

TECHNICAL REPORT NO. 518

Quaternion Gauss Maps
and
Optimal Framings of Curves and Surfaces

by

Andrew J. Hanson

October 1998

COMPUTER SCIENCE DEPARTMENT
INDIANA UNIVERSITY
Bloomington, Indiana 47405-4101

Quaternion Gauss Maps and Optimal Framings of Curves and Surfaces

Andrew J. Hanson
Department of Computer Science
Lindley Hall 215
Indiana University
Bloomington, IN 47405 USA

October 1998

Abstract

We propose a general paradigm for generating optimal coordinate frame fields that may be exploited to annotate and display curves and surfaces. Parallel-transport framings, which work well for open curves, generally fail to have desirable properties for cyclic curves and for surfaces. We suggest that minimal quaternion measure provides an appropriate generalization of parallel transport. Our fundamental tool is the “quaternion Gauss map,” a generalization to quaternion space of the tangent map for curves and of the Gauss map for surfaces. The quaternion Gauss map takes 3D coordinate frame fields for curves and surfaces into corresponding curves and surfaces constrained to the space of possible orientations in quaternion space. Standard optimization tools provide application-specific means of choosing optimal, e.g., length- or area-minimizing, quaternion frame fields in this constrained space. We observe that some structures may have distinct classes of minimal quaternion framings, e.g, one disconnected from its quaternion reflection, and another that continuously includes its own quaternion reflection. We suggest an effective method for visualizing the geometry of quaternion maps that is used throughout. Quaternion derivations of the general moving-frame equations for both curves and surfaces are given; these equations are the quaternion analogs of the Frenet and Weingarten equations, respectively. We present examples of results of the suggested optimization procedures and the corresponding tubings of space curves and sets of frames for surfaces and surface patches.

1 Introduction

We propose a general framework for selecting optimal systems of coordinate frames that can be applied to the study of geometric structures such as curves and surfaces in three-dimensional space. The methods contain “minimal-turning” parallel-transport framings of curves as a special case,

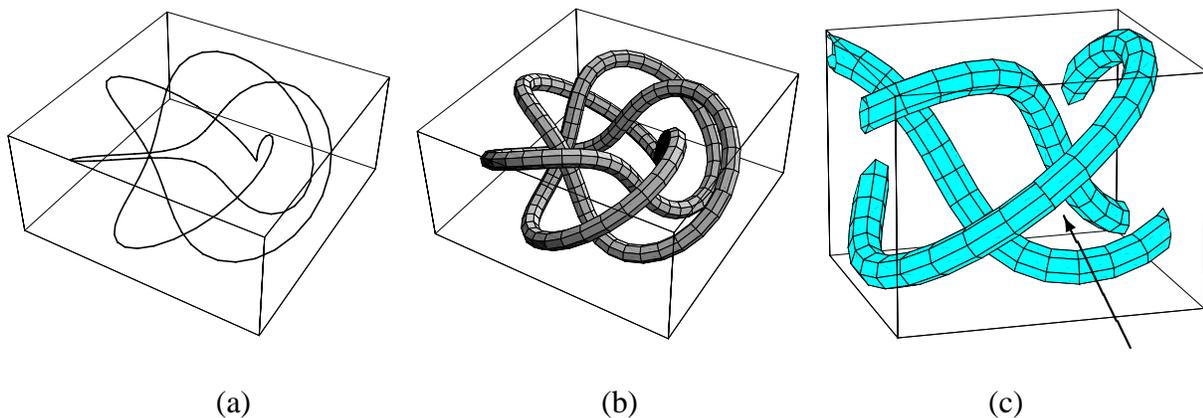


Figure 1: The (3,5) torus knot, a complex periodic 3D curve. (a) The line drawing is nearly useless as a 3D representation. (b) A tubing based on parallel transporting an initial reference frame produces an informative visualization, but is not periodic. (c) The arrow in this closeup exposes the subtle but crucial non-periodic mismatch between the starting and ending parallel-transport frames; this would invalidate any attempt to *texture* the tube. The methods of this paper provide robust parameterization-invariant principles for resolving such problems.

and extend naturally to situations where parallel-transport is not applicable. This article presents additional details of the IEEE Visualization '98 paper by the author [16].

Motivation. Many graphics problems require techniques for effectively displaying the properties of curves and surfaces. The problem of finding appropriate representations can be quite challenging. Representations of space curves based on single lines are often inadequate for graphics purposes; significantly better images result from choosing a “tubing” to display the curve as a graphics object with spatial extent. Vanishing curvature invalidates methods such as the Frenet frame, and alternative approaches such as those based on parallel transport involve arbitrary heuristics to achieve such properties as periodicity. Similar problems occur in the construction of suitable visualizations of complex surfaces and oriented particle systems on surfaces. If a surface patch is represented by a rectangular but nonorthogonal mesh, for example, there is no obvious local orthonormal frame assignment; if the surface has regions of vanishing curvature, methods based on directions of principal curvatures break down as well.

While we emphasize curves and surfaces in this paper to provide intuitive examples, there are several parallel problem domains that can be addressed with identical techniques. Among these are extrusion methods and generalized cones in geometric modeling, the imposition of constraints on a camera-frame axis in key-frame animation, and the selection of a 2D array of camera-frame axis choices as a condition on a constrained-navigation environment (see, e.g., Hanson and Wernert [20]).

Figure 1 summarizes the basic class of problems involving curves that will concern us here. The line drawing (a) of a (3,5) torus knot provides no useful information about the 3D structure.

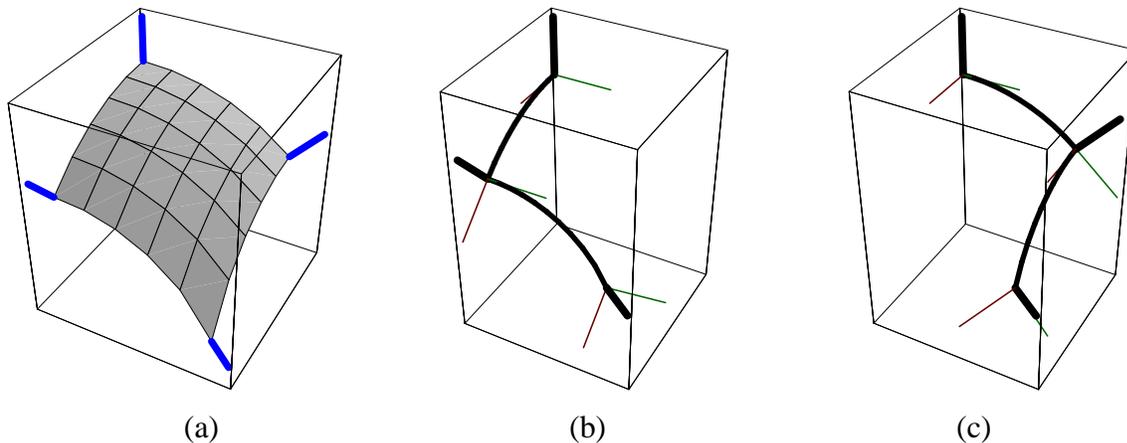


Figure 2: (a) A smooth 3D surface patch having a non-orthogonal parameterization, along with its geometrically-fixed normals at the four corners. No unique orthonormal frame is derivable from the parameterization. If we imitate parallel transport for curves to evolve the initial frame at the top corner to choose the frame at the bottom corner, we find that paths (b) and (c) result in incompatible final frames at the bottom corner. This paper addresses the problem of systematically choosing a compatible set of surface frames in situations like this.

Improving the visualization by creating a tubing involves a subtle dilemma that we attempt to expose in the rest of the figure. We cannot use a periodic Frenet frame as a basis for this tubing because inflection points or near-inflection points occur for many nice-looking torus knot parameterizations, and in such cases the Frenet frame is undefined or twists wildly. The parallel-transport tubing shown in (b) is well-behaved but not periodic; by looking carefully at the magnified portion next to the arrow in Figure 1(c), one can see a gross mismatch in the tessellation due to the non-periodicity which would, for example, preclude the assignment of a consistent texture. While it would be possible in many applications to ignore this mismatch, it has been the subject of a wide variety of previous papers (see, e.g., [24, 36, 5]), and must obviously be repaired for many other applications such as those requiring textured periodic tubes.

Figure 2 illustrates a corresponding problem for surface patches. While the normals to the four corners of the patch are always well-defined (a), one finds two different frames for the bottom corner depending upon whether one parallel transports the initial frame around the left-hand path (b) or the right-hand path (c). There is no immediately obvious right way to choose a family of frames covering this surface patch.

Our goal is to propose a systematic family of optimization methods for resolving problems such as these.

Methodology. We focus on unit quaternion representations of coordinate frames because of the well-known natural structure of unit quaternions as points on the three-sphere S^3 , which admits a natural distance measure for defining optimization problems, and supports in addition a variety of regular frame-interpolation methods (see, e.g., [37, 35, 31, 23]). We do not address the related

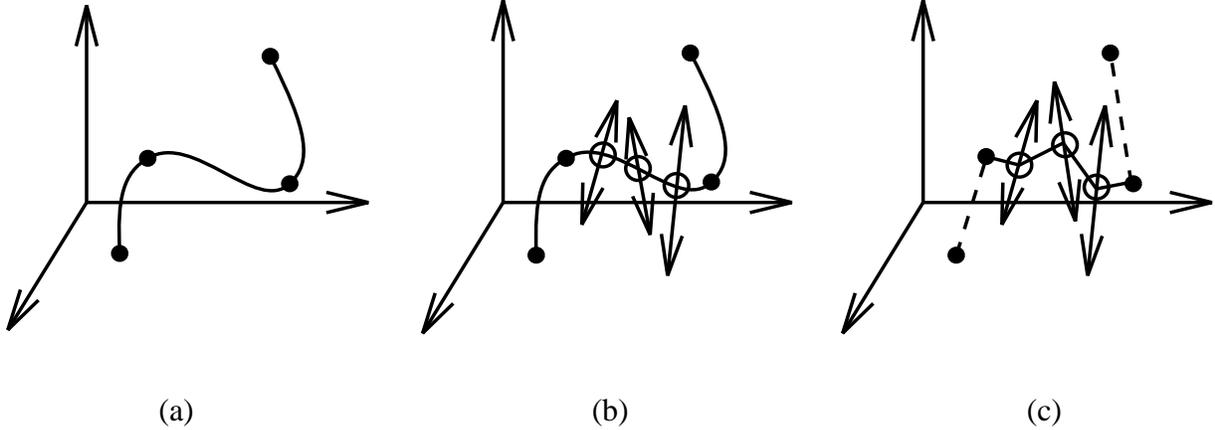


Figure 3: (a) The camera frame interpolation problem is analogous to the problem of finding a minimal-bending spline curve through a series of fixed key points. (b) The optimal curve frame assignment problem is analogous to fixing the end points of a curve segment and choosing *in addition* a family of lines along which the intermediate points are constrained to slide during the optimization process; in 3D, the spline path need not pass through the constraint lines. (c) In this paper, our sample points are generally close enough together that we apply the constraints to piecewise linear splines analogous to those shown here.

question of optimal freely moving frames treated by the minimal-tangential-acceleration methods (see, e.g., [2, 34, 11]); we are instead concerned with closely-spaced points on curves and surfaces where one direction of the frame is already fixed, and the chosen functional minimization in quaternion space must obey the additional constraint imposed by the fixed family of directions. Additional references of interest, especially regarding the treatment of surfaces, include [22, 32]. Figure 3 provides a visualization of the difference between the general interpolation problem and our constrained problem: a typical spline minimizes the bending energy specified by the chosen anchor points; requiring intermediate points to slide on constrained paths during the minimization modifies the problem. In particular, 3D spline curves need not intersect any of the constraint paths. In addition, we note that we typically have already sampled our curves and surfaces as finely as we need, so that piecewise linear splines are generally sufficient for the applications we discuss.

Our solution to the problem is to transform the intrinsic geometric quantities such as the tangent field of a curve and the normal field of a surface to quaternion space and to construct the quaternion manifold corresponding to the one remaining degree of rotational freedom in the choice of coordinate frame at each point. Curves and surfaces in these *spaces of possible frames* correspond to specific choices of the *quaternion Gauss map*, a subspace of the space of possible quaternion frames of the object to be visualized. Mathematically speaking, the space of possible frames is the circular *Hopf fiber* lying above the point in S^2 corresponding to each specific curve tangent or surface normal (see, e.g., [39, 3]).

For space curves, specifying a frame assignment as a quaternion path leads at once to tubular surfaces that provide a “thickened” representation of the curve that interacts well with lighting and rendering models. For surface patches, the approach results in a structure equivalent to that of

an anisotropic oriented particle system whose pairs of tangent vector fields in the surface produce natural flow fields that characterize the local surface properties and are easy to display. We will see that certain complex features of surfaces that are well-known in manifold theory arise naturally and can be clearly visualized using the quaternion Gauss map.

In the course of the discussion, we introduce a useful method of visualizing the geometry of the space of quaternions in which quaternion Gauss maps and the spaces of possible quaternion frames are represented. We show how to compute the required subspaces of frames in practice, and how to express this information in a form that can be used to optimize an energy measure, thereby leading to optimal frame choices. We also outline in the appendix a treatment of the curve and surface frame differential equations expressed directly in quaternion coordinates using the quaternion Lie algebra; these methods expose essential fundamental features of the quaternion frame methodology that are analogous to the Frenet and Weingarten equations.

Parallel Transport and Minimal Measure. Constraining each quaternion point (a frame) to its own circular quaternion path (the axial degree of rotational freedom), we then minimize the quaternion length of the frame assignment for curves and the quaternion area of the frame assignment for surfaces to achieve an optimal frame choice; this choice reduces to the parallel-transport frame for simple cases. Our justification for choosing minimal quaternion length for curves is that there is a unique rotation in the plane of two neighboring tangents that takes each tangent direction to its next neighbor along a curve: this is the geodesic arc connecting the two frames in quaternion space, and is therefore the minimum distance between the quaternion points representing the two frames. The choice of minimal area for surface frames is more heuristic, basically a plausibility argument that the generalization of minimal length is minimal area; no doubt this could be made more rigorous.

By imposing other criteria such as endpoint derivative values and minimal bending energy (see Barr et al. [2, 34]), the short straight line segments and polygons that result from the simplest minimization could be smoothed to become generalized splines passing through the required constraint rings; since, in practice, our curve and surface samplings are arbitrarily dense, this was not pursued in the current investigation.

For space curves, specifying a frame assignment as a quaternion path leads at once to tubular surfaces that provide a “thickened” representation of the curve that interacts well with texturing, lighting, and rendering models. For surface patches, the approach results in a structure equivalent to that of an anisotropic oriented particle system (also a species of texture) whose pairs of tangent vector fields in the surface produce natural flow fields that characterize the local surface properties and are easy to display.

Background. General questions involving the specification of curve framings have been investigated in many contexts; for a representative selection of approaches, see, e.g., [24, 36, 5, 28]. The quaternion Gauss map is a logical extension of the quaternion frame approach to visualizing space curves introduced by Hanson and Ma [19, 18]. The formulation of the quaternion form of the differential equations for frame evolution was introduced as early as the 1890’s by Tait [41].

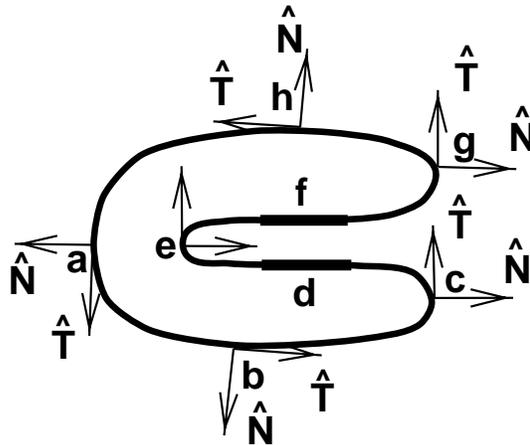


Figure 4: A smooth 2D curve with its normal and tangent frame fields. The segments d and f are intended to be straight.

For basic information on orientation spaces and their relationship to quaternions, see, e.g., [1, 33, 25]. Our own conventions are summarized in the appendix. The task of visualizing quaternions is also important, and we will describe our own approach below; for an interesting alternative, see Hart, Francis, and Kauffman [21]. Additional background on the differential geometry of curves and surfaces may be found in sources such as the classical treatise of Eisenhart [8] and in Gray's MATHEMATICA-based text [12], which inspired a number of the illustrations in this paper.

2 The Differential Geometry of Coordinate Frames

Our first goal is to define moving coordinate frames that are attached to curves and surfaces in 3D space. We will assume that our curves and surfaces are defined in practice by a discrete set of sample points connected by straight line segments, so that numerical derivatives can be defined at each point if analytic derivatives are not available. We begin with a pedagogical presentation of the properties of 2D curves, and then extend the surprisingly rich concepts that arise to 3D curves and surfaces.

2.1 Orientation Maps of 2D Curves

Suppose we have a smooth, arbitrarily differentiable 2D curve $\mathbf{x}(t) = x(t)\hat{\mathbf{x}} + y(t)\hat{\mathbf{y}}$. The curve itself generates a continuous set of changing tangents and normals of the form

$$\mathbf{T}(t) = d\mathbf{x}(t)/dt = x'\hat{\mathbf{x}} + y'\hat{\mathbf{y}} \quad (1)$$

$$\mathbf{N}(t) = y'\hat{\mathbf{x}} - x'\hat{\mathbf{y}}. \quad (2)$$

We choose this relative orientation convention so that in any dimension the tangent vector is expressible as the positive-signed cross-product of the normal(s); see [15] for further details. Unit

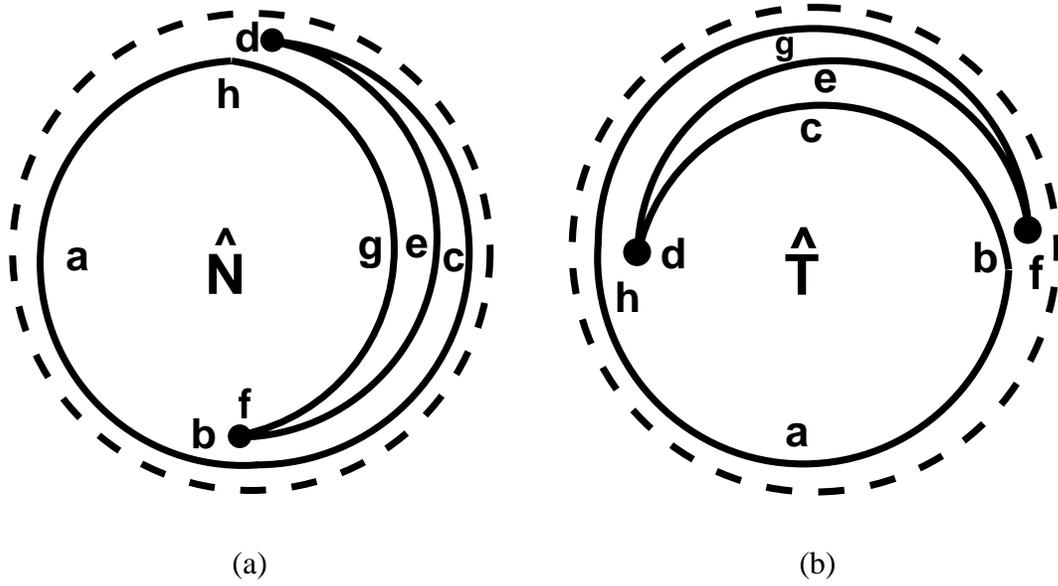


Figure 5: 2D Gauss map sketches of (a) the normal directions and (b) the tangent directions corresponding to the U-shaped curve in Figure 4. All these vectors lie on the unit circle in 2D. The straight line segments along d and f in Figure 4 correspond to single points in both maps.

length vectors will hereafter be distinguished with the conventional notation $\hat{v} = \mathbf{v}/\|\mathbf{v}\|$, so the normalized tangent and normal directions are denoted by $\hat{\mathbf{T}}$ and $\hat{\mathbf{N}}$.

In Figure 4, we show an example of a 2D curve with its tangent and normal fields. The *normalized* tangent and normal fields have only one degree of freedom, which we denote by the angle $\theta(t)$; the column vectors $\hat{\mathbf{N}}$ and $\hat{\mathbf{T}}$ then represent a moving orthonormal coordinate frame that may be expressed in the form

$$\begin{bmatrix} \hat{\mathbf{N}} & \hat{\mathbf{T}} \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}. \quad (3)$$

We may derive a 2D version of the frame equations by differentiating the frame to get

$$\hat{\mathbf{N}}'(t) = +v\kappa\hat{\mathbf{T}} \quad (4)$$

$$\hat{\mathbf{T}}'(t) = -v\kappa\hat{\mathbf{N}}, \quad (5)$$

where $\kappa(t)$ is the curvature and $v(t)$ is the “velocity” relative to the infinitesimal measure of curve length $ds^2 = \mathbf{dx}(t) \cdot \mathbf{dx}(t)$, that is, $d\theta(t)/dt = (ds/dt)(d\theta(s)/ds) = v(t)\kappa(t)$.

Note: we will find sign choices to be a subtle exercise throughout this paper. In Figure 4, the fact that the normal $\hat{\mathbf{N}}$ is chosen to point to the *outside* of a curve encircling an enclosed area in the right-hand sense makes the system inequivalent to the Frenet frame of the corresponding 3D curve, which would have $\hat{\mathbf{N}}$ pointing *inwards* everywhere except around the point e , and would be undefined along the straight segments d and f .

2D Tangent Map and Gauss Map. A 2D version of the Gauss map [8, 12] used in the classical differential geometry of surfaces follows when we discard the original curve in Figure 4 and restrict our view to show *only* the path of the normalized normals, as in Figure 5(a), or the normalized tangents, as in Figure 5(b); both vector fields take values only in the unit circle. We note that any sufficiently small open neighborhood of the curve has unique tangent and normal directions, up to the possibility of a shared limit point for straight segments such as d and f in Figure 4; over the whole curve, however, particular neighborhoods of directions may be repeated many times, resulting in an overlapping, non-unique 2D map, as indicated schematically in Figure 5. We will accept this as a feature, not necessarily a deficiency, of the construction.

2D “Quaternions.” In the appendix, we present the details of a derivation of a quaternion-like approach to the representation of 2D frames that may be informative to some readers. A brief summary begins by noting that the normal and tangent vectors can be parameterized by a quadratic form in the two variables a and b as

$$\begin{bmatrix} \hat{\mathbf{N}} & \hat{\mathbf{T}} \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} = \begin{bmatrix} a^2 - b^2 & -2ab \\ 2ab & a^2 - b^2 \end{bmatrix}, \quad (6)$$

where imposing the constraint $a^2 + b^2 = 1$ guarantees orthonormality of the frame.

By taking derivatives and extracting common factors, we find that the single matrix equation

$$\begin{bmatrix} a' \\ b' \end{bmatrix} = \frac{1}{2}v(t) \begin{bmatrix} 0 & -\kappa \\ +\kappa & 0 \end{bmatrix} \cdot \begin{bmatrix} a \\ b \end{bmatrix} \quad (7)$$

in the two variables with one constraint contains *both* the frame equations $\hat{\mathbf{T}}' = -v\kappa\hat{\mathbf{N}}$ and $\hat{\mathbf{N}}' = +v\kappa\hat{\mathbf{T}}$. The $\hat{\mathbf{N}}'$ and $\hat{\mathbf{T}}'$ equations are superficially a more complex set of two vector equations in four variables with three constraints. Equation (7) is effectively the *square root* of the frame equations. Rotations may be realized as complex multiplication in $(a + ib)$, and the pair $(a, b) = (\cos(\theta/2), \sin(\theta/2))$ parameterizes any rotation. Since $(a, b) \sim (-a, -b)$, the variables give a double covering of the space of rotations if we take the angular range from $0 \rightarrow 4\pi$ instead of 2π . These are precisely the properties we expect of quaternion representations of rotations.

2.2 3D Space Curves

We now move on to three-dimensional space curves. The fundamental difference in 3D is that, while the tangent direction is still determinable directly from the space curve, there is an additional degree of rotational freedom in the normal plane, the portion of the frame perpendicular to the tangent vector. This is indicated schematically in Figure 6.

