Modeling 3D trajectories using Bézier curves with application to hand motion

Julian J. Faraway†
University of Bath, Bath, United Kingdom
Matthew P. Reed
University of Michigan, Ann Arbor, USA.
Jing Wang
Bank of America, Atlanta, USA.

[Received September 2006]

Summary. A modelling approach for 3D trajectories with particular application to hand reaching motions is described. Bézier curves are defined by control points which have a convenient geometrical interpretation. A fitting method for the control points to trajectory data is described. These fitted control points are then linked to covariates of interest using a regression model. This allows the prediction of new trajectories and the ability to model the variability in trajectories.
The methodology is illustrated with an application to hand trajectory modelling for ergonomics. Motion capture was used to collect a total of about 2000 hand trajectories performed by 20 subjects to a variety of targets. A simple model with strong predictive performance and interpretablility is developed. The use of hand trajectory models in the digital human models for virtual manufacturing applications is discussed.

1. Introduction

When you take a book from a shelf, your hand travels along a three-dimensional trajectory. The shape of this trajectory may depend on the location of the book, your height and age, the weight of the book and potentially many other variables. If you repeat the task, the trajectory will not be exactly the same. Furthermore, another person, even if they are quite similar to you, will trace out a different hand trajectory. The purpose of this article is to describe a modeling approach to show how these trajectories depend in a systematic way on some predictors and how they vary from repetition to repetition.

The particular application that motivated this article comes from Ergonomics. Ergonomics applies models of human capability and limitations to improve the interaction between people and products or workplaces. Historically, industrial ergonomists reacted to problems as they arose, redesigning tasks and equipment to improve productivity and reduce injuries. Now, however, ergonomists aim to intercept these problems before they occur. For example, a new manufacturing plant can be designed in a virtual world. Virtual human workers are then needed to test the design to ensure that the proposed tasks can be performed safely and efficiently. These digital human models (DHM) represent variation in human size, shape, and movement. The hand trajectory model described here was developed as part of an effort to improve the realism of the movement simulation

†*Address for correspondence:* Julian Faraway, Department of Mathematical Sciences, University of Bath, Bath BA2 7AY, United Kingdom
E-mail: jjf23@bath.ac.uk

capability in DHM. 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 the variance that can be expected.

Models for hand trajectories are useful in other areas. In neuroscience and motor control, researchers have long been interested in how and why we move in a particular way. Various theories have been proposed for predicting hand trajectories motivated by optimizing some characteristic. Flash and Hogan (1985) proposed minimizing jerk, which is the rate of change of acceleration, as the organizing principle for multi-joint arm motions. This results in a straight line hand trajectory between the start and finish with a symmetric bell-shaped velocity. Although this model has been confirmed in limited and controlled circumstances, hand trajectories are generally curved for larger scale motions. Uno et al. (1989) propose minimum torque change as the criterion. Various other optimization criteria have been proposed in the literature. However, the methods we describe in this article are purely empirical. We are not advocating any organizing principle for hand trajectories — we just wish to model how potential predictors affect the trajectory. Hand trajectories are sometimes considered as the inputs rather than outputs of models in neurosciences – see for example, Paninski et al. (2004). There is also considerable interest in explaining why hand trajectories should vary when the same task is repeated by the same person. See Harris and Wolpert (1998), Todorov and Jordan (2002) and Osborne et al. (2005). This work suggests that the residual variation in our models is not solely attributable to unmeasured variables, but is naturally occurring.

Hand trajectory modeling has also attracted interest in the computer science literature. Simulating human motion for CGI effects in movies and video games has motivated the development of a wide range of techniques. In some cases, the trajectory is manually specified by an artist in a subjective manner. While this is adequate for entertainment applications, it is time-consuming to do well. Perhaps the most common approach in practice is to take a single observed motion and modify it to meet some new constraint. For example, suppose we collect data on how a person takes a book from a shelf. We might wish to modify this to represent a somewhat taller person taking the book from a somewhat lower shelf. Gleicher (2001) and Park et al. (2004) provide examples of this. Optimization based approaches may also be found in Marler et al. (2005) and Zacher and Bubb (2005).

An interest in hand trajectories may be found in other areas such as Psychology — see Breteler et al. (1998) and Jackson and Husain (1997). In these and other papers such as Blackmon et al. (1997), certain characteristics of the trajectory such as maximum velocity are analyzed rather than the complete curve.

Although our particular application of interest is in hand trajectories, the methods described here could be used for modeling trajectories formed by other objects. Of course, physical law is most appropriate for predicting the motion of ballistic objects, but other objects follow trajectories that are not so easily described. See for example, Moore (1988) and Small and Samson (1983).

The trajectory data we will model is described in Section 2. This is used to motivate the statistical approach for modeling 3D 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 used to model the hand trajectories was collected at the Human Motion Simulation Laboratory (HuMoSim) at the University of Michigan. One set of experiments concentrated on the motions of standing people performing reaches to 30 different shelves. The shelves are located on 3 towers (in front of the subject, to the right side and 45° to the right), at 5 levels ranging from eye to

