

12

A Differential Equation for Juggling

12.1 Introduction

Chapter 11 introduced the notion of modeling functional data using a differential equation. We introduced this type of model by relating it to ordinary least squares regression analysis. In this chapter we tackle a similar problem, but in a more challenging context that requires some extensions of the approach.

We saw that a sample of handwriting could be described by a linear differential equation of the form

$$x'''(t) = \beta_1(t)x'(t) + \beta_2(t)x''(t) + f(t) , \quad (12.1)$$

where $x'(t)$ is velocity, $x''(t)$ is acceleration, $x'''(t)$ is the third derivative or jerk taken along a specific coordinate direction, and $f(t)$ is a forcing or residual term that we hope is small. In the handwriting data, we had a separate equation of this sort for each of three coordinates, with the X -axis along the line of writing, the Y -axis vertically on the paper, and the Z -coordinate measuring a lift off the paper.

What makes a differential equation model particularly interesting is its capacity to link the observed position function $x(t)$ in a particular coordinate direction with the behavior of the velocity, acceleration and jerk functions, which we must derive from the observed data. It is at the level of acceleration especially that we can expect to see the influence of the body's motor control system, as contracting muscles apply forces to the body's framework, which in turn change acceleration directly as a con-

sequence of Newton's Second Law. However, while the visual system can feed position information back to the brain to modify the control process, in rapid and highly automated tasks such as handwriting and juggling, the time delays involved in neural transmission and central processing imply that visual feedback is playing only a limited role. Instead, the brain probably uses information from the strains applied to the body as it moves and interacts with its environment, and these messages are translated into control over acceleration through muscle contractions. In short, it is at the acceleration level that most of the action is to be found, and a model that accurately couples acceleration with the observable data is vital.

The X -, Y -, and Z -axes defined above constitute the coordinate system used to describe handwriting. Is this coordinate system natural? It is obvious that the X - Y plane should be the plane of the paper, and writing has a preferred direction on the paper, so defining the X -axis to be parallel to the lines on the paper is a natural coordinate system for handwriting. The presence of a natural coordinate system and the lack of external forces (other than the constraint that writing has to be on the paper itself!) must surely simplify the modeling process.

The process of juggling a ball seems more challenging than handwriting, and certainly far fewer people master it. The motion of the hand is in all three dimensions in space, however these dimensions are defined. Moreover, once the ball leaves the hand, the laws of physics take over, and the brain must anticipate where these laws are going to deliver the ball back into the juggler's hand. Since no two throws can be exactly the same, there is inevitably an interaction between brain processes and the external world that may complicate the situation. Finally, it is much less obvious how we should arrange the coordinate axes—or even whether we should use rigid Euclidean coordinates at all. This means that the model must be sufficiently invariant with respect to choice of coordinate frame that the fit will still work even when the statistical analysis has used the “wrong” coordinates.

12.2 The data and preliminary analyses

The data were collected from Professor Michael Newton of the Department of Biostatistics at the University of Wisconsin who, in addition to being a fine statistician, is an expert juggler. He juggled three balls for 10 sequences of 10 seconds each. Within a sequence, there were 12 to 13 cycles of throwing a ball and catching another, with a total of 123 cycles.

Small infrared emitting diodes (IREDs) were placed on the tip of Michael's forefinger, his wrist, and three locations on his chest. The positions of these IREDs were tracked 200 times a second by an OPTOTRAK camera system. Our main concentration is on the data for the tip of the forefinger. The chest IREDs provided a point of reference for slow move-

ments of his entire body, and by subtracting their positions from the other two, these drifts were removed from the data. We used a coordinate system with the X -axis sideways to Michael's right, the Y -axis outward from his body, and the Z -axis vertically upward. In most of our discussion we refer to the X -, Y -, and Z -coordinates as coordinates 1, 2, and 3 respectively.

