

Functional Data Analysis of Amplitude and Phase Variation

J. S. Marron, James O. Ramsay, Laura M. Sangalli and Anuj Srivastava

Abstract. The abundance of functional observations in scientific endeavors has led to a significant development in tools for functional data analysis (FDA). This kind of data comes with several challenges: infinite-dimensionality of function spaces, observation noise, and so on. However, there is another interesting phenomena that creates problems in FDA. The functional data often comes with lateral displacements/deformations in curves, a phenomenon which is different from the height or amplitude variability and is termed *phase variation*. The presence of phase variability artificially often inflates data variance, blurs underlying data structures, and distorts principal components. While the separation and/or removal of phase from amplitude data is desirable, this is a difficult problem. In particular, a commonly used alignment procedure, based on minimizing the L^2 norm between functions, does not provide satisfactory results. In this paper we motivate the importance of dealing with the phase variability and summarize several current ideas for separating phase and amplitude components. These approaches differ in the following: (1) the definition and mathematical representation of phase variability, (2) the objective functions that are used in functional data alignment, and (3) the algorithmic tools for solving estimation/optimization problems. We use simple examples to illustrate various approaches and to provide useful contrast between them.

Key words and phrases: Functional data analysis, registration, warping, alignment, elastic metric, dynamic time warping, Fisher–Rao metric.

1. INTRODUCTION

1.1 A First Look at Phase Variation in Functional Data

Experimental units of data that are distributed over lines and areas, known as functional data, are best represented as curves and surfaces, respectively; and we expect that these will vary in height over any particular point. But we often notice that

J. S. Marron is Professor, Department of Statistics and Operation Research, University of North Carolina, 318 Hanes Hall, CB# 3260, Chapel Hill, North Carolina 27599-3260, USA e-mail: marron@unc.edu. James O. Ramsay is Professor Emeritus, Department of Psychology, McGill University, 1205 Dr Penfield Avenue Montreal, QQuebec H3A 1B1, Canada e-mail: ramsay@psych.mcgill.ca. Laura M. Sangalli is Associate Professor, MOX, Dipartimento di Matematica, Politecnico di Milano, Piazza L. da Vinci 32, 20133 Milano, Italy e-mail: laura.sangalli@polimi.it. Anuj Srivastava is Professor, Department of Statistics, Florida State University, Tallahassee, Florida 32306, USA e-mail: anuj@stat.fsu.

This is an electronic reprint of the original article published by the [Institute of Mathematical Statistics](#) in *Statistical Science*, 2015, Vol. 30, No. 4, 468–484. This reprint differs from the original in pagination and typographic detail.

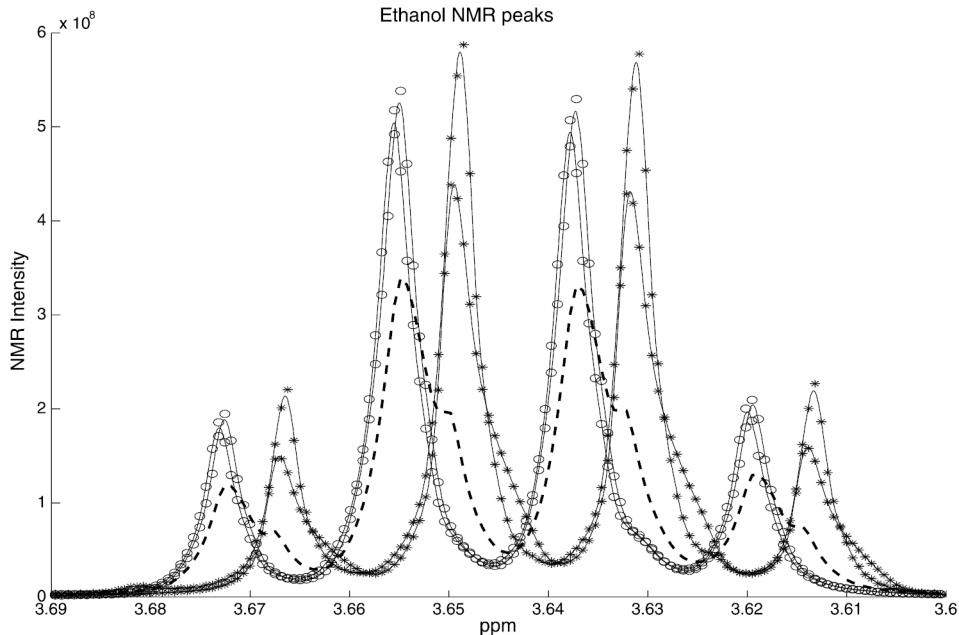


FIG. 1. Circles correspond to intensities over an ethanol region of the NMR spectrum for two typical red wines, and asterisks indicate a white and a rosé. The light solid lines are smooth fits of the data using order 6 B-spline basis functions with a knot at every sampling and a light penalty ($\lambda = 10^4$) on the fourth derivative. The heavy dashed line is the mean intensity across 31 reds, 7 whites, and 2 rosés. The mean has a far different shape from the quite similar shape of all the data curves, due to registration issues.

the continuous substrate of the data seems itself to be transformable, and that these transformations vary across functional observations.

Figure 1 displays four peaks for each of four samples of wines in the part of the nuclear magnetic resonance (NMR) spectrum corresponding to ethanol. Two of these wines are red, one is white, and one is a rosé. We notice that most of the variation across these four samples is due to the peaks of the white and rosé wines being displaced to the right relative to those for the red wines. It is known that the pH level in a solution has this effect on the location of the couplets, triplets, and m-tuplets that NMR generates; and also that red wines have pH's from 3.3 to 3.5, while white pH's are in the range 3.0–3.3. Moreover, the effects of pH and other factors are known to vary from one location in the spectrum to another, with displacements in opposing directions not being unusual.

The functional data analysis (FDA) literature refers to lateral displacements in curve features as *phase variation*, as opposed to *amplitude variation* in curve height. As in music, we imagine that time can be compressed or stretched over different intervals in a single performance. Consequently, we distinguish between measured clock time and related

but different time scales. Relative to human growth time, for example, puberty for girls occurs on average at the age of 11.7 years, but hormonal and other physiological factors shift this age forward and backward to the variable clock times that parents actually see.

Few time-varying events are more important than the weather. Figure 2 allows us to explore phase variation in Montreal's daily temperature variation over three winters, winter being the most dynamic period in the Canadian climate