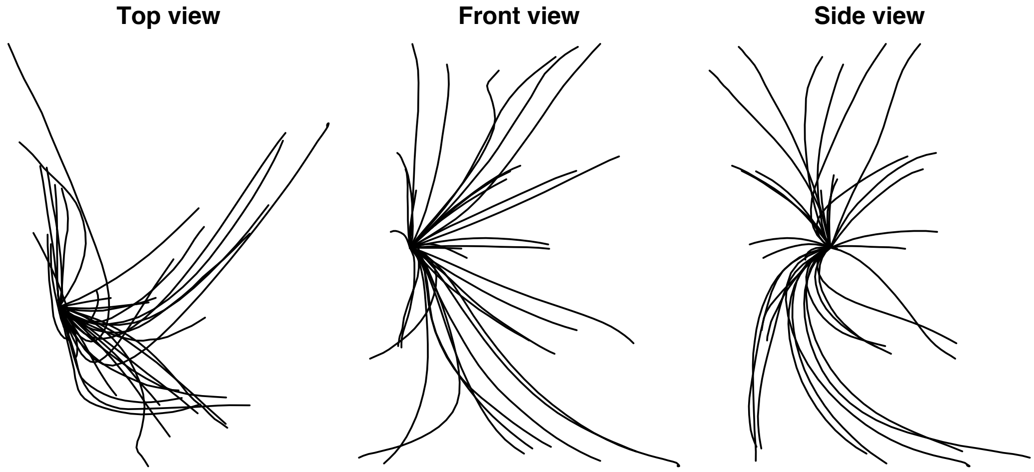

Fig. 1. Three views of reaches by one subject to 30 targets delivering a vertically-oriented cylinder. 6 reaches are replicated for 36 curves in all. All curves have been translated to start from the same location.

ankle level and 2 depths (near and far). The subjects were asked to move a box with both hands, a vertical cylinder with the right hand and a horizontal cylinder with the right hand. The hand motion began in front of a person on a small home shelf with the subject delivering the object to the target shelf.

The subjects were selected to provide a means to assess the effects of anthropometry, gender, and age on the motions. The subjects ranged from very short to very tall and from 20 to 70 years of age. Ten were male and ten were female. All were right handed. A total of 2031 motions were performed by the group of 20 subjects. Some motions were replicated. Some motions were lost due to data collection errors. The data discussed here is a subset of a larger sequence of experiments.

The motion capture was achieved with an optical reflective marker system (Qualysis MacReflex) recording at 25Hz. Our particular interest is in the movement of a marker attached to the right wrist, although 21 locations on the body were tracked in total. We choose to model the motion of the right wrist rather than hand because the hand motion will contain some rotation mixed in with the translation. It is easier to model the wrist first and then attach the rotating hand later — see Choe (2006).

All the reaches of one subject performing a specific task, namely delivering a vertically-oriented object from a home position in front of the subject to one of several nearby shelves, are shown in Figure 1. We observe points along the trajectories every 0.04 second. We have joined these points. The smoothness of these trajectories reflects the relative lack of measurement error and the inherent smoothness of human motion. We see a wide range of shapes in the trajectories depending on the target location. The trajectories are clearly not straight. We are interested in how the shape depends on the location of the target.

The reaches of all subjects to one location while delivering a vertically-oriented cylinder are shown in Figure 2. Here we see that the curves have approximately the same shape, but there is also substantial variation. Can this variation be attributed to the varying characteristics of the subjects? How much natural variation is to be expected? We aim to provide answers to these questions.

Although the reaches are to the same nominal target (shelf), the subjects were not required to be particularly precise in their placement of the cylinder. There is also some variation in the grips

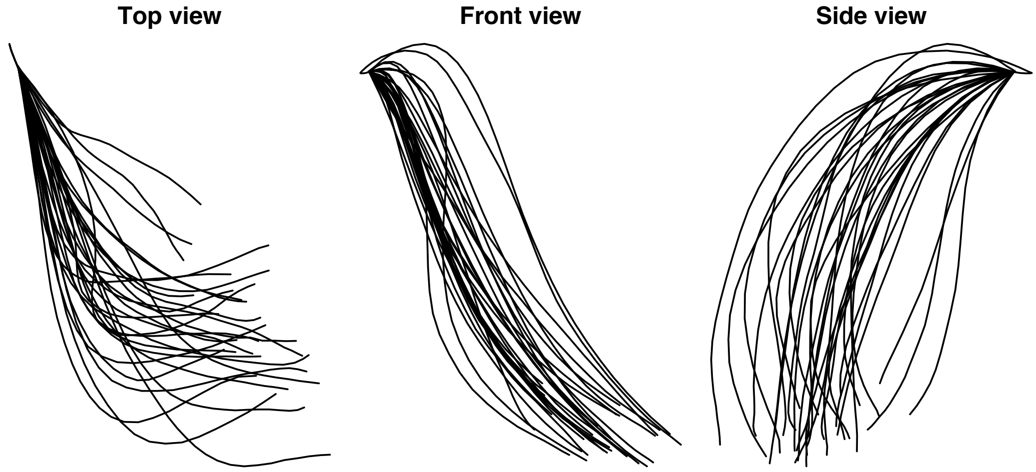

Fig. 2. Reaches by 20 subjects to the same target delivering a vertically-oriented cylinder. The target is at ankle-height to the right and somewhat to the rear of the subject. All curves have been translated to start from the same location.

used. Furthermore, the start and end of a reach is not easily defined as subjects are never completely stationary and may perform small adjustments of grip and position just before and after the main portion of the reach. These adjustment motions have been removed from the data. All this means that the curves do not all terminate at exactly the same point. There has long been an interest in the accuracy and timing of motions in the motor control literature — see for example, the well known Fitt's law (Fitts (1954)). However, our experiments were not designed to investigate such matters. Our interest here lies in the shapes of the trajectories and not the final location. For the purposes of our analysis, we shall condition on the final location being known. Variation in the final location is interesting, of course, but can be studied using simpler statistical methods.

The type of models 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 some of 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 specify the starting and ending location of the hand. The model must return a predicted hand trajectory given these inputs.

The 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.

Stable extrapolation is another 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 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.

Users will also wish to explore how design changes will affect virtual humans of different types.

This can be achieved by modifying the inputs. Users may also legitimately want to manually modify the output trajectory. For example, they may wish to avoid an intermediate obstacle. Therefore, it is helpful if the functional form of the predicted trajectory is of a type that allows it to be modified, whether programmatically or by manual user input, while preserving the underlying structure.

