Chapter 1
Introduction to Functional Data Analysis

The main characteristics of functional data and of functional models are introduced. Data on the growth of girls illustrate samples of functional observations, and data on the US nondurable goods manufacturing index are an example of a single long multilayered functional observation. Data on the gait of children and handwriting are multivariate functional observations. Functional data analysis also involves estimating functional parameters describing data that are not themselves functional, and estimating a probability density function for rainfall data is an example. A theme in functional data analysis is the use of information in derivatives, and examples are drawn from growth and weather data. The chapter also introduces the important problem of registration: aligning functional features.

The use of code is not taken up in this chapter, but R code to reproduce virtually all of the examples (and figures) appears in files ”fdarm-ch01.R” in the ”scripts” subdirectory of the companion ”fda” package for R, but without extensive explanation in this chapter of why we used a specific command sequence.

1.1 What Are Functional Data?

1.1.1 Data on the Growth of Girls

Figure 1.1 provides a prototype for the type of data that we shall consider. It shows the heights of 10 girls measured at a set of 31 ages in the Berkeley Growth Study (Tuddenham and Snyder, 1954). The ages are not equally spaced; there are four measurements while the child is one year old, annual measurements from two to eight years, followed by heights measured biannually. Although great care was taken in the measurement process, there is an average uncertainty in height values of at least three millimeters. Even though each record is a finite set of numbers, their values reflect a smooth variation in height that could be assessed, in principle, as

J.O. Ramsay et al., *Functional Data Analysis with R and MATLAB*, Use R,
DOI: 10.1007/978-0-387-98185-7_1,

often as desired, and is therefore a height *function*. Thus, the data consist of a sample of 10 *functional* observations $\texttt{Height}_i(t)$.

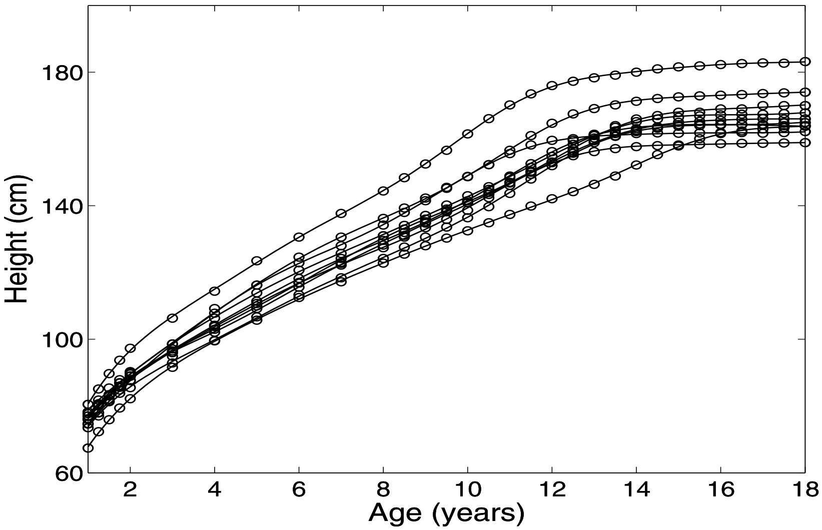

Fig. 1.1 The heights of 10 girls measured at 31 ages. The circles indicate the unequally spaced ages of measurement.

There are features in these data too subtle to see in this type of plot. Figure 1.2 displays the acceleration curves $D^2\texttt{Height}_i$ estimated from these data by Ramsay et al. (1995a) using a technique discussed in Chapter 5. We use the notation D for differentiation, as in

$$D^2\texttt{Height} = \frac{d^2\texttt{Height}}{dt^2}.$$

The pubertal growth spurt shows up as a pulse of strong positive acceleration followed by sharp negative deceleration. But most records also show a bump at around six years that is termed the midspurt. We therefore conclude that some of the variation from curve to curve can be explained at the level of certain derivatives. The fact that derivatives are of interest is further reason to think of the records as functions rather than vectors of observations in discrete time.

The ages are not equally spaced, and this affects many of the analyses that might come to mind if they were. For example, although it might be mildly interesting to correlate heights at ages 9, 10 and 10.5, this would not take account of the fact that we expect the correlation for two ages separated by only half a year to be higher than that for a separation of one year. Indeed, although in this particular example the ages at which the observations are taken are nominally the same for each girl, there is no real need for this to be so. In general, the points at which the functions are observed may well vary from one record to another.

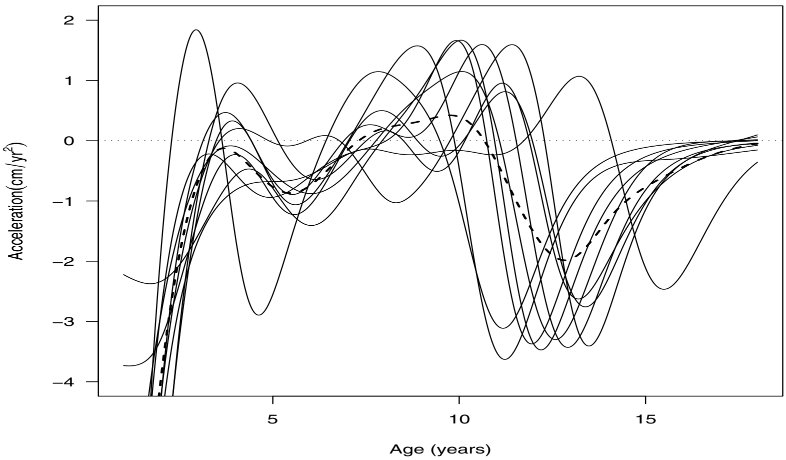

Fig. 1.2 The estimated accelerations of height for 10 girls, measured in centimeters per year per year. The heavy dashed line is the cross-sectional mean and is a rather poor summary of the curves.

