

Modelling three-dimensional trajectories by using Bézier curves with application to hand motion

Julian J. Faraway,
University of Bath, UK

Matthew P. Reed
University of Michigan, Ann Arbor, USA

and Jing Wang
Bank of America, Atlanta, USA

[Received September 2006. Revised April 2007]

Summary. A modelling approach for three-dimensional 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 by 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 interpretability is developed. The use of hand trajectory models in the digital human models for virtual manufacturing applications is discussed.

Keywords: Digital human model; Functional data analysis; Splines; Virtual manufacturing

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 paper is to describe a modelling 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 paper comes from ergonomics. Ergonomics applies models of human ability 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 to 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 tasks proposed can be performed safely and efficiently. These digital

Address for correspondence: Julian Faraway, Department of Mathematical Sciences, University of Bath, Claverton Down, Bath, BA2 7AY, UK.
E-mail: jjf23@bath.ac.uk

human models represent variation in human size, shape and movement. The hand trajectory model that is described here was developed as part of an effort to improve the realism of the movement simulation ability in digital human models. 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 multijoint 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) proposed minimum torque change as the criterion. Various other optimization criteria have been proposed in the literature. However, the methods that we describe in this paper 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 modelling has also attracted interest in the computer science literature. Simulating human motion for computer-generated imagery effects in films 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. Although 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 to modify it to meet some new constraint. For example, suppose that 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) provided 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 Bretelet *et al.* (1998) and Jackson and Husain (1997). In these and other references such as Blackmon *et al.* (1997), certain characteristics of the trajectory such as maximum velocity were analysed rather than the complete curve.

Although our particular application of interest is in hand trajectories, the methods that are described here could be used for modelling trajectories that are 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 that we shall model are described in Section 2. These are used to motivate the statistical approach for modelling three-dimensional trajectories in Section 3. We demonstrate the application to modelling hand trajectories in Section 4 and close with a discussion in Section 5.

The data that are analysed in the paper can be obtained from

<http://www.blackwellpublishing.com/rss>

2. Data and motivation

The data that are used to model the hand trajectories were collected at the Human Motion Simulation Laboratory 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 three towers (in front of the subject, to the right-hand side and 45° to the right), at five levels ranging from eye to ankle level and two 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. 10 were male and 10 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 because of data collection errors. The data that are discussed here are a subset of a larger sequence of experiments.

The motion capture was achieved with an optical reflective marker system (Qualysis Mac-Reflex) recording at 25 Hz. Our particular interest is in the movement of a marker that was 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 to 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 Fig. 1. We observe points along the trajectories every 0.04 s. 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.




Fig. 1. Three views of reaches by one subject to 30 targets delivering a vertically oriented cylinder (six reaches are replicated for 36 curves in all; all curves have been translated to start from the same location): (a) top view; (b) front view; (c) side view




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): (a) top view; (b) front view; (c) side view

The reaches of all subjects to one location while delivering a vertically oriented cylinder are shown in Fig. 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 (a shelf), the subjects were not required to be particularly precise in their placement of the cylinder. There is also some variation in the grips that were used. Furthermore, the start and end of a reach are 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 by using simpler statistical methods.

The type of models that are 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 digital human model software that is 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 that are 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, 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 modify the output trajectory manually. 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 a brief introduction here. Kim *et al.* (1999) and subsequent references have considered the use of Bézier curves in smoothing, but this is unrelated to the work that is presented here.

The Bernstein polynomials of degree d for $i = 0, \dots, d$ 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 end points of the curve and the first and last control points coincide: $P_0 = C(0)$ and $P_d = C(1)$. Fig. 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 in Fig. 3. The geometric construction for Bézier curves using the De Casteljau algorithm




Fig. 3. Cubic ($d = 3$) Bézier curve construction in two dimensions: the four control points P_i are shown; the end points 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 end points

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 several interesting properties, some of which are particularly relevant for the application that is considered here.

