Citation for published version:

Faraway, J & Choe, SB 2009, 'Modelling orientation trajectories', *Statistical Modelling*, vol. 9, no. 1, pp. 51-68.
<https://doi.org/10.1177/1471082X0800900104>

DOI:

[10.1177/1471082X0800900104](https://doi.org/10.1177/1471082X0800900104)

Publication date:

2009

Document Version

Peer reviewed version

[Link to publication](#)

Faraway, J. J., & Choe, S. B. (2009). Modelling orientation trajectories. *Statistical Modelling*, 9(1), 51–68.
Copyright © 2009 SAGE Publications. Reprinted by permission of SAGE Publications.

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Modeling orientation trajectories

Julian J. Faraway* and Su Bang Choe
University of Bath and University of Michigan

January 3, 2008

Abstract

We describe a modeling approach for orientation trajectories. The initial and final orientations of the object are taken as known and the problem of how the object will transition between the endpoints is considered. Orientations are represented as quaternions and mapped to a tangent space. The deviation from the geodesic connecting the endpoints is introduced and called the slerp residual. Modeling the slerp residuals substantially reduces the distortion caused by mapping to the tangent space. Bézier curves are used to compactly model the trajectories and provide a linkage to potential predictors.

Data from an experiment to study human motion while reaching to perform a wide variety of grasps is considered. The orientation trajectories of the hand are modeled in several ways resulting in a simple yet interpretable model.

Keywords: Functional data analysis, Bézier curves, quaternions

1 Introduction

Imagine that, while seated at your desk, your pen drops to the floor. You turn your head to locate the pen and reach down to retrieve it. The orientation of your head will follow some trajectory during this task. Your hand, torso and pelvis will also rotate in space during your motion. The shapes of these orientation trajectories will depend on the location of the pen, perhaps your height and age, and potentially many other variables. If you repeat the task, the trajectories will not be exactly the same. Furthermore, another person, even if they are quite similar to you, will trace out different trajectories. The purpose of this article is to describe a modeling approach to show how these trajectories might depend in a systematic way on some predictors and how they vary from repetition to repetition.

The methodology developed in this article is motivated by problems in Ergonomics. Ergonomics applies models of human capability and limitations to improve the interaction between people and products or workplaces. In current practice, ergonomists aim to anticipate problems with new vehicle or workplace designs before they are deployed. Of course, it is possible to test prototype designs on real humans, but this is time consuming and expensive. Virtual humans can

*Department of Mathematical Sciences, University of Bath, BA2 7AY, United Kingdom, jjf23@bath.ac.uk

be used to test the design in a virtual world to ensure that the proposed tasks can be performed safely and efficiently. Digital human models (DHM) represent variation in human size, shape, and movement. The orientation trajectory model described here was developed as part of an effort to improve the realism of the movement simulation capability in DHMs. The goal is to predict movements as a function of task and human characteristics. The model should yield not only typical or average movements, but also describe the variation that can be expected.

The modeling of orientation data has attracted a moderate amount of interest in Statistics. Downs (1972) develops a distribution on rotations and uses a tangent space approach to modeling. Khatri and Mardia (1977) follow an alternative parametric approach to orientation modeling, while Prentice (1986) used a quaternion-based approach. Orientation data can be viewed as a subset of directional data as described in the book Mardia and Jupp (2000).

Practical applications showing statistical modeling of orientation data can be found in a variety of fields: Biomechanics in Rancourt et al. (2000), Vectorcardiography in Downs et al. (1971) and Geophysics in Chang (1986). The research presented here differs in that it relates to orientation curves or trajectories rather than point observations. Prentice (1987) showed how to fit smooth curves to a sequence of orientation data, but our focus is on modeling the curves rather than fitting them.

There is substantial interest in orientations in the human animation literature. The article of Shoemaker (1985) has been influential in introducing the use of quaternions in orientation modeling within this field. There is also interest in generating, interpolating and filtering orientation curves — see for example, Park and Ravani (1997) and Lee and Shin (2002). However, the focus of the field has generally been non-statistical, where sometimes motion predictions are based on little or no direct data. This is not a criticism as the approach reflects the needs and objectives of the field. Quaternion-based approaches to orientation modeling may be found in Ergonomics and Biomechanics. Coburn and Crisco (2005) is a recent example where the interest is in interpolation rather than statistical methods.

The orientation trajectory data we will model is described in Section 2. This is used to motivate the statistical approach for modeling orientation trajectories in Section 3. We demonstrate the application to modeling hand trajectories in Section 4 and close with a discussion in Section 5.

