

Regression modeling of motion with endpoint constraints

Julian Faraway*

A statistical model is described for the prediction of reaching motions using motion capture data on a variety of individuals performing reaches to a range of targets. The modeling approach allows for various inputs such as the stature, age and the location of the target to be specified and then computes the predicted trajectories of the kinematic chains of body markers necessary to place an object exactly at the specified target. Functional regression methods for modeling time-varying angles and other quantities as well as trajectories are described. A new parameterization of posture is described that facilitates the satisfaction of specific endpoints such as placing an object at a target. The methodology is illustrated with an application to two-handed standing lifts. Copyright © 2003 John Wiley & Sons, Ltd.

Received: 25 March 2002; Revised 15 July 2002

KEY WORDS: inverse kinematics; postural constraints; analytical algorithms; computer animation

Introduction

Motion capture technology allows the collection of databases of human motions. We can collect data on a variety of individuals performing a range of motions. Even so, no matter how large our database may be, we may wish to see the motion of an individual with specified anthropometry perform a motion to some specified target and find that such a particular example does not exist within our database. In this article I describe how to use the observed data to build a statistical model that can be used to predict motions under a range of input conditions. Motion editing techniques take a single motion and modify it to meet requirements—see Gleicher¹ for a survey of such methods. I will show how to combine information from a larger number of observed motions to produce new predicted motions that will have superior statistical properties to a prediction based on a single motion. See Wiley and Hahn² for another approach to combining captured motions to produce new motions.

The application which motivated this work comes from ergonomics. Many vehicle interiors and workplaces are first designed using a CAD system. Physical prototypes are useful for evaluating the ergonomic

characteristics of the design but these are expensive and time consuming, especially when a design must proceed through several iterations. The ability to place an authentically moving virtual human within software such as *Jack*³ helps the designer detect problems with the layout that can be rapidly corrected, thus speeding the design process. In some other applications, it is sufficient to produce motions that appear to be correct, but our application demands that our predicted motions be close to how people really move. Such demands impose greater costs in terms of the quantity of motion capture data required. Even so, the methodology presented below is generalizable and would be useful across a wide range of applications.

Human motion is quite variable. The same individual performing the same task will not move identically on each repetition. Furthermore, even greater variation will be observed between the motions of different individuals performing the same task, even if they have the same height, weight or other personal characteristics of interest. Given specified characteristics and a task, I wish to predict the average motion. Regression modeling is a well-known statistical tool for predicting an output in terms of several inputs. The inputs here are quantities such as the height of the individual and the coordinates of the target of a reach. The output, however, is not the usual scalar or vector seen in ordinary regression models, but the whole motion. I describe an extension of regression models to handle such a complicated output.



*Correspondence to: J. Faraway, Department of Statistics, University of Michigan, Ann Arbor, MI 48109, USA.
E-mail: faraway@umich.edu

Although the main purpose of this article is to present a statistical methodology for modeling motion, I also demonstrate that the methods have been successfully applied to build predictive models based on a large amount of data collected on standing full body lifts with two hands where an object is moved to a wide range of locations. I describe how this data was collected below.

The dynamic posture can be described in terms of angles that vary over time and the trajectories of the body joints which can be represented as univariate functions and curves in 3D respectively. I describe new regression techniques that can predict such quantities in terms of inputs such as the stature, age and gender of the subject. I introduce a new representation of the posture, called stretch pivot coordinates, which will allow us to combine these angle and trajectory predictions to satisfy precise endpoint constraints such as placing an object in a specified location. These stretch pivot coordinates replace the joint angles used in most motion modeling. I then demonstrate how these methods were applied to our data and discuss how well they perform.

One important characteristic of my method is that no numerical optimization is required—it is analytic in the sense of Tolani *et al.*⁴ The construction of the predicted motion requires no iteration or solution of sets of equations and, as such, can be computed very rapidly compared to optimization and iterative inverse kinematic methods. In contrast to Tolani *et al.*⁴, my method allows the prediction of endpoint-constrained kinematic chains with an arbitrarily large number of links, but does not incorporate information about joint flexibility and limits.

Data

In this section, I describe the data to which the methods were applied. Since the primary objective of this article is to describe the motion-modeling methodology, I provide only enough detail to understand the nature of the data. There are many factors that may affect motion and the precise details of the conditions of the experiments are important *if* conclusions about these factors are required. However, this particular article is about how to model this type of data and does not aspire to specific conclusions about the effect of age etc. on motion. I emphasize that the methodology presented here could be applied to data collected in different ways for different linkages.