The replication of these height curves invites an exploration of the ways in which the curves vary. This is potentially complex. For example, the rapid growth during puberty is visible in all curves, but both the timing and the intensity of pubertal growth differ from girl to girl. Some type of principal components analysis would undoubtedly be helpful, but we must adapt the procedure to take account of the unequal age spacing and the smoothness of the underlying height functions.

It can be important to separate variation in *timing* of significant growth events, such as the pubertal growth spurt, from variation in the *intensity* of growth. We will look at this in detail in Chapter 8 where we consider *curve registration*.

1.1.2 Data on US Manufacturing

Not all functional data involve independent replications; we often have to work with a single long record. Figure 1.3 shows an important economic indicator: the nondurable goods manufacturing index for the United States. Data like these often show variation as multiple levels.

There is a tendency for the index to show geometric or exponential increase over the whole century, and plotting the logarithm of the data in Figure 1.4 makes this trend appear linear while giving us a better picture of other types of variation. At a finer scale, we see departures from this trend due to the depression, World War II, the end of the Vietnam War and other more localized events. Moreover, at an

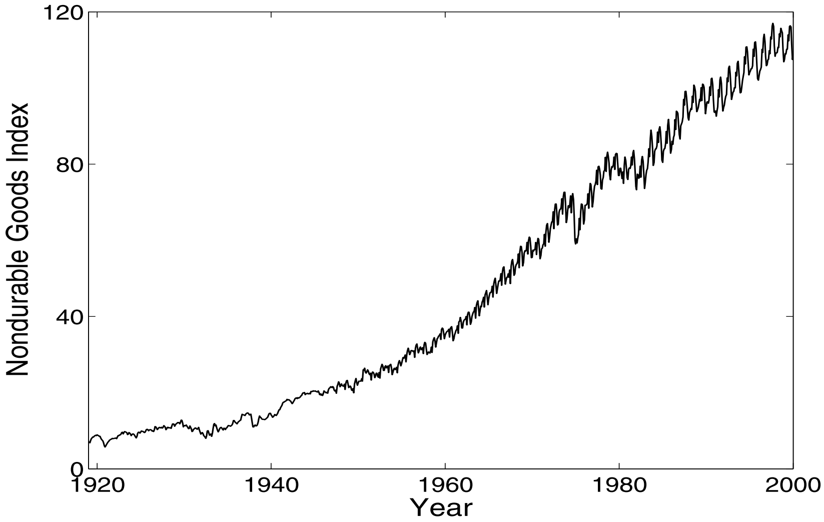

Fig. 1.3 The monthly nondurable goods manufacturing index for the United States.

even finer scale, there is a marked annual variation, and we can wonder whether this *seasonal trend* itself shows some longer-term changes. Although there are no independent replications here, there is still a lot of repetition of information that we can exploit to obtain stable estimates of interesting curve features.

1.1.3 Input/Output Data for an Oil Refinery

Functional data also arise as input/output pairs, such as in the data in Figure 1.5 collected at an oil refinery in Texas. The amount of a petroleum product at a certain level in a distillation column or cracking tower, shown in the top panel, reacts to the change in the flow of a vapor into the tray, shown in the bottom panel, at that level. How can we characterize this dependency? More generally, what tools can we devise that will show how a system responds to changes in critical input functions as well as other covariates?

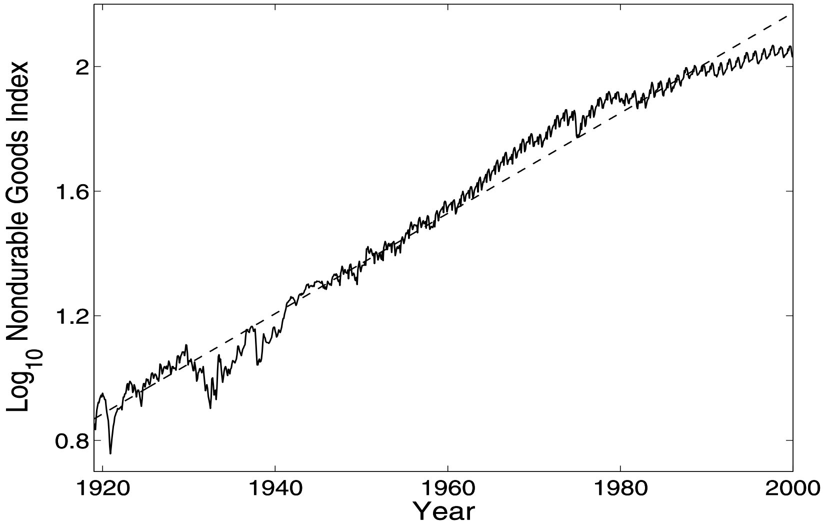

Fig. 1.4 The logarithm of the monthly nondurable goods manufacturing index for the United States. The dashed line indicates the linear trend over the whole time period.

1.2 Multivariate Functional Data

1.2.1 Data on How Children Walk

Functional data are often multivariate. Our third example is in Figure 1.6. The Motion Analysis Laboratory at Children's Hospital, San Diego, CA, collected these data, which consist of the angles formed by the hip and knee of each of 39 children over each child's gait cycle. See Olshen et al. (1989) for full details. Time is measured in terms of the individual gait cycle, which we have translated into values of t in $[0,1]$. The cycle begins and ends at the point where the heel of the limb under observation strikes the ground. Both sets of functions are periodic and are plotted as dotted curves somewhat beyond the interval for clarity. We see that the knee shows a two-phase process, while the hip motion is single-phase. It is harder to see how the two joints interact: The figure does not indicate which hip curve is paired with which knee curve. This example demonstrates the need for graphical ingenuity in functional data analysis.

