

CavePainting: A Fully Immersive 3D Artistic Medium and Interactive Experience

Daniel F. Keefe Daniel Acevedo Feliz Tomer Moscovich
David H. Laidlaw Joseph J. LaViola Jr.

Department of Computer Science
Brown University Providence, RI 02912
{dfk,daf,tm,dhl,jjl}@cs.brown.edu

Abstract

CavePainting is an artistic medium that uses a 3D analog of 2D brush strokes to create 3D works of art in a fully immersive Cave environment. Physical props and gestures are used to provide an intuitive interface for artists who may not be familiar with virtual reality. The system is designed to take advantage of the 8 ft. x 8 ft. x 8 ft. space in which the artist works. CavePainting enables the artist to create a new type of art and provides a novel approach to viewing this art after it has been created. In this paper, we describe CavePainting's 3D brush strokes, color pickers, artwork viewing mode, and interface. We also present several works of art created using the system along with feedback from artists. Artists are excited about this form of art and the gestural, full-body experience of creating it.

CR Categories and Subject Descriptors: I.3.6 [Computer Graphics]: Methodology and Techniques - Interaction Techniques; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism - Virtual Reality; J.5 [Arts and Humanities]: Fine Arts

Additional Key Words: 3D painting, 3D modeling, gestures, tangible user interface, Cave

1 Introduction

Paintings in which individual brush strokes are clearly visible are often described as "loosely painted." Impressionistic paintings, for example, often fit this description [26]. When we examine the brush strokes in these paintings closely, we find both subtle and striking variations in color, size, shape, and texture. If we move close enough, our eyes can only see the individual brush strokes. At this level, the amazing variation in the type of stroke used, even within a single painting, is apparent. The layering of strokes on top of each other is also apparent at this close level. As we move away from the painting, we stop seeing individual strokes. Our mind is able to fuse the strokes together and comprehend a complex scene.

Our goal in this project was to create a system which uses 3D brush strokes that function in the same way as the 2D brush strokes described above. That is, we wanted our 3D brush strokes to be clearly visible when viewed up close, but we wanted a viewer to be able to step back from a Cave-painting and see a recognizable scene. Since the artwork is created entirely from these basic stroke elements, the



Figure 1 In CavePainting, the artist arranges and layers 3D brush strokes in space to create a scene. *Wedding Day* - Daniel Keefe.

artist must be provided with a great deal of control over the type of strokes used, variation within a stroke type, and 3D placement of the strokes. The variation and placement of the brush strokes is what allows us to discern a meaningful scene from a collection of strokes.

CavePainting provides the artist with a natural interface for creating a virtual 3D scene. The system runs in a four wall immersive virtual reality system called a Cave [9]. The Cave-painter's interaction with the system is very similar to that of a traditional painter working on a large canvas. The artist is free to create long expressive brush strokes and then step back to observe the work from different angles. Interaction with the computer is accomplished through the use of simple gestures and props that are commonly used in painting and positioned on a table inside the Cave. Scenes are created by layering and arranging virtual 3D brush strokes in space.

We have invited many artists to try this new way of working. Approximately forty art students along with several instructors from the Rhode Island School of Design and Brown University have used

the system in its current state. In addition, feedback from many artists and art students who used the system in earlier stages helped to guide the development of this application and its interface. Recently, we have begun working with several artists who are interested in using the system for serious artistic compositions. Artists have had an overwhelmingly positive reaction to CavePainting as a new artistic medium.

In the remainder of this paper, we first discuss related work. Then, we introduce the different types of paint strokes and describe user interaction in the system. Next, we present artwork and feedback from artists who have used the system and conclude with a discussion of the system and future work.

2 Related Work

In Surface Drawing [23] [24], Schkolne presented a free-form 3D surface construction tool. Surface Drawing is run on a Responsive Workbench. 3D shape is created by moving a tracked CyberGlove™ on the users hand through the air. Schkolne's system has only one type of stroke, although the CyberGlove™ input allows the user to control this stroke to obtain interesting variations as it is created. The focus of Schkolne's work is on fusing these strokes together automatically to form smooth surfaces.

Brody and Hartman presented a sketch at SIGGRAPH 2000 [5] on a system called Body Language User Interface (BLUI). This is another free-form modeling system that runs on a workbench. Lines, point clouds, and extruded surfaces may be created by moving a tracked wand through space. A menu is used to select drawing operations. The distinctive feature of this system is that the geometry created during this process is saved, imported into and rendered in Maya, and then printed as a large panorama.

Other related projects such as 3DM [6], HoloSketch [10], and 3-Draw [21] have used a tracked wand to create free-form lines or geometry. HoloSketch, for example, was able to create free-form "toothpaste" geometry, wire-frame lines, and clouds of triangle particles.