In 1999, the Human Motion Simulation Laboratory (HUMOSIM) at the University of Michigan conducted a set of experiments concentrated on the motions of

standing people performing reaches to a dispersed set of targets. In one experiment, subjects were required to move a small box with both hands from a position in front of the body at waist height to a shelf at the specified target. The subject then released the box and returned to the rest position, then reached back to the box, grasped it and returned it to the rest position. The left foot was held fixed but the rest of the body was free to move.

A total of about 3000 motions were performed by a group of 20 subjects reaching to 30 targets, where the four parts of the motion described above are counted separately. The subjects, from whom informed consent was obtained, were selected to provide a means to assess the effects of anthropometry (height in particular), gender, and age on the motions. The subjects ranged from very short to very tall and from 20 to 60 years of age. Two different motion capture systems were simultaneously used to estimate joint center locations throughout the motions: an optical reflective marker system and an electromagnetic one. The whole posture was described by 23 joint center locations located at the left and right feet, ankles, knees, hips, shoulders, elbows, wrist, front and back of hand together with the L5/S1 (small of the back), the C7/T1 (back of the neck), the sternoclavicular joint (supersternale), the mid-tragion (between the ears) and the nasion (between the eyes). Some of these locations were computed by projecting down from the externally placed markers.

I extracted the portions of the recorded motion where the subject was reaching based on the motion of the right hand. Because the data was collected at 20 Hz and each motion lasted around 1 second on average, each motion was represented by the 3D coordinates of the 23 markers at around 20 timepoints on average. Due to occlusion or other data collection problems, some markers are missing during some parts of some motions. Where the amount of missing data was not substantial, I interpolated these cases, otherwise the motion was discarded.

Methods

Functional Regression Analysis

Many individual components of the motion, such as angles between body segments as they change over time, can be described as functions. For example, consider an axis joining the initial and final location of the hand. I can compute the orthogonal distance of the hand from this axis during motion, which I will call the *radial*




Figure 1. The radial deviation of the hand from a straight-line path when reaching with the right hand to a location on the left for 20 subjects.

deviation. Because I observe the data only at discrete timepoints, I have a sequence of observed values from the start to the end of the motion. These sequences are of different lengths because some targets are further away than others and people reach at different speeds. We rescale all these motions so that $t=0$ is the start of the motion and $t=1$ is the end of the motion. I can save the actual time taken as a possible predictor of the motion.

Plots of this distance for 20 subjects reaching with the right hand to a location somewhat to the left and front of the body and about the same height as the initial position of the hand are shown in Figure 1. Since there are 30 targets and some reaches were replicated, a plot of the complete data for this measure would show considerably more curves.

I can describe each motion in terms of a set of functions. Some of these functions may describe how given angles change over time, while others may describe how other quantities like the radial deviation change. There are many choices of angles and other quantities that could be used to describe the motion, but I leave the question of which choices are better until later. I aim to build a set of models that can predict each of these functions, which can then be combined to predict the complete motion.

Suppose the rescaled functions for some chosen quantity are given by $\mathbf{y}(t) = (y_1(t), \dots, y_n(t))$, where the subscript runs over the n collected motions. These functions might be expected to depend on certain covariates such as the location of the target being reached, the age and anthropometry of the subject and other factors. For the i th curve, I collect these predictors in a vector x_i . Typically, the first term in this vector is one. We then propose a functional linear model:

$$y_i(t) = x_i^T \beta(t) + \varepsilon_i(t)$$

Notice that this is similar to a standard regression model but the response is now a function, as is the error term $\varepsilon_i(t)$. The regression coefficients $\beta(t)$ are now a vector of functions. A general introduction to the area of functional data analysis may be found in Ramsay and Silverman.⁵ The particular coefficient function for a given covariate will now represent the effect on the response of that covariate over the duration of the reach. I can now estimate $\beta(t)$ using least squares by

$$\hat{\beta}(t) = (X^T X)^{-1} X^T \mathbf{y}(t)$$

where X is the matrix whose rows are given by the x_i 's.

This formula cannot be directly applied since I cannot observe a $y_i(t)$ at all possible t . One approach is to approximate the functions on a grid of values. This was done in Faraway.⁶ A fine grid of values is necessary for accurate representation, which is somewhat inefficient. So, instead, I represent the curves as linear combinations of m cubic B-spline basis functions, $\Psi_j(t)$. A curve $y_i(t)$ is represented as $y_i(t) \approx \sum_{j=1}^m y_{ij} \Psi_j(t)$, where the coefficients y_{ij} are estimated using least squares over the points at which $y_i(t)$ is observed.

Given that human motion is usually quite smooth, it is not necessary to have a large number of basis functions. In our particular application, I found that eight basis functions were sufficient. So each observed curve is represented by eight coefficients and the functional response is thereby converted into a multivariate response.