- (a) The curve is tangent to the line segments P_0P_1 and $P_{d-1}P_d$ at the end points. In two dimensions and higher, one can visualize the length of these line segments as controlling the approach to the end points. The longer the line segments are, the closer the curve follows the tangent on moving away from the end points. 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 end points.
- (b) The Bézier curve lies within the convex hull of the control points. In the context of hand motion modelling, 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 co-ordinate system.
- (d) Bézier curves are a special case of B -splines where the first $d + 1$ knots are at 0 and the second $d + 1$ knots are at 1, with no internal knots. In contrast with B -splines, the Bézier basis functions, i.e. the Bernstein polynomials, are supported on the whole interval $[0, 1]$. So Bézier curves lack the local fitting and numerical stability properties that are valued in many B -spline applications in statistics. However, in our application, we are interested only in small d , usually $d = 3$, and the geometric interpretation of the Bézier control points aids the modelling, in contrast with 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, \dots, Z_n . We wish to fit a Bézier curve $C(t)$ of degree d for $t \in [0, 1]$ such that $C(0) = Z_0$ 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, $0 \leq s_1 \leq \dots \leq s_{n-1} \leq 1$, 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

$$s_i = \frac{\sum_{j=1}^i a_j}{\sum_{j=1}^n a_j}$$

where $a_i = |Z_i - Z_{i-1}|$. 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

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




Fig. 4. Cubic Bézier curve fit to some three-dimensional trajectory data (the curve starts from the origin; ●, control points): (a) top view; (b) front view; (c) side view

where B is an $(n - 1) \times (d - 1)$ matrix whose (i, j) th element is $B_j^d(s_i)$ and Y is an $(n - 1) \times q$ 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 Fig. 4. The position of the interior control points relates to the initial and final directions of the data.

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 that are used for fitting curves by using B -splines by using methods such as cross-validation 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 that are associated with higher order polynomials. Secondly, we plan to use the coefficients of the Bézier fit as the responses for the next stage of the modelling. 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 modelling process can then be pursued except that interpreting the fit will become more difficult.

3.3. Modelling

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

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

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

In the end point 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 to the




Fig. 5. Interior control points P_1 and P_2 can be measured relatively to the end points P_0 and P_3 or relatively to axis points that are equispaced along the axis, A_1 and A_2

ending point, $Y_{d-1} = P_{d-1} - P_d$. Other interior control points could be parameterized relatively to either end point.

In the axis parameterization, we define *axis points* A_1, \dots, 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 $Y_i = \mathbf{0}, \forall i$. 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 Fig. 5 for a cubic Bézier curve.

We can then model this multivariate response Y in q dimensions by using standard multivariate linear modelling techniques (Johnson and Wichern, 2002). Alternatively, for specifically three-dimensional trajectories, we could use spherical co-ordinates. The radial component can then be modelled by using ordinary linear modelling methods whereas the spherical component can be modelled by using techniques for spherical response regression (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 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 modelled by 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 analysed 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 modelling requires us to fit Bézier curves to the trajectories. We must choose the degree of these curves. We used the fitting procedure that was described in the previous section for all degrees up to 4. The results are shown in Table 1. We used an arc length spacing, as even spacing gives much poorer results. The curve with degree 1 is simply a straight line between the end points.

We shall 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 end points must coincide. We shall then use the median of these maxima to summarize the overall

Table 1. Approximation error measured by the maximum distance of the fitted Bézier curve to the observed points of the trajectory†

<i>Approximating curve</i>	<i>Median maximum error (cm)</i>
Linear	10.5
Quadratic Bézier	3.6
Cubic Bézier	1.7
Quartic Bézier	1.0

†The median of these maxima is reported in the table.

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 10 cm square so errors that are less than this would mean that the predicted and observed hand locations would always overlap.

The linear fit has been used in some human modelling implementations because of its simplicity, but much better fits are possible by 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 modelled. 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 distribution.

