision, while drawing too slowly makes it hard to avoid jitter. For the drag and tape techniques, the key was to pay close attention to the guideline and to work deliberately by advancing the drawing along the guideline only after it appears to have reached the right orientation. The freehand techniques were always introduced before drag and tape because they served as an easier-to-understand introduction to VR and 3D drawing. The two scoring measures “position” and “direction”, discussed in more detail in the Results section, were also introduced during training.

The second stage of training was a miniature version of the entire experiment. Participants did five tracing trials with each of the four input conditions. The order of the conditions was the same as for the rest of the experiment. To make spatial judgments a bit easier in this training stage, additional depth cues were added by displaying 10 bull’s eyes evenly spaced along the length of the prompts. Participants were shown their position and direction scores after completing each trial in the training stage.

After these initial twenty training trials, participants did one block of twenty trials

for each of the four input conditions, for a total of eighty non-training trials.

4.4 Participants

There were twelve compensated participants in the study, six male and six female. All had significant experience in drawing with physical media. All except one were enrolled in a leading design school and reported drawing with physical media daily on a post questionnaire. The participant not at the design school also had significant collegiate-level artistic training and reported drawing with physical media at least monthly. Seven of the participants had never experienced VR before, three had experienced it one to five times before, and two had experienced it more than twenty times. Five participants had never used a 3D modeling program before, three had used such programs one to five times before, and four had used them more than 20 times before. All participants were right-handed.

4.5 Results

Two primary measures of error were used to describe performance on the task. The first, “position”, computes a mean of closest distances for the prompt P and the drawn curve D :

$$pos(P, D) = \frac{d_m(P, D) + d_m(D, P)}{2}, \quad (4.1)$$

$$\text{where } d_m(A, B) = \text{mean}(\min_{a \in A} \min_{b \in B} |a - b|). \quad (4.2)$$

The second measure, “direction”, computes the average angle between the tangents of the two curves at corresponding samples:

$$dir(P, D) = \text{mean}_{d \in D}(\arccos(d' \cdot (p' \text{ for the } p \in P \text{ closest to } d))). \quad (4.3)$$

Before computing the metrics, both curves are resampled at a constant interval of 0.3 millimeters.

Data from 20 of the 960 total non-training trials (2%) were considered outliers and removed from the analysis. The measures for the remaining trials were averaged to find per participant means. Mean scores for position, direction, and time were

Measure	Condition	Mean	SD
position (mm)	drag	1.45	0.45
	tape	1.81	0.69
	sand	2.37	0.65
	free	2.68	0.75
direction (degrees)	drag ^A	7.37	1.61
	tape ^A	7.38	1.98
	sand ^B	18.01	2.32
	free ^B	19.25	1.60
time (seconds)	drag ^C	23.88	6.03
	tape ^C	19.89	6.97
	sand ^D	13.57	5.34
	free ^D	12.75	4.92

Table 4.1: Experimental results. Differences in values with the same superscript are not statistically significant.

analyzed with an analysis of variance with input technique (drag, tape, sand, free) as a within-subjects factor. The sphericity assumption was met for position, but not for the other measures. Huynh-Feldt corrections were applied in the latter cases. The main effect of input technique was significant, for position $F(3, 33) = 37.78, p < 0.01$, for direction $F(1.62, 83.41) = 201.67, p < 0.01$, and for time $F(2.51, 27.60) = 34.69, p < 0.01$.

Post-hoc comparisons were performed using the Bonferroni adjustment for multiple comparisons and $p = 0.05$. The results are summarized in Table 4.1 and Figures 4.2, 4.3, and 4.4.

In a post questionnaire, participants were asked to rank the four drawing techniques in order from best to worst for control of position, control of direction, and control of both position and direction combined. The sand technique was always ranked third, and the free technique was always last. For position, tape received nine first place votes to drag’s three. For direction, tape received five first place votes to drag’s seven, and for control over both position and direction, tape received eight first place votes to four for drag.

Participants were also asked to rate how likely they would be to use each of the techniques if they were to create a 3D medical illustration with the Drawing on Air tool. On a scale of one to seven, with one being “not likely” and seven “very likely”,




Figure 4.2: Mean positional error across the four input conditions.

their mean responses were: drag 6.5, tape 6.5, sand 4.0, and free 2.0.

4.6 Analysis of Results

The Drawing on Air (drag and tape) techniques outperformed sand and free on both positional and directional measures, with mean errors that were roughly half that of the two freehand-based techniques. Thus, the data support our first hypothesis. In artistic practice, we see that this difference in error makes a real difference in style and subject matter.

Drag had less positional error than tape, but tape was favored in a post-questionnaire for control of position and overall control. The difference in drawing time between drag and tape was not statistically significant. However, we observed a trend that is consistent with the data collected: drag seems to be faster than tape for drawing approximate shapes, but slower for drawing very exact shapes. The difference can




Figure 4.3: Mean directional error across the four input conditions.




Figure 4.4: Mean drawing time across the four input conditions.

probably be attributed to the separation of the two tasks of setting the drawing direction and advancing along it. Once the difficulty reaches a certain threshold, it may be faster to assign one of these tasks to each hand than to overload a single hand.

The tape-drawing technique does take longer than drag to learn, and, from our experience with artists who have used the tool for more extended periods of time, we hypothesize that the slight difference we see here in performance between drag and tape will diminish over time and perhaps even reverse itself. Nevertheless, we conclude that both drag and tape techniques are valuable parts of a controlled 3D drawing suite. User preference given a particular line to draw may be the best way to select a drawing approach, and thus the tight integration of both techniques into Drawing on Air makes sense given the level of control participants exhibited with each.

Figure 4.5 indicates the types of differences we see in the lines drawn with each technique. Shown here are the four *best* tracing results obtained by one participant with each of the techniques. The prompt curve in these results is inspired by an anatomical feature on the human scapula. Thus, if we imagine this line to be part of a medical illustration, the various inflection points and shape changes are important to capture because they mark regions where particular muscles of the shoulder attach.

In the versions drawn with tape and drag, the participant has followed along the path of the line quite precisely. This is somewhat hard to see in these 2D projections but much clearer when viewed in stereo. Notice that the two lines overlap significantly, and we see from the shadow that this is also true in the Z direction. The best of the sand and free drawings are unable to capture the shape accurately. There is considerable error: sometimes, it appears as jaggy bumps in the line, sometimes the shape is just completely off. For example, the shadow in the sand result reveals a large error in the Z direction. If we imagine a more complete drawing formed by many lines like this, we can see a problem typical of freehand input: the drawing quickly becomes loose and imprecise. We get a sense of what the artist means, but not the clarified understanding we demand in applications such as medical illustration.

More sophisticated input filtering techniques might improve some of the jaggy-type error in the sand and free results. However, for anatomically inspired lines, small kinks and shape changes in line are often used to indicate important features. Since




Figure 4.5: One participant's best tracings of a line inspired by an anatomical feature. The prompt is the dotted blue line, and the user drew the solid orange line.

these variations regularly occur at the same scale as the muscular error seen with the sand and free techniques, it would be difficult for automatic data filtering to separate user error from artistic intention. The two Drawing on Air techniques successfully avoid this issue by putting the user in continuous control of the drawing direction, a strategy we can think of as a user-guided filter.

The drawing times for both the drag and tape techniques were significantly longer than for the freehand techniques. This raises the question: Would performance with sand and free be better if participants spent more time in drawing with these techniques? In practice, we find the answer is no. The sand and free techniques have a “sweet spot” in terms of drawing speed. If the drawing is done too quickly, it is difficult to capture the shape of the curve correctly; if the drawing is done too slowly, it is difficult to maintain smooth hand movement and thereby control directional error.

In contrast, the Drawing on Air techniques make deliberate drawing possible. There is no motor-control penalty associated with drawing slowly and carefully, and at a high level, guidelines built into the techniques allow the user to continuously check the position and orientation in space. As a result, we find the familiar speed-accuracy tradeoff we desire in a drawing tool. Performance only increases when artists wish to invest more time in the drawing.

Of final note is the significant difference found between sand and free techniques in the mean positional error. The difference found supports our second hypothesis. The addition of haptic frictional and viscosity forces appears to aid control in this 3D task, although not to the level of the more sophisticated Drawing on Air techniques.

4.7 Discussion

In this section, we present discussion of issues related to the design of the experiment. User-guided drawing is also introduced as a general class of controlled 3D input techniques.

4.7.1 Appropriateness of the Task

Clearly, not all lines in an illustration are tracings, raising the question of whether a tracing task is the most appropriate for testing control of the various drawing

interfaces. In fact, we pilot tested other tasks, such as replicating a line seen in the distance. One of the main obstacles in non-tracing approaches is making sure that participants have an accurate 3D understanding of the shape they are about to draw. This is very difficult to achieve across various participants with anywhere near the level of certainty attained with tracing. Thus, with tracing, the participant makes fewer errors due to lack of understanding of 3D shape, and our error measures reflect more accurately how much control the participant has over the particular technique.

Tracing is also not so different from what illustrators typically do, as we learned by working closely with illustrators and by doing our own serious illustrations. When artists work at an intricate level, the work precisely and deliberately. In anatomical drawing, for example, a line with a particular bump on it conveys exactly where a tendon attaches to a bone. When working at this level, artists are certainly not sketching. The exact shape of the line is extremely important, just as in tracing. Lines are often drawn relative, or even parallel, to other lines, and in these situations the act of drawing is almost exactly tracing.

The five prompt curves used in the tracing trials were chosen to be representative of curves found in anatomical illustration, ranging from the simple bend of a tendon going over a knuckle to a tracing of the spine of the scapula, an extremely important anatomical feature in figure drawing and illustrations of flying bats. All the prompt curves contain variation in all three dimensions.

4.7.2 Importance of Training and Depth Cues

Most participants encountered many new concepts during this study. The various drawing techniques were of course new, but so was the very idea of virtual reality and interacting with a 3D stereo display. In pilots, we found that understanding depth relationships in VR was one of the most important and challenging hurdles for novices. In normal use of Drawing on Air, artists build up an entire drawing checking depth relationships and even drawing guidelines or scaffolding as they go. In this study, however, the prompt curve is seen completely in isolation, so there are no “relative size” visual cues and very few “occlusion” cues. Cutting and Vishton provide an overview of these and other relevant cues in perceiving spatial layout in near visual space [19].

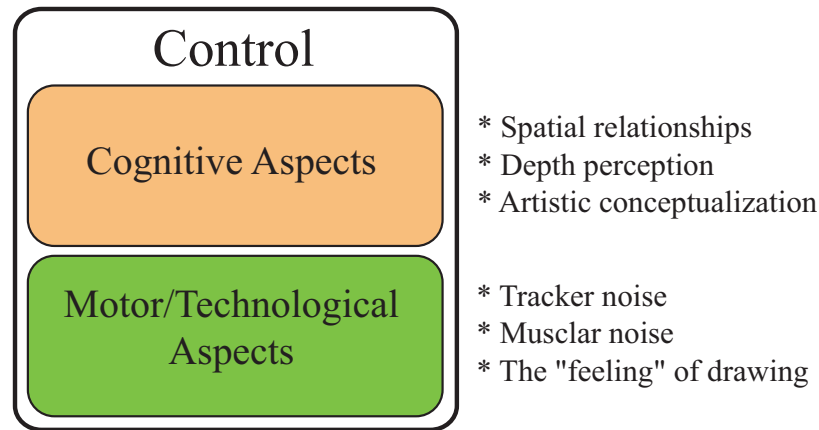


Figure 4.6: Precision in 3D drawing tasks requires both high-level cognitive aspects and lower-level motor/technological aspects of control.

To help participants learn to judge depth within this new environment, we paid special attention to the rendering of the experimental scene, shown in Figure 4.1. Everything in the scene is textured, which helps with shape perception. Shadows, a ground plane, and the bull's-eye forms used extensively during the training session were also added to provide additional cues for clarifying depth relationships. Even after these cues were added, a few participants still complained that they were having difficulty getting used to working in 3D space, but after the training trials, they became comfortable enough to perform the task accurately.

4.7.3 User-Guided Drawing and Controlled 3D Input

In this section, we discuss controlled 3D drawing in terms of requirements that must be met in order to achieve control. We then describe a class of interaction techniques, called user-guided drawing techniques, that address these control requirements at a general level. Both modes of Drawing on Air can be described as user-guided drawing techniques, and additional techniques may also be possible.

Figure 4.6 presents an overview of controlled 3D drawing. The key observation brought forward here is that two different aspects of control must be achieved in order to have precise input: high-level cognitive aspects and lower-level motor/technological aspects. These observations may seem obvious, but what is not obvious is how to achieve high-level control for a general class of drawing styles and whether solutions

to the high-level control problem can be leveraged to address the low-level control problem. We describe each area of control next, before presenting the concept of user-guided drawing, which addresses these points in more detail.

Contained within the cognitive aspects of control are issues of understanding spatial relationships and achieving effective depth perception. In order to draw with precision, users must be able to build an accurate mental model of the space in which they work. While they are drawing a 3D curve, this mental model is constantly updated based on visual feedback. Controlled approaches to 3D drawing need to facilitate building these mental models so that artists can correctly conceptualize the shape they are trying to draw and respond accurately in real time to visual feedback presented while drawing.

During drawing, motor and technological aspects of control become important. The goal is to make it possible for the user's current conception of the drawn shape to be conveyed precisely to the computer. This task is confounded by issues of noise in sending the shape signal from the user's brain through the hand to the computer. Part of the noise in the system is due to technological limitations in measuring the user's motion; we call this tracker noise, since motion trackers are used to capture this data. Another portion of the noise comes from the user's own inability to move his hand with the precision desired. We use the term muscular noise to describe this uncontrollable error and jitter in the motion of the hands. Controlled approaches to 3D drawing must be able to eliminate the negative effects of both of these causes of noise. For users with an artistic background, another important aspect of motor control is the feeling of drawing. Artists spend years developing fine motor skills used in drawing. It may be possible to leverage this existing skill in 3D drawing techniques if the techniques are designed to engage feedback mechanisms and muscle memory used in traditional drawing.

User-Guided Drawing

Techniques based in "user-guided drawing" address both of the aspects of control specified above. User-guided drawing techniques are characterized by two key properties:

1. Noise-removing filtering of the input signal is achieved through a user-driven

constraint.

2. The constraint acts as a drawing guide, helping the user to measure space and judge positioning while drawing.

The first property concerns filtering of the drawing input signal. In order to eliminate both muscular and tracker noise from the system, controlled drawing techniques must involve some aspect of interactive filtering, and to do this filtering well, the filtering algorithm must include some notion of prediction. For some special, constrained cases, it may be possible to predict automatically where the drawing will go next, at least robustly enough to filter out most noise in the signal, but for general cases of drawing input predictive filters are quite challenging to implement. The alternative solution, proposed here as a key component of user-guided drawing, is to place the predictive filtering in the hands of the user by building it explicitly into the drawing technique.

The second key property of user-guided drawing techniques is that they incorporate the filtering constraint described above into a visual guide manipulated by the user. The guide acts as a preview for the line to be drawn, and by moving the guide through space, the user constantly builds a more refined mental map of the space, addressing issues of control at the higher, cognitive level.

Drawing on Air Examples

The drag and tape modes of Drawing on Air are two examples of techniques that fit within this user-guided drawing framework. In both cases, the particular filtering constraint used is based on a drawing tangent direction. Drawing proceeds by repeatedly setting the drawing constraint and then advancing the line along it. Since the line may advance only along the tangent constraint, both tracker and muscular noise is effectively filtered out of the input signal. From a higher-level standpoint, the directional guide, visible to the user while drawing, provides a precise preview of where the drawing will go next. The user can actually stop drawing and move the guide around in space to measure and plan the next move before continuing to draw.

Potential Additional Examples

While Drawing on Air uses a tangent constraint, other constraints might also be possible. A user-guided technique based on a perpendicular constraint might lend itself to drawing circular shapes. Other more sophisticated constraints might also work, for example, French curves, constraints defined by previously drawn curves, pathways through 3D scientific data. Tangent-based constraints are a good place to begin because they work for a wide variety of shapes, but future exploration of alternative user-guided drawing techniques may lead to a more complete toolset of methods for precise, continuous 3D computer input.

4.8 Summary and Conclusions

This chapter presented the drawing control experiment, a formal user study of controlled 3D interaction comparing Drawing on Air's two drawing modes to two free-hand input techniques. Results indicated statistically significant reductions in error when using the Drawing on Air techniques, supporting the claim that Drawing on Air is a more appropriate means of achieving precise, continuous 3D computer input. Also presented in this chapter is a discussion of the experimental design of the user study, including conclusions from pilot studies and other lessons learned about robust studies of computer input in virtual reality. This chapter also presents requirements for controlled 3D input, including a discussion of user-guided drawing, a notion that provides a framework within which Drawing on Air and future input techniques may be described and evaluated.

Chapter 5

Models for User Performance in 3D Tracing Tasks

Here, we derive two models that describe user performance in 3D tracing tasks. Our work builds on previous models describing human performance, such as the Steering Law [3] and the Power Law [100, 102]. We extend these to address the unique situation of controlled, continuous 3D input. Analysis resulting from this work leads to a better understanding of what makes 3D drawing difficult and highlights differences between Drawing on Air and freehand drawing techniques.

Newell and Card have argued that the advancement of HCI lies in “hardening” the field with quantitative models grounded in psychological theory from which user behavior can be predicted [68]. One of the few areas where models of this type have been able to take hold is in describing human motor control, so-called action laws. Several examples are described in the next section.

In general, models of this sort are limited to describing fairly low-level phenomena. It is difficult to capture the complexities and nuances of tasks that require higher-level cognition on the part of the user. 3D drawing lies on the edge of this complexity boundary. Clearly, motor control is an important part of drawing, especially in 3D. However, successful illustration depends on much more than motor control. Choosing the “right” lines to draw, working with correct proportion, seeing and understanding form, refining and reworking: all these high-level tasks are an important part of drawing. Thus, it makes sense to examine drawing tasks at several levels. In this




Figure 5.1: Target touching task relating to Fitts' Law.




Figure 5.2: Steering through a tunnel.

chapter, we focus on low-level, motor control analysis in the style of human action laws.

5.1 Review of Action Laws in HCI Literature

Despite Newell and Card's call two decades ago for quantitative models to guide user interface developers, few robust models actually exist today. The most notable exception is Fitts' Law [29], which has rigorously stood the test of time and spurred vast amounts of follow-on work. Fitts' Law predicts the time T needed to point to a target of width W at a distance A , as seen in Figure 5.1. The task can be modeled with the relationship

$$T = a + b \log_2\left(\frac{A}{W} + c\right), \quad (5.1)$$

where a and b are empirically determined constants and c is 0, 0.5, or 1 [63].

