

Constrained Optimal Framings of Curves and Surfaces using Quaternion Gauss Maps

Andrew J. Hanson *

Computer Science Department
Indiana University
Bloomington, IN 47405 USA

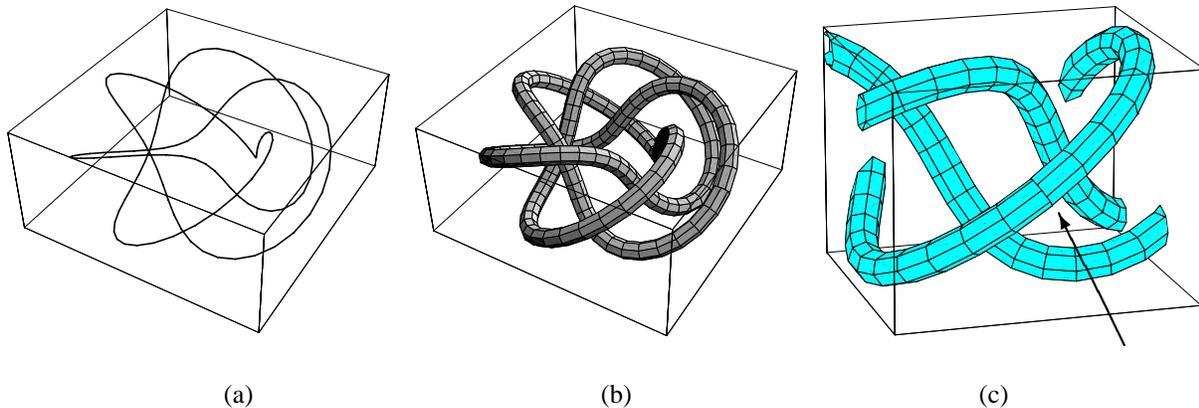


Figure 1: The (3,5) torus knot, a complex periodic 3D curve. (a) The line drawing is nearly useless as a 3D representation. (b) A tubing based on parallel transporting an initial reference frame produces an informative visualization, but is not periodic. (c) The arrow in this closeup exposes the subtle 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.

Abstract

We propose a general paradigm for computing optimal coordinate frame fields that may be exploited to visualize 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 heuristic generalization of parallel transport. Our approach differs from minimal-tangential-acceleration approaches due to the addition of “sliding ring” constraints that fix one frame axis, but allow an axial rotational freedom whose value is varied in the optimization process. 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.

CR Categories: I.3.6 [Computer Graphics]: Methodology and Techniques. I.3.8 [Computer Graphics]: Applications.

Keywords: Quaternions; Frames; Tubing; Curves; Surfaces

1 Introduction

We propose a general framework for selecting optimal systems of coordinate frames that can be applied to visualizing 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, are independent of parameterization, and extend naturally to situations where parallel transport is not applicable.

Motivation. Many visualization 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 to tubing involve heuristics unrelated to parameterization-invariant optimization measures in order to achieve such properties as periodicity. Similar problems occur in the construction of suitable visualizations of complex surfaces and oriented particle systems on surfaces, since the intrinsic orientation properties may be poorly exposed by the original representation. If a surface patch is represented by a rectangular but nonorthogonal mesh, for example, there is no obvious way to choose among alternative local orthonormal frame assignments; if the surface has regions of vanishing curvature, methods based on directions of principal curvatures break down as well.