Thus I can write the model in the form

$$Y_{n \times m} \Psi_{m \times 1}(t) = X_{n \times p} B_{p \times m} \Psi_{m \times 1}(t) + \varepsilon_{n \times 1}(t)$$

or factoring out the $\Psi(t)$, I can write it in the simpler form

$$Y_{n \times m} = X_{n \times p} B_{p \times m} + \varepsilon_{n \times m}$$

which is now a multivariate multiple regression model where the coefficient matrix B may be estimated using least squares:

$$\hat{B} = (X^T X)^{-1} X^T Y$$

I may then use the standard methods of statistical inference using this modeling approach. Details of such methods may be found in texts such as Johnson and Wichern.⁷ For prediction and interpretation

purposes, it is necessary to transform back from this basis function representation to the original form by

$$\hat{\beta}_{p \times 1}(t) = \hat{B}_{p \times m} \Psi_{m \times 1}(t)$$

I can predict future responses given a new predictor value x_0 by

$$\hat{y}_0(t) = x_{01 \times p} \hat{B}_{p \times m} \Psi_{m \times 1}(t)$$

More details on this approach may be found in Faraway.⁸

Trajectory Prediction

In addition to predicting univariate functions, such as the radial deviation, I may wish to predict the 3D curves formed by the trajectory of the hand or other body markers. I could simply model each of the three Cartesian coordinates of the trajectory using the technique described above, but this seems unsatisfactory as it is not invariant to rotations of the coordinate system. I have chosen a different parameterization that has more interpretable components and assumes that I have specified the beginning and end of the trajectory (as would be the case for predicting the movement of the hand for specified tasks).

I define $r(t)$ as the radial deviation at time t describing the orthogonal distance from the axis joining the endpoints, $p(t) \in [0, 1]$ as the proportionate progress along the axis at time t and $v(t) = dp/dt$ as the relative axial velocity.

See the left panel of Figure 2 for a depiction of these quantities.

Let $\phi(t)$ be the angle describing the position of the hand at time t on the circle orthogonal to the axis of the

reach and whose center lies on this axis. I define $\phi(t) = 0$ to be the projection of the unit vertical vector onto this circle.

I model the trajectory using the triplet

$$(v(t), r(t), \phi(t))$$

I model each of these components using a functional regression. Since $v(0) = v(1) = r(0) = r(1) = 0$ by definition, I can accommodate this by omitting the first and last cubic B-spline basis function which, since these are, respectively, the only non-zero basis functions at $t = 0$ and $t = 1$, will ensure the desired property. Furthermore, I should not directly model $\phi(t)$ using a functional regression because it is an angle. Angles $2\pi - \epsilon$ and 0 are only ϵ apart but if I enter the numerical values into the regression, this property will not be respected. For this reason, I model the responses $\cos \phi(t)$ and $\sin \phi(t)$ and then use the relation $\phi = \tan^{-1}(\sin \phi / \cos \phi)$ to predict ϕ , which does respect the appropriate continuity properties of an angle.

Greater detail on this method of trajectory prediction may be found in Faraway.⁹ Note also that a different method of predicting the trajectory could be freely inserted without changing the rest of the approach. For example, the trajectory could be modified to avoid an obstacle using other techniques.

Stretch Pivot Coordinates

In the previous sections, I have described how I model functions in terms of predictors and how I can model the trajectories of selected body markers. In this section, I describe how these methods may be combined to predict the motion of the whole using a representation that I




Figure 2. On the left, I show a side view of the reach with the axis drawn as a straight line connecting the start and end of the reach. On the right, the view is down the axis, looking from the start towards the end.

am calling the *stretch pivot* coordinates. The method is appropriate for predicting motions with constrained endpoints—for example, that the hand may be required to start and end at specified locations. Before I do this, I will describe the nature of the problem, a number of candidate solutions and why these are unsatisfactory.

The Problem. I have a closed kinematic chain. The endpoints are either fixed (for example, the foot when the subject remains in a standing position) or predicted using the trajectory method (for example, the hand during a reaching motion). There is a kinematic chain linking these two endpoints. In our example, the chain will stretch from the left foot to the right hand and contain 10 links. Of course, other linkages may branch off this chain. In our example, these would be the right leg, left arm and the head. I will model the main linkage first. These secondary linkages can be modeled in a similar fashion once the main linkage has been predicted. See Figure 3.

The problem is to predict the location of the interior markers of the chain.