3. Methods

3.1. Properties of Bézier curves

Bézier curves are widely used in graphics, but have attracted little attention in statistics. Hence, we present brief introduction here. Kim et al. (1999) and subsequent articles consider the use of Bézier curves in smoothing, but this is unrelated to the work presented here.

The Bernstein polynomials of degree d for $i = 0, \ldots, d$ 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. The endpoints of the curve and the first and last control points coincide: $P_0 = C(0)$ and $P_d = C(1)$. Figure 3 shows the construction for $d = 3$ in two dimensions. The shape of the curve is determined by the interior control points: P_1 and P_2

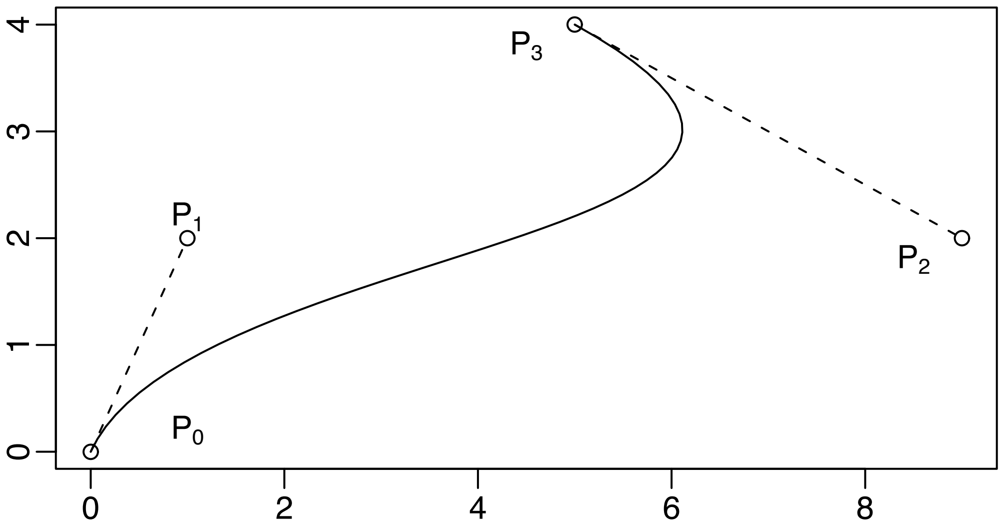

Fig. 3. Cubic $(d = 3)$ Bézier curve construction in two dimensions. The four control points P_i are shown. The endpoints of the curve coincide with the first and last control points. The first and last line segments between the control points are tangential to the curve at the endpoints.

in Figure 3. 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). Bézier curves in three dimensions, as needed for hand trajectories, use three dimensional control points.

Bézier curves have a number of interesting properties, some of which are particularly relevant for the application considered here.

(a) The curve is tangent to the line segments P_0P_1 and $P_{d-1}P_d$ at the endpoints. In 2D and higher, one can visualize the length of these line segments as controlling the approach to the endpoints. The longer the line segments are, the closer the curve follows the tangent on moving away from the endpoints. In the context of hand motion, the direction and length of these line segments describe the nature of the initial and final motion of the hand near the endpoints.
(b) The Bézier curve lies within the convex hull of the control points. In the context of hand motion modeling, this provides a simple way of bounding the trajectory and thus avoiding extreme extrapolations.
(c) The Bézier representation is affinely invariant meaning that we do not need to be concerned about dependence on the choice of coordinate system.
(d) Bézier curves are a special case of B-splines where the first $d+1$ knots are at zero and the second $d+1$ knots are at one, with no internal knots. In contrast to B-splines, the Bézier basis functions, that is the Bernstein polynomials, are supported on the whole interval [0,1]. So Bézier curves lack the local fitting and numerical stability properties valued in many B-spline applications in statistics. However, in our application, we are only interested in small d, usually $d = 3$, and the geometric interpretation of the Bézier control points aids the modeling in contrast to the more opaque combination of B-spline knots and control points (usually called coefficients by statisticians).

3.2. *Fitting Bézier Curves*

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

The question arises of which value of t along the curve $C(t)$ should correspond to an observed point Z_i. In some applications, it might be reasonable to use evenly spaced t, but for hand trajectories, we know that the hand does not move at constant speed. We need to allow for non-constant speeds and so we define a sequence of relative times, $0 \le s_1 \le \cdots \le s_{n-1} \le 1$, representing the relative progress of the hand along the trajectory. We define the sum of squares for the fit as:

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

We want to choose the control points of C to minimize SS. Evenly spaced times, $s_i = i/n$, might be best in some cases, but for our application, we used an arc length-based spacing of:

$$s_i = \sum_{j=1}^{i} a_j / \sum_{j=1}^{n} a_j \qquad \text{where} \qquad a_i = |Z_i - Z_{i-1}|$$

This allows us to scale out some of the variation in trajectories due to variable relative speed. For other applications, an appropriate choice will need to be made.

The control points $\hat{P}$ that minimize SS are given by:

$$B^T B\hat{P} = B^T Y$$

where B is an $(n-1)\times(d-1)$ matrix whose $(i,j)^{th}$ element is $B_j^d(s_i)$ and Y is a $(n-1)\times q$ matrix whose i^{th} row is given by:

$$Y_i = Z_i^T - B_0^d(s_i)Z_0^T - B_d^d(s_i)Z_n^T$$

An example of a cubic Bézier fit to some trajectory data is shown in Figure 4. The position of the interior control points relates to the initial and final directions of the data.