Figure 1.7 shows the gait cycle for a single child by plotting knee angle against hip angle as time progresses round the cycle. The periodic nature of the process implies that this forms a closed curve. Also shown for reference purposes is the same relationship for the average across the 39 children. An interesting feature in this plot is the cusp occurring at the heel strike as the knee momentarily reverses its extension to absorb the shock. The angular velocity is clearly visible in terms of the spacing between numbers, and it varies considerably as the cycle proceeds.

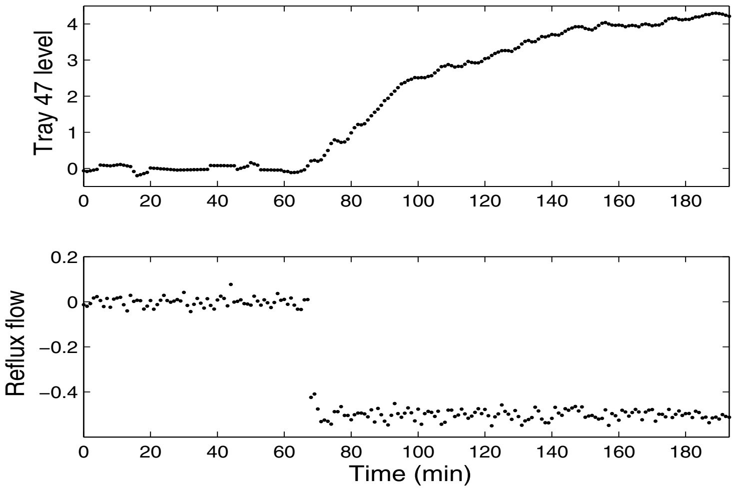

Fig. 1.5 The top panel shows 193 measurements of the amount of petroleum product at tray level 47 in a distillation column in an oil refinery. The bottom panel shows the flow of a vapor into that tray during an experiment.

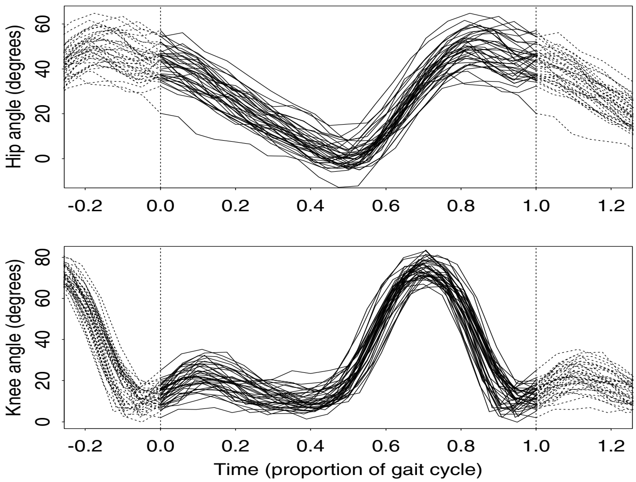

Fig. 1.6 The angles in the sagittal plane formed by the hip and knee as 39 children go through a gait cycle. The interval $[0,1]$ is a single cycle, and the dotted curves show the periodic extension of the data beyond either end of the cycle.

The child whose gait is represented by the solid curve differs from the average in two principal ways. First, the portion of the gait pattern in the C–D part of the cycle shows an exaggeration of movement relative to the average. Second, in the part of the cycle where the hip is most bent, this bend is markedly less than average; interestingly, this is not accompanied by any strong effect on the knee angle. The overall shape of the cycle for this particular child is rather different from the average. The exploration of variability in these functional data must focus on features such as these.

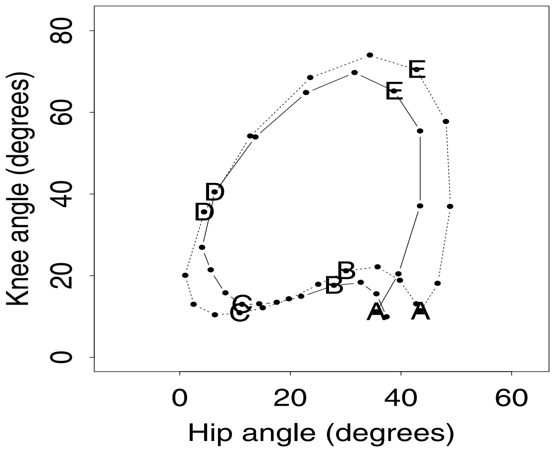

Fig. 1.7 Solid line: The angles in the sagittal plane formed by the hip and knee for a single child plotted against each other. Dotted line: The corresponding plot for the average across children. The points indicate 20 equally spaced time points in the gait cycle. The letters are plotted at intervals of one fifth of the cycle with A marking the heel strike.

1.2.2 Data on Handwriting

Multivariate functional data often arise from tracking the movements of points through space, as illustrated in Figure 1.8, where the X-Y coordinates of 20 samples of handwriting are superimposed. The role of time is lost in plots such as these, but can be recovered to some extent by plotting points at regular time intervals.

Figure 1.9 shows the first sample of the writing of “statistical science” in simplified Chinese with gaps corresponding to the pen being lifted off the paper. Also plotted are points at 120-millisecond intervals; many of these points seem to coincide with points of sharp curvature and the ends of strokes.