In the Steering Law, Accot and Zhai derive a model for trajectory-based interactions, such as navigating through nested menus and drawing curves, from the basis formed by Fitts' Law [3]. The task used in their studies is steering through a tunnel. Initially, this took the form of drawing a line using a tablet or mouse through tunnels displayed on the screen, as seen in Figure 5.2; later, they also explored steering a car in a VR driving simulation [109]. Until our work, this was the only steering experiment performed in VR. However, the driving simulator was limited to a track lying

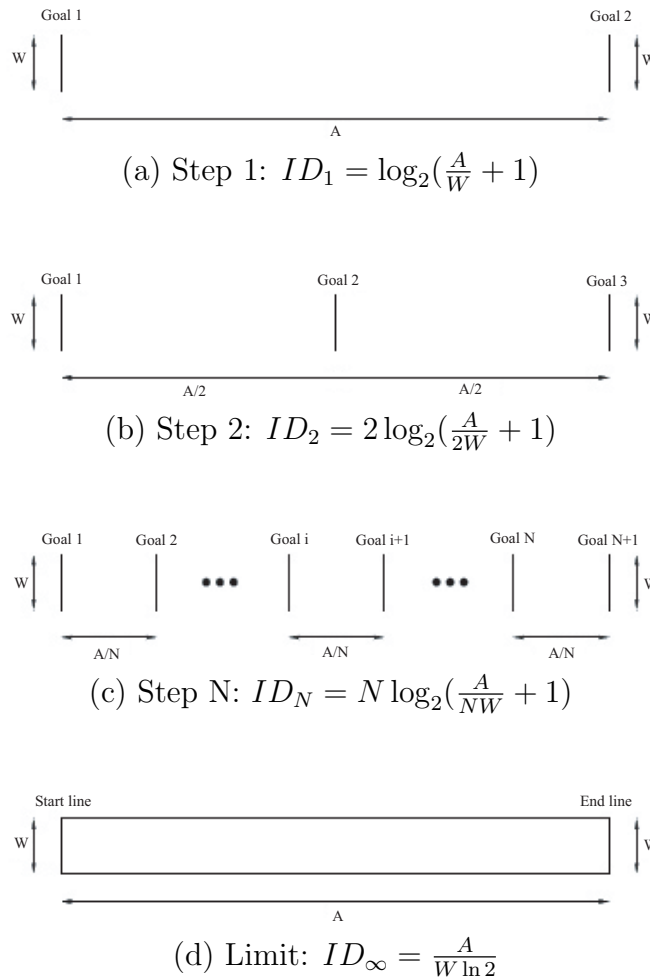


Figure 5.3: Derivation of the Steering Law from a single goal-passing task that follows Fitts' Law to a straight tunnel-steering task. Adapted from Accot and Zhai.

within a plane. Thus, despite the VR environment, the task was rather 2D in nature. A 3D hand-based task (described in more detail in section 5.5.2) was also proposed in this work, but to our knowledge, no study using this task was ever performed.

The Steering Law is derived by considering the elementary task of drawing a line through a goal of width W at distance A , as seen in step 1 in Figure 5.3. Accot and Zhai begin by showing experimentally that this task follows Fitts' Law. If the tunnel steering problem is posed as a series of these goal passing tasks, as outlined in Figure 5.3, then we can find the total drawing time by adding the contribution from each. As the number of discrete tasks goes to infinity, we take the limit and arrive at

the Steering Law for a straight tunnel:

$$T = a + b \left(\frac{A}{W \ln 2} \right) \quad (5.2)$$

which simplifies to

$$T = a + b \left(\frac{A}{W} \right). \quad (5.3)$$

Accot and Zhai also demonstrate experimentally that the law holds for the more general cases of tunnels that vary in width and that follow the path of a curve, such as a spiral. For a curved path C with varying width defined by $W(s)$, the general form of the Steering Law is

$$T_C = a + b \int_C \frac{ds}{W(s)} \quad (5.4)$$

Like Fitts' Law, the Steering Law was shown in Accot and Zhai's experiments to be robust, with correlations described by $r^2 = 0.97$ on average. In their methodology, the trial was canceled whenever the line drawn by the user exited the tunnel. The data from these trials were removed from the regression analysis and simple error rates were reported for each experiment. Accot and Zhai suggest modeling error as an important next step. The work in this chapter begins to take that step.

5.2 Review of Related Theory in Neuroscience and Motor Control

Our formulation draws upon work in the neuroscience and motor control literature. In this section, we briefly review two key areas of research. First, we review the Power Law, which describes the velocity of rhythmic drawing motions as a function of the curvature of the shape drawn. Then, we discuss work in characterizing accuracy in both perception and movement in 3D space.

5.2.1 The Power Law

The Power Law has been the subject of numerous studies in the experimental neuroscience literature. Viviani and Terzuolo first observed a systematic relationship in handwriting and unconstrained rhythmic drawing between the velocity of the end

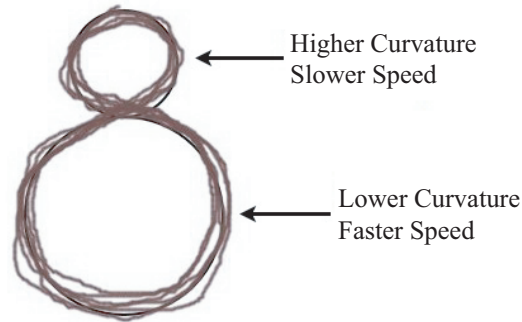


Figure 5.4: The Power Law describes a relationship between drawing speed and curvature of the drawing trajectory in rhythmic drawing motions. Drawing is slower in areas of high curvature.

effector trajectory and the geometric path it describes [102]. This can be seen clearly in rhythmic, repeated drawing of shapes, such as the figure eight shown in Figure 5.4: drawing is consistently slower in areas of high curvature.

The Power Law arising from this work states that the angular velocity $a(t)$ of the endpoint is proportional to the curvature $c(t)$ of the end effector path according the following relation:

$$a(t) = kc(t)^{2/3} \quad (5.5)$$

which can be rewritten in terms of the tangential velocity $v(t)$ and the radius of curvature $r(t)$ as:

$$v(t) = kr(t)^\beta \text{ where } \beta = \frac{1}{3} \quad (5.6)$$

The k in Equation (5.6) is an experimentally determined proportionality constant referred to as the “velocity gain factor”.

The Power Law has gained support from numerous studies [102, 100, 99, 101] and has become widely accepted as “an important invariant in biological movement trajectories”. However, recent work has demonstrated that at least the exponent β in the law, may not be as “invariant” as once believed. In particular, β has been shown to exhibit some variation depending on the scale and shape of the movements studied [74].

The tasks typically modeled with the Power Law are rhythmic motions, such as repeatedly drawing circles, ellipses, figure eights, or handwriting. Many studies allow for 3D movement, but the shapes drawn are often limited in their use of depth. The

use of 3D seems to be more important for eliminating a planar constraint from the motion rather than for exploring the possibilities of motion that varies roughly equally along three axes.

One recent study comparing sensory-based and memory-based guidance mechanisms in hand and arm movement did make fairly significant use of three dimensional space, and it is closely related to our work in that it includes a 3D tracing component [31]. Three motion conditions were tested in a VR environment. First, subjects tracked a moving target along an unseen path. Then, the entire path was displayed and retraced. Finally, all visual feedback was removed and the subject had to redraw the path from memory. The second phase, where the subject traced the path while it was displayed, is quite similar to our experimental task. One key difference is that the drawing was repeated continuously several times in a row for each condition to establish a sense of rhythm.

The authors found that subjects adhered to the Power Law most closely in the tracing condition that resembles our experimental task. Another finding of interest was that only in the tracing condition did the value for the exponent β change significantly due to the size and spatial orientation of the trajectory.

The literature in this area is ultimately concerned with understanding the organizational principles of biological motor control and the physiology of the nervous system. The fairly recent realization that at least the exponent term in the Power Law is not as robust across changes in shape and scale as previously thought has spurred a debate in the literature about the potential biological mechanisms responsible for this style of motion planning.

Despite the debate, it remains clear that a relationship does exist between the velocity of a drawing implement and the curvature of its trajectory, at least within the context of this drawing style. As there is also specific evidence for this relationship in 3D tracing tasks, we anticipate that curvature will play a role in the velocity, and consequently the measured drawing times, in our analysis.

5.2.2 Anisotropy in Control of Visually Guided 3D Motion

The differences in human perception along different axes in space are well documented in the psychological literature. The horizontal-vertical illusion is one of the most

popular examples for study [26]. The illusion arises when two lines of equal length are presented in a L or an upside down T shape: the vertical line appears to be longer.

How do distortions in human perception of space, as demonstrated by the horizontal-vertical illusion, affect motor control? Studies disagree about the extent to which the visuomotor system responds to horizontal-vertical space distortions. Some work reports that the visuomotor system is not affected by the illusion [84]. Other work shows different responses depending on the particular motor behavior studied. For example, different styles of grasping produce different results even when the same visual information is presented [98]. It is highly unlikely, in fact, that a single model can describe how perception of space relates to controlling motion.

For drawing-specific tasks, differences in accuracy between drawing in vertical, horizontal, and oblique directions has been demonstrated in several settings. In children, drawing accuracy has been shown to be highly dependent upon the context provided by the shape of the paper [7, 8]: oblique lines are difficult for children to draw on rectangular paper, but they can draw triangles accurately on triangular paper.

In adults this particular effect is non-existent or at least far less apparent; however, accuracy of drawing movements still depends significantly upon drawing direction. Strong evidence for this can be seen in nearly any animation lab or artist studio. Artists routinely adjust the orientation of the artwork to achieve the best arrangement relative to their bodies for drawing lines in different directions. Orienting artwork has also been shown to be important in 2D computer-based systems [30].

Reorienting artwork during 2D work is largely driven by optimizing the arm and hand mechanics rather than by perceptual issues. However, perception appears to play a much more important role in three-dimensional work, especially in VR environments where many of the visual cues we expect in the real world are absent.

VR has recently served as a testbed for investigating issues of perception of space [93, 107]. Most of this work has focused on depth perception with respect to relatively long (typically walkable) distances, but a few studies have examined visuomotor performance in VR within close proximity to the participant.

The most influential findings with respect to our work are for hand-controlled VR tasks. This has been studied for a pointing task [27], a continuous target-matching

task involving both translation and rotation [110], and a 3D straight-line-drawing task [92]. All studies report anisotropy in control along the horizontal, vertical, and depth dimensions. The horizontal is always the easiest to control; vertical is slightly harder, and depth is the hardest. For pointing tasks an average error ratio of 5:6:11 is reported. A similar ratio was found for the continuous target matching task, although this study found that with considerable practice, vertical error was reduced to be roughly the same as horizontal. Zhai suggests that a plausible explanation for the reduction in vertical error over time is attention priority [110]. When subjects are first learning a task, they tend to concentrate on the horizontal dimension, but as they become more comfortable, they pay more and more attention to the vertical and depth dimensions. While errors in the vertical dimension decrease with time to roughly the same rate as horizontal errors, depth errors decrease somewhat, but never reach the same level as the other dimensions because of the inherent difficulty in perceiving depth. Tano et al. [92] report similar findings for drawing straight lines in 3D space at different orientations: in general, error increases with depth away from the viewer.

5.3 Local Models for 3D Tracing Performance

In this section, three local models for 3D tracing performance, inspired by the Steering Law, are derived on the basis of the theories of motor control and perception introduced in the previous sections and are then evaluated with respect to data from the drawing control experiment described in chapter 4. Small (nearly instantaneous) segments of the 3D curves are sampled from the raw data to provide local data for this analysis. In section 5.4, we integrate the local formulations derived here to obtain models appropriate for describing user performance over an entire curve-tracing action.

The next two sections provide background on the structure, methodology, and data used in the local analysis. Then three models are derived, and it is hypothesized that each successive model improves significantly over the previous one. This hypothesis is evaluated with respect to experimentally collected data, and several conclusions are drawn from the analysis.

5.3.1 Structure and Methodology

The action laws from which we derive our inspiration begin by defining an “Index of Difficulty” for a task and then demonstrate that a relationship exists between this index and the predicted variable. In the Steering Law, for example, the $\frac{A}{W}$ term is the Index of Difficulty, and the law describes a linear relationship between this index and drawing time (see Equation (5.3).)

We follow a similar structure. In the sections below, we derive a Local Index of Difficulty for 3D tracing, D , and investigate a relationship among D and the measured values: task completion time T , positional error E_p , and directional error E_d . As in Fitts’ Law and the Steering Law, the relationship we anticipate is a linear one,

$$T = a_1 + b_1 D \quad (5.7)$$

$$E_p = a_2 + b_2 D \quad (5.8)$$

$$E_d = a_3 + b_3 D \quad (5.9)$$

where a_n and b_n are empirically determined constants.

The success of the linear model in the Steering Law, which is used to describe a similar drawing-style task, motivates our hypothesis that the drawing-style relationships we study will also exhibit a linear rather than some alternative nonlinear relationship with D . As we proceed with the analysis, this hypothesis is reevaluated several times, both by statistical analysis and by visual inspection of scatter plots of the data. Several scatter plots of the data are reproduced to illustrate the linear trends exhibited in the data.

5.3.2 Connection to the “Tunnel-Width” Term of the Steering Law

The Steering Law states that expected drawing time decreases with the width of the tunnel through which we steer. For controlled drawing, rather than the “as fast as you can” style of drawing used in steering tasks, there is no tunnel. In a controlled tracing task, the goal is to get as close to the prompt curve as possible. In Steering Law terms, this means the width of the tunnel, W , is zero, which makes the task infinitely hard according to the Steering Law’s Index of Difficulty. The question is

raised, does “tunnel width” apply to more controlled drawing situations, such as tracing?

The purpose of the tunnel in steering tasks is to describe explicitly acceptable error, a notion that, intuitively, should also apply to more controlled drawing situations. In tracing, we hypothesize that each subject has his or her own mental model of acceptable error. Drawing time is likely to decrease with increases in this user specific value, just as with increases in tunnel width. Unfortunately, this value is nearly impossible to quantify, as it is likely to vary both between subjects and within subjects, according to practice, fatigue, attention, and drawing skill.

Rather than measure this challenging acceptable-error term, our approach is to remove its effects from the data by calculating the mean drawing times and errors between subjects. This provides a way to examine error and drawing time as a function of the properties of the curve being traced without the confounding influence of participants’ interpretation or skill in performing the task.

5.3.3 Model 1: Local Orientation

The next three sections introduce three different models for an index of difficulty appropriate for describing local performance in 3D tracing tasks. The first local model is derived from the theory of anisotropy in perception and motion errors in 3D manipulations, as discussed in section 5.2.2 above. It follows from this theory that both position in space and drawing direction (local curve orientation) are likely to play a role in the difficulty of drawing a local segment of 3D curve. This suggests that $D(s)$, the difficulty of 3D tracing at point s along a curve, is a function of both the 3D position and the orientation at point s :

$$D(s) = f(p(s), \vec{d}(s)). \quad (5.10)$$

Our data do not contain a regular enough sampling of 3D positions to confirm that difficulty is indeed affected by spatial positioning. Thus, despite its potentially important role in 3D input problems, spatial positioning is not factored into our model and is left as a potential future extension to the work presented here. Local curve orientation, on the other hand, is quite significant in our analysis and forms the basis for the first model for task difficulty.




Figure 5.5: Relationships among x , y , and z components of the local orientation of the prompt curve and measured directional error.

Orientation is described based on its components in the horizontal (+ X to the right), vertical (+ Y up), and depth (+ Z out of the screen) dimensions relative to the viewer. For a sample along the curve being traced, orientation, $\vec{d}(s)$, is a normalized vector,

$$\vec{d}(s) = \langle d_x(s), d_y(s), d_z(s) \rangle. \quad (5.11)$$

There are likely to be some differences between positive and negative movement along some of these axes. However, related work reports small differences for horizontal and vertical movement. Thus, a reasonable measure for difficulty of movement along X and Y dimensions is a weighted combination of the absolute value of the components of the orientation: $w_x|d_x(s)| + w_y|d_y(s)|$. In the Z dimension, movement into the screen ($-Z$) is harder than movement out of the screen ($+Z$). Therefore, $-d_z(s)$ is a more accurate measure of error than $|d_z(s)|$. At orientations approaching $+Z$, this approximation breaks down: movement directly out of the screen is likely to cause an increase in difficulty. A model intended to capture 3D input tasks at a more general level should take this case into account. Our data do not contain such sharp angles, however, so $-d_z(s)$ is used as an approximation in the model that follows. Figure 5.5 shows data for directional error as a function of orientation that support the general observations described here.

This derivation leads to the first model for local difficulty in 3D tracing,

$$D_1(s) = w_y|d_y(s)| + w_z(-d_z(s)). \quad (5.12)$$

Two additional models are defined in the next sections, each adding an additional term to the structure presented here. The fit of each model to experimental data is

analyzed in section 5.3.6.

5.3.4 Model 2: Incorporating Local Curvature

From the Power Law, we expect instantaneous drawing speed $v(s)$ to be related to the radius of curvature $r(s)$ of the curve drawn,

$$v(s) = kr(s)^\beta. \quad (5.13)$$

Rewritten in terms of curvature, $K = 1/r$, we have

$$v(s) = k \left(\frac{1}{K(s)^\beta} \right) \quad (5.14)$$

For an incremental piece of the curve ds , the time dt taken to draw ds is

$$dt = \frac{ds}{v(s)}, \quad (5.15)$$

which leads to a local instantaneous formulation of drawing time as a function of curvature,

$$dt = \frac{K(s)^\beta ds}{k}. \quad (5.16)$$

Integrating over a curve or section of curve, C , leads to an expected drawing time for the task expressed as a function of local curvature,

$$t = \int_C \frac{1}{k} K(s)^\beta ds. \quad (5.17)$$

Following the structure of the Steering Law, we reason that task completion time is proportional in this context to task difficulty. It follows that local 3D tracing difficulty as suggested by the Power Law can be written as

$$D_K(s) = K(s)^\beta, \quad (5.18)$$

where the velocity gain factor k is assumed to be absorbed in the empirically determined constants of the linear relation between D_K and the performance measures, and $K(s)$ is a measure of the local curvature of C at point s .

The second model for local difficulty adds this D_K term to the first model:

$$D_2(s) = w_y |d_y(s)| + w_z (-d_z(s)) + w_k K(s)^\beta. \quad (5.19)$$




Figure 5.6: Comparison of results for a range of values for the exponent β in Equation (5.18). Plotted values are adjusted coefficients of determination for regressions with measured directional error.

The first hypothesis to be evaluated in section 5.3.6, which contains a comparative analysis of the local models, is that D_2 , which takes both anisotropy in 3D error and the effects of local curvature into account, will be more highly correlated with experimental results than the first model D_1 , which lacks the contribution of curvature.

Before introducing the final local model, it is useful to consider the role of the exponent β in the model described above. Recall from the discussion of related work that the value first proposed for β of $1/3$ has recently been shown to vary considerably based on the scale and style of drawing trajectory studied. To investigate the proper value for β for this tracing task, linear regression analysis was used to compare the correlation between D_K and experimental data for a range of values of β . Results are summarized in Figures 5.6, 5.7, and 5.8. For directional error and drawing time, the data suggest clear trends and clear optimal values for the exponent β . For positional error, the trend is less clear, as local curvature appears to be a far less reliable measure for describing positional error, especially for the freehand drawing techniques. Table 5.1 summarizes of the best β for each combination of performance measure and drawing interface. Values under the “best” column are used in the analysis of the local models in section 5.3.6. Scatter plots of measured performance against local curvature raised to the best β as reported in table 5.1 are shown in Figures 5.9, 5.10,




Figure 5.7: Comparison of results for a range of values for the exponent β in Equation (5.18). Plotted values are adjusted coefficients of determination for regressions with measured positional error.

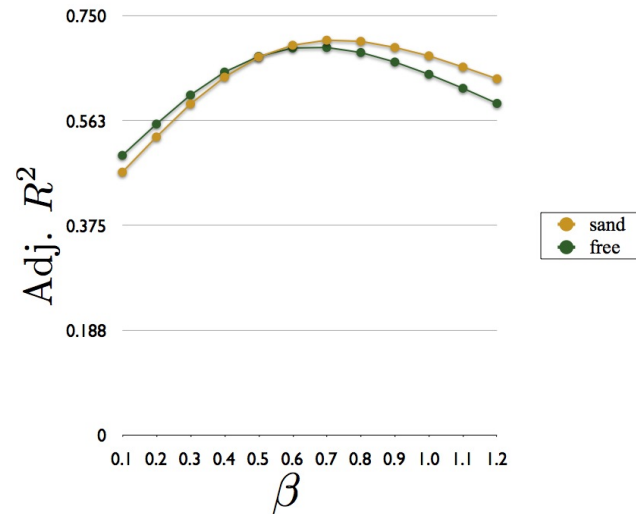


Figure 5.8: Comparison of results for a range of values for the exponent β in Equation (5.18). Plotted values are adjusted coefficients of determination for regressions with measured drawing time.