2 Data and Motivation

The data described here was collected at the Human Motion Simulation Laboratory (HuMoSim) at the University of Michigan. The particular experiment was designed to study the dynamic postures of subjects who were asked to grasp blocks with a variety of hand grips. The experimental apparatus consisted of three pods of five blocks each. The subject was told to grasp a block with the right hand, placing the thumb on one of the four sides of a cube and the fingers on the opposite side. The subject was also instructed to either push or pull on the block once it had been grasped. The pod of five blocks consisted of one centrally placed block facing the subject and four others angled away to the right, left, top and bottom of the central block. Three pods were attached to a centrally-placed tower, with the lowest at seat height, the middle one around neck level and the highest slightly above head height. Thus there were three pods by five blocks by four thumb grips for a total of 60 possible tasks. However, some of these motions are physically impossible without

getting up from the seat while some are merely difficult. Thus only 35 of the 60 possible motions were requested of the subjects. The apparatus is shown in Figure 1.

Figure 1: The first picture shows a side view of a pod of five blocks. The subject is instructed to grip a chosen block using the right hand with the thumb on a specified side with the fingers gripping the opposing side. Subjects were told to push or pull on the cube once they had grasped it. The second picture shows the arrangement of the three pods of blocks on a tower relative to the seat.

The subjects were selected to provide a means to assess the effects of anthropometry, gender, and age on the motions. The subjects ranged from short to tall and from 20 to 70 years of age. Ten were male and ten were female. All were right handed. A total of 1088 motions were performed by the group of 20 subjects. Some motions were lost due to data collection errors and some were not included by design for all subjects. The data discussed here is a subset of a larger sequence of experiments and our objective in this article is to introduce the modeling methodology rather than make general claims about hand trajectories. Modeling the orientation trajectories is only one part of the larger objective of modeling the whole motion of the body in performing these tasks.

The motion capture (i.e. data collection) was achieved with an optical reflective marker system recording at 25Hz combined with a magnetic system. The motion of the whole body was recorded, but for the purposes of this example, we shall focus on the right arm only. We may obtain the 3D trajectories of the right shoulder, elbow and wrist from the optical system. The externally placed markers are projected to estimated joint centers. A magnetic marker attached to the back of the hand is able to capture both the location and orientation of the hand. Substantial effort is required to merge the two data collection systems. We did not attempt to capture the motion of the fingers as this is difficult in these circumstances.

The type of model presented here could be used in a wide range of applications, but the needs of the particular field of ergonomics for industrial design and manufacturing have motivated our choices. We envisage that a user of these models has, for example, a specific industrial workstation of known dimension in mind. The user will specify a virtual human of known size and shape and

provide other pertinent information such as age, weight and sex. The user will also specify the starting and ending location and orientation of the hand. This corresponds to the circumstances of our experiment where the task constrains the final location and orientation of the hand. The subjects are instructed to start with the hand palm down on the right thigh, so this is also fixed.

Thus the specific application of the proposed methodology is to model how the hand transitions from one known orientation to another known orientation. One can certainly envisage other problems where the starting and ending orientations are not known. These would need to be predicted before our method for orientation trajectory prediction could be applied.

Orientation trajectory prediction is only one part of a complete model for human motion. An overview of the human motion modeling project of which this methodology forms a part can be found in Faraway and Reed (2007). Consider the problem of predicting the hand orientation trajectory while performing a task such as those described above. One must decide whether to predict the orientation with respect to a global or local coordinate system. The use of global coordinates has the disadvantage that the orientation of the hand is affected not only by its rotation around the wrist, but by the motion of the arm and the rest of body. We have found that the use of a local coordinate system, that computes the orientation of the hand with respect to the triad of points formed by the wrist, elbow and shoulder, is preferable. Within this local system of coordinates, the range of possible orientations is limited by the flexibility of the wrist joint. This also requires prediction or knowledge of the location trajectory of the wrist and the orientation of the arm. We have developed methods for predicting these components which are described elsewhere — Faraway (2004) and Faraway et al. (2007).

Of course, we wish to develop as simple a model as possible for the orientation trajectory while still achieving a good fit to the data. However, there are two specific reasons why we shall have a stronger preference towards simplicity as opposed to fit: Firstly, the orientation trajectory model is only a small part of the complex DHM software used for simulations of virtual workers. Because the model must be readily integrated into existing software, it is essential that the model be easy to implement from a written description of the algorithm. Secondly, stable extrapolation is an important concern. Future users of the design software may attempt to use the model under circumstances quite different from the data from which the model was derived. Ideally, proper caution would be taken, but, in practice, this cannot be guaranteed. Therefore, it is essential that when the model is extrapolated, that stable and credible predicted orientation trajectories are obtained. Given natural human variation, precise predictions are not to be expected or indeed required. However, it is important that predictions not be badly wrong. Simpler models are therefore preferred because their extrapolation properties are easier to anticipate.