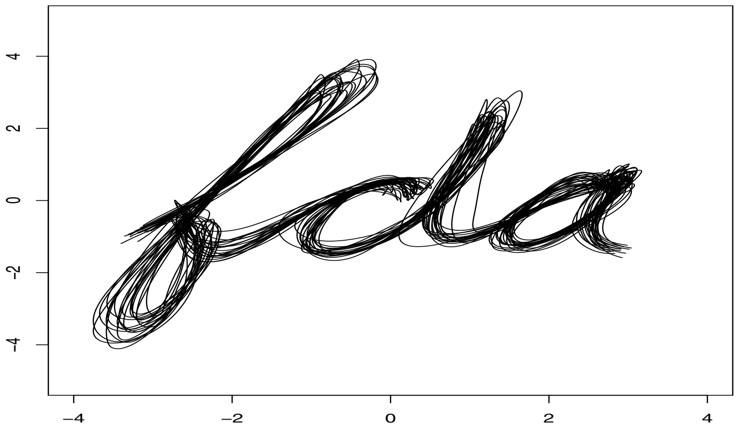

Fig. 1.8 Twenty samples of handwriting. The axis units are centimeters.

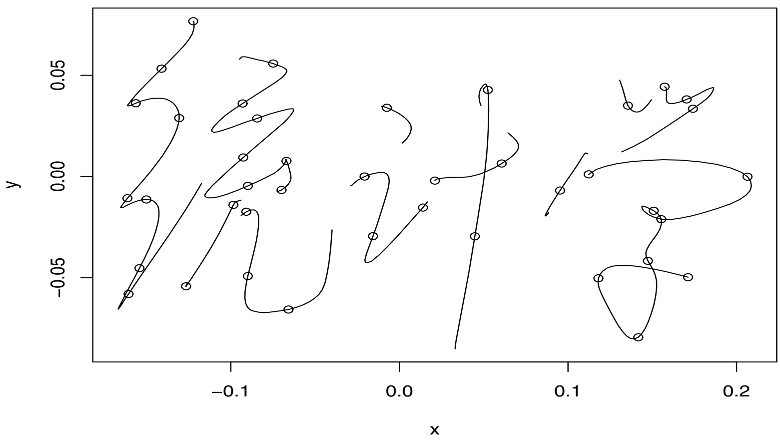

Fig. 1.9 The first sample of writing “statistical science” in simplified Chinese. The plotted points correspond to 120-millisecond time steps.

Finally, in this introduction to types of functional data, we must not forget that they may come to our attention as full-blown functions, so that each record may consist of functions observed, for all practical purposes, everywhere. Sophisticated online sensing and monitoring equipment now routinely used in research in fields such as medicine, seismology, meteorology and physiology can record truly functional data.

1.3 Functional Models for Nonfunctional Data

The data examples above seem to deserve the label "functional" since they so clearly reflect the smooth curves that we assume generated them. Beyond this, functional data analysis tools can be used for many data sets that are not so obviously functional.

Consider the problem of estimating a probability density function p to describe the distribution of a sample of observations $x_1, \ldots, x_n$. The classic approach to this problem is to propose, after considering basic principles and closely studying the data, a *parametric model* with values $p(x|\theta)$ defined by a fixed and usually small number of parameters in the vector θ. For example, we might consider the normal distribution as appropriate for the data, so that $\theta = (\mu, \sigma^2)'$. The parameters themselves are usually chosen to be descriptors of the shape of the density, as in location and spread for the normal density, and are therefore the focus of the analysis.

But suppose that we do not want to assume in advance one of the many textbook density functions. We may feel, for example, that the application cannot justify the assumptions required for using any of the standard distributions. Or we may see features in histograms and other graphical displays that seem not to be captured by any of the most popular distributions. *Nonparametric density* estimation methods assume only smoothness, and permit as much flexibility in the estimated $p(x)$ as the data require or the data analyst desires. To be sure, parameters are often involved, as in the density estimation method of Chapter 5, but the number of parameters is not fixed in advance of the data analysis, and our attention is focused on the density function p itself, not on parameter estimates. Much of the technology for estimation of smooth *functional parameters* was originally developed and honed in the density estimation context, and Silverman (1986) can be consulted for further details.

Psychometrics or mental test theory also relies heavily on functional models for seemingly nonfunctional data. The data are usually zeros and ones indicating unsuccessful and correct answers to test items, but the model consists of a set of *item response functions*, one per test item, displaying the smooth relationship between the probability of success on an item and a presumed latent ability continuum. Figure 1.10 shows three such functional parameters for a test of mathematics estimated by the functional data analytic methods reported in Rossi et al. (2002).

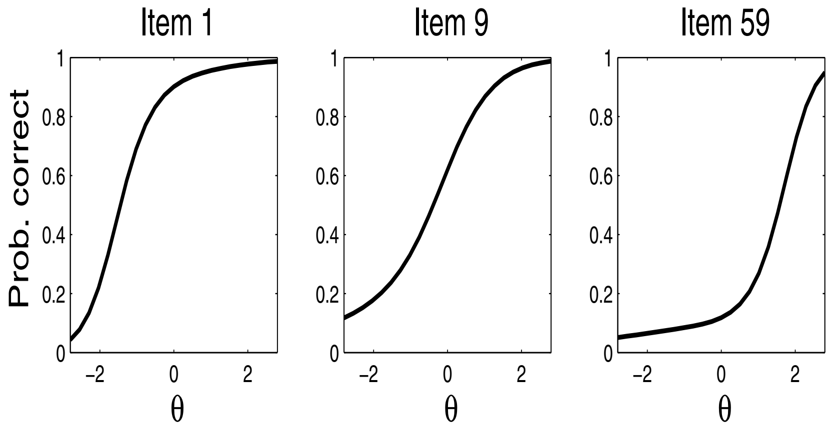

Fig. 1.10 Each panel shows an item response function relating an examinee's position θ on a latent ability continuum to the probability of a correct response to an item in a mathematics test.

