

Multimodal Exploration of the Fourth Dimension

Andrew J. Hanson*

Hui Zhang†

Computer Science Department, Indiana University




Figure 1: Multimodal perceptual exploration of the 3D projection of the spun trefoil, a knotted sphere embedded in 4D.

ABSTRACT

We present a multimodal paradigm for exploring topological surfaces embedded in four dimensions; we exploit haptic methods in particular to overcome the intrinsic limitations of 3D graphics images and 3D physical models. The basic problem is that, just as 2D shadows of 3D curves lose structure where lines cross, 3D graphics projections of smooth 4D topological surfaces are *interrupted* where one surface intersects another. Furthermore, if one attempts to trace *real* knotted ropes or a plastic models of self-intersecting surfaces with a fingertip, one inevitably collides with parts of the physical artifact. In this work, we exploit the free motion of a computer-based haptic probe to support a continuous motion that follows the *local continuity* of the object being explored. For our principal test case of 4D-embedded surfaces projected to 3D, this permits us to follow the full local continuity of the surface as though in fact we were touching an actual 4D object. We exploit additional sensory cues to provide supplementary or redundant information. For example, we can use audio tags to note the relative 4D depth of illusory 3D surface intersections produced by projection from 4D, as well as providing automated refinement of the tactile exploration path to eliminate jitter and snagging, resulting in a much cleaner exploratory motion than a bare uncorrected motion. Visual enhancements provide still further improvement to the feedback: by opening a view-direction-defined cutaway into the interior of the 3D surface projection, we allow the viewer to keep the haptic probe continuously in view as it traverses any touchable part of the object. Finally, we extend the static tactile exploration framework using a

dynamic mode that links each stylus motion to a change in orientation that creates at each instant a maximal-area screen projection of a neighborhood of the current point of interest. This minimizes 4D distortion and permits true metric sizes to be deduced locally at any point. All these methods combine to reveal the full richness of the complex spatial relationships of the target shapes, and to overcome many expected perceptual limitations in 4D visualization.

CR Categories: I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques; H.5.2 [Information Interface and Presentation]: User Interfaces—Haptic I/O; G.4 [Mathematics of computing]: Mathematical Software—User Interfaces; I.2.9 [Artificial Intelligence]: Robotics—Kinematics and Dynamics

Keywords: multimodal, haptics, visualization

1 INTRODUCTION

“Words and language, whether written or spoken, do not seem to play any part in my thought processes. The psychological entities that serve as building blocks for my thought are certain signs or images, more or less clear, that I can reproduce and recombine at will. The elements that I have mentioned are, in my case, visual and sometimes motor.”

— Albert Einstein, in a letter to mathematician Jacques Hadamard [10].

We now possess interactive graphics tools incorporating touch-responsive features that make it possible for a computer interface to implement and even to improve upon the multisensory mental paradigm described by Einstein. By exploiting such tools, we feel that we can make a non-trivial contribution to building intuition about classes of geometric problems whose intuitive, non-symbolic comprehension lies beyond the reach of the unaided human intellect. The challenge is to find specific examples where, in fact, we

*e-mail: hanson@cs.indiana.edu

†e-mail: huizhang@cs.indiana.edu

have problems that are simple enough to formulate clearly, yet complex enough that even Einstein would have had trouble working out the details in “images that [one] can reproduce and recombine at will.”

In this paper we take the first steps towards this goal by designing and implementing novel multimodal methods for exploring surfaces (two-manifolds) embedded in 4D Euclidean space.

2 PREVIOUS WORK ON 4D VISUALIZATION

The idea of cross-dimensional understanding has long been a subject of fascination, starting with *Flatland*'s conundrum of how two-dimensional creatures might attempt to understand three-dimensional space [1, 7]. Banchoff's pioneering work on the corresponding question of how 3D computer-based projections can be used to study 4D objects [2, 3] has been particularly influential. Other representative efforts include a variety of ways to render 4D objects (see, e.g., Noll [20], Hollasch [14], Banks [4], Roseman [25], and Egli, Petit, and Stewart [8]), and to extend lighting model techniques to 4D (see, e.g., [6, 28] and [12]).

Typically, 4D visualization methods employ a projection to 3D as a fundamental step; this helps the viewer to identify salient global features of the four-dimensional object, and provides structural continuity when rotating either the object's (rigid) 3D projection, or, with more difficulty, the 4D orientation matrix, which causes non-rigid deformations in the 3D projection. While the resulting imagery supplies a sense of the object's overall shape, structural continuity is often difficult to discern. What is really needed is an enhancement of the 3D visual representation that allows intuition-building exploration of the shape itself. Thus the question arises, “How can we touch the fourth dimension?”