The data were centered so that the average position was zero for each coordinate. Because we knew that we would need a smooth estimate of the third derivative, the jerk function $x'''(t)$, the data were smoothed using a roughness penalty method penalizing for the integral of the squared fifth derivative of each coordinate function. Software details, together with the data themselves, are given on the Web site corresponding to this chapter. The OPTOTRAK measurements are accurate to within 0.5 mm, and it was appropriate to choose a rather small penalty parameter value of $\lambda = 10^{-12}$.

Here is our notation. The function $x_j(t)$, $j = 1, 2, 3$ indicates the position of the forefinger in the X , Y , and Z directions, respectively. Corresponding velocities are denoted by $x'_j(t)$, accelerations by $x''_j(t)$, and jerks by $x'''_j(t)$. We also need the tangential velocity and the tangential acceleration

$$\begin{aligned}\|x'(t)\| &= \sqrt{x'_1(t)^2 + x'_2(t)^2 + x'_3(t)^2} \\ \|x''(t)\| &= \sqrt{x''_1(t)^2 + x''_2(t)^2 + x''_3(t)^2} .\end{aligned}\tag{12.2}$$

Partial cycles at the beginning and end of each record were trimmed off to obtain records that were comparable across trials, and we had to find a suitable landmark or curve feature to separate one cycle from another. The tangential velocity for the finger IRED showed a deep and stable minimum within each cycle, corresponding to the lowest point in the forefinger's trajectory, and the beginning of the launch of the ball. The beginning of a cycle was therefore defined by the location of this minimal value of $\|x'(t)\|$. The average duration of the cycles was 719 msec, and half the durations fell in the interval from 696 to 736 msec. The cycles did show some phase variation, so we applied the continuous registration method described in Section 7.6.1 to the tangential velocity functions $\|x'(t)\|$. The registration process means that averages of the curves and their derivatives will give good summaries of what is happening.

12.3 Features in the average cycle

Figure 12.1 shows the average position of the index finger in terms of coordinates 1 and 3, as could be observed by someone seeing through the juggler from behind. Figure 12.2 displays the view of the cycle from the right of the juggler. It is clear that there is a substantial twist in the cycle, and so all three coordinates are essential.

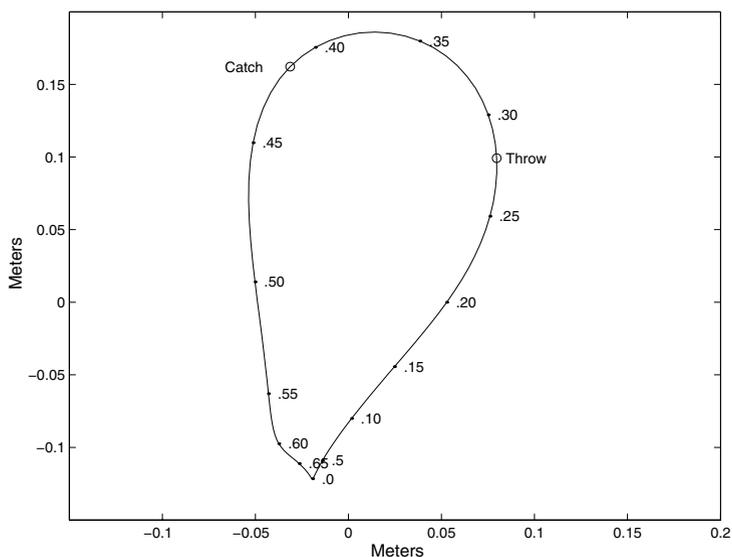


Figure 12.1. The average juggling cycle in the X - Z plane, or as seen from the juggler's perspective facing forward. The points on the curve indicate times in seconds, and the total cycle takes 0.711 seconds. The time when the ball leaves the hand and the time of the catch are shown as circles.