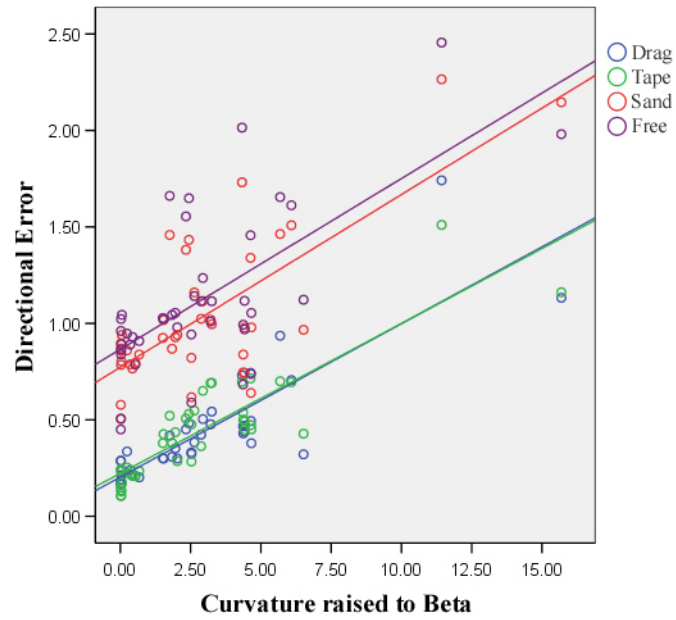


Figure 5.9: Scatter plot of local curvature and measured directional error.




Figure 5.10: Scatter plot of local curvature and measured positional error.




Figure 5.11: Scatter plot of local curvature and measured drawing time.

	Best β				best
	drag	tape	sand	free	
Dir Err	0.8	0.6	1.0	0.7	0.8
Pos Err	0.2	0.1	0.1	0.1	0.2
Time	-	-	0.7	0.7	0.7

Table 5.1: The best values for the exponent β for each drawing condition. Values under the “best” column are used for all cases in the comparative analysis that follows.

and 5.11. As in Figure 5.5 these plots provide some sense of the “raw” data and an initial assessment of the appropriateness of a linear model.

5.3.5 Model 3: Incorporating Interactions

Model D_2 (Equation (5.19)) assumes that the two terms of the model do not interact: that is, that drawing difficulty from local curvature is independent of local orientation. Intuitively, it seems unlikely that this assumption is valid. For example, the effects of curvature may be greater when the drawing orientation points away from the viewer as compared to the “easier” situation of drawing to the left. The third local model examined in this chapter builds on model D_2 by incorporating a cross-product term

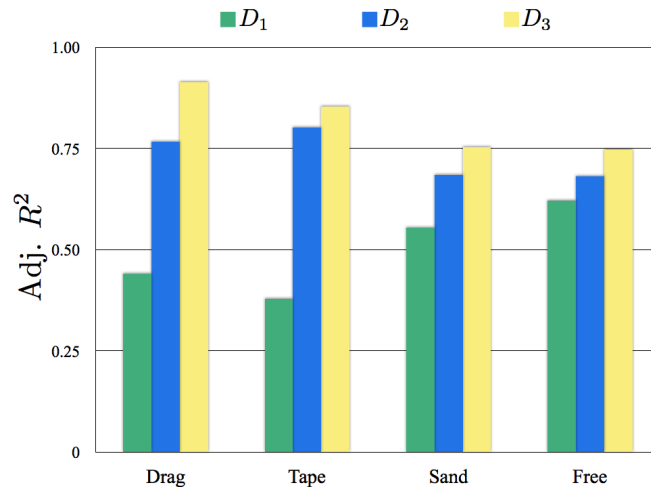


Figure 5.12: Comparison of local models for directional error. Differences between models are significant. (Hierarchical multiple regression, F-Test on R-square change, $p < .05$.)

to capture interactions between orientation and curvature:

$$D_3(s) = w_y|d_y(s)| + w_z(-d_z(s)) + w_k K(s)^\beta + K(s)^\beta(w_i|d_y(s)| + w_j(-d_z(s))). \quad (5.20)$$

This leads to a second hypothesis to be evaluated in section 5.3.6: the addition of the interaction terms in model D_3 will improve the correlation of the model with experimental data as compared to model D_2 .

5.3.6 Comparison and Analysis of Local Models

This section evaluates the following two hypotheses with respect to experimentally collected data and sets forth conclusions on the most appropriate local model for 3D tracing tasks.

Hypothesis 1: The addition of the curvature term in model D_2 , significantly improves the correlation of the model with the experimental data as compared to D_1 .

Hypothesis 2: The addition of the interaction terms in model D_3 , significantly improves the correlation of the model with the experimental data as compared to D_2 .

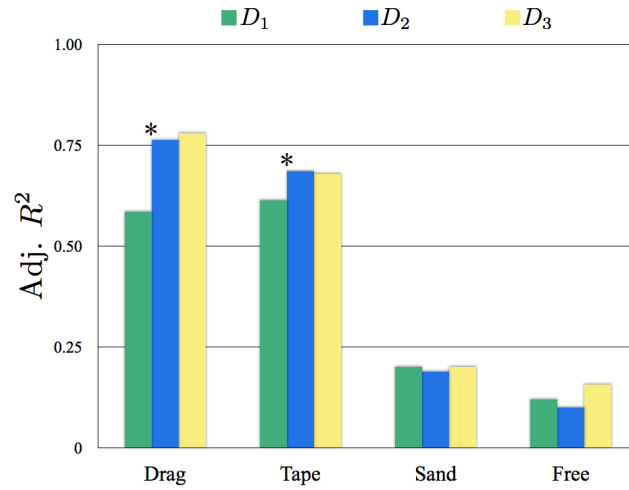


Figure 5.13: Comparison of local models for positional error. Asterisks denote significant differences. (Hierarchical multiple regression, F-Test on R-square change, $p < .05$.)




Figure 5.14: Comparison of local models for drawing time. Differences between models are significant. (Hierarchical multiple regression, F-Test on R-square change, $p < .05$.)

Model Comparison

Figures 5.12, 5.13, and 5.14 report results from comparison of the models. Due to the nested structure of the three models, hierarchical multiple linear regression is used to analyze the data, and significance tests are performed on the R-square change between the models ($p < .05$). Adjusted coefficients of multiple determination are reported to account for the difference in the number of terms in each model.

For directional error (Figure 5.12), each successive model adds significantly to the correlation with experimental data. The data support both Hypothesis 1 and 2. Results also indicate an additional trend that is consistent across much of the analysis. The importance of curvature in the combined models is greater for the two Drawing on Air techniques than for the freehand techniques.

For positional error (Figure 5.13), the interactions between local curvature and local orientation never significantly improve the model, but the addition of curvature in D_2 is a significant improvement over D_1 for the two Drawing on Air techniques. The data support Hypothesis 1 only for the Drawing on Air techniques and do not support Hypothesis 2 for any of the techniques.

For drawing time (Figure 5.14), each model represents a significant improvement over the previous one. Both Hypotheses 1 and 2 are supported by the data.

Final Local Model Structure

Analysis suggests that each term in model D_3 (Equation (5.20)) is important in capturing some aspect of a local measure of difficulty, but, as indicated by the results presented here, not all terms are significant in each input technique/performance measure situation. Figures 5.15, 5.16, and 5.17 demonstrate the strong linear trends seen in the experimental data when plotted against the local index of difficulty defined by D_3 . We conclude that Equation (5.20) gives an appropriate structure for the local index of difficulty and that, following the analysis of significance reported in the previous section, the complete form of the model can be simplified depending on the particular drawing task and performance measure in question.




Figure 5.15: Local measured directional error plotted against the local index of difficulty, D_3 .

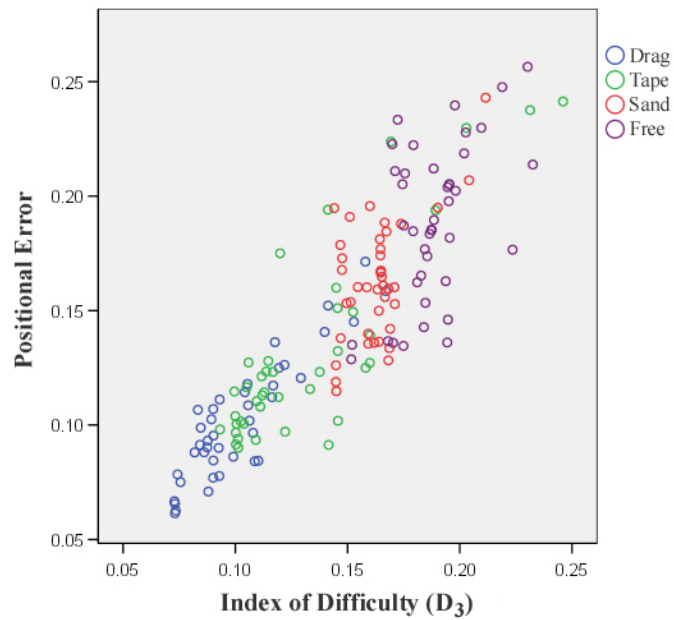


Figure 5.16: Local measured positional error plotted against the local index of difficulty, D_3 .




Figure 5.17: Local measured drawing times plotted against the local index of difficulty, D_3 .

5.4 Global Models for 3D Tracing Performance

The previous sections have developed local models to describe performance in tracing a small section of 3D curve. In this section, these local formulations are integrated along entire 3D curves to arrive at global models to describe performance in tracing a full curve.

The following three global models correspond to the local ones proposed in the previous sections and are derived by integrating the local formulations along a 3D curve C :

$$G_1(C) = \int_C (w_y |d_y(s)| + w_z (-d_z(s))) ds, \quad (5.21)$$

$$G_2(C) = \int_C (w_y |d_y(s)| + w_z (-d_z(s))) ds + \int_C w_k K(s)^\beta ds, \quad (5.22)$$

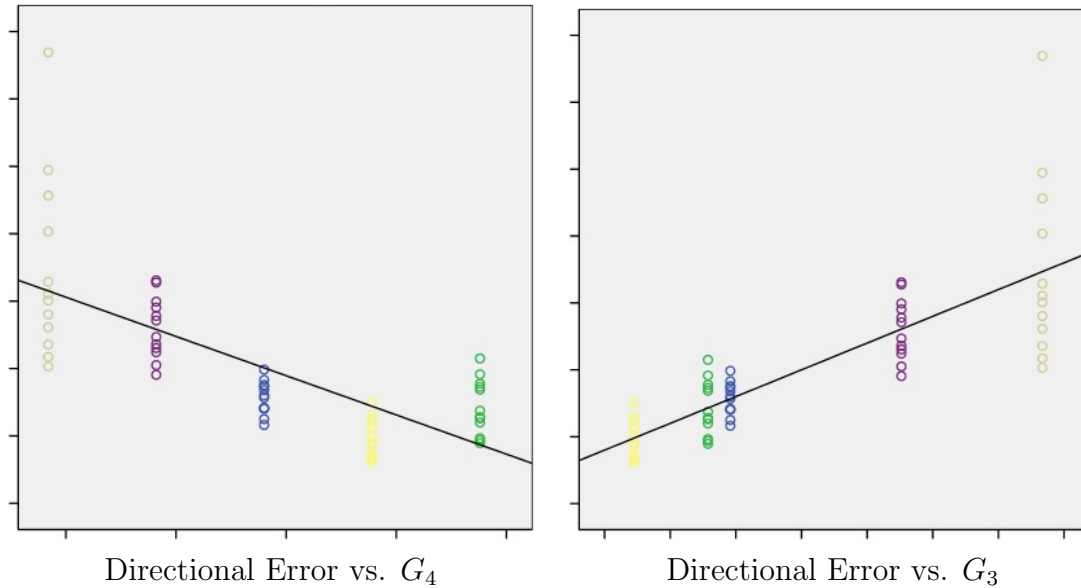


Figure 5.18: Trends in directional error versus global models for difficulty (left: G_4 , right: G_3). In both cases, we expect measured error to increase with increases in the measure of difficulty. Surprisingly, we see a decreasing trend for G_4 , which is derived from a strict interpretation of the Steering Law.

$$G_3(C) = \int_C (w_y |d_y(s)| + w_z (-d_z(s))) ds + \int_C w_k K(s)^\beta ds + \int_C K(s)^\beta (w_i |d_y(s)| + w_j (-d_z(s))) ds. \quad (5.23)$$

In the analysis that follows, these models are compared with a model suggested by a strict interpretation of the Steering Law:

$$G_4(C) = \int_C \frac{ds}{W(s)}. \quad (5.24)$$

In controlled drawing, the $W(s)$ term of this Steering Law model reduces to a user-specific notion of acceptable error, as discussed in section 5.3.2. This constant term is absorbed in the constant parameters of the linear relationship, leaving an index of difficulty defined simply by the arc length of the curve:

$$G_4(C) = \int_C ds. \quad (5.25)$$

The model inspired by the Steering Law, G_4 , presents an interesting case for study. Even in this simplified form, the Steering Law makes intuitive sense. Difficulty should

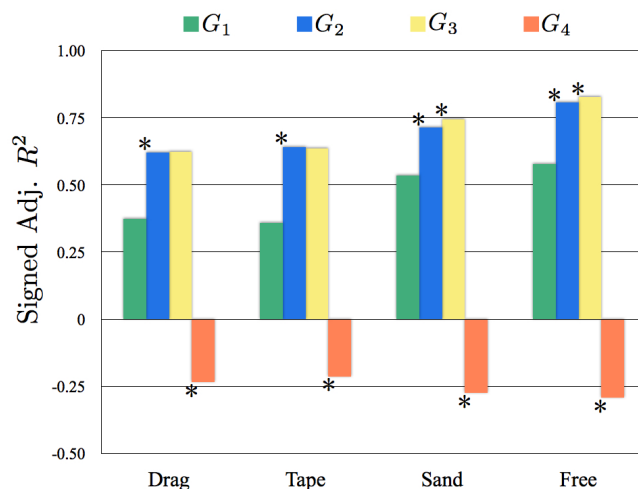


Figure 5.19: Comparison of global models for directional error. Asterisks denote significant differences. (Hierarchical multiple regression, F-Test on R-square change, $p < .05$.)

increase as the arc length of the curve drawn increases. Interestingly, our data describe the opposite trend: error decreases as arc length increases. Figure 5.18 compares this trend to the more expected trend found using the most successful global model, G_3 .

A probable explanation for the usual trend found with G_4 is that the differences in arc length in the data examined are not great enough to overwhelm the more subtle contributions of orientation and curvature. The prompts used in the drawing control experiment of the previous chapter ranged from roughly 16 to 22 cm in length. The difference of 6 cm between the longest and shortest prompt curve is roughly twenty percent of the extent of working volume for the task. With more drastic variation in the arc length of the curves tested, we hypothesize that differences in arc length will dominate the effects of curvature and orientation that appear to be more closely correlated with difficulty in data examined here. Finding the threshold for this hypothesized shift is an interesting area for future investigation.

Complete comparisons of the four global models applied to directional error measures are reported in Figure 5.19. The addition of curvature by G_2 makes a significant contribution in each case, while the interaction term included in G_3 significantly improves the correlation only in the freehand cases.

For analysis of drawing times and positional error, regression fits for the global

models consistently reported nonsignificant trends. We hypothesize that these trends are too subtle to identify given the limited number of prompt curves in the data. Recall that just five different prompt curves were used in the drawing control experiment. In the local analysis described earlier, small samples were drawn from these curves to provide a large number of data points for analysis. However, for global trends, analysis is based on just five distinct values for the differential properties and placement of the curves, making trends based on these values difficult to interpret from the data.

5.5 Suggested Follow-On Investigations

This section describes two proposed follow-on studies that build on the findings of this work and create a bridge back to the Steering Law. The first examines the importance of curvature in 2D steering tasks, in effect revising the Steering Law to provide a tighter fit for curved paths and situations where there is little variation in path length. The second builds upon the drawing control experiment of this dissertation and work proposed by Zhai and Woltjer [109] to explore 3D steering rather than controlled 3D tracing tasks.

5.5.1 Bending the Steering Law: Curved Paths and Paths of Equal Length

Both the results in this chapter and those of researchers studying the Power Law suggest that curvature is indeed an important factor in predicting drawing time, even in 2D. This suggests that the Steering Law might be improved in two areas: 1. It currently does not handle cases of comparisons between curves of equal length and width. 2. A tighter fit for curved paths, even when they do vary in length, is likely to be achieved by including curvature in the relation.

Accot and Zhai note the potential benefits of curvature for creating an improved local version of the Steering Law, citing research on the Power Law [3]. However, the experimental results reported are so highly correlated with the Steering Law Index of Difficulty, which does not include curvature, that little motivation remains to examine curvature as an important contributor to drawing time for these tasks.

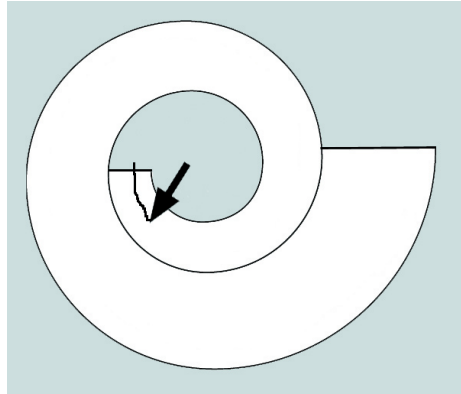


Figure 5.20: In the spiral tunnel task tested by Accot and Zhai [3], the width of the tunnel increases as curvature decreases. In addition to the change in width of the tunnel, changes in curvature may also have played a role in the trends observed in the experimental data.

Further inspection of the experimental task used to evaluate the curved path scenario in the original Steering Law investigations reveals a potential confounding factor that may explain the apparent unimportance of curvature: the tunnel used in the spiral drawing task (similar to the one in Figure 5.20) increases in width at the same rate that it decreases in curvature. Without controlling for these two potential contributors to difficulty, it is difficult to attribute the high correlation with the Index of Difficulty seen in the experimental results to either the change in width (Accot and Zhai’s interpretation) or the change in curvature (left relatively unexplored).

A spiral is probably not the most easily generalized test case for paths that vary in curvature. Gestures, handwriting, or the trajectory of an animated bouncing ball all seem more closely related to drawing paths typically found in computer-based tasks. However, since the spiral has already been established as a testbed for Steering Law experiments, an interesting thought experiment arises from considering a related experimental design that controls for width and curvature while steering through a spiral form.

Two sketches of additional spiral steering tasks for such a design are shown in Figure 5.21. The original task considers the case where width increases as curvature decreases. The task on the left of Figure 5.21 covers the case of constant width and decreasing curvature, and the spiral on the right covers decreasing width and decreasing curvature. When the spiral is traced from the outside in, we cover the




Figure 5.21: Left: A spiral tunnel with constant width. Right: A spiral tunnel that decreases in width as curvature decreases.

three corresponding cases for increasing curvature while drawing.

These three spiral configurations will produce data points for a wide variety of curvature and tunnel-width combinations. Local analysis of drawing speed with respect to the resulting curvature and tunnel width pairs is likely to provide considerable insight into the relative importance of tunnel-width and curvature in 2D steering tasks.

5.5.2 Considerations in Testing 3D Path Steering

The analysis of this chapter focuses on tracing tasks, which are quite representative of controlled drawing in general.¹ An equally important direction of inquiry, although less applicable to the controlled drawing interfaces explored here, is the investigation of 3D steering tasks as motivated by the Steering Law. Our work is likely to generalize in this direction, but since our experimental task is tracing rather than tunnel steering, we can only speculate on the relation of path width to curvature and length in 3D steering tasks.

Zhai and Woltjer propose a ring and wire task for investigating path steering in 3D hand-based tasks [109]. This is a reasonable choice for a representative task, but

¹See the discussion of the appropriateness of the task in the previous chapter (section 4.7.1) for justification.

it is not exactly a drawing task, making it less applicable to drawing-style interfaces than the related task of drawing a line through a tube-shaped tunnel. The difficulty with drawing through a tube is appropriate rendering of the tube to allow for accurate perception of both the tube form and the drawing implement that is maneuvered inside it. This issue was investigated in pilot studies that preceded the drawing control experiment of the previous chapter. If implemented with attention to rendering details, both these tasks should produce reasonable results for 3D Steering Law-style analysis.