*Email: hanson@cs.indiana.edu

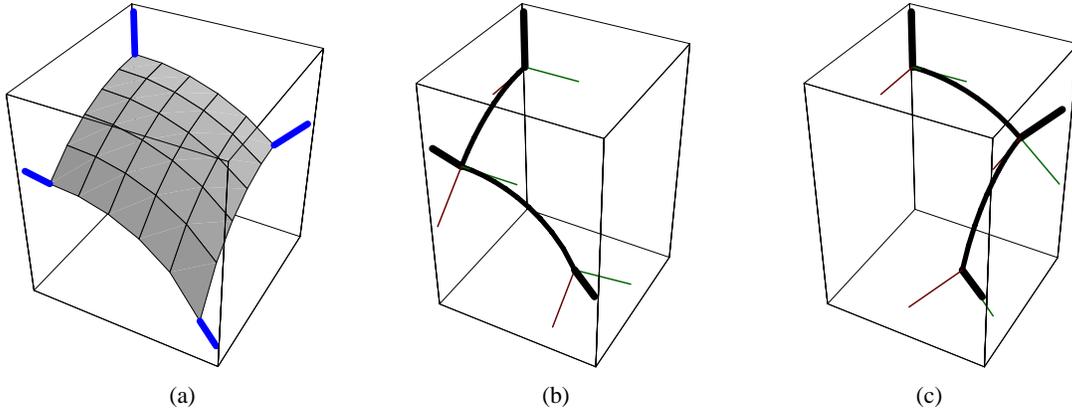


Figure 2: (a) A smooth 3D surface patch having a non-orthogonal parameterization, along with its geometrically-fixed normals at the four corners. No unique orthonormal frame is derivable from the parameterization. If we imitate parallel transport for curves to evolve the initial frame at the top corner 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.

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 [13]).

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. 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 nonperiodicity. 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., [16, 24, 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., [25, 23, 19, 15]). We do not address the related question of optimal freely moving frames treated by the minimal-tangential-acceleration methods (see, e.g., [2, 22, 8]); 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 [14, 20]. 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. Paths in this *space 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., [26, 3]).

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

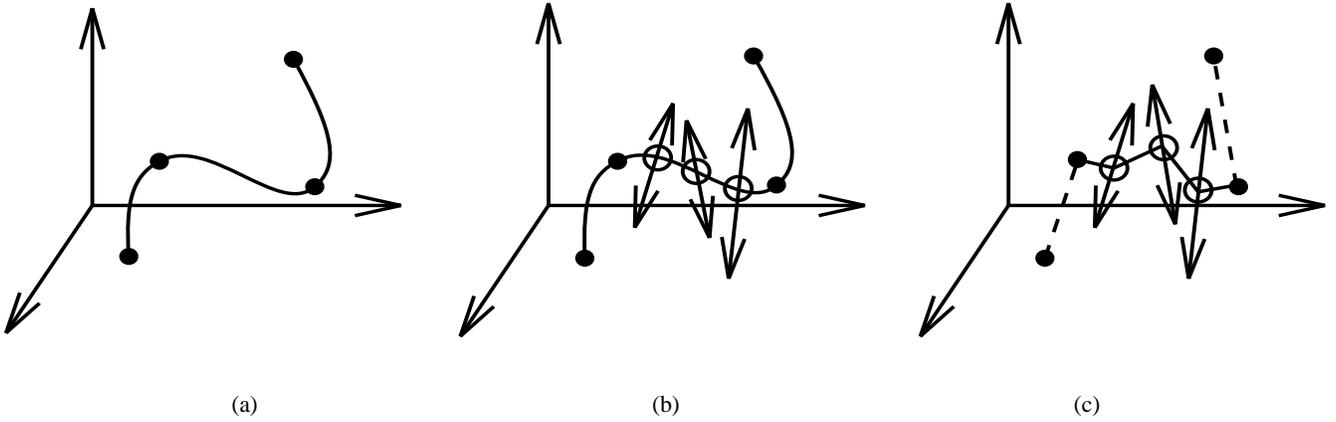


Figure 3: (a) The camera frame interpolation problem is analogous to the problem of finding a minimal-bending spline curve through a series of fixed key points. (b) The optimal curve frame assignment problem is analogous to fixing the end points of a curve segment and choosing *in addition* a family 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.

minimal bending energy (see Barr et al. [2, 22]), 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., [16, 24, 5]. The quaternion Gauss map is a logical extension of the quaternion frame approach to visualizing space curves introduced by Hanson and Ma [11, 12]. For basic information on orientation spaces and their relationship to quaternions, see, e.g., [1, 21, 19].

Background on the differential geometry of curves and surfaces may be found in sources such as the classical treatise of Eisenhart [7] and in Gray’s *MATHEMATICA*-based text [9], which inspired a number of the illustrations in this paper. The classical Frenet frame is defined and studied in these texts. The frame we refer to as the parallel-transport frame was first described carefully by Bishop [4], and has been commonly used in graphics (see, e.g., [5, 24, 17]). A significant difference between these two methods is that the Frenet frame is locally defined but possibly discontinuous, whereas the parallel-transport frame is continuous but non-local, corresponding to the solution of a differential equation.