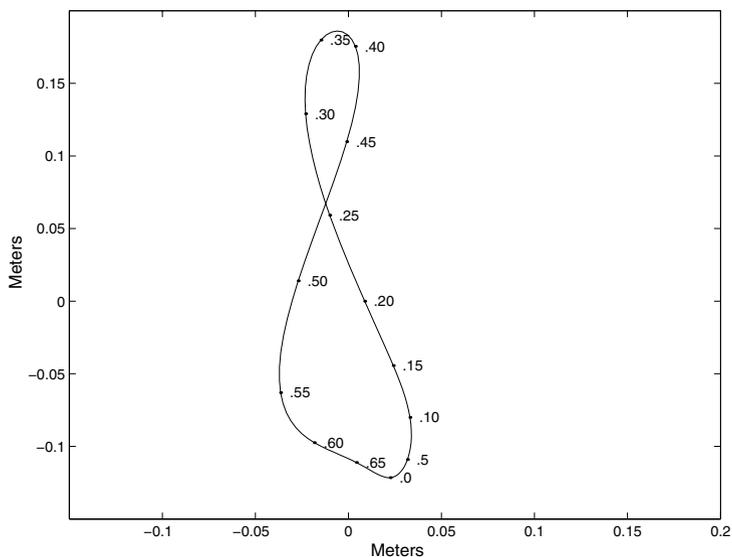


Figure 12.2. The average juggling cycle in the Y - Z plane, as seen from the perspective of someone standing to the juggler's right.

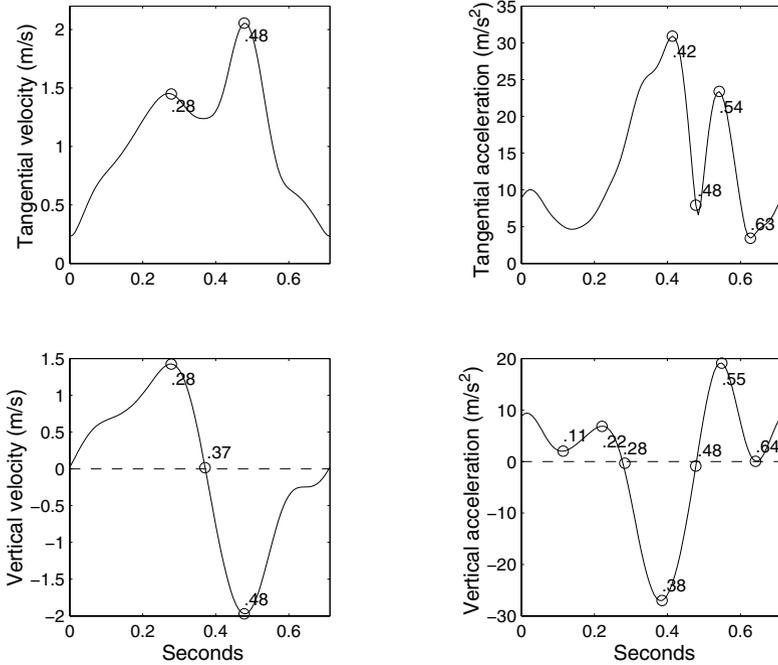


Figure 12.3. The tangential velocity $\|x'(t)\|$ of the tip of the forefinger is shown in the top left panel, and the tangential acceleration $\|x''(t)\|$ is shown in the top right panel. The vertical velocity $x'_3(t)$ and acceleration $x''_3(t)$ of the tip of the forefinger are shown in the bottom left and right panels, respectively. Important features in these curves are indicated by circles, and their times are indicated in seconds.

The cycle begins at the bottom of the trajectory, where the forefinger is poised to begin the launch of the ball. The ball leaves the hand at the point where the motion is nearly vertical, at 0.28 seconds. The catch occurs at 0.42 seconds, after the hand has moved through the top of the arc, and a little before the motion becomes vertical. The ball is lowered to a position a little to the left of the beginning of the cycle, and then transferred laterally to the point where the cycle begins again.