CavePainting differs from these virtual reality free-form modeling applications in several key ways. The first can be attributed to the fact that our system runs in a fully immersive Cave environment. The Cave provides the artist with enough space to stand up and walk around while working. This directly affects the type of work that the artist creates, as well as the way in which the artist works. Additionally, since the user wears shutter glasses in the Cave, he or she is able to see both the real world and the virtual world at the same time. Thus, we are able use physical props and gestures that a painter uses every day to interact with the system, eliminating the need for a menu of drawing operations.

Second, CavePainting provides the artist with fine control over color and a large, varied set of brush strokes with which to work. CavePainting does not attempt to be a modeling system in a traditional sense, where the user is often concerned with exact coordinate representations for the size or shape of objects. Rather, CavePainting aspires to be an extension of painting to three dimensions. Just as an oil painter builds up a painting with layers of varied brush strokes, the Cave-painter creates many different 3D strokes to convey the impression of a 3D scene. Finally, CavePainting promotes the idea that art created by this dynamic 3D tool is meant to be viewed in an interactive 3D display environment, since a static 2D print, no matter how large, cannot truly convey the 3D nature of this type of work. CavePainting presents a viewing mode of its own which takes this notion a step further by providing the observer with additional insight into the artistic process that produced each work.

Several others have worked on volumetric modeling using a tracker as input [13] [25]. CavePainting is fundamentally different from these systems because it creates varied 3D strokes of paint rather than modifying voxel data which represents a solid, such as wood or marble.

Additionally, there are several systems which use 2D input to generate 3D forms [8] [16] [28]. Though these can be used to obtain 3D artistic results, they are limited in their ability to create dramatically different 3D strokes and lack the fully immersive experience that is so critical to CavePainting. Maya [2] is a more advanced artistic package where 2D input can be used to generate 3D form with a painterly quality.

Adding pigment or texture maps to a 3D surface [4] [14] is a technique often referred to as 3D painting. CavePainting defines 3D painting differently. In our case, 3D painting is a process that generates 3D form.

3 Painting with 3D Strokes

A CavePainting is composed of many 3D paint strokes. These individual strokes are layered and arranged in space to produce a scene. The artist can choose between several stroke types. The current stroke types in the system are line, ribbon, tube, bumpy tube, trail of any type of geometry, Jackson_Pollock++, splat, extrusion, and bucket. The artist picks a stroke type to indicate the general characteristics of the stroke. This is analogous to choosing to apply oil paint with a large flat brush, a small round brush, a sponge, or a palette knife, since the artist can obtain considerable variation in a stroke, even after a stroke type has been chosen. The artist actually applies the virtual paint, by moving a tracked paint brush prop around in the cave. The virtual strokes respond to fine variations in the position and orientation of the paint brush prop. The immediacy of the response of the virtual paint to the artist's movements is very important to the artists that use CavePainting. The direct control over the 3D paint is what allows them to create expressive variations in strokes. We found that it was a mistake to try to program too much expressiveness into a stroke type. Rather, artists seem to produce the most expressive strokes when given a simple stroke type that they are able to easily control and immediately see the results of their movements.

The following sections explain the way one artist used CavePainting's 3D paint strokes to create the complex 3D painting shown in Figure 9, in the Results section. This Cave-painting was inspired by an oil pastel drawing the artist did while in a vineyard in Florence. The original drawing is shown in Figure 2a.

As seen in Figure 2b, the artist started by defining a ground plane and a wall rising out of it. This was done in a loose, abstract style with the Jackson_Pollock++ and splat strokes. These strokes are animated as they leave the artist's brush. They fall in the direction that the brush points until they reach one of the walls of the Cave, where they splatter in the virtual world, as if they had actually hit the physical wall of the Cave. Both strokes provide an interesting link between the physical space the user occupies and the virtual world in which he finds himself immersed. We find that this is important in CavePainting, since the artist is essentially defining the space around him or her. The Jackson_Pollock++ stroke drips a line of virtual paint, reminiscent of the drip paintings done by the great expressionist artist [22]. In virtual reality, we are free from some of the limitations imposed on Jackson Pollock. For example, our paint does not need to drip according to gravity. We take advantage of this, and are able to drip on all six sides of the cube defined by the Cave. The same is true of the splat stroke. This stroke drips small droplets of paint that create a random splatter shape when they hit a wall of the Cave.




Figure 2 The progression of a Cave-painting.

In Figure 2c, the artist has started to define some of the structure from which the grape vines will hang. A tube stroke was used to create the posts coming out of the wall. The cross section of the tube stroke follows the orientation of the brush as it is created. By turning the brush while painting, the artist can create a flattened thin stroke or a tube with a circular cross section. The tube stroke is an example of a relatively simple geometry that changes dynamically as the artist moves and twists the brush.