3 Methods

3.1 Geometry

Orientation is the attitude of a rigid body in space and is expressed as a rotation with respect to a fixed coordinate frame. An orientation can be represented as an orthogonal 3×3 matrix with unit determinant and thus lies in $SO(3)$. It can also be represented (with some redundancy) as a point on the surface of a unit sphere in four dimensions, S^3 . It is difficult to do statistics with data on the surface of a sphere because the space is non-Euclidean. Thus a convenient approach is to take

Figure 2: A depiction of tangent mapping at the origin to the space of quaternions. In reality, orientations lie on the surface of a four-dimensional unit sphere S^3 while the tangent space has three dimensions. Geodesics, starting from the origin, map to straight lines in the tangent space.

the tangent at some chosen origin on the sphere and map all data onto this tangent space. This would enable the use of standard statistical methods. Any predictions or other conclusions could be mapped back to the orientation space. However, if the orientations vary far from the origin, the approximation error may be unacceptably large. For the hand trajectory data analyzed here, there is much too wide a range of orientations for this approach to work.

Now consider two distinct orientations and the shortest trajectory linking them. This geodesic was named the *slerp* or *spherical linear interpolator* by Shoemake (1985) in the context of animation problems. Take one of these two orientations as the origin in S^3 and construct the tangent space at that point. The slerp in this tangent space will then be the linear interpolant between the two orientations mapped into that space. The situation is depicted in Figure 2. For slerps between orientations away from the origin, the linear interpolation no longer holds in the tangent space. The distortion increases as we move away from the origin.

We use the idea of a *slerp residual* to substantially reduce the distance from the origin. Compare the observed trajectory and the slerp trajectory. For any point on the slerp trajectory, we compute the rotation that maps to the corresponding point on the observed trajectory. (We shall clarify what is meant by corresponding later.) These rotations can be connected along the whole trajectory to give the slerp residual trajectory. The slerp residual trajectory begins and ends at the origin. For our data, the observed trajectories are not too far from a slerp, and so the slerp residual trajectory typically does not vary far from the origin, thus avoiding the worst of the distortion. The slerp residual is depicted in Figure 3.

The geometry may provide some intuition into the method we propose, but it does not indicate how these quantities may actually be computed. We shall now show that quaternions provide a convenient way to do this.

Figure 3: The slerp residual is the rotation from the slerp to the observed trajectory. The slerp residual is itself an orientation trajectory.

3.2 Quaternions

There are several ways to parametrize orientations. Three methods in common use are 3×3 rotation matrices, Euler angles, sometimes called the roll, pitch and yaw angles in nautical or aeronautical parlance and the *axis-angle* parametrization that derives from a direct reading of Euler's rotation theorem. See Zatsiorsky (1998) for their application to human motion modeling. However, all three methods have mathematical disadvantages.

Quaternions provide a more elegant approach to the modeling of orientation. First introduced by Hamilton in 1843, they are a four dimensional extension of complex numbers. We provide a brief introduction here. We write a quaternion $\mathbf{q}$ as:

$$\mathbf{q} = [\mathbf{v}, w] = \mathbf{i}x + \mathbf{j}y + \mathbf{k}z + w$$

where $x, y, z, w \in \mathbb{R}$ and $\mathbf{v} \in \mathbb{R}^3$. $\mathbf{i}^2 = \mathbf{j}^2 = \mathbf{k}^2 = \mathbf{ijk} = -1$. $\mathbf{v}$ is known as the vector and w as the scalar.

Two quaternions, $\mathbf{q}_1$ and $\mathbf{q}_2$, may be added simply by adding the corresponding vector and scalar components. Their product is:

$$\mathbf{q}_1 * \mathbf{q}_2 = [\mathbf{v}_1 \times \mathbf{v}_2 + w_1 \mathbf{v}_2 + w_2 \mathbf{v}_1, w_1 w_2 - \mathbf{v}_1 \cdot \mathbf{v}_2]$$

where $\times$ represents cross product between vectors. Notice that multiplication is not commutative.

The dot product is given by $\mathbf{q}_1 \cdot \mathbf{q}_2 = w_1 w_2 + \mathbf{v}_1 \cdot \mathbf{v}_2$. The conjugate, $\mathbf{q}^*$, is $[-\mathbf{v}, w]$ while the norm is defined by $\|\mathbf{q}\|^2 = \mathbf{q} * \mathbf{q}^* = w^2 + x^2 + y^2 + z^2$. The inverse is then $\mathbf{q}^{-1} = \mathbf{q}^* / \|\mathbf{q}\|$. Like complex numbers, quaternions can be put in polar form:

$$\mathbf{q} = re^{\hat{\mathbf{v}}\theta/2} = \|\mathbf{q}\|(\hat{\mathbf{v}} \sin \frac{\theta}{2}, \cos \frac{\theta}{2})$$

where $\|\mathbf{q}\|$ is called the magnitude of the quaternion, $\theta/2 = \arccos(w/\|\mathbf{q}\|)$ is called the *angle* and $\hat{\mathbf{v}} = \mathbf{v}/\|\mathbf{v}\|$ is the *axis*.