The top two panels in Figure 12.3 display the average tangential velocity and acceleration. We need to also look at the vertical component of acceleration separately since it is the upward movement of the ball that is the key to a successful juggling cycle, and these are found in the bottom two panels. The throw accelerates the ball from near rest to a velocity sufficient to carry it into the air for enough time that it can be caught on the third cycle. Consequently, we see that both the tangential and vertical velocity increase steadily to a maximum at 0.28 seconds. The vertical velocity and acceleration show a transitional phase within this throw portion at 0.11

seconds, probably due to the upward motion being transferred from the forefinger and wrist to the more slowly accelerating arm. After the throw, the forefinger then slows down slightly to permit the ball to clear the hand, and while it is moving across the top of the arc.

The catch shows up as a sharp negative minimum in vertical acceleration at 0.38 seconds as the downward force of the moving ball is transferred to the hand and finger. Moreover, since the hand is also moving laterally at this point, and must transfer this motion to the ball, we see a strong peak in the tangential acceleration at 0.42 seconds. The hand then accelerates downward, reaching its maximum velocity at 0.48 seconds, when the ball is falling nearly vertically.

At this point we enter the setup phase where the ball is positioned for its launch. A sharp positive peak in vertical acceleration is caused by the arm muscles contracting to slow the ball prior to transferring it back across the body to the launch position. We see in Figure 12.1 that this transfer takes around 0.11 seconds, and is comparatively slow compared to the speed that we see in the postcatch phase between 0.42 and 0.52 seconds.

In summary, we see something of note happening at intervals as small as 0.06 seconds in a cycle of length 0.7 seconds. As in many biomechanical processes, such as speaking, writing, and playing the piano, the brain is able to control muscular systems on very short time scales.

12.4 The linear differential equation

As with handwriting, we model juggling via a second-order linear differential equation in velocity rather than in position. In other words, the velocity function $x'(t)$ is the basic function to be modeled. The model then remains unchanged if we change the origin of the measurements. Since our decision to make the average spatial coordinates equal to zero was rather arbitrary, and certainly not related to any intrinsic structure of the motor control system that we are aware of, having a model that is invariant under translations seems essential.

Unfortunately, the coordinate system that we are using is not likely to be “natural” from a motor control point of view, unlike the handwriting situation where lateral, vertical, and lifting movements have a good chance of being controlled independently. Indeed, why should we even assume that the brain uses rigid Cartesian coordinates at right angles that do not change with time? Certainly, there may be cross-talk between coordinates, so that what is happening for the lateral X -coordinate may depend on what is also going on for the vertical Z -coordinate, for example. We need, therefore, a more general form that will not change if we alter coordinate systems at a later point in our research when we have some better ideas about