Tangent Map. The tangent direction of a 3D curve at each point is given simply by taking the algebraic or numerical derivative of the curve at each sample point and normalizing the result. Each tangent direction thus has two degrees of freedom and lies on the surface of the two-sphere S^2 . The curve resulting from joining the ends of neighboring tangents is the *tangent map* of the curve. As

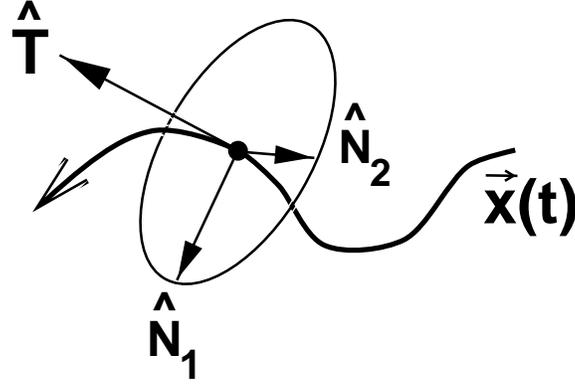


Figure 6: General form of a moving frame for a 3D curve $\mathbf{x}(t)$, with the tangent direction $\hat{\mathbf{T}}$ determined directly from the curve derivative, and the exact orientation of the basis $(\hat{\mathbf{N}}_1, \hat{\mathbf{N}}_2)$ for the normal plane determined only up to an axial rotation about $\hat{\mathbf{T}}$.

in the 2D case treated above, the tangent map of a 3D curve is not necessarily single-valued except in local neighborhoods, and may have limit points (e.g., if there are straight segments). In Figure 7(a,b), we show examples of two classic 3D curves, one a closed knot, the (2,3) trefoil knot lying on the surface of a torus, and the other the open helix:

$$\begin{aligned} \mathbf{x}_{\text{torus}}(p, q)(a, b, c)(t) &= (a + b \cos(qt)) \cos(pt) \hat{\mathbf{x}} + (a + b \cos(qt)) \sin(pt) \hat{\mathbf{y}} + c \sin(qt) \hat{\mathbf{z}} \\ \mathbf{x}_{\text{helix}}(a, b, c)(t) &= a \cos(t) \hat{\mathbf{x}} + b \sin(t) \hat{\mathbf{y}} + ct \hat{\mathbf{z}} . \end{aligned}$$

Differentiating these curves yields the tangent maps in Figure 7(c).

General Form of Curve Framings in 3D. The evolution properties of all possible frames for a 3D curve $\mathbf{x}(t)$ can be written in a unified framework. The basic idea is to consider an arbitrary frame to be represented in the form of columns of a 3×3 orthonormal rotation matrix,

$$\text{Curve Frame} = [\hat{\mathbf{N}}_1 \hat{\mathbf{N}}_2 \hat{\mathbf{T}}] . \quad (8)$$

Here $\hat{\mathbf{T}}(t) = \mathbf{x}'(t)/\|\mathbf{x}'(t)\|$ is the normalized tangent vector determined directly by the curve geometry, and which is thus unalterable; $(\hat{\mathbf{N}}_1(t), \hat{\mathbf{N}}_2(t))$ is a pair of orthonormal vectors spanning the plane perpendicular to the tangent vector at each point of the curve. Since $\|\hat{\mathbf{T}}\|^2 = \|\hat{\mathbf{N}}_1\|^2 = \|\hat{\mathbf{N}}_2\|^2 = 1$ and all other inner products vanish by definition, any change in a basis vector must be orthogonal to itself and thereby expressible in terms of the other two basis vectors. Thus the most general possible form for the frame evolution equations is

$$\begin{bmatrix} \hat{\mathbf{N}}_1'(t) \\ \hat{\mathbf{N}}_2'(t) \\ \hat{\mathbf{T}}'(t) \end{bmatrix} = v(t) \begin{bmatrix} 0 & +k_z(t) & -k_y(t) \\ -k_z(t) & 0 & +k_x(t) \\ +k_y(t) & -k_x(t) & 0 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{N}}_1(t) \\ \hat{\mathbf{N}}_2(t) \\ \hat{\mathbf{T}}(t) \end{bmatrix} , \quad (9)$$

where $v(t) = \|\mathbf{x}'(t)\|$ is the velocity of the curve if we are not using a unit speed parameterization.

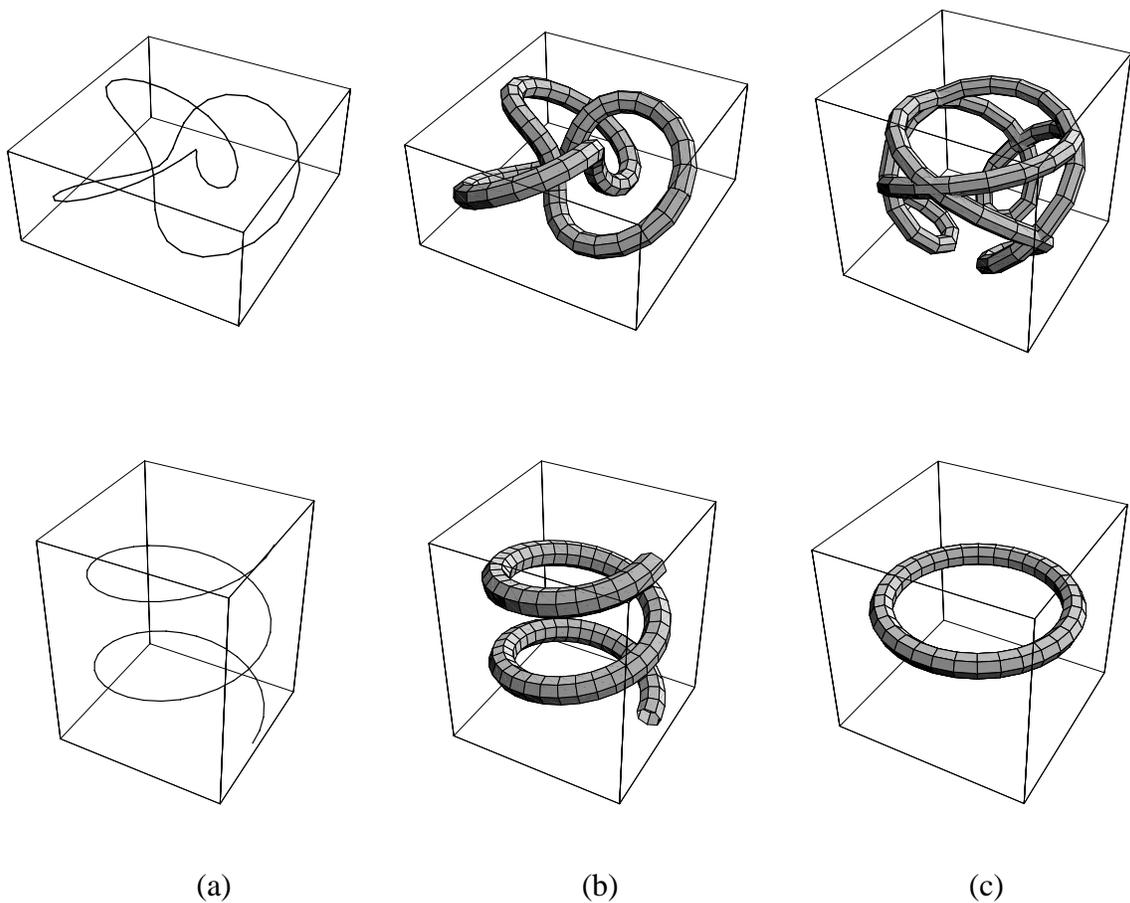


Figure 7: Tangent maps. (a) The (2,3) torus knot and the helix as 3D line drawings. (b) Illustrating an application of tubing to make the 3D curves more interpretable. (c) The corresponding normalized tangent maps determined directly from the curve geometry. These are curves on the two-sphere, and have also been tubed to improve visibility.

The particular choice of notation and signs for the curvatures \mathbf{k} in Eq. (9) is compellingly motivated by the quaternion Lie algebra treatment in the appendix, and its natural properties are also exposed using the Darboux form of the equations,

$$\begin{aligned}\hat{\mathbf{N}}'_1 &= v(t) \mathbf{F} \times \hat{\mathbf{N}}_1 \\ \hat{\mathbf{N}}'_2 &= v(t) \mathbf{F} \times \hat{\mathbf{N}}_2 \\ \hat{\mathbf{T}}' &= v(t) \mathbf{F} \times \hat{\mathbf{T}} ,\end{aligned}\tag{10}$$

where \mathbf{F} generalizes the Darboux vector field (see, e.g., Gray [12], p. 205):

$$\mathbf{F} = k_x \hat{\mathbf{N}}_1 + k_y \hat{\mathbf{N}}_2 + k_z \hat{\mathbf{T}} .\tag{11}$$

The square magnitude of the total “force” acting on the frame is $\|\mathbf{F}\|^2 = k_x^2 + k_y^2 + k_z^2$, and we will see below that this is a minimum for the parallel-transport frame.

The arbitrariness of the basis $(\hat{\mathbf{N}}_1(t), \hat{\mathbf{N}}_2(t))$ for the plane perpendicular to $\hat{\mathbf{T}}(t)$ can be exploited as desired to eliminate any one of the (k_x, k_y, k_z) (see, e.g., [4]). For example, if

$$\begin{aligned}\hat{\mathbf{M}}_1 &= \hat{\mathbf{N}}_1 \cos \theta - \hat{\mathbf{N}}_2 \sin \theta \\ \hat{\mathbf{M}}_2 &= \hat{\mathbf{N}}_1 \sin \theta + \hat{\mathbf{N}}_2 \cos \theta ,\end{aligned}\tag{12}$$

differentiating and substituting Eq. (9) yields

$$\hat{\mathbf{M}}'_1 = \hat{\mathbf{M}}_2(k_z - \theta') - \hat{\mathbf{T}}(k_x \sin \theta + k_y \cos \theta)\tag{13}$$

$$\hat{\mathbf{M}}'_2 = -\hat{\mathbf{M}}_1(k_z - \theta') + \hat{\mathbf{T}}(k_x \cos \theta - k_y \sin \theta) .\tag{14}$$

Thus the angle $\theta(t)$ may be chosen to cancel the angular velocity k_z in the $(\hat{\mathbf{N}}_1(t), \hat{\mathbf{N}}_2(t))$ plane. The same argument holds for any other pair. Attempting to eliminate additional components produces new mixing, leaving at least two independent components in the evolution matrix.

Tubing. For completeness, we note that to generate a ribbon or tube such as those used to display curves throughout this paper, one simply sweeps the chosen set of frames through each curve point $\mathbf{p}(t)$ to produce a connected tube,

$$\mathbf{x}(t, \theta) = \mathbf{p}(t) + \cos \theta \hat{\mathbf{N}}_1(t) + \sin \theta \hat{\mathbf{N}}_2(t) .$$

The resulting structure is sampled in t and over one full 2π period in θ to produce a tessellated tube. Arbitrary functions of (t, θ) can be introduced instead of the cosine and sine to produce ribbons and general linear structures.

Classical Frames. We now note a variety of approaches to assigning frames to an entire 3D space curve, each with its own peculiar advantages. Figure 8 compares the tubings of the (2,3) trefoil knot and the helix for each of the three frames described below.

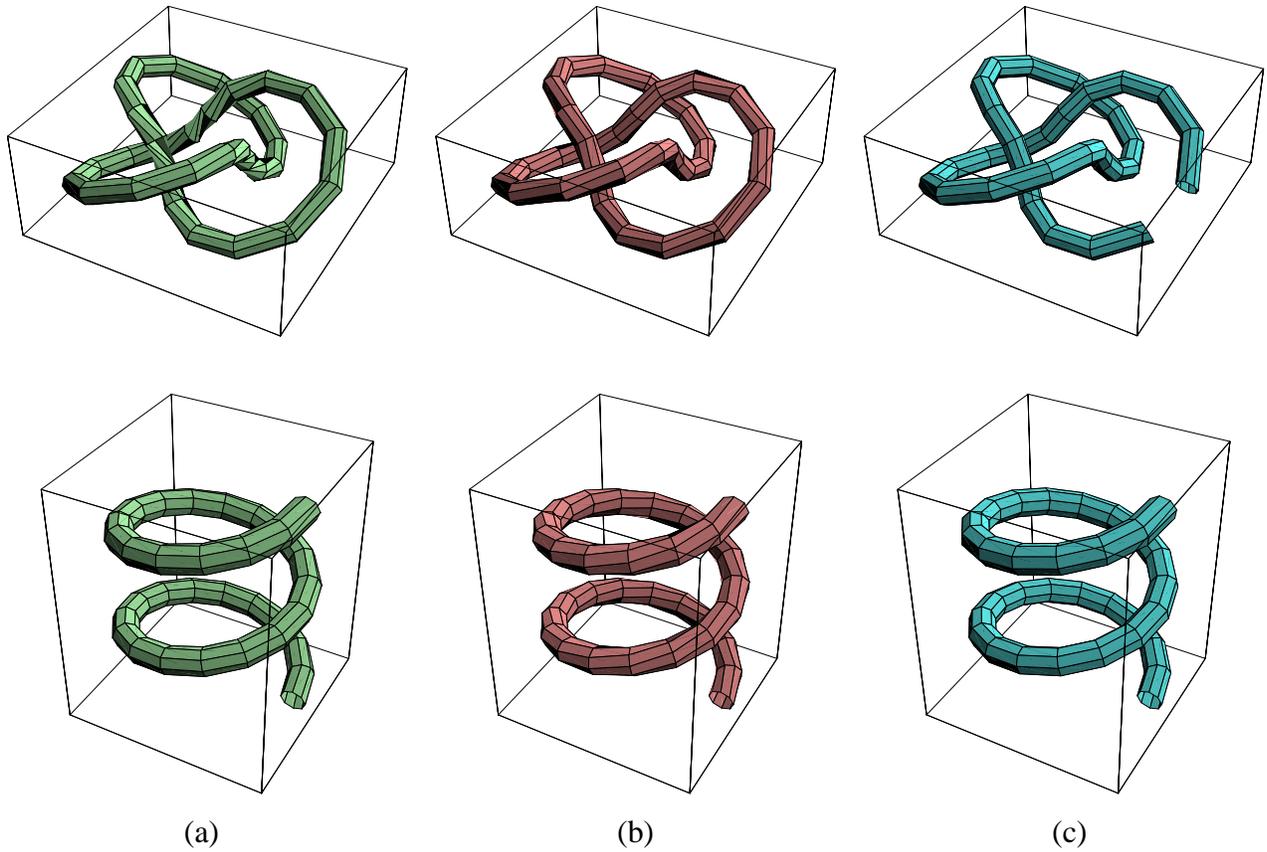


Figure 8: Curve framings for the (2,3) torus knot and the helix based on (a) Frenet frame, (b) Geodesic Reference frame (minimal tilt from North pole), and (c) Parallel Transport frame, which is not periodic like the other frames.

- **Frenet-Serret Frame.** This classical frame is determined by local conditions at each point of the curve, but is undefined whenever the curvature vanishes (e.g., when the curve straightens out or has an inflection point). For the Frenet frame, $k_x = 0$, k_y is the inverse radius of curvature, i.e., the curvature $\kappa(t)$, and $k_z(t)$ is the torsion $\tau(t)$, which mixes the two normal vectors in their local plane. This choice produces the usual equations

$$\begin{bmatrix} \hat{\mathbf{T}}'(t) \\ \hat{\mathbf{N}}'(t) \\ \hat{\mathbf{B}}'(t) \end{bmatrix} = v(t) \begin{bmatrix} 0 & \kappa(t) & 0 \\ -\kappa(t) & 0 & \tau(t) \\ 0 & -\tau(t) & 0 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{T}}(t) \\ \hat{\mathbf{N}}(t) \\ \hat{\mathbf{B}}(t) \end{bmatrix}. \quad (15)$$

Note that the squared Darboux vector is thus $\|\mathbf{F}\|^2 = \kappa^2 + \tau^2 \geq \kappa^2$.

If $\mathbf{x}(t)$ is any thrice-differentiable space curve, we can identify the triad of normalized Frenet frame vectors directly with the local derivatives of the curve,

$$\begin{aligned} \hat{\mathbf{T}}(t) &= \frac{\mathbf{x}'(t)}{\|\mathbf{x}'(t)\|} \\ \hat{\mathbf{N}}_1 = \hat{\mathbf{N}}(t) &= \hat{\mathbf{B}}(t) \times \hat{\mathbf{T}}(t) \\ \hat{\mathbf{N}}_2 = \hat{\mathbf{B}}(t) &= \frac{\mathbf{x}'(t) \times \mathbf{x}''(t)}{\|\mathbf{x}'(t) \times \mathbf{x}''(t)\|}, \end{aligned} \quad (16)$$

with $\kappa = \|\mathbf{x}'(t) \times \mathbf{x}''(t)\|/\|\mathbf{x}'(t)\|^3$, $\tau = \mathbf{x}'(t) \times \mathbf{x}''(t) \cdot \mathbf{x}'''(t)/\|\mathbf{x}'(t) \times \mathbf{x}''(t)\|^2$. For further details, see [8, 12].

- **Parallel-Transport Frame.** This frame is equivalent to a heuristic approach that has been frequently used in graphics applications (see, e.g., [24, 36, 5, 28]). A careful mathematical treatment by Bishop [4] presents its differential properties in a form that can be easily compared with the standard features of the Frenet frame. The parallel transport frame is distinguished by the fact that it uses the smallest possible rotation at each curve sample to align the current tangent vector with the next tangent vector. The current orientation of the plane normal to the tangent vector depends on the history of the curve, starting with an arbitrary initial frame, and so one is essentially integrating a differential equation for the frame change around the curve. The frame depends on the initial conditions, and cannot be determined locally on the curve like the Frenet frame. The algorithm with the best limiting properties [27] for computing this frame involves determining the normal direction $\hat{\mathbf{N}} = \mathbf{T}_i \times \mathbf{T}_{i+1}/\|\mathbf{T}_i \times \mathbf{T}_{i+1}\|$ to the plane of two successive tangents to the curve, finding the angle $\theta = \arccos(\hat{\mathbf{T}}_i \cdot \hat{\mathbf{T}}_{i+1})$, and rotating the current frame to the next frame using the 3×3 matrix $R(\theta, \hat{\mathbf{N}})$ or its corresponding quaternion (see appendix)

$$q(\theta, \hat{\mathbf{N}}) = q(\arccos(\hat{\mathbf{T}}_i \cdot \hat{\mathbf{T}}_{i+1}), \mathbf{T}_i \times \mathbf{T}_{i+1}/\|\mathbf{T}_i \times \mathbf{T}_{i+1}\|). \quad (17)$$

If the successive tangents are collinear, one leaves the frame unchanged; if the tangents are anti-collinear, a result can be returned, but it is not uniquely determined.

To identify the parallel transport frame with Eq. (9), we set $k_y \Rightarrow k_1$, $-k_x \Rightarrow k_2$, and $k_z = 0$ to avoid unnecessary mixing between the normal components (effectively the definition of parallel transport); this choice produces Bishop's frame equations,

$$\begin{bmatrix} \hat{\mathbf{N}}'_1(t) \\ \hat{\mathbf{N}}'_2(t) \\ \hat{\mathbf{T}}'(t) \end{bmatrix} = v(t) \begin{bmatrix} 0 & 0 & -k_1(t) \\ 0 & 0 & -k_2(t) \\ k_1(t) & k_2(t) & 0 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{N}}_1(t) \\ \hat{\mathbf{N}}_2(t) \\ \hat{\mathbf{T}}(t) \end{bmatrix}. \quad (18)$$

Since $\|\hat{\mathbf{T}}'\|^2 = (k_1)^2 + (k_2)^2$ is an invariant independent of the choice of the normal frame, Bishop identifies the curvature, orientation, and angular velocity

$$\begin{aligned} \kappa(t) &= \left((k_1)^2 + (k_2)^2 \right)^{1/2} \\ \theta(t) &= \arctan \left(\frac{k_2}{k_1} \right) \\ \omega(t) &= \frac{d\theta(t)}{dt}. \end{aligned}$$

k_1 and k_2 thus correspond to a Cartesian coordinate system for the ‘‘curvature polar coordinates’’ (κ, θ) with $\theta = \theta_0 + \int \omega(t) dt$; $\omega(t)$ is effectively the classical torsion $\tau(t)$ appearing in the Frenet equations. Note that the squared Darboux vector $\|\mathbf{F}\|^2 = \|\hat{\mathbf{T}}'\|^2 = k_1^2 + k_2^2 = \kappa^2$ is now a frame invariant. It is missing the torsion component present for the Frenet frame, and thus assumes its minimal value.

- **Geodesic Reference Frame.** In this paper, we will often need a frame that is guaranteed to have a particular axis in one direction, but we will not care about the remaining axes because they will be considered as a space of possibilities. A convenient frame with these properties can always be constructed starting from the assumption that there exists a canonical reference frame in which, say, the $\hat{\mathbf{z}}$ axis corresponds to the preferred direction. Thus if $\hat{\mathbf{v}}$ is the desired direction of the new axis, we can simply tilt the reference axis $\hat{\mathbf{z}}$ into $\hat{\mathbf{v}}$ along a minimal, geodesic curve using an ordinary rotation $R(\theta, \hat{\mathbf{n}})$ or its corresponding quaternion (see appendix):

$$q(\theta, \hat{\mathbf{n}}) = q(\arccos(\hat{\mathbf{z}} \cdot \hat{\mathbf{v}}), \hat{\mathbf{z}} \times \hat{\mathbf{v}} / \|\hat{\mathbf{z}} \times \hat{\mathbf{v}}\|). \quad (19)$$

Clearly any reference frame, including frames related to the viewing parameters of a moving observer, could be used instead of $\hat{\mathbf{z}}$. This frame has the drawback that it is ambiguous whenever $\hat{\mathbf{v}} = -\hat{\mathbf{z}}$; sequences of frames passing through this point will not necessarily be smoothly varying since only a single instance of a one-parameter family of frames can be returned automatically by a context-free algorithm. Luckily, this is of no consequence for our application. As we will discuss later in the quaternion framework, this property is directly related to the absence of a global vector field on the two-sphere.

- **General Frames.** We will henceforth work with the general framework for coordinate frames of arbitrary generality, rather than choosing conventional frames or hybrids of the

frames described so far (see, e.g., Klock [24]). While the classical frames have many fundamentally appealing mathematical properties, we are not in fact restricted to use any one of them. Keeping the tangent vector field intact, we may modify the angle of rotation about the tangent vector at will to produce an application-dependent frame assignment. An example of such an application is a closed curve with inflection points: the Frenet frame is periodic but not globally defined, the parallel transport frame will not be periodic in general, and the Geodesic Reference frame will be periodic but may have discontinuities for antipodal orientations. Thus, to get a satisfactory smooth global frame, we need something close to a parallel transport frame but with a periodic boundary condition; an example of an ad hoc solution is to take the Parallel Transport frame and impose periodicity by adding to each vertex's axial rotation a fraction of the angular deficit of the parallel transport frame after one circuit. But this is highly heuristic and depends strongly on the chosen parameterization. In the following sections, we introduce a more comprehensive approach.

2.3 3D Surfaces

If we are given a surface patch $\mathbf{x}(u, v)$ with some set of non-degenerate coordinates (u, v) , we may determine the normals at each point by computing

$$\mathbf{N}(u, v) = \mathbf{x}_u \times \mathbf{x}_v, \quad (20)$$

where $\mathbf{x}_u = \partial \mathbf{x} / \partial u$ and $\mathbf{x}_v = \partial \mathbf{x} / \partial v$. For surfaces defined numerically in terms of vertices and triangles, we would choose a standard procedure such as averaging the normals of the faces surrounding each vertex to determine the vertex normal. Alternatively, if we have an implicit surface described by the level-set function $f(\mathbf{x}) = 0$, the normals may be computed directly from the gradient at any point \mathbf{x} satisfying the level set equation:

$$\mathbf{N}(\mathbf{x}) = \nabla f(\mathbf{x}).$$

The normalized normal is defined as usual by $\hat{\mathbf{N}} = \mathbf{N} / \|\mathbf{N}\|$.

For 3D curves, the geometry of the curve determined the tangent vector $\hat{\mathbf{T}}$ and left a pair of normal vectors $(\hat{\mathbf{N}}_1, \hat{\mathbf{N}}_2)$ with one extra degree of freedom to be determined in the total frame $[\hat{\mathbf{N}}_1 \hat{\mathbf{N}}_2 \hat{\mathbf{T}}]$. The analogous observation for surfaces is that the geometry fixes the *normal* at each surface point, leaving a pair of *tangent* vectors $(\hat{\mathbf{T}}_1, \hat{\mathbf{T}}_2)$ with one extra degree of freedom to be determined in the total surface frame,

$$\text{Surface Frame} = [\hat{\mathbf{T}}_1 \hat{\mathbf{T}}_2 \hat{\mathbf{N}}]. \quad (21)$$

When a (u, v) surface parameterization is available, the surface partial derivatives \mathbf{x}_u and \mathbf{x}_v can in principle be used to assign a frame $[\hat{\mathbf{T}}_1 \hat{\mathbf{T}}_2 \hat{\mathbf{N}}]$ (using Gram-Schmidt if $\mathbf{x}_u \cdot \mathbf{x}_v \neq 0$), but there is no reason to believe that this frame has any special properties in general. In practice, it is extremely convenient to define a rectangular mesh on the surface patch, and a grid parameterized by (u, v) typically serves this purpose.