3 MOTIVATION

People learn about the everyday world by combining sensory modalities, and knowledge of shape comes from a combination of sight, touch, and exploration. By combining computer graphics with computer haptics, which imitate the 3D sense of touch, we can provide multimodal exploration tools that can in principle improve on real life. This improvement is possible because, for example, the image of a knotted rope is interrupted where one part crosses another, and if we try to trace a *real* knotted rope with a fingertip, we will eventually collide with some part of the rope and have to interrupt our smooth passage. With a touch-based computer interface, there are no physical obstructions to the motion of the pictured “computer hand,” and the entire object can be traced without the interruptions imposed by sight or a real physical model. We can therefore use the same methods to help us understand and manipulate the much more complicated case of shapes with self-intersecting *surfaces* (which arise naturally when projecting 4D to 3D) using the touch-based multimodal paradigm. If sound cues are added to describe the passage of the computer hand across a visual obstruction, one can explore a knot or a surface without necessarily having to use vision at all; when used in combination with a visual representation, the auditory cues provide additional redundant feedback enabling improved intuitive perception of spatial structure.

3D Example. To create a visual representation of a 4D-embedded surface, we typically project the entire object to 3D, construct a standard 3D computer graphics representation of the result, and render that to a 2D screen image, possibly as a stereo pair. We may then, in effect, depend on the second order effects of our practical experience with 3D structures to reconstruct a 3D shape in our minds. Full 3D information can be obtained from stereography or motion parallax, or, given the resources, an actual physical model. However, while the physical model for the type of problem we are

considering can be very interesting, they may not be as useful as one might think. In Figure 2, for example, we see images of an actual surface corresponding to a particular cubic polynomial projected from 4D to 3D [11]; although the shape of the surface is completely without self-intersections in 4D, the model of the 3D projection has numerous very complex intersections, and, even with the physical model in hand, the self-intersections make it very difficult to trace and comprehend the intrinsic shape of the surface.




Figure 2: Two views of a physical model representing a highly self-intersecting surface, constructed by 3D projection from the four-dimensional mathematical description. It is nearly impossible to trace the shape continuously and with clear comprehension even when you are holding the physical object in your hands. (Model courtesy of Stewart Dickson.)

We are thus motivated to consider how 4D shapes projected to 3D might be explored to advantage with a 3D haptic probe. Figure 3(a) shows a pair of intersecting surfaces that apparently pass right through one another, so that we would naturally think of a path on the green surface as blocked by the brown “wall.” However, if these two surface patches were distinct in 4D and the 3D intersection only an artifact of projection, one surface might be “above” and the other “below;” using either an analogy to the cutaway method used to draw a knot on a blackboard, or a corresponding analogy to depth-buffered rendering, we could erase the clashing section of the brown wall as in Figure 3(b), and continue our path along the 4D green surface without interruption. Another alternative is to add an additional visual cue by assigning a surface color keyed to 4D depth relative to the projection center, as in Figure 3(c).

Finally, as suggested in Figure 3(d) as well as Figure 1, we can explore a more complex continuous 4D surface using a 3D haptic probe that avoids all physical entanglement: this is the main innovation that we shall present in this paper — a 3D probe that respects a surface's local 4D continuity despite the self-intersections of its 3D projection. By creating auditory cues that mark crossing events we achieve a full multisensory modality that can be used with or without supplementary visual representations.

Overriding Conflicting Evidence. The key ideas of the overall scenario should now be clear. The logical series of modeling steps, the problems they induce, and the ultimate resolution of the problems are as follows:

- *Create a model of a smoothly embedded object.* Examples are knotted curves embedded in 3D and knotted surfaces embedded in 4D, though there are obviously many other shapes that could be used.
- *Project to one lower dimension.* Various parts of the shape appear to touch each other when the smooth (non-intersecting) original shape is projected. What is seen is essentially the $(N - 1)$ -dimensional “shadow” of the original N -dimensional object.




Figure 3: (a) Two intersecting surface segments, typically resulting from the projection of a 4D embedded object. (b) Crossing diagram of the surfaces corresponding to the green surface being closer to the 4D projection point than the brown surface. (c) A continuous 4D surface, where the pieces of the single surface do not touch in reality, but appear to intersect in the projection shown; depth is color-coded to show the absence of intersection. (d) Schematic summary of a tactile path with auditory cues enabling the meaningful 3D exploration of the spun trefoil, a knotted sphere embedded in 4D.

- The “shadow” object is visually discontinuous. The projected image has self-intersections.
- Tracing the “shadow” object is discontinuous. If you attach a haptic probe to the projection or “shadow” object, you must detach the probe wherever self-intersections occur, pull away, and re-attach on the other side of the intersection in order to trace the logically continuous surface. Alternatively, you can allow the probe to stay attached at the intersections and choose a direction, but this does not maintain consistent local continuity.
- Create a crossing-diagram object. By performing a depth-buffer operation in $(N - 1)$ dimensions, you can do slightly better than the shadow; the part of the self-intersecting collision that is “in front” (nearer the projection point) can be made continuous (as in a depth-buffered rendering), and the part that is farther from the projection point can be interrupted.
- Tracing the crossing-diagram object is discontinuous. A haptic probe attached to the crossing-diagram object is continuous as long as we cross the self-intersection on the near segment, but once again we hit a discontinuity when we come around from the other direction and try to cross at the far segment; at this point, one again must detach the probe and move away to find the matching piece of the continuous object.
- Solution: Haptically override the conflicting evidence. To answer the question “How do we touch the fourth dimension?” we *override* the apparent visual collisions, conflicts, and gaps in the rendered image of the projection, and keep the haptic probe anchored to the *higher-dimensional* continuity that underlies the whole structure, regardless of whether it is above or below another conflicting part relative to the projection point.
- Exploit and manipulate the “phantom effect” to advantage. As noted, e.g., by Massie and Salisbury [17], users may initially find it disturbing to have a single point on a haptic probe constrained to a surface, while the rest of the virtual hand is passing, ghost-like, through other parts of the object. Our paradigm focuses on learning to *take advantage* of the lack of physical obstruction, which, when present, makes even the most elaborate physical model unexplorable. In addition, we have experimented with one additional feature, a *dynamic cut-away* algorithm that opens a path from the viewpoint to the haptic probe contact point, thus making the tip of the haptic probe visible to the eye even when buried behind additional layers of the surface.