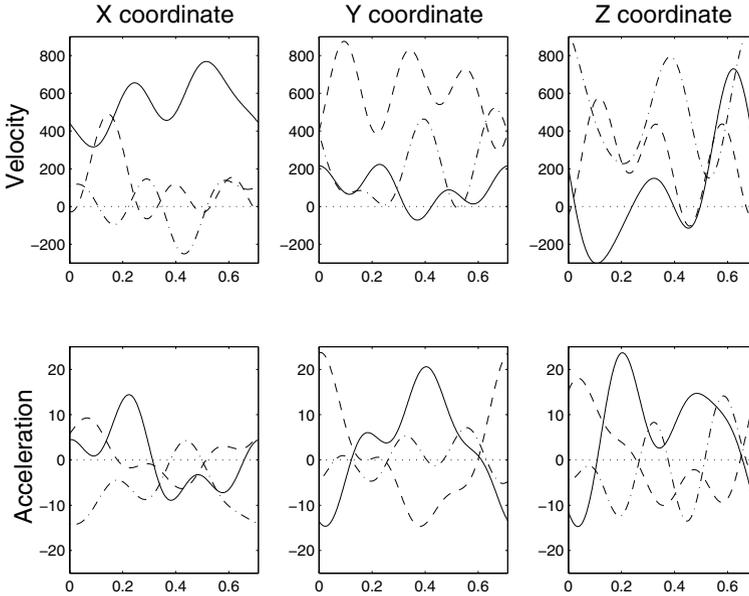


Figure 12.4. The top panels display the weight functions $\beta_{jk1}(t)$ on the three coordinate velocities for each coordinate. The bottom three panels show the acceleration weight functions $\beta_{jk2}(t)$. Within each panel, the X-coordinate weight function is the solid line, the Y-coordinate weight is the dashed line, and the Z-coordinate weight is the dashed-dotted line.

coordinates intrinsic to motor control, and that will allow for properties of one coordinate of velocity to be influential on another.

Consequently, we move to a *coupled* differential equation, while retaining linearity. This means that the change in each coordinate and its derivatives is considered to involve counterpart changes in each other coordinate. Here is the more general equation that we used:

$$x_{ij}'''(t) = \sum_{k=1}^3 [\beta_{jk1}(t)x'_{ik}(t) + \beta_{jk2}(t)x''_{ik}(t)] + f_{ij}(t) \quad \text{for } j = 1, 2, 3. \quad (12.3)$$

Note that i indexes replications and both j and k index coordinates. For coordinate j , regression coefficient weight functions $\beta_{jj1}(t)$ and $\beta_{jj2}(t)$ correspond to those given in the model (12.1) above. But for this j th coordinate we also have the four cross-coordinate regression coefficient weight functions $\beta_{jk1}(t)$ and $\beta_{jk2}(t)$, $k \neq j$. There are, therefore, a grand total of 18 weight functions to be estimated. This might seem like a lot, but remember that we have 123 juggling cycles at our disposal.

Figure 12.4 shows the weight functions $\beta_{jk1}(t)$ and $\beta_{jk2}(t)$ that we estimated. The estimation method is outlined in Section 12.6 below. The resulting estimated forcing functions correspond to residuals in standard

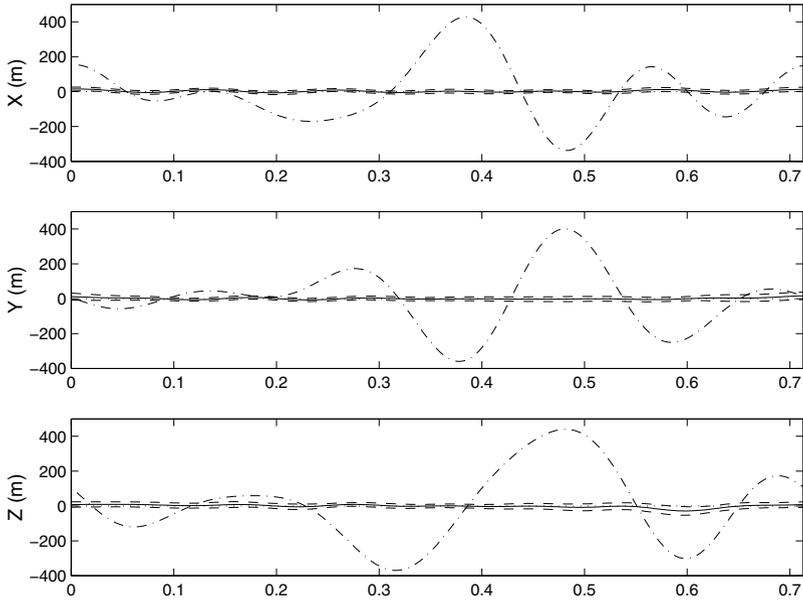


Figure 12.5. The panels display the mean forcing function $\bar{f}_j(t)$, 95% pointwise confidence limits for this function, and, for reference purposes, the mean jerk function \bar{J}_j for each coordinate. The solid line close to zero is the mean forcing function, the dashed lines on either side are 95% pointwise confidence limits, and the dashed-dotted line is the mean third derivative, displayed to indicate the relative size of the forcing function.

statistical modeling, and Figure 12.5 gives one assessment of the fit of the model, showing that the mean forcing function is much closer to zero than the mean jerk function for each of the coordinates. Another measure of fit is obtained by noting that, for each coordinate and for all t , over 99% of the variability in the jerk function is explained by the model.