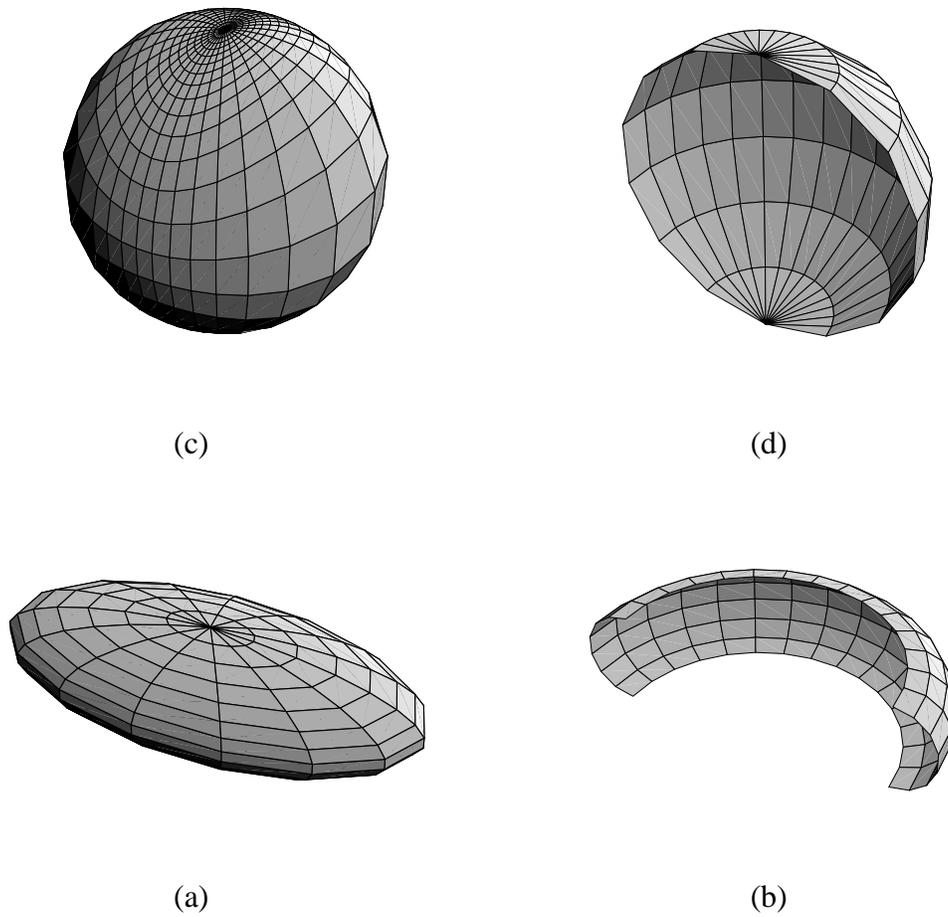


Figure 9: Classical Gauss maps of surfaces. (a) An ellipsoid and (b) a portion of a torus. (c,d) The corresponding standard Gauss maps of the normal vectors onto the sphere. Patches with coincident normals (e.g., for the full torus) would overlap in this representation.

Classical Gauss Map. The surface analog of the tangent map of a curve is the Gauss map, which takes a selection of points on the surface, typically connected by a mesh of some sort, and associates to each point its normalized surface normal. The Gauss map is then the plot of each of these normals in the coordinate system of a unit sphere S^2 in \mathbb{R}^3 . The Gauss map is guaranteed to be unique in some sufficiently small open set of each point of a regular surface, but may be arbitrarily multiple valued for the entire surface; note also that many nearby surface points can be mapped to a single point in the Gauss map, e.g., for certain types of planar curves in the surface or a planar area patch.

In Figure 9, we show a coordinate mesh on an elliptical surface and its single-valued Gauss map, as well as a quarter of a torus and its Gauss map; the Gauss map of the entire torus would cover the sphere twice, and there are two entire circles on the torus that correspond to single points, the North and South poles, in the Gauss map.

Surface Frame Evolution. The equations for the evolution of a surface frame follow the same basic structure as those of a space curve, except the derivatives are now directional, with two linearly independent degrees of freedom corresponding to the tangent basis $(\hat{\mathbf{T}}_1, \hat{\mathbf{T}}_2)$ in the surface. Typically (see [8, 12]), one assumes a not-necessarily-orthogonal parameterization (u, v) that permits one to express the tangent space in terms of the partial derivatives $(\mathbf{x}_u, \mathbf{x}_v)$, giving the normals $\hat{\mathbf{N}}(u, v)$ of Eq. (20). Then one can express the local curvatures in terms of any linearly independent pair of vector fields (\mathbf{U}, \mathbf{V}) as

$$D_{\mathbf{U}}\hat{\mathbf{N}} \times D_{\mathbf{V}}\hat{\mathbf{N}} = K (\mathbf{U} \times \mathbf{V}) \quad (22)$$

$$D_{\mathbf{U}}\hat{\mathbf{N}} \times \mathbf{V} + \mathbf{U} \times D_{\mathbf{V}}\hat{\mathbf{N}} = 2H (\mathbf{U} \times \mathbf{V}) . \quad (23)$$

With $\mathbf{U} = \mathbf{x}_u \cdot \nabla$ and $\mathbf{V} = \mathbf{x}_v \cdot \nabla$, we get the classical expressions. As Gray succinctly notes, since all the derivatives of $\hat{\mathbf{N}}$ are perpendicular to $\hat{\mathbf{N}}$, the whole apparatus amounts to constructing the tangent map of the Gauss map.

If we try to build the geometry of surfaces from a parametric representation, then each directional derivative has a vector equation of the form of Eq. (9). Thus we may write equations of the general form

$$\frac{\partial}{\partial u} \begin{bmatrix} \hat{\mathbf{T}}_1(u, v) \\ \hat{\mathbf{T}}_2(u, v) \\ \hat{\mathbf{N}}(u, v) \end{bmatrix} = \begin{bmatrix} 0 & +a_z(u, v) & -a_y(u, v) \\ -a_z(u, v) & 0 & +a_x(u, v) \\ +a_y(u, v) & -a_x(u, v) & 0 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{T}}_1(u, v) \\ \hat{\mathbf{T}}_2(u, v) \\ \hat{\mathbf{N}}(u, v) \end{bmatrix} \quad (24)$$

and

$$\frac{\partial}{\partial v} \begin{bmatrix} \hat{\mathbf{T}}_1(u, v) \\ \hat{\mathbf{T}}_2(u, v) \\ \hat{\mathbf{N}}(u, v) \end{bmatrix} = \begin{bmatrix} 0 & +b_z(u, v) & -b_y(u, v) \\ -b_z(u, v) & 0 & +b_x(u, v) \\ +b_y(u, v) & -b_x(u, v) & 0 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{T}}_1(u, v) \\ \hat{\mathbf{T}}_2(u, v) \\ \hat{\mathbf{N}}(u, v) \end{bmatrix} . \quad (25)$$

The last lines of each of Eqs. (24) and (25) are typically combined in textbook treatments to give

$$\begin{bmatrix} \frac{\partial \hat{\mathbf{N}}(u, v)}{\partial u} \\ \frac{\partial \hat{\mathbf{N}}(u, v)}{\partial v} \end{bmatrix} = [\mathcal{K}] \begin{bmatrix} \hat{\mathbf{T}}_1(u, v) \\ \hat{\mathbf{T}}_2(u, v) \end{bmatrix} . \quad (26)$$

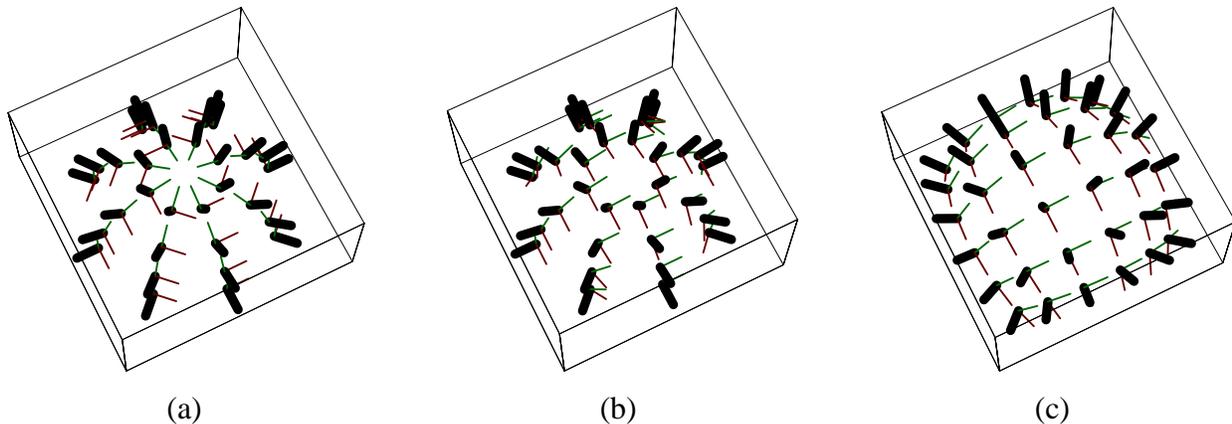


Figure 10: Examples of frame choices for the upper portion of an ordinary sphere. (a) Frames derived from standard polar coordinates on sphere. (b) Geodesic Reference frame for the sphere; each frame is as close as possible to the canonical coordinate axes at the North pole. (c) Frames derived from projective coordinates on the sphere, which turn out to be the same frame field as the Geodesic Reference frame.

where the matrix $[\mathcal{K}]$ has eigenvalues that are the principal curvatures k_1 and k_2 , and thus

$$K = \det [\mathcal{K}] = k_1 k_2 \quad (27)$$

is the Gaussian curvature and

$$H = \frac{1}{2} \text{tr} [\mathcal{K}] = \frac{1}{2}(k_1 + k_2) \quad (28)$$

is the mean curvature.

Examples of Surface Framings. If we are given a description of a surface, we can compute normals and choices of the corresponding frames by various means. In Figure 10, we illustrate three of these for the sphere. The first is derived from the standard orthonormal polar coordinate system, and the second is the extension to surfaces of the Geodesic Reference frame, which assigns the frame closest to a standard reference axis at the North Pole. The third is a frame based on polar projective coordinates for the sphere,

$$\begin{aligned} x(u, v) &= \frac{2u}{1 + u^2 + v^2} \\ y(u, v) &= \frac{2v}{1 + u^2 + v^2} \\ z(u, v) &= \frac{1 - u^2 - v^2}{1 + u^2 + v^2}, \end{aligned} \quad (29)$$

which map the real plane into the unit sphere with $x^2 + y^2 + z^2 = 1$ except for the point at infinity corresponding to the South pole. In fact, the polar projective coordinates generate the

same assignments as the Geodesic Reference frame does, so, except for the difference in locations of the grid sampling, these are the same framings.

Note: Do not be confused by alternate *samplings* of the same *framings*; if a parameterization $\mathbf{x}(u, v)$ gives a frame with $\mathbf{T}_1 = \partial\mathbf{x}(u, v)/\partial u$ and $\mathbf{T}_2 = \partial\mathbf{x}(u, v)/\partial v$, we can change to a polar sampled *mesh*, ($r = (u^2 + v^2)^{1/2}$, $\theta = \arctan(v, u)$), yet still retain the same frames at the same points $\mathbf{x}(r, \theta) = \mathbf{x}(u = r \cos \theta, v = r \sin \theta)$.

3 Quaternion Frames

In Section 2.1, we discussed the nature of 2D frames and noted a means of re-expressing the four equations with three constraints of the conventional frame system more efficiently; we showed a transformation into an equivalent set of two equations involving a single pair of variables obeying a unit length constraint and whose rotation transformation properties were realized by complex multiplication. Quaternions accomplish exactly this same transformation for 3D rotations: they permit the nine coupled frame equations with six orthonormality constraints in 3D to be succinctly summarized in terms of four quaternion equations with the single constraint of unit length. Detailed derivations along with other basic properties of quaternions are provided for reference in the appendix. A brief summary is given below.

Quaternion Frame Equations. Our task is now to rephrase the general properties of curve and surface frames in quaternion language so that, for example, we have a sensible space in which to consider optimizing frame assignments.

We begin with the standard definition for the correspondence between 3×3 matrices R^i_j and quaternions q :

$$R_q(\mathbf{V})^i = \sum_j R^i_j V^j = q * (0, V^i) * q^{-1}. \quad (30)$$

Henceforth, we will use the notation “*” to distinguish quaternion multiplication, and will use “.” when necessary to denote ordinary Euclidean inner products. Next, we express each orthonormal frame component as a column of R^i_j by using an arbitrary quaternion to rotate each of the three Cartesian reference axes to a new, arbitrary, orientation:

$$\begin{aligned} \hat{\mathbf{N}}_1 \text{ or } \hat{\mathbf{T}}_1 &= q * (0, \hat{\mathbf{x}}) * q^{-1} \\ \hat{\mathbf{N}}_2 \text{ or } \hat{\mathbf{T}}_2 &= q * (0, \hat{\mathbf{y}}) * q^{-1} \\ \hat{\mathbf{T}} \text{ or } \hat{\mathbf{N}} &= q * (0, \hat{\mathbf{z}}) * q^{-1}. \end{aligned} \quad (31)$$

(Technically speaking, in the above equation $\hat{\mathbf{T}}$ really means the quaternion $(0, \hat{\mathbf{T}})$ with only a vector part, etc.) All this can be transformed into the following explicit representation of the frame vectors as columns of a matrix of quaternion quadratic forms:

$$\begin{aligned} \left[\begin{array}{ccc} [\hat{\mathbf{N}}_1] & [\hat{\mathbf{N}}_2] & [\hat{\mathbf{T}}] \end{array} \right] &= \\ \left[\begin{array}{ccc} [\hat{\mathbf{T}}_1] & [\hat{\mathbf{T}}_2] & [\hat{\mathbf{N}}] \end{array} \right] &= \end{aligned}$$

$$\begin{bmatrix} q_0^2 + q_1^2 - q_2^2 - q_3^2 & 2q_1q_2 - 2q_0q_3 & 2q_1q_3 + 2q_0q_2 \\ 2q_1q_2 + 2q_0q_3 & q_0^2 - q_1^2 + q_2^2 - q_3^2 & 2q_2q_3 - 2q_0q_1 \\ 2q_1q_3 - 2q_0q_2 & 2q_2q_3 + 2q_0q_1 & q_0^2 - q_1^2 - q_2^2 + q_3^2 \end{bmatrix}. \quad (32)$$

Note: MATHEMATICA users should remind themselves that matrices are stored as lists of rows in MATHEMATICA, so one must *transpose* a standard matrix to easily retrieve column vectors from Eq. (32) and avoid mysterious sign errors.

Taking differentials of Eq. (31), we generate expressions of the form

$$dq = q * (q^{-1} * dq) = q * \frac{1}{2} (0, \mathbf{k}) \quad (33)$$

$$\begin{aligned} dq^{-1} &= (dq^{-1} * q) * q^{-1} \\ &= -(q^{-1} * dq) * q^{-1} \\ &= -\frac{1}{2} (0, \mathbf{k}) * q^{-1} \end{aligned} \quad (34)$$

where

$$\mathbf{k} = 2(q_0 \mathbf{dq} - \mathbf{q} dq_0 - \mathbf{q} \times \mathbf{dq}).$$

Substituting these expressions into the the calculation for the first column, we immediately find the expected commutators of quaternion multiplication:

$$\begin{aligned} d\hat{\mathbf{N}}_1 &= dq * (0, \hat{\mathbf{x}}) * q^{-1} + q * (0, \hat{\mathbf{x}}) * dq^{-1} \\ &= \frac{1}{2} q * ((0, \mathbf{k}) * (0, \hat{\mathbf{x}}) - (0, \hat{\mathbf{x}}) * (0, \mathbf{k})) * q^{-1} \\ &= q * (0, \mathbf{k} \times \hat{\mathbf{x}}) * q^{-1}. \end{aligned}$$

The rest of the columns are computed similarly, and a straightforward expansion of the components of the cross products proves the correspondence between Eq. (33) and Eq. (9).

To relate the derivative to a specific curve coordinate system, for example, we would introduce the curve velocity normalization $v(t) = \|\mathbf{x}'(t)\|$ and write

$$q' = v(t) \frac{1}{2} q * (0, \mathbf{k}). \quad (35)$$

One of our favorite ways of rewriting this equation follows directly from the full form for the quaternion multiplication rule given in the appendix; since this multiplication can be written as an orthogonal matrix multiplication on the 4D quaternion space, we could equally well write

$$\begin{bmatrix} q'_0 \\ q'_1 \\ q'_2 \\ q'_3 \end{bmatrix} = v(t) \frac{1}{2} \begin{bmatrix} 0 & -k_x & -k_y & -k_z \\ +k_x & 0 & +k_z & -k_y \\ +k_y & -k_z & 0 & +k_x \\ +k_z & +k_y & -k_x & 0 \end{bmatrix} \cdot \begin{bmatrix} q_0 \\ q_1 \\ q_2 \\ q_3 \end{bmatrix}. \quad (36)$$

This is the 3D analog of Eq. (7).

At this point, there are many other directions we could carry this basic structure, but we will not pursue the general theory of quaternion differential geometry further here. We will conclude with a short summary of the quaternion treatment of the classical surface equations. Starting from Eq. (33), we are led immediately to the quaternion analogs of Eqs. (24) and (25):

$$q_u \equiv \partial q / \partial u = \frac{1}{2} q * (0, \mathbf{a}) \quad (37)$$

$$q_v \equiv \partial q / \partial v = \frac{1}{2} q * (0, \mathbf{b}) . \quad (38)$$

But how shall we express the curvatures in a way similar to the classical formula in Eq. (26)? An elegant form follows by pursuing the quaternion analog of the vector field equations given in Eqs. (22,23). We write

$$\begin{aligned} q_u * q_v^{-1} &= -\frac{1}{4} q * (0, \mathbf{a}) * (0, \mathbf{b}) * q^{-1} \\ &= -\frac{1}{4} q * (-\mathbf{a} \cdot \mathbf{b}, \mathbf{a} \times \mathbf{b}) * q^{-1} \\ &= -\frac{1}{4} \left[-\mathbf{a} \cdot \mathbf{b} \hat{\mathbf{I}} + (\mathbf{a} \times \mathbf{b})_x \hat{\mathbf{T}}_1 + (\mathbf{a} \times \mathbf{b})_y \hat{\mathbf{T}}_2 + (\mathbf{a} \times \mathbf{b})_z \hat{\mathbf{N}} \right] , \end{aligned} \quad (39)$$

where we use the quaternion forms in Eq. (31) with the addition of the quaternion identity element $\hat{\mathbf{I}} = (1, \mathbf{0}) = q * (1, \mathbf{0}) * q^{-1}$ for the frame vectors. We see that the projection to the normal direction gives precisely the determinant $(\mathbf{a} \times \mathbf{b})_z = K$ identified in Eq. (26) as the scalar curvature. The mean curvature follows from an expression similar to Eq. (23),

$$\begin{aligned} q * (0, \hat{\mathbf{x}}) * q_u^{-1} + q * (0, \hat{\mathbf{y}}) * q_v^{-1} &= -\frac{1}{2} q * (-\hat{\mathbf{x}} \cdot \mathbf{a} - \hat{\mathbf{y}} \cdot \mathbf{b}, \hat{\mathbf{x}} \times \mathbf{a} + \hat{\mathbf{y}} \times \mathbf{b}) * q^{-1} \\ &= -\frac{1}{2} \left[-(a_x + b_y) \hat{\mathbf{I}} + b_z \hat{\mathbf{T}}_1 - a_z \hat{\mathbf{T}}_2 + (a_y - b_x) \hat{\mathbf{N}} \right] \end{aligned} \quad (40)$$

where again the coefficient of the normal, $(a_y - b_x) = \text{tr}[\mathcal{K}] = 2H$, is the desired expression. Similar equations can be phrased directly in the 4D quaternion manifold using the forms of Eq. (36).

3.1 Visualizing Quaternion Frames

Seeing the parameters of a single quaternion. Any (unit) quaternion is a point on S^3 and therefore is described by three parameters incorporated in the standard parameterization

$$q(\theta, \hat{\mathbf{n}}) = \left(\cos \frac{\theta}{2}, \hat{\mathbf{n}} \sin \frac{\theta}{2} \right) , \quad (41)$$

where $0 \leq \theta < 4\pi$, and the eigenvector of the rotation matrix (unchanged by the rotation), is a point on the two-sphere S^2 representable as $\hat{\mathbf{n}} = (\cos \alpha \cos \beta, \sin \alpha \cos \beta, \sin \beta)$ with $0 \leq \alpha < 2\pi$ and $-\pi/2 \leq \beta \leq \pi/2$.

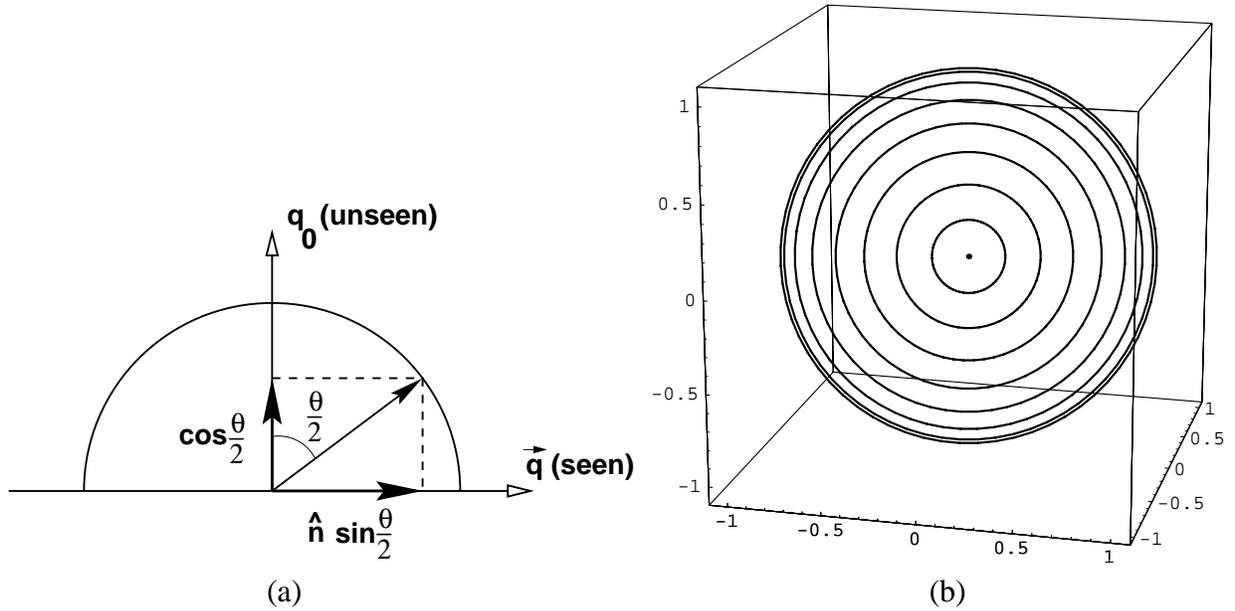


Figure 11: Illustration of how the q_0 part of a quaternion is “known” if we have a 3D image of the vector part $\mathbf{q} = \hat{\mathbf{n}} \sin \frac{\theta}{2}$ of the quaternion. (b) Schematic representation of the concentric-sphere uniform distance scales needed to form a mental model of the metric distances in quaternion space between two points in the parallel 3D projection. Distances are roughly Euclidean near the origin ($\mathbf{q} \approx 0$ in (a)) and equal-length lines appear increasingly compressed as the radius approaches unity.

An informative visualization of quaternions can be constructed by examining their properties carefully. If we simply make a 3D display of the vector part of the quaternion, $\hat{\mathbf{n}} \sin \frac{\theta}{2}$, we see that the scalar element of the quaternion is redundant, since, for each θ ,

$$q_0 = \cos \frac{\theta}{2} = \pm \left(1 - \left| \hat{\mathbf{n}} \sin \frac{\theta}{2} \right|^2 \right)^{1/2}. \quad (42)$$

That is, q_0 is just the implicitly known height of the 4D unit vector in the unseen projection direction, as illustrated in Figure 11(a). In Figure 11(b), we schematize the mental model of metric distance required to complete the interpretation of the visualization. If we imagine dividing the arc of the semi-circle in Figure 11(a) into equal angular segments, the arc lengths are all the same distance apart in spherical coordinates. Projected onto the \mathbf{q} plane, however, the projected spacing is non-uniformly scaled by a factor of $\sin \theta$. Thus to keep our vision of distance consistent, we imagine the space to be like 3D graph paper with concentric spheres drawn at equal distances in the special scale space; such 3D graph paper would look like Figure 11(b). Distances are essentially Euclidean near the 3D origin, for small 3D radii, and are magnified as the radius approaches unity to make the marked spheres equidistant in conceptual space.