2 The Space of Frames

We begin by introducing 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 $\mathbf{x}(u, v)$.

We now 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

$$F(\theta, \hat{\mathbf{v}}) = \begin{bmatrix} f_0^2 + f_1^2 - f_2^2 - f_3^2 & 2f_1f_2 - 2f_0f_3 & 2f_1f_3 + 2f_0f_2 \\ 2f_1f_2 + 2f_0f_3 & f_0^2 - f_1^2 + f_2^2 - f_3^2 & 2f_2f_3 - 2f_0f_1 \\ 2f_1f_3 - 2f_0f_2 & 2f_2f_3 + 2f_0f_1 & f_0^2 - f_1^2 - f_2^2 + f_3^2 \end{bmatrix}, \quad (1)$$

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

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

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

$$b(\hat{\mathbf{v}}) = q(\arccos(\hat{\mathbf{z}} \cdot \hat{\mathbf{v}}), (\hat{\mathbf{z}} \times \hat{\mathbf{v}}) / \|\hat{\mathbf{z}} \times \hat{\mathbf{v}}\|). \quad (3)$$

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 4. If $\hat{\mathbf{v}} = \hat{\mathbf{z}}$, there is no problem, since we just take $b(\hat{\mathbf{v}})$ to be the quaternion $(1, 0)$; if $\hat{\mathbf{v}} = -\hat{\mathbf{z}}$,

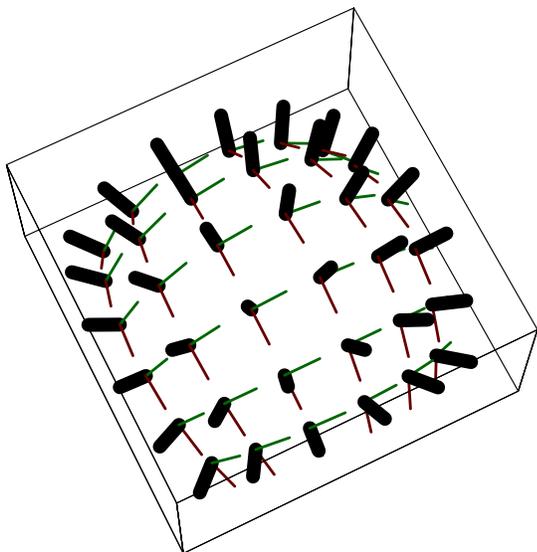


Figure 4: Example of the Geodesic Reference Frame: on the northern hemisphere of a 2-sphere, the Geodesic Reference Frame tilts the \hat{z} axis of the north pole's identity frame along the shortest arc to align with a specified reference direction.

we may choose any 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

$$f(\theta, \hat{v}) = q(\theta, \hat{v}) * b(\hat{v}), \quad (4)$$

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.

Surface Patch Example. Figure 5 shows how the frame choice problem of Figure 2 can be visualized in the quaternion space of frames. We choose a quaternion projection that shows only the 3-vector part of the quaternion, dropping q_0 . A frame choice is achieved by moving a point around the *sliding ring constraint* defined by Eq. (4) to the desired position. The constraint rings in Figure 5 are the generalizations to quaternion space of the constraint lines symbolized in Figure 3(b). The vertex A admits a family of frames $f(\theta, \hat{z})$ that is a circle in quaternion space, but projects “edge-on” to a vertical bar in our default projection. The spaces of frames at the other vertices project as ellipses. The outer ring in Figure 5(b) is touched by two paths, corresponding to the clockwise and counterclockwise parallel transport routes in Figure 5(a); the gap between the intercepts in the outer ring corresponds to the inequivalence of the two frames at the bottom vertex of Figure 5(a).

Closed Curve Example. In Figure 6, we show a simple closed curve, the trefoil knot, the quaternion plot of its periodic Frenet frame, and, just to show we can do it, the entire constraint surface in which the Frenet frame and all other possible quaternion framings of the trefoil must lie. In the next section, we show the results of optimizing a continuous family of frames lying within this remarkable surface.