The quaternion representation of an orientation is obtained by restricting to unit quaternions ($\|\mathbf{q}\| = 1$) with the axis and angle of the orientation being applied in the polar form. Suppose $\mathbf{x}$ is a vector in $\mathbb{R}^3$ which we want to rotate about an axis $\mathbf{a}$ by an angle θ . The rotated vector is then:

$$\mathbf{q} * \tilde{\mathbf{x}} * \mathbf{q}^{-1} \quad \text{where} \quad \mathbf{q} = (\mathbf{a} \sin(\theta/2), \cos(\theta/2))$$

where $\tilde{\mathbf{x}}$ is the pure quaternion form of $\mathbf{x}$, that is $\tilde{\mathbf{x}} \equiv [\mathbf{x}, 0]$. The double multiplication explains the use of $\theta/2$ rather than θ in the polar form. Thus unit quaternions, lying in S^3 , may be used to represent an orientation.

Consider the tangent space at the identity quaternion, $[\mathbf{0}, 1]$. The *exponential map* allows a map from a vector $\mathbf{v} \in \mathbb{R}^3$, lying in the tangent space to the space of unit quaternions by:

$$\text{expmap}(\mathbf{v}) = \begin{cases} [\mathbf{0}, 1] & \text{if } \mathbf{v} = \mathbf{0} \\ [\hat{\mathbf{v}} \sin(\theta/2), \cos(\theta/2)] & \text{otherwise} \end{cases}$$

where $\theta = \|\mathbf{v}\|$ and $\hat{\mathbf{v}} = \mathbf{v}/\theta$. The map can be inverted using the logarithm map:

$$\text{logmap}(\mathbf{q}) = \mathbf{v}/\text{sinc}(\theta/2)$$

where $\mathbf{v}$ is the vector part of $\mathbf{q}$ and $\text{sinc}(x) = \sin x/x$. For x close to zero, $\text{sinc}(x) \approx 1 + x^2/6$.

There are some drawbacks to this mapping. First, we must restrict θ to $[0, \pi]$. Furthermore, the antipodal points $\mathbf{q}$ and $-\mathbf{q}$ represent the same orientation.

We use the angle metric which is also the geodesic distance. The distance between a given unit quaternion, representing a rotation, and the identity quaternion is preserved as the Euclidean distance in the tangent space. However, in general the distance between two points in S^3 is not preserved in the tangent space. The difference between the two distances increases the further they are from the identity. Thus, one can avoid some distortion caused by performing statistics in the tangent space by centering the data on the origin. For more on the exponential map, see Grassia (1998).

The slerp interpolates along the shortest geodesic path between two quaternions, $\mathbf{q}_1$ and $\mathbf{q}_2$:

$$\text{slerp}(\mathbf{q}_1, \mathbf{q}_2, \alpha) = \frac{\sin((1 - \alpha)\theta/2)}{\sin(\theta/2)} \mathbf{q}_1 + \frac{\sin(\alpha\theta/2)}{\sin(\theta/2)} \mathbf{q}_2$$

where $\cos(\theta/2) = \mathbf{q}_1 \cdot \mathbf{q}_2$ and $\alpha \in [0, 1]$. Linear interpolation in the tangent space is not equivalent to the slerp unless one of the endpoints is the origin. The slerp could be regarded as a null model for predicting the orientation trajectory between one orientation and another. Indeed, this is a common approach in computer animation.

3.3 Slerp Residuals

The steps necessary to compute a slerp residual are:

1. Compute a sequence of quaternions representing the observed trajectory: For each frame of data, we compute a 3×3 rotation matrix representing the orientation of the hand relative to

the arm. There will be a choice of coordinate systems and origin, but the main concern is to pick one consistently. The `orientlib` R package of Murdoch (2006) provides a convenient way to compute the quaternions from the rotation matrix representation.

2. Remove any discontinuities due to antipodal symmetry: The two unit quaternions, $\mathbf{q}$ and $-\mathbf{q}$, represent the same orientation. Due to the vagaries of trigonometry in converting to quaternions, it sometimes happens that the angular distance between successive quaternions on the trajectory is close to π instead of being quite small, representing a switch between antipodes. We remove such apparent discontinuities by reflecting the second quaternion to lie close to the first. We sequentially process all curves in this manner to remove such discontinuities.
3. Compute the observed speed profile along the trajectory from spacing of the observed orientations: We observe that the hand does not rotate along its trajectory at a constant speed. We compute the angular speed profile along the trajectory. Consider a sequence of observed orientations: $\mathbf{q}_0, \dots, \mathbf{q}_n$ and compute the step lengths:

$$s_i = \arccos |\mathbf{q}_i \cdot \mathbf{q}_{i-1}|$$

In practice, the speed profile can be adequately modeled with a Beta density. We would need such a model to make predictions of orientation trajectories.