What features do the estimated weight functions display? These functions were estimated using a Fourier basis with seven basis functions, which permits precisely three cycles, and in most cases the variation at this scale is clear. Allowing more cycles produces almost no improvement in fit, but on the other hand the fit deteriorates if fewer basis functions are allowed. This suggests that there is genuine detail in the brain's control mechanism at cycle lengths of order a quarter of a second.

Were we right in allowing cross-talk between coordinates? Looking at the effect of velocity (the top three panels in Figure 12.4) we see that the jerk in each coordinate is clearly influenced by that coordinate's own velocity. However, the Z -velocity has a clear influence on the jerks in the X - and Y -coordinates, and all three velocities seem to affect the jerk in the vertical direction. The acceleration effects are less clearcut, but there is

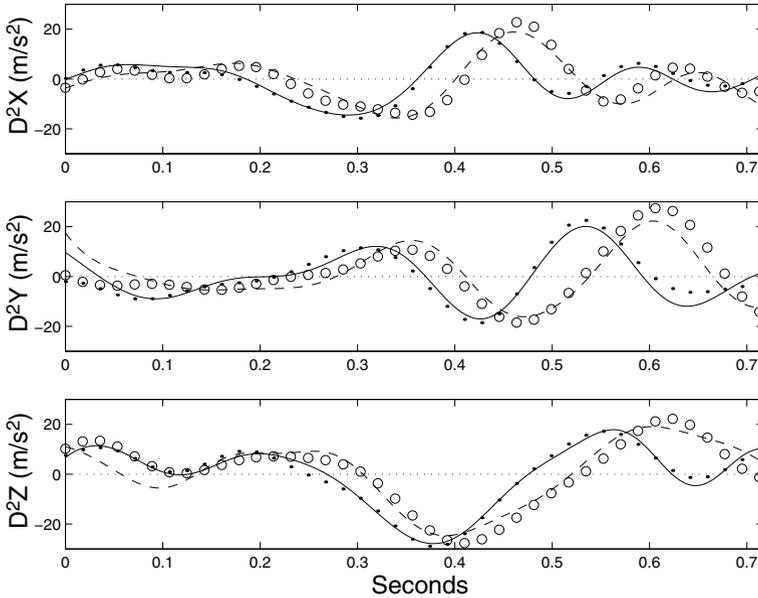


Figure 12.6. The fits to two sets of the coordinate acceleration functions based on the homogeneous linear differential equation. The fit and the actual data for the first record are plotted as a solid line and dots, respectively; and the results for a record from the middle of the juggling sequence as a dashed line and open circle, respectively.

no sense in which the effects of the different coordinates are disentangled. Confirmatory evidence of the need to allow influence between coordinates was provided by attempting to fit separate differential equation models to each coordinate; the quality of fit was much lower.

The fit of the equation to the data can be explored further by solving the homogeneous version of differential equation (12.3) for each coordinate, using the estimated weight functions β_{jk1} and β_{jk2} . There are six linearly independent solutions for the velocity functions, and every solution can be expressed as a linear combination of these. The solutions can each be thought of as a basis function, or mode of variation, in juggling cycle space, in rather the same way as the harmonics in principal components analysis. We may then approximate each of the 123 actual cycle curves or their derivatives by expanding them in terms of these functions. In approximating the curves themselves, we use the condition that the mean position is zero to recover the absolute position from the velocity. If a cycle is well modeled by the equation we would expect it to be well approximated by these basis functions.