In summary, we facilitate perception of the intuitive structure of an object such as a 4D-embedded surface projected to 3D by emphasizing the *contrast* between apparent discontinuities in the *visual* representation; the interface enforces continuity in the haptic representation that emphasizes the true topology, even if this conflicts with what we see, while at the same time keeping visual contact with what is being touched as much as possible.

Overview of Interface Elements. Our interface design includes a variety of methods for providing redundant representation of 4D shape information, including the following:

- Local continuity of the haptic freedom of movement throughout the entire higher dimensional shape; visual conflicts that are artifacts of the projection have no haptic effects.
- Intelligent force-guided mechanisms for assisting the user in the haptic exploration. This assists the user in concentrating on navigating the shape itself, rather than being distracted by snags and discontinuities in the haptic path.
- 4D visual cues, including crossing diagram depth ordering methods and depth color coding.
- Attention-driven cutaways to “see through” occluding layers intersecting the line of sight between the viewpoint and the probe contact point.
- “Rolling manifold” methods that enable the user to optimize the local projection of the current local segment of the shape being touched by the cursor, providing accurate local metric information in the screen plane.
- Auditory cues that reinforce the occlusion information available from visual cues, and can potentially enable the exploration of an entire space without depending on visual representations.

4 HAPTIC METHODS

Haptic interfaces can be effectively exploited to improve the sense of realism and to enhance the manipulation of virtual objects (see, e.g., [5]). The haptic exploration of unknown objects by robotic fingers (see, e.g., [21, 22], [16] and [30]) is another variant that has requirements similar to ours. While we of course exploit many techniques of force-feedback and haptic user assistance that have been widely used in other interfaces (see, e.g., the work of [31],

[15], [9], [23], and [24]), we have found that many of the problems we encounter have been fairly unique, and thus have required customized hybrid approaches.

In this section, we introduce a “predictive navigational assistance” model for traversing a 4D polygonal object using a virtual proxy. To be effective, the navigation assistance must meet two goals: on the one hand, it must constrain the probe to the surface of the 4D object to give the impression of a “magnetic object” to which the haptic probe is drawn; on the other hand, the force should be applied in a way that facilitates exploring the local continuity of 4D objects such as knotted spheres, whose visual pictures are *interrupted* by massive self-intersections.

4.1 Simulated Sticky Stylus

Our basic force model implements a “sticky” stylus using a conventional damped-spring force to attach the stylus to the object surface. The proxy, a graphics point that closely follows the position of the haptic device, is constrained to the surface in question. The haptic rendering engine [26] continually updates the position of the proxy, attempting to move it to match the haptic device position, and applying the damped-spring force between the haptic device position and the proxy position.

We use the following model to compute the force in the normal direction constraining the 3D cursor or proxy to the neighborhood of the object surface,

$$\|\mathbf{f}_n\| = \alpha r^{1+\beta}, \quad (1)$$

where α is a constant and $\beta = 0$ for the standard case of an ideal (linear) spring. No force is applied in the direction tangent to the surface, thus allowing free motion and facilitating exploration of the structure. The damping force

$$\mathbf{f}_d = -K_d \mathbf{V}, \quad (2)$$

where V is the radial velocity, is used to smooth the force feedback.

4.2 Constrained Navigation on the Local Surface Model

The force model presented so far in fact extends trivially to arbitrary dimensions of the vertex coordinates. The only essential difference is that in a 4D projection to 3D, each vertex has a 4D “eye-coordinate,” or depth w , in addition to the coordinates (x, y, z) of the 3D projection. To constrain the proxy to a continuous surface embedded in 4D, we introduce local submodels of the geometry that cover a neighborhood of the current contact point in the 3D projection and serve as the current active haptic contact domain.

The following steps describe the haptic servo loop model:

1. *Get coordinates.* Let the current proxy coordinate data be denoted as $(x_{\text{cur}}, y_{\text{cur}}, z_{\text{cur}}, w_{\text{cur}})$, and the previous proxy coordinate as $(x_{\text{pre}}, y_{\text{pre}}, z_{\text{pre}}, w_{\text{pre}})$.
2. *Estimate local continuity.* In order to construct a continuous local structure to guide the user’s movements, we first compute the depth range for a single local update period as

$$\Delta_w = w_{\text{cur}} - w_{\text{pre}}. \quad (3)$$

The estimated local continuity of 4D depth is then given by:

$$w_{\text{cur}} - \Delta_w \leq w \leq w_{\text{cur}} + \Delta_w \quad (4)$$

3. *Construct local model.* The search for candidate facets begins at the current proxy location and ends when Eq. (4) no longer holds. These facets form the local model for computing forces. An alternate implementation could directly exploit mesh continuity in parameter space, while the current method uses only 4D proximity.