4. Compute a slerp between the two endpoints of the observed trajectory with the points on the slerp placed with the observed speed profile: We compute a slerp between the observed endpoints with same speed profile as the observed trajectory. Let $\alpha_i = \sum_{j=1}^i s_j / \sum_{j=1}^n s_j$ (with $\alpha_0 = 0$) and define the empirical slerp sequence $\mathbf{s}_i = \text{slerp}(\mathbf{q}_0, \mathbf{q}_n, \alpha_i)$.
5. Compute the slerp residual as the sequence of rotations from points on the slerp to points on the observed trajectory using:

$$\mathbf{r}_i = \mathbf{s}_i^* * \mathbf{q}_i$$

The slerp residual orientation trajectories begin and end at the origin.

6. Convert the slerp residual sequence to a sequence of 50 points using interpolation: In our raw data, the trajectories are of different lengths since they take varying amounts of time to complete. To allow for comparison between trajectories, we convert them all to a common length using linear interpolation. The clock time taken to complete the trajectory is saved for separate modeling and potentially for use as a predictor of the trajectory.
7. Map the slerp residual to the tangent space at the origin using the log map.

Some variations on this procedure are worth considering. We have chosen to match points along the observed trajectory and slerp using the relative arc length. Thus we are using the relative time as the index for our curves. We could have done this indexing using the location of the hand as it translates. However, this latter choice means would require that we link the models for the orientation and location of the hand, thus adding more complexity to the overall modeling.

We now wish to describe how these residual trajectories vary and how they might depend on predictors. We will fit the trajectories with Bézier curves and then link the fitted parameters (control points) to the predictors.

3.4 Slerp Residual Trajectory Modeling

Bézier curves are widely used in graphics, but have attracted little attention in statistics. Hence, we present a brief introduction here. The Bernstein polynomials of degree d for $i = 0, \dots, n$ for $t \in [0, 1]$ are defined as:

$$B_i^d(t) = \binom{d}{i} t^i (1-t)^{d-i}$$

A Bézier curve of degree d has the form:

$$C(t) = \sum_{i=0}^d P_i B_i^d(t)$$

The control points, P_i , determine the shape of the curve and can be multivariate if desired; we require three dimensions in our case. The appearance of these curves for $d = 3$ in two dimensions is illustrated in Figure 4: The start and endpoints, P_0 and P_d are coincident with the origin in

Figure 4: Cubic ($d = 3$) Bézier curve approximation of a slerp residual trajectory. Line segments between the control points, P_0P_1 and P_2P_3 are tangential at the origin. P_0 and P_3 coincide with the origin.

the case of a slerp residual trajectory. The shape of the curve is then determined by the interior control points. The geometric construction for Bézier curves using the De Casteljau algorithm that motivates their use in graphic design can be found in many texts such as Prautzsch et al. (2002).

For the case that $d = 1$, the residual trajectory does not leave the origin and could be regarded as a null model since it corresponds to the slerp trajectory. For $d = 2$, there is a single control point. The residual trajectory moves towards this control point, reaching half of the distance from the origin to the control point at $t = 1/2$ and then returning along the same line. For $d = 3$, the residual trajectory lies in the plane formed by the origin and the two control points, as seen in Figure 4. The tangents at the origin for the outgoing and incoming curves are given by the line segments from the origin to the respective control points. This tangent property is shared for curves for $d > 3$ where the line segments run to the first and last non-origin control points.

Consider an observed sequence of points along a slerp residual trajectory : $Z_0, \dots, Z_n$. We wish to fit a Bézier curve $C(t)$ of degree d for $t \in [0, 1]$ such that $C(0) = Z_0 = \mathbf{0}$ and $C(1) = Z_n = \mathbf{0}$. We

define the sum of squares for the fit as:

$$SS = \sum_{i=1}^{n-1} |Z_i - C(\alpha_i)|^2$$

We choose the control points of C to minimize SS which is a straightforward least squares problem.

Let the vector of predictors associated with each trajectory be $(x_1, \dots, x_p)$, which in our example might represent variables such as age, height and gender. We aim to find models of the general form:

$$g(\hat{P}_i) = f(x_1, \dots, x_p) + \varepsilon \quad i = 1, \dots, d-1$$

We can then model this multivariate response $g(P_i)$ in three dimensions. Textbook multivariate linear modeling techniques can be applied. Alternatively, we could use spherical coordinates. The radial component can then be modeled using ordinary linear modeling methods while the spherical component can be modeled using regression techniques for a spherical response. See Mardia and Jupp (2000) for example.

It would be possible to model the trajectories with more complex and flexible methods such as B-splines. However, this would make it more challenging to link the control points to the predictors in an interpretable and compact manner.

4 Application

The data we shall use consists of the 3D coordinates of three optical markers on the right arm and a magnetic marker on the right hand that can record orientation. The data are collected at the rate of 25Hz. There are 1088 successful motions in database. Because the cameras are turned on continuously, the data contains some frames of motion from before and after the actual task. There is no single obvious way to trim the unwanted frames because different parts of the body start and stop at different times. We trimmed the data based on the proximity of the wrist to the starting and ending locations. Next, the orientation of the hand was computed relative to a coordinate system defined by the three markers on the wrist, elbow and shoulder and expressed in quaternion form.