3D steering is likely to follow several of the trends established by previous 2D steering studies and by the 3D controlled drawing work presented here. For example, it is likely that, just as increased tunnel width in 2D decreases drawing time, the same will be true in 3D. The increased difficulty in areas of high curvature found in the analysis of this chapter is likely to carry over to less controlled, steering tasks, as are the correlations between drawing direction and difficulty.

The remaining questions to be addressed in modeling 3D steering have to do with the interactions and relative importance of the difficulty factors: 1. At what point does the length of the drawing trajectory overwhelm the contributions to difficulty of curvature and orientation? 2. Where do width, curvature, and orientation lie in a relative ranking of difficulty factors? 3. Does width exhibit a similar trend to length in that, within some range of similar widths, curvature and orientation are more important than width, but for extremely different values, width dominates the other factors? If so, where does the cutoff point lie?

The challenge of the experimental design required to address these questions will be to sufficiently explore the space of interactions among the different factors (length, width, curvature, 3D orientation) while avoiding an overwhelming number of experimental trials. A key to realizing such a design may be to assume that the general trends identified by the analysis in this chapter and in previous steering experiments hold in this new 3D context and to limit the focus of the experiment to the specific remaining questions identified in the previous paragraph. Such an experiment may reveal important insights for extending the work presented here from highly controlled input scenarios to more loosely controlled tasks, useful for gesture and marking-style interaction techniques.

5.6 Summary and Conclusions

This chapter presented local and global statistical models for predicting user performance (drawing time, positional error, and directional error) in 3D tracing tasks.

The models are grounded in related human action laws presented in the human-computer interaction, neuroscience, and motor control literature. The relationship between drawing amplitude and acceptable error posed in the Steering Law, the curvature-velocity relationship described by the Power Law, and the anisotropy in perception and motor control along different 3D axes relative to the user found in the experimental psychology literature are all incorporated in the models. Orientation, curvature, and interactions between orientation and curvature were all found to be significant factors in modeling user performance in 3D tracing. A formal structure for local (Equation (5.20)) and global (Equation (5.23)) indices of difficulty for 3D tracing tasks was presented on the basis of these findings. For some combinations of drawing interface and performance measure, some of the terms in these models are insignificant, and the structure may be simplified in these situations.

The performance models of this chapter have two important implications for researchers and designers of 3D applications. First, like other human action laws, they provide a structure for predicting the difficulty of particular design decisions that involve controlled 3D input. Second, they establish common metrics for comparison of future controlled, continuous, 3D computer input strategies. Our work highlights the second application area, as one of the contributions of this analysis is a better understanding of differences between the Drawing on Air and freehand drawing techniques. In particular, it was found that difficulty is far more highly correlated with curvature with Drawing on Air than it is with freehand drawing techniques. This highlights the inherent bias toward straight lines imposed by the tangent constraint used in Drawing on Air and suggests targeting more effective drawing in areas of high curvature as an important future direction. Two follow-on studies are also suggested that would create bridges from this work back to the Steering Law, which has provided inspiration for the investigation. Our results suggest interesting future investigations in both 2D and 3D path steering.

Chapter 6

Designing VR Visualizations with Artists: Experiments and Methodology

This chapter reports on applications of 3D drawing to problems of visual design in virtual reality (VR). A new methodology, made possible by the controlled 3D drawing techniques presented previously, is proposed for design of VR scientific visualizations. One of the most challenging aspects of developing scientific visualizations is designing effective visual codings and abstractions for the data. Unlike technical challenges in simulation, data processing, and developing interactive rendering algorithms, this is best described as a *visual design* problem, and it is made particularly challenging by the unusual visual characteristics of several of our most prominent visualization media, including VR.

As is customary in visualization, we turn to visual guidelines [95], insights on human perception [104], and the study of time-tested artistic techniques [61] for direction in solving these visual problems. Guidelines are often difficult to interpret, however, and they rarely describe how to handle the conflicting requirements imposed by multivariate visualizations. Even when we find relevant guidelines to direct our work, applying them to the unusual immersive visual space we find in VR is rarely straightforward. We are thus left with a difficult visual design problem that typically requires an iterative solution of the form: design, evaluate, redesign, reevaluate, etc.



Figure 6.1: Designs for VR scientific visualizations in traditional artistic media. While artists can clearly contribute to visualization problems using these media, they can contribute additional insight if given appropriate tools for creating and refining visual ideas within target visualization environments, such as VR.

The underlying premise of the work presented in this chapter is that successful collaborations with visual experts, such as artists, illustrators, and designers, have great potential for addressing these challenging *visual* problems. After several successful collaborations between our visualization research lab and artists at the Rhode Island School of Design, we have become convinced that this collaboration can be an important aid to science. Indeed, this idea is not without precedent. Donna Cox pioneered the development of “renaissance teams” where experts in art, science, and technology come together to make effective illustrations of science [17]. Vibeke Sorensen’s work describing an artist’s contribution to scientific visualization presents a noteworthy account of artistic collaboration for solving visual problems in scientific domains [87].

In this chapter, we build on these collaborative approaches to scientific visualization along with insights gathered from our own experiences. Our contributions fall into two areas. First, we present a formal methodology for the style of artistic collaboration that we have found beneficial. In particular, we describe appropriate roles for artists, scientists, and computer/visualization experts within this process. We also discuss the pitfalls and rewards that are likely to arise and provide a clear analysis of the design task.

Our second contribution lies in insight gained from experimentation with design tools to support artistic collaboration in visualization. In particular, we present results of four experiments used to refine 3D drawing tools in “sketching” visualization ideas directly in VR. We chose to pursue tools based on drawing metaphors because the act of drawing plays such an integral role in almost all design processes [14]. Other toolsets may also be useful, however, and we hope to contribute to a body of knowledge characterizing the strengths and weaknesses of these tools with respect to their ability to design subjects of importance to science.

Our focus is VR-based visualization. While we believe many of the insights presented here will transcend the VR medium, VR is a particularly challenging and interesting case to study. The following observations have been true of our group’s experience with VR visualization [86, 103, 111] and have helped to motivate our approach:

1. Designs, including prototypes, for new visual VR techniques are difficult to

evaluate in any medium other than VR.

2. Implementation via programming for immersive VR tends to be even slower than other computer-based visualization media.

While drawings and quick paintings, like those in Figure 6.1, are clearly important tools for exploring visual ideas, it is difficult to evaluate how well even the most sophisticated concept drawings will work when translated to a VR environment. Traditional design media, such as drawing, painting, Photoshop, and desktop-based 3D modelers, simply do not capture the unique VR experience: head-tracked stereo vision, multimodal interaction, and sense of presence and immersion. As such, it is difficult to critique these traditional designs accurately without seeing them at least partially realized in VR.

As suggested by our second observation, evaluating visual designs in series with traditional software development can result in long design-iteration times. Constraints on software development for VR, including rendering graphics at interactive rates, synchronizing rendering across multiple displays, and synchronizing and responding to input from multiple sources, often make implementation slow.

It stands to reason that tools for rapid prototyping of visual ideas in VR can be valuable in developing effective visualizations if they let us explore visual ideas within VR in less time than is necessary to program these ideas. Thus, our approach is to collaborate with visual experts to create VR-based visual prototypes that we can critique and refine. Only once we have converged upon a refined visual design do we turn to serious, time-consuming implementation to arrive at a fully data-driven visualization application.

In the sections below, we begin with an overview of related work. We then describe a series of four experiments in collaborative VR visualization design, each yielding conclusions about the most appropriate tools to support artistic involvement in the process. We present a methodology we call Scientific Sketching, that combines the conclusions from our experiments into a coherent framework to guide future collaborative efforts. Finally, we give an analysis and discussion of the work.

6.1 Related Work

In this section, we describe related approaches to collaboration with artists, and we compare our work with techniques in software engineering and prototyping.

6.1.1 Artistic Collaboration for Scientific Visualization

Many researchers in the visualization community have recognized the important role that art and artists can play in informing effective visualization strategies. One important subarea of this research involves developing visualization techniques from the study of successful artistic technique. For example, Kirby et al. provide an overview of painting technique applied to 2D multi-variate visualization [58, 59]. Many other techniques for art-based or non-photorealistic rendering have been demonstrated in both the visualization [42, 47, 61, 62] and graphics communities [34].

Other approaches that are more applicable to our methodology involve significant collaboration with practicing artists rather than study of artistic technique. Many of these follow a renaissance-team model, by which experts from art, science, and technology work together to produce scientific imagery [17]. The distinguishing characteristic of much of this work is the particular role that artists play in the collaboration. Sorensen outlines several possibilities for these roles in a collaborative scientific process [87]. We often think of artists only when we reach the dissemination stage of scientific research. While artists can certainly contribute at this point, this is a too limited use of artistic insight. As Sorensen explains, artists can play key roles throughout the scientific process, notably in the design and conceptualization stages that come early in the scientific process. We hope to make possible more significant contributions from artists in early visualization design stages.

Another recent research area in artistic involvement in visualization is in evaluating visualization techniques. Jackson et al. show that expert visual designers can predict user performance with different visualization techniques on tasks required for analysis of 2D fluid flow [48]. This work makes a quantitative case for the efficacy of incorporating artistic critiques of visualizations in an evaluation process. Our goal is to take this role for artists a step further. In addition to helping us evaluate visualizations, we want experts trained in art and design to collaborate in posing new

visualization designs and refining existing ones.

6.1.2 Software Rapid Prototyping

In software and usability engineering terms, our methodology is closely related to development via rapid prototyping, which also embraces an iterative approach to design and recognizes the costly nature of implementation via programming [12]. Learning by evaluating rough (not completely functional) prototypes early in the development process is the premise of this approach.

One of the most successful application areas for this style of software development is in user interfaces. The benefits of incorporating feedback from user testing have been well documented in this context. In some cases, the prototypes are minimally functional and may even be constructed from paper.

In a related approach, functionality of prototypes can be faked for the purpose of user testing in what has been termed a “Wizard-of-Oz” approach [23, 57]. Here a technician or “wizard,” who is typically hidden from the user, controls the system so that it responds to user feedback, simulating the effects of features that are challenging or costly to implement — speech recognition, for example. Our methodology incorporates Wizard-of-Oz techniques for prototyping aspects of the VR visualizations that respond to user interaction.

While the idea of this design style is not new, several researchers have recently called for a renewed focus on design strategies, particularly in visual and interface-centric applications. Notably, Buxton, in a soon-to-appear book on the importance of design techniques in interface development, cites Wizard-of-Oz techniques as one of the chief means of achieving something akin to a design sketch, specific to interactive situations [14]. Similarly, a recent special issue of IEEE Pervasive Computing was devoted to rapid prototyping for ubiquitous computing [20]. Ubiquitous computing and VR pose similar problems for software development in that in both the user experience is simply very difficult to capture via traditional design media.

Our main technical contribution in this area is combining these interactive prototyping techniques with tools for drawing out design ideas in VR, thus allowing us to explore-by-drawing within the complex space of possible interactive VR visualizations and leverage the artistic skills of our collaborators.

6.1.3 Toolkits for Rapid Visualization Development

Several tools based on visual programming [96] and more conventional programming [80] can facilitate rapid development of visualizations. While these ease the burden of programming visualizations, they are limited in their ability to directly support artistic involvement in design. Tools like these are likely to fit nicely into the later (implementation) stages of our framework.

6.2 Background in Artistic Critique

Critiques are a primary teaching tool used in art and design education. They are also used outside academia in a variety of fields where visual design is important. The critique is a careful, critical, detailed group discussion evaluates a visual artifact. Discussion is oriented around specific aspects of the visual being studied and how well each “works” to support the goals of the piece. Good critiques involve detailed comments backed up with a basis for each evaluation. The comment, “I don’t like the colors...” must be followed with an explanation, “... because the use of primary colors dominates the visual field, leaving little room for perception of the subtleties that are really the most important concept in this piece” [94].

In our use of critique, scientists, visualization experts, and artists participate together. As is customary in illustration, art, and other visually oriented fields, our visualizations have a purpose behind them — in our case, it is science. Thus, the visual questions explored during critique serve to evaluate how well the visualization design functions in effectively representing the science. Scientists must be involved in making this evaluation. Together, the design team works toward posing the scientific problem as a visual problem, allowing visual concepts, such as color, texture, form, composition, metaphor, and narrative to be discussed as tools in service of the science.

6.3 Experiments in Visualization with Artists

When we began to collaborate on VR visualization problems, we learned very quickly two things. First, artists have visual insight that is clearly valuable to VR scientific visualization problems. Second, conveying this insight within VR and exploring the

visual space of VR is too difficult for artists not fluent in programming. Implementing visual ideas in VR can take weeks. Artists want to be able to create, critique, and refine designs in VR within hours or at most days. Thus, we have been motivated from the beginning of our collaborations to explore novel ways for artists to work with scientific visual ideas in VR.

In this section, we describe a series of four experiments in developing appropriate scientific design tools for artists. These were conducted over the course of several years. During this time, our research collaborations grew into an interdisciplinary class co-taught by professors of computer science, biomedical engineering, and evolutionary biology from Brown and of illustration from RISD. Several of the results pictured in this chapter have come from students that began in this class and have continued on to become involved in research. Each experiment led to an important refinement in artist-accessible visualization design tools; collectively, they led to insight that informs the Scientific Sketching methodology presented in section 6.4.

6.3.1 Experiment 1: CavePainting Visualizations

Experiment 1 explores using a free-form 3D modeling tool called CavePainting [54] for visualization design. CavePainting provides several styles of 3D “paint strokes” (ribbons, tubes, and other shapes) with which the artist draws in space using 3D tracker-based input. Like sketching or painting, the complexity of form in CavePainting comes directly from sweeping movements of the hand. Like other VR modeling systems based on 3D drawing-style interactions [78], CavePainting is easily understood and adopted by artists for artistic work. Thus, there are reasonable expectations that artists will quickly engage with this system, but it is unclear how well tools like this translate into letting one create scientific rather than artistic models.

Hypothesis and Methodology

The hypothesis is that, using CavePainting, artists will be able quickly to sketch out prototype visualizations that can then be critiqued directly in VR, eventually leading to visual insight and quick VR design iteration times.

Four artists involved in our collaborations were asked to create designs for one of the 3D fluid-flow visualization problems described below. Some of these were




Figure 6.2: Experiment 1: A free-form 3D modeling tool, CavePainting, was used to sketch designs for visualizations directly in VR.

quick initial designs, others advanced to more refined states. In all cases, the process began with hand-drawn sketches on paper or by searching for inspiration in paintings and photography that exhibit patterns of fluid flow. In some cases, elements of this preparatory work were scanned in and imported into VR to be used as textures in the CavePainting program.

To guide the design task, two active scientific visualization research problems within our group were targeted. The first visualization scenario is examining blood flow through a branching coronary artery [86]. Scientists are studying the depositing of plaque on the arterial walls. To investigate this phenomenon they need to understand time-varying pulsatile blood flow in various conditions. Variables such as velocity, vorticity, pressure, shear stress, and residence time are of importance, particularly near the arterial walls. The visual challenge is designing a visualization that highlights local relationships among these variables while preserving a global sense of the time varying flow.

The second visualization problem is investigating airflow around a bat's wings during flight [91]. Scientists are studying the evolution of flight in bats and its potential implications for future unmanned aircraft design. The 3D complexity of this problem is astounding, as bats have as many degrees of freedom in wing movement as the human hand, and their flexible wing membrane changes shape drastically during flight. The challenging visual problem is depicting the complex geometry of the bat's anatomy along with the flow detail that may hold keys to understanding the formation of lift.

Results

Images from the work of two artists are shown in Figure 6.2. In the left image, an artist is working with the CavePainting system on a design for the bat problem. Her work is inspired by the Miró painting, *The Gold of the Azure* [32], which contains some interesting flow-like visual structures. The right images are snapshots from within VR. The viewer is standing inside a scaled up version of the artery model. Two different representations for flow data have been sketched in 3D using CavePainting.

Experiment 1 Conclusions

While we confirmed that artists were able to quickly adopt and work with CavePainting, this experiment helped us to establish two key limitations of the CavePainting-based approach to visualization design. First, a lack of animated views connected to data makes it difficult to evaluate designs based on flowing icons or glyphs, as in the proposed artery visualizations. Second, a lack of control over the form drawn using CavePainting makes it hard to create illustrations that look scientific. While the loose quality of CavePainted designs is exciting for artistic purposes (it makes possible a hand-crafted aesthetic that is rare in computer graphics), it is inappropriate for depicting scientific subjects that demand precision. Follow-up experiments and tool refinements will address these two issues.



Figure 6.3: Artists designed several data-driven flow glyphs for use with the artery problem. This one, inspired by the natural forms of sea creatures, changes shape in response to both velocity and pressure as it moves through the flow.

6.3.2 Experiment 2: Data-Driven Glyph Sketching

In experiment 2, a new software system was developed to incorporate a tighter connection to the underlying scientific data within the artistic design tool. The scope of visualizations was limited to the design of glyph-based flow visualizations like the ones proposed in experiment 1 for the artery problem. The goal of this experiment was to address directly the animation limitation discovered in experiment 1.

Hypothesis and Methodology

The hypothesis is that animating artists' drawings in response to real data will improve our ability to evaluate the success of glyph-based visual techniques for 3D flow visualization.

A designer works with the revised software by drawing a legend describing how a glyph should change in response to data. Based on the specification the legend provides, the system automatically produces a visualization of animated data-driven glyphs moving through a flow volume. The artist begins by drawing several instances of a 3D glyph using CavePainting. These are then associated with specific data values. For example, to indicate a change in the geometry of the glyph in response to the variable *flow speed*, the artist would draw what the glyph should look like at low speeds and link this drawing to the slow end of the *flow speed* legend. Then, she would draw a second representation for high speeds and link it to the fast end of the legend. The system computes a 3D morph between the two representations that is




Figure 6.4: Experiment 2: An animated, data-driven visualization sketch can be automatically created from a few key-frame glyphs drawn in 3D and attached to a data legend.

used to continuously vary the form of the glyph as it moves within the flow volume. Multivariate glyphs are constructed by adding additional legends to the specification. For example, additional drawings could be used to make the color or texture of the glyph change in response to *flow pressure*. Once the specification is complete, a seeding algorithm for particle-based visualization of time-varying flows [86] is used to distribute and advect a set of glyphs through a visualization of the flow dataset. This visualization is then critiqued directly in VR.

Working together with one illustration student and several of the designs posed by her peers, we evaluated this system by creating test cases targeted at the artery problem. One driving example of a data-driven glyph that we tried to capture using the system is shown in Figure 6.3. In this design proposed (here as initial design sketches) by illustration students in our class, an organic squid-like glyph changes shape in response to velocity and pressure as it moves through the flow.

Results

The complexity of the motion available through this method is impressive and hard to capture in any other way. A view of the artist-created legend and resulting data-driven visualization for the squid case is shown in Figure 6.4. While it is difficult to convey in the static images here, when seen in VR, designs created using this method clearly capture a sense of flow that is entirely lacking in the results of experiment 1.

Despite this advantage, we found significant limitations in this approach. Establishing correct correspondences between multiple drawings of a glyph composed of an unlimited number of 3D “brushstrokes” is a challenging algorithmic problem. In order to simplify the problem, our implementation required artists to keep the glyphs used in the system very simple, severely limiting their power to use CavePainting as intended. A good CavePainted glyph would be suggested, as in traditional painting, by many tiny brushstrokes oriented in space, but the geometries we could explore were of the form of those in Figure 6.4, relatively simple geometric forms.