1.4 Some Functional Data Analyses

Data in many fields come to us through a process naturally described as functional. Consider Figure 1.11, where the mean temperatures for four Canadian weather stations are plotted as smooth curves. Montreal, with the warmest summer temperature, has a temperature pattern that appears to be nicely sinusoidal. Edmonton, with the next warmest summer temperature, seems to have some distinctive departures from sinusoidal variation that might call for explanation. The marine climate of Prince Rupert is evident in the small amount of annual variation in temperature. Resolute has bitterly cold but strongly sinusoidal temperatures.

One expects temperature to be periodic and primarily sinusoidal in character and over the annual cycle. There is some variation in the timing of the seasons or phase, because the coldest day of the year seems to be later in Montreal and Resolute than in Edmonton and Prince Rupert. Consequently, a model of the form

$$\texttt{Temp}_i(t) \approx c_{i1} + c_{i2}\sin(\pi t/6) + c_{i3}\cos(\pi t/6) \tag{1.1}$$

should do rather nicely for these data, where $\texttt{Temp}_i$ is the temperature function for the ith weather station, and (c_{i1}, c_{i2}, c_{i3}) is a vector of three parameters associated with that station.

In fact, there are clear departures from sinusoidal or simple harmonic behavior. One way to see this is to compute the function

$$L\texttt{Temp} = (\pi/6)^2 D\texttt{Temp} + D^3\texttt{Temp}. \tag{1.2}$$

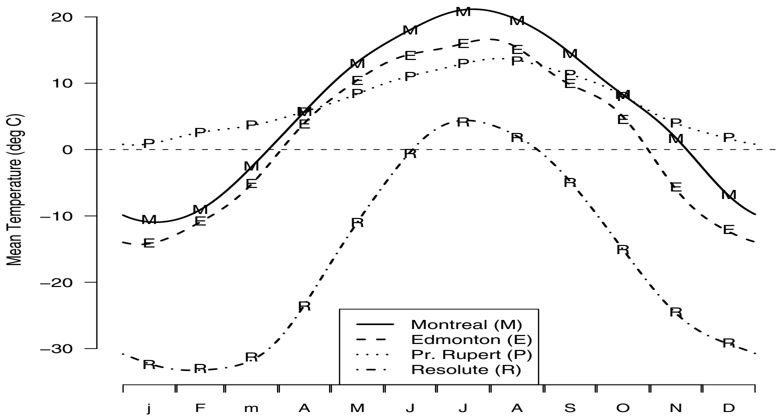

Fig. 1.11 Mean temperatures at four Canadian weather stations.

The notation D^mTemp means "take the mth derivative of function Temp," and the notation LTemp stands for the function which results from applying the *linear differential operator* $L = (\pi/6)^2 D + D^3$ to the function Temp. The resulting function, LTemp, is often called a *forcing function.* If a temperature function is truly sinusoidal, then LTemp should be exactly zero, as it would be for any function of the form (1.1). That is, it would conform to the *differential equation*

$$L\texttt{Temp} = 0 \ \text{ or } \ D^3\texttt{Temp} = -(\pi/6)^2 D\texttt{Temp}.$$

But Figure 1.12 indicates that the functions $L\texttt{Temp}_i$ display systematic features that are especially strong in the summer and autumn months. Put another way, temperature at a particular weather station can be described as the solution of the *nonhomogeneous* differential equation corresponding to $L\texttt{Temp} = u$, where the forcing function u can be viewed as input from outside of the system, or as an exogenous influence. Meteorologists suggest, for example, that these spring and autumn effects are partly due to the change in the reflectance of land when snow or ice melts, and this would be consistent with the fact that the least sinusoidal records are associated with continental stations well separated from large bodies of water.

Here, the point is that we may often find it interesting to remove effects of a simple character by applying a differential operator, rather than by simply subtracting them. This exploits the intrinsic smoothness in the process. Long experience in the natural and engineering sciences suggests that this may get closer to the underlying driving forces at work than just adding and subtracting effects, as is routinely done in multivariate data analysis. We will consider this idea in depth in Chapter 11.

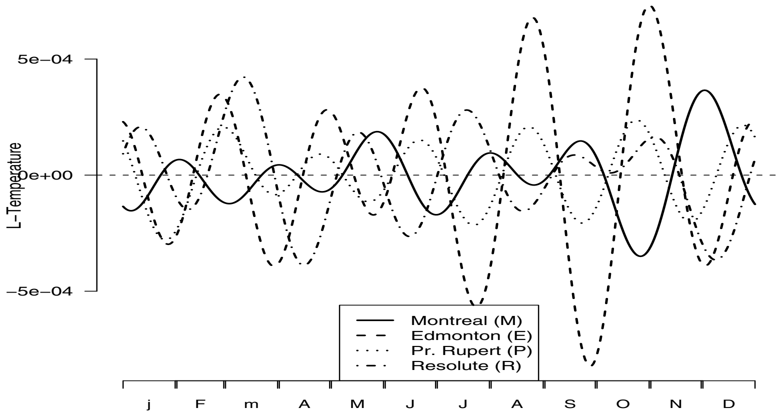

Fig. 1.12 The result of applying the differential operator $L = (\pi/6)^2 D + D^3$ to the estimated temperature functions in Figure 1.11. If the variation in temperature were purely sinusoidal, these curves would be exactly zero.