It is difficult to display an orientation trajectory, that is a sequence of hand orientations, compactly. Displaying much more than one on the same plot directly seems futile. However, we can compute the empirical slerp residuals and display these since these are 3D trajectories, beginning and ending at the origin. We computed these for all the reaches for a specific task combination as seen from the three axis directions in Figure 5. Although there is substantial variation, we can see some consistency, particularly in the first and third views. The trajectories move away from and back to the origin in a somewhat similar way. Random trajectories in all directions would be expected for all three plots if the trajectories varied about a slerp in a random manner. The relatively high overall level of noise indicates that it will be difficult to find highly predictive models and yet there is some structure to capture. Plots for other task combinations are qualitatively similar.

In Figure 6, we show all the reaches of one subject to one pod. We might expect there to be greater variation in this set of plots than in Figure 5 because tasks might be expected to be more variable than subjects performing the same task. However, we see the general appearance of the plots is similar, with some consistency underlying the noise. Note that we are not claiming that

Figure 5: Empirical slerp residual trajectories for all reaches to one task combination. 29 trajectories are shown. All trajectories begin and end at the origin. The maximum distance from the origin for any of the trajectories is 1.0 radians. Three orthogonal equiscaled axis views are shown.

Figure 6: Empirical slerp residuals for all reaches of one subject to a single pod of five blocks. 24 trajectories are shown in all three coordinate directions.

this subject positions his or her hand in the same way, merely that they diverge in a similar way from a slerp in this positioning towards different final hand orientations.

Our objective is to find a good predictive model for the hand orientation trajectory. We have a strong preference for a simple model because this will be easier for users to implement and produce stable results in a wide range of conditions. In the analysis to follow, we will reject more complex models unless they achieve a significantly better fit for practical purposes. We are much less interested in questions of statistical significance for this application.

Due to natural human variation, we expect there to be some unavoidable error in the prediction, no matter how good the model is. Our first task is to estimate this lower bound on the error. Unfortunately, there are no exact replicates in the database. However, subjects were often asked to perform the same grasp to the same location, the only difference being whether they were to push or pull on the block once grasped. Using the methods described in this article and other studies of other aspects of this data revealed no difference in the motion due to the pushing or pulling action. We shall therefore treat these as replicates. There were 414 such replicates.

Let us denote the slerp residual residual trajectories by $r_{ijk}(t)$ where i varies over subjects, j

varies over tasks and k varies over replicates (at most two). Note that not all combinations exist in the data. The index t varies over $[0, 1]$ and describes the proportion of time taken. The trajectories are each represented by a sequence of 50 points in 3D where the first and last point lie at the origin.

Our measure of within subject error is:

$$\text{median}_{i,j} \max_t |r_{ijk}(t) - r_{ij\bullet}(t)|$$

where averaging is denoted by the $\bullet$ in the subscript. Only those trajectories with a replicate were considered and note that the maximum will be the same for either member of the pair. The value obtained was 0.125 radians. We do not expect that any model can do better than this. An alternative measure would be to consider the maximum distance between replicate pairs and compute the median over the data. This could be viewed as an estimate of the prediction error where the first member of the pair is used to predict the second. This would result in a value exactly twice the first criterion i.e. 0.250 radians. We will use the first measure to be consistent with the results to follow.

In our experience with demonstrating predicted human motion, the viewer tends notice to momentary larger errors much more than moderate errors over the whole course of the motion. This explains our use of the maximum in the criterion.

Our measure of between subject error is:

$$\text{median}_j \text{median}_{i,k} \max_t |r_{ijk}(t) - r_{\bullet j\bullet}(t)|$$

Here we are considering all the reaches of different subjects to perform the same task. There were 35 different task combinations that were performed out of the 60 potentially possible. The value obtained was 0.242 radians.

As we might expect, this between subject variation is larger than the within subject variation because of differences between the subjects. It remains to be seen whether any of this variation can be explained by observable characteristics such as height or gender.

The simplest approach to predicting orientation trajectories is to use the slerp. Thus we might measure the performance of such a prediction model with:

$$\text{median}_{ijk} \max_t |r_{ijk}(t)|$$

This gives a value of 0.397 radians. However, this is not a practical prediction model as it uses the observed speed profile for each slerp corresponding to the observed orientation trajectory. In practice, we would need to use single speed profile. We averaged across the observed speed profiles as computed in step 3 in Section 3.3 and used this average in the subsequent slerp residual calculation. This resulted in an error of 0.448 radians. This is the simplest possible model for the data and thus represents an upper bound on the acceptable error.