Another simple stroke which has this quality and is used quite often in this painting is the ribbon stroke. It was used to create the leaves on the grape vine shown in Figures 2c and 2d. The actual vines were created using the bumpy tube stroke. This is a tube stroke with a radius that varies automatically. The radius changes based on time, so a long drawn out bump can be created by painting quickly, as was done for the vines, or short round bumps can be created by moving the brush slowly. This was used to create the grapes in this scene. All of these strokes are simple geometric shapes, but the interaction in the system gives the artist the power to vary their appearance as they are applied. This makes them very powerful in terms of their expressive quality to the artist.

The extrusion stroke was used twice in this painting. The petals of the red flower in the flower box and the flower pot on the right side of the image in Figure 2e were created using the extrusion stroke. Creating an extrusion stroke is a three step process. First, a free form line is drawn using the brush. This will be the cross section for the extruded surface. Next, the user positions the brush at an anchor point somewhere near the line and presses the button on the brush. This attaches the line to the brush so that it moves along with the tracked brush that the user holds in his or her hand. The final step is to paint normally with the cross section line attached to the brush. The result is a smooth surface swept out along the path that the user follows with the brush.

The cross section line that is swept out to form an extruded surface is represented as a list of 3D points. We create a mesh by connecting these points from frame to frame as the user moves them around by painting with the brush. A challenge in implementing this stroke was to avoid creating prohibitively large meshes that would

grind rendering to a halt. To preserve the features of the cross section while reducing the number of samples, we compute the Haar wavelet transform of the list of points. We remove small coefficients that have little effect on the curve, and then invert the transform to produce a new list of points. Since we reduce the number of samples, the 3D mesh created by connecting these points together has fewer polygons.

The final stroke illustrated in this artwork is the geometry stroke. This is used to orient and place down a trail of any predefined geometric model. In this example, the artist has imported a model of a leaf and has placed colored instances of this model in several locations to complete the scene. They can be seen on the ground and on the bench in Figure 2f.

4 Interaction

The first goal for the interaction in this system was to make it very natural for an artist to use. We noticed early on that painting in 3D in an area large enough to stand up and walk around was very similar to how a painter works on a large canvas. For example, when working on a large canvas, painters often make long gestural paint strokes. After a few strokes are placed on the canvas, the painter steps back to see the effect the entire painting. As the work progresses, there is a continued motion of the painter between the canvas and some area an optimal viewing distance away from the work. This is, in fact, how many artists, especially painters, work in our Cave.

In addition to making the interaction natural for an artist, a secondary goal was to make the system intuitive for novices and artists of all ages. We wanted CavePainting to scale to different ages and levels of artistic ability just as real painting does.

We also wanted to take advantage of the physical space provided to the user when working in a Cave. When the user first enters this cube, he or she knows that there are projection screen walls approximately 4 feet away on all sides. Rather than trying to completely hide these walls, we tried to take advantage of the fact that the user




Figure 3 The painting table interface.

knows roughly where these walls are. We do this with several of our paint strokes, which interact with the walls as if they were part of the virtual world. We also do this by placing physical props on a real painting table [12] [15][17] (see Figure 3) that is located inside the Cave along one of the walls, and by placing virtual controls in the world so that they always appear to be attached to the physical walls of the Cave.

4.1 Painting



Figure 4 An artist at work.

Strokes are created in 3D space by moving a tracked paint brush through the air, as seen in Figure 4. A 3D path with orientations within the Cave is sampled from the tracker on the brush. For most strokes, the tracker is sampled once per frame. When the single button on the brush is depressed, a stroke is interactively drawn through the sampled 3D points. As the tracked physical paint brush

moves around, its position and orientation are reflected on a 3D virtual brush icon. The size of the stroke being created is represented by the size of the brush icon and is controlled by rotating a knob on the painting table or by using a tracked pinch glove [1] [19] worn on the non-dominant hand.

4.2 Assigning a Stroke Type to the Brush

To select the type of stroke to paint, the artist simply dips the physical paint brush into a cup (located on the painting table shown in Figure 3) that “contains” the desired stroke. This interface is similar in style to the tool tray Fitzmaurice used in his GraspDraw application [11]. When the brush touches the cup, the stroke type is changed. This is implemented by placing conductive cloth on the tip of the brush and along the inside of the paint cups. Audio feedback is given to indicate that the brush’s stroke type has changed. This is an interface which even young artists have been able to readily understand.

4.3 Picking Colors




Figure 5 A 3D color picker represents the Hue, Lightness, and Saturation color space. A large sphere is moved around inside the color space to select a color. Hue is selected by moving around the circumference of the color space representation. Saturation is proportional to the distance from the vertical axis, and the vertical position of the sphere within the space controls the lightness of the color.