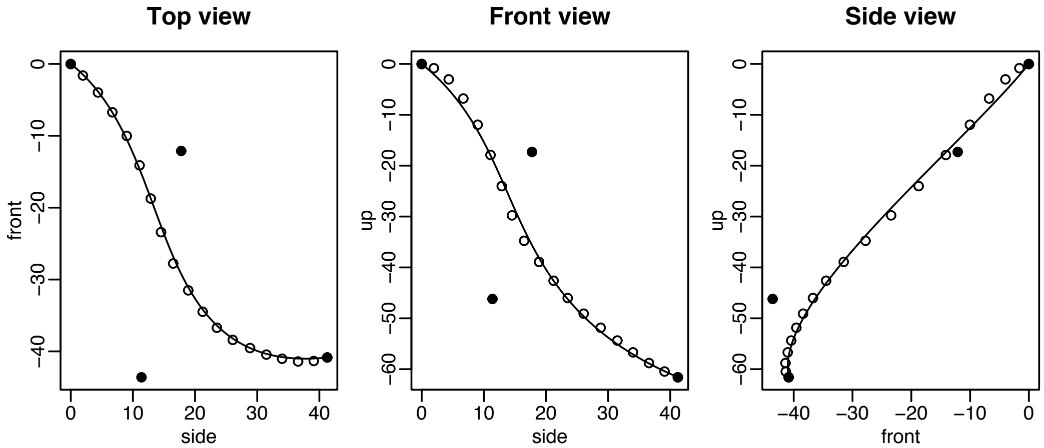

Fig. 4. Cubic Bézier fit to some 3D trajectory data. The four control points are shown as solid dots. The curve starts from the origin.

In practice, we need to select the degree of the Bézier curves that we shall use. If our sole concern was to fit a single curve, we could readily adapt the techniques used for fitting curves using B-splines using methods such as crossvalidation to choose the degree. However, we need to fit a large number of curves all with the same degree. Furthermore, there are two compelling reasons to choose as small a degree as possible. Firstly, unlike B-splines, the Bézier basis functions have a global support, so using a high degree will lead to problems with numerical stability and other such problems associated with higher order polynomials. Secondly, we plan to use the coefficients of the Bézier fit as the responses for next stage of the modeling. Interpretation of the model will be easier if there are a smaller number of coefficients.

For trajectories that tend to contain a larger number of twists and turns, a Bézier curve fit of low degree will be inadequate. Rather than increase the degree, it will be better to use either piecewise Bézier curves of low degree or B-splines. Essentially the same modeling process can then be pursued except that interpreting the fit will become more difficult.

3.3. Modeling

Let the vector of predictors associated with each trajectory be $(x_1,\ldots,x_p)$. We aim to find models of the general form:

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

P_0 and P_d are the known endpoints so modeling the position of the other P_i relative to these points is likely to be helpful. We will consider two parameterizations.

In the endpoint parameterization, we model the location of the first internal control point relative to the starting point: $Y_1 = P_1 - P_0$ and the last internal control point relative the ending point: $Y_{d-1} = P_{d-1} - P_d$. Other interior control points could be parameterized relative to either endpoint.

In the axis parameterization, we define *axis points* $A_1, \ldots, A_{d-1}$ that are evenly spaced along the axis joining P_0 and P_d. We may then model $Y_i = P_i - A_i$. This makes it easier to understand the model in terms of divergence from a straight line reach where $Y_i = \mathbf{0} \quad \forall i$. It also removes some of the dependence of the shape of the trajectory on the relative locations of the initial and final positions. The choice of parameterizations is depicted in Figure 5 for a cubic Bézier curve.

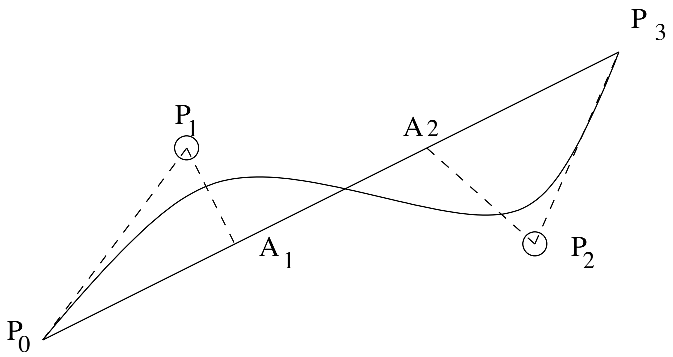

Fig. 5. The interior control points P_1 and P_2 can be measured relative to the endpoints P_0 and P_1 or relative to axis points equispaced along the axis, A_1 and A_2.

We can then model this multivariate response Y in q dimensions using standard multivariate linear modeling techniques (Johnson and Wichern (2002)). Alternatively, for specifically 3D trajectories, we could use spherical coordinates. The radial component can then be modeled using ordinary linear modeling methods while the spherical component can modeled using techniques for spherical response regression techniques (Mardia and Jupp (2000)). Using these known methods, we could, for example, investigate the statistical significance of potential covariates and generate confidence bands for predicted trajectories.

When some form of data dependent spacing s_i is used, a model may need to be constructed for these spacings in order to construct a predicted trajectory for given inputs. For example, in the hand trajectory example, the hand accelerates at the beginning of the motion and decelerates towards the end of the motion. We proposed an arc length-based spacing above. We found that the spacing could be adequately modeled using a Beta distribution. If a fixed spacing were used, such a model would be unnecessary.

The residuals from the fit, $\hat{\varepsilon} = Y - \hat{Y}$ can be analyzed to investigate the pattern of variation in the trajectories. Again, standard techniques from multivariate analysis, such as principal components analysis, are available and no special new methods are necessary.

4. Application

We aim here only to illustrate several aspects of the statistical analysis for this particular application. The first stage of the modeling requires us to fit Bézier curves to the trajectories. We must choose the degree of these curves. We used the fitting procedure described in the previous section for all degrees up to four. The results are shown in Table 1. We used an arc-length spacing, as even

Table 1. Approximation Error is measured by the maximum distance of the fit Bézier curve to the observed points of the trajectory. The median of these maxima is reported in the table.