3 Minimal Frames

We have computed a wide selection of examples using the Evolver of K. Brakke [6] as our optimization tool. The Evolver is a public-domain, extensively documented system with a huge range of constraint-solving capabilities, widely used in mathematics and certain engineering problems. It has a very simple interface for handling parametric constraints like our sliding ring constraints, and can also handle a wide variety of energy functionals and boundary specifications. Most of the examples shown here take only a few seconds to stabilize in the Evolver; more complex geometries will of course take longer.

Two enhancements to the Evolver handle the specific issues related to quaternion optimization; the symmetry specification `symmetry_group "central_symmetry"` identifies the quaternion q with $-q$ if desired during the variation to prevent reflected double traversals from varying independently, and the system is able to use the pullback 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 quaternion three-sphere. Computation using this metric, however, is very slow, and so in practice we have used the Euclidean R^4 chord approximation, which works quite well for closely spaced samples and is much faster. (There are other choices of three-sphere parameterizations and quaternion distance measures that we have not yet attempted that could be more efficient still.) The energy functional that we chose to specify for the Evolver (or that would be implemented in a dedicated system) is thus simply the sum of the Euclidean lengths of each line segment in R^4 :

$$d = \sum_{i,j} |x_i - x_j|$$

where $|q| = \sqrt{q \cdot q} = \sqrt{q_0 q_0 + q_1 q_1 + q_2 q_2 + q_3 q_3}$. For surface areas, the Evolver breaks polygons into triangles, computes their areas, and minimizes the total sum as the vertex positions vary.

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

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

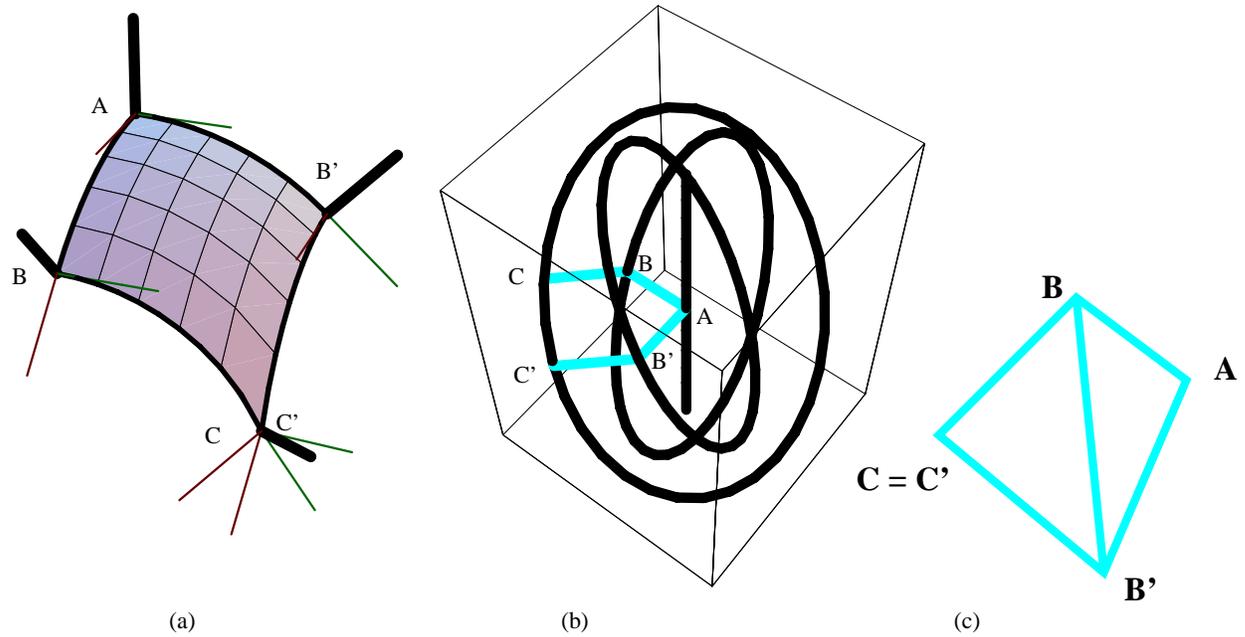


Figure 5: A different viewpoint of the mismatch problem of Figure 2. (a) Choosing different routes to determine the frame at the bottom point results in the incompatible frames shown here in 3D space. (b) The same information is presented here in the quaternion space-of-frames picture. We use throughout 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.