To select a specific color, the artist uses a 3D color picker that represents the Hue, Lightness, and Saturation color space. (see Figure 5) This is similar Deering’s 3D RGB cube color picker [10]. There are two ways to activate the color picker. The first is to make a circular gesture with the brush pointed straight up. This makes the color picker appear at the center of the circle with the brush location controlling the color selection. As the user moves the brush through the color space, a large sphere moves with the brush and changes color based on the location within the space. To select the color, dismiss the color picker, and continue painting, the user tilts the brush so that it no longer points up. In practice, we find that advanced users want to switch colors quite often. These users often choose to wear a tracked pinch glove on the non-dominant hand that gives them quick access to the color picker and other options. When the glove is worn, the color picker is activated by pinching and holding down the thumb and index finger. This brings up the color space at the location of the hand when the pinch was made. The position of the non-dominant hand now controls the color selection, and the color is updated immediately on the brush icon. To select a color and continue painting, the user releases the thumb and index finger.

In addition to selecting a specific color for the stroke, the user can

select a color mapping to apply to the space of the Cave. In this mode, the color of the stroke is determined by the position of the brush within the physical space of the Cave. The artist can pick between several mappings of this space. The first is a straight mapping of the red, green, and blue components of the color to the x, y, z coordinates of the brush in the Cave. This mapping is primarily used for abstract paintings. A more useful mapping for the artist is our warm to cool color mapping. In this mapping we selected a range of warm and cool earth tones. We interpolate between these colors to assign a warm color to the brush while it is in the front of the Cave and a progressively cooler color as it moves to the back. We can also assign a “hot spot” in the cave and interpolate between warm and cool colors based on the distance from this spot. We have a large sphere hanging from the ceiling in our cave as part of our Polhemus tracking system. We often use this physical object as the “hot spot”.

These automatic color modes are used mostly for quick gestural paintings that tend to explore the concepts of depth and spatial relations in three-dimensions. The way that artists tend to work within these color modes is to do 3 or 4 quick paintings in a row. Each of these will generally fill the whole space of the cave. After each one, they translate the virtual world aside to clear the space of the cave. Essentially, they turn to a blank canvas.

4.4 Painting with the Bucket



Figure 6 Virtual paint can be dripped out of a real bucket that has a tracker attached to it. This paint can be thrown or dripped onto the walls and floor of the Cave.

In addition to using the brush to paint, users can pick up a real bucket and drip or throw paint out of it and onto the walls and floor of the cave, as shown in Figure 6. When paint is thrown out of the bucket, it flies out in the direction that the bucket is moving and falls with an acceleration equal to gravity. When the bucket is tilted, a stream of paint flows from the lip of the bucket. The size of the flow is calculated based on the angle at which the bucket is held. The flow changes dynamically as the user moves the bucket and adjusts its tilt angle. Thus, this type of painting is another example, of a simple stroke that is very responsive to the artist’s movements.

When not in use, the bucket is kept on two metal hooks on the side of the painting table. When the metal handle of the bucket rests on these two hooks, it creates a circuit. We detect that the user has picked up the bucket and begin tracking it to determine its position and orientation by checking to see whether this circuit is broken.

4.5 Two Handed Interaction

Advanced users have the option of wearing a tracked pinch glove on the non-dominant hand to gain quick access to several features. As mentioned previously, pinching the thumb and the index finger activates the color picker. To set the brush size, the artist pinches and holds down the middle finger and the thumb. The size of the brush is proportional to the distance of the non-dominant hand from the body. This method of setting the brush size offers much finer control than using the knob on the painting table. Both of these tasks can be carried out using the non-dominant hand while the dominant hand (holding the brush prop) is used to paint. Thus, the artist is able to change the width of a stroke or its color as it is being created.¹

Deering used a mouse and a tracker to explore a similar two-hand interaction technique [10]. Others, have also explored two-handed interaction and noted several advantages over one-handed input when the two-handed interface is designed properly [7] [18] [27]. The most compelling advantage to using two hands in this case comes when the glove and the brush are used simultaneously to paint one stroke that contains a smooth transition between multiple colors. Typically, artists use this technique by starting a stroke with the non-dominant hand held steady at a certain color or size value, and then adjusting this value while painting with the dominant hand. Users typically focus on the brush rather than the color picker and develop a feel for how moving the non-dominant hand up or down causes a change in the lightness of the color and moving the hand in other directions causes changes in hue or saturation.

Pinching the ring finger and the thumb activates the translation mode, described in the navigation section below. This can also be used while painting with the dominant hand.

Pinching the pinky and the thumb activates a painting scaling mode, described below.

To help the user remember the function of each pinch, we present 3D text next to the appropriate finger on a 3D icon of the hand, whenever the user turns his or her hand so that the palm faces up.

4.6 Navigation

For small translations, a tracked pinch glove is used on the non-dominant hand. The user grabs the world by pinching together the thumb and the ring finger, and then drags the world around by moving the hand [1] [19]. This is often used for quickly positioning an object in the correct area to continue painting it, although, it can also be used, as mentioned above, to move the world while painting to create very long strokes.