We use Bézier curves to approximate the slerp residual trajectories. For each slerp residual trajectory, $r_l(t)$ where l runs over all curves, we fit a Bézier curve obtaining control points, P_l^m where $m = 1, \dots, d - 1$ for degree d . Let $\hat{r}(t, P_l)$ be the predicted curve produced using control points P_l . We assess the approximation error using:

$$\text{median}_l \max_t |r_l(t) - \hat{r}(t, P_l)|$$

For $d = 2, 3, 4$, that is using 1, 2 or 3 control points, these errors were 0.126, 0.083 and 0.049 respectively. Of course, these are not predictive models, just approximations to the curve, but it does indicate that we do not need to use many control points to obtain a satisfactory fit. Use of higher order curves does allow a better fit to the individual curves as one might expect from Figures 5 and 6, but does not improve the fit in the actual prediction models discussed below.

Next we consider a series of null models, that is models without any predictors, by averaging over the estimated control points thus giving a single predicted slerp residual over all trajectories. Let $P_\bullet$ be the vector of averaged control points, then our measure of error is:

$$\text{median}_l \max_t |r_l(t) - \hat{r}(t, P_\bullet)|$$

The errors for $d = 2, 3, 4$ were 0.296, 0.371 and 0.294 respectively. Higher order fits were worse. We believe the poor fit obtained for $d = 3$ (which means two control points) is due to the lack of a central control point, which is present for $d = 2$ and $d = 4$. This might be apparent from the shape of the curves seen in Figures 5 and 6.

A related idea is to simply average the slerp residuals to obtain a single predicted trajectory. We measure this error as:

$$\text{median}_l \max_t |r_l(t) - r_\bullet(t)|$$

This gives a value of 0.288. We would not expect even higher order fits to better this and given that the $d = 2$ fit comes close, we might choose this for simplicity.

Given the difference in between and within subject variation, there is some hope that we might find significant subject specific effects. To investigate this, we average the control points within subject, obtaining median maximums of 0.281, 0.375 and 0.298 for $d = 2, 3, 4$ respectively. Although the value for $d = 2$, at least, is a small improvement over the corresponding null model, we were unable to find any relationship between the control points and measurable characteristics such as height, weight, age or gender. So the small improvement would only apply to predictions for the particular subjects of our experiment. We suspect that individuals have personal movement characteristics that cannot be predicted from observables. Thus we prefer the null model.

Similarly, we can investigate the possibility of task effects by averaging control points within tasks. The corresponding median maximums were 0.251, 0.377 and 0.278. Here there is a somewhat greater improvement for $d = 2$. Fitting a linear model to predict the sole internal control point for $d = 2$ reveals significant task effects. The orientation of the five blocks on the pod of five blocks has the single largest effect on the orientation trajectory followed by the location of the pod itself among the three locations available. The thumb position has the least effect. However, there are significant and substantial interactions between these variables and so such a task-based model is difficult to describe succinctly. Furthermore, it would be difficult to build a model that would extrapolate reliably to task conditions outside those in the experiment. In our judgment, the null model is preferable due its simplicity and stability whereas any task-based model would be more complex while fitting only somewhat better than the null, even in the best of circumstances.

Thus our preferred model uses a common slerp residual fitted with a Bézier curve defined by a single control point ($d = 2$) giving an error of 0.296 radians. Thus predicted orientation trajectories will diverge from the slerp in a common way defined by this control point, P_1 , and passing through $P_1/2$. We show this orientation with respect to the initial position in Figure 7. Thus, we see that the hand tilts downwards from the wrist with respect to a slerp. The reader might try a reach where the

Model	Error	$d = 2$	$d = 3$	$d = 4$
Within Subject	0.125			
Between Subject	0.265			
Slerp – Empirical	0.397			
Slerp — Average	0.448			
Approximation		0.126	0.083	0.049
Null		0.296	0.371	0.294
Null Average	0.288			
Subject Predictor		0.281	0.375	0.298
Task Predictor		0.251	0.377	0.278

Table 1: Median of maximum prediction errors for the orientation trajectory models considered. Models involving Bézier curve fits have errors given for degrees $d = 2, 3, 4$ while other models report a single error.

Figure 7: The picture on the left shows the hand in an (arbitrary) initial position while the picture on the right shows the hand in the orientation relative to this initial position as it passes through the point of maximum deviation from the slerp under the preferred model.

initial and final hand orientations are the same so that the slerp would keep the hand in the same posture throughout. Particularly if the reach is performed rapidly, the reader may find that the hand tends to trail the motion of the wrist causing the downward tilt that we are proposing here. Note that in reaches where the hand must change posture, the downward tilt is imposed on top of the slerp movement.

It may seem disappointing that we have been unable to find substantial significant task or subject effects, however, this is rather convenient from a practical point of view because the prediction model is simple and robust to implement and we have the satisfaction that it cannot be substantially improved upon.

5 Discussion

We have demonstrated a modeling approach for hand orientation trajectories that has good predictive performance and yet has interpretable elements. The model is simple to specify and implement. It is common practice to use slerp trajectories for orientations in human motion modeling, but even this limited dataset demonstrates that a substantial improvement may be achieved over the slerp without undue addition complexity. We would hesitate to claim that the specific model

for hand orientation trajectories applies very generally. In particular, when the hand must progress from one awkward posture to another, it would tend to pass through a more comfortable intermediate posture in practice. The modeling techniques presented here could be used to investigate such an effect.