Forward Kinematics Using Angle Predictions. Each segment of the linkage may be described using approximately two angles depending on the choice of linkage. There are many ways in which these angles may be specified but the same difficulties I describe below will arise. The angles may be modeled using functional regression analysis in terms of the desired predictors as described above. The posture must be reconstructed starting from one end of the chain, building up the linkage using the predicted angles using forward kinematics equations. However, when I reach the other end of the chain, the predicted location of the terminal marker will almost certainly not coincide with the required endpoint. When the linkage is long, the variation accumulates in the final marker and the discrepancy between the forward kinematic predicted location of the endpoint and the required location may




Figure 3. Schematic of human figure. The critical kinematic chain linking the endpoints at the left ball of the foot which, in our example, does not move, and the right grip, which moves along a trajectory that I have already predicted. I need to predict the locations of the nine interior markers of the 10-link chain as labeled in the diagram. When this is done, the three smaller chains consisting of the right leg, left arm and head can be predicted. For two-handed moves, the trajectory of the left hand can also be predicted.

be substantial. In our particular example, this error had a median value of around 25 cm across all the motions and subjects. This amount of error is unacceptable for many applications.

Unrestricted forward kinematic prediction is still useful for other situations. For example, in our case I need to predict the location/orientation of the head. This can be described using a two-segment linkage originating at $c7/t1$, proceeding to the mid-tragion (center of the head) and ending at the nasion (between the eyes). Since the position of the nasion is unconstrained, I can simply predict it using this method.

Inverse Kinematics. For an unrestricted kinematic chain with l links, fully flexible joints and with only one endpoint fixed, we need $2l$ angles to describe the position of the chain. When the other endpoint is fixed, three degrees of freedom are used and thus only $2l - 3$ angles (or other quantities) would be necessary to describe the location of the chain. Thus, the forward kinematics approach suffers from having more angles than I need, while from the inverse kinematics point of view there are an infinite number of postures that would satisfy the endpoint constraints. Many authors have described various criteria for selecting the 'best' solution from this infinite number of possibilities. See, for example, Zhang *et al.*¹⁰ or Hsiang and Ayoub¹¹ in the biomechanics literature, while Gleicher and Litwinowicz¹² is an example of an animation-motivated approach which also reviews previous approaches. These criteria have been drawn from various considerations regarding comfort, energy, joint mobility restrictions etc. In Faraway *et al.*¹³ I proposed selecting the solution that was closest to the unconstrained forward kinematics prediction.

These approaches require the solution of an optimization problem with non-linear constraints. This is feasible for short linkages but the methods rapidly become increasingly time consuming, complex and unreliable as the number of links increases. Tolani *et al.*⁴ describe the disadvantages of these numerical approaches.

Stretch Pivot. Only $2l - 3$ parameters are necessary to describe a closed l -link kinematic chain. The key to success is selecting these parameters in a clever way. As a starting point, consider a closed two-link chain. Only one parameter is need to describe this linkage, since the midpoint of the chain is constrained to lie on a circle whose center lies on and is orthogonal to an axis joining the endpoints. I need only specify the angle on this circle. I call this midpoint the *pivot* and I call this




Figure 4. The pivot angle describes the location of the midpoint on the circle of its possible positions.

angle the *pivot angle*. Such an angle was used by Kor-ein¹⁴ Wang^{15,16} and Tolani *et al.*⁴ Some authors have called this a *swivel* angle. The angle is illustrated in Figure 4.

For more than two links, one might consider modeling $l - 2$ of the links using forward kinematics and then use a pivot angle to describe the orientation of the remaining two links. However, the predicted position of the endpoint of the $l - 2$ links could be further away from the other endpoint than the combined distance of the two remaining links. For a long linkage, this condition may easily occur and is the reason why this solution was not satisfactory for our particular example. Nevertheless, the pivot angle idea does provide the hint towards a more general solution.

Consider an l -link chain and pick a marker in the middle of this chain. The selected marker does not need to be the middle in the sense of median, but assume, for now, that there are at least two links on each side of the selected marker. Let us arbitrarily call the two endpoint markers the proximal and the distal, and the selected midpoint the medial. Let l_p and l_d be the distances between the proximal and the medial and the distal and the medial respectively. If l_p and l_d are considered fixed, then the position of the medial relative to the proximal and distal may be described in terms of a pivot angle, θ_m , lying on the circle orthogonal to, and whose center lies on, the axis joining the proximal and distal.

Let m_p and m_d respectively represent the total length of all the links joining the proximal and the medial, and the distal and the medial. So

$$0 \leq l_p \leq m_p \quad \text{and} \quad 0 \leq l_d \leq m_d$$

Define $p_p = l_p/m_p$ and $p_d = l_d/m_d$. I call the p 's the *stretch* parameters. I may describe the position of the medial in terms of the three parameters (θ_m, p_p, p_d) . Hence the name *stretch pivot* (see Figure 5).