If we assume the positive root is always taken for q_0 , then we effectively restrict ourselves to a single hemisphere of S^3 and eliminate the two-fold redundancy in the correspondence between

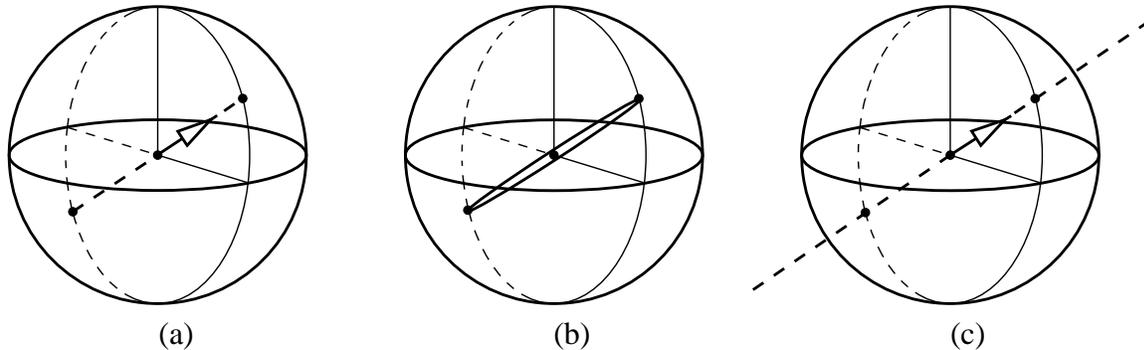


Figure 12: (a) This image represents the 3D vector part of the quaternion $q = (\cos \frac{\theta}{2}, \hat{\mathbf{t}} \sin \frac{\theta}{2})$ representing a single instance of the one-parameter family of possible rotations leaving invariant the tangent vector $\hat{\mathbf{t}}$ at one point of a space curve. (b) A representation of the entire one-parameter space of possible frames having the same tangent vector $\hat{\mathbf{t}}$. The vector part of the quaternion must lie on the diameter of the two-sphere in 3D depicted here. We depict the diameter as a very skinny ellipse, because it is in fact a degenerate projection of a circle in 4D, which could be exposed as shown by making a small 4D rotation before projecting to 3D. (c) A polar projection of the same object removes the doubling by projecting the circle to a line through infinity in R^3 .

quaternions and the rotation group. Alternatively, despite the fact that quaternions with both signs of q_0 map to the same point in this projection, we can indicate the simultaneous presence of both hemispheres using graphical cues; one possible method is to use saturated colors in the “front” hemisphere, and faded colors (suggesting distance) for objects in the “back” hemisphere.

Hemispheres in S^3 . To clarify the terminology, we note that a projected hemisphere for S^2 is a filled disk (a “two-ball”) in the plane, and the full surface of the sphere consists of two such disks joined at the outer circular boundary curve; for S^3 , we use the word hemisphere to indicate a filled solid two-sphere (technically a “three-ball”), and imagine the full volume of the three-sphere to consist of *two* such spherical solids joined on the skin (a two-sphere) of the surface enclosing both.

The family of possible values of Eq. (41) projects to a double-valued line (actually an “edge-on” projection of a circle) which is a directed diameter of the unit two-sphere, in the direction of $\hat{\mathbf{n}}$; in a polar projection, this circle becomes a line to infinity through the origin. These representations of a unit quaternion as a vector from the origin to a point inside the solid two-sphere (the three-ball) are illustrated schematically in Figure 12.

Any particular 3D rotation is represented twice, since the quaternion circle is parameterized by $0 \leq \theta < 4\pi$. A simple parallel projection thus produces two solid balls on top of each other in the 3D projection, one the analog of the “North pole disk” of a two-sphere parallel projected from 3D to a screen, the other the analog of the “South pole disk” of a two-sphere. The analog of a polar projection, which for a two-sphere sends the North pole to infinity of R^2 , flattens the three-sphere out to fill R^3 , as shown in Figure 12(c), and eliminates the double-valued properties of the parallel projection.

CURVE LENGTHS	(2,3) Torus Knot	Helix
Frenet Frame	14.3168	6.18501
Geodesic Reference Frame	14.6468	7.82897
Parallel Transport Frame	10.1865	6.06301

Table 1: Relative lengths (in radians) of the quaternion frame maps for various frame choices describing the (2,3) torus knot and the helix. The Parallel Transport frame is the shortest possible frame map.

3.2 Quaternion Frames for Curves

We now can produce quaternion frames for space curves directly by several techniques.

Quaternions from Local 3D Frames. In the case of the Frenet frame, we have no choice but to consider each frame as totally independent of the others. Each is locally computable, and there is in principle no relation between them, since the curvature could vanish at any point. In this case, we compute the frames directly from Eq. (16), thus deriving a 3×3 orthogonal matrix $R(t)$ at each point of the curve. We then apply standard inversion algorithms [37, 35, 31] to obtain the corresponding quaternion up to a sign. Finally, we apply a simple operator that checks the local continuity of the corresponding frames. If two quaternion vectors representing neighboring frames have a dot product near negative one, we change the sign of one to keep it near its neighbors. If two neighbors are excessively far apart in terms of the 4D angle between them, and are not simply near-negatives of one another, then the Frenet frame probably is poorly defined and should be tagged as such until continuity resumes. Figure 13 shows the Frenet frame tubing of a torus knot and the corresponding trajectory of these frames in the vector subspace of quaternion space.

Note: Forcing close quaternion Frenet frames on closed curves such as torus knots results in a very interesting phenomenon. Depending on the parameters of the curve, the path in quaternion space may close after a single traversal of the curve, or it may require two or more traversals, as in the case shown in Figure 13. We have checked this feature on a wide range of torus knots, and found that there are generally “jumps” between needing different numbers of circuits at those parameter values that imply inflection points (zero curvature) in the curve.

Direct Quaternion Frames. The Geodesic Reference frame and the Parallel Transport frame, in contrast to the Frenet frame, can be defined directly in terms of quaternions if desired, as indicated in Section 2.2; all that is needed is an initial quaternion reference frame, and then the geometry of the curve specifies enough at each point to express the needed rotation in quaternion form.

Comparison of Tubings and Quaternion Frames. Previously, in Figure 8, we compared the tubings for the (2,3) torus knot and for the helix based on the Frenet, Geodesic Reference, and Parallel Transport frames. The corresponding quaternion paths are illustrated together in Figure 14. The Parallel Transport frame shown uses the initial Frenet frame as a starting point; we could

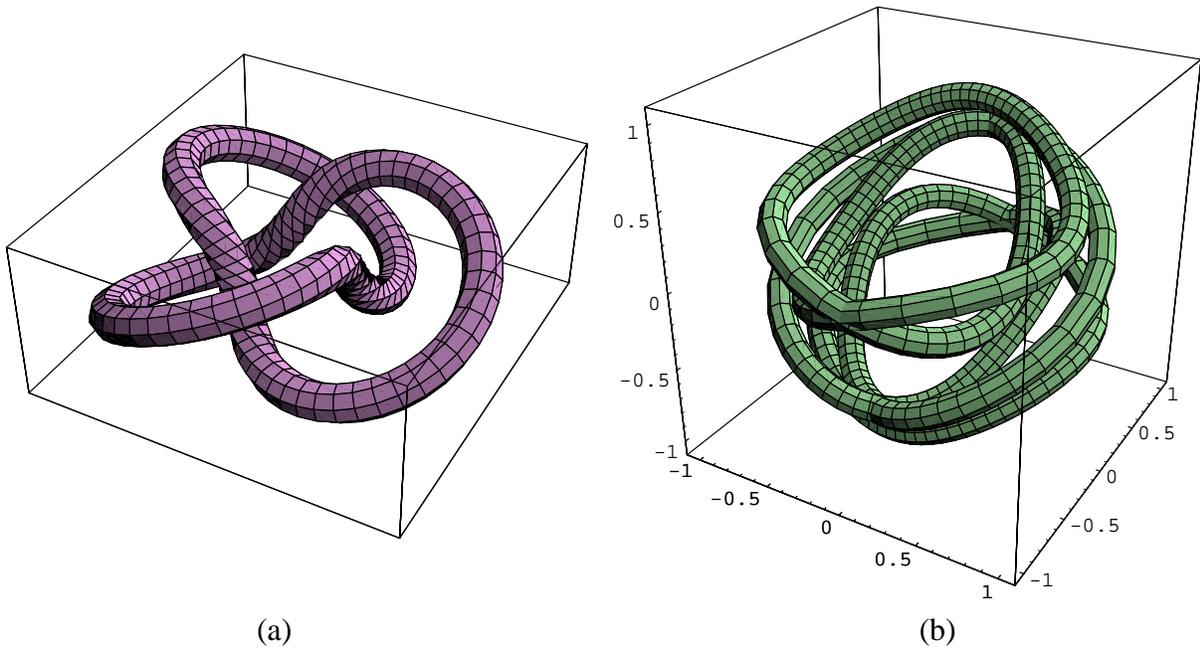
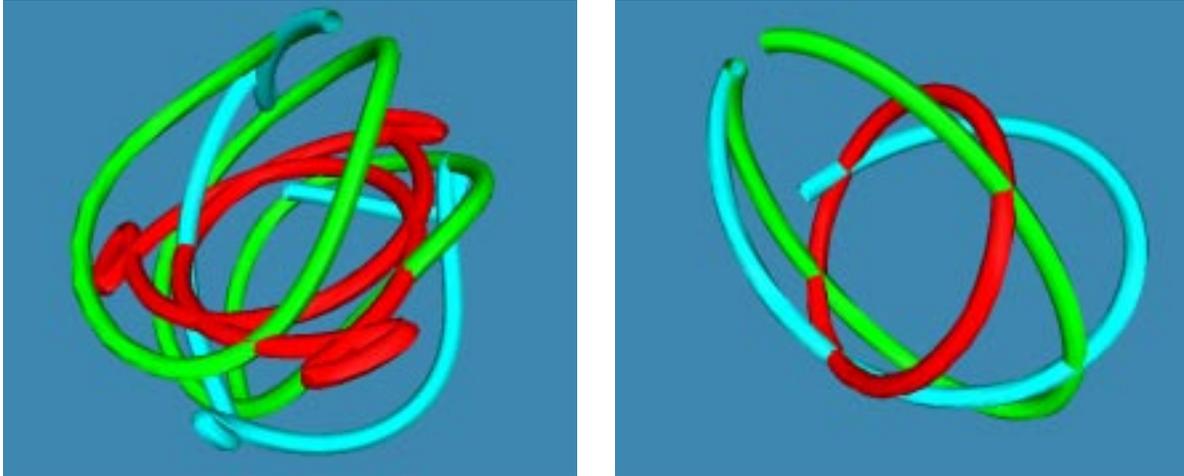


Figure 13: (a) A trefoil torus knot. (b) Its quaternion Frenet frame projected to 3D. For this trefoil knot, the frame does not close on itself in quaternion space unless the curve is traversed twice, corresponding to the double-valued “mirror” image of the rotation space that can occur in the quaternion representation. Observe the longer segments in (b): these correspond to the three high-torsion segments observable in (a).



(a)

(b)

Figure 14: (a) Quaternion frames in “standard” 3D vector visualization projection for the (2,3) torus knot: Red—Geodesic Reference: this is planar by construction, since all 3D points must lie in the plane perpendicular to the reference axis; the 3D origin is at the centroid of the red curve. Green—Frenet: the Frenet frame is actually cyclic, but to see this easily for this 2,3 torus knot, the mirror image of the current frame must be added, giving effectively a double traversal of the curve as shown in Figure 13. Cyan—Parallel Transport: the PT frame must be given a starting value, which here is seen at the top center of the image to coincide with the (green) Frenet frame. The PT frame is not cyclic, but is the shortest path, with three very noticeable tight loops. (b) The same selection of quaternion frames for the helix. Again, the red Geodesic Reference curve is planar (and cycles back on itself twice for this helix); the green Frenet frame takes a longer path that will return to its original orientation, and the cyan Parallel Transport frame, seen starting at the same orientation as the Frenet, will not ordinarily return to the same orientation, but will have the shortest 4D path length. (The hidden double circuit of the Geodesic Reference frame for this helix in fact makes it longer.)

use any starting quaternion with the correct tangent vector. The relative path lengths of the curves in Figure 14 are summarized in Table 1.

We note the following properties:

- **Frenet.** Periodic for periodic non-singular curves, has a tendency to twist a bit too much (where the torsion is high), leaving long jumps between neighboring samples in quaternion space; undefined at inflection points and zero curvature segments.
- **Geodesic Reference.** Also guaranteed to be periodic for periodic curves, but has the odd property that it always lies in a plane perpendicular to the reference axis in our preferred 3D quaternion projection. Ambiguous and therefore potentially not smooth for frames opposing the reference frame direction.
- **Parallel Transport.** This is the quaternion frame with minimal 4D length, though it may be difficult to see this feature immediately in our standard projection. It is not in general a periodic path. Different choices of starting frame produce curves of identical length differing by rigid (possibly reflecting) 4D motions (see Eq. (56)).

3.3 Quaternion Gauss Map for Surfaces

The quaternion Gauss map extends the Gauss map to include a representation of the entire coordinate frame at each surface point, introducing a number of new issues. In particular there is a useful, but mathematically suspect, approach that we might call an “engineering” approach to the quaternion Gauss map that lets us quickly get informative visualizations for those special cases where we are given a locally orthogonal parameterization of the surface except perhaps for isolated singularities of the coordinate system.

For these cases, we may construct the precise quaternion analog of the Gauss map by lifting the surface’s coordinate mesh into the space of quaternions at each value of the orthonormal coordinatization (u, v) of the surface or surface patch. The correspondence of this map to the Gauss map is *not* directly visible, since (see Eq. (32)) the normal directions of the Gauss map are non-trivial quadratic forms constructed from all the quaternion components; however, a projection to a subspace of the quaternion space based on the bilinear action of quaternions on pure vectors may be constructed by imitating the projection of the Hopf fibration of S^3 (see, e.g., Shoemake [39, 3]). In Figure 15, we show two such cases, an ellipsoid with orthonormal polar coordinates singular at the poles and a torus with global, nonsingular, coordinates, using our now-standard projections of the quaternion Gauss map to 3D. In each of these cases, a single circuit of the surface generates only one-half of the quaternion surface shown; the symmetric quaternion figure results from traversing the surface twice to adjoin the reflected image of the single-circuit quaternion surface. That is, each point on the 3D surfaces appears twice, once at q , and once at $-q$, in these periodic quaternion Gauss maps.

We see that the singular coordinate system typically used for the ellipsoid is topologically a cylinder; the circles corresponding to the singularities of the coordinate system (circles of normal directions) at the North and South poles correspond to *boundaries* of the quaternion Gauss map.

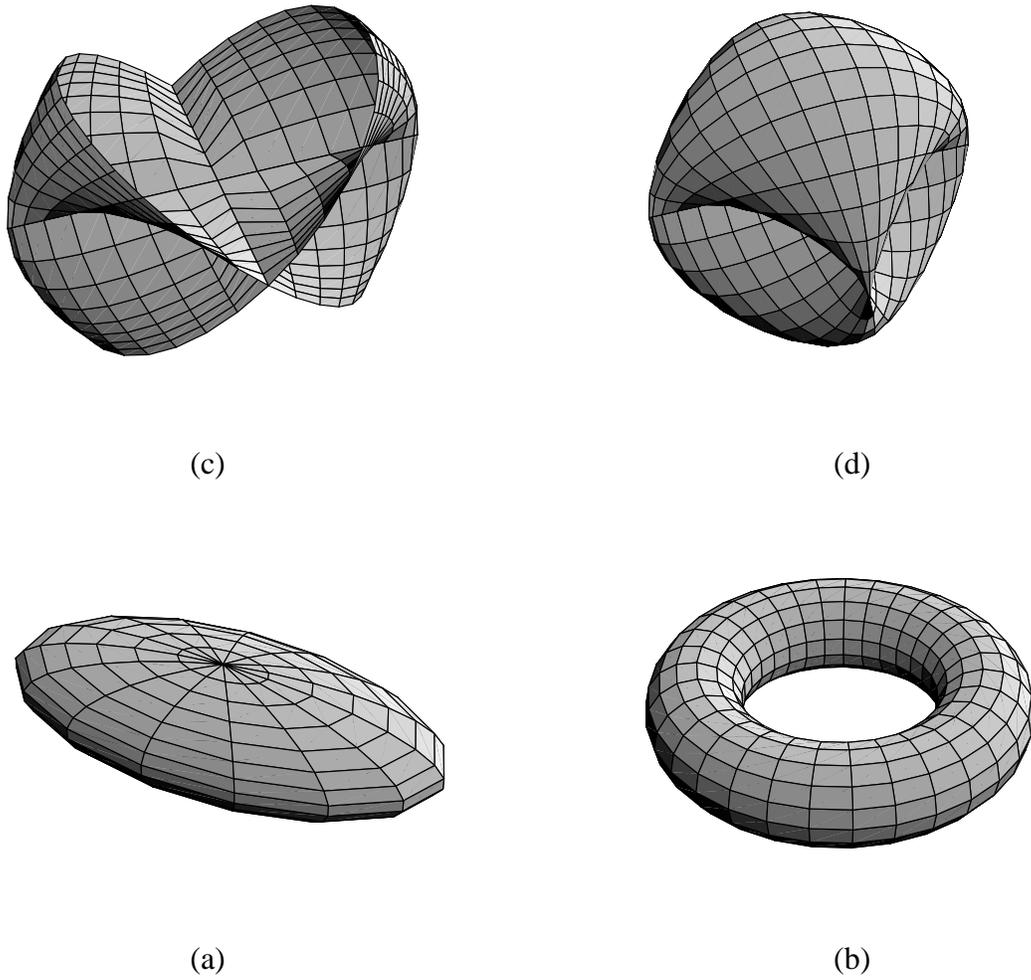


Figure 15: Examples of quaternion Gauss maps for surfaces. (a) The ellipsoid and (b) the torus. (c,d) The corresponding Quaternion Gauss maps, projected from the three-sphere in 4D. The equatorial direction has been traversed twice in order to get a closed path in the map; the singular poles in the ellipsoid coordinate system correspond to the edges or boundaries of the quaternion-space ribbon.

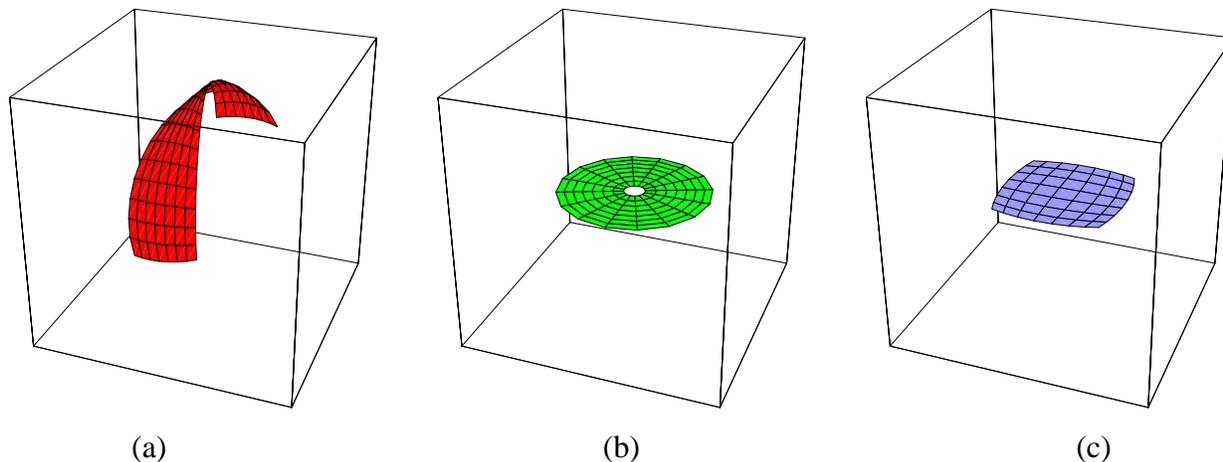


Figure 16: Examples of quaternion Gauss maps for the frame choices for the upper portion of an ordinary sphere given originally in Figure 10. (a) Frames derived from standard polar coordinates on sphere. (b) Geodesic reference frame for the sphere; each frame is as close as possible to the canonical coordinate axes at the North pole. (c) Frames derived from projective coordinates on the sphere.

PATCH AREAS	Hemispherical patch
Polar Coordinates	2.1546
Geodesic Reference Frame	1.9548

Table 2: Areas (in steradians) of the quaternion frame maps for the polar coordinate and Geodesic Reference frame choices on the hemispherical patches of Figure 16.

The torus, which has the extremely unusual feature that it possesses a global regular coordinate system, has a (reflection doubled) quaternion Gauss map which is another, four-dimensional, torus embedded in the quaternion S^3 space.

Quaternion Maps of Alternative Sphere Frames. In Figure 10, we showed three alternate sets of frames for the upper half of an ordinary sphere. The assigned coordinate systems may be converted directly into quaternion frames and coerced into consistency in the usual manner. In Figure 16, we show the results. The Geodesic Reference frames and the projective coordinates are in fact the same space of frames computed in different ways: both are planes perpendicular to the \hat{z} axis. The coordinate systems used to compute the quaternion Gauss maps in parts (a) and (b) of the figure are commensurate, so we may compare the areas, computed using solid angle on the three-sphere in units of steradians; the results are shown in Table 2.

Covering the Sphere and the Geodesic Reference Frame South Pole Singularity. The Geodesic Reference frame for a surface patch has the peculiarity that it has an ambiguity whenever the vector to be assigned is exactly opposite the reference frame. As we show in Figure 17, the tilting from the

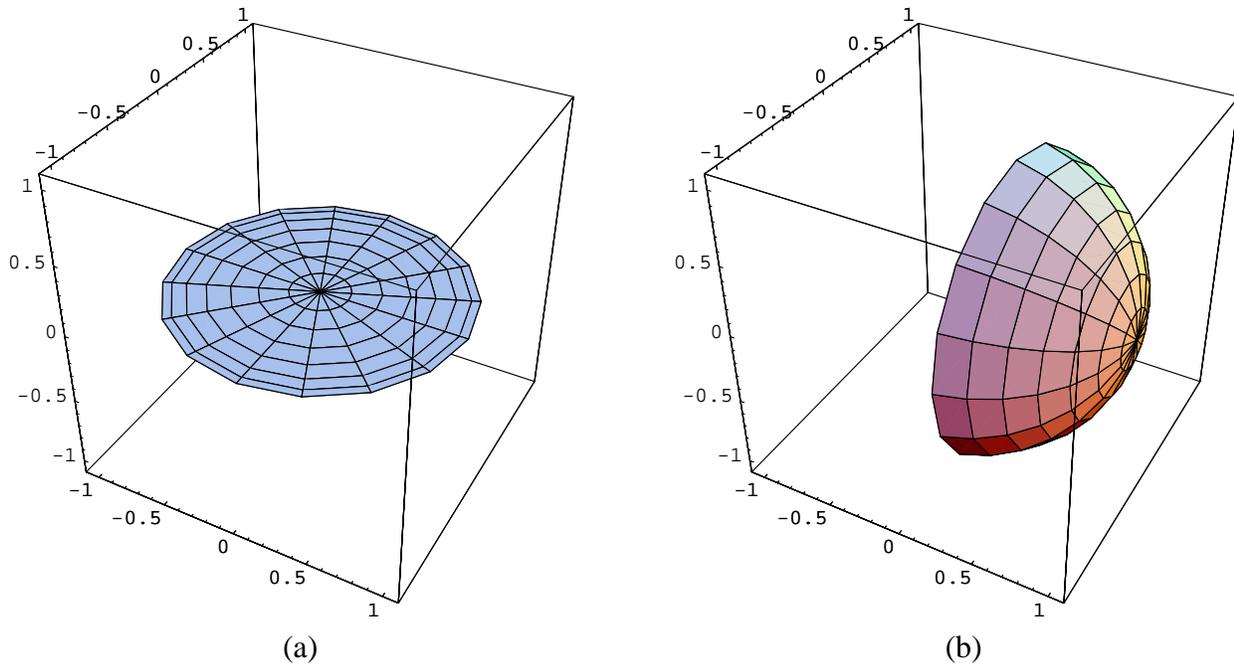


Figure 17: The Geodesic reference frame tilts to an ambiguous result as the tilt angle approaches π , the inverted direction of the chosen reference frame. We see two different 3D projections of the quaternion surface, (a) giving the vector coordinates (q_1, q_2, q_3) , and (b) the coordinates (q_0, q_1, q_2) . The center is the North pole, the middle ring is the equator, and outer circle is in fact the space of possible frames at the South pole of the sphere: there is no unique way to tilt the North pole to the orientation of the South pole, as there is a full circle of arbitrariness in the choice.

reference frame in quaternion space (easily seen in ordinary 3D space as well) eventually reaches a quaternion circle representing the ambiguous orientation of the frame with reference direction along the $-\hat{z}$ axis. This phenomenon is a practical consequence of the fact that the two-sphere does not admit a global vector field: according to classical manifold theory (see, e.g., Milnor or Grimm and Hughes [29, 13]), one needs at least two separate patches, one for the North pole and one for the South pole, to place a complete set of coordinates (or equivalently, for our problem, a set of frames) on a sphere.

The more mathematical approach requires that interesting surfaces be defined as a collection of patches [13, 7], and the spaces of frames for each patch must be matched up and sewn together by assigning a transition function along the boundaries. There are a variety of ways one can approach the problem of taking a manifold and associating fiber bundles with it; the most relevant fiber bundle for the context of the current problem is the *space of moving frames* of the space \mathbb{R}^3 in which the surface is embedded [7, 40]. We in fact move as usual from the space of frames to the space of associated quaternions. Then at each point x of a patch we have frames that are matrix-valued functions from the patch into the group $SU(2)$ of quaternion frames (which we treat as the topological space S^3). We can express the relationship between the frames q and q' of two neighboring patches U and U' , represented as quaternions, via quaternion multiplication by a transition function t :

$$q' = t * q .$$

We may in fact explicitly construct the transition functions between the two patches as quaternion maps, giving a quaternion version of one of the classical procedures of manifold theory. In Figure 18(a), we show the the projective coordinates on the sphere that produced the set of coordinate frames in Figure 10(c), which are essentially equivalent to those in (b) sampled at polar coordinate values. Using the polar coordinate sampling, so that we can easily identify the equator, we show in Figure 18(b) the quaternion maps corresponding to the coordinate frames derived from this orthonormal coordinate system covering the North pole (disk in center) and the South pole (smashed side view of a hemisphere in the 3D projection with its $q \rightarrow (-q)$ partner). These coordinate systems agree at exactly *one* point on the equator, which is (almost) evident from the figure; note that we have chosen to display the coordinate systems only up to the equator, unlike the patches of Figure 17, which cover the entire sphere except for one pole.