Experiment 2 Conclusions

While the tight connection to the data gave artists a powerful tool, it also forced them to work within a framework that often limited their ability to convey the new insights and sophisticated visual thinking that motivated our collaborations. We also discovered that these relatively large, multivariate glyphs are less appropriate for flow visualization than we had originally thought. Dense, simple particles yield a much more understandable representation of flow patterns in VR. In the following experiments, alternative methods for connecting designs with underlying scientific data were explored.

6.3.3 Experiment 3: Appropriate Artistic Control for Science

Having gained some insight into designing time-varying visualizations, we return to the other limitation identified in experiment 1, the inability of artists to control 3D drawing tools with the precision required for science. In addressing this issue, we developed the far more precise, haptic-aided 3D drawing tool called Drawing

on Air, that is described in detail in chapter 3. Drawing on Air increases control while maintaining the immediate, exploratory qualities of modeling tools based on 3D drawing. Thus, it remains accessible to artists with minimal computer training. This experiment explores the use of this new, more controllable tool for depicting complex scientific problems.

Hypothesis and Methodology

The hypothesis is that the improvement in precision seen with Drawing on Air is significant enough to let artists combine the strengths of modeling tools based on a 3D drawing paradigm with the precision required to address difficult visual subjects in science. If this requirement is met, these tools should be useful for generating finished 3D scientific and medical *illustrations* as well as intermediate visualization designs.

Two artists within our collaborative group worked closely with us and with the evolutionary biologist leading the bat-flight project to create 3D illustrations of bats posed in flight using the Drawing on Air tools. In some of the illustrations, such as the top two in Figure 6.5, drawings were created “on top of” experimentally collected bat-flight data. The blue sphere markers at the joint positions in these images were positioned inside the 3D drawing system to correspond with motion-capture data collected from flying bats in a wind tunnel. The artists used these guides to create anatomical illustrations within the reference frames provided by the experimental data. Together with the evolutionary biologist guiding this project, we hypothesize this style of 3D illustration will be far superior to traditional 2D representations for several purposes, including studying which muscles are likely to be active at particular points in the wing-beat cycle.

Results

In general, significant improvements in the 3D drawing precision using this new technique were observed compared to previous results created with tools in the spirit of CavePainting. The additional control clearly has significant ramifications for depicting scientific subjects with precision. For example, the smooth curves of the wing

bones in Figure 6.5 would be impossible to draw with CavePainting or similar free-hand tools.

Experiment 3 Conclusions

With tools of this level of artistic control and expression, artists can combine the benefits of hand-drawn 3D modeling with the control needed to address scientific subjects. Even for intermediate design, achieving this level of control is important: without it, the science is confused rather than clarified and visual critiques are less accurate.




Figure 6.5: Experiment 3: The additional control provided by the Drawing on Air haptic-aided drawing interface lets artists address scientific subjects with the necessary level of precision.

6.3.4 Experiment 4: Time-Varying, Interactive Prototyping

Experiment 4 explores the issue of designing for time-varying, interactive scenarios that are a key part of VR visualization.

Hypothesis and Methodology

The hypothesis is that a substantial number of the time-varying and interactive properties of VR visualizations can be captured through two extensions of the tool: 1. support for multiple drawing frames that can be cycled through, as in stop-motion animation, and 2. support for Wizard-of-Oz-style interaction prototyping.

To provide a role for time-varying data, drawing on top of motion capture markers is extended to support multiple frames of data. Wizard-of-Oz prototyping is implemented by adding drawing layers that can be turned on and off using keys on the keyboard. Designs are created in such a way that during critique a “wizard” operates the keyboard to cycle through frames of an animated drawing and toggle drawing layers on/off to simulate visual changes in the scenario that result from user interaction.

The resulting system was incorporated into the toolset used to teach the class Virtual Reality Design for Science mentioned in the introduction to this section. Results from this experiment come from the various final projects that art and computer science students prepared as part of that class. Each design went through several sessions of artistic and scientific critique and revision, initially in paper form and later in multiple VR evaluation sessions.

Results

Fourteen designs for VR visualizations were created as part of this experiment; we describe one of these in detail here. Figure 6.6 shows results from several stages of the design of one successful project, “Cutting Mirrors for 3D Flow Visualization”, that resulted from this experiment. These cutting mirrors are an interactive visualization tool for comparing multiple bat-flight datasets. It is hypothesized that such a tool will be useful for comparing the flight of different species of bats as well as the same species performing different maneuvers. Comparison of these situations should provide scientific insight on flight mechanics and evolution.

The leftmost image in Figure 6.6 is an initial sketch from a design storyboard for the project. During critique, this sketch sparked a discussion about the visual effectiveness of the technique as presented. It was agreed that several changes would improve the readability of the data, and these were incorporated into the design before it was investigated in 3D. One of these changes was a move from displaying both datasets within a shared coordinate system to displaying the flight data in separate but correlated axes.

In the next stage of design, (center image of Figure 6.6), a 3D drawing of the technique was created. Multiple drawing layers were used in conjunction with a Wizard-of-Oz approach to investigate the interactive scenario of placing a cutting mirror within the visualization and moving it around near the wing in order to compare a portion of the wing as seen during the upstroke to the same portion of the wing during a downstroke. Again, discussion at the critique identified a visual problem, also relating to the coordinates of the two visualization spaces.

A final 3D design (rightmost image of Figure 6.6) was again presented for critique. At this point, the design had been significantly refined from a visual and interactive standpoint in a matter of weeks. An interesting observation was made during critique: since the coordinate space in which the data was visualized had evolved from a simple regular grid to a space that could be cut with a mirror and then opened up like a book for close visual inspection, several probable assumptions in the strategies for programming the original design were now no longer valid. Had the technique been implemented after the first set of sketches, incorporating the later refinements would have involved changing major assumptions and likely forced rewriting large portions of the code.

Experiment 4 Conclusions

By drawing on top of motion-capture data, artists created stop-motion animations that are tied to real scientific data. However, unlike the very tight data constraints in experiment 2, the strategy used in this experiment allows considerable artistic freedom in the design. We conclude that linking design sketches to data can be an important aid in capturing a more realistic view of the goal visualization, but the constraints imposed by the data must be balanced with the goal of artistic freedom. Gradual

introduction of data into the design is likely to encourage a productive balance.

The ability to cycle through drawing layers to prototype user interactions in a Wizard-of-Oz style was also explored in this experiment. Prototyping these interactions allowed important design exploration that could not be captured by our previous tools.




Figure 6.6: Experiment 4: The addition of Wizard-of-Oz interaction prototyping techniques and key-framed animation controls enables artists to sketch the *experience* of VR as well as the visual artifacts. A design for a cutting-mirror technique for exploring multiple bat-flight datasets was proposed first on paper (left) and then refined during multiple VR sessions (middle and right) in which the interactive as well as visual characteristics were critiqued.

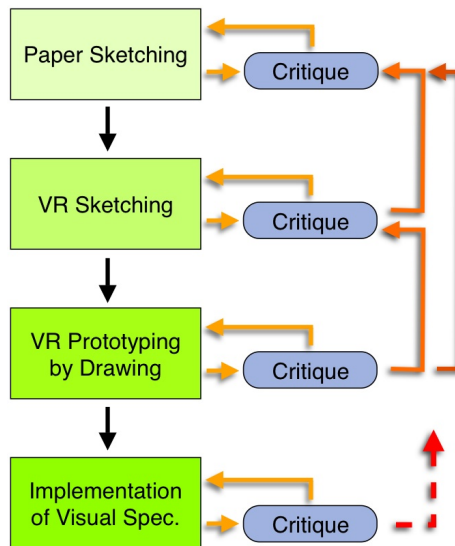


Figure 6.7: Overview of the Scientific Sketching methodology.

6.4 Scientific Sketching: A Collaborative VR Design Methodology

In this section, we present a formalization and refinement of the design process explored in the experiments of the previous section. As indicated by the orange arrows in Figure 6.7, progression through the stages of this process can involve iterative loops; however, we try to avoid the loop from a partially implemented idea to the earlier design stages because it is so costly.

Each of the stages described below involves both individual and collaborative work. Design sketching and programming are typically done in individual sessions, although sometimes collaboration is helpful here as well. The key collaborative process in the methodology is the critique, as described in section 6.2. Critiques serve as both a beginning and an end, since during critique designs are evaluated and future directions are also established. All the team members bring unique insights to this critical process, which is repeated multiple times within each design stage.

Design Stage	Goals	Output
Paper Sketching	Explore design space with few limitations. Understanding of the scientific problem.	Many sketches (low overhead) using traditional media.
<i>Keys for Transition:</i> Prep for VR design by preparing imagery to import as billboards. Prepare plan views of the spatial layout.		
VR Sketching	Transition to VR. Explore spatial and color relationships and stereo effects that can only be seen in VR.	3D designs that can be critiqued by all collaborators. Sketch-like in their quickness and low cost.
<i>Keys for Transition:</i> Incorporate motion via key frame animation. Draw on top of simple data representations.		
VR Prototyping by Drawing	Move much closer to an actual visualization. Critique and refine color, form, and metaphor at a low-level.	Design "drawings" rather than sketches. Wizard-of-Oz placeholders for interactive techniques and walkthroughs of scenarios.
<i>Keys for Transition:</i> Use elements of the hand-drawn designs as placeholders. Actively maintain collaborations.		
Implementation of Visual Specification	Leverage the now refined VR visual specification to create an effective visualization.	Iterative refinement leading to a fully data-driven visualization application.

Figure 6.8: Goals and deliverables for each stage of design.

6.4.1 Paper Sketching

Almost all successful design processes begin with quick sketching on paper. Paper is quick, easy, accessible, and disposable, so it is easy to explore many possibilities quickly and engage in visual thinking.

Role of Artists: The artist's role is to present many visual ideas. Artists should ask questions to learn enough about the science in order to pose the problem visually.

Role of Visualization Experts: Visualization/computer experts play a key role in facilitating discussion between the artists and scientists. Sketching ideas is also important. In this early stage, sketching good visual ideas without regard for the difficulty of implementing the ideas is often useful.

Role of Scientists: The scientists' chief role is to explain the scientific problem and data sufficiently for others to understand at a level of potential cause-and-effect relationships and the relative importance of variables. If one variable can be identified

as the first thing to look for in a visualization, that information is important to convey. Landmarks in the data are also useful to identify: for example, in brain visualization, the ventricles often provide a landmark for understanding spatial orientation.

6.4.2 VR Sketching

After several paper designs have been proposed and refined, it is important to begin to evaluate the ideas within VR. Initial VR sketches should be almost as disposable, and almost as quick, as paper sketches. Easing the transition to VR is an important goal at this stage. Importing into the VR design tool scanned initial design sketches and photographs of textures made with traditional media can be a great help in beginning a design.

Role of Artists: While working in VR to “sketch” 3D versions of the designs from the previous stage, the artist should work as she would with any new brush and color palette. Learn the medium through experimentation via drawing. How do colors interact, and how does computer-generated lighting on the forms interact with artist-specified color? What stands out? What fades to the background?

Role of Visualization Experts: Initial 3D design sketches should also be created. One issue to explore in particular is visual simplification. Often designs that pose significant implementation challenges can be simplified visually making them easier to implement, and this often increases visual effectiveness as well.

Role of Scientists: Speak with reference to specific visuals during critiques. How would the data be interpreted if presented exactly as they are sketched now? How well do the visual forms reflect underlying scientific concepts? Discussion of visual metaphor and narrative is helpful.

6.4.3 VR Prototyping: Drawing, Wizardry, and Connections to Data

The switch in terminology from “sketching” to “drawing” is deliberate in the name of this stage of the process. Many of the same 3D drawing tools are used, but the drawing becomes deliberate and exact, rather than quick and sketchy. Prototypes, especially those including interactive scenarios, may be created over several days,

rather than the hours needed for paper and VR sketching. Only the most promising one or two designs should advance to this level. This is the stage at which Wizard-of-Oz interaction techniques should be explored. Additionally, some effort should be made to establish connections to the underlying data.

Role of Artists: The artist should focus on refined drawings. Picking a very specific hypothesis and then creating a view that is useful for investigating it is a focused way to proceed.

Role of Visualization Experts: While the visualization/computer expert should do some of the same drawing and visual refinement as the artists, a need for limited programming may also emerge. In experiment 3, importing the motion-capture data into the drawing system required less than three hours of programming and made a tremendous difference in the team's ability to critique time-varying designs with confidence. Another possibility for simple programming is adding controls to features the artists have drawn. If a feature is intended to rotate in response to data, simple script-like additions to the program can be used to mock up these visual effects for critique. Making a design element rotate may involve an hour of programming, while making it rotate in response to vorticity values in a pulsating, time-varying flow may take days or weeks. At this stage, the non-data-driven version is likely to be nearly as valuable as the data-driven version in determining how well the design functions visually.

Role of Scientists: The team is converging upon a specification for a fully implemented visualization. It is imperative to address whether the necessary quantities are visible in this visualization in order to investigate the driving scientific questions.

6.4.4 Implementation of the Resulting Visual Specification

Most visualizations, especially exploratory ones, target a final result of a fully data-driven visual display. At some point, a programmer must take the lead in building this type of application, but the transition to implementation can be difficult to navigate.

Role of Artists: The artist must advocate for the visual decisions made in the earlier design stages. It is easy for artists to be left out of the process as the programmer assumes responsibility for what actually ends up displayed on the screen. Implementing the design is bound to necessitate some design changes. The artist

needs to stay involved in discussion and redesign through smaller-scale repetition of earlier drawing-based design stages.

Role of Visualization Experts: The visualization/computer experts lead the implementation of the visual specification devised in previous stages. A conscious effort must be made to maintain collaborations. One of the best ways to do this with respect to the artists is to build the implementation on top of the prototypes established in previous stages. This way, hand-drawn elements of the design can live on in the “final” presentation as placeholders or annotations. The artists should be encouraged to continue to draw on top of the latest state of the visualization to continue to refine the specification for the yet-to-be-implemented portions.

Role of Scientists: Scientists also must take care to stay involved in the process during implementation. An important role is to help determine intermediate goals and set priorities for features to implement. What is the next hand-drawn placeholder that should be replaced with a data-driven visual element? Estimates of the relative difficulty of implementing features should be used collectively to determine the most important next steps to take towards the ultimate goal, scientific insight.

6.5 Analysis

Here, we provide a theoretical analysis of the utility of our collaborative design methodology along with a characterization of design tools, including four axes by which tools can be measured.

6.5.1 The Value of Collaborative Design

In this section, we provide a formal description of collaborative design for visualization and explain how Scientific Sketching works to optimize this process. We were inspired by van Wijk’s recent work “The Value of Visualization” [97], including his call for further study of methodological issues. We adopt his model and notation for the visualization process in this analysis.

Ultimately, the goal of visualization is insight, which van Wijk describes as an increase in knowledge, dK/dt , given a time varying image I seen by the user. The knowledge conveyed is based on specific properties of perception and cognition P of

the user and is a function of I and the current knowledge K of the user:

$$\frac{dK}{dt} = P(I, K). \quad (6.1)$$

The visualization process V generates image I as a function of data D and a specification S :

$$I(t) = V(D, S, t). \quad (6.2)$$

In van Wijk's formulation, S is composed of a hardware specification, algorithms to apply, and specific parameters. We introduce a more refined notion of S that provides a link to the design process used to generate a visualization *method*, M .

In our formulation, S is a function of a method M that results from a design process, such as that described in this chapter. M contains a specification for a visual language that maps data to visual form and a specification for human-computer interactions that map user input to changes in data or visuals. The full specification S is a function of M , a specific hardware and software implementation, and state information as determined by user interactions.

The focus of the work in this chapter, then, is the incremental design process for the visualization method, dM/dt , which can be modeled as a function F :

$$\frac{dM}{dt} = F(E, A, T), \quad (6.3)$$

where E is the expertise of the design team and A is the level of accuracy with which the design will be expressed. Intuitively, A is a measure of how "sketchy" the design will be in this iteration, ranging from a back-of-the-napkin sketch to significant programmatic implementation. Implicit in the current method under design M are two additional variables: 1. The critiqueability of this design, and 2. the current level of design refinement. In general, as A increases, the design process results in a M with increased critiqueability. That is, the more representative of an actual implementation the design is, the more accurate we can be in our evaluation of it. The current level of refinement also affects the process. As we proceed with more and more design refinements, the ability to advance the design significantly through additional iterations decreases.

Given a function U that lets us measure the utility of a design, we expect that as the process proceeds, future designs are better and no worse than previous ones:

$$U(M(t + \tau)) \geq U(M(t)). \quad (6.4)$$

If we do derive a design that is actually worse, we can always throw it away during the evaluation stage of the design loop and keep the previous design.

Therefore, we expect the utility of a new design at time $t + \tau$ to be that of the previous design plus the utility of the particular design process taken:

$$U(M(t + \tau)) = U(M(t)) + U(F(A, E, M, t)). \quad (6.5)$$

From this formulation, we derive two clear goals for the design process. First, we want to minimize the duration of each design iteration τ because we expect significant new insight to come from each iteration. The second goal, maximizing $U(F)$, must be balanced with the first. We want each iteration to be quick and also take a significant step forward.

In these terms, Scientific Sketching is a valuable design process because it takes very clear steps to minimize τ and maximize $U(F)$. To minimize τ we avoid using programming as a design tool because it is so costly and time consuming. To maximize $U(F)$ we 1. increase the accuracy A with which we represent our designs (and consequently the critique-ability of resulting design M) by working with appropriate design tools, and 2. we maximize the expertise of our design team E by working with both scientific experts and experts in visual design, illustration, and art.

In addition to critiquing/evaluating the success $U(M)$ of a design, visual experts tend to be very good at telling us *why* M works as it does and also identifying the best trajectory M' to take to improve M . We believe that tools and processes that let us exploit this directional information are likely to have great scientific impact.

6.5.2 Characterization of Design Tools

In this section, we present four axes along which design tools for visualization may be measured along with an analysis of the specific tools we developed for collaborative VR design. Our work explores an important region in the high-dimensional space of tools that can be described according to the metrics presented here. Additional exploration within this space is certainly warranted. We hope this characterization serves as a guide to future efforts.

Axes of Tool Utility

Axis 1 — Accuracy of Visual Critique/Evaluation: The first axis assesses the design team’s ability to accurately critique and evaluate designs created using the tool. This is often equivalently described as the ability of the tool to accurately depict the final visualization. Important in this assessment are: 1. whether designs can be viewed in the target visualization medium, 2. whether there is support for animated views, 3. whether user interactions may be evaluated, 4. whether the data drives the visual display, and 5. expressiveness: whether the tool’s modeling primitives can accurately capture existing and new designs.

Axis 2 — Cost: The second axis is cost. This is assessed in terms of the amount of effort that goes into the design in time, people, and other resources. Related to this is the cost of abandoning a design created with the tool. As with any optimization process, we want to find a global rather than a local minimum. Design with some tools requires such a high cost that we are unwilling to abandon a design once begun. If we start design with these tools near a global minimum, there is no problem, but if we start near a local minimum, we may miss an altogether better design.

Axis 3 — Accessibility: Of chief importance in our discussion of tools is their accessibility to different users. Clearly for novel systems, the tool of programming, which is in general not accessible by our collaborators, must be used to some extent. However, as we demonstrate, other tools that are much more accessible can play an important role in increasing the value of our collaborators’ contributions.

Axis 4 — Learning by Doing: Like sketching, tools such as VR-based 3D drawing techniques lend themselves to exploration. Because they operate within the target visualization medium, these tools have also the advantage that designers can learn about the medium as they work. With each new design, additional exposure to the medium helps artists gauge VR’s possibilities and constraints. These tools support a clear path to becoming an expert in the medium.

Tool Analysis

Before this work, tools based on 3D drawing approaches scored well on the axes of Cost, Accessibility, and Learning by Doing but, as illustrated in experiment 1, they were inappropriate for investigating scientific subjects, leading to poor performance

on our final evaluation axis, Accuracy of Visual Critique/Evaluation.