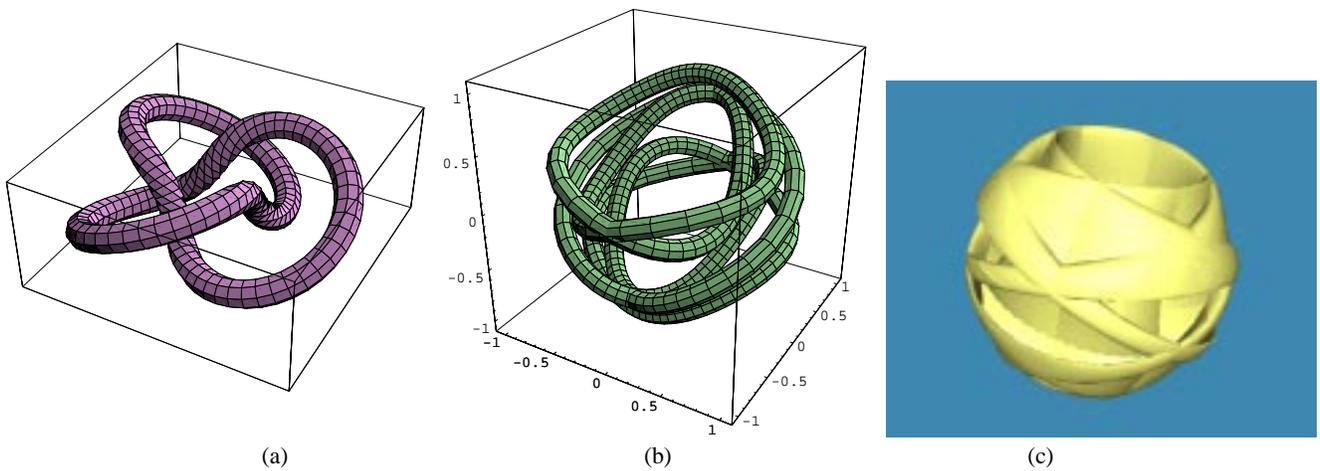


Figure 6: (a) A trefoil torus knot. (b) Its quaternion Frenet frame projected to 3D. For this trefoil knot, the frame does not close on itself in quaternion space unless the curve is traversed twice, corresponding to the double-valued “mirror” image of the rotation space that can occur in the quaternion representation. The longer segments in (b) correspond to the three high-torsion segments observable in (a). (c) The full constraint space for the trefoil knot is a very complex surface swept out by the constraint rings. All quaternions are projected to 3D using only the vector part.

Here `Qprod` and `makeQfromVec` perform the quaternion product and create the quaternion corresponding to Eq. (4) with `P1` replacing θ . Note that, since the 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.

With these preliminaries, the Evolver can easily be used to minimize the length of the total piecewise linear path among sliding ring constraints for selected curves, and the total area spanned by analogous sliding rings for surfaces. We made no attempt to go beyond piecewise linear curves. One interesting result is that there appear to be families of topologically distinct minima: depending on the conditions imposed, one may find either two disjoint curves (surfaces), one the $q \rightarrow (-q)$ image of the other, or a single quaternion curve (surface) that contains its own reflection, such as that in Figure 6(b). The families of frame manifolds containing their own reflected images have minima distinct from the disjoint families.

We now present some simple examples to give a feeling for the process.

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 7, we can compute the ring constraints on the possible quaternion frames at each sample point and let the Evolver automatically find the optimal framing. The results 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.

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 8, 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.

The types of solutions we find are 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 could be similarly specified (see, e.g., Barr et al. [2, 22]).

Surface Patch Framings. A classic simple example of a surface patch framing problem was presented in the discussion of Figures 2 and 5. This problem can also be handled naturally by the Evolver: we choose an initial quaternion frame for the mesh and minimize the area in quaternion space subject to the constraints that the normals remain unchanged. That is, the frame choices may only slide around constraint rings such as those depicted in Figure 5(b) for the frames at the corners. The results are shown in Figures 10 and 9. As a test, we started one case in a random initial state with a range of 2π in the starting values. All converged to the same optimal final framing. While more complex examples could be given, all the essential features of the method short of dealing rigorously with non-trivial topological manifolds are illustrated by this surface patch example.