In order to establish a mapping covering the complete sphere, we must write down an explicit correspondence between the quaternion frames for each patch at each shared point on the equator. In Figure 19(a), we show the geodesic arcs on S^3 symbolizing the transition rotation

$$t(\theta) = q_{\text{south}}(\theta) * q_{\text{north}}^{-1}(\theta)$$

at each point on the equatorial circle parameterized by θ . Note carefully the order of quaternion multiplication; with our conventions a different order will not work. The arcs themselves are actually segments of the space of possible frames, since the simplest rotation between two frames with the same normal (at the same point on the equator) is a geodesic rotation about that normal. Figure 19(b) and (c) shows the transition functions $q(\theta)$ sampled at regular intervals in θ and referred to the origin $(1, 0, 0, 0)$ in quaternion space. Each quaternion point at the end of an arc

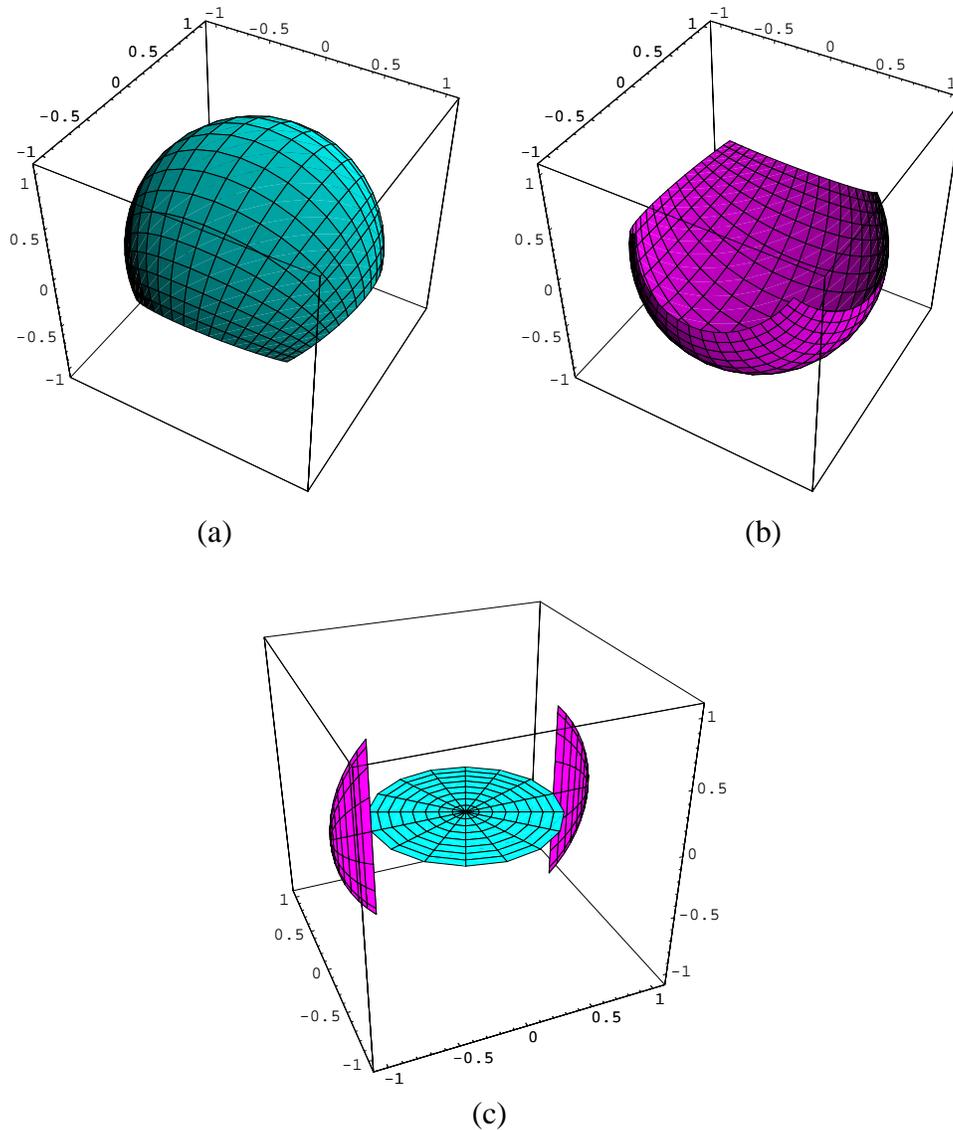


Figure 18: (a) The North pole projective coordinatization of the sphere; (b) a similar regular patch for the South pole. Because of the “no-hair” theorem, no single regular patch can cover the entire sphere. (c) The quaternion mappings of the systems of frames given by the North and South pole coordinate patches, sampled in polar coordinates. The $q \rightarrow (-q)$ reflected images are included, though the North pole’s images both have the same projection and are thus indistinguishable here. The maps in (c) extend only to the equator, unlike the patches given in Figure 17.

represents a rotation to be applied to a point on the North pole patch equator to obtain the coordinate frame at the corresponding point on the South pole patch equator. One point is in fact the identity, and there is some degeneracy due to reflection symmetry across the equator.

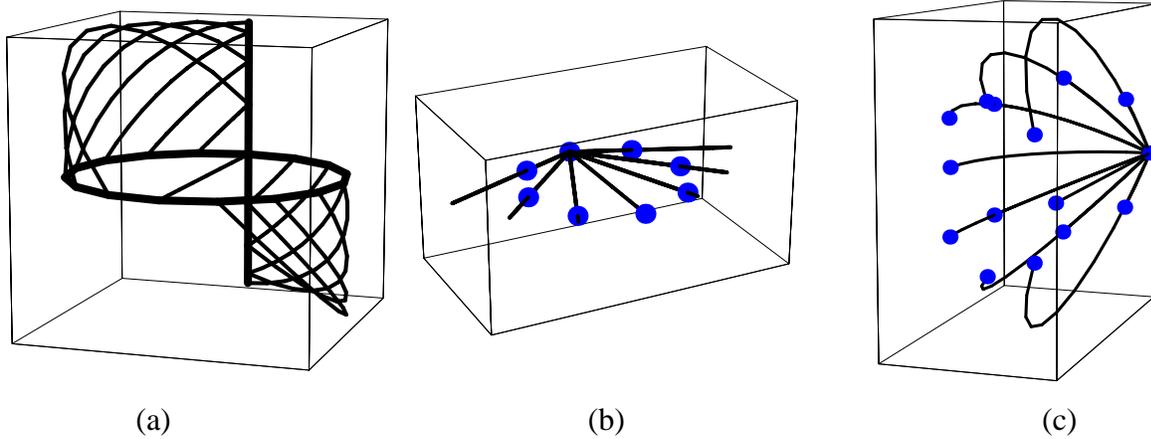


Figure 19: (a) The transition functions from the North pole frame to the South pole frame as arcs in the three-sphere. These arcs are pieces of the space of possible frames with a given normal on the equatorial point. (b) A representation of the transition functions as arcs from the origin in rotation space (the pole $(1, 0, 0, 0)$ in quaternion space) common to all the arcs here. The ends of the arcs thus represent the actual rotation needed to match the coordinate systems at each point on the equator. (c) A different projection from 4D to 3D, showing more details of the structure of the transition function arcs, which have a two-fold degeneracy in the standard projection (b).

4 The Space of Frames

We are at last ready to introduce the key concept of the *space of possible frames*.

Suppose at each sample point $\mathbf{x}(t)$ of a curve, we are given a unit tangent vector, $\hat{\mathbf{T}}(t)$, computed by whatever method one likes (two-point sampling, five-point sampling, analytic, etc.). Then one can immediately write down a one-parameter family describing all possible choices of the normal plane orientation: it is just the set of rotation matrices $R(\theta, \hat{\mathbf{T}}(t))$ (or quaternions $q(\theta, \hat{\mathbf{T}}(t))$) that leave $\hat{\mathbf{T}}(t)$ fixed.

For surfaces, the analogous construction follows from determining the unit normal $\hat{\mathbf{N}}(u, v)$ at each point $\mathbf{x}(u, v)$ on the surface patch. The needed family of rotations $R(\theta, \hat{\mathbf{N}}(u, v))$ (or quaternions $q(\theta, \hat{\mathbf{N}}(u, v))$) now leaves $\hat{\mathbf{N}}(u, v)$ fixed and parameterizes the space of possible *tangent* directions completing a frame definition at each point $\mathbf{x}(u, v)$.

However, there is one slight complication: the family of frames $R(\theta, \hat{\mathbf{v}})$ leaving $\hat{\mathbf{v}}$ fixed does not have $\hat{\mathbf{v}}$ as one column of the 3×3 rotation matrix, and so does not actually describe the desired family of frames. Therefore we proceed as follows:

We define $f(\theta, \hat{\mathbf{v}}) = (f_0, f_1, f_2, f_3)$ to be a quaternion describing the family of frames for which the direction $\hat{\mathbf{v}}$ is a preferred fixed axis of the frame, such as the tangent or normal vectors. The orthonormal triad of 3-vectors describing the desired frame is

$$F(\theta, \hat{\mathbf{v}}) = \begin{bmatrix} f_0^2 + f_1^2 - f_2^2 - f_3^2 & 2f_1f_2 - 2f_0f_3 & 2f_1f_3 + 2f_0f_2 \\ 2f_1f_2 + 2f_0f_3 & f_0^2 - f_1^2 + f_2^2 - f_3^2 & 2f_2f_3 - 2f_0f_1 \\ 2f_1f_3 - 2f_0f_2 & 2f_2f_3 + 2f_0f_1 & f_0^2 - f_1^2 - f_2^2 + f_3^2 \end{bmatrix}, \quad (43)$$

where one column, typically the 3rd column, must be $\hat{\mathbf{v}}$.

The standard rotation matrix $R(\theta, \hat{\mathbf{v}})$ leaves $\hat{\mathbf{v}}$ fixed but does not have $\hat{\mathbf{v}}$ as one column of the 3×3 rotation matrix, and so we have more work to do. To compute $f(\theta, \hat{\mathbf{v}})$, we need the following:

- A base reference frame $b(\hat{\mathbf{v}})$ that is guaranteed to have, say, the 3rd column exactly aligned with a chosen vector $\hat{\mathbf{v}}$, which is either the tangent to a curve or the normal to a surface.
- A one-parameter family of rotations that leaves a fixed direction $\hat{\mathbf{v}}$ invariant.

The latter family of rotations is given simply by the standard quaternion

$$q(\theta, \hat{\mathbf{v}}) = \left(\cos \frac{\theta}{2}, \hat{\mathbf{v}} \sin \frac{\theta}{2} \right), \quad (44)$$

for $0 \leq \theta < 4\pi$, while the base frame can be chosen as

$$b(\hat{\mathbf{v}}) = q(\arccos(\hat{\mathbf{z}} \cdot \hat{\mathbf{v}}), (\hat{\mathbf{z}} \times \hat{\mathbf{v}}) / \|\hat{\mathbf{z}} \times \hat{\mathbf{v}}\|). \quad (45)$$

We refer hereafter to the frame $b(\hat{\mathbf{v}})$ as the *Geodesic Reference Frame* because it tilts the reference vector $\hat{\mathbf{z}}$ along a geodesic arc until it is aligned with $\hat{\mathbf{v}}$; see Figure 20. If $\hat{\mathbf{v}} = \hat{\mathbf{z}}$, there is no problem, since we just take $b(\hat{\mathbf{v}})$ to be the quaternion $(1, \mathbf{0})$; if $\hat{\mathbf{v}} = -\hat{\mathbf{z}}$, we may choose any

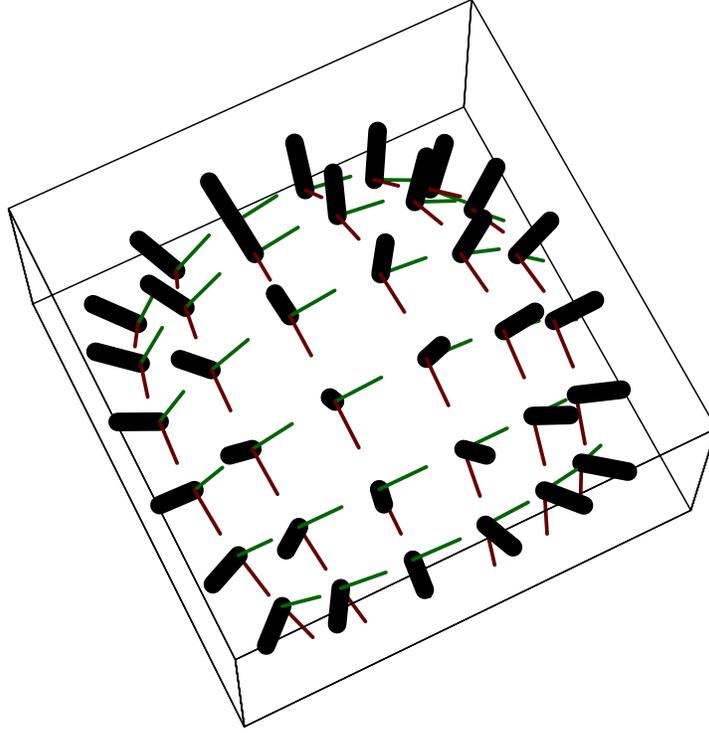


Figure 20: Example of the Geodesic Reference Frame: on the northern hemisphere of a 2-sphere, the Geodesic Reference Frame tilts the \hat{z} axis of the north pole's identity frame along the shortest arc to align with a specified reference direction.

compatible quaternion such as $(0, 1, 0, 0)$. We escape the classic difficulty of being unable to assign a global frame to all of S^2 because we need a parameterization of *all possible* frames, not any one particular global frame. If one wants to use a reference frame that is not the identity frame, one must premultiply $b(\hat{v})$ on the right by a quaternion rotating from the identity into that reference frame; this is important when constructing a nonstandard Geodesic Reference Frame such as that required to smoothly describe a neighborhood of the southern hemisphere of S^2 .

We can thus write the full family of possible quaternion frames keeping \hat{v} as a fixed element of the frame triad to be the quaternion product

$$f(\theta, \hat{v}) = q(\theta, \hat{v}) * b(\hat{v}), \quad (46)$$

where $*$ denotes quaternion multiplication and all possible frames are described twice since $0 \leq \theta < 4\pi$. To summarize, if we specify a frame axis \hat{v} to be fixed, then the variable θ in $f(\theta, \hat{v})$ serves to parameterize a *ring* in quaternion space, each point of which corresponds to a particular 3D frame, and each frame has a diametrically opposite twin.

We argue that, since optimization will typically be done in the full quaternion space, the fact that two opposite-sign quaternions map to the same physical three-space rotation is not a detriment; in fact, it potentially permits an additional stability in the variational process since rotations by $+\pi$

and $-\pi$ are not close to each other in quaternion space as they are in ordinary rotation matrices. In principle, any quaternion Gauss map can be replaced by its exact negative, and the variational process could converge from an ambiguous starting point to either one; the frames would be the same. In our standard projection, the two reflection-equivalent maps are inversions of one another about the 3D origin; their unseen opposite q_0 values can of course cause an additional large separation of the maps in 4D space.

4.1 Full Space of Curve Frames

We can now construct the space of frames step by step using the method above. In Figure 21, we illustrate various views of the construction of the space of frames for the trefoil knot, beginning with a few tangent vectors and the quaternion basis frames corresponding to quaternions that tilt the reference axis into this tangent direction. The circular curve of quaternions representing the space of normal frames is drawn for each tangent; each basis frame touches this curve once. Then the family of these circular curves sweeps out a cylindrical two-manifold, the full space of invariant frames for a 3D curve.

This space has several nontrivial properties. One is that, given one circular ring of frames, a neighboring ring that is a parallel-transported version of the first ring is a so-called “Clifford parallel” of the first ring: the distance from any point on one ring to the nearest point on the second ring is the same. This is nontrivial to visualize and is a feature of the 4D space we are working in. Another property is that the intervals between rings in the quaternion space directly indicate the curvature. This comes about because the magnitude of \hat{T}' is related to the parallel transport transition between any two sample points, given by Eq. (17); since the parallel transport frames are legal frames, and since the starting frame is arbitrary, each full ring is a parallel transport of its predecessor, with the angular distance of the transition rotation providing a measure of the curvature relative to the sampling interval.

4.2 Full Space of Surface Maps

The full space of frames for a surface patch is even more complex to visualize, since it is a “hyper-cylindrical” 3-manifold, formed by the direct product of patches of surface area with the rings of possible frames through each surface point.

As a very simple case of a surface, consider the patch introduced at the beginning of the paper in Figure 2(a). The coordinate system used does not provide a unique tangent frame, and so one cannot immediately determine a logical frame choice.

In Figure 22, we show spaces of possible frames for the four corners as four rings of quaternion values compatible with the normals at the patch corners. Parallel transporting the initial frame along the two different routes in Figure 2(b,c) produces incompatible frames at the final corner; we represent this situation in Figure 22 by drawing the routes in quaternion space between the initial frame (the degenerate circle appearing as a central vertical line) and the final frame; the mismatch between the two final frames is illustrated by the fact that the two paths meet at different points on the final ring specifying the frame freedom for the bottom corner’s frame.

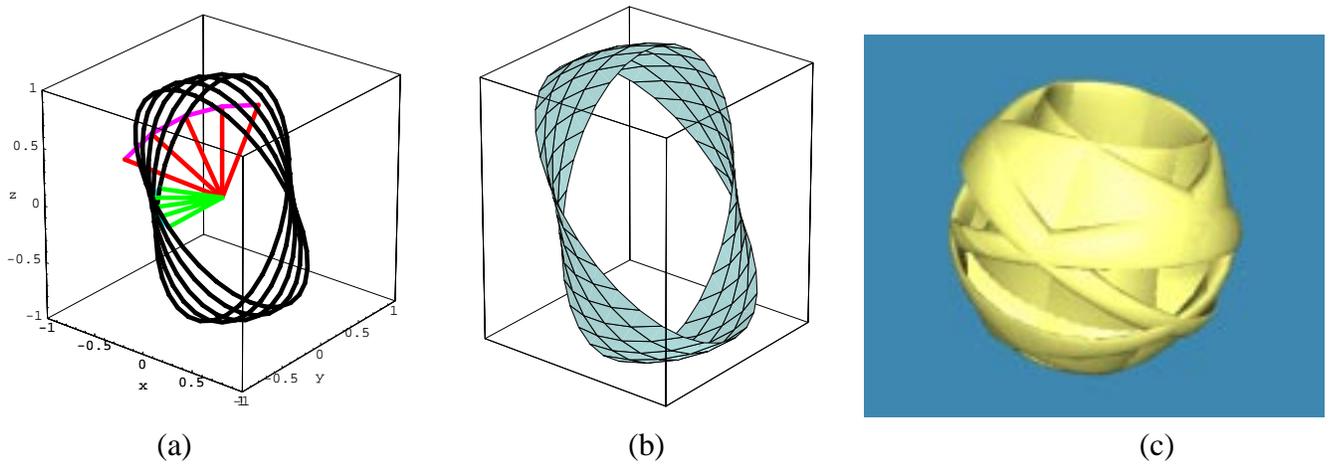


Figure 21: (a) The first several pieces of the construction of the invariant quaternion space for the frames of the trefoil knot. The red fan of vectors shows the first several elements of the tangent map, represented as vectors from the origin to the surface of the two-sphere and connected by a line. Each green vector points from the origin to the Geodesic Reference element of the quaternion space $q(\arccos(\hat{\mathbf{t}} \cdot \hat{\mathbf{z}}), \hat{\mathbf{t}} \times \hat{\mathbf{z}} / \|\hat{\mathbf{t}} \times \hat{\mathbf{z}}\|)$ guaranteed to produce a frame with the tangent $\hat{\mathbf{t}}$. The black curves are the first several elements of the one-parameter space of quaternions representing *all possible* quaternion frames with the tangent $\hat{\mathbf{t}}$. (b) This piece of the space of possible frames represented as a continuous surface, where a circle on the surface corresponds to the space of frames for one point on the curve. (c) The rest of the full constraint space for the trefoil knot. All quaternions are projected to 3D using only the vector part.

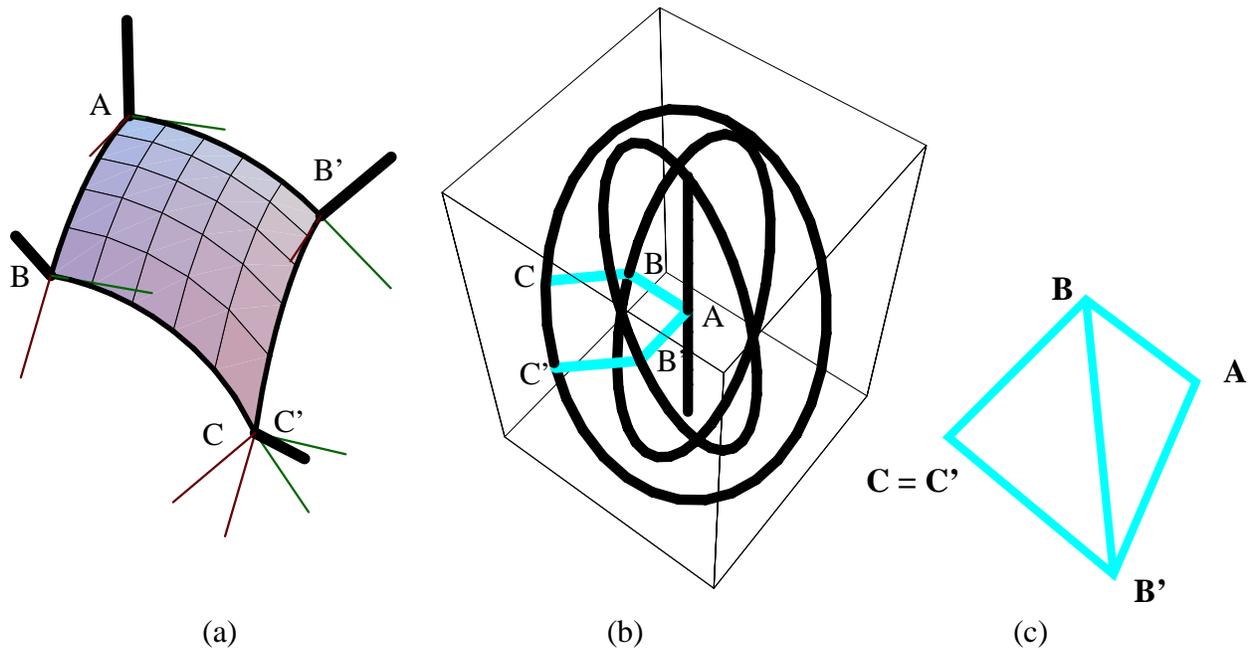


Figure 22: A different viewpoint of the mismatch problem of Figure 2. (a) Choosing different routes to determine the frame at the bottom point results in the incompatible frames shown here in 3D space. (b) The same information is presented here in the quaternion space-of-frames picture. We use throughout a quaternion projection that shows only the 3-vector part of the quaternion, dropping q_0 ; this is much like projecting away z in a polar projection of the 2-sphere. Each heavy black curve is a ring of possible frame choices that keep fixed the normals in (a); the labels mark the point in quaternion space corresponding to the frames at the corners in (a), so the gap between the labels C and C' represents the frame mismatch in quaternion space on the same constraint ring. (The apparent vertical line is the result of drawing a squashed circle of frames at vertex A in this projection.) (c) The method proposed in this paper to resolve this conflict is to fix one point, say A , divide the polygon $ABCB'$ into triangles, and slide B , C , and B' along the constraint rings until the total triangle areas are minimized, and some compromise with $C = C'$ is reached.