4. *Haptically override the conflicting evidence.* Now the sticky stylus can be applied trivially to the local model, as shown in Figure 4. The virtual proxy is constrained to a locally continuous domain to fully exploit the model’s topological information, and haptic operations on this local model are fast and efficient.

5. *Update.* Set $(x_{\text{pre}}, y_{\text{pre}}, z_{\text{pre}}, w_{\text{pre}})$ to the values of $(x_{\text{cur}}, y_{\text{cur}}, z_{\text{cur}}, w_{\text{cur}})$.




Figure 4: Despite apparent surface conflicts, only the local model with 4D continuity constraints is actively involved in the haptic process.

This rendering rule allows us to slide through the visual intersections as though they were ghost images while constraining motion to the smooth surface that exists in 4D but cannot be seen without interruption in 3D. As shown in Figures 5(a) and (b), we can feel the four-dimensional continuity even though it is not apparent in the visual representation. Figure 5(c) depicts the continuous depth-encoded 4D path on the 3D projection with the shaded surface removed. The path inherits the 4D depth cues from the original 4D shape and reveals the fact that the user is navigating the true continuous underlying surface and exploring a “feelable” structure.




Figure 5: (a) This 4D object contains massive self-intersections in its 3D projection. (b) The proxy slides through the visual intersections and explores the continuous 4D structure. (c) The projected 3D path with color-coded 4D depth is exposed in wireframe mode.

4.3 Adding Force Suggestions

The sticky stylus *in principle* is sufficient to allow the viewer to explore the full continuity of a given surface. However, in practice, many interesting geometric objects such as 4D knotted surfaces produce complex images that snag the stylus in sharp bends and obscure the structural continuity and local surface features (see, e.g., Figure 2). This presents a significant challenge to the user if only the bare attractive force interface is supplied. Our purpose in this section is to present a predictive navigation interface that assists and guides the user towards the neighboring surface features without depending on visual feedback concerning the direction the hand should move.

Assuming we know the desired position and desired velocity for the stylus, we propose the following supplementary force

$$\mathbf{f}_{\text{pd}} = K_p (\mathbf{P}_{\text{des}} - \mathbf{P}) - K_v (\mathbf{V}_{\text{des}} - \mathbf{V}) \quad (5)$$

where \mathbf{f}_{pd} is the force due to the predictive navigation assistance, \mathbf{P} and \mathbf{V} are the actual (current) contact point positions and velocities, and \mathbf{P}_{des} and \mathbf{V}_{des} are the desired contact point positions specified by the user if the default coordinates are to be overridden. K_p and K_v serve as stiffness and damping matrices. Starting from this point, we next investigate how Kalman Filters can be used as *predictors* to assist navigation control.

▷ **Noise Analysis.** The position and velocity data obtained from haptic sensors are typically noisy due to both the motion of the human hand and the construction of the sensors themselves, as illustrated by the data in Figure 6; the measured velocity data (blue) are obtained by sliding the probe over a curved surface. The noise present in the measured data makes position and velocity prediction difficult.

▷ **Kalman Filtering.** In order to predict the desired position and velocity from noisy data, we need to apply an appropriate filter. A logical candidate is the Kalman Filter (KF), which is essentially a recursive solution of the least-squares problem [29, 18]. Among applications similar to ours, we note Simon’s [27] application of Kalman filtering to vehicle navigation, and Negenborn’s [19] study of the robot localization problem.

The KF and discrete stochastic dynamical equations for our system can be written as

State equation: $\hat{\mathbf{x}}_{k+1} = \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix} \hat{\mathbf{x}}_k + \begin{bmatrix} T^2/2 \\ T \end{bmatrix} \mathbf{u}_k + \mathbf{w}_k$

Output equation: $\hat{\mathbf{y}}_k = \begin{bmatrix} 1 & 0 \end{bmatrix} \hat{\mathbf{x}}_k + \mathbf{z}_k$

where $\hat{\mathbf{x}}_k = \begin{bmatrix} \mathbf{P}_k \\ \mathbf{V}_k \end{bmatrix}$, and \mathbf{u}_k is the random, time-varying acceleration and T is the time between step k and step $k+1$. We assume that the process noise \mathbf{w}_k is white Gaussian noise with noise covariance $\mathbf{S}_w = \mathbf{E}(\mathbf{w}_k \mathbf{w}_k^T)$, and the measurement noise \mathbf{z}_k is white Gaussian noise with noise covariance matrix $\mathbf{S}_z = \mathbf{E}(\mathbf{z}_k \mathbf{z}_k^T)$, and that it is not correlated with the process noise. \mathbf{Q} is called the estimation error covariance, which is initialized as \mathbf{S}_w . The KF estimation is then formulated as follows:

Gain: $\mathbf{K}_k = \mathbf{A}\mathbf{Q}_k\mathbf{C}^T(\mathbf{C}\mathbf{Q}_k\mathbf{C}^T + \mathbf{S}_z)^{-1}$

Prediction: $\hat{\mathbf{x}}_{k+1} = (\mathbf{A}\hat{\mathbf{x}}_k + \mathbf{B}\mathbf{u}_k) + \mathbf{K}_k(\hat{\mathbf{y}}_{k+1} - \mathbf{C}\hat{\mathbf{x}}_k)$

Update: $\mathbf{Q}_{k+1} = \mathbf{A}\mathbf{Q}_k\mathbf{A}^T + \mathbf{S}_w - \mathbf{A}\mathbf{Q}_k\mathbf{C}^T\mathbf{S}_z^{-1}\mathbf{C}\mathbf{Q}_k\mathbf{A}^T$

$(\mathbf{A} = \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix}; \mathbf{B} = \begin{bmatrix} T^2/2 \\ T \end{bmatrix}; \mathbf{C} = \begin{bmatrix} 1 & 0 \end{bmatrix})$

The KF provides optimal (minimum variance, unbiased) estimation of the state $\hat{\mathbf{x}}_{k+1}$ with the given observed data. In Figure 6, we have plotted 30 time units of true velocity (red), measured velocity (blue), and the estimated velocity (green). We notice that the measured data introduce significant jitter (human and mechanical), and thus cause difficulty determining the desired position and velocity for predictive force rendering. We see explicitly that the noisy observed signal can be improved using the Kalman Filter to accurately predict the velocity for the haptic stylus (green).

▷ **Steering towards desired position** With the desired position \mathbf{P}_{k+1} and velocity \mathbf{V}_{k+1} , the force suggestion can be computed using Equation 5 at each servo loop. This force is then projected to local facet of the contact point, and applied to the haptic probe.

4.4 Experiments and Results

The proposed predictive “power assisted” force model was tested on a 2.8 GHz Pentium 4 PC running Windows XP with NVIDIA’s

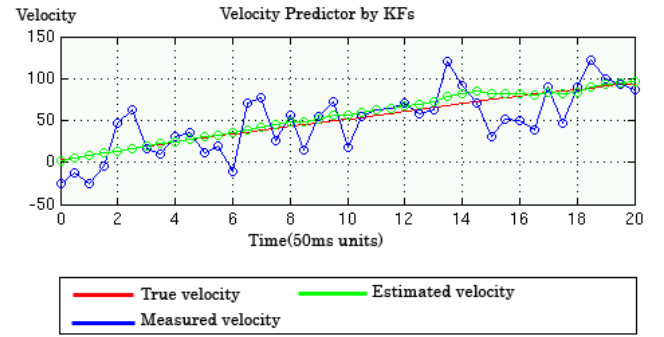


Figure 6: The KF helps to predict the desired velocity in the presence of noise. The desired position is calculated based on the desired velocity and provides an accurate prediction of the user’s behavior.

GeForce FX 5200 graphics card. The graphics process routines were updated at about 20 Hz, whereas the haptics process ran at a 1 kHz update rate. As illustrated in Figure 1 and Figure 5, the motion conforms to the local continuity constraints so that the user’s probe effectively follows the shape of the projected 4D object.

Next, a series of tests were carried out to verify the effectiveness of our proposed method for constructing the local model. Haptic rendering of large and detailed 4D objects requires a significant computational load; by choosing to render only a local subset of the polygons in the haptic process, we can significantly reduce this load.

Table 1 illustrates the typical saving in polygon count that we are able to achieve without any degradation of the user’s perception of the interface; the local model strategy thus improves the system response at no observable cost.

Table 1: Local Model vs Overall Model

4D Model	Polygon Number in Overall Model	Polygon Number in Local Model
4-Torus	900	≈ 220
Spun Trefoil Knot	2800	≈ 600
Steiner’s Roman Surface	1600	≈ 360
Twist Spun Knot	2800	≈ 600

On the other hand, our algorithm exploits Kalman Filtering for more accurate estimations of a *noisy dynamic system*. The KF estimates the state of a noisy system using noisy measurements. In this way, our algorithm significantly reduces the effect of noise in obtaining clean path prediction, and this facilitates smooth navigation on the surface.

5 AUDITORY CUES

The addition of sound cues to an environment is well-known to improve presence and realism (see, e.g., [31]). Audible annotation of a user’s exploration of a 4D object can be used to coordinate and accent particular events. Creating auditory “advice” to supplement the haptic exploration is significant because it makes it possible for users to build mental models of 4D geometry without necessarily depending on sight, though the visual feedback can be very useful as a redundant information source. We remark that using the multimodal interface with one’s eyes closed creates an intense experience that could conceivably improve one’s internal mental model because, paradoxically, it removes potentially confounding visual distractors.