Approximating Curve	Median Maximum Error
Linear	10.5cm
Quadratic Bezier	3.6cm
Cubic Bezier	1.7cm
Quartic Bezier	1.0cm

spacing gives much poorer results. The curve with degree one is simply a straight line between the endpoints.

We will need a general way of assessing the fit for this method and those to follow. We shall use the maximum error, $\max_i |Z_i - \hat{C}(s_i)|$, to represent the difference between the observed and fitted trajectories. The maximum will occur near the middle of the trajectory because the endpoints must coincide. We will then use the median of these maxima to summarize the overall fit to the data. The distribution of the maximum is asymmetric with the first quartile typically being about 2/3 of the median and the third quartile being about 5/3 of the median. In a few observed cases, subjects follow quite unusual trajectories and so the maximum error is relatively large. It is difficult to say how much error is acceptable for practical use since the answer would depend on the application. However, bear in mind that the palm of the hand is roughly 10cm square so that errors less than this would mean that the predicted and observed hand locations would always overlap.

The linear fit has been used in some human modeling implementations due to its simplicity, but much better fits are possible using the Bézier approximation. We see that the error decreases rapidly with the degree of the fit. We decided to use the cubic fit. The cubic fit results in two internal control points which can be used to represent the initial and final directions of travel along the trajectory. The quartic fit would introduce an additional control point between these two whose location is somewhat more difficult to interpret. Of course, the quartic fit is clearly better than the cubic fit but an additional control point would then need to be modeled. We have chosen to favour simplicity and stability over fit in this case.

We also need to model the relative speed profile for these data because we are using the arc length parameterization. We use a Beta distribution function with parameters α and β. We used maximum likelihood to estimate these parameters for each curve, by treating the observed values as a sample from a Beta.

Across the whole dataset, we found the median $\hat{\alpha} = 1.63$ and $\hat{\beta} = 1.75$. The median location of maximum relative speed occured at $t = 0.48$ — slightly less than half way through the motion. Several previous authors have proposed a symmetric relative speed distribution, see Flash and Hogan (1985) for example. Although the difference is statistically significant here, we recognize that the result is sensitive to how the beginning and end of the motion are determined which would affect the location of the maximum.

Some error in prediction can attributed to the modeling of the relative speed. For each observation, we computed $(\hat{\alpha}, \hat{\beta})$ and used these values to fit a cubic Bézier curve to the observed data. The median maximum error was 3.2cm, worse than the 1.7cm when using the observed relative speed. The difference is due to the choice of the Beta density model. One might hope to do better than this, but we have been unable to find a simple relative speed model that has a superior fit.

Furthermore, in practice we must provide the (α, β) to construct a predicted trajectory. Using the median values of $(\hat{\alpha}, \hat{\beta})$ gives an error of 6.9cm. Using the Flash–Hogan relative speed model which has density $f(t) = 30t^2(1-t)^2$ gives 11.4cm. Note that the Flash–Hogan speed density is just a Beta density with $\alpha = 2$ and $\beta = 2$, which is somewhat different from the optimal values estimated from our data.

The fitted control points provide convenient measures to explore the distribution of trajectories. We can plot the trajectories themselves as seen in Figures 1 and 2, but if we add more than a few curves, the plots become too tangled to easily interpret. In contrast, we can plot the control points more conveniently. As an example, consider Figure 6. We take the view from above and compute the direction of the first control point relative to the starting point. A rose diagram showing the distribution of these directions is shown in the first plot of Figure 6. We see that the subjects tend to move their hands backwards and to the right of the starting point. However, these directions must depend on where the object is eventually to be placed. In the second plot, we see the directions of the targets. The three towers of targets are clearly visible. The third plot shows the difference in the angle between the initial direction of the hand and the direction of the target. We can see that the subjects tend to veer somewhat to the right of a straight line reach initially. Further plots with different views and considering the final approach to the target gain more insight into the distribution of the trajectories.

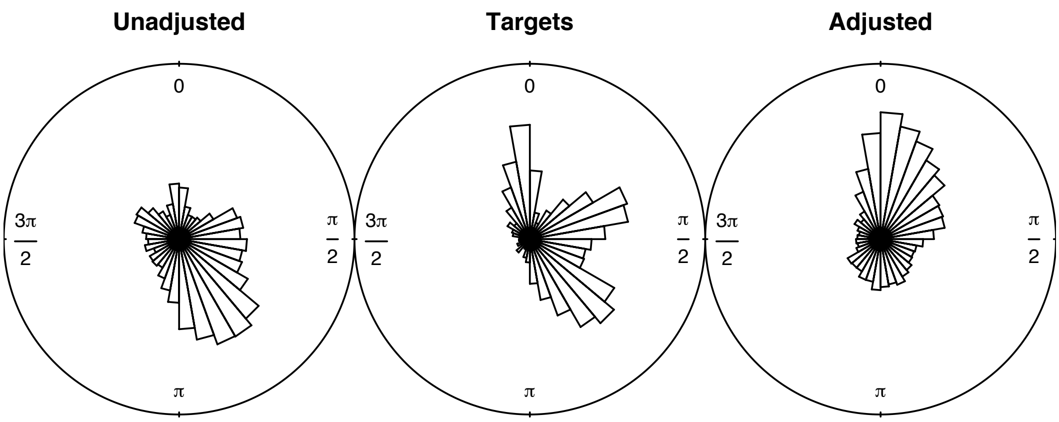

Fig. 6. Rose diagrams showing the distribution of the initial direction of the trajectory. All views are from above and zero degrees is the forward direction in the first two plots. The first plot shows the direction of the first control point relative to the starting point. The second plot shows the direction from the starting to the ending point, where the three groups of targets may be seen. The third plot shows the angle between these two directions and thus represents the deviation from a straight line trajectory.