Figure 5. Stretch parameters illustrated. The distance between the proximal and medial if that part of the chain were fully extended is m_p , while the corresponding distance for the medial to the distal is m_d . I define $p_p = l_p/m_p$ and $p_d = l_d/m_d$.

Once the position of the pivot has been determined, the problem is reduced to two smaller problems. I may repeat the same procedure on the two halves of the chain recursively until I have reduced it to a collection of two- or three-link chains. For the two-link chains, only a pivot angle is needed to describe the middle marker. For a three-link chain, I may describe the position of one of the midpoints in terms of one pivot angle and one stretch parameter, thereby reducing the problem to a two-link chain requiring only one more pivot angle for description. This parameterization uses exactly $2l - 3$ parameters, equal to the number of degrees of freedom in the chain.

I may model the stretch pivot parameters using functional regression analysis. Any predicted pivot angle will be valid in the sense that it will lead to a possible

configuration of the chain. Of course, anatomy further restricts what angles are physically possible. I will discuss such concerns later.

The stretch parameters must lie between 0 and 1 so it is simpler to model $\log p/(1 - p)$ as the response so that the predicted p will always lie between 0 and 1. Depending on the particular configuration, further restrictions will apply to the pair (p_p, p_d) because the two halves of the chain must neither be too stretched or not stretched enough so that the halves might not join in the middle. Let m_{pd} be the distance from the proximal to the distal. The triangle inequality and their being proportions means that (p_p, p_d) must satisfy the following constraints:

$$\begin{aligned} p_d m_d + p_p m_p &\geq m_{pd} \\ p_d m_d + m_{pd} &\geq p_p m_p \\ m_{pd} + p_p m_p &\geq p_d m_d \\ 0 &\leq p_p \leq 1 \\ 0 &\leq p_d \leq 1 \end{aligned}$$

Depending on the relative values of (m_p, m_d, m_{pd}) , this results in a range of potential values for (p_p, p_d) , some of which are depicted in Figure 6.

I need to model a bivariate response, (p_p, p_d) , but the irregular shape of the domain being dependent on the relative values of (m_p, m_d, m_{pd}) makes it difficult to build a prediction equation that will always produce valid values. I propose the following solution: I will work with $(p_p + p_d, p_d - p_p)$, where the first parameter




Figure 6. Allowable range for (p_p, p_d) lies within the unit square, further restricted by the triangle inequality constraints shown as dotted lines. In all three panels, the proximal is at $(0, 0, 0)$ and the distal is at $(1, 1, 1)$. In the first panel, $m_p = 4, m_d = 4$; in the second $m_p = 2, m_d = 2$; and in the third $m_p = 1, m_d = 4$. The points in all three plots are the values of (p_p, p_d) corresponding to four different values of (q_s, q_d) : $(0.1, 0.9)$, $(0.1, 0.1)$, $(0.9, 0.1)$ and $(0.9, 0.9)$.

measures how stretched out the configuration is, while the second measures the extent to which the midpoint is towards the proximal or distal.

For any given (m_p, m_d, m_{pd}) , I may determine the maximum, q_{smax} , and the minimum, q_{smin} , values that can be taken by the sum, $s = p_p + p_d$. I then define a new parameter q_s as

$$q_s = (s - q_{smin}) / (q_{smax} - q_{smin})$$

So q_s is a number between 0 and 1 representing the sum of the stretch parameters scaled within its allowable range.

Now treating, $s = p_p + p_d$ as fixed, I compute the range of the difference $d = p_d - p_p$ as (q_{dmin}, q_{dmax}) , and define q_d as

$$q_d = (d - q_{dmin}) / (q_{dmax} - q_{dmin})$$

So q_d is again a number between 0 and 1 representing the difference of the stretch parameters conditional on their sum, scaled within its allowable range. So any combination of (q_s, q_d) within the unit square may be mapped to a (p_p, p_d) for any (m_p, m_d, m_{pd}) combination. Importantly, the predicted configurations for a given (q_s, q_d) for different (m_p, m_d, m_{pd}) combinations will be roughly homologous in the sense of placing the medial in about the same location to the extent possible given the values of (m_p, m_d, m_{pd}) . This feature is important if I am to use (q_s, q_d) as regression responses.

The values of $q_{smin}, q_{smax}, q_{dmin}, q_{dmax}$ for given (m_p, m_d, m_{pd}) are given in the Appendix.

There remains the problem of the three-link chains which require only one stretch parameter to describe. Again the specific configuration of (m_p, m_d, m_{pd}) will determine the range of this stretch parameter, which will be a subset of $[0, 1]$. The problem is just a simpler version of the one solved above. I may compute the range of this stretch parameter and then compute an appropriate $q \in [0, 1]$ as above.