1.5 First Steps in a Functional Data Analysis

1.5.1 Data Representation: Smoothing and Interpolation

Assuming that a functional datum for replication i arrives as a finite set of measured values, $y_{i1}, \ldots, y_{in}$, the first task is to convert these values to a function x_i with values $x_i(t)$ computable for any desired argument value t. If these observations are assumed to be errorless, then the process is *interpolation*, but if they have some observational error that needs removing, then the conversion from (finite) data to functions (which can theoretically be evaluated at an infinite number of points) may involve *smoothing.*

Chapter 5 offers a survey of these procedures. The *roughness penalty* smoothing method discussed there will be used much more broadly in many contexts throughout the book, and not merely for the purpose of estimating a function from a set of observed values. The daily precipitation data for Prince Rupert, one of the wettest places on the continent, is shown in Figure 1.13. The curve in the figure, which seems to capture the smooth variation in precipitation, was estimated by penalizing the squared deviations in *harmonic acceleration* as measured by the differential operator (1.2).

The gait data in Figure 1.6 were converted to functions by the simplest of interpolation schemes: joining each pair of adjacent observations by a straight line segment. This approach would be inadequate if we required derivative information. However,

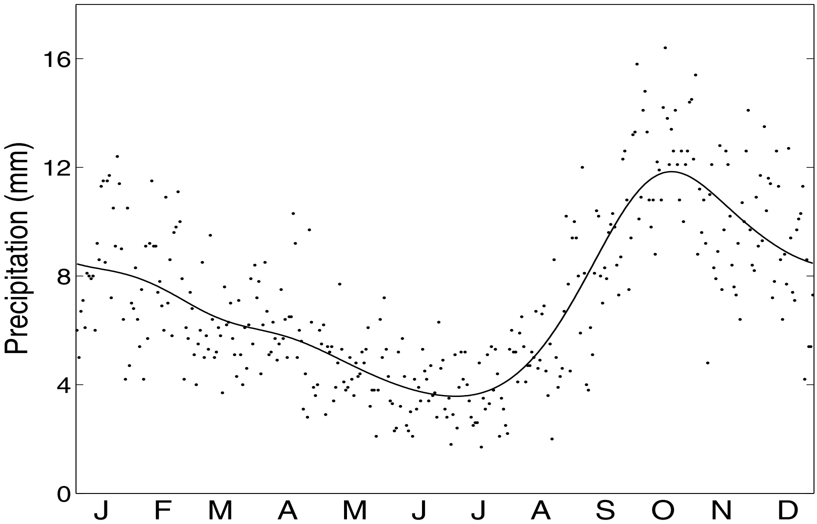

Fig. 1.13 The points indicate average daily rainfall at Prince Rupert on the northern coast of British Columbia. The curve was fit to these data using a roughness penalty method.

one might perform a certain amount of smoothing while still respecting the periodicity of the data by fitting a Fourier series to each record: A constant plus three pairs of sine and cosine terms does a reasonable job for these data. The growth data in Figures 1.1, 1.2, and 1.15 were fit using smoothing splines. The temperature data in Figure 1.11 were fit smoothing a finite Fourier series. This more sophisticated technique can also provide high-quality derivative information.

There are often conceptual constraints on the functions that we estimate. For example, a smooth of precipitation such as that in Figure 1.13 should logically never be negative. There is no danger of this happening for a station as moist as Prince Rupert, but a smooth of the data in Resolute, the driest place that we have data for, can easily violate this constraint. The growth curve fits should be strictly increasing, and we shall see that imposing this constraint results in a rather better estimate of the acceleration curves that we saw in Figure 1.2. Chapter 5 shows how to fit a variety of constrained functions to data.

1.5.2 Data Registration or Feature Alignment

Figure 1.14 shows some biomechanical data. The curves in the figure are 20 records of the force exerted on a meter during a brief pinch by the thumb and forefinger. The subject was required to maintain a certain background force on a force meter and then to squeeze the meter aiming at a specified maximum value, returning af-

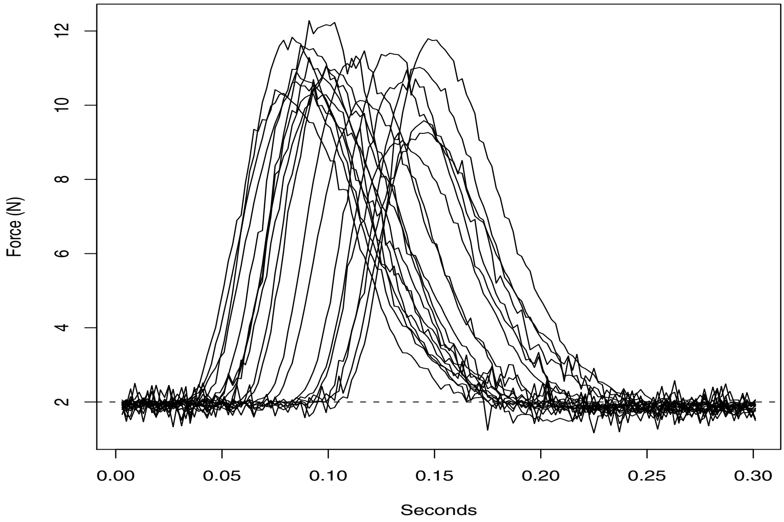

Fig. 1.14 Twenty recordings of the force exerted by the thumb and forefinger where a constant background force of 2 newtons was maintained prior to a brief impulse targeted to reach 10 newtons. Force was sampled 500 times per second.

terwards to the background level. The purpose of the experiment was to study the neurophysiology of the thumb–forefinger muscle group. The data were collected at the MRC Applied Psychology Unit, Cambridge, by R. Flanagan (Ramsay et al. 1995b).