Across the whole data set, we found the median $\hat{\alpha} = 1.63$ and $\hat{\beta} = 1.75$. The median location of maximum relative speed occurred at $t = 0.48$ —which is slightly less than half way through the motion. Several previous researchers 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 be attributed to the modelling 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 3.2 cm, which is worse than the 1.7 cm when using the observed relative speed. The difference is due to the choice of the beta density model. We 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 6.9 cm. Using the Flash–Hogan relative speed model which has density $f(t) = 30t^2(1 - t)^2$ gives 11.4 cm. 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 that were estimated from our data.




Fig. 6. Rose diagrams showing the distribution of the initial direction of the trajectory (all views are from above and 0° is the forward direction in (a) and (b)): (a) direction of the first control point relative to the starting point (unadjusted); (b) direction from the starting- to the end point, where the three groups of target may be seen; (c) angle between these two directions and thus representing the deviation from a straight line trajectory (adjusted)

The fitted control points provide convenient measures to explore the distribution of trajectories. We can plot the trajectories themselves as seen in Figs 1 and 2 but, if we add more than a few curves, the plots become too tangled to interpret easily. In contrast, we can plot the control points more conveniently. As an example, consider Fig. 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 Fig. 6(a). 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 Fig. 6(b), we see the directions of the targets. The three towers of targets are clearly visible. Fig. 6(c) 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.

Owing to the natural variability in human motion, we cannot expect perfect predictions. 352 motions in the data set 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 6.7 cm. 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, i.e. all reaches to the same target and the same object handled. Within each group we compute the median trajectory and then consider the maximum deviation from this median trajectory for the individual reaches. The median of these maxima was 7.6 cm. 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 7.6 cm is little more than the within-individual error of 6.7 cm, 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 to predict the trajectory. The full Flash–Hogan straight trajectory model gives an error of 12.6 cm on this data set. Since this model is well established and very easy to implement, we need to do better than this.

We must make some choices relating to the parameterization and type of co-ordinates. The end point 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 shall investigate both parameterizations. We choose to model these as Cartesian co-ordinates rather than spherical co-ordinates. Cartesian co-ordinates are natural given the particular meaning of 'up', 'forward' and 'lateral' for the user. Furthermore, spherical co-ordinates require modelling angular responses which introduce further difficulties.

There is a question of what is an appropriate null model for these data—i.e. a model that does not use any predictor information. A null model often corresponds to the notion of an average response over the complete data set. With this in mind, we can average over control points in their end point 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 end point parameterization, the maximum median error was 11.2 cm but, for the axis parameterization, the error was 8.2 cm. It is possible to mix the parameterizations, i.e. to use the end point for the first control point and the axis for the second or vice versa. The error for these two choices was 12.3 cm and 11.9 cm respectively.

The axis method is clearly superior and approaches the between-subject error of 7.6 cm. 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 8.0 cm. 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 7.6 cm. However, this would require that the software calculates the appropriate duration for the reach or that the user provides a value. In some human modelling 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 30 targets in the experiment and three objects moved. If we average just over the targets, the error for the end point parameterization is 8.2 cm, and for the axis parameterization it is 8.3 cm. If we average over targets and objects, the error is reduced to 7.7 cm. 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 co-ordinates 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 that was described above.

To look for systematic subject effects, we averaged the distance-scaled control points over each subject by using the axis parameterization. The error by using this model was 8.1 cm, 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 who were used in our experiment, we can see that there is little evidence for substantial systematic subject effects.

A summary table of the maximum median errors for the models chosen 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 $(-0.015, -0.139, -0.002)$. For the median reach distance of 56.7 cm, this gives offsets of $(-1.4, -8.4, 4.3)$ cm and

Table 2. Maximum median error for the models considered†

<i>Model type</i>	<i>Model form</i>	<i>Error (cm)</i>
Natural variation	Within subject	6.7
	Between subject	7.6
Speed profile	Observed (α, β)	3.2
	Median (α, β)	6.9
Null	Flash–Hogan speed	11.4
	Flash–Hogan	12.6
	End point parameterization	11.2
	Axis parameterization	8.2
	Axis with distance scaling	8.0
Predictor	Axis with distance–time scaling	7.6
	Target mean with end point	8.2
	Target mean with axis	8.3
	Target–object mean with axis	7.6
	Subject mean with axis and distance scaling	8.1

†For a full description and discussion, see the text.