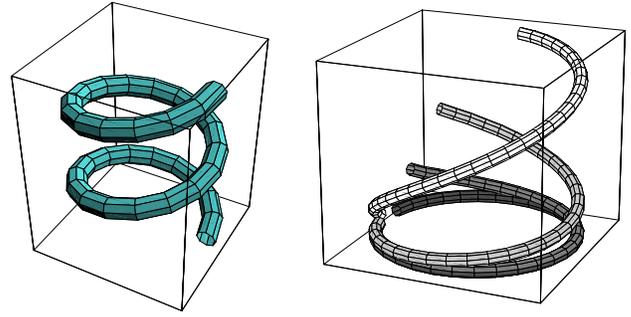


Figure 7: Helix (left) and its evolving quaternion frames (right). Starting from the Geodesic Reference quaternion frame for a single turn of the helix, the very dark gray circle, the Evolver produces these intermediate steps while minimizing the total quaternion curve length subject to the constraints 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; note that the *quaternion length* of the white curve is the shortest, even though in this projection that is not obvious. The numerical energies of the four 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.

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., [18, 10]). 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 surfaces. 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, 22]) to smoothly compute intermediate frames subject to the constraints. For more general constrained navigation methods like those described by Hanson and Wernert [13]), 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.

4 Conclusion

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 en-

gines with texturing. 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 texturing based on 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. Work remaining to be done in the future includes applying the method to other domains such as geometric modeling and viewpoint interpolation, studying more carefully the topologically distinct minimal quaternion area solutions found for certain surface framings, and studying more challenging problems in the surface domain, e.g., topological tori with various numerical bumps and deformations are known to admit global frames, but little is known about how to compute good ones, and this method is a logical candidate.

Acknowledgments

The author gratefully acknowledges the cordial hospitality of Claude Puech and the members of the iMAGIS laboratory, a joint project of CNRS, INRIA, Institut National Polytechnique de Grenoble, and Université Joseph Fourier, where this research was initiated. Thanks are also due to Ken Brakke for help with the Evolver. This research was made possible in part by NSF infrastructure grant CDA 93-03189.