So I am now able to describe the posture in terms of pivot angles and transformed stretch parameters, q , which can each take any value within their allowable range independent of the value of any of the other parameters. I can independently manipulate any of these parameters while maintaining the endpoint constraint. This *crucial property* allows us to independently model each parameter using a functional regression equation. Given any predicted pivot and stretch parameters, I can construct a motion for any choice of link lengths long enough to reach the target.

Defining the Pivot Angles

In defining a pivot angle, I should define an origin (zero degrees). Two problems may occur. Firstly, when $m_p + m_d = m_{pd}$, i.e. the two links are collinear, the angle cannot be defined at all. When this condition is close to occurring, there will be some instability in the pivot angle in that small changes in the midpoint may result in large changes in the pivot angle. However, given that the ultimate aim is to predict the midpoint and not the pivot angle, this problem is not serious.

The second problem is more troublesome. One way to define zero is to pick a direction, say the vertical, and project this direction onto the circle of the pivot angle to define a zero. However, if the circle lies in the horizontal plane, the zero will be undefined. Of course, this is unlikely to occur exactly in practice, but, more seriously, there will be much instability in the pivot angle when the motion is such that the axis joining proximal and distal passes close to vertical. (Zero could also be defined using a plane, but the same problems will arise.)

Of course, one could just choose a different polar direction but this will simply move the problem elsewhere. For some combinations of body markers, it is possible to pick a good polar direction since we might know that any axis that is likely to occur will not be close to the polar direction. However, some combinations, such as the wrist, elbow and shoulder, can clearly have axes in all directions. Some adaptive choice is necessary.

Our approach is to avoid this zero problem entirely by defining the angle in terms of the normal vector to the plane passing through the pivot and the two endpoints. I then model the three coordinates of the normal vector using three functional regressions. Admittedly, three parameters are introduced where only one appears to be needed, but the polar instabilities are avoided and the normal vector does contain relevant information about the orientation of the endpoints and the pivot which can be usefully modeled.

Results

In this section, I describe the particular implementation I used for modeling the data described above and how well the models performed. These choices would need to be reconsidered and recombined for other types of motion but it is worth describing our choices here to understand the considerations involved.

The selection of pivots is essential to the implementation of the stretch pivot coordinates. First consider the

primary kinematic chain linking the left ball of the foot to the right hand. The pivots selected were c7/t1, the left hip, the right shoulder, the right elbow and the left knee, where the order indicates the order of precedence. The remaining markers on this chain required only a pivot angle to describe.

The right ball of the foot was free to move in our experiment. I modeled the Cartesian coordinates $(x(t), y(t), z(t))$ directly for this marker using a separate functional regression for each component. Once the position of the right foot was determined, the right knee was used as a pivot, with the remaining markers on this chain each requiring only a pivot angle to describe.

When the left hand was holding the box, the angles describing the orientation of the box were modeled using functional regression equations. Thus the position of the left grip was derived from the right grip together with the predicted box angles. For the motions where the left grip was empty (returning to the rest position or reaching to grasp the box on the shelf), the trajectory of the left hand was independently modeled in the same way as the right hand.

The two-link head chain did not have a fixed endpoint, so that the angles describing the orientation were simply modeled using functional regressions.

In order to make a prediction, it is necessary to provide the segment lengths (the distances between adjacent markers). Various studies have been made concerning these lengths and how they vary as a function of anthropometry. In our case, I fit simple regression models to describe these lengths as a linear function of stature. For our small number (20) of subjects, there was no evidence that more than stature alone would be helpful in predicting these lengths.

It is possible to accommodate variable segment lengths. For example, I might model the c7/t1 to l5/s1 spine link as a function varying in time depending on the predictors. This allows for some additional flexibility and biomechanical fidelity. However, all segment lengths were fixed in the model presented below.

All the functional regressions were fit using the robust regression method of Huber.¹⁷ I found that this method gave better fits than least squares. There are many potential predictors that could be used in the functional regression equations and still more ways that these can be combined. I needed to choose among these possibilities. If I consider the modeling of a single stretch or pivot angle parameter, I find, of course, that adding more predictors always improves the fit. Although the well-known techniques of variable selection could be applied, we may simply observe that fitting the

individual stretch and pivot parameters is not our primary objective. For any individual data point (which is a complete reach) and given model I compute a fitted value which is the predicted reach for the predictors associated with that data point. I can compute the distance between the observed data and the predicted value over the markers and over the time of the reach. I can repeat this over the complete data set and derive some measure of fit. I can use this to evaluate the competing models.