(−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. towards himself or herself) and up (i.e. lifts the object) relative to a straight line path. This is consistent with the pattern that is seen in Fig. 6(c). 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 Fig. 7. This should be compared with Fig. 1. The predicted trajectories have shapes that are comparable with 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 model preferred, i.e. the null axis parameterization with distance scaling, and computed the residuals with respect to the




Fig. 7. Predicted reaches to 30 targets by using the null axis model with distance scaling (compare this with the observed data in Fig. 1): (a) top view; (b) front view; (c) side view

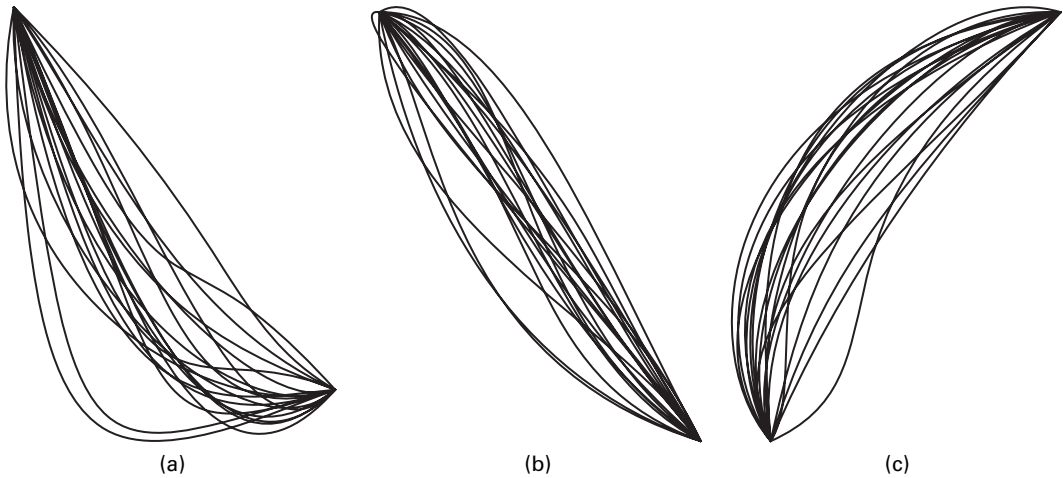


Fig. 8. Simulated reaches to the same target as in Fig. 2 (in this case, both the start and the end of the trajectories have been fixed): (a) top view; (b) front view; (c) side view

control points, resulting in a 2031×6 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 could not find any strong relationship between the residuals and the predictors, suggesting that no significant improvement to this model is possible within this general modelling approach.

We can use the covariance structure of the residuals to generate simulated trajectories that vary around the mean prediction. To illustrate this, we generated 20 observations from a multivariate normal distribution to construct the simulated reaches that are shown in Fig. 8. This can be compared with Fig. 2 where 20 actual trajectories are observed, albeit with somewhat varying end points. None-the-less, the similarity in the variation structure may be observed.

5. Discussion

We have demonstrated a modelling 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 that is 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 that are 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 that is 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 fitted well by the Bézier model that was 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 direction of motion 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 end points 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 modelling hand trajectories, it may also be applicable to other trajectories that are generally smooth and have curvatures that are affected by kinematic constraints at the end points. In the human modelling 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 co-ordinate system that is defined by a global vertical and a ‘forward’ direction that is defined by the initial posture and position of the subject. The definition of forward for these data was conveniently defined by the test set-up, but for other tasks, particularly those with concurrent pelvis motion or stepping, the appropriate co-ordinate system for the modelling 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-hand side of the body might be modelled by reversing the lateral co-ordinate signs in the model, but of course data from left-handed reaches will be necessary to assess this.