For large translations, the user taps a foot pedal located on the floor of the Cave near the painting table. This places the system in translation mode and causes a wire-frame floor plane to be displayed in the virtual world. The artist then drips a trail of paint out of the brush and onto the floor plane. After painting out the trail to follow, the user controls the translation of the Cave forward and backward along the path painted in the virtual world using a lever attached to the painting table.

¹It is possible to change both the width and color of the stroke at the same time while painting, but it is difficult to control both parameters at the same time, with the same hand.

We have found that all of our users who choose to wear the pinch glove prefer to use it for both short and long translations. The way that our most comfortable artists work is to almost continuously hold the world with the non-dominant hand. They tend to apply a few strokes, step back and take a look at them, then grab the world and move it a few feet. Then they take another look at their work and keep proceeding in this fashion. The ability to quickly translate the world as they work is essential to these artists. Those that do not wear the pinch glove generally work on less serious compositions and tend to fill the cave with paint and then want to move to a new area to start something new. Dripping a path of paint to translate along is appropriate for these artists.



Figure 7 To rotate a painting, the user makes a circular gesture with the brush that activates this rotation widget.

To rotate the world, a simple gesture is used. The user points the tip of the brush down and makes a circular gesture, as if mixing a witch's brew in a big cauldron. After one full rotation around in a circle, a rotation widget (shown in Figure 7) pops up in the center of the circle. The widget is a compass. Moving the brush around the circumference of the compass causes both the compass and the entire world to rotate as if it were attached to the brush and rotating around the center of the compass.

4.7 Scaling

The user can activate a scaling mode by pinching together the thumb and pinky. When in scaling mode, the brush icon changes shape to indicate that the brush now controls the scale. To scale up, the user points the brush up and holds down the button on the brush. To scale down, the user does the same while pointing the brush down. The entire painting is scaled around the location of the brush. To exit scaling mode and continue painting, the user pinches the thumb and pinky again.

4.8 Interactive Viewing Mode

CavePainting records the state of the artwork as it progresses. Once the work is finished, an interactive viewing mode provides the observer with an interface for accessing this data. When it is launched, the interactive viewer first recreates the painting by moving the virtual brush icon according to the saved data produced by the artist. This creates an exact reproduction of the artist's brush work, and the final result is the painting that the artist produced. After this, the observer is in control of the display. A timeline widget (shown in Figure 8) is displayed on the floor of the Cave. When the user stands on the timeline, the position on the line determines the state of the painting to display. The left end of the timeline is the initial state,



Figure 8 By moving around on a timeline widget placed on the floor of the Cave, an observer controls the state of the 3D painting currently displayed to see the artistic process that created the painting.

which is a blank canvas. The right end of the timeline is the finished state of the painting. When the user steps off of the timeline widget, the painting is held at the current state, so the user can walk around it in the Cave and examine the details of the brushwork.

5 Results



Figure 9 *Florentine Vineyard* - Daniel Keefe.

Figures 9 through 12 are images of 3D scenes created with our system. Note that these are snapshots of actual 3D models which we prefer to present in an immersive viewing system.

6 Discussion

6.1 User Feedback

We have hosted two art classes at the Cave. One was a Brown University class, and one was a Rhode Island School of Design (RISD) class. In addition, we have worked with several painters and art students of varied backgrounds.

Painters, who used the system in its final state, expressed their surprise at how easy the system was to understand and use. Two of these artists had minimal experience with computers and had no trouble learning the prop and gesture based interface. The overwhelming comments from the artists were about their excitement regarding the new things they could do with this medium. All of them asked if they could come back again after their initial experiences because they wanted to return with sketches of ideas for serious compositions.




Figure 10 *Number 3* - Eileen Vote.



Figure 11 *Sailing* - Daniel Keefe.

While working in the Cave, users said that they liked the control they felt when picking colors and the responsiveness of the painting process. They also said that movement is essential in this new medium, and they liked that the painting was created by moving the entire body. They enjoyed watching the recreation of their work in the interactive viewing mode and then walking through their painting. Many of the artists commented that virtual sculpture may be a better metaphor for the experience, although they noted that the looseness of the line quality is like drawing or painting.

The artists did find it hard to be as accurate as they wanted when lining up the brush to paint fine details or paint on top of another stroke, although the ability to scale a painting helps with these tasks. Another comment was that the strokes are fairly final once they are placed. The user can backup and undo strokes in order, but unlike oil painting, there is no way to move paint around or scrape it off the canvas once it has been placed. One artist commented that she could not reproduce the earth tones with which she usually painted, since the projectors in the Cave could not faithfully reproduce all of these colors.



Figure 12 *Madman in a car with really bad exhaust* - Leora Maltz.