Based on insight from our experiments and subsequent design tool development, we have now reached a level where design via 3D drawing is practical for science. Tool development has greatly enhanced our ability to represent and critique complex subjects in VR. This makes possible a collaborative design methodology that leverages a great deal of artistic insight and skill.

Common design tool alternatives, such as quick programming, scripting, or visual programming tools, should also be evaluated by these metrics. In general, these do well in terms depicting and evaluating a design, but they have a high cost and are very limited in their accessibility to artists. With these tools, the involvement of artists in a collaborative design process is limited to critique and visual development outside of the target medium. As we adopt more artist-accessible design tools, artists can play a much more important role in addressing visual problems posed by science.

6.6 Discussion

The benefits, limitations, and possible future extensions of Scientific Sketching are described below.

6.6.1 Toward Renaissance Visualization Designers

A byproduct of collaborative visualization design that pulls in artists, illustrators, designers, and scientists is that everyone in the group learns a bit about the other fields. By learning about process, technique, and language in these other fields, we become better suited for the next collaboration and bring additional insight to the next noncollaborative project as well. Establishing and maintaining these collaborations takes time and energy and is clearly not appropriate for every visualization problem; however, our belief is that the benefits extend well beyond each particular collaborative design process.

Some visualization researchers have expressed reservations about this approach with comments such as, “I got into this field because I enjoy the visual and artistic parts of the research — why would I want to hand all these over to an artist?” This is the wrong way to conceive of this collaborative process. The goal is not to hand over

all visual decisions to artists, but rather, to recognize and spotlight the importance of visual design. If we truly want to engage in research and debate at a visual level, then we cannot ignore the importance of practitioners in these fields. We need to find ways to draw these experts into our processes. Ultimately, this has the potential to make our work in computer-based techniques even more enjoyable in terms of visual problem solving.

6.6.2 Limitations and Future Directions

The most challenging stage of the collaborative design process is the transition from design to implementation. One way to maintain productive collaborations at this stage is to structure implementation so that the application under development can be linked with the library in which the VR designs were created, thus letting the design refine and grow along with the implementation. Design elements can remain inside the implementation as placeholders, and the VR design tools can be used to annotate or augment the current status of the implementation. To date, our experience with this approach is limited and we have no use-based conclusions to relay. Validation of this approach and investigation of alternatives is a valuable area for future research.

6.7 Summary and Conclusions

In this chapter, we present four experiments in developing VR visualization design tools and propose a new methodology for collaborative visualization design called Scientific Sketching. Through our experiments, we learn that effective artist-accessible design tools must support sufficient artistic control, animated views, and some connection to underlying scientific data. We propose a refined methodology for collaborative design based on these insights and tools for prototyping visualizations by sketching in VR. In analyzing this methodology, we build on van Wijk's definition of the visualization process to provide a link to design. We describe optimal design as striking a balance between minimizing time between design-cycle iterations and maximizing the effectiveness of each iteration. Scientific Sketching very explicitly targets these two goals and effectively leverages the visual insight of our collaborators. In looking toward the future, we believe artists, illustrators, and designers can make important

contributions to science, and refining the methodologies and tools to bring these visual experts into our scientific processes is critical to future artist-scientist collaborations.

Chapter 7

Critique-Based Evaluation of Art Applications

In this chapter, we present results and conclusions from artistic investigations with the CavePainting and Drawing on Air free-form modeling techniques. The results in this chapter reflect the work of more than ten artists in both collaborative and individual settings. Results have typically been very well received in the art community and pieces have been displayed in several juried exhibitions and academic conferences on digital art. Work with our tools has also been included in several RISD illustration courses, including independent study and group study six-week courses and individual student projects contributing to senior portfolios.

In addition to these contributions in art, directed art projects have led to several computer science insights: 1. a refined framing of the control problem that is the focus of this dissertation; 2. conclusions about appropriate use of the Cave VR environment, including cognitive and emotional responses to virtual stimuli; and 3. conclusions about effective new strategies for depicting complex 3D forms in head-tracked, stereo VR displays. These insights prompted several follow-up scientific investigations, including a study of changes in perception of 3D form in VR under different visual background conditions [49] and an exploration of body-scale 3D drawing as an interaction metaphor for scientific visualization applications [86]. Thus, in our experience, there is a tight connection between directed, serious art practice and scientific discovery.

Our most thorough guided critique-based evaluation is the result of a several-month collaborative study between Fritz Drury, professor of illustration at RISD, and myself. We begin by presenting hypotheses, methodology, results, and analysis of this study, and then describe results and analysis from additional artistic investigations.

The evaluations in this chapter are based in artistic critique, a primary teaching tool in art and design schools. Critique is an important evaluation tool in many visually oriented disciplines [94]. In computer science, use of critique for evaluation of visual work has been gaining momentum, largely as a result of work in our lab at Brown [48, 55]. Other researchers have now also adopted the strategy, notably Colin Ware’s group at the University of New Hampshire. One important benefit of this style of evaluation over more quantitative analysis is that visual experts, such as artists, can often tell us *why* a visual technique works as it does and also *how* to make it better. We generally do not get this information from standard quantitative analysis. Several of the conclusions drawn in this chapter come from artistic insight in answering these “how” and “why” questions. Further background on the goals and process of artistic critique, along with considerations in applying it to scientific domains, can be found in section 6.2.

7.1 Introduction to Evaluations in Artistic Anatomical Illustration

Our most in-depth critique of an application in art was a several-month long investigation involving more than ten evaluative critique sessions in both the Cave and fishtank VR environments. Having a strong art background myself, I was a key participant in these investigations. I created a series of works targeting anatomical subjects in art, and then Fritz Drury led critiques of each work in VR. Our methodology is described in more detail in section 7.3.

In focusing on a specific and difficult subject for a prolonged period of time, this experiment pursues a standard method of inquiry. Our goal is to achieve new visual insights about appropriate use of free-form modeling in depicting challenging subjects. Among our concerns are: effective means of suggesting form with 3D strokes rather than with the full surface meshes commonly found in computer graphics, use of color,

and effectiveness of different drawing interfaces and viewing environments for this style of depiction.

7.2 Hypothesis

Our primary hypothesis is that careful, directed artistic study of a challenging subject will yield insights on effective modes of visual depiction for free-form modeling based on 3D drawing. Further, we believe such insight is likely to be of value in evaluating current tools and refining appropriate directions for future research.

7.3 Methodology

The methodology is modeled on the course Artistic Anatomy taught in RISD's Illustration Department; the key differences are the tools used to produce the illustrations and that these tools work in virtual 3D spaces rather than on paper with pencil and charcoal.

In keeping with the structure of the RISD course, Drury took the lead in our investigation in assigning different subjects for each stage of the work. We targeted the torso, arms, hands, full figures, and faces in succession. One of the important teaching tools in Drury's traditional anatomy course is working on two-part illustration problems. The first part is a descriptive illustration that presents the bones, muscles, and tendons responsible for placing a figure in a particular pose. The second part is a dynamic illustration of the body in that pose. An example of one of these drawings was shown in Figure 1.2. In these dynamic illustrations, we should see evidence of the underlying human anatomy. When this evidence is clear, our perception of the action is also clear. We adopted this same structure for much of our investigation because it directs the study and provides a clear path for critique-based evaluation.

Biweekly critique in the Cave and fishtank environments served to evaluate the work and establish an appropriate next direction of inquiry. A final critique at the end of the study, in which each work was revisited, served to wrap individual insights into more general principles of depiction for the medium and thoughts on useful future directions for our tools.




Figure 7.1: A descriptive depiction of a torso.

7.4 Results and Individual Critique Analysis

The following sections present results and a summary of the evaluation that came out of the critique for each work. In section 7.5 we draw some overarching conclusions from the individual analyses presented here.

7.4.1 Torsos: Descriptive and Dynamic

Descriptive Torso

Figure 7.1: 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 to improve their drawing skill, 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 are




Figure 7.2: A dynamic depiction of a torso.

intended to function as shading information and are applied 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: do these marks represent form or lighting? They seem to do a poor job of indicating lighting. One of our first observations from this work follows from this and the next critique: defining the appropriate role for use of color in hand-drawn computer form is difficult because artist-assigned color is mixed with computer graphics-assigned color from lighting equations. The ribs are drawn with single ribbon strokes and are too flat compared to the pelvis. Drury suggests using a tube form or several ribbons close together to indicate more volume in the ribs.

Dynamic Torso

Figure 7.2: This is a loose, gestural work that suggests a sense of twisting movement. It shows somewhat less of the body than what Drury normally demands for a dynamic pose, and that motivates the decision to investigate other parts of the body next. Torsos have a beautiful and magical aspect, and the challenge is to get the same sense of gesture and movement with an arm or with a whole body. An elegant and provocative part 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 in stereo viewing for depicting forms with minimal use of line, an area worth exploring in our visualization techniques where visual space is always at a premium.

The holes in the form (the ability to see through parts of the surface suggested by the hand-drawn ribbons) are of course one of the most distinguishing characteristics of this model. Drury agrees that this technique works well: it is interesting (not necessarily a problem, as we tend to assume in computer graphics work) to have these areas blank.

The proportion in the model is exceptional. The use of color is the weakest aspect. Drury suggested relying much more heavily upon the built-in computer graphics lighting rather than specifying it with color changes in the form. Color may be useful here in suggesting lighting, but probably should have a more limited role than it does in this piece. Tying the marks drawn to specific anatomical structures is a main focus




Figure 7.3: A descriptive depiction of the scapula, shoulder, and arm bones.

of the piece and Drury notes that some of the marks are not tied closely enough. Particularly when working with this semi-transparent style, he suggests looking for features that aid in depiction, like the muscles on the side of the torso that pull the eye around the figure.

7.4.2 Arms: Descriptive and Dynamic

Descriptive Shoulder and Arm

Figure 7.3: The scapula, the main accomplishment in this model, was created by zooming in and defining it with dozens of small marks. When we step back and look at it at a normal scale, it has the feel of a solid form. This is a more sculptural style than many of the other works. It functions well in evoking a real sense of a solid form, much more than the descriptive torso model above. Tube marks form the




Figure 7.4: Sketches of arms in action.

rough shape of the bones, and then ribbons 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 in this drawing is not as good as the dynamic torso and the form of the clavicle is drastically oversimplified.

Sketches of Arms in Action

Figure 7.4: These sketches reveal a definite struggle. The pose is problematic in each and there are serious artistic questions about how much of the form to try to represent through the ribbon surface patches. Drury's assessment is that selecting the right lines to draw is best thought of as a process of first defining the correct pose, then clarifying the important parts of that pose.

There is a nice parallel in these action poses to our visualization work. The image on the right, for example, depicts 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? The first step is to get the pose right. This fisherman is far too upright. He is not struggling against the weight of the net and the rope. Instead, he looks as if 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




Figure 7.5: Hands pulling in a fishing net.

and in the muscles of his arms and torso that struggle to haul his catch. This level of detailed information is what we try to convey in artistic illustration. In pushing the limits of this frontier with computer tools, we learn visual strategies that can be applied to similar visual problems in science.

7.4.3 Hands: Descriptive and Dynamic

Hands Pulling in a Fishing Net

Figure 7.5: This work is a second exploration into a more sculpted aesthetic, building from the descriptive shoulder and arm presented above. Drury’s comments on the posing of the hands were that the pose of the hand about to grab the rope is fairly good, but 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 and pulling it with




Figure 7.6: Hand holding a baseball.




Figure 7.7: Hands wringing out a towel.




Figure 7.8: Anatomical drawing of a hand.

force. One interesting point is that drawing out the skeletal structure of the fingers was easy. For anatomical forms, a hand-drawn underlying skeletal system might be useful as an artistic control (as in traditional animation packages) for situations in which manipulations to the skeleton affect the surrounding surfaces.

Hand Holding a Baseball

Figure 7.6: This work is a final attempt at a sculptural style of depiction. The entire form is specified with ribbon-like minisurfaces. The knuckles are successful here, but again this style of working is problematic because it mixes metaphors: we are using something like a line-drawing tool and trying to create sculpture. Suggesting form with characteristic, controlled lines may be more appropriate with this tool than these more sculptural approaches.

Hands Wringing a Towel

Figure 7.7: This model is the most difficult to understand without the aid of stereo or animated viewing. It takes a completely different approach from the previous work, and is based on the following realizations about working with this style in VR. When seen in stereo, a suggestion as small as a dot can suffice to specify the location of a key anatomical feature (the elbow for example), because with stereo we can easily tell the exact spatial location of that dot. In this work, the hands are suggested with as little use of line as possible. To provide an intriguing contrast, the towel inside the hands is specified in a far more sculptural style. Thus, we see through the hands to the towel that they are deforming. The idea is an interesting one, but as Drury notes, the implementation is less successful than the idea, in part because the representation for the towel is clumsy and jagged completely dominates the visual field.

Anatomy of a Hand

Figure 7.8: Several interesting ideas are at play in this work. Again, there is an attempt at economy of line, and again, the piece works well from some angles, not as well from others. Transparency is used in this model to define the muscles, and the mix of transparent forms with solid ribbon forms works well in the muscle areas. The insertion points for tendons and muscles could be more clearly specified. The




Figure 7.9: Seated woman.




Figure 7.10: Woman undressing.

translation from stereo to mono viewing drastically degrades the interpretation of the model. To make 3D illustrations useful in 2D projections on paper, we need to follow more of the basic guidelines for suggesting depth that we learn from study of traditional drawing. For example, the color of the marks should be adjusted with the respect to the background color to help clarify depth relationships. A fruitful avenue for future research may be to combine form created with 3D drawing with NPR rendering schemes.

7.4.4 Figures: Gesture and Economy of Line

Seated Woman

Figure 7.9: Drury notes that this model is the most successful example of minimal use of line: it looks good from many viewpoints and the model is depicted with only a few lines. Unfortunately, this is difficult to see outside VR because the lines are so thin. Drury points out that the “correct” lines were drawn to suggest the form, but interestingly they weren’t always the same. For example, a different representative feature is used to depict each leg: on one it is the outline of a quad muscle, on the other the kneecap. This style works 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, which pose a difficult problem for this style of representation because they protrude from the form. A contour line along the center of the form 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, omitting this type of indication from the picture, as is done here, destroys the impression of 3D form from several important viewpoints. In this work, we have converged upon a successful style of depiction and also identified a key limitation.

The obvious tool extension 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 to get a representation that truly works from all angles for that last ten percent of the form, we may need to be able to draw the model differently depending on the orientation of the viewer.

Gesture: Taking Off a Robe

Figure 7.10: This model illustrates one of the advantages of working in the Cave: it is a 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 inspired by Picasso’s light pen drawings and achieved via real time tone-mapping. The glowing rendering effect is intriguing, but the challenge is to control it while creating the form. This type of 2D effect may help our projected images appear less flat and comparatively boring when printed on paper.



Figure 7.11: A Swahili bride wearing a green veil.

7.4.5 Faces

Swahili Bride

Figure 7.11: This work is a face that has correct 3D proportion and thus looks good from most angles. Most of the lines are in the right place spatially, but because the use of line is fairly minimal, we find that the lines are not necessarily the best ones to draw from all angles. The nose is good, the eyes fairly good. The simplified color palette works well for this model.

Bearded Man

Figure 7.12: This model has a well defined nose. The guideline that is part of the Drawing on Air tape-mode technique was useful in creating this form because so many of the curves are drawn by tracing along next to existing curves. Drury liked the 3D placement of the brow of the head: when looking from above we get a sense



Figure 7.12: A bearded man.

that it jets out of the form. The right eye is off in its bilateral symmetry. Again the sparseness of representation, for example the indication of the hair, worked well.

7.5 Conclusions from Anatomical Studies

Several conclusions arise from critique of the results presented above.

7.5.1 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 triangle-mesh surfaces. 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 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 and effective.

This style of depiction may hold promise for future scientific investigations as well. For example, in creating an illustration of neural fiber tracts in the brain, one of the challenges is providing sufficient contextual information in order for scientists to maintain spatial bearings while examining data. Since this style of depiction is based on suggestions of form, it could be particularly useful for describing features like the skull or eyes in ways that provide contextual information but do not dominate the visual field, allowing the focus to be on the dataset presented within the surrounding context.

7.5.2 Economy of Line and Occlusion

Embracing the idea of suggesting form with few lines as described above, means that we can accurately specify some highly complex shapes with minimal use of line, especially when the drawings are viewed in stereo. Unfortunately, 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/back-facing problem in our software, for example through view-dependent rendering, we may have a much more powerful new aesthetic with which to work.

7.5.3 Use of Color Together with CG Lighting

In traditional drawing or painting, color often functions to indicate lighting on an object that helps define its form. In computer graphics this type of shading is a routine part of the graphics pipeline. The question in free-form modeling via 3D drawing is when to rely on computer graphics shading and when to rely on adjusting color by hand. Working with color seems to have a role here, but probably a much more limited one than we imagined before this investigation.

7.5.4 Stereo Verses 2D Viewing

The faintest hint of the tip of an elbow can 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 lost. This speaks to the importance of stereo and head-tracked rendering, but it also points to an area we may be able to improve. In critiquing these 2D projections, we find that many fundamental drawing principles for indicating depth are ignored in the rendering, and this gives rise to several ideas for incorporating art-based rendering techniques into the drawing of these forms.

7.5.5 Comparison of Cave and Fishtank Environments

Using Drawing on Air within the fishtank environment lent itself to working with fewer marks per model and trying to make each mark the right one. On the other hand, CavePainting in 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. Guide lines were required as a strategy for establishing correct working proportions. Thus, both environments




Figure 7.13: The CavePainting on the left was inspired by the oil painting on the right.

have strengths and weaknesses, and future research that captures the strengths of both environments will be useful.

7.6 Additional Artistic Investigations

In this section, we present results from additional artistic investigations that have been less rigorously critiqued but were nevertheless successful in their use of 3D drawing and exploration of VR for artistic purposes.

7.6.1 Contrasts in Styles of Depiction: Painting vs. Sculpting

Here we return to the theme of appropriate use of drawing-style modeling for depicting form. Comparing the work in Figures 7.11, 7.13, and 7.14 points out some very interesting artistic contrasts. In Figure 7.13, an early work with CavePainting, we see the tightest connection to a traditional painting process in the way the 3D form




Figure 7.14: CavePainting of a lion fish, by Helen Zhu.




Figure 7.15: CavePainting of a falcon, by Helen Zhu.

is suggested by many overlapping 3D brushstrokes. In Figure 7.11, the style is much more similar to drawing: the 3D ribbons of form are long and precise and there is a great deal of empty space, left for the viewer to fill in.

In contrast to these approaches, some artists have chosen to approach CavePainting more as 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 in this style, primarily with a tapered tube-shape brush stroke rather than ribbons. The lighting effects produced by this primitive contribute to the sense of volume evoked in her CavePaintings, shown in Figures 7.14 and 7.15.

Each of these styles is valid for artistic creation. While the work in Figures 7.14 and 7.15 is quite appealing visually, works like these are very difficult to create with CavePainting and require great patience due to the inherent loose quality of the medium. CavePainting lends itself to styles that embrace a loose quality, such as we find in Figure 7.13 or in gesture drawings, such as Figure 7.10. In contrast, Drawing on Air lends itself to much more controlled illustration, as seen in Figure 7.11 and in more scientifically oriented work presented in the bat illustrations in chapter 3. Still, with Drawing on Air, creating something as sculptural as Zhu's work would be very time consuming. We conclude that 3D drawing techniques are most effectively utilized for modeling situations that require more minimal use of line than that found in these sculptural forms. However, artists find this immediate style of interaction with computers exciting, and as investigations with these tools advance, additional unforeseen artistic applications of the techniques presented here are likely to be developed.

7.6.2 Interactive Environments Utilizing the Space of the Cave

We have also been involved in several Cave art environments, and we discuss insights from one of these, Hiding Spaces, in this section. These works are intended to be interactively experienced in the Cave, and the investigations focused on viewer perception of form and sense of presence within the interactive environment.