We have chosen to separate the modelling 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 modelling 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 smoothness of velocity 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 Human Motion Simulation Laboratory at the Center for Ergonomics at the University of Michigan and provided many years of support to the research that is presented in this paper. We are also grateful for the support of our colleagues at the Human Motion Simulation Laboratory.

References

- Blackmon, T., Cavusoglu, M., Fuji, L. and Stark, L. (1997) Human hand trajectory analysis in point-and-direct telerobotics. In *Proc. 8th Int. Conf. Advanced Robotics*, pp. 927–932. New York: Institute of Electrical and Electronics Engineers.
- Breteler, M., Meulenbroek, R. and Gielen, S. (1998) Geometric features of workspace and joint-space paths of 3d reaching movements. *Acta Psychol.*, **100**, 37–53.
- Choe, S. (2006) Statistical analysis of orientation trajectories via quaternions with application to human motion. *PhD Thesis*. Department of Statistics, University of Michigan, Ann Arbor.
- Fitts, P. (1954) The information capacity of the human motor system in controlling the amplitude of movement. *J. Exptl Psychol.*, **47**, 381–391.
- Flash, T. and Hogan, N. (1985) The coordination of arm movements: an experimentally confirmed mathematical model. *J. Neurosci.*, **5**, 1688–1703.
- Gleicher, M. (2001) Comparing constraint-based motion editing methods. *Graph. Mod.*, **63**, 107–134.
- Harris, C. and Wolpert, D. (1998) Signal-dependent noise determines motor planning. *Nat. Neurosci.*, **394**, 780–784.
- Jackson, S. and Husain, M. (1997) Visual control of hand action. *Trends Cognit. Sci.*, **1**, 310–317.
- Johnson, R. and Wichern, D. (2002) *Applied Multivariate Statistical Analysis*, 5th edn. Englewood Cliffs: Prentice Hall.
- Kim, C., Kim, W., Hong, C., Park, B.-U. and Jeong, M. (1999) Smoothing techniques via the bezier curve. *Commun. Statist. Theory Meth.*, **28**, 1577–1596.
- Mardia, K. and Jupp, P. (2000) *Directional Statistics*. Chichester: Wiley.
- Marler, T., Rahmatalla, S., Shanahan, M. and Abdel-Malek, K. (2005) A new discomfort function for optimization-based posture prediction. *Technical Paper 2005-01-2680*. Society of Automotive Engineers, Allendale.
- Moore, M. (1988) Predicting iceberg trajectories. *Chance*, **1**, no. 2, 30–36.
- Osborne, L., Lisberger, S. and Bialek, W. (2005) A sensory source for motor variation. *Nature*, **437**, 412–416.
- Paninski, L., Shoham, S., Fellows, M., Hatsopoulos, N. and Donoghue, J. (2004) Superlinear population encoding of dynamic hand trajectory in primary motor cortex. *J. Neurosci.*, **24**, 8551–8561.
- Park, W., Chaffin, D. and Martin, B. (2004) Toward memory-based human motion simulation: development and validation of a motion modification algorithm. *IEEE Trans. Syst. Man Cybernet. A*, **34**, 376–386.
- Prautzsch, H., Boehm, W. and Paluszny, M. (2002) *Bezier and B-spline Techniques*. New York: Springer.
- Small, M. and Samson, P. (1983) Stochastic simulation of atmospheric trajectories. *J. Clim. Appl. Meteorol.*, **22**, 266–277.
- Todorov, E. and Jordan, M. (2002) Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.*, **5**, 1226–1235.
- Uno, Y., Kawato, M. and Suzuki, R. (1989) Formation and control of optimal trajectory in human multijoint arm movement minimum torque-change model. *Biol. Cybernet.*, **61**, 89–101.
- Wang, J. (2006) Statistical modeling for 3-D trajectories. *PhD Thesis*. Department of Statistics, University of Michigan, Ann Arbor.
- Zacher, I. and Bubb, H. (2005) Strength-based discomfort model of posture and movement. *Technical Paper 2004-01-2139*. Society of Automotive Engineers, Allendale.