6.2 Young Artists

A group of 21 ninth-grade students also came to the Cave to try the program. None of them had ever been in virtual reality before. The interesting observation about this experience was that all of these novice users immediately understood the paint brush interface. Since there is only one button on the brush, there was no confusion about how to begin drawing a stroke. Dipping the brush into paint cups was also immediately understood as the way to change brush strokes. In fact, these young artists found the props so convincing that several of them took extra care moving the brush back and forth in the bottom of cup, as if to make sure the brush was fully loaded with paint.

6.3 Table Interface

We adopted a table interface for our props in the Cave for several reasons. The first is that if we place the table in one corner of the Cave, we can light it individually and cast a minimal amount of light on the projection surfaces. This is important when working with props in a Cave because the room lights must generally be kept very low so that there is adequate contrast for the projector screens. In a previous version of our interface, we used a different paint brush for each stroke. The user carried these brushes around as he or she moved in the Cave by attaching them with velcro to an apron. When we brought art students who were unfamiliar with virtual reality to the Cave to try our system, they thought holding a brush to create the strokes felt very natural, but their biggest complaint was that they had a hard time picking up the correct brush from the apron since it can be very difficult to see the props in the Cave when the lights are turned down and the user is wearing shutter glasses.

Another advantage to using our table based prop interface is that it scales much better than our previous interface to the different ages and abilities of our users. In targeting a range of users, we had hoped for an interface in which the most inexperienced user could walk into the Cave, pick up a paintbrush, and start using the system right away. Our original interface required putting on two pinch gloves and an apron. It was intimidating and restricting, especially to users who had never experienced virtual reality before, to have to put on these devices, all of which had wires coming out of them, before even stepping into the Cave. In the current interface, a novice user begins by picking up the paintbrush and starting to paint with whatever stroke is selected. As users become comfortable, we introduce the table props, and then the color picker and navigation controls.

We chose to put brush color and size controls on the pinch glove because we noticed that artists use these options frequently. They

tend to adjust the color and size of a stroke much more than they pick a different type of stroke. It would probably be distracting to have to walk to the table to perform a task that occurs as often as picking a color, but even artists working on serious compositions did not mind taking a step or two to get to the painting table to switch brush type. Given this, we chose to use gestures and the pinch glove to activate controls that are used very frequently, and the table for less frequent operations. We think that requiring the user to walk over to a physical table to perform these less frequent operations actually offers several enhancements to a user's Cave experience.

6.4 Using the Space of the Cave

For an application where the painting results are much more dramatic when the artist uses the entire space of the Cave, it was important to make even novice users understand that they could walk around the Cave, and to have them feel comfortable doing so. Our interface was designed to do both of these things. Users were required to interact with props placed along the walls of the Cave and virtual controls that were stationed on the floor surface. One advantage to doing this is that the user is able to take advantage of spatial memory to remember where all his or her tools are in the world. Just as when we enter a room, we remember where the light switch is, when the user enters the Cave, it is easy to remember which wall to walk to for a particular control panel. Additionally, if that control panel happens to be a physical table or prop attached to a wall, interacting with this physical object creates a link between the physical space of the Cave and the virtual world. Retaining some link to the real world while in the Cave seemed to have a positive effect on our users, especially those that were new to virtual reality. We think a large part of this could be attributed to their increased comfort level while working in the Cave. For an artist who is not too familiar with computers in general, walking into a dark room, called the Cave, that is filled with computers, wires, and technology you have never seen before, can be a rather intimidating experience by itself. Something as simple as reducing the fear of bumping into a wall and falling through an expensive projection surface by providing some indication of where the walls of the Cave are can make a big difference in making a user comfortable enough in the Cave to really walk around and use the whole space of the Cave.

6.5 Viewing Mode

We were interested in a special viewing mode for two reasons. First, because our computer-based medium has great potential for archiving the artist's process. Second, our observations and discussions with artists lead us to consider the process of creating a CavePainting as part of the finished artistic result. That is, a CavePainting is an artistic happening. Viewing a CavePainting without some sense of the dynamic movement of the brush and body that created it is only looking at part of the art.

The art world is often extremely interested in the process used to create a work. Unfortunately, this is usually very difficult to determine by looking at a finished work. For example, a skilled painter or art historian can often tell the order in which portions of a painting were created, but it is impossible to completely peel back each brush stroke of a masterpiece to see what lies beneath. There are a few examples of successful artworks in which the artist gives us a rare glimpse into the artistic process. For example, Pablo Picasso's series of lithographs of a bull [3]. In this series, Picasso made 18 prints of 11 different states of his lithography stone. The first states are detailed realistic representations of a bull. As we look at the next states in order, we see the image on Picasso's lithography stone progress from an intricate, realistic image of a bull to an abstract line representation. Printmakers find this work particularly interesting

because each of the 18 prints is a record of the state of Picasso's stone. 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 a way that anyone who goes to look at the art can understand.