In Hiding Spaces [71, 72] (details in Figure 7.16), CavePainted tree forms twist up




Figure 7.16: In *Hiding Spaces* the imagery on the floor and walls of the Cave shifts in response to the user's movement. Shown here is a progression of three snapshots from the front wall of the Cave. As a viewer walks from left to right within the space of the Cave a door appears on the wall and then disappears again as the viewer comes closer to it.

from the floor and into the walls of the Cave, where they meet carefully constructed imagery textured on 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 Cave space. 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.

Shifts in imagery are triggered by the viewer's body movement within the space, which is measured by projecting the head position onto the plane defined by floor of the Cave. This plane is divided into a grid of 2 ft. x 2 ft. squares and textures to display on the Cave walls are determined by 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, producing a simple view-dependent rendering technique.

This work illustrates the utility of art processes for exploring visual concepts in VR. *Hiding Spaces* resulted from a thorough months-long collaborative investigation

with artist Cynthia Rubin of RISD. We focused our investigation on the ambiguity and tension present when the physical walls of the Cave are defined in VR by virtual form that encloses the Cave space. This is an unusual concept, since in traditional Cave use, the walls are made to disappear and we look beyond them to the surrounding virtual scene. Projecting form onto the walls suddenly makes our virtual experience much more closely resemble our physical experience. *Hiding Spaces* explores this notion in detail, constantly working to promote user exploration of the space and question our “real-world” assumption that the form we view is static.

The compelling visual effects explored in this work prompted a more scientifically structured investigation into the concept of matching virtual form to the physical walls of the Cave and its use in scientific visualization [49].

7.7 Discussion of Cave-Based Artistic Media

CavePainting and the series of artistic investigations that resulted from it led to interesting insights regarding effective and nontraditional use of the Cave environment. Many of these have implications for scientific as well as artistic Cave applications.

7.7.1 Movement in the Space of 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 is a feature of nearly all VR applications built for the Cave, but very few applications actually encourage this sort of movement. Artistic applications do a better job of this in general than do scientific ones [71, 72]. Yet scientific applications almost always provide the motivation for establishing and maintaining costly Cave facilities. In *CavePainting*, the movement-heavy style of interaction is surprisingly effective, so much so that it seems that an application that does not leverage this characteristic of the medium to explore and become proficient with the data displayed is not effectively utilizing the power of the Cave medium. Perhaps this style 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 that exploit 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.

We hypothesize that the more users move within the Cave space, the more spatial understanding they develop. It will be interesting to test this hypothesis with respect to scientific data. Is encouraging movement within VR actually beneficial to scientific discovery? If so, then the implications for how we structure our presentations of visual information in VR are significant.

7.7.2 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 bodies as a tool for making more controlled movements. Locking one's elbow in place to create a rigid lever hinged at the shoulder in order to make 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 make certain types of strokes in the Cave.

7.7.3 Prop-Based Interaction

CavePainting uses the space of the Cave in another nontraditional way by placing a table that holds interaction props along one wall of the Cave. There is no question




Figure 7.17: This table of props in the Cave mixes physical interaction with the virtual Cave medium.

that the interactions performed through this table could be performed more efficiently through virtual pop-up menus located near the artist or with 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 7.17, 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, user interactions are limited. A typical interactive scenario consists of standing near 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 is performed 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 key use of the defining characteristic of the Cave, its large walkable size.

A second prop-based interaction also leverages the space and walls of the Cave environment. As an alternative to sweeping brush movements through space, form



Figure 7.18: A physical bucket prop is used to throw virtual paint onto the walls and floor of the Cave.

can be painted in the Cave by pouring virtual paint out of a bucket prop, as seen in Figure 7.18. In practice, this technique is most valuable as an exploratory tool rather than a method for serious design or modeling; however, it builds on the idea that our presence in the environment is enhanced when virtual forms coexist with physical forms, and, in doing so, it provides a compelling example of presence in VR that is rarely seen in scientific applications.

7.7.4 Importance of Artistic Process

Some of the most interesting feedback from artists regarding CavePainting concerns the artistic process: many felt that this process is just as interesting an artistic contribution as the work that results from it. We were encouraged to examine ways of capturing some sense of this process within the presentation of finished works.

The art world is often 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. In a few successful artworks, however, the




Figure 7.19: A timeline widget on the floor of the Cave controls a visualization of the painting process.




Figure 7.20: Visualization of the painting process for *La Guitarrista Gitana*.

artist gives us a rare glimpse into the artistic process, for example, Pablo Picasso's series of lithographs of a bull [2]. 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 the special viewing mode shown in Figure 7.19, in which a timeline widget is displayed on the floor of the Cave and an observer's movement along the timeline controls the state of the painting displayed. At the far left, we see a blank canvas, but as the viewer moves to the right, more and more strokes are displayed in the order created by the artist.

This visualization of process provides insight into the use of these new tools. Figure 7.20 shows an interesting example of the development of a human figure. Initially a chair is created to define the space for the form. Then, a scaffolding is set up as a way of establishing correct proportion. Finally, the scaffolding is covered by textured brush strokes that define the form of the final work.

The walking-based interaction described above that drives the display of the process brings us back to the theme of making use of the space of the Cave and creating a connection between the physical and the virtual. Here, movement within the Cave space is encouraged — in fact, it is required to drive the display. Also, we see again a connection to the physical surface of the Cave, as the timeline widget exists directly on the floor of the Cave display.

7.8 Conclusions

Artists learn through creative experimentation. Directed study of visual problems like these lead, by design, to important insights about effective visual representation. We conclude that while these have clear implications for artistic depiction, they also make scientific contributions by helping us better understand depiction in an scientific medium and suggesting future directions in computer research.

Chapter 8

Discussion and Conclusions

In this chapter, we present a final discussion of the dissertation, including directions for future research and a review of our primary contributions and conclusions.

8.1 Discussion and Future Work

Several limitations of our work provide additional insight and suggest important future research directions.

8.1.1 Drawing Circles and Other Specific 3D Input Tasks

The Drawing on Air methods provide a general-purpose technique for controlled 3D input in that smooth trajectories of a variety of shapes can be captured with the tool. We can gain some additional insight, however, by considering what appears to be a limitation of the approach that is especially apparent in the two-handed tape drawing technique. It is very difficult to draw a circle with Drawing on Air.

Intuitively, this seems like a fairly big problem. A circle is a simple shape, and these are controlled drawing techniques. If we think about it a bit more, an interesting analogy comes to mind: the pencil. Drawing a circle with a pencil is hard too. It can be done, but it often requires first drawing some guidelines and then connecting the dots with a smooth, controlled stroke. Drawing on Air works the same way: It makes possible that second step of the process, the smooth, controlled stroke that would be impossible to create with freehand 3D drawing. Drawing on Air is a controlled




Figure 8.1: A proposed system for surgical-path planning requires controlled 3D input. In data-driven applications like this one, it may be possible to use the data to help steer user-supplied 3D input in the right direction, thereby further increasing the control and utility of continuous 3D interactions.

drawing technique in the same sense that artists control drawing with a pencil on paper. If what we want is a perfectly controlled, specific shape, like a circle, then other special-purpose techniques are likely to be faster and more appropriate.

One interesting direction for future research is to investigate alternative drawing tools that lend themselves to controlled drawing of different classes of shape. Tools based on perpendicular rather than tangent drawing constraints may be better at targeting circular forms, for example. We lay the groundwork for this future investigation with the discussion of user-guided 3D drawing in section 4.7.3.

A related line of inquiry targets specific controlled 3D input problems in non-artistic domains. Consider, for example, a neurosurgeon using a VR surgical path-planning application such as that pictured in Figure 8.1. Clearly, 3D precision is important in this application. How does the surgeon reliably input a 3D pathway through the neural fibers of the brain? The techniques presented in this thesis would be one solution, but taking advantage of the specific data being visualized may allow even more controlled 3D input. For example, if the pathway entered by the surgeon should not intersect vital neural structures, input techniques should help steer the surgeon safely around these areas as he draws through the dataset. Data-driven interactions of this style have great potential increase the utility of 3D computer

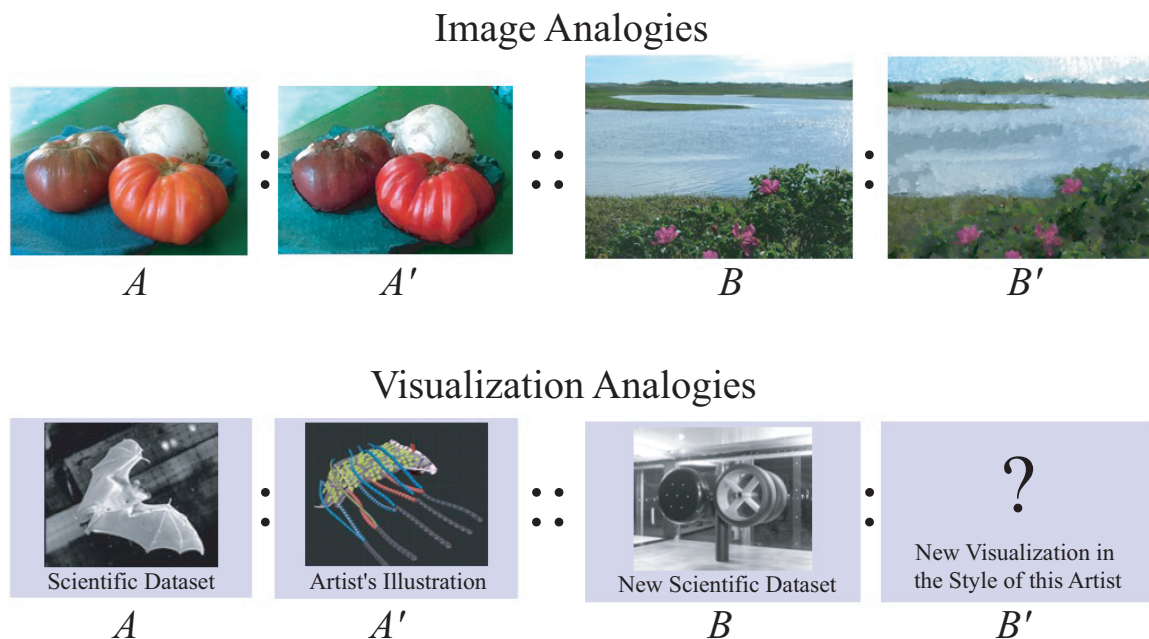


Figure 8.2: In related work in image synthesis, the result B' is automatically generated after learning the transformation from A to A' . If we can characterize an artist's style of 3D illustration, then similar algorithms may let us reapply this style to new scientific datasets.

input for science.

8.1.2 Reuse of Artist-Created Visualization Designs: Visualization by Analogy

Our work enables artists to create 3D illustrations of science. We found these were useful both as expository illustrations and as design prototypes for more exploratory, data-driven visualizations. This new ability to create refined 3D illustrations opens a door to an exciting new direction. If an artist's 3D illustration *style* can be characterized by analyzing his illustration of some well-understood phenomena, then it may be possible to apply this style to new scientific data to automatically generate a visualization that captures the visual artist's visual decisions.

This future direction is motivated by a technique called Image Analogies, a related approach in texture synthesis [43]. Figure 8.2 outlines the proposed system. The top row of images is a result from work in Image Analogies. Given a photograph A and

an artist's watercolor rendering of that photograph A' , a new photograph B is input, and the result B' is automatically generated by a style-mapping function learned from the input images.

The useful visualization analogy illustrated in the bottom row of Figure 8.2 is envisioned. Given a scientific dataset A and an artist's 3D illustration of the data A' , the illustration style of the artist is learned by comparing the 3D illustration to the underlying scientific data. Then, given a new scientific dataset B , a new visualization B' is generated in the style of the artist. Some restrictions would need to govern appropriate pairs of datasets. The style of a brain-tumor visualization is unlikely to provide a useful mapping to a fluid-flow problem, but two different experimental fluid flows may be close enough to provide a useful mapping. Figure 8.2 suggests one example is suggested: the dataset A is from the bat-flight problem discussed earlier in the dissertation and the dataset B describes flow around landing gear. Both sets of data were collected experimentally in wind tunnels, and in both the flow structure in the wake of an object is important to visualize.

Approaches like this one that harness artistic insight and reproduce it in new situations have great potential to broaden the impact of artists in scientific domains. As we look for ways to leverage artistic insight in visual problems in science, we need also to examine how to characterize this insight, either implicitly through learning and reproducing effective styles, as in Visualization by Analogy, or explicitly through critique, evaluation, and dissemination of effective guidelines for visual depiction.

8.1.3 Additional New Directions

Many other new directions are made possible by the techniques, methodologies, and theoretical foundations explored in this dissertation. We anticipate significant advances in tools that capture the richness of artistic interactions in the physical world, yet let artists explore media that are not possible without the aid of computer technology. As these efforts mature, we expect that they will continue to benefit artist-scientist collaborations and contribute to significant discoveries in art, science, and many related fields.

8.2 Summary of Primary Contributions and Conclusions

This dissertation investigates using interactive 3D drawing to control free-form modeling for art and scientific visualization. In this section we provide a summary and conclusions from each of the main areas of research contributions, followed by conclusions from the entire body of work.

8.2.1 Interactive Algorithms for 3D Drawing

Our work on interactive algorithms for 3D drawing began with the CavePainting system, from which it was concluded that 3D drawing has some important benefits for artistic free-form modeling but also suffers from a lack of control, making it difficult to apply to modeling subjects that demand more precision, such as are found in science and in artistic illustration.

To address the issue of control, a toolset of interactive 3D drawing algorithms called Drawing on Air was developed. Within Drawing on Air, two modes for input of controlled, stylized, continuous 3D trajectories were presented: one-handed drag drawing and two-handed tape drawing. Each provides for rich, high-degree-of-freedom input in the form of 3D position, 3D orientation, and 1D line-weight parameters. Presented in this work were haptic-aided algorithms for tangent-preserving transitions between the two drawing modes and for tangent-preserving redrawing of 3D curves. A brush model for 3D pigment was also presented.

Drawing on Air was evaluated with driving illustration applications and with the formal Drawing Control Experiment described in chapter 4. Quantitative results indicated that both one-handed drag drawing and two-handed tape drawing could be controlled by artists with roughly the same precision, while both techniques showed a significant improvement in precision over freehand and haptic-aided-freehand alternatives. To explain these results, a discussion of user-guided drawing methods, of which Drawing on Air is a particular example, was presented. Within this discussion, we explained how this style of interaction might be extended to other input scenarios.

We conclude that user-guided drawing methods, such as Drawing on Air, provide

significant advances in control over state-of-the-art freehand alternatives. The difference measured in our quantitative analysis is clearly evident in artistic use of the tool, and the additional control provided by these techniques has a real impact for artists in the difficulty of subjects that they can address effectively with 3D drawing-based techniques.

8.2.2 Statistical Models for 3D Tracing Tasks

Statistical models for user performance in 3D tracing tasks were developed based on theory in psychology, neuroscience, and human-computer interaction. This work extended the Steering Law to scenarios in which its current form is too simplified to be accurate. In particular, these models address cases where the drawing trajectories are of roughly similar length and study controlled, 3D drawing scenarios rather than 2D draw-as-fast-as-you-can scenarios. In contrast to the Steering Law, drawing error as well as drawing times was also investigated.

The first set of models works at a local level. They describe expected performance (drawing time, positional error, and directional error) in tracing a small section of a 3D curve based on the local orientation, curvature, and interactions between orientation and curvature of the curve traced. Results indicate that each of these factors is important in characterizing the difficulty of the task, leading to a formalization of a local index of difficulty for 3D tracing. For some combinations of input technique and performance measure, this index may be simplified by removing non-significant terms.

A second set of models for global performance in tracing an entire 3D curve was derived as an integration of the local models along a 3D curve. Again, regression analysis with experimental data indicated significant trends, but only for directional error. Alternative experimental designs are required to examine global trends in positional error and drawing times.

We conclude that adding terms to the framework established by the Steering Law is important for describing controlled 3D input, particularly when comparing 3D input trajectories of roughly similar length. Analysis has already led to better understanding of the problem of 3D drawing and helped to target important future research directions. In particular, the finding that error is more closely correlated

with curvature for Drawing on Air techniques than for freehand techniques has suggested alternative designs for high-curvature input tasks. In the future, we expect these models to be useful as tools for comparative evaluation of controlled 3D input techniques and as a common evaluative framework for future 3D interface development.

8.2.3 Artistic Collaboration in Scientific Visualization

Our work toward supporting artistic collaboration in scientific visualization was motivated by a lack of artist-accessible tools for visual design in visualization media. A series of four experiments was conducted to refine an appropriate toolset to address this limitation. The tools developed were based in 3D drawing-style interactions and addressed issues in developing appropriate artistic control for science, establishing connections to underlying scientific data in hand-drawn visualization designs, and prototyping scientific interactive scenarios in VR. On the basis of conclusions from these experiments, Scientific Sketching, a methodology for iterative, collaborative design of scientific visualizations for VR was developed. Scientific Sketching builds on related work in rapid prototyping within the software engineering literature and draws significantly upon artistic design processes, especially in its use of critique as an evaluative design tool. This methodology describes specific roles for visual experts, scientists, and technologists as the design team advances through for a four-stage process that includes a significant role for design and evaluation within VR, as made possible by 3D drawing tools.

A theoretical analysis of Scientific Sketching builds on van Wijk's formal description of the visualization process [97] and extends it to include visualization design. By posing the design process as an optimization problem minimizing time between iterations while simultaneously maximizing the utility of each design-evaluation iteration, it is demonstrated that Scientific Sketching takes clear and deliberate steps towards an optimal global solution to the visualization design task.

We conclude from these results, and our numerous collaborations with artists and scientists over several years, that visual experts such as artists, illustrators, and designers can play important roles in scientific visualization, even in challenging media, such as VR. However, in order to maximize the utility of the contributions of these

experts, their visual design process must be supported. If VR is to be used as an effective visualization medium, then visual design in VR must be supported. Scientific Sketching is a powerful step in this direction that is made possible, in large part, by the VR drafting pencil provided by controllable 3D drawing techniques.

8.2.4 Computer Tools for 3D Artistic Illustration

In addition to scientific applications, results from several serious art investigations were described. Many of these grew from collaboration with artists at the Rhode Island School of Design (RISD). Results of this work (interactive art virtual spaces, videos, academic papers, and other artifacts) have been well received at juried art exhibitions and conferences, and within critiques in the RISD Illustration Department.

Conclusions from a several-month-long guided artistic study, including a series of more than ten evaluative critique sessions designed to investigate CavePainting and Drawing on Air as applied to artistic anatomical illustration, were also presented. Detailed insight resulted from these evaluations in the areas of sculpting vs. suggesting form, economy of line and occlusion, use of color together with lighting, stereo vs. 2D viewing, and differences between Cave and fishtank VR environments.

This area of our work contains a mix of contributions to both art and computer science literature. One of the important conclusions to be drawn from these endeavors is that serious visual investigations in art can play a role in guiding computer research. Insight from this work has led to: 1. a refined framing of the control problem that is the focus of this dissertation; 2. conclusions about appropriate use of the Cave VR environment, including cognitive and emotional responses to virtual stimulus; and 3. conclusions about effective new strategies for depicting complex 3D forms. These results prompted several follow-up scientific investigations, including a study exploring changes in perception of 3D form in VR under different visual background conditions [49] and body-scale 3D drawing as an interaction metaphor for use in scientific visualization applications [86].

Supporting the thesis of this dissertation, conclusions from these art investigations confirm that the same lack of control limiting scientific depiction is also a barrier to controlled artistic depictions. With the introduction of Drawing on Air, artists are able to overcome much of this limitation to create more effective illustrations

of challenging subjects in art, such as anatomical forms. The scientific community has much to learn from artists and other experts in visual depiction. In addition to scientifically motivated collaborations, driving problems in art practice, such as the artistic illustrations explored here, have potential to elicit fundamental knowledge of importance to the fields of computer graphics and visualization.