Due to the natural variability in human motion, we cannot expect perfect predictions. 352 motions in the dataset were made twice i.e. the same person performed the same task two times. We can compute the maximum distance between each pair of replicated trajectories over relative time. The median of these maxima is 6.7cm. We cannot expect any real trajectory prediction method to better this. Furthermore, we should expect some variation from individual to individual that will add to this lowest achievable error. To assess this error, we group the data by task, that is all reaches to the same target and the same object handled. Within each group we compute the median trajectory and then consider maximum deviation from this median trajectory for the individual reaches. The median of these maxima was 7.6cm. This is not a lower bound on the error because better predictive performance may be obtained if personal characteristics like height or sex provide some information about the trajectory. However, given that 7.6cm is little more than the within individual error of 6.7cm, there is not much scope for these measurable individual predictors to explain the total variation.

Now we turn to a consideration of how much task and individual factors can help predict the trajectory. The full Flash–Hogan straight trajectory model gives an error of 12.6cm on this dataset. Since this model is well-established and very easy to implement, we need to do better than this.

We must some choices relating to the parameterization and type of coordinates. The endpoint parameterization is likely to be most appropriate if the beginning and ending parts of the trajectory are similar no matter where the target is. The axis parameterization will be better if the trajectory is more dependent on the target. We will investigate both parameterizations. We choose to model these as Cartesian coordinates rather than spherical coordinates. Cartesian coordinates are natural given the particular meaning of "up", "forward" and "lateral" for the user. Furthermore, spherical coordinates require modeling angular responses which introduce further difficulties.

There is a question of what is an appropriate null model for this data — that is a model that does not use any predictor information. A null model often corresponds to the notion of an average response over the complete dataset. With this in mind, we can average over control points in their endpoint or axis parameterization and use these averages to construct predicted trajectories. Note that the predicted trajectories will be different for different targets, but the manner in which they curve around the axis will be the same depending on the choice of parameterization. For the endpoint parameterization, the maximum median error was 11.2cm, but for the axis parameterization, the error was 8.2cm. It is possible to mix the parameterizations, that is use the endpoint for the first control point and the axis for the second or vice versa. The error for these two choices was 12.3cm and 11.9cm respectively.

The axis method is clearly superior and approaches the between subject error of 7.6cm. A further improvement is possible by noticing that longer reaches are likely to have greater curvature. If we scale by d_a, the distance between P_0 and P_1, and then average the responses $Y_1 = (P_1 - A_1)/d_a$ and $Y_2 = (P_2 - A_2)/d_a$, the prediction error is reduced to 8.0cm. The distance scaling will be important when predicting short reaches as it ensures that these will be close to linear, as is observed in practice.

Even better performance is obtained by also using some time scaling. We set t as the time of the reach and average over $Y_1 = (P_1 - A_1)/\sqrt{td_a}$ and $Y_2 = (P_2 - A_2)/\sqrt{td_a}$. The error is further reduced to 7.6cm. However, this would require that the software calculate the appropriate duration for the reach or that the user provide a value. In some human modeling applications, standardized movement time systems and other time prediction algorithms are available, but even in these situations the user would need to provide a value for complex movements.

There are thirty targets in the experiment and three objects moved. If we average just over the targets, the error for the endpoint parameterization is 8.2cm and for the axis parameterization, it is 8.3cm. If we average over targets and objects, the error is reduced to 7.7cm. However, note that these models could only be used for prediction to the targets that were used in our experiment and could not be extrapolated to new targets. To do this, we would need, for example, to construct a function of the target coordinates to use as a predictor. This would inevitably be more complex and perform worse. We have been unable to find such a model that is any better than the distance scaled axis null model described above.

To look for systematic subject effects, we averaged the distance-scaled control points over each subject using the axis parameterization. The error using this model was 8.1cm, which is slightly higher than when we just averaged over the whole data. Especially considering that this model would only be useful for predicting trajectories for the subjects used in our experiment, we can see

Table 2. Maximum median error in cm. for models considered. For full description and discussion, see the text.

Model Type	Model Form	Error
Natural Variation	Within Subject	6.7
	Between Subject	7.6
Speed profile	Observed (α, β)	3.2
	Median (α, β)	6.9
	Flash–Hogan speed	11.4
Null	Flash–Hogan	12.6
	Endpoint parameterization	11.2
	Axis parameterization	8.2
	Axis with distance scaling	8.0
	Axis with distance/time scaling	7.6
Predictor	Target mean with endpoint	8.2
	Target mean with axis	8.3
	Target/Object mean with axis	7.6
	Subject mean with axis and distance scaling	8.1

there is little evidence for substantial systematic subject effects.

A summary table of the maximum median errors for the chosen models is shown in Table 2

Our chosen model is the null axis model with distance scaling. This model performs competitively and yet is simple to specify and implement. The observed means for the distance-scaled control points offset from the axis points are (-0.025, -0.148, 0.077) and (-0.015, -0.139, -0.002) respectively. For the median reach distance of 56.7cm, this gives offsets of (-1.4, -8.4, 4.3) cm and (-0.8, -7.9, -0.1) cm. We can see that the first control point is somewhat to the rear and above the corresponding axis point as the subject initially moves the object back (i.e., toward him or herself) and up (i.e., lifts the object) relative to a straight line path. This is consistent with the pattern seen in the third plot of Figure 6. The second control point is mostly to the rear of the axis point as the subject prepares to place the object on the shelf. The prediction error in this model tends to be larger for curves with larger Bézier approximation error and/or greater curvature.