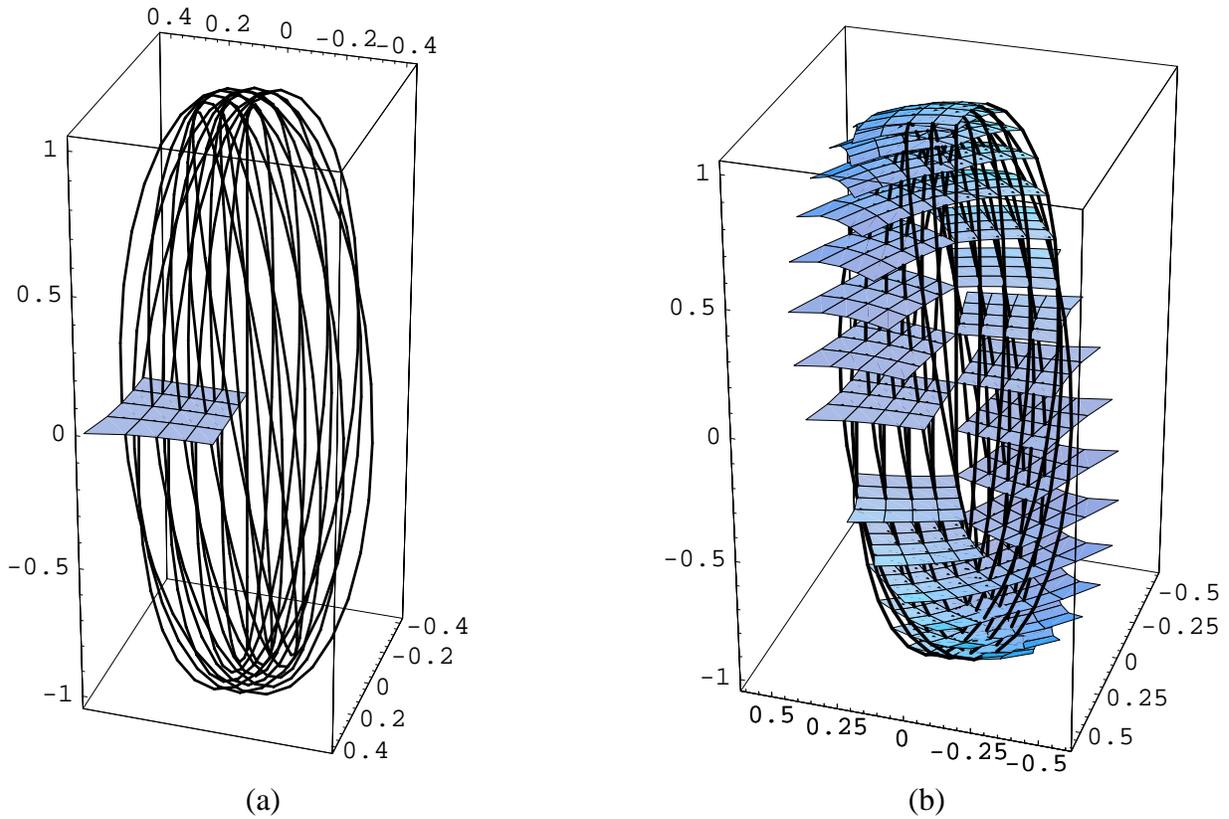


Figure 23: (a) A more complete picture of the space of frames for this surface patch; the surface shown is a sparse quaternion frame choice for the surface, and we show a subset of the rings of constraints. Each ring passes through one quaternion point on the frame map, the point specifying the current frame choice. Variations must keep each vertex on its ring. (b) An equivalent set of frames is formed by applying a rotation to the entire set of frames. All points follow their own ring of constraints to keep the same normal, These pictures represent the *three-manifold* in quaternion space swept out by the possible variations.

Sliding and Overall Rotational Freedom. In Figure 23(a), We go one step further, and first show how the quaternion Gauss map of an entire patch is situated relative to the ring space; keeping one corner fixed and sliding the rest of the frames around the circular rings takes us to distinct families of frames, which obviously have different areas in the quaternion space. Finally, in Figure 23(b), we keep the fundamental space of frames the same, but exercise the freedom to choose the single parameter describing the basis for the overall orientation; rotating the basis sweeps out both the three-manifold describing the space of frames for this patch, and the family of equivalent frames differing by an insignificant orientation change in the basis vector.

In order to resolve the frame choice ambiguity, one needs a systematic approach; we propose in the next section to accomplish this by optimizing appropriate quantities, e.g., by minimizing the area of the quaternion Gauss map in quaternion space.

We remark that the general features of the surface curvature can in principle be noted from the space of possible frames in a similar manner to that for curves. The family of curves through any point spanning the surfaces tangent space at that point possesses a family of rings parallel to the space of frames at the point, allowing estimates of the rates of change in different directions; the principal curvatures then correspond to the maxima and minima.

5 Choosing Paths in Quaternion Space

We have now seen that the space of possible frames at any point of a curve or surface thus takes the form of a great circle on the unit three-sphere representing the unit quaternions in 4D Euclidean space. While diametrically opposite points on this circle represent the same frame in 3D space, the two-fold redundancy can actually be an advantage, since it helps avoid certain types of wrap-around problems encountered when trying to find paths in the space. Our task then is to select a set of values of the parameter on each of these great circles.

The advantage of looking at this entire problem in the space of quaternions is that one can clearly compare the intrinsic properties of the various choices by examining such properties as length and smoothness in the three-sphere. We note the following issues:

- **Frame-frame distance.** Suppose we are given two neighboring tangents, $\hat{\mathbf{t}}_1$ and $\hat{\mathbf{t}}_2$, and two corresponding candidate frame choices parameterized by θ_1 and θ_2 . What is the “distance” in frame space between these? The simplest way to see how we should define the distance is by observing that, by Euler’s fundamental theorem, there is a single rotation matrix $R(\theta, \hat{\mathbf{n}})$ That takes one frame to the other; if $R_1(\theta_1, \hat{\mathbf{t}}_1)$ and $R_2(\theta_2, \hat{\mathbf{t}}_2)$ are the two frames, then one can write $R = (R_2 \cdot (R_1)^{-1})$ and solve for θ and $\hat{\mathbf{n}}$. Clearly the value of θ gives a sensible measure of the closeness of the two frames.
- **Quaternion distance.** We remark that essentially the same procedure is required to obtain the parameters of R directly or to find the value of the equivalent quaternion. If we work in quaternion space, we compute $q_1(\theta_1, \hat{\mathbf{t}}_1)$ and $q_2(\theta_2, \hat{\mathbf{t}}_2)$, and then find rather more straightforwardly an equivalent result by noting that the zeroth component of $q = q_2 * (q_1)^{-1}$ is identical to the rotation-invariant scalar product of the two quaternions, $q_1 \cdot q_2$, and thus provides the

needed angle at once:

$$\theta = 2 \arccos(q_1 \cdot q_2) .$$

- **Approximation by Euclidean distance.** One can in principle compute quaternions in polar coordinates and use the induced metric on the sphere to compute precise arc-length distance integrals. However, one generally can expect to be dealing with fine tessellations of smoothly varying geometric objects; in this case, it may be sufficient for numerical purposes to estimate frame-to-frame distances using the Euclidean distance in R^4 , since the chord of an arc approximates the arc length well for small angles.

Optimal Path Choice Strategies. Why would one want to choose one particular set of values of the frame parameters over another? The most obvious is to keep a tubing from making wild twists such as those that occur for the Frenet frame of a curve with inflection points. In general, one can imagine wanting to minimize the total twisting, the aggregate angular acceleration, etc., subject to a variety of boundary conditions. A bewildering variety of energy functions to minimize can be found in the literature (see, e.g., [6]). We summarize a selection of such criteria for choosing a space of frames below, with the caveat that one certainly may think of others!

- **Minimal Length and Area.** The most obvious criterion is to minimize the total turning angle experienced by the curve frames. Fixing the frames at the ends of a curve may be required by periodicity or external conditions, so one good solution would be one that minimizes the sum total of the turning angles needed to get from the starting to the ending frame. The length to minimize is just the sum of the angles rotated between successive frame choices, as noted above, either exact or approximate. Similar arguments apply to the area of a surface’s quaternion Gauss map.
- **Parallel Transport along Geodesics.** Given a particular initial frame, and no further boundary constraints, one may also choose the frame that uses the minimum *local* distance to get between each neighboring frame. Since the parallel transport algorithm corresponding to the Bishop frame uses precisely the smallest possible rotation to get from one frame to the next, this gives the minimal free path that could be computed frame-by-frame. On a surface, the resulting paths are essentially geodesics, but, as noted in Figure 2, there is no obvious analog of a global parallel transport approach to surface framing.
- **Minimal Acceleration.** Barr, Currin, Gabriel, and Hughes [2] proposed a direct generalization of the “no-acceleration” criterion of cubic Euclidean splines for quaternion curves constrained to the three-sphere; the basic concept was to globally minimize the squared tangential acceleration experienced by a curve of unit quaternions. Though the main application of that paper was animation, the principles are entirely valid for numerically computing optimal frames for curves and surfaces in our context.
- **Keyframe splines and constraints.** If for some reason one must pass through or near certain specified frames with possible derivative constraints, then a direct spline construction in the quaternion space may actually be preferred; see, e.g., [37, 35, 31, 39, 23]. Most splines can

be viewed in some sense as solving an optimization problem with various constraints and conditions, and so the keyframe problem essentially reverts again to an optimization.

General Remarks. For both curves and surfaces, there is a single degree of freedom in the frame choice at each point where we have sampled the tangent or normal direction, respectively. This degree of freedom corresponds to a relatively common “sliding ring” constraint that occurs frequently in minimization problems. General packages for solving constraints are mentioned in Barr, et al. [2], who chose MINOS [30]. For our own experiments, we have chosen Brakke’s Surface Evolver package [6], which has a very simple interface for handling parametric constraints as “boundary” conditions, and can be used for a wide variety of general optimization problems. Two enhancements to the Evolver have recently been added to handle the specific issues related to quaternion optimization; a new symmetry “`symmetry_group "central_symmetry" ”` identifies the quaternion q with $-q$ if desired during the variation to prevent reflected double traversals like that in Figure 13 from varying independently, and the system is now able to use the pull-back metric on the sphere

$$ds^2 = \sum_{i,j} dx_i dx_j r^{-4} (r^2 \delta_{i,j} - x_i x_j)$$

to compute distances directly on the three-sphere. Computation using the metric, however, is very slow, and so in practice we have used the Euclidean \mathbb{R}^4 chord approximation, which works quite well for closely spaced samples and is much faster. Yet another alternative proposed by Brakke (private communication) is to use periodic coordinates on \mathbb{S}^3 of the form

$$(x_1 = \sin r \cos s, x_2 = \sin r \sin s, x_3 = \cos r \cos t, x_4 = \cos r \sin t),$$

and to vary directly on an \mathbb{R}^3 space with $(x = r, y = s, z = t)$ and the metric

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & \sin^2 x & 0 \\ 0 & 0 & \cos^2 x \end{bmatrix}.$$

Our own use of the Evolver required only changing the parameter “`#define BDRYMAX 20`” in `skeleton.h` to the desired (large) value and recompiling. Then, remembering to set “`space_dimension 4`” when working in \mathbb{R}^4 , one needs in addition a piece of code similar to the following MATHEMATICA fragment to define the boundary constraints for each fixed vector (tangent or normal) and the chosen initial quaternion frame:

```
Do[ringeqn = Qprod[makeQfromVec[veclist[[i]],P1],
    q0list[[i]]//Chop;
  Write[file, " boundary ",i," parameters 1"];
  Write[file, "x1: ", CForm[ ringeqn[[2]]];
  Write[file, "x2: ", CForm[ ringeqn[[3]]];
  Write[file, "x3: ", CForm[ ringeqn[[4]]];
  Write[file, "x4: ", CForm[ ringeqn[[1]]],
    {i,1,Length[veclist]}]
```

Note that, since Evolver displays only the first three coordinates, we have moved the scalar quaternion to the end; then the Evolver will display our preferred projection automatically.

General Remarks on Optimization in Quaternion Space. Numerical optimization remains a bit of an art, requiring patience and resourcefulness on the part of the investigator. We found, for example, that curve optimization was relatively more stable than surface optimization because single curve outliers add huge amounts to the length, whereas single surface points stuck in a far away crevice may contribute only a tiny amount to the area of a large surface. Although the Evolver in principle handles spherical distances, we used the default 4D Euclidean distance measure as an approximation; this generally corresponded well to explicit area calculations using solid angle performed on the same data sets. However, we did find that extremely random initial conditions (unrealistic for most applications), could produce isolated points stuck in local minima diametrically across quaternion space, at $q \rightarrow -q$, from where they should be. This type of problem can be largely avoided simply by running a consistency preprocessor to force nearby neighbors to be on the same side of the three-sphere. Another useful technique is to organize the data into hierarchies and optimize from the coarse scale down to the fine scale. In other cases when things seem unreasonably stuck, a manual “simulated annealing” procedure like that afforded by the Evolver’s `jiggle` option often helps.

6 Examples

We now present some examples of frame choices computed using the Evolver to minimize the length of the total path among sliding ring constraints for selected curves, and the total area spanned by analogous sliding rings for surfaces. One interesting result is that there appear to be families of distinct minima: if the initial data for a periodic surface, for example, are set up to return directly to the same point in quaternion space after one period, one has two disjoint surfaces, one the $q \rightarrow (-q)$ image of the other; if the data do not naturally repeat after one cycle, they must after two, since there are only two quaternion values that map to the same frame. The family of frame surfaces containing their own reflected images have a minimum distinct from the disjoint family.

Minimal Quaternion Frames for Space Curves. The helix provides a good initial example of the procedure we have formulated. We know that we can always find an initial framing of a curve based on the Geodesic Reference algorithm; however, suppose we wish to impose minimal length in quaternion space on the framing we select, and we do not know whether this frame is optimal with respect to that measure. Then, as illustrated in Figure 24, we can send the ring constraints on the possible quaternion frames at each sample point to the Surface Evolver and let it automatically find the optimal framing. The results and energies for several stages of this evolution are shown in the figure; the final configuration is indistinguishable from the Parallel Transport frame, confirming experimentally our theoretical expectation that parallel transport produces the minimal possible twisting.

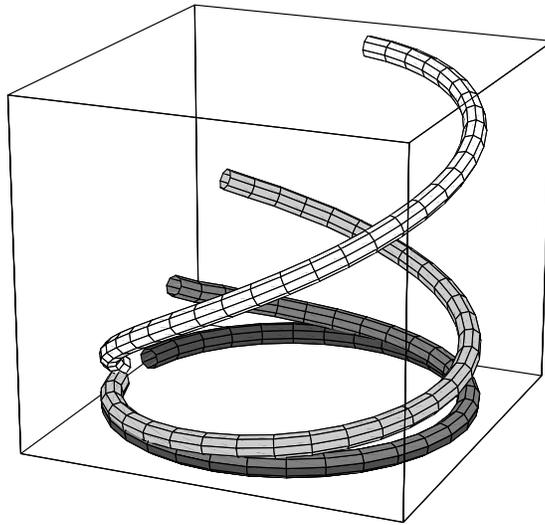


Figure 24: Starting from the Geodesic Reference quaternion frame for a single turn of the helix, the very dark gray circle, the Evolver produces these intermediate steps while minimizing the total quaternion curve length subject to the constraints in the space of frames. The final result is the white curve, which is identical to several decimal points with the Parallel Transport quaternion frame for the same helix. The numerical energies of the curves, from dark to light in color, are 3.03, 2.91, 2.82, and 2.66 for the Parallel Transport frame. The individual tubings used to display these curves are in fact created using the Parallel Transport frame for each individual curve.

In Figure 1, we introduced the question of finding an optimal framing of a particular (3,5) torus knot whose almost-optimal Parallel Transport framing was not periodic. In Figure 25, we show the solution to this problem achieved by clamping the initial and final quaternion frames to coincide, then letting the Evolver pick the shortest quaternion path for all the other frames. It would be possible, as in the case of the (2,3) torus knot framing shown in Figure 13, to have different conditions produce a framing solution containing its own reflected image rather than having a distinct reflected image as is the case for Figure 25.

The types of solutions we find are remarkable in that they should be essentially the same for all reparameterizations of the curve; regardless of the spacing of the sampling, the continuous surface of possible frames is geometrically the same in quaternion space, so paths that are minimal for one sampling should be approximately identical to paths for any reasonable sampling. On the other hand, if we *want* special conditions for certain parameter values, it is easy to fix any number of particular orientations at other points on the curve, just as we fixed the starting points above; derivative values and smoothness constraints leading to generalized splines can be similarly specified (see [2]).

Surface Patch Framings. A classic simple example of a surface patch framing problem was presented in the discussion of Figures 2 and 22. This problem can also be handled by the Evolver:

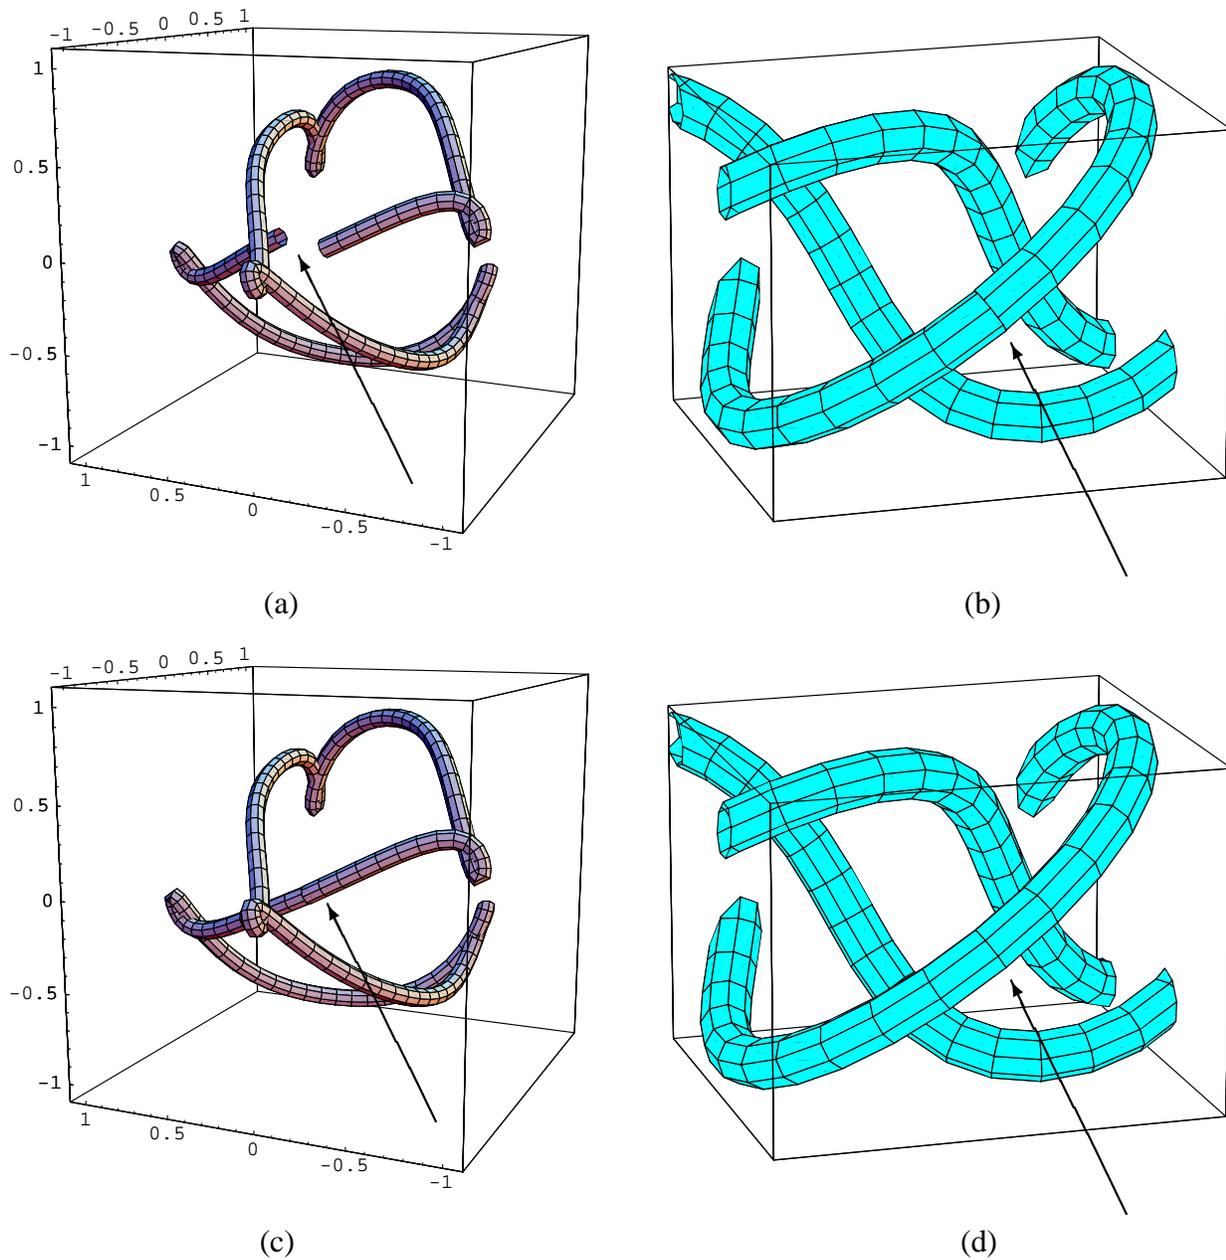


Figure 25: Optimization of the non-periodic parallel transport frame of the (3,5) torus knot introduced in Figure 1 to produce a nearby periodic framing. (a) The original quaternion parallel transport frame used to produce the tubing in Figure 1(b,c). (b) The frame mismatch, repeated for completeness. (c) The result of fixing the final frame to coincide with the initial frame, leaving the other frames free to move on the constraint rings, and minimizing the resulting total length in quaternion space. The length of the original curve was 13.777 and that of the final was 13.700, not a large difference, but noticeable enough in the tube and the quaternion space plot. (d) Closeup of the corresponding framing of the knot in ordinary 3D space, showing that the mismatch problem has been successfully resolved. This tube can *now* be textured, since the frames match exactly.

we choose an initial quaternion frame for the mesh corresponding to one of the arbitrary choices noted, and minimize the area in quaternion space subject to the constraints that the normals remain unchanged, and hence the frame choices may only slide around the rings depicted in Figure 22(b). The results are shown in Figures 26, and 27. As a test, we started one case with a random initial state with a range of 2π in the starting values. All converged to the same optimal final framing. A basic observation is that while none of the standard guesses appeared optimal, the Geodesic Reference frame is very close to optimal for patches that do not bend too much.

Minimal Surfaces. Minimal surfaces possess many special properties following from the fact that the mean curvature is everywhere the vanishing sum of two canceling local principal curvatures [12]. We present a family of classic examples here that is remarkable for the fact that the usual framings are already very close or exactly optimal; thus we do not have much work to do except to admire the results, though there may be some interesting theorems implicit that would be beyond the scope of this paper to pursue.

In Figure 28(a,b,c), we present the following classical minimal surfaces:

$$\mathbf{x}_{\text{catenoid}}(u, v) = \cos u \cosh v \hat{\mathbf{x}} + \sin u \cosh v \hat{\mathbf{y}} + v \hat{\mathbf{z}} \quad (47)$$

$$\mathbf{x}_{\text{helicoid}}(u, v) = v \cos u \hat{\mathbf{x}} + v \sin u \hat{\mathbf{y}} + u \hat{\mathbf{z}} \quad (48)$$

$$\mathbf{x}_{\text{enneper}}(u, v) = (u - u^3/3 + uv^2) \hat{\mathbf{x}} + (v - v^3/3 + vu^2) \hat{\mathbf{y}} + (u^2 - v^2) \hat{\mathbf{z}} \quad (49)$$

The quaternion Gauss map choices determined by these parameterizations and by the Geodesic Reference algorithm are shown in Figure 29. The coordinate-based catenoid map and helicoid map are 4π double coverings, while Enneper's surface curiously has a coordinate system map that is exactly identical to the Geodesic Reference framing. For the periodic framings of the catenoid and helicoid, we find the noteworthy result that the Geodesic Reference frame, which has a disjoint quaternion reflected image, is a minimum under variations of the surface that is distinct from the quaternion frames derived from the coordinate systems which are *also* minima, but contain their own $q \rightarrow (-q)$ reflected images. The Enneper surface quaternion frames, which are the same, appear to move very slightly around the borders under minimization, but it is not clear to what extent this is significant as opposed to a numerical border effect in the variation. The resulting 3D frame triads are shown in Figure 30 for comparison. A theoretical analysis of the general properties of quaternion Gauss maps for minimal surfaces is beyond the scope of this paper, but experimentally we see that there could be very interesting general properties.

Manifolds. For general manifolds, one must treat patches one at a time in any event, since global frames may not exist at all. Although the locally optimal patches cannot be globally joined to one another, we conjecture that some applications might benefit from the next best thing: matching boundary frames of neighboring patches using transitional rotations (see, e.g., [29, 13]). We have carried this out explicitly for simple cases, but omit it here for brevity.