One of the great artistic questions that surrounds working in virtual reality as a medium is how to display one's artwork. Some artists choose to display their art in VRML files on the web. Others prefer to capture 3D scenes at different angles and arrange these snapshots into a print that is more appropriate for display in a traditional museum or exhibition. While it is possible to do either of these with CavePainting, we realized while working with it, that a great deal of the art is lost when it is viewed in this manner. Walking through an artist's work, and watching the dance-like process that the artist used to create it, is really essential to getting a feel for the depth and dynamic nature of the scene and the new way that art is created around an artist with this way of working. The interactive viewing mode attempts to provide some of this sense to the viewer. Currently, we succeed in capturing and presenting the dynamic movements of the brush that were used to create a painting. These movements by themselves are very dance-like, and of course, they hint at the actual pose of the artist while painting. In the future, we would like to be able to capture much more information about the pose and movement of the artist's entire body. Perhaps, with a system similar to [20].

7 Conclusion and Future Work

Painting in its traditional form almost always produces a 2D result. Paint strokes can have some 3D qualities since they may be laid down with thick blobs of paint and layered on top of each other, but traditionally, they are arranged on a 2D surface. Despite this, we chose painting as a model for our new way of creating art in the Cave. The reason is that virtual reality has given us the power to create 3D strokes that mirror the function of a brush stroke in a 2D painting. A single brush stroke is a tiny element in a completed painting. Yet, when we look at the painting, these strokes are able to work together to form a coherent scene. One of the ways that the painter produces a meaningful scene from a collection of brush strokes is by varying the size, shape, color, and texture of the stroke. Strokes with different characteristics evoke different responses, and the skilled artist uses this knowledge to create a work which clearly conveys a message to the observer.

Our 3D brush strokes function in the same way as 2D brush strokes. By providing the artist with a wide range of choices in the type of brush stroke to apply as well as complete control over the colors and size of the stroke, we present the artist with numerous ways to vary brush strokes within a work. The resulting strokes may be used together and layered in three dimensions in a way which is analogous to a 2D painting.

We are interested in continuing to advance the look and behavior of our paint strokes by adding textures to them and creating strokes that can produce animated paintings. Another area of future work is developing an immersive viewing environment that provides at least some coarse tracking, but is still feasible for setup in a museum or an exhibition. Perhaps the most interesting area to explore in the future is to embrace the idea of CavePainting as a performance art. We have started to experiment with painting to music and bringing dancers into the Cave.

Our observations have shown that CavePainting provides an intuitive interface for the skilled artist working to create a meaningful piece of art as well as for the novice user or young artist. In addition, we have found CavePainting to be a very successful Cave

application in terms of enticing a novice user to move around freely and explore an environment in the Cave. As such, it is an excellent application for introducing people to the world of virtual reality. The brush strokes provided by the system are sufficient for experienced artists to create complex and meaningful 3D works using the system. CavePainting allows these artists to create a type of art that was not possible before and present this art in a novel viewing environment.

8 Acknowledgments

Many thanks to the artists that have worked with us: Eileen Vote, Leora Maltz, and students and professors at the Rhode Island School of Design and Brown University. This work originated as a final project for an Interactive Computer Graphics class taught by John Hughes. We would like to thank him for his support. Also, thanks to Robert Zeleznik for helpful discussions and comments on the paper, and Andrew Forsberg for his invaluable help working in the Cave. This work was partially supported by NSF (CCR-0086065, CCR-9996209).