These data illustrate a common problem in functional data analysis. The start of the pinch is located arbitrarily in time, and a first step is to align the records by some shift of the time axis. In Chapter 8 we take up the question of how to estimate this shift and how to go further if necessary to estimate record-specific linear or nonlinear transformations of the argument.

1.5.3 Graphing Functional Data

Displaying the results of a functional data analysis can be a challenge. With the gait data in Figures 1.6 and 1.7, we have already seen that different displays of data can bring out different features of interest, and the standard plot of $x(t)$ against t is not necessarily the most informative. It is impossible to be prescriptive about the best type of plot for a given set of data or procedure, but we shall give illustrations of various ways of plotting the results. These are intended to stimulate the reader's imagination rather than to lay down rigid rules.

1.5.4 Plotting Pairs of Derivatives: Phase-Plane Plots

Let us look at a couple of plots to explore the possibilities opened up by access to derivatives of functions. Figure 1.15 contains *phase-plane plots* of the female height curves in Figure 1.1, consisting of plots of the accelerations or second derivatives against their velocities or first derivatives. Each curve begins in the lower right in infancy, with strong positive velocity and negative acceleration. The middle of the pubertal growth spurt for each girl corresponds to the point where her velocity is maximized after early childhood. The circles mark the position of each girl at age 11.7, the average midpubertal age. The pubertal growth loop for each girl is entered from the right and below, usually after a cusp or small loop. The acceleration is positive for a while as the velocity increases until the acceleration drops again to zero on the right at the middle of the spurt. The large negative swing terminates near the origin where both velocity and acceleration vanish at the beginning of adulthood.

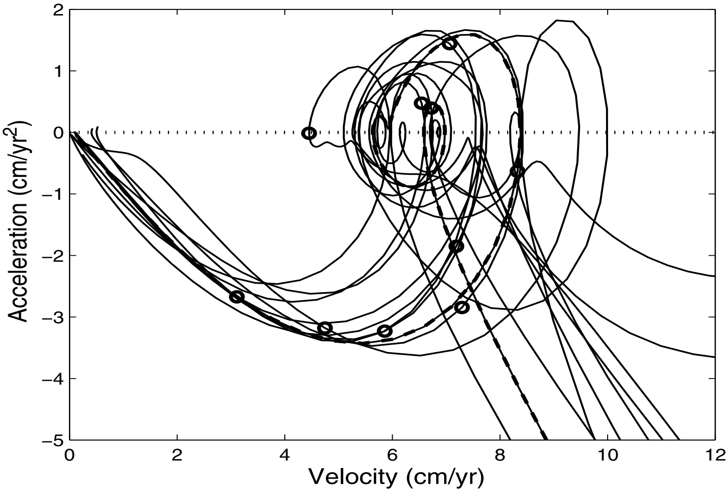

Fig. 1.15 The second derivative or acceleration curves are plotted against the first derivative or velocity curves for the ten female growth curves in Figure 1.1. Each curve begins in time off the lower right with the strong velocity and deceleration of infant growth. The velocities and accelerations at age 11.7 years, the average age of the middle of the growth spurt, are marked on each curve by circles. The curve is highlighted by a heavy dashed line is that of a girl who goes through puberty at the average age.

Many interesting features in this plot demand further consideration. Variability is greatest in the lower right in early childhood, but it is curious that two of the 10 girls have quite distinctive curves in that region. Why does the pubertal growth spurt show up as a loop? What information does the size of the loop convey? Why are the larger loops tending to be on the right and the smaller on the left? We see from the shapes of the loop and from the position of the 11.7 year marker that girls with early pubertal spurts (marker point well to the left) tend to have very large loops, and late-spurt girls have small ones. Does interchild variability correspond to something like growth energy? Clearly there must be a lot of information in how velocity and acceleration are linked together in human growth, and perhaps in many other processes as well.

1.6 Exploring Variability in Functional Data

The examples considered so far offer a glimpse of ways in which the variability of a set of functional data can be interesting, but there is a need for more detailed and sophisticated ways of investigating variability. These are a major theme of this book.

1.6.1 Functional Descriptive Statistics

Any data analysis begins with the basics: estimating means and standard deviations. Functional versions of these elementary statistics are given in Chapter 7. But what is elementary for univariate and classic multivariate data turns out to be not always so simple for functional data. Chapter 8 returns to the functional data summary problem, and shows that *curve registration* or feature alignment may have to be applied in order to separate *amplitude variation* from *phase variation* before these statistics are used.

1.6.2 Functional Principal Components Analysis

Most sets of data display a small number of dominant or substantial modes of variation, even after subtracting the mean function from each observation. An approach to identifying and exploring these, set out in Chapter 7, is to adapt the classic multivariate procedure of principal components analysis to functional data. Techniques of smoothing are incorporated into the functional principal components analysis itself, thereby demonstrating that smoothing methods have a far wider rôle in functional data analysis than merely in the initial step of converting a finite number of observations to functional form.

1.6.3 Functional Canonical Correlation

How do two or more sets of records covary or depend on one another? While studying Figure 1.7, we might consider how correlations embedded in the record-to-record variations in hip and knee angles might be profitably examined and used to further our understanding the biomechanics of walking.