Extensions to Other Domains. We have focussed for expository purposes in this paper on frames with intrinsic natural constraints imposed by the tangents to curves and normals to sur-

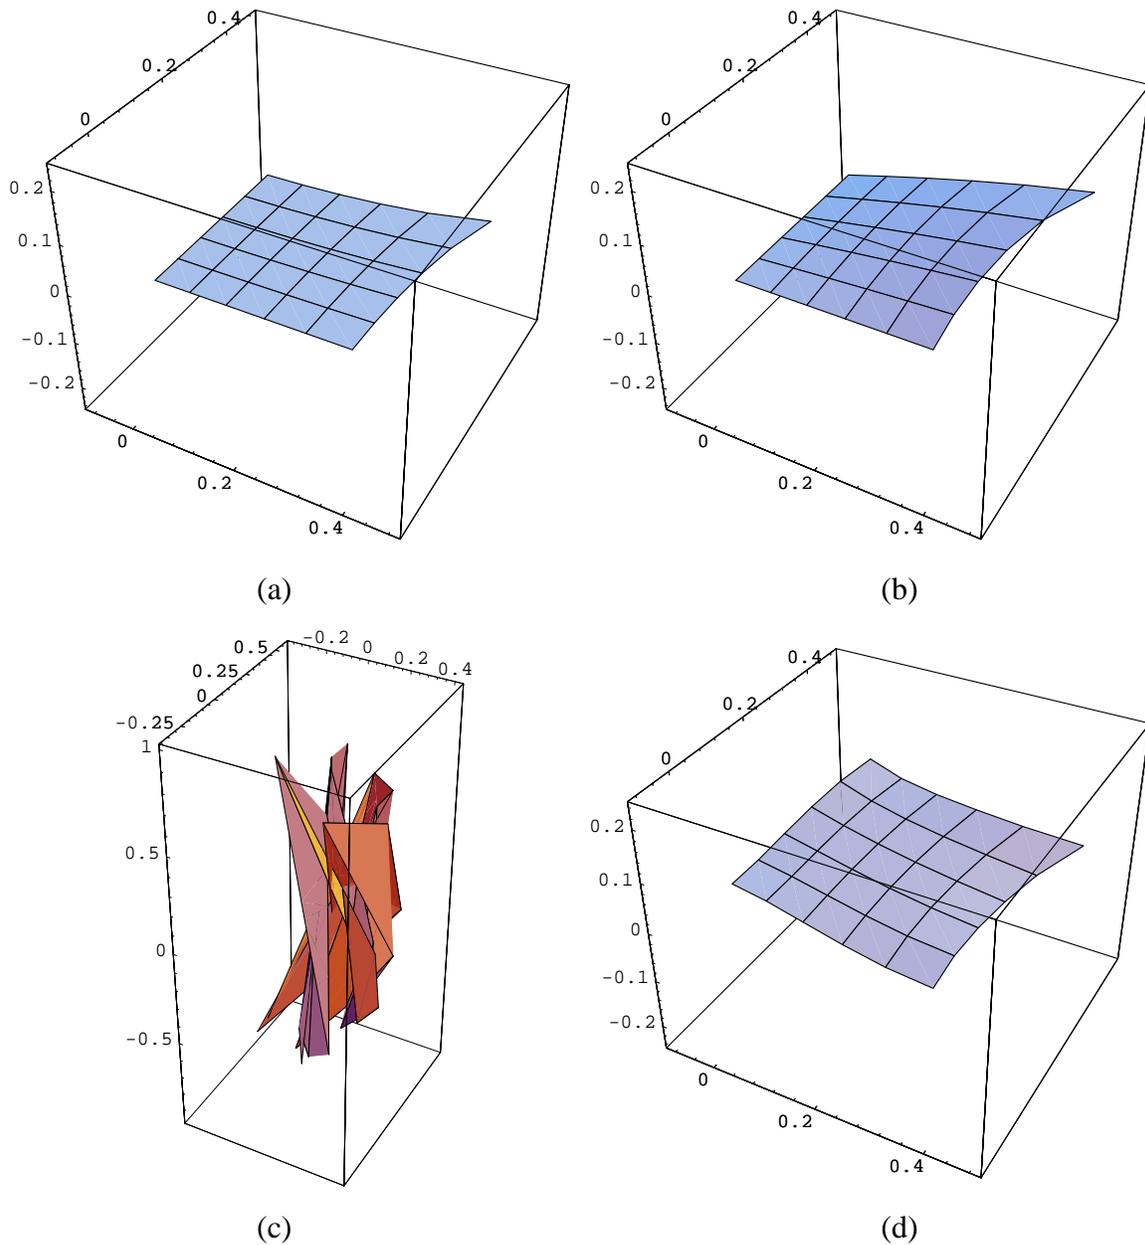


Figure 26: (a) The initial Geodesic Reference quaternions for the small patch shown in Figure 2. (b) Initial quaternions from parallel transporting the vertex frame down one edge, and then across line by line. (c) A random starting configuration with the single same fixed corner point as (a) and (b) and a range of $-\pi$ to $+\pi$ relative to the Geodesic Reference frame. (d) The result of minimization of the quaternion area is the same for all starting points. The relative areas are: 0.147, 0.154, 0.296, and 0.141, respectively. Thus the Geodesic Reference is very close to optimal.

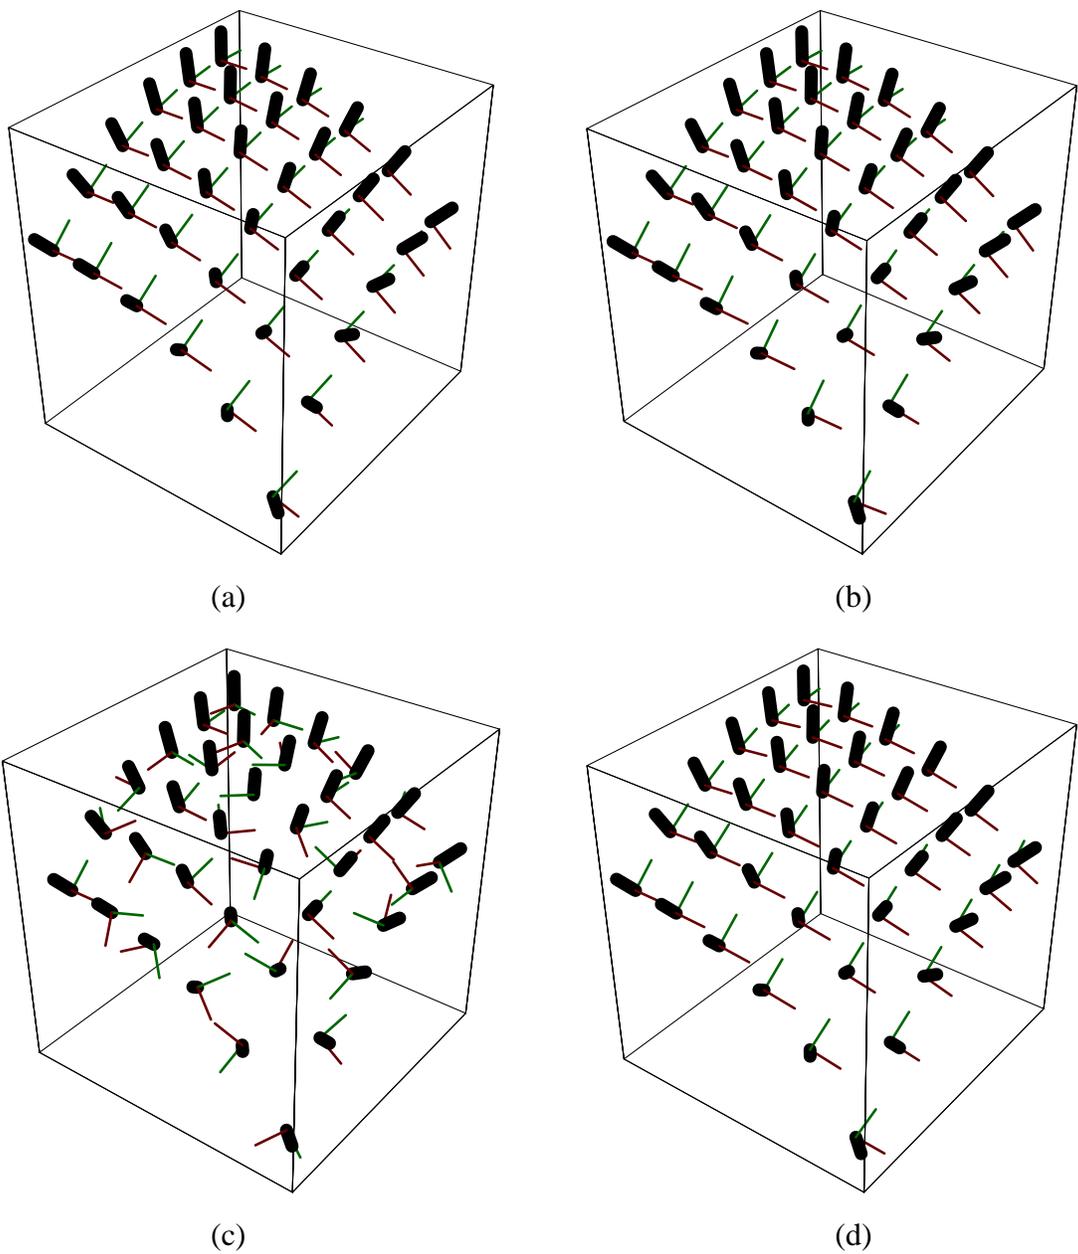


Figure 27: The 3D frame configurations corresponding to the quaternion fields in Figure 26. (a) The Geodesic Reference frame. (b) Two-step parallel transport frame. (c) Random frames. (d) The frame configuration resulting from minimizing area in quaternion space.

faces. However, the method extends almost trivially to applications involving externally specified constraints on frames. Geometric construction algorithms based on extrusions reduce to the tubing problem. For ordinary camera control interpolation, one could constrain any direction of the camera frame to be fixed by calculating its appropriate constraint ring in the quaternion Gauss map, and then extend a method like that of Barr et al. [2, 34]) to smoothly compute intermediate frames subject to the constraints. For more general constrained navigation methods like those described by Hanson and Wernert [20]), the camera vertical direction could be fixed at chosen points over the entire constraint manifold, and the remaining frame parameters determined by optimization within the manifold of ring constraints, possibly subject to fixing entire key-frames at selected locations or boundaries.

7 Conclusion and Future Directions

We have introduced a general framework derived from the quaternion Gauss map for studying and selecting appropriate families of coordinate frames for curves and surface patches in 3D space. Minimizing length for quaternion curve maps and area for surfaces is proposed as the appropriate generalization of parallel transport for the selection of optimal frame fields. These smooth frames can be used to generate tubular surfaces based on the space curves, thus allowing their effective display on polygon-based graphics engines. The analogous results for surface patches allow the selection of optimal local coordinate systems that may be adapted for display purposes and related applications such as oriented particle systems. Our principal new tool is the space of all possible frames, a manifold of constraints immersed in the space of quaternion frames. By defining energies and boundary conditions in this space one can produce a rich variety of application-adapted criteria for specifying optimal families of frames.

Topics for future investigation include the treatment of manifolds in higher dimensional spaces, improved interfaces for visualizing the quaternion optimization process and its results, and further analysis of the pure mathematics implied by the general framework. N -dimensional generalizations of the Frenet frame equations have been studied in the literature (see, e.g., Forsyth [10]), but the analogs of quaternions in higher dimensions are much more complex and involve Clifford algebras and the corresponding Spin groups (see, e.g., Lawson and Michelson [26]). Special simplifications do occur for the 4D case, however, allowing a treatment in terms of pairs of unit quaternions (see, e.g., [15]); this case must in fact be investigated to produce a more rigorous formulation of the 4D surface tubings proposed in [17]. Among other applications that may be approached by the quaternion formulation of coordinate frames we note the description of anisotropic surfaces (see, e.g., Kajiya [22]), the quaternion generalization of bump-mapping, and the dynamics of anisotropic particle systems. Another possible application could be the determination of optimal configurations for long-chain molecules and similar 1D and 2D structures. There are thus ample challenges for future work.

Acknowledgments

The author gratefully acknowledges the cordial hospitality of Claude Puech and the members of the iMAGIS project, a joint project of CNRS, INRIA, Institut National Polytechnique de Grenoble, and Université Joseph Fourier. We were also fortunate to have the collaboration of Hui Ma in developing the early foundations of this work. Finally, we are deeply indebted to Ken Brakke for his readiness to provide a wealth of mathematical insight as well as assistance in effectively using the Evolver system. This research was made possible in part by NSF infrastructure grant CDA 93-03189.

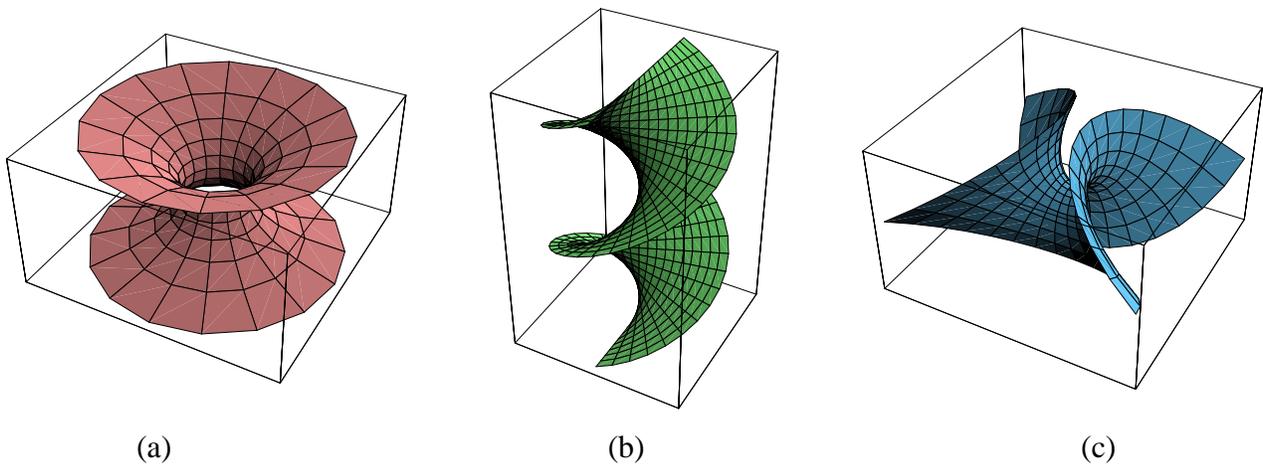


Figure 28: (a) The catenoid, a classic minimal surface in 3D space with a natural orthonormal parameterization. (b) The helicoid. (c) Enneper's surface.

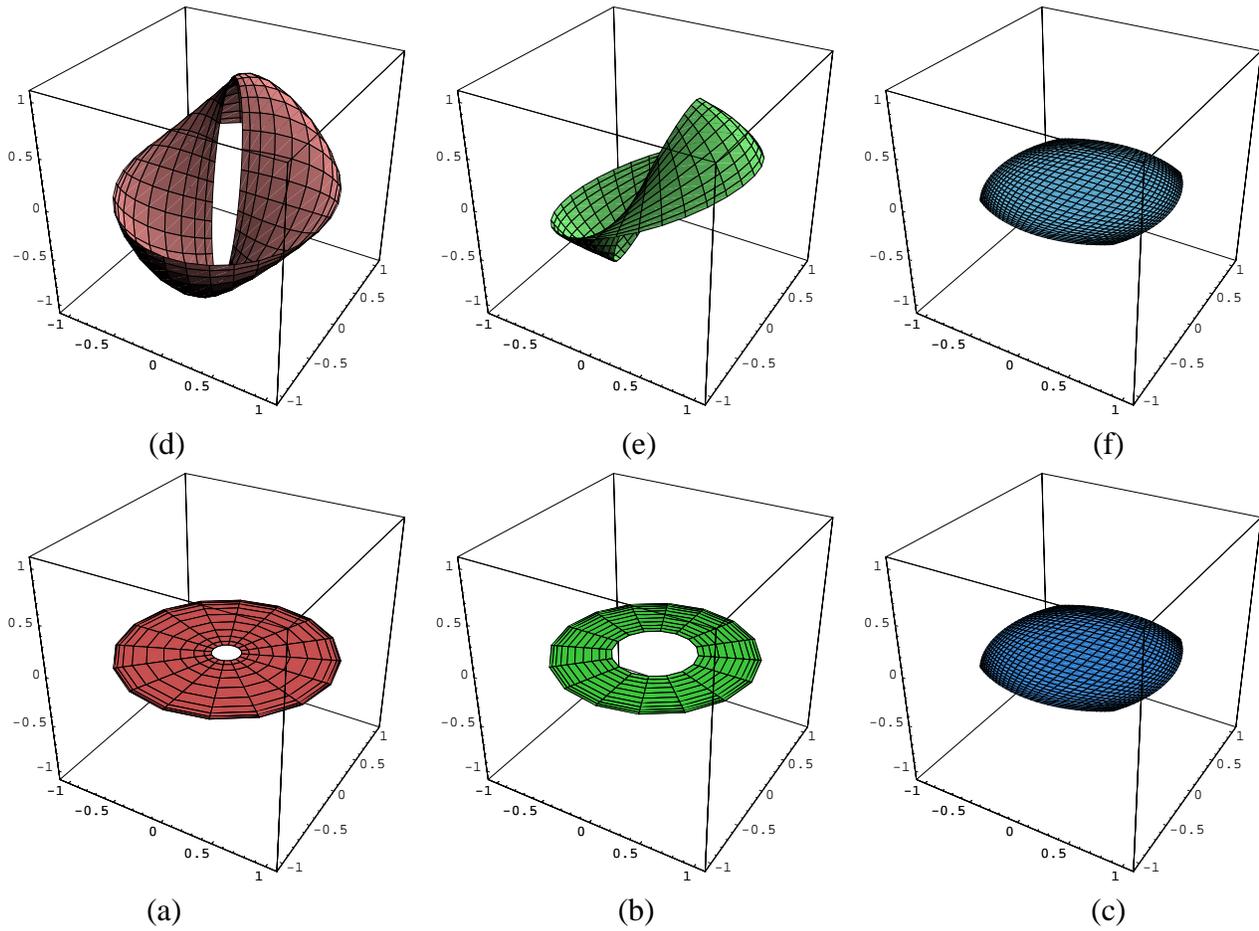


Figure 29: The Geodesic Reference quaternion frames of (a) the catenoid, (b) the helicoid, and (c) Enneper's surface. (d, e, f) The corresponding quaternion Gauss maps determined directly from the parameterization. Both the catenoid and the helicoid fail to be cyclic in quaternion space without a 4π turn around the repeating direction, so these are doubled maps. The Enneper's surface framing turns out to be identical to its Geodesic Reference frame.

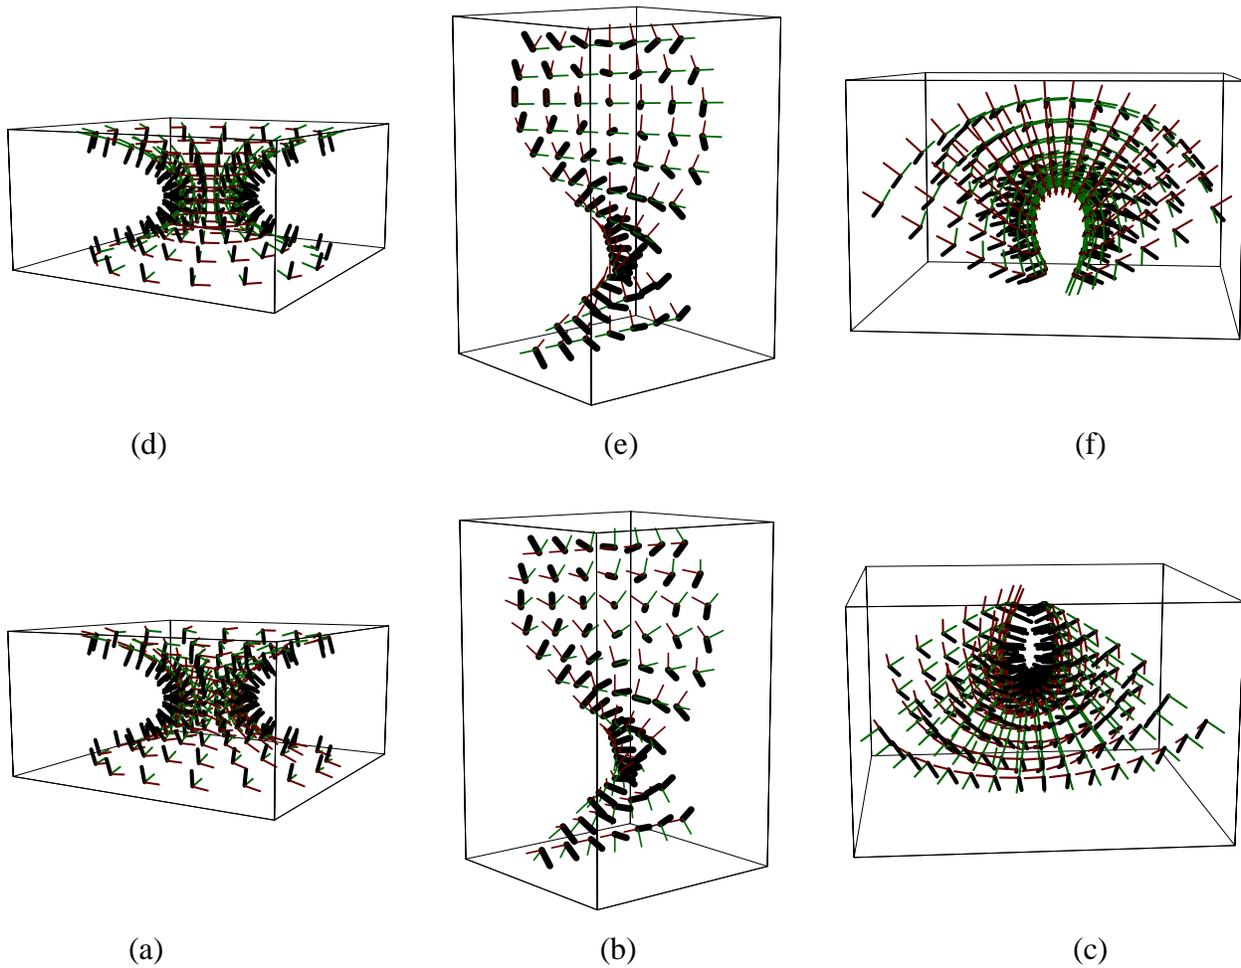


Figure 30: The 3D Geodesic Reference frames displayed directly on the surfaces of (a) the catenoid, (b) the helicoid, and (c) Enneper's surface. (d,e,f) The 3D frames computed directly from the standard parameterizations; since Enneper's surface is the same, we show in (f) a different viewpoint of the same frames.

Appendix: Quaternion Frames

2D “Quaternion” Frames

We provide below an exercise that may give some insight into the quaternion world. We show that we may express the 2D frame equations in terms of a new set of variables exactly analogous to quaternions in 3D. We begin by guessing a double-valued quadratic form for the frame:

$$\begin{bmatrix} \hat{\mathbf{N}} & \hat{\mathbf{T}} \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} = \begin{bmatrix} a^2 - b^2 & -2ab \\ 2ab & a^2 - b^2 \end{bmatrix}. \quad (50)$$

We can easily verify that if $a^2 + b^2 = 1$, this is an orthonormal parameterization of the frame, and that θ is related to (a, b) by the half-angle formulas:

$$a = \cos(\theta/2), \quad b = \sin(\theta/2).$$

If desired, the redundant parameter can be eliminated locally by using projective coordinates such as $c = b/a = \tan(\theta/2)$ to get the form (compare Eq. (30))

$$\begin{bmatrix} \hat{\mathbf{N}} & \hat{\mathbf{T}} \end{bmatrix} = \frac{1}{1+c^2} \begin{bmatrix} 1-c^2 & -2c \\ 2c & 1-c^2 \end{bmatrix}. \quad (51)$$

If we now define $W_1 = \begin{bmatrix} a & -b \\ b & a \end{bmatrix}$ and $W_2 = \begin{bmatrix} -b & -a \\ a & -b \end{bmatrix}$, then we may write

$$2W_1 \cdot \begin{bmatrix} a' \\ b' \end{bmatrix} = \hat{\mathbf{N}}', \quad 2W_2 \cdot \begin{bmatrix} a' \\ b' \end{bmatrix} = \hat{\mathbf{T}}'$$

and we may also express the right-hand side of the 2D frame equations as

$$W_1 \cdot \begin{bmatrix} 0 & -\kappa \\ +\kappa & 0 \end{bmatrix} \cdot \begin{bmatrix} a \\ b \end{bmatrix} = +v\kappa\hat{\mathbf{T}}.$$

The analogous expression for W_2 yields $\hat{\mathbf{T}}' = -v\kappa\hat{\mathbf{N}}$. Matching terms and multiplying by $W_i^T = W_i^{-1}$, we find that the equation

$$\begin{bmatrix} a' \\ b' \end{bmatrix} = \frac{1}{2}v \begin{bmatrix} 0 & -\kappa \\ +\kappa & 0 \end{bmatrix} \cdot \begin{bmatrix} a \\ b \end{bmatrix} \quad (52)$$

contains *both* the frame equations $\hat{\mathbf{T}}' = -\kappa\hat{\mathbf{N}}$ and $\hat{\mathbf{N}}' = +\kappa\hat{\mathbf{T}}$, but now in 2D “quaternion” space! If we take the angular range from $0 \rightarrow 4\pi$ instead of 2π , we have a 2 : 1 quadratic mapping from (a, b) to $(\hat{\mathbf{N}}, \hat{\mathbf{T}})$ because $(a, b) \sim (-a, -b)$ in Eq. (50).