5.1 The Problem

Traditional techniques for perceiving visual pictures of 4D surfaces intersecting in a 3D projection rely on associating 4D depth with visual cues such as color coding, texture density, etc. Images are undoubtedly important for understanding complex spatial relationships and structures, but they may also serve to confuse the user when the visual evidence is difficult to interpret. There are currently only limited methods for presenting information non-visually, e.g., haptics and 3D hardcopy methods such as those used to produce the models in Figure 2. Physical models obstruct efforts at tactile exploration of continuous but self-intersecting surfaces by definition, leaving us with haptic methods; however, haptics methods alone do not necessarily convey the *context* of an exploration such as the location of a crossing or intersection. (Note that in 3D knot theory, the marking of over/under crossings in the 2D projection is an essential step.) While the location of a crossing can be flagged by the visual interruption itself, or by adding a slight jitter or jump to the haptic interface at the transition point, extending the modality to include auditory cues for these events is perhaps more versatile, like the narration of a “teacher” who is guiding us through a touch-based world. We are thus led to exploit sound, using word labels such as “over” and “under” to perceptually mark those points at which the continuous topological motion of the haptic device passes a visual disruption in the graphics image, as illustrated in Figure 7(a).




Figure 7: (a) With visual feedback, the user is acutely aware of sliding through visual interruptions caused by the 3D projection. Without using visuals, however, one would be totally unaware of encountering a significant interruption unless supplementary cues such as sound feedback are supplied. (b) Sound cues supplement or replace visual cues to assist in building a clear mental model of a 4D surface.

5.2 The Solution

We can thus choose to supplement our design by allowing the user to explore the 4D surfaces visually and haptically, but accompanied in addition by appropriate sounds triggered, e.g., by passing the locations of illusory 3D surface intersections. Using these multiple sensory modalities to present information can overcome contradictions in the visual feedback, or even make the visual feedback superfluous.

Example. When performing a haptic exploration revealing local 4D continuity, users can feel the shape of the object via the haptic interface as described above. They can trace out the shape of a path with their computer-idealized finger, hear, for example, a change in pitch to indicate 4D depth, and have specific auditory token sounds or spoken narration triggered at special locations such as visual obstruction transitions. A schematic picture for the audition-based supplement to the mental-model building exploratory interface is shown in Figure 7(b).

As the user interacts with the 4D object in its 3D projection, the visual, haptic, and auditory stimuli seamlessly contribute to a new mental picture corresponding to 4D multimodal visualization. Thus our paradigm provides an interface that assimilates the three sensory modalities of our perceptual model and makes them all work together to achieve the goal of perceiving the nature of 4D shape.

6 4D HAPTIC ROLLING MANIFOLD

There are some tasks that can be performed better by automated manipulation. For example, we can explore all points of a ball by rolling it around until the desired point comes up to face us; as advocated by Hanson and Ma[13], the presentation of arbitrary curves and surfaces to the user’s view can be automatically optimized by “walking” along a local (typically geodesic) path. With the aid of the haptic probe and a 3D projection, we can adapt this family of methods to produce not only a maximal *viewable* but also a maximal *touchable* local region with each step of the walk; this permits the user with a haptic stylus to continuously sense a maximally presented portion of a neighborhood around the central point, and thus to get a richer tactile sense of the surface shape.

6.1 Maximizing Viewable and Touchable Aspects at Each Step

Consider a surface represented by quadrilateral facets in ordinary 3D space. Our basic maximal projection procedure then requires that the center of the facet of interest lie completely in the screen plane, thus facilitating both viewing and touching. As shown in the Figure 8(a), we begin by choosing a point of interest P in the facet, and then use an algorithm such as shortest-path motion to transform the relation between the object and the viewpoint to maximize viewable area as we move to a new interest point P' . If P' lies in the current face, no rotation is needed. If P' lies in an adjacent face, we rotate as prescribed below and then translate P' to the screen center. Translation and rotation are applied in tandem in the haptics and graphics rendering threads.




Figure 8: (a) Surface facet with interest point P , maximally projected onto the 2D screen to facilitate viewing and touching. (b) 3D facet rotation.

6.2 Rotating Between Facets

We now compute the transition across a given edge between two facets, assuming we know the unit vectors T (the probe’s current incremental motion vector), and T' , the projection of T straight down onto the target facet.

With $\hat{N} = T \times T' / |T \times T'|$ giving the rotation axis, and $\cos \theta = T \cdot T'$ giving the rotation angle, the axis-angle rotation matrix $\text{Rotate}(\theta, \hat{N})$ re-oriens the user to match the projected geodesic direction, as shown in Figure 8(b).

6.3 Tunneling the Probe to Target the Surface

During the rolling process, visual perception of the interest point can be further assisted by opening a visual window into the interior of the 3D projection; this permits the viewer to keep the haptic probe continuously in view as it traverses any touchable part of the object, thus enabling both *touching* and *seeing* the true continuous topology.

6.4 Fixing the Stylus Using a “Rubber Band Line”

Since, in the paradigm we are describing, the facet of interest is fixed to the screen center both graphically and haptically, one would in principle imagine fixing the haptic probe to the screen center where the interest facet is projected; however, this is not necessarily a useful interface choice, since the user gets no tactile feedback. Instead, we implement a supplementary interface feature based on a simulated “rubber band line” from the current haptic position to the screen center, and then compute a haptic force rendered using Eq. (1) with $\beta = 3$, and r the distance from proxy position to screen center. This pulls the displaced haptic probe gently back to the screen center (like a weak rubber band tether), and attaches it to the newly projected interest facet. In this way, we can keep the facet of interest maximally projected, and yet allow the user to feel a limited continuous neighborhood of the central focus of attention, thus getting a tactile sense of the surface shape. Figure 9 illustrates a 4D “rolling manifold” with the proxy constrained to the maximal viewable and touchable region with each step of the walk.