The functional linear modeling framework approaches this question by considering one of the sets of functional observations as a covariate and the other as a response variable. In many cases, however, it does not seem reasonable to impose this kind of asymmetry. We shall develop two rather different methods that treat both sets of variables in an even-handed way. One method essentially treats the pair $(\texttt{Hip}_i, \texttt{Knee}_i)$ as a single vector-valued function, and then extends the functional principal components approach to perform an analysis. Chapter 7 takes another approach, a functional version of canonical correlation analysis, identifying components of variability in each of the two sets of observations which are highly correlated with one another.

For many of the methods we discuss, a naïve approach extending the classic multivariate method will usually give reasonable results, though regularization will often improve these. However, when a linear predictor is based on a functional observation, and also in functional canonical correlation analysis, imposing smoothness on functional regression coefficients is not an optional extra, but rather an intrinsic and necessary part of the analysis; the reasons are discussed in Chapters 7 and 8.

1.7 Functional Linear Models

The techniques of linear regression, analysis of variance, and linear modeling all investigate the way in which variability in observed data can be accounted for by other known or observed variables. They can all be placed within the framework of the linear model

$$y = \mathbf{Z}\beta + \varepsilon \tag{1.3}$$

where, in the simplest case, y is typically a vector of observations, β is a parameter vector, $\mathbf{Z}$ is a matrix that defines a linear transformation from parameter space to observation space, and ε is an error vector with mean zero. The design matrix $\mathbf{Z}$ incorporates observed covariates or independent variables.

To extend these ideas to the functional context, we retain the basic structure (1.3) but allow more general interpretations of the symbols within it. For example, we might ask of the Canadian weather data:

- If each weather station is broadly categorized as being Atlantic, Pacific, Continental or Arctic, in what way does the geographical category characterize the detailed temperature profile `Temp` and account for the different profiles observed? In Chapter 10 we introduce a functional analysis of variance methodology, where both the parameters and the observations become functions, but the matrix $\mathbf{Z}$ remains the same as in the classic multivariate case.

- Could a temperature record `Temp` be used to predict the logarithm of total annual precipitation? In Chapter 9 we extend the idea of linear regression to the case where the independent variable, or covariate, is a function, but the response variable (log total annual precipitation in this case) is not.
- Can the temperature record `Temp` be used as a predictor of the entire precipitation profile, not merely the total precipitation? This requires a fully functional linear model, where all the terms in the model have more general form than in the classic case. This topic is considered in Chapter 10.
- We considered earlier the many roles that derivatives play in functional data analysis. In the functional linear model, we may use derivatives as dependent and independent variables. Chapter 10 is a first look at this idea, and sets the stage for the following chapters on differential equations.

1.8 Using Derivatives in Functional Data Analysis

In Section 1.4 we have already had a taste of the ways in which derivatives and linear differential operators are useful in functional data analysis. The use of derivatives is important both in extending the range of simple graphical exploratory methods and in the development of more detailed methodology. This is a theme that will be explored in much more detail in Chapter 11, but some preliminary discussion is appropriate here.

Chapter 11 takes up the question, unique to functional data analysis, of how to use derivative information in studying components of variation. An approach called *principal differential analysis* identifies important variance components by estimating a linear differential operator that will annihilate them (if the model is adequate). Linear differential operators, whether estimated from data or constructed from external modeling considerations, also play an important part in developing regularization methods more general than those in common use.

1.9 Concluding Remarks

In the course of the book, we shall describe a considerable number of techniques and algorithms to explain how the methodology we develop can actually be used in practice. We shall also illustrate this methodology on a variety of data sets drawn from various fields, including where appropriate the examples introduced in this chapter. However, it is not our intention to provide a cookbook for functional data analysis.

In broad terms, our goals are simultaneously more ambitious and more modest: more ambitious by encouraging readers to think about and understand functional data in a new way but more modest in that the methods described in this book are hardly the last word in how to approach any particular problems. We believe that readers will gain more by experimenting with various modifications of the principles described herein than by following any suggestion to the letter. To make this easier, script files like ”fdarm-ch01.R” in the ”scripts” subdirectory of the companion ”fda” package for R can be copied and revised to test understanding of the concepts. The ”debug” function in R allows a user to walk through standard R code line by line with real examples until any desired level of understanding is achieved.

For those who would like access to the software that we have used, a selection is available on the website:

`http://www.functionaldata.org`

and in the `fda` package in R. This website will also be used to publicize related and future work by the authors and others, and to make available the data sets referred to in the book that we are permitted to release publicly.

1.10 Some Things to Try

In this and subsequent chapters, we suggest some simple exercises that you might consider trying.

1. Find some samples of functional data and plot them. Make a short list of questions that you have about the processes generating the data. If you do not have some data laying around in a file somewhere, here are some suggestions:
   a. Use your credit card or debit/bank card transactions in your last statement. If you keep your statements or maintain an electronic record, consider entering also the statements for five or so previous months or even for the same month last year.
   b. Bend over and try to touch your toes. Please do not strain! Have someone measure how far your fingers are from the floor (or your wrist if you are that flexible). Now inhale and exhale slowly. Remeasure and repeat for a series of breath cycles. Now repeat the exercise, but for the person doing the measuring.
   c. Visit some woods and count the number of birds that you see or the number of varieties. Do this for a series of visits, spread over a day or a week. If over a week, record the temperature and cloud and precipitation status as well.
   d. Visit a weather website, and record the five-day temperature forecasts for a number of cities.
   e. If you have a chronic concern like allergies, brainstorm a list of terms to describe the severity of the condition, sort the terms from mild to severe and assign numbers to them. Also brainstorm a list of possible contributing factors and develop a scale for translating variations in each contributing factor into numbers. Each day record the level of the condition and each potential contributing factor. One of us solved a serious allergy problem doing this.