Although we have considered just hand trajectories here, the same methodology can be applied to other body parts such as the head and the pelvis. For these parts, the final orientation may not be completely specified by the task. However the task will impose constraints which make predicting this final orientation much easier.

The method could be applied more generally to other orientation trajectories. The main limitation is that the observed trajectories not diverge too far from a slerp so that too much distortion due to the use of the tangent space can be avoided. In some cases, the divergence from a slerp may be too much — for example, consider an extended hand gesture as when using sign language. We recommend estimating the angular distance of the largest slerp residuals likely to be encountered. Generate two separate quaternions with this distance from the origin. Compute the slerp between these two both directly in the quaternion space and then via the tangent space. Compare these two slerps to judge the likely size of errors due to the tangent space approximation.

For more complex trajectories, more control points will be necessary and B-spline rather than Bézier modeling of the curves may be preferable. In our data example, the use of so few control points results in an essentially parametric analysis. However, with more control points, the method would approach a functional data analysis. In the example considered here, we were pleased to find an adequate constant model. In other cases, the regression modeling approach suggested here may reveal interesting predictors. The use of predicted control points provides a direct means of visualizing the predictor effect and avoiding unexpected extreme predictions.

As an alternative to the slerp residual and tangent space approach, we might attempt to perform statistics directly on the quaternion data. Unfortunately, even simple operations such as averaging, are problematic because naively averaging the coordinates of the quaternions will not result in a unit quaternion. Even so, some progress has been made — see some of the papers referenced in the introduction and Choe (2006) for some applications in human motion modeling.

Acknowledgments

We thank Don Chaffin and Matt Reed of the HuMoSim laboratory at the Center for Ergonomics at the University of Michigan for support and substantive advice. We thank the referees for a substantial improvement in the presentation of this paper.

References

- Chang, T. (1986). Spherical regression. *Annals of Statistics* 14, 907–924.
- Choe, S. (2006). *Statistical analysis of orientation trajectories via quaternions with application to human motion*. Ph. D. thesis, Department of Statistics, University of Michigan.
- Coburn, J. and J. Crisco (2005). Interpolating three-dimensional kinematic data using quaternion splines and hermite curves. *Journal of Biomechanical Engineering* 127, 311–317.

- Downs, T. (1972). Orientation statistics. *Biometrika* 59, 665–676.
- Downs, T., J. Liebman, and W. Mackay (1971). Statistical methods for vectorcardiogram orientations. In I. Hoffman, R. Hamby, and E. Glassman (Eds.), *Vectorcardiography 2: Proc. XIth International Symposium on Vectorcardiography*, Amsterdam, pp. 216–222. North-Holland.
- Faraway, J. (2004). Human animation using nonparametric regression. *Journal of Graphical and Computational Statistics* 13, 537–553.
- Faraway, J. and M. Reed (2007). Statistics for digital human motion modeling in ergonomics. *Technometrics* 47, 262–276.
- Faraway, J., M. Reed, and J. Wang (2007). Modeling 3D trajectories using Bézier curves with application. *Applied Statistics* 56, 571–585.
- Grassia, F. (1998). Practical parameterization of rotations using the exponential map. *The Journal of Graphics Tools* 3, 29–48.
- Khatri, C. and K. Mardia (1977). The Von-Mises Fisher distribution in orientation statistics. *Journal of the Royal Statistical Society, Series B* 39, 95–106.
- Lee, J. and S. Shin (2002). General construction of time-domain filters for orientation data. *IEEE Transactions on Visualization and Computer Graphics* 08(2), 119–128.
- Mardia, K. and P. Jupp (2000). *Directional Statistics*. Chichester: Wiley.
- Murdoch, D. (2006). *Orientlib: Support for orientation data*. R package version 0.8.3.
- Park, F. and B. Ravani (1997). Smooth invariant interpolation of rotations. *ACM Transactions on Graphics* 16, 277–295.
- Prautzsch, H., W. Boehm, and M. Paluszny (2002). *Bezier and B-Spline Techniques*. New York: Springer.
- Prentice, M. (1986). Orientation statistics without parametric assumptions. *JRSS-B* 48, 214–222.
- Prentice, M. (1987). Fitting smooth paths to rotation data. *Applied Statistics* 36, 325–331.
- Rancourt, D., L.-P. Rivest, and J. Asselin (2000). Using orientation statistics to investigate variations in human kinematics. *Applied Statistics* 49, 81–94.
- Shoemake, K. (1985). Animating rotation with quaternion curves. *ACM SIGGRAPH* 19, 245–254.
- Zatsiorsky, V. (1998). *Kinematics of Human Motion*. Champaign, IL: Human Kinetics.