I experimented with a variety of different models and discovered that better results were obtained with more local rather than global models. One would not expect that information about a reach to a high shelf would be useful in predicting reaches to low shelves and yet a single predictor equation for the whole target area does just this. Instead, I found that better results were obtained by subsetting the data by target area and fitting separate models. This is not surprising. Consider, for example, reaches to a range of vertically stacked target shelves. For the higher shelves, it is not necessary to bend the knees, but, as the target shelf moves lower, at some point the subject will bend his knees during the reach. The transition from one pattern of motion to another is abrupt and not continuous in the target shelf—hence the need for more local fitting. A better understanding of such transition points would be valuable.

It is quite time consuming to fit and evaluate models so I was unable to further explore the wide variety of potential models that might be considered. Note also that I used the same predictors for all the functional regressions, although in fact different predictors could be used for different stretch and pivot parameters.

For my final model, evaluated and fit on the same data, I found a median error over all reaches averaged over time of 6.9 cm for the right grip and 7.9 cm for c7/t1. I used predicted (based on height) segment lengths rather than the actual segment lengths. Had we used the actual segment lengths for each individual, lower figures would have been obtained but these would not reflect the performance that one might expect in actual use as most users will only be able to supply the height and not all the segment lengths. When considering these figures, it is important to realize that there is a lower bound to what may be achieved given the natural variability in motion beneath which one cannot hope to predict. The same person repeating the same task will not do it identically. Even greater variation might be expected between reaches performed by two people of the same height, age and gender (and whichever other predictors might have been used) doing the same reach.

If I build a regression model with an indicator variable for each target and subject, I can get a notion of the best possible predictive performance (within the class of methods that I am using). The median errors for the right hand and c7/t1 for this model are 6.3 cm and 6.7 cm respectively. Of course this model is not practical since it could only be used for predicting motions to our targets made by our subjects, but it does indicate that the model I have chosen could not be greatly improved upon. This does not preclude the possibility that someone else may present a completely different modeling technique that outperforms mine, but, given the inherent natural variability in motion, the scope for significant improvement is not substantial.

A functional demonstration of the motion prediction may be obtained from my web site at <http://www.stat.lsa.umich.edu/~faraway>.

One important advantage of using stretch pivot coordinates over other methods that require optimization is that it is fast. To predict all four phases of the motion required only 0.01 seconds on an 850 Mhz Pentium III PC. The algorithm was implemented in C but without any special efforts at optimization. This speed makes it practical for applications where rapid animation is important.

Discussion

I mention here several shortcomings in the current method and discuss potential improvements.

The performance of any empirical model is only as good as the data it is based upon. In our case, the predictive performance declines as we specify inputs that are far from those observed in the experiment. For example, if I input a target far to the left of the subject, I cannot expect a good prediction since all the targets in the experiment were to the right of the sagittal plane. Such failings are inevitable and can only be reliably rectified by collecting more data with the required input conditions.

For a complex model such as this, it is inevitable that some small proportion of inputs will produce unrealistic predicted motions. One can try to reduce the occurrence of such errors with model improvements, but complete success seems unattainable. Some errant predicted motions can be detected because they violate joint mobility constraints or other numerical criteria. Corrections may be applied. Nevertheless, some motions may be numerically very close to real motions and yet appear unnatural. In some applications, the user might be willing and able to intervene to either correct or reject the motion. When

fully automated prediction is required, such errors are problematic. Another pitfall, particularly for ergonomic applications, is predicted motions which appear natural but have some hidden defect.

The current method contains no notion of obstacle avoidance. If the hand must avoid some obstacle then the trajectory prediction must be suitably modified. This seems feasible enough, but more problematic are obstacles that intersect with other parts of the body. An even more difficult problem occurs when the predicted position of the body intersects with itself. For example, on reaches to targets near the feet, I need to avoid predictions where the hand passes through the knees.

Some motions should not be averaged. For example, the typical head motion during a lift involves a fairly rapid glancing motion from looking at the box to the shelf. The timing of this rapid motion is variable. If several such motions are averaged, a smooth, slower move of the head from looking at the box to the target results. Such a move is not representative and so a different way of combining such observed head motions is required.

Most joints have mobility restrictions but the method does not respect these restrictions. This is a weakness since the joint angles, through which such restrictions are typically expressed, are not used by our method and so it not possible to explicitly enforce these restrictions. Since the method is based on empirical data, these limitations are not often exceeded, but nonetheless the problem may occur. In such cases, the dynamic posture predicted by our method may serve as an initial estimate requiring some correction. For example, elbow angle mobility violations can be rectified with a rotation upper-lower arm linkage without changing the rest of the posture.

Our model was built for implementation in Digital Human Modeling software for ergonomic design such as Jack. Joint angles may be directly computed from our predicted joint centers. However, the posture representation in Jack and other animation software varies—different joint centers, degrees of freedom and other constraints are used. Mapping our predicted motions into forms suitable for other software is not always straightforward but further discussion is beyond the scope of this article.