References

- [1] SmartScene™ is a product of Multigen, Inc. More information on SmartScene™ is available from MultiGen's website at <http://www.multigen.com>.
- [2] Maya. AliasWavefront. <http://www.aw.sgi.com/>.
- [3] *Picasso the Printmaker: Graphics from the Marina Picasso Collection*. Dallas Museum of Art, 1983.
- [4] Maneesh Agrawala, Andrew C. Beers, and Marc Levoy. 3d painting on scanned surfaces. *1995 Symposium on Interactive 3D Graphics*, pages 145–150, April 1995. ISBN 0-89791-736-7.
- [5] Bill Brody and Chris Hartman. Painting space with BLUI. *Conference Abstracts and Applications of SIGGRAPH 00*, page 242, August 2000.
- [6] Jeff Butterworth, Andrew Davidson, Stephen Hench, and T. Marc Olano. 3dm: A three dimensional modeler using a head-mounted display. *1992 Symposium on Interactive 3D Graphics*, 25(2):135–138, March 1992. ISBN 0-89791-467-8.
- [7] W. Buxton and B. A. Myers. A study in two-handed input. In *Proceedings of ACM CHI 86 Conference on Human Factors in Computing Systems*, pages 321–326, New York, NY, USA, April 1986. ACM Press.
- [8] Jonathan M. Cohen, John F. Hughes, and Robert C. Zeleznik. Harold: A world made of drawings. In *Proceedings of the First International Symposium on Non-Photorealistic Animation and Rendering*, pages 83–90, New York, NY, USA, May 2000. ACM Press.
- [9] Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the cave. *Proceedings of SIGGRAPH 93*, pages 135–142, August 1993. ISBN 0-201-58889-7. Held in Anaheim, California.
- [10] M. Deering. Holosketch: A virtual reality sketching/animation tool. *ACM Transactions on Computer-Human Interaction*, 2(3):220–238, 1995.
- [11] George W. Fitzmaurice. Graspable user interfaces. Ph.D. Thesis, University of Toronto, 1996.
- [12] George W. Fitzmaurice, Hiroshi Ishii, and William Buxton. Bricks: Laying the foundations for graspable user interfaces. In Irvin R. Katz, Robert Mack, Linn Marks, Mary Beth Rosson, and Jakob Nielsen, editors, *Proceedings of ACM CHI 95 Conference on Human Factors in Computing Systems*, pages 442–449, New York, NY, USA, May 1995. ACM Press.
- [13] Tinsley A. Galyean and John F. Hughes. Sculpting: An interactive volumetric modeling technique. *Proceedings of SIGGRAPH 91*, pages 267–274, July 1991. ISBN 0-201-56291-X. Held in Las Vegas, Nevada.
- [14] Pat Hanrahan and Paul E. Haeberli. Direct wysiwyg painting and texturing on 3d shapes. *Proceedings of SIGGRAPH 90*, pages 215–223, August 1990. ISBN 0-201-50933-4. Held in Dallas, Texas.
- [15] Ken Hinckley, Randy Pausch, John C. Goble, and Neal F. Kassell. Passive real-world interface props for neurosurgical visualization. In *Proceedings of ACM CHI 94 Conference on Human Factors in Computing Systems*, volume 2 of *PAPER ABSTRACTS: Interacting in 3-D*, page 232, 1994.
- [16] Takeo Igarashi, Satoshi Matsuoka, and Hidehiko Tanaka. Teddy: A sketching interface for 3d freeform design. *Proceedings of SIGGRAPH 99*, pages 409–416, August 1999. ISBN 0-20148-560-5. Held in Los Angeles, California.
- [17] Hiroshi Ishii and Brygg Ullmer. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proceedings of ACM CHI 97 Conference on Human Factors in Computing Systems*, volume 1 of *PAPERS: Beyond the Desktop*, pages 234–241, 1997.
- [18] Paul Kabbash, William Buxton, and Abigail Sellen. Two-handed input in a compound task. In Beth Adelson, Susan Dumais, and Judith Olson, editors, *Proceedings of ACM CHI 94 Conference on Human Factors in Computing Systems*, pages 417–423, New York, NY, USA, April 1994. ACM Press.
- [19] D. P. Mapes and J. M. Moshell. A two-handed interface for object manipulation in virtual environments. *Presence*, 4(4):403–416, 1995.
- [20] Simon Penny, Jeffrey Smith, and Andre Bernhardt. Traces: Wireless full body tracking in the cave. In *ICAT 99 Conference Proceedings*, 1999.
- [21] Emanuel Sachs, Andrew Roberts, and David Stoops. 3-draw: A tool for designing 3D shapes. *IEEE Computer Graphics and Applications*, 11(6):18–26, November 1991.
- [22] Irving Sandler. *Abstract Expressionism: The Triumph of American Painting*. Pall Mall P., London, 1970.
- [23] Stephen Schkolne. Surface drawing: The perceptual construction of aesthetic form. Master's Thesis, Caltech, May 1999. <http://www.cs.caltech.edu/ss/msthesis.html>.
- [24] Steven Schkolne and Peter Schroder. Surface drawing. Technical Report CS-TR-99-03, CalTech Department of Computer Science, 1999.
- [25] Sidney W. Wang and Arie E. Kaufman. Volume sculpting. *1995 Symposium on Interactive 3D Graphics*, pages 151–156, April 1995. ISBN 0-89791-736-7.
- [26] Barbara Ehrlich White. *Impressionists side by side: their relationships, rivalries, and artistic exchanges*. A.A. Knopf, New York, 1996.
- [27] Robert C. Zeleznik, Andrew S. Forsberg, and Paul S. Strauss. Two pointer input for 3d interaction. *1997 Symposium on Interactive 3D Graphics*, pages 115–120, April 1997. ISBN 0-89791-884-3.
- [28] Robert C. Zeleznik, Kenneth P. Herndon, and John F. Hughes. Sketch: An interface for sketching 3d scenes. *Proceedings of SIGGRAPH 96*, pages 163–170, August 1996. ISBN 0-201-94800-1. Held in New Orleans, Louisiana.