References

- [1] S. L. Altmann. *Rotations, Quaternions, and Double Groups*. Oxford University Press, 1986.
- [2] Alan H. Barr, Bena Currin, Steven Gabriel, and John F. Hughes. Smooth interpolation of orientations with angular velocity constraints using quaternions. In Edwin E. Catmull, editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 313–320, July 1992.
- [3] M. Berger. *Geometry I, II*. Springer Verlag, Berlin, 1987.
- [4] Richard L. Bishop. There is more than one way to frame a curve. *Amer. Math. Monthly*, 82(3):246–251, March 1975.
- [5] Jules Bloomenthal. Calculation of reference frames along a space curve. In Andrew Glassner, editor, *Graphics Gems*, pages 567–571. Academic Press, Cambridge, MA, 1990.
- [6] Kenneth A. Brakke. The surface evolver. *Experimental Mathematics*, 1(2):141–165, 1992. The “Evolver” system, manual, and sample data files are available by anonymous ftp from geom.umn.edu, The Geometry Center, Minneapolis MN.
- [7] L. P. Eisenhart. *A Treatise on the Differential Geometry of Curves and Surfaces*. Dover, New York, 1909 (1960).
- [8] S. Gabriel and James T. Kajiya. Spline interpolation in curved space. In *State of the Art Image Synthesis*, 1985. Siggraph '85 Course notes.
- [9] Alfred Gray. *Modern Differential Geometry of Curves and Surfaces with Mathematica*. CRC Press, Inc., Boca Raton, FL, second edition, 1998.
- [10] Cindy M. Grimm and John F. Hughes. Modeling surfaces with arbitrary topology using manifolds. In *Computer Graphics Proceedings, Annual Conference Series*, pages 359–368, 1995. Proceedings of SIGGRAPH '95.
- [11] A. J. Hanson and H. Ma. Quaternion frame approach to streamline visualization. *IEEE Trans. on Visualiz. and Comp. Graphics*, 1(2):164–174, June 1995.
- [12] A. J. Hanson and Hui Ma. Visualizing flow with quaternion frames. In *Proceedings of Visualization '94*, pages 108–115. IEEE Computer Society Press, 1994.
- [13] A. J. Hanson and E. Wernert. Constrained 3d navigation with 2d controllers. In *Proceedings of Visualization '97*, pages 175–182. IEEE Computer Society Press, 1997.
- [14] J. T. Kajiya. Anisotropic reflection models. In *Computer Graphics*, volume 19, pages 15–21, 1985. Proceedings of SIGGRAPH '85.
- [15] Myoung-Jun Kim, Myung-Soo Kim, and Sung Yong Shin. A general construction scheme for unit quaternion curves with simple high order derivatives. In *Computer Graphics Proceedings, Annual Conference Series*, pages 369–376, 1995. Proceedings of SIGGRAPH '95.
- [16] F. Klock. Two moving coordinate frames for sweeping along a 3d trajectory. *Computer Aided Geometric Design*, 3, 1986.
- [17] N. L. Max. Computer representation of molecular surfaces. *IEEE Computer Graphics and Applications*, 3(5):21–29, Aug 1983.
- [18] J. Milnor. *Topology from the Differentiable Viewpoint*. The University Press of Virginia, Charlottesville, 1965.
- [19] G. M. Nielson. Smooth interpolation of orientations. In N.M. Thalmann and D. Thalmann, editors, *Computer Animation '93*, pages 75–93, Tokyo, June 1993. Springer-Verlag.
- [20] John C. Platt and Alan H. Barr. Constraint methods for flexible models. In John Dill, editor, *Computer Graphics (SIGGRAPH '88 Proceedings)*, volume 22, pages 279–288, August 1988.
- [21] D. Pletincks. Quaternion calculus as a basic tool in computer graphics. *The Visual Computer*, 5(1):2–13, 1989.
- [22] Ravi Ramamoorthi and Alan H. Barr. Fast construction of accurate quaternion splines. In Turner Whitted, editor, *SIGGRAPH 97 Conference Proceedings, Annual Conference Series*, pages 287–292. ACM SIGGRAPH, Addison Wesley, August 1997. ISBN 0-89791-896-7.
- [23] John Schlag. Using geometric constructions to interpolate orientation with quaternions. In James Arvo, editor, *Graphics Gems II*, pages 377–380. Academic Press, 1991.
- [24] Uri Shani and Dana H. Ballard. Splines as embeddings for generalized cylinders. *Computer Vision, Graphics, and Image Processing*, 27:129–156, 1984.
- [25] K. Shoemake. Animating rotation with quaternion curves. In *Computer Graphics*, volume 19, pages 245–254, 1985. Proceedings of SIGGRAPH 1985.
- [26] Ken Shoemake. Fiber bundle twist reduction. In Paul Heckbert, editor, *Graphics Gems IV*, pages 230–236. Academic Press, 1994.

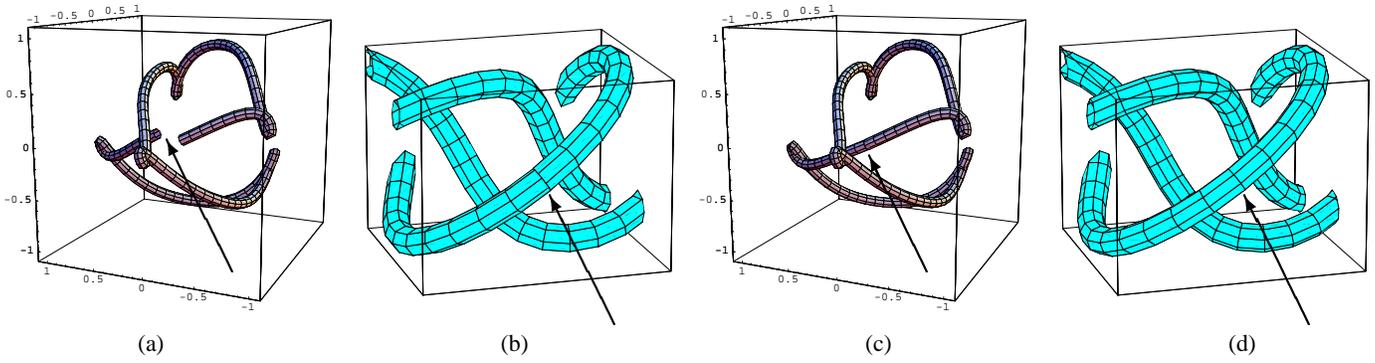


Figure 8: Optimization of the non-periodic parallel transport frame of the (3,5) torus knot introduced in Figure 1 to produce a nearby periodic framing. (a) The original quaternion parallel transport frame used to produce the tubing in Figure 1(b,c). (b) The frame mismatch, repeated for completeness. (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.