Predicted reaches to all 30 targets are shown in Figure 7. This should be compared to Figure 1. The predicted trajectories have shapes comparable to the observed trajectories although there is a small amount of additional curvature in the observed data.

Next we turn to analysis of variation in the trajectories. We took the preferred model, that is the null axis parametrization with distance scaling, and computed the residuals with respect to the control points resulting in an 2031x6 matrix. A principal components analysis of this matrix revealed that variation could not reasonably be reduced to a lower dimension as even the sixth principal component was relatively large. However, there was some mild correlation and the residual variation is clearly not spherical. We were unable to find any strong relationship between the residuals and the predictors suggesting that no significant improvement to this model is possible within this general modeling approach.

We can use the covariance structure of the residuals to generate simulated trajectory that vary around the mean prediction. To illustrate this, we generated 20 observations from a multivariate normal to construct the simulated reaches shown in Figure 8. This can be compared to Figure 2 where 20 actual trajectories are observed, albeit with somewhat varying endpoints. Nonetheless, the similarity in the variation structure may be observed.

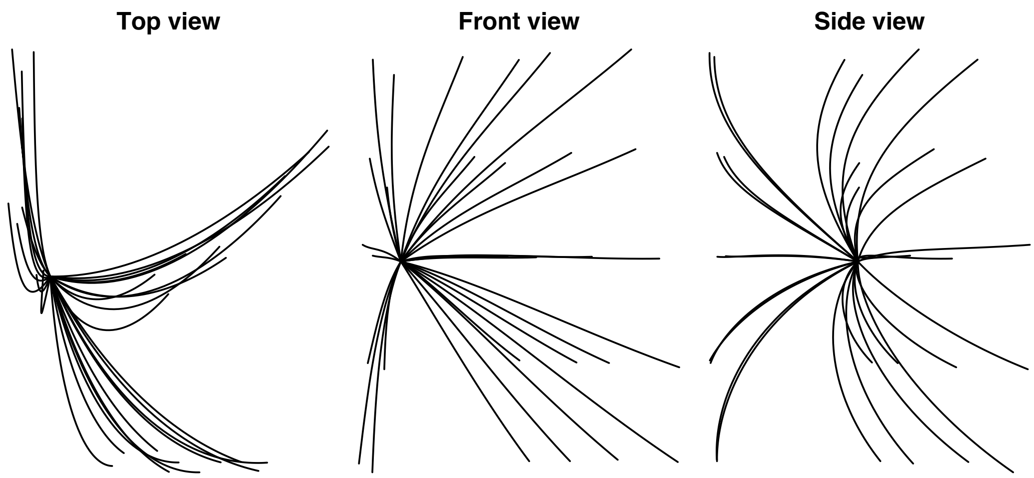

Fig. 7. Predicted reaches to 30 targets using the null axis model with distance scaling. Compare this to the observed data in Figure 1.

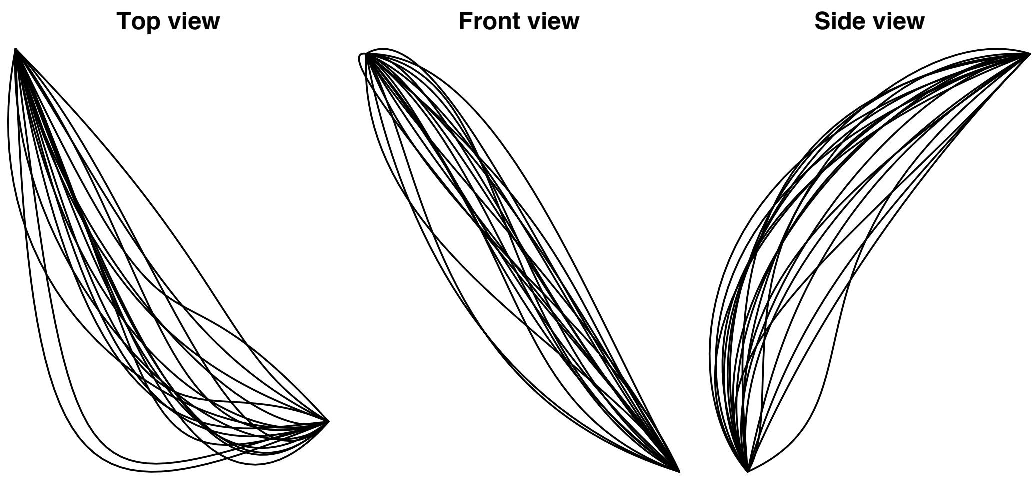

Fig. 8. Simulated reaches to the same target as in Figure 2. In this case, both the start and end of the trajectories have been fixed.

5. Discussion

We have demonstrated a modeling approach for hand trajectories that has good predictive performance and yet retains interpretable elements. The model is simple to specify and implement. The performance of the model approaches the lower bound determined by natural human variation and is significantly better than the Flash–Hogan linear trajectory model.

The convex hull property of Bézier curves means that we can control extrapolations so that extremely divergent trajectories will not be predicted. This property is potentially valuable in the human motion simulation domain to provide either automated or manual obstacle avoidance. The geometric interpretation of the Bézier control points may simplify the design of user interfaces for situations in which the model predictions must be adjusted by the user to conform to task constraints. Indeed, Bézier curves are already widely used in computer graphics programs.

The data from this relatively simple set of hand motions deviated substantially from the linear trajectories produced by other models. Our interpretation of the curvature in the trajectories is based on an assumption that the kinematic constraints that operate at the beginning and end of the motion may have effects that are more local than global. In this regard, the current analysis is limited by the single task represented in the data (transferring a cylinder between shelves). Tasks that require greater manual precision or complex hand motions may produce hand trajectories that are not fit well by the Bezier model introduced here. However, Wang (2006) previously found that the Bézier formulation provided a good fit to data from an experiment with a wider range of hand target orientations, increasing confidence that this model may have general utility for human motion simulation.