8.2.5 General Conclusions

This research investigates the potential of combining rich, controllable interactions found in traditional artistic media with 3D computer tools. Computer input techniques based on sweeping 3D movement of the hands are adopted because they capture the intuitiveness, immediacy, and high-degree-of-freedom input characteristic of traditional artistic media like drawing and painting. In overcoming a control limitation commonly found in this style of computer input, research advances presented here have succeeded in increasing the expressive power of artists in presenting challenging, important subjects in both science and in art through free-form 3D modeling techniques.

As we look to the future, the symbiosis between art and science, as made possible by computers and as exhibited in this dissertation, will continue to be a theme of important research advances in many fields. As our ability to collect vast amounts of complex scientific data increases, one of the most important challenges will be analyzing this data, and as humans, one of our most effective analytic tools is our visual system. In endeavors to advance effective visual communication using computer tools, we cannot afford to ignore fields such as art, illustration, and design that have focused for years on related visual problems. Rather, we need to engage with visual experts and use our role as computer toolsmiths wisely to facilitate more effective collaborations. We believe that the contributions of this thesis are an important step in this promising direction.

Bibliography

- [1] Autodesk Maya software. <http://www.autodesk.com>.
- [2] *Picasso the Printmaker: Graphics from the Marina Picasso Collection*. Dallas Museum of Art, 1983.
- [3] Johnny Accot and Shumin Zhai. Beyond Fitts' law: Models for trajectory-based HCI tasks. In *Proceedings of CHI '97*, pages 295–302, 1997.
- [4] John Alex. *Hybrid Sketching: A New Middle Ground Between 2- and 3-D*. PhD thesis, Massachusetts Institute of Technology, 2005.
- [5] Ravin Balakrishnan, George Fitzmaurice, Gordon Kurtenbach, and William Buxton. Digital tape drawing. In *Proceedings of UIST '99*, pages 161–169, 1999.
- [6] Bill Baxter, Vincent Scheib, Ming C. Lin, and Dinesh Manocha. DAB: interactive haptic painting with 3D virtual brushes. In Eugene Fiume, editor, *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, pages 461–468. ACM, ACM Press / ACM SIGGRAPH, New York, 2001.
- [7] Phyllis W. Berman. Young children's use of the frame of reference in construction of the horizontal, vertical, and oblique. *Child Development*, 47(1):2259–263, 1976.
- [8] Phyllis W. Berman, Joseph G. Cunningham, and John Harkulich. Construction of the horizontal, vertical, and oblique by young children: Failure to find the "oblique effect". *Child Development*, 45(2):474–478, 1974.

- [9] David Bourguignon. *Interactive Animation and Modeling by Drawing — Pedagogical Applications in Medicine*. PhD thesis, Institut National Polytechnique de Grenoble, 2003.
- [10] David Bourguignon, Marie-Paule Cani, and George Drettakis. Drawing for illustration and annotation in 3D. In *EUROGRAPHICS Computer Graphics Forum*, pages C114–C122, 2001.
- [11] Bill Brody and Chris Hartman. Painting space with BLUI. In *SIGGRAPH '00 Conference Abstracts and Applications*, page 242, 2000.
- [12] Frederick P. Brooks, Jr. *The Mythical Man-Month*. Addison-Wesley, anniversary edition, 1995.
- [13] 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.
- [14] Bill Buxton. *Sketching User Experiences: Getting the Design Right and the Right Design*. Morgan Kaufmann, 2007. In press.
- [15] William Buxton, George Fitzmaurice, Ravin Balakrishnan, and Gordon Kurtenbach. Large displays in automotive design. *IEEE Computer Graphics and Applications*, 20(4):68–75, 2000.
- [16] James H. Clark. Designing surfaces in 3-D. *Communications of the ACM*, 19(8):454–460, 1976.
- [17] Donna Cox. Using the supercomputer to visualize higher dimensions: An artist's contribution to scientific visualization. *Leonardo*, 21(3):233–242, 1988.
- [18] Cassidy J. Curtis, Sean E. Anderson, Joshua E. Seims, Kurt W. Fleischer, and David H. Salesin. Computer-generated watercolor. In *Proceedings of SIGGRAPH '97*, pages 421–430, 1997.
- [19] James E. Cutting and Peter M. Vishton. *Handbook of Perception and Cognition*, volume 5: Perception of Space and Motion, chapter Perceiving layout

and knowing distances: The integration, relative potency, and contextual use of different information about depth, pages 69–117. Academic Press, New York, 1995.

- [20] N. Davies, J. Landay, S. Hudson, and A. Schmidt. Guest editors' introduction: Rapid prototyping for ubiquitous computing. *IEEE Pervasive Computing*, 4(4):15–17, 2005.
- [21] Michael Deering. High resolution virtual reality. In *SIGGRAPH '92 Conference on Computer Graphics and Interactive Techniques*, pages 195–202, 1992.
- [22] Michael F. Deering. Holosketch: A virtual reality sketching/animation tool. *ACM Transactions of Computer-Human Interaction*, 2(3):220–238, 1995.
- [23] Steven Dow, Blair MacIntyre Jaemin Lee, Christopher Oezbek, Jay David Bolter, and Maribeth Gandy. Wizard of Oz support throughout an iterative design process. *IEEE Pervasive Computing*, 4(4):18–26, 2005.
- [24] Munro Ferguson. June, 2003. Film created with the SANDDE animation system. National Film Board of Canada <http://www.nfb.ca>.
- [25] Munro Ferguson. Falling in love again, 2006. Film created with the SANDDE animation system. National Film Board of Canada <http://www.nfb.ca>.
- [26] M. Fineman. *The Nature of Visual Illusion*. New York: Dover, 1996.
- [27] Michele Fiorentino, G. Monno, P. A. Renzulli, and A. E. Uva. 3D pointing in virtual reality: Experimental study. In *Proceedings of Ingeggraf International Conference on Tools and Methods Evolution in Engineering Design*, Napoli, 2003.
- [28] Michele Fiorentino, Giuseppe Monno, Pietro Alexander Renzulli, and Antonio E. Uva. 3D sketch stroke segmentation and fitting in virtual reality. In *Proceedings of GRAPHICON*, pages 188–191, 2003.
- [29] Paul Fitts. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47:381–391, 1954.

- [30] George Fitzmaurice, Ravin Balakrishnan, Gordon Kurtenbach, and William Buxton. An exploration into supporting artwork orientation in the user interface. In *Proceedings of CHI '99*, pages 167–174, 1999.
- [31] Martha Flanders, Leigh A. Mrotek, and C.C.A.M. Gielen. Planning and drawing complex shapes. *Experimental Brain Research*, 171(1):116–128, 2005.
- [32] Fundació Joan Miró. Joan Miró, the gold of the azure, 1967, 2007. http://www.bcn.fjmiro.es/angles/_coleperm/_salaanys6070/lordelatzur.html.
- [33] Tinsley A. Galyean and John F. Hughes. Sculpting: An interactive volumetric modeling technique. In *Computer Graphics (Proceedings of ACM SIGGRAPH 91)*, volume 25, pages 267–274. ACM, 1991.
- [34] Bruce Gooch and Amy Gooch. *Non-Photorealistic Rendering*. A.K. Peters Ltd., 2001.
- [35] Arthur D. Gregory, Stephen A. Ehmman, and Ming C. Lin. inTouch: Interactive multiresolution modeling and 3D painting with a haptic interface. In *Proceedings of IEEE VR 2000*, pages 45–52, 2000.
- [36] Jen Grey. Human-computer interaction in life drawing, a fine artist’s perspective. In *Proceedings of the Sixth International Conference on Information Visualisation*, 2002.
- [37] 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.
- [38] Tovi Grossman, Ravin Balakrishnan, Gordon Kurtenbach, George Fitzmaurice, Azam Khan, and Bill Buxton. Creating principal 3D curves with digital tape drawing. In *Proceedings of CHI '02*, pages 121–128, 2002.
- [39] Tovi Grossman, Ravin Balakrishnan, and Karan Singh. An interface for creating and manipulating curves using a high degree-of-freedom input device. In *Proceedings of CHI '03*, pages 185 – 192, 2003.

- [40] Song Han and G. Medioni. 3dsketch: modeling by digitizing with a smart 3D pen. In *ACM International Conference on Multimedia*, pages 41–49, 1997.
- [41] 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.
- [42] Christopher G. Healey and James T. Enns. Perception and painting: A search for effective, engaging visualizations. *IEEE Computer Graphics and Applications*, 22(2):10–15, 2002.
- [43] Aaron Hertzman, Charles E. Jacobs, Nuria Oliver, Brian Curless, and David H. Salesin. Image analogies. In *Proceedings of SIGGRAPH 2001*, pages 327–340, 2001.
- [44] Jing Hua and Hong Qin. Haptics-based dynamic implicit solid modeling. *IEEE Transactions on Visualization and Computer Graphics*, 10(5):574–586, 2004.
- [45] Takeo Igarashi, Satoshi Matsuoka, and Hidehiko Tanaka. Teddy: a sketching interface for 3D freeform design. In *Proceedings of ACM SIGGRAPH 99, Computer Graphics Proceedings, Annual Conference Series*, pages 409–416. ACM, ACM Press / ACM SIGGRAPH, New York, 1999.
- [46] Victoria Interrante, Henry Fuchs, and Stephen Pizer. Conveying the 3D shape of smoothly curving transparent surfaces via texture. *IEEE Transactions on Visualization and Computer Graphics*, 3(2):98–117, 1997.
- [47] Victoria Interrante and Chester Grosch. Strategies for effectively visualizing 3D flow with volume LIC. In *Proceedings of IEEE Visualization 1997*, pages 421–424, 1997.
- [48] Cullen Jackson, Daniel Acevedo, David H. Laidlaw, Fritz Drury, Eileen Vote, and Daniel Keefe. Designer-critiqued comparison of 2D vector visualization methods: A pilot study. *ACM SIGGRAPH 2003 Sketches and Applications*, July 2003.

- [49] 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.
- [50] Robert D. Kalnins, Lee Markosian, Barbara J. Meier, Michael A. Kowalski, Joseph C. Lee, Philip L. Davidson, Matthew Webb, John F. Hughes, and Adam Finkelstein. WYSIWYG NPR: drawing strokes directly on 3D models. *ACM Transactions on Graphics*, 21(3):755–762, 2002.
- [51] Olga Karpenko, John F. Hughes, and Ramesh Raskar. Free-form sketching with variational implicit surfaces. *Computer Graphics Forum*, 21(3), 2002.
- [52] Olga Karpenko, John F. Hughes, and Ramesh Raskar. Epipolar methods for multi-view sketching. In *Eurographics Workshop on Sketch-Based Interfaces*, 2004.
- [53] Daniel F. Keefe. The making of La Guitarrista Gitana. In *SIGGRAPH 2002 Art Gallery Catalogue*, 2002.
- [54] 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 I3D 2001*, pages 85–93, 2001.
- [55] 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, 2005.
- [56] Daniel F. Keefe, Robert C. Zeleznik, and David H. Laidlaw. Drawing on air: Input techniques for controlled 3d line illustration. *IEEE Transactions on Visualization and Computer Graphics*, 2007. In press.
- [57] J.F. Kelly. An empirical methodology for writing user-friendly natural language computer applications. In *Proceedings of ACM SIG-CHI '83*, pages 193–196, 1983.

- [58] 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.
- [59] Michael M. Kirby, Daniel F. Keefe, and David H. Laidlaw. *The Visualization Handbook. Painting and Visualization*, pages 873–891. Elsevier Inc., 2005.
- [60] Rick Komerska and Colin Ware. Haptic state-surface interactions. *IEEE Computer Graphics and Applications*, 24(6):52–59, 2004.
- [61] Aidong Lu, Christopher J. Morris, Joe Taylor, David S. Ebert, Charles Hansen, Penny Rheingans, and Mark Hartner. Illustrative interactive stipple rendering. *IEEE Transactions on Visualization and Computer Graphics*, 9(2):127–138, 2003.
- [62] 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.
- [63] Scott MacKenzie. Fitts’ law as a research and design tool in human-computer interaction. *Human-Computer Interaction*, 7(1):91–139, 1992.
- [64] Wille Mäkelä. Working 3D meshes and particles with finger tips, towards an immersive artists’ interface. In *IEEE VR 2005 Workshop Proceedings*, pages 77–80, 2005.
- [65] Lee Markosian, Jonathan M. Cohen, Thomas Crulli, and John Hughes. Skin: a constructive approach to modeling free-form shapes. In *Proceedings of ACM SIGGRAPH 99*, Computer Graphics Proceedings, Annual Conference Series, pages 393–400. ACM, ACM Press / ACM SIGGRAPH, New York, 1999.
- [66] Barbara J. Meier. Painterly rendering for animation. In *Proceedings of SIGGRAPH ’96*, pages 477–484, 1996.
- [67] Paul Morstad. Moon man, 2004. Film created with the SANDDE animation system. National Film Board of Canada <http://www.nfb.ca>.

- [68] Allen Newell and Stuart K. Card. The prospects for psychological science in human-computer interaction. *Human-Computer Interaction*, 1(3):209–242, 1985.
- [69] 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.
- [70] Gonzalo Ramos, Matthew Boulos, and Ravin Balakrishnan. Pressure widgets. In *Proceedings of CHI '04*, pages 487–494, 2004.
- [71] Cynthia B. Rubin and Daniel F. Keefe. Hiding spaces: A cave of elusive immateriality. ACM SIGGRAPH 2002 Conference Abstracts and Applications, July 2002.
- [72] 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.
- [73] 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.
- [74] Stefan Schaal and Dagmar Sternad. Origins and violations of the 2/3 power law in rhythmic three-dimensional arm movements. *Experimental Brain Research*, 136(1):60–72, 2004.
- [75] Steven Schkolne. *3D Interfaces for Spatial Construction*. PhD thesis, California Institute of Technology, 2003.
- [76] Steven Schkolne, Hiroshi Ishii, and Peter Schröder. Immersive design of DNA molecules with a tangible interface. In *Proceedings of IEEE Visualization*, pages 227–234, 2004.
- [77] Steven Schkolne and Daniel F. Keefe. Visual space after virtual reality. In *Proceedings of Young VR 2001*, Korea, December 2001.

- [78] Steven Schkolne, Michael Pruett, and Peter Schröder. Surface drawing: creating organic 3D shapes with the hand and tangible tools. In *Proceedings of CHI '01*, pages 261–268, 2001.
- [79] 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.
- [80] Will Schroeder, Kenneth M. Martin, and William E. Lorensen. *The Visualization Toolkit*. Prentice-Hall, Inc., 2nd edition, 1998.
- [81] Mark Schroering, Cindy Grimm, and Robert Pless. A new input device for 3D sketching. In *Vision Interface*, pages 311–318, 2003.
- [82] Jürgen Schulze, Andrew Forsberg, Alexander Kleppe, Robert Zeleznik, and David H. Laidlaw. Characterizing the effect of level of immersion on a 3D marking task. In *Proceedings of HCI International*, 2005.
- [83] SensAble Technologies, Inc. *FreeForm Concept Product Brochure*, 2005.
- [84] P. Servos, H. Carnahan, and J. Fedwick. The visuomotor system resists the horizontal-vertical illusion. *Journal of Motor Behavior*, 32(4):400–404, 2000.
- [85] 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.
- [86] Jason S. Sobel, Andrew S. Forsberg, 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.
- [87] Vibeke Sorensen. The contribution of the artist to scientific visualization, 1989. <http://visualmusic.org/text/scivi1.html>.
- [88] Irving Stone, editor. *Dear Theo: The Autobiography of Vincent Van Gogh*. Plume, 1995.

- [89] Ivan Sutherland. *Sketchpad: A Man Machine Graphical Communication System*. PhD thesis, Massachusetts Institute of Technology, 1963.
- [90] Ivan Sutherland. A head-mounted three dimensional display. In *Proceedings of the AFIPS Fall Joint Computer Conference*, volume 33, pages 757–764, 1968.
- [91] Sharon M. Swartz, Kristin L. Bishop, and Maryem-Fama Ismael Aguirre. *Functional and Evolutionary Ecology of Bats*, chapter Dynamic complexity of wing form in bats: implications for flight performance. Oxford University Press, 2005. In press.
- [92] Shun'ichi Tano, Toshiko Matsumoto, and Mitsuru Iwata. Quantitative analysis of human behavior and implied user interface in 3d sketching. *Lecture Notes in Computer Science*, 3101:481–490, 2004.
- [93] William B. Thompson, Peter Willemsen, Amy A. Gooch, Sarah H. Creem-Regehr, Jack M. Loomis, and Andrew C. Beall. Does the quality of computer graphics matter when judging distances in visually immersive environments? *Presence: Teleoperators and Virtual Environments*, 13(5):560–571, 2004.
- [94] Jean Trumbo. The process of critique in visual communication. *Journalism Educator*, 52(2), 1997.
- [95] Edward R. Tufte. *Envisioning Information*. Graphics Press, 1990.
- [96] Craig Upson, Thomas Faulhaber Jr., David Kamins, David H. Laidlaw, David Schlegel, Jeffrey Vroom, Robert Gurwitz, and Andries van Dam. The application visualization system: A computational environment for scientific visualization. *IEEE Computer Graphics and Applications*, 9(4):30–42, 1989.
- [97] Jarke J. van Wijk. The value of visualization. In *Proceedings of IEEE Visualization 2005*, pages 79–86, 2005.
- [98] P. M. Vishton, J. G. Rea, J. E. Cutting, and L. N. Nunez. Comparing effects of the horizontal-vertical illusion on grip scaling and judgment: Relative versus absolute, not perception versus action. *Journal of Experimental Psychology. Human Perception and Performance*, 25(6):1659–1672, 1999.

- [99] P. Viviani and M. Cenzato. Segmentation and coupling in complex movements. *Journal of Experimental Psychology. Human Perception and Performance*, 11(6):828–845, 1985.
- [100] P. Viviani and G. McCollum. The relation between linear extent and velocity in drawing movements. *Neuroscience*, 10(1):211–218, 1983.
- [101] P. Viviani and R. Schneider. A developmental study of the relationship between geometry and kinematics in drawing movements. *Journal of Experimental Psychology. Human Perception and Performance*, 17(1):198–218, 1991.
- [102] P. Viviani and C. Terzuolo. Trajectory determines movement dynamics. *Neuroscience*, 7(2):431–437, 1982.
- [103] Eileen Vote, Daniel Acevedo, David H. Laidlaw, and Martha Joukowsky. Discovering Petra: Archaeological analysis in VR. *IEEE Computer Graphics and Applications*, 22(5):28–50, 2002.
- [104] Colin Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann Publishers, 2000.
- [105] Colin Ware and Glenn Franck. Evaluating stereo and motion cues for visualizing information nets in three dimensions. *ACM Transactions on Graphics*, 15(2):121–139, 1996.
- [106] Gerold Wesche and Hans-Peter Seidel. Freedrawer: a free-form sketching system on the responsive workbench. In *Proceedings of VRST 2001*, pages 167–174, 2001.
- [107] P. Willemsen, A. A. Gooch, W. B. Thompson, and S. H. Creem-Regehr. Effects of stereo viewing conditions on distance perception in virtual environments. *Presence: Teleoperators and Virtual Environments*, 2007. In press.
- [108] 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.

- [109] Shuman Zhai, Johnny Accot, and Rogier Woltjer. Human action laws in electronic virtual worlds — an empirical study of path steering performance in VR. *Presence: Teleoperators and Virtual Environments*, 13(2):113–127, 2004.
- [110] Shumin Zhai. *Human Performance in Six Degree of Freedom Input Control*. PhD thesis, University of Toronto, 1995.
- [111] Song Zhang, Cagatay Demiralp, Daniel F. Keefe, Marco DaSilva, Benjamin D. Greenberg, Peter J. Basser, Carlo Pierpaoli, E. A. Chiocca, T. S. Deisboeck, and David H. Laidlaw. An immersive virtual environment for DT-MRI volume visualization applications: a case study. In *Proceedings of IEEE Visualization 2001*, pages 437–440, 2001.