Equation (52) is the *square root* of the frame equations (note the factor of $(1/2)$). The curvature matrix is basically $g^{-1}dg$, an element of the Lie algebra for the 2D rotation “spin group,” and takes the explicit form,

$$\begin{bmatrix} a & b \\ -b & a \end{bmatrix} \begin{bmatrix} a' & -b' \\ b' & a' \end{bmatrix} = \begin{bmatrix} aa' + bb' & -ab' + ba' \\ ab' - ba' & aa' + bb' \end{bmatrix}.$$

Here $aa' + bb' = 0$ due to the constraint $a^2 + b^2 = 1$, and

$$\begin{aligned} ab' - ba' &= \cos \frac{\theta}{2} \left[\frac{\theta'}{2} \cos \frac{\theta}{2} \right] - \sin \frac{\theta}{2} \left[-\frac{\theta'}{2} \sin \frac{\theta}{2} \right] \\ &= \frac{\theta'}{2}, \end{aligned}$$

giving the identification $v\kappa = \theta'$ when we pull out the factor of $1/2$ as in Eq. (52). The actual group properties in (a, b) space follow from the multiplication rule (easily deduced from the formulas for the trigonometric functions of sums of angles)

$$(a, b) * (\tilde{a}, \tilde{b}) = (a\tilde{a} - b\tilde{b}, a\tilde{b} + b\tilde{a}),$$

which is in turn isomorphic to complex multiplication with $(a, b) = a + ib$. This is no surprise, since $\text{SO}(2)$ and its double covering spin group are subgroups of the corresponding 3D rotation groups, and complex numbers are a subset of the quaternions.

3D Quaternion Frames

We next outline the basic features of quaternion frames; see, e.g., [1] for a nice textbook treatment of quaternions and their properties.

A quaternion frame is a four-vector $q = (q_0, q_1, q_2, q_3) = (q_0, \vec{q})$ with the following features:

- **Unit Norm.** If we define the inner product of two quaternions as

$$q \cdot p = q_0 p_0 + q_1 p_1 + q_2 p_2 + q_3 p_3, \quad (53)$$

then the components of a quaternion frame obey the constraint

$$q \cdot q = (q_0)^2 + (q_1)^2 + (q_2)^2 + (q_3)^2 = 1, \quad (54)$$

and therefore lie on S^3 , the three-sphere embedded in four-dimensional Euclidean space \mathbb{R}^4 .

- **Multiplication rule.** The quaternion product of two quaternions p and q is defined to give a positive cross-product in the vector part, and may be written as

$$p * q = (p_0 q_0 - \mathbf{p} \cdot \mathbf{q}, p_0 \mathbf{q} + q_0 \mathbf{p} + \mathbf{p} \times \mathbf{q}),$$

or more explicitly in component form as

$$p * q = \begin{bmatrix} [p * q]_0 \\ [p * q]_1 \\ [p * q]_2 \\ [p * q]_3 \end{bmatrix} = \begin{bmatrix} p_0 q_0 - p_1 q_1 - p_2 q_2 - p_3 q_3 \\ p_1 q_0 + p_0 q_1 + p_2 q_3 - p_3 q_2 \\ p_2 q_0 + p_0 q_2 + p_3 q_1 - p_1 q_3 \\ p_3 q_0 + p_0 q_3 + p_1 q_2 - p_2 q_1 \end{bmatrix}. \quad (55)$$

This rule is isomorphic to left multiplication in the group $\text{SU}(2)$, the double covering of the ordinary 3D rotation group $\text{SO}(3)$. What is more useful for our purposes is the fact that it

is also isomorphic to multiplication by a member of the group of (possibly sign-reversing) orthogonal transformations in \mathbb{R}^4 :

$$p * q = [P]q = \begin{bmatrix} p_0 & -p_1 & -p_2 & -p_3 \\ p_1 & p_0 & -p_3 & p_2 \\ p_2 & p_3 & p_0 & -p_1 \\ p_3 & -p_2 & p_1 & p_0 \end{bmatrix} \begin{bmatrix} q_0 \\ q_1 \\ q_2 \\ q_3 \end{bmatrix}, \quad (56)$$

where $[P]$ is an orthogonal matrix, $[P]^t \cdot [P] = I_4$; since $[P]$ has only 3 free parameters, it does not itself include all 4D rotations. However, we may recover the remaining 3 parameters by considering transformation by right multiplication to be an independent operation, resulting in a similar matrix but with the signs in the lower right-hand off-diagonal 3×3 section reversed. (This corresponds to the well-known decomposition of the 4D rotation group into two 3D rotations; see, e.g., [14].)

If two quaternions a and b are transformed by multiplying them by the same quaternion p , their inner product $a \cdot b$ transforms as

$$(p * a) \cdot (p * b) = (a \cdot b)(p \cdot p) \quad (57)$$

and so is invariant if p is a unit quaternion frame representing a rotation. This also follows trivially from the fact that $[P]$ is orthogonal.

The *inverse* of a unit quaternion satisfies $q * q^{-1} = (1, \mathbf{0})$ and is easily shown to take the form $q^{-1} = (q_0, -\mathbf{q})$. The relative quaternion rotation t transforming between two quaternions may be represented using the product

$$t = p * q^{-1} = (p_0 q_0 + \mathbf{p} \cdot \mathbf{q}, q_0 \mathbf{p} - p_0 \mathbf{q} - \mathbf{p} \times \mathbf{q}).$$

This has the convenient property that the zeroth component is the invariant 4D inner product $p \cdot q = \cos(\theta/2)$, where θ is the angle of the rotation in 3D space needed to rotate along a geodesic from the frame denoted by q to that given by p . In fact, the 4D inner product reduces to

$$p * q^{-1} + q * p^{-1} = (2q \cdot p, \mathbf{0}),$$

while the 3D dot product and cross product arise from the symmetric and antisymmetric sums of quaternions containing only a 3-vector part:

$$\begin{aligned} \mathbf{p} * \mathbf{q} &\equiv (0, \mathbf{p}) * (0, \mathbf{q}) = (-\mathbf{p} \cdot \mathbf{q}, \mathbf{p} \times \mathbf{q}) \\ \mathbf{p} * \mathbf{q} + \mathbf{q} * \mathbf{p} &= -2 \mathbf{p} \cdot \mathbf{q} \\ \mathbf{p} * \mathbf{q} - \mathbf{q} * \mathbf{p} &= 2 \mathbf{p} \times \mathbf{q} \end{aligned}$$

- **Mapping to 3D rotations.** Every possible 3D rotation R (a 3×3 orthogonal matrix) can be constructed from either of two related quaternions, $q = (q_0, q_1, q_2, q_3)$ or $-q =$

$(-q_0, -q_1, -q_2, -q_3)$, using the quadratic relationship $R_q(\mathbf{V}) = q * (0, \mathbf{V}) * q^{-1}$, written explicitly as

$$R = \begin{bmatrix} q_0^2 + q_1^2 - q_2^2 - q_3^2 & 2q_1q_2 - 2q_0q_3 & 2q_1q_3 + 2q_0q_2 \\ 2q_1q_2 + 2q_0q_3 & q_0^2 - q_1^2 + q_2^2 - q_3^2 & 2q_2q_3 - 2q_0q_1 \\ 2q_1q_3 - 2q_0q_2 & 2q_2q_3 + 2q_0q_1 & q_0^2 - q_1^2 - q_2^2 + q_3^2 \end{bmatrix}. \quad (58)$$

The signs here result from choosing the left multiplication convention $R_p R_q(\mathbf{V}) = R_{pq}(\mathbf{V}) = (p * q) * (0, \mathbf{V}) * (p * q)^{-1}$. Algorithms for the inverse mapping from R to q require careful singularity checking, and are detailed, e.g., in [31, 38].

The analog for Eq. (58) of the projective coordinates for 2D rotations noted in Eq. (51) is obtained by converting to the projective variable $\mathbf{c} = \mathbf{q}/q_0 = \tan(\theta/2) \hat{\mathbf{n}}$ and factoring out

$$(q_0)^2 = \frac{(q_0)^2}{(q_0)^2 + \mathbf{q} \cdot \mathbf{q}} = \frac{1}{1 + \mathbf{q} \cdot \mathbf{q}/(q_0)^2} = \frac{1}{1 + \|\mathbf{c}\|^2}.$$

We then find

$$R = \frac{1}{1 + \|\mathbf{c}\|^2} \begin{bmatrix} 1 + c_1^2 - c_2^2 - c_3^2 & 2c_1c_2 - 2c_3 & 2c_1c_3 + 2c_2 \\ 2c_1c_2 + 2c_3 & 1 - c_1^2 + c_2^2 - c_3^2 & 2c_2c_3 - 2c_1 \\ 2c_1c_3 - 2c_2 & 2c_2c_3 + 2c_1 & 1 - c_1^2 - c_2^2 + c_3^2 \end{bmatrix}. \quad (59)$$

- **Rotation Correspondence.** When we substitute $q(\theta, \hat{\mathbf{n}}) = (\cos \frac{\theta}{2}, \hat{\mathbf{n}} \sin \frac{\theta}{2})$ into Eq. (58), where $\hat{\mathbf{n}} \cdot \hat{\mathbf{n}} = 1$ is a unit three-vector lying on the two-sphere \mathbb{S}^2 , $R(\theta, \hat{\mathbf{n}})$ becomes the standard matrix for a rotation by θ in the plane perpendicular to $\hat{\mathbf{n}}$. The quadratic form ensures that the two distinct unit quaternions q and $-q$ in \mathbb{S}^3 correspond to the *same* $\text{SO}(3)$ rotation. For reference, the explicit form of $R(\theta, \hat{\mathbf{n}})$ is [9]

$$\mathbf{R}(\theta, \hat{\mathbf{n}}) = \begin{bmatrix} c + (n_1)^2(1 - c) & n_1n_2(1 - c) - sn_3 & n_3n_1(1 - c) + sn_2 \\ n_1n_2(1 - c) + sn_3 & c + (n_2)^2(1 - c) & n_3n_2(1 - c) - sn_1 \\ n_1n_3(1 - c) - sn_2 & n_2n_3(1 - c) + sn_1 & c + (n_3)^2(1 - c) \end{bmatrix}, \quad (60)$$

where $c = \cos \theta$, $s = \sin \theta$, and $\hat{\mathbf{n}} \cdot \hat{\mathbf{n}} = 1$. For example, choosing the quaternion $q = (\cos \frac{\theta}{2}, 0, 0, \sin \frac{\theta}{2})$ yields the rotation matrix

$$\mathbf{R} = \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

producing a right-handed rotation of the basis vectors $\hat{\mathbf{x}} = (1, 0, 0)$ and $\hat{\mathbf{y}} = (0, 1, 0)$ around the $\hat{\mathbf{z}}$ axis.

- **Quaternion Frame Evolution.** All 3D coordinate frames can be expressed in the form of quaternions using Eq. (58). If we assume the columns of Eq. (58) are the vectors $(\hat{\mathbf{N}}_1, \hat{\mathbf{N}}_2, \hat{\mathbf{T}})$,

respectively, one can explicitly express each vector in terms of the following matrices

$$[W_1] = \begin{bmatrix} q_0 & q_1 & -q_2 & -q_3 \\ q_3 & q_2 & q_1 & q_0 \\ -q_2 & q_3 & -q_0 & q_1 \end{bmatrix} \quad (61)$$

$$[W_2] = \begin{bmatrix} -q_3 & q_2 & q_1 & -q_0 \\ q_0 & -q_1 & q_2 & -q_3 \\ q_1 & q_0 & q_3 & q_2 \end{bmatrix} \quad (62)$$

$$[W_3] = \begin{bmatrix} q_2 & q_3 & q_0 & q_1 \\ -q_1 & -q_0 & q_3 & q_2 \\ q_0 & -q_1 & -q_2 & q_3 \end{bmatrix}, \quad (63)$$

with the result that $\hat{\mathbf{N}}_1 = [W_1] \cdot [q]$, $\hat{\mathbf{N}}_2 = [W_2] \cdot [q]$, and $\hat{\mathbf{T}} = [W_3] \cdot [q]$, where $[q]$ is the column vector with components (q_0, q_1, q_2, q_3) . Differentiating each of these expressions and substituting Eq. (9), one finds that factors of the matrices $[W_i]$ can be pulled out and a single universal equation linear in the quaternions remains:

$$\begin{bmatrix} q'_0 \\ q'_1 \\ q'_2 \\ q'_3 \end{bmatrix} = \frac{v}{2} \begin{bmatrix} 0 & -k_x & -k_y & -k_z \\ +k_x & 0 & -k_z & +k_y \\ +k_y & +k_z & 0 & -k_x \\ +k_z & -k_y & +k_x & 0 \end{bmatrix} \cdot \begin{bmatrix} q_0 \\ q_1 \\ q_2 \\ q_3 \end{bmatrix}. \quad (64)$$

The first occurrence of this equation that we are aware of is in the works of Tait [41]. Here $v(t) = \|\mathbf{x}'(t)\|$ is the scalar magnitude of the curve derivative if a unit-speed parameterization is not being used for the curve. One may consider Eq. (64) to be in some sense the *square root* of the 3D frame equations. Alternatively, we can deduce directly from $R_q(\mathbf{V}) = q * (0, \mathbf{V}) * q^{-1}$, $dq = q * (q^{-1} * dq)$, and $(q^{-1} * dq) = -(dq^{-1} * q)$, that the 3D vector equations are equivalent to the quaternion form

$$q' = \frac{1}{2}v q * (0, k_x, k_y, k_z) = \frac{1}{2}v q * (0, \mathbf{k}) \quad (65)$$

$$(q^{-1})' = -\frac{1}{2}v (0, \mathbf{k}) * q^{-1}, \quad (66)$$

where $\mathbf{k} = 2(q_0 d\mathbf{q} - \mathbf{q} dq_0 - \mathbf{q} \times d\mathbf{q})$, or, explicitly,

$$\begin{aligned} k_0 &= 2(dq_x q_x + dq_y q_y + dq_z q_z + dq_0 q_0) = 0 \\ k_x &= 2(q_0 dq_x - q_x dq_0 - q_y dq_z + q_z dq_y) \\ k_y &= 2(q_0 dq_y - q_y dq_0 - q_z dq_x + q_x dq_z) \\ k_z &= 2(q_0 dq_z - q_z dq_0 - q_x dq_y + q_y dq_x). \end{aligned}$$

Here $k_0 = 0$ is the diagonal value in Eq. (64).

The quaternion approach to the frame equations exemplified by Eq. (64) (or Eq. (65)) has the following key properties:

- $q(t) \cdot q'(t) = 0$ by construction. Thus all unit quaternions remain unit quaternions as they evolve by this equation.
- The number of equations has been reduced from nine coupled equations with six orthonormality constraints to four coupled equations incorporating a single constraint that keeps the solution vector confined to the three-sphere.

- **Quaternion Surface Evolution.** The same set of equations can be considered to work on curves that are paths in a surface, thus permitting a quaternion equivalent to the Weingarten equations for the classical differential geometry of surfaces as well. Explicit forms permitting the recovery of the classical equations follow from re-expressing Eqs. (22) and (23) in quaternion form. The curvature equation is essentially the cross-product of two derivatives of the form of Eqs. (65,66), and thus obtainable by a quaternion multiplication:

$$\begin{aligned}
q_u * q_v^{-1} &= -\frac{1}{4} q * (0, \mathbf{a}) * (0, \mathbf{b}) * q^{-1} \\
&= -\frac{1}{4} q * (-\mathbf{a} \cdot \mathbf{b}, \mathbf{a} \times \mathbf{b}) * q^{-1} \\
&= -\frac{1}{4} \left[-\mathbf{a} \cdot \mathbf{b} \hat{\mathbf{I}} + (\mathbf{a} \times \mathbf{b})_x \hat{\mathbf{T}}_1 + (\mathbf{a} \times \mathbf{b})_y \hat{\mathbf{T}}_2 + (\mathbf{a} \times \mathbf{b})_z \hat{\mathbf{N}} \right]. \quad (67)
\end{aligned}$$

Here Eq. (31) defines the quaternion frame vectors, $\hat{\mathbf{N}} \equiv (0, \hat{\mathbf{N}}) = q * (0, \hat{\mathbf{z}}) * q^{-1}$, etc., and we have introduced the identity element $\hat{\mathbf{I}} = (1, \mathbf{0}) = q * (1, \mathbf{0}) * q^{-1}$ as a fourth quaternion basis vector. The mean curvature equation has only one derivative and a free vector field; an expression producing the right combination of terms is

$$\begin{aligned}
\hat{\mathbf{T}}_1 * q * q_u^{-1} + \hat{\mathbf{T}}_2 * q * q_v^{-1} &= -\frac{1}{2} q * (-\hat{\mathbf{x}} \cdot \mathbf{a} - \hat{\mathbf{y}} \cdot \mathbf{b}, \hat{\mathbf{x}} \times \mathbf{a} + \hat{\mathbf{y}} \times \mathbf{b}) * q^{-1} \\
&= -\frac{1}{2} \left[-(a_x + b_y) \hat{\mathbf{I}} + b_z \hat{\mathbf{T}}_1 - a_z \hat{\mathbf{T}}_2 + (a_y - b_x) \hat{\mathbf{N}} \right] \quad (68)
\end{aligned}$$

Projecting out the $\hat{\mathbf{N}}$ component of these equations recovers the scalar and mean curvatures:

$$\begin{aligned}
K &= \det \begin{bmatrix} +a_y(u, v) & -a_x(u, v) \\ +b_y(u, v) & -b_x(u, v) \end{bmatrix} = a_x b_y - a_y b_x \\
H &= \frac{1}{2} \text{tr} \begin{bmatrix} +a_y(u, v) & -a_x(u, v) \\ +b_y(u, v) & -b_x(u, v) \end{bmatrix} = \frac{1}{2} (a_y - b_x).
\end{aligned}$$

References

- [1] ALTMANN, S. L. *Rotations, Quaternions, and Double Groups*. Oxford University Press, 1986.
- [2] BARR, A. H., CURRIN, B., GABRIEL, S., AND HUGHES, J. F. Smooth interpolation of orientations with angular velocity constraints using quaternions. In *Computer Graphics (SIG-GRAPH '92 Proceedings)* (July 1992), E. E. Catmull, Ed., vol. 26, pp. 313–320.

- [3] BERGER, M. *Geometry I, II*. Springer Verlag, Berlin, 1987.
- [4] BISHOP, R. L. There is more than one way to frame a curve. *Amer. Math. Monthly* 82, 3 (March 1975), 246–251.
- [5] BLOOMENTHAL, J. Calculation of reference frames along a space curve. In *Graphics Gems*, A. Glassner, Ed. Academic Press, Cambridge, MA, 1990, pp. 567–571.
- [6] BRAKKE, K. A. The surface evolver. *Experimental Mathematics* 1, 2 (1992), 141–165. The “Evolver” system, manual, and sample data files are available by anonymous ftp from geom.umn.edu, The Geometry Center, Minneapolis MN.
- [7] EGUCHI, T., GILKEY, P., AND HANSON, A. Gravitation, gauge theories and differential geometry. *Physics Reports* 66, 6 (December 1980), 213–393.
- [8] EISENHART, L. P. *A Treatise on the Differential Geometry of Curves and Surfaces*. Dover, New York, 1909 (1960).
- [9] FOLEY, J., VAN DAM, A., FEINER, S., AND HUGHES, J. *Computer Graphics, Principles and Practice*, second ed. Addison-Wesley, 1990. page 227.
- [10] FORSYTH, A. R. *Geometry of Four Dimensions*. Cambridge University Press, 1930.
- [11] GABRIEL, S., AND KAJIYA, J. T. Spline interpolation in curved space. In *State of the Art Image Synthesis* (1985). Siggraph ’85 Course notes.
- [12] GRAY, A. *Modern Differential Geometry of Curves and Surfaces with Mathematica*, second ed. CRC Press, Inc., Boca Raton, FL, 1998.
- [13] GRIMM, C. M., AND HUGHES, J. F. Modeling surfaces with arbitrary topology using manifolds. In *Computer Graphics Proceedings, Annual Conference Series* (1995), pp. 359–368. Proceedings of SIGGRAPH ’95.
- [14] HANSON, A. J. The rolling ball. In *Graphics Gems III*, D. Kirk, Ed. Academic Press, Cambridge, MA, 1992, pp. 51–60.
- [15] HANSON, A. J. Rotations for n-dimensional graphics. In *Graphics Gems V*, A. Paeth, Ed. Academic Press, Cambridge, MA, 1995, pp. 55–64.
- [16] HANSON, A. J. Constrained optimal framings of curves and surfaces using quaternion gauss maps. In *Proceedings of Visualization ’98* (1998), IEEE Computer Society Press, pp. 375–382. Missing pages 2 and 4 in original Proceedings; correct version in second printing of Proceedings, on conference CDROM, or <ftp://ftp.cs.indiana.edu/pub/hanson/Vis98/vis98.quat.pdf>.
- [17] HANSON, A. J., AND HENG, P. A. Illuminating the fourth dimension. *Computer Graphics and Applications* 12, 4 (July 1992), 54–62.

- [18] HANSON, A. J., AND MA, H. Visualizing flow with quaternion frames. In *Proceedings of Visualization '94* (1994), IEEE Computer Society Press, pp. 108–115.
- [19] HANSON, A. J., AND MA, H. Quaternion frame approach to streamline visualization. *IEEE Trans. on Visualiz. and Comp. Graphics 1*, 2 (June 1995), 164–174.
- [20] HANSON, A. J., AND WERNERT, E. A. Constrained 3d navigation with 2d controllers. In *Proceedings of Visualization '97* (1997), IEEE Computer Society Press, pp. 175–182.
- [21] HART, J. C., FRANCIS, G. K., AND KAUFFMAN, L. H. Visualizing quaternion rotation. *ACM Trans. on Graphics 13*, 3 (1994), 256–276.
- [22] KAJIYA, J. T. Anisotropic reflection models. In *Computer Graphics* (1985), vol. 19, pp. 15–21. Proceedings of SIGGRAPH '85.
- [23] KIM, M.-J., KIM, M.-S., AND SHIN, S. Y. A general construction scheme for unit quaternion curves with simple high order derivatives. In *Computer Graphics Proceedings, Annual Conference Series* (1995), pp. 369–376. Proceedings of SIGGRAPH '95.
- [24] KLOCK, F. Two moving coordinate frames for sweeping along a 3d trajectory. *Computer Aided Geometric Design 3* (1986).
- [25] KUIPERS, J. *Quaternions and Rotation Sequences*. Princeton University Press, 1999.
- [26] LAWSON, H. B., AND MICHELSON, M.-L. *Spin Geometry*. Princeton University Press, 1989.
- [27] MA, H. *Curve and Surface Framing for Scientific Visualization and Domain Dependent Navigation*. PhD thesis, Indiana University, February 1996.
- [28] MAX, N. L. DNA animation, from atom to chromosome. *Journal of Molecular Graphics 3*, 2 (1985), 69–71.
- [29] MILNOR, J. *Topology from the Differentiable Viewpoint*. The University Press of Virginia, Charlottesville, 1965.
- [30] MURTAGH, B. A., AND SAUNDER, M. A. MINOS 5.0 user's guide. Technical Report SOL 83-20, Dept. of Operations Research, Stanford University, 1983.
- [31] NIELSON, G. M. Smooth interpolation of orientations. In *Computer Animation '93* (Tokyo, June 1993), N. Thalmann and D. Thalmann, Eds., Springer-Verlag, pp. 75–93.
- [32] PLATT, J. C., AND BARR, A. H. Constraint methods for flexible models. In *Computer Graphics (SIGGRAPH '88 Proceedings)* (Aug. 1988), J. Dill, Ed., vol. 22, pp. 279–288.
- [33] PLETINCKS, D. Quaternion calculus as a basic tool in computer graphics. *The Visual Computer 5*, 1 (1989), 2–13.

- [34] RAMAMOORTHY, R., AND BARR, A. H. Fast construction of accurate quaternion splines. In *SIGGRAPH 97 Conference Proceedings* (Aug. 1997), T. Whitted, Ed., Annual Conference Series, ACM SIGGRAPH, Addison Wesley, pp. 287–292. ISBN 0-89791-896-7.
- [35] SCHLAG, J. Using geometric constructions to interpolate orientation with quaternions. In *Graphics Gems II*, J. Arvo, Ed. Academic Press, 1991, pp. 377–380.
- [36] SHANI, U., AND BALLARD, D. H. Splines as embeddings for generalized cylinders. *Computer Vision, Graphics, and Image Processing* 27 (1984), 129–156.
- [37] SHOEMAKE, K. Animating rotation with quaternion curves. In *Computer Graphics* (1985), vol. 19, pp. 245–254. Proceedings of SIGGRAPH 1985.
- [38] SHOEMAKE, K. Animation with quaternions. Siggraph Course Lecture Notes, 1987.
- [39] SHOEMAKE, K. Fiber bundle twist reduction. In *Graphics Gems IV*, P. Heckbert, Ed. Academic Press, 1994, pp. 230–236.
- [40] STEENROD, N. *The Topology of Fibre Bundles*. Princeton University Press, 1951. Princeton Mathematical Series 14.
- [41] TAIT, P. *An Elementary Treatise on Quaternions*. Cambridge University Press, 1890.