Figure 12.6 shows how well the acceleration curves $x''_{ik}(t)$ are fit for the first cycle and for a cycle drawn from the middle of the sequence. The fits

shown are fairly typical for all cycles. Similar quality of approximation is also achieved for both position and velocity. Thus, the equation does a fine job of capturing both the curve shape for an individual record and its first two derivatives. Moreover, the six basis functions seem to do a good job of following the variation in the shape of the observed functional data from replication to replication.

12.5 What have we seen?

We saw in Section 12.3, and especially in Figure 12.3, that there are three main phases in the juggling cycle: throwing, catching, and setting up the next throw. Each phase lasts around 0.24 seconds. The tangential acceleration curves seem to display some cyclical features with approximate cycle lengths multiples of 0.12 seconds. This quasiperiodic character of acceleration has been observed in a wide variety of situations in neurophysiology, for example by Ramsay (2000) in the study of handwriting. It leads us to suspect that the motor control system uses a basic clock cycle to synchronize the contractions of the large numbers of muscles involved in complex tasks.

Our main modeling tool was the linear differential equation discussed in Section 12.4. We used this type of model because we already saw how important the acceleration curves were in describing the juggling process, and we wanted an approach that provided a good model of velocity and acceleration as well as the observed position data. Also, linear differential equation systems are the backbone of models in mechanics and other branches of engineering and science, and they should prove useful for describing biomechanical systems such as this. All the Fourier cycles used in the fitting of the weight functions shown in Figure 12.4 have cycle lengths that are multiples of 0.12, but this would not have been the case if a richer Fourier basis were used. The good fit of the model with this property is certainly consonant with the motor control clock cycle hypothesis.

The data were fitted extremely well by a second-order linear homogeneous differential equation, without any forcing function or nonlinear effects. The six modes of variability corresponding to the solutions of this equation fit individual juggling cycles extremely well and also allowed for the variation from one juggling cycle to another. In a certain sense, there is no variation between cycles; they are all controlled by the same differential equation, suggesting that the process of learning to juggle is one of “programming” a suitable differential equation into the person’s motor system.

It is beyond the scope of this chapter to attempt to discern what coordinate system the brain is using to plan movement. Preliminary investigations involving eigenvalue analyses of the matrices of coefficients β suggest that

the coordinate system remains relatively stable during parts of the cycle, and then changes as different muscle groups come into play. This, and several other aspects of our model fitting, are fascinating topics for future research.

12.6 Notes and references

The juggling study was carried out in collaboration with Dr. Paul Gribble of the University of Western Ontario in the motor control laboratory of Prof. David Ostry at McGill University.

Chapter 14 of Ramsay and Silverman (1997) gives more detail of the underlying methodology of this chapter, but only for the case of a one-dimensional variable rather than a space curve. We fit the model (12.3) by an integrated least squares procedure, the natural extension of the method set out in their Section 14.2. The criterion of fit of the functions β is to minimize the integrated residual sum of squares

$$\begin{aligned} \text{IRSE} &= \int \sum_{i,j} [f_{ij}(t)]^2 dt \\ &= \int \sum_{i,j} \left\{ x'''_{ij}(t) - \sum_{k=1}^3 [\beta_{jk1}(t)x'_{ik}(t) + \beta_{jk2}(t)x''_{ik}(t)] \right\}^2 dt \quad (12.4) \end{aligned}$$

The fit is regularized by constraining each β to have an expression in terms of a fairly small set of basis functions. In the juggling context, a seven-term Fourier expansion was used because of the periodicity of the problem; an alternative would be a B-spline basis on a fairly coarse knot sequence. The number of basis functions controls the degree of regularization, and other regularization approaches are possible.

In the present context, there are 18 β functions to be estimated, and hence $7 \times 18 = 126$ basis coefficients altogether. Substituting the basis expansions into (12.4) gives an expression for IRSE as a quadratic form in these 126 coefficients. The matrix and vector defining this quadratic form are found by numerical integration, and standard numerical techniques then yield the estimated coefficients.