Our current approach to fitting the interior control points that modifies the motion direction at the ends of the trajectory is purely empirical and hence our current model may not perform well for other tasks. However, we expect that specific information about the kinematic constraints that are active at the endpoints will allow greater generality in the model and better extrapolation. For example, including information about the required approach direction for a target (say, inserting a part from the right or from the bottom) is likely to improve the ability to predict the hand trajectory.

Although we have applied the Bézier formulation for modeling hand trajectories, it may also be applicable to other trajectories that are generally smooth and have curvatures affected by kinematic constraints at the end points. In the human modeling domain, we are also exploring the use of Bézier curves to represent foot and pelvis trajectories. The kinematic and dynamic constraints on human motion tend to produce smooth trajectories for volitional motions, encouraging the selection of relatively simple and readily implemented models. However, we have not explored whether the use of the Bézier curve to represent hand trajectories can be motivated by biomechanical or motor-control considerations.

Our model operates in a coordinate system defined by global vertical and a "forward" direction defined by the initial posture and position of the subject. The definition of forward for these data was conveniently defined by the test setup, but for other tasks, particularly those with concurrent pelvis motion or stepping, the appropriate coordinate system for the modeling is not as apparent. Some of the curvature that is observed in these data is probably due to the kinematic limitations and comfortable patterns of motion of the upper extremity. This suggests that the trajectories of reaches with the left hand to similar targets mirrored to the left side of the body might be modeled by reversing the lateral-coordinate signs in the model, but of course data from left-handed reaches will be necessary to assess this.

We have chosen to separate the modeling of the spatial trajectory from the temporal progression along the trajectory. For the current application, this approach is justified by our understanding of what constitutes "error" in trajectory prediction. For digital human modeling applications, errors in spatial location are much more important than temporal errors. Moreover, the requirement that the hand be stationary at the beginning and end of the motion, coupled with velocity smoothness imposed by the biomechanics of the upper extremity, means that the range of possible velocity profiles is substantially constrained and relatively easily predicted. However, in other domains, a positional error along a trajectory at a particular time may be of greater importance.

Acknowledgements

We acknowledge the vital support of Don Chaffin who founded the HuMoSim laboratory at the Center for Ergonomics at the University of Michigan and provided many years of support to the research presented in this paper. We are also grateful for the support of our colleagues at HuMoSim.

References

Blackmon, T., M. Cavusoglu, L. Fuji, and L. Stark (1997). Human hand trajectory analysis in point-and-direct telerobotics. *8th International Conference on Advanced Robotics.*, 927–32.

Breteler, M., R. Meulenbroek, and S. Gielen (1998). Geometric features of workspace and joint-space paths of 3d reaching movements. *Acta Psychologica 100*, 37–53.

Choe, S. (2006). *Statistical analysis of orientation trajectories via quaternions with application to human motion.* Ph. D. thesis, Department of Statistics, University of Michigan.

Fitts, P. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology 47*(6), 381–391.

Flash, T. and N. Hogan (1985). The coordination of arm movements: an experimentally confirmed mathematical model. *Journal of Neuroscience 5*, 1688–1703.

Gleicher, M. (2001). Comparing constraint-based motion editing methods. *Graphical Models 63*, 107–134.

Harris, C. and D. Wolpert (1998). Signal-dependent noise determines motor planning. *Nature Neuroscience 394*, 780–784.

Jackson, S. and M. Husain (1997). Visual control of hand action. *Trends in Cognitive Sciences 1*, 310–317.

Johnson, R. and D. Wichern (2002). *Applied Mulivariate Statistical Analysis* (5th ed.). New Jersey: Prentice Hall.

Kim, C., W. Kim, C. Hong, B.-U. Park, and M. Jeong (1999). Smoothing techniques via the bezier curve. *Communications in Statistics — Theory and Methods 28*, 1577–1596.

Mardia, K. and P. Jupp (2000). *Directional Statistics*. Chichester: Wiley.

Marler, T., S. Rahmatalla, M. Shanahan, and K. Abdel-Malek (2005). A new discomfort function for optimization-based posture prediction. *SAE Technical Paper 2005-01-2680*.

Moore, M. (1988). Predicting iceberg trajectories. *Chance 1*(2), 30–36.

Osborne, L., S. Lisberger, and W. Bialek (2005). A sensory source for motor variation. *Nature 437*, 412–416.

Paninski, L., S. Shoham, M. Fellows, N. Hatsopoulos, and J. Donoghue (2004). Superlinear population encoding of dynamic hand trajectory in primary motor cortex. *The Journal of Neuroscience 24*(39), 8551–8561.

Park, W., D. Chaffin, and B. Martin (2004). Toward memory-based human motion simulation: Development and validation of a motion modification algorithm. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans 34*(3), 376–386.

Prautzsch, H., W. Boehm, and M. Paluszny (2002). *Bezier and B-Spline Techniques*. New York: Springer.

Small, M. and P. Samson (1983). Stochastic simulation of atmospheric trajectories. *Journal of Climate and Applied Meteorology 22*, 266–277.

Todorov, E. and M. Jordan (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience 5*(11), 1226–1235.

Uno, Y., M. Kawato, and R. Suzuki (1989). Formation and control of optimal trajectory in human multijoint arm movement minimum torque-change model. *Biological Cybernetics 61*, 89–101.

Wang, J. (2006). *Statistical modeling for 3-D trajectories*. Ph. D. thesis, Department of Statistics, University of Michigan, Ann Arbor.

Zacher, I. and H. Bubb (2005). Strength-based discomfort model of posture and movement. *SAE Technical Paper 2004-01-2139*.