ACKNOWLEDGEMENTS

The author thanks the HuMoSim project and its industrial sponsors at the University of Michigan, and Don Chaffin in

particular. The Department of Mathematical Sciences at the University of Bath, which hosted the author during the period in which this article was written, is acknowledged.

References

- Gleicher M. Comparing constraint-based motion editing methods. *Graphical Models* 2001; **63**: 107–134.
- Wiley D, Hahn J. Interpolation synthesis of articulated figure motion. *IEEE Computer Graphics and Applications* 1997; **17**: 39–45.
- Badler N, Phillips C, Webber B. *Simulating Humans: Computer Graphics Animation and Control*. Oxford University Press: New York, 1993.
- Tolani D, Goswami A, Badler N. Real-time inverse kinematic techniques for anthropomorphic limbs. *Graphical Models* 2000; **62**: 353–388.
- Ramsay J, Silverman B. *Functional Data Analysis*. Springer: New York, 1997.
- Faraway J. Regression analysis for a functional response. *Technometrics* 1997; **39**: 254–261.
- Johnson R, Wichern D. *Applied Multivariate Statistical Analysis* (3rd edn). Prentice-Hall: Englewood Cliffs, NJ, 1992.
- Faraway J. Modeling reach motions using functional regression analysis. Society of Automotive Engineers Technical Paper 2000-01-2175, 2000.
- Faraway J. Modeling hand trajectories during reaching motions. Technical Report #383, Department of Statistics, University of Michigan.
- Zhang X, Kuo A, Chaffin D. An optimization-based differential inverse kinematics approach for modeling moderately complex three-dimensional dynamic seated postures. *Journal of Biomechanics* 1998; **31**: 1035–1042.
- Hsiang S, Ayoub M. Development of methodology in biomechanical simulation of manual lifting. *International Journal of Industrial Ergonomics* 1995; **17**: 271–288.
- Gleicher M, Litwinowicz P. Constraint-based motion adaptation. *Journal of Visualization and Computer Animation* 1998; **9**(2): 65–94.
- Faraway J, Zhang X, Chaffin D. Rectifying postures reconstructed from joint angles to meet constraints. *Journal of Biomechanics* 1999; **32**: 733–736.
- Korein J. *A Geometric Investigation of Reach*. MIT: Cambridge, MA, 1985.
- Wang X. Three-dimensional kinematic analysis of influence of hand orientation and joint limits on the control of arm postures and movements. *Biological Cybernetics* 1999; **80**: 449–463.

- Wang X. A behavior-based inverse kinematics algorithm to predict arm prehension postures for computer-aided ergonomic evaluation. *Journal of Biomechanics* 1999; **32**: 453–460.
- Huber P. *Robust Statistics*. Wiley: New York, 1981.

Appendix: Range of Sum and Difference of Stretch Parameters

q_{smax} :

Set $q_{smax} = 2$ but

- If $m_d - m_p \geq m_{pd}$ set $q_{smax} = (m_{pd} + m_p + m_d)/m_d$
- If $m_p - m_d \geq m_{pd}$ set $q_{smax} = (m_{pd} + m_p + m_d)/m_p$

q_{smin} :

- If $m_{pd} < m_d$ then set $a = m_{pd}/m_d$ else set $a = (m_{pd} - m_d + m_p)/m_p$
- If $m_{pd} < m_p$ then set $b = m_{pd}/m_p$ else set $b = (m_{pd} - m_p + m_d)/m_d$
- Set $q_{smin} = \min(a, b)$.

Let $x = p_p + p_d$.

q_{dmax} :

- Set $a = 2(m_{pd} + xm_p)/(m_p + m_d) - x$
- If $m_p > m_d$ set $b = x - 2(m_{pd} - xm_d)/(m_p - m_d)$ else set $b = 2$
- Set $q_{dmax} = \min(x, 2 - x, a, b)$

q_{dmin} :

- Set $a = x - 2(m_{pd} + xm_d)/(m_p + m_d)$
- If $m_p < m_d$, set $b = x - 2(m_{pd} - xm_d)/(m_p - m_d)$ else set $b = -2$.
- Set $q_{dmin} = \min(-x, x - 2, a, b)$

Author's biography:

Julian Faraway is a Professor of Statistics at the University of Michigan. He received a BA in Mathematics from Cambridge University in 1982 and a PhD in Statistics from the University of California, Berkeley, in 1987. He is currently Associate Director of the Human Motion Simulation Laboratory in the Center for Ergonomics at the University of Michigan. He is interested in the applications of statistics to human motion modeling of both the body and the face.