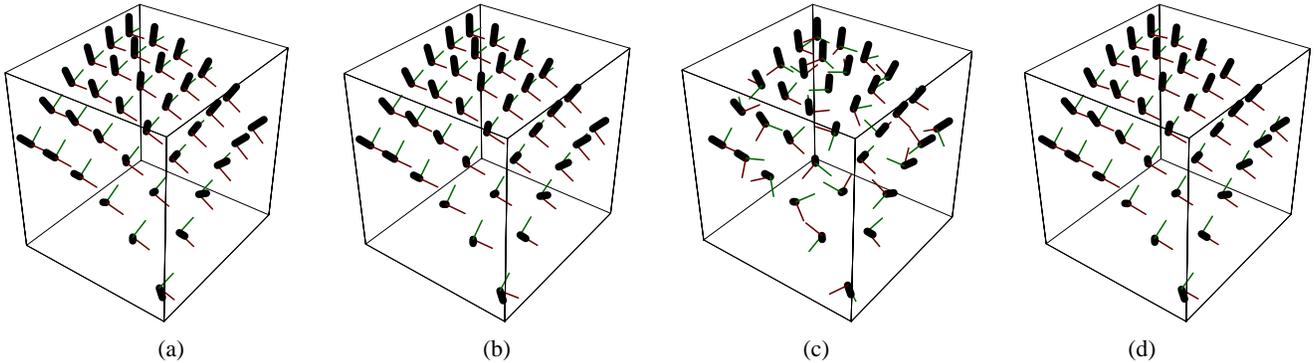


Figure 9: Study of *possible* and *optimal* reference frames on a surface patch; the corresponding quaternion fields are given in Figure 10. (a) The Geodesic Reference frames for the small patch of Figure 2. (b) Two-step parallel transport frames. (c) Random frames. (d) The unique frame configuration resulting from minimizing area in quaternion space with the upper corner fixed.

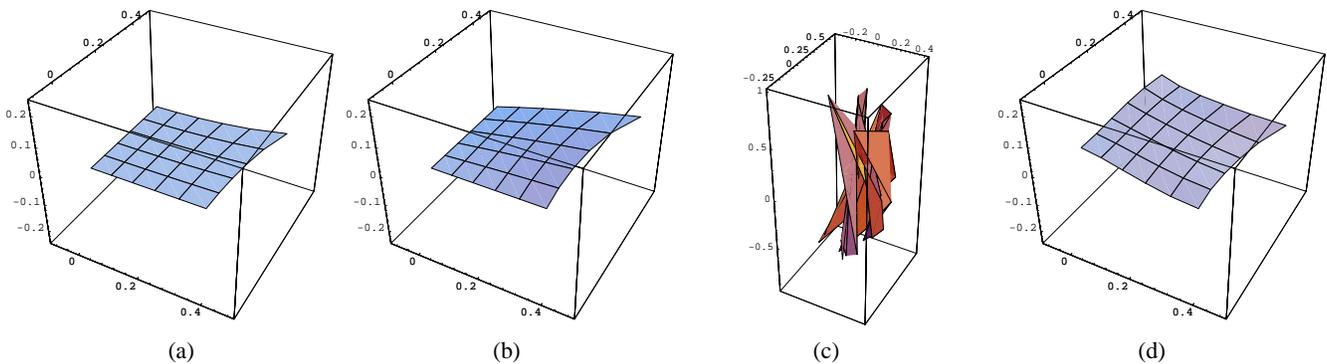


Figure 10: Quaternion areas corresponding to the frame assignments in Figure 9. (a) The initial Geodesic Reference quaternions for the small patch shown in Figure 2. (b) Initial quaternions from parallel transporting the vertex frame down one edge, and then across line by line. (c) A random starting 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, but is distinct.