7 USER ENVIRONMENT AND FURTHER EXAMPLES

Our implementation is based on a standard OpenGL graphics system with a high-performance graphics card and sound card, combined with SensAble Technology’s Omni PHANTOM force-feedback haptic device. Our user interface is based on OpenGL, SensAble’s OpenHaptics toolkit, and a locally customized GLUT API. Various options provide haptic forces and selected auditory signals to enhance the interactive feedback. Figure 10 shows a representative view of the interface.

The user can load or create a variety of geometric objects for investigation, and is also presented with a wide selection of options to help maintain an optimal presentation of the object being studied.

Figure 11 shows the user exploring a 2-torus (technically the two-manifold $S^1 \times S^1$) embedded in Euclidean 4-space and projected orthographically to 3D. We can see that the user can effectively slide through the visual interruption and *walk* on the continuous 4D structure.

In Figure 12, we show a geodesic path followed by the haptics-based walking algorithm on an ordinary torus embedded in 3D (which happens to coincide with a particular perspective projection of the 2-torus from 4D).

8 CONCLUSION AND FUTURE WORK

Current computer interfaces can support multimodal representations that integrate visual information with haptic feedback and interaction. Exploiting these capabilities permits us to build a particular type of kinesthetic intuition about continuous complex geometric surfaces, and is especially suited to the families of such surfaces that result naturally from projecting 4D-embedded surfaces to 3D. In accordance with the best visualization science practice, different features can be applied to the overall images, using attention-driven cutaway algorithms to reduce the cognitive load still further.

In this paper, we discuss a set of methods to *touch* and *see* the geometric structures with a combination of haptic and visual sensing. We can begin to imagine the way that Einstein might have thought

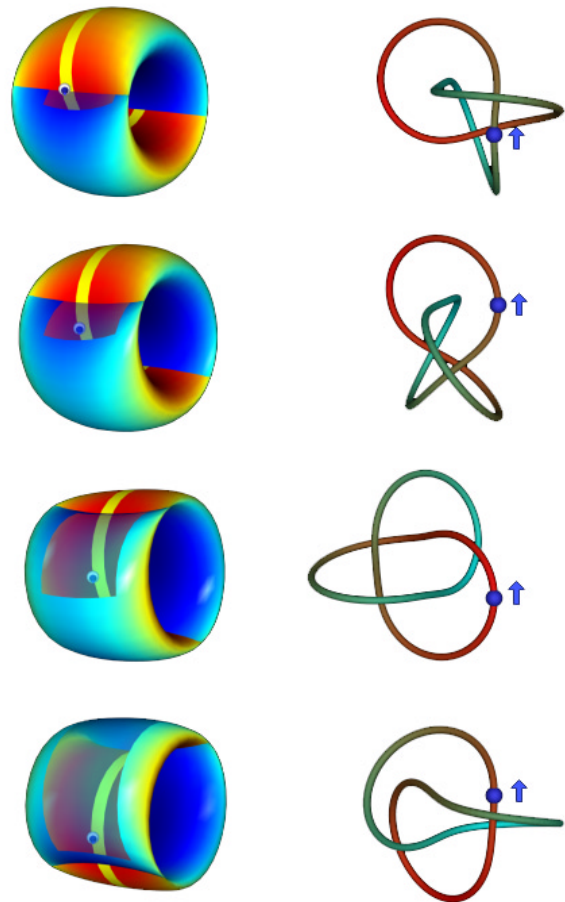


Figure 9: A haptic interface for constrained motion keyed to automate the manifold/knot orientation toward the viewer. The viewer can then see and touch the surface feature of interest locally. Perceiving the touched interest point is furthermore assisted for surfaces using the automated cutaway algorithm to suppress obscuring polygons.

about the description of higher spaces by combining “visual and sometimes motor” sensations. With the multimodal approach advocated here, many facets of our intuitive real-life experiences can in principle be transferred to the fourth dimension.

ACKNOWLEDGMENTS

This work was supported in part by National Science Foundation grants IIS-0430730 and CCR-0204112. We gratefully acknowledge Yin Wu’s contribution to the utility software packages, particularly the 4D cutaway and transparent blending tools, and Sidharth Thakur’s contribution to the concepts and design of the haptic interface that was extended to support the functionality used in this paper.

REFERENCES

- [1] Edwin A. Abbott. *Flatland*. Dover Publications, Inc., 1952.
- [2] T. F. Banchoff. Visualizing two-dimensional phenomena in four-dimensional space: A computer graphics approach. In E. Wegman and D. Priest, editors, *Statistical Image Processing and Computer Graphics*, pages 187–202. Marcel Dekker, Inc., New York, 1986.




Figure 10: Overview of the multimodal 4D exploration interface.




Figure 11: Exploration by following the continuous structure of the torus embedded in 4D, ignoring or overriding the disruptive self-intersections of the 3D projection.




Figure 12: Following an approximate geodesic path, illustrating the enhanced navigation mode on the ordinary torus.

[3] T. F. Banchoff. Beyond the third dimension: Geometry, computer graphics, and higher dimensions. *Scientific American Library*, 1990.

[4] David Banks. Interactive display and manipulation of two-dimensional surfaces in four dimensional space. In *Symposium on Interactive 3D Graphics*, pages 197–207. ACM, 1992.

[5] William V. Baxter, Vincent Scheib, and Ming C. Lin. dAb: Interactive haptic painting with 3D virtual brushes. In Eugene Fiume, editor, *SIGGRAPH 2001, Computer Graphics Proceedings*, pages 461–468. ACM SIGGRAPH, ACM, 2001.

[6] S. A. Carey, R. P. Burton, and D. M. Campbell. Shades of a higher dimension. *Computer Graphics World*, pages 93–94, October 1987.

[7] A. K. Dewdney. *The Planiverse: Computer Contact with a Two-Dimensional World*. Poseidon Press, 1984.

[8] R. Egli, C. Petit, and N. F. Stewart. Moving coordinate frames for representation and visualization in four dimensions. *Computers and Graphics*, 20(6):905–919, November–December 1996.

[9] Ben Forsyth. Intelligent support of interactive manual control: Design, implementation and evaluation of look-ahead haptic guidance. Master's thesis, The University of British Columbia, 2004.

[10] Jacques Hadamard. *The Mathematician's Mind: The Psychology of Invention in the Mathematical Field*. Princeton University Press, Princeton, NJ, 1996. Revised from original 1945 edition.

[11] A. J. Hanson. A construction for computer visualization of certain complex curves. *Notices of the Amer.Math.Soc.*, 41(9):1156–1163, November/December 1994.

[12] A. J. Hanson and P. A. Heng. Illuminating the fourth dimension. *Computer Graphics and Applications*, 12(4):54–62, July 1992.

[13] A. J. Hanson and H. Ma. Space walking. In *Proceedings of Visualization '95*, pages 126–133. IEEE, IEEE Computer Society Press, 1995.

[14] S. Hollasch. Four-space visualization of 4D objects. Master's thesis, Arizona State University, 1991.

[15] J. M. Kennedy. Optics and haptics: The picture. In *The conference*

on Multimodality of Human Communication: Theory, Problems and Applications. University of Toronto, 2002.

[16] L. Kim, A. Kyriku, G. Sukhatme, and M. Desbrun. An implicit-based haptic rendering technique, 2002.

[17] Thomas H. Massie and J. Kenneth Salisbury. The PHANTOM haptic interface: a device for probing virtual objects. *ASME Dynamic Systems and Control*, 55(1):295–301, 1994.

[18] Peter S. Maybeck. *Stochastic Models, Estimation, and Control, Volume 1*. Academic Press, New York, NY, 1979. The first chapter is available at www.cs.unc.edu/~welch/Kalman/index.html.

[19] Rudy Negenborn. Robot localization and kalman filters. Master's thesis, Utrecht University, 2003.

[20] Michael A. Noll. A computer technique for displaying n-dimensional hyperobjects. *Communications of the ACM*, 10(8):469–473, August 1967.

[21] Allison M. Okamura. *Haptic Exploration of Unknown Objects*. PhD thesis, Stanford University, Department of Mechanical Engineering, California, USA, June 2000.

[22] Allison M. Okamura and Mark R. Cutkosky. Feature detection for haptic exploration with robotic fingers. *The International Journal of Robotics Research*, 20(12):925–938, December 2001.

[23] June Gyu Park and Günter Niemeyer. Haptic rendering with predictive representation of local geometry. In *HAPTICS*, pages 331–338, 2004.

[24] Craig W. Reynolds. Steering behaviours for autonomous characters. In *Game Developers Conference 1999*, 1999.

[25] Dennis Roseman. Twisting and turning in four dimensions, 1993. Video animation, Department of Mathematics, University of Iowa, and the Geometry Center.

[26] SensAble, Inc. *3D Touch SDK OpenHaptics Toolkit Programmer's Guide*, 2004.

[27] Dan Simon. Kalman filtering. *Embedded Systems Programming*, 14(6):72–79, June 2001.

[28] K. V. Steiner and R. P. Burton. Hidden volumes: The 4th dimension. *Computer Graphics World*, pages 71–74, February 1987.

[29] Greg Welch and Gary Bishop. An introduction to the Kalman filter. Technical Report TR 95-041, University of North Carolina at Chapel Hill, Department of Computer Science, Chapel Hill, NC 27599-3175, 1995.

[30] W. Yu, R. Ramloll, and S.A. Brewster. Haptic graphs for blind computer users. In *First International Workshop on Haptic Human-Computer Interaction*, pages 102–107. British HCI Group, Springer, 2000.

[31] Mihaela A. Zahariev and Christine L. MacKenzie. Auditory, graphical and haptic contact cues for a reach, grasp, and place task in an augmented environment. In *Proc. of 5th Intl. Conf on Multimodal Interfaces (ICMI)*, pages 273–276, New York, 2003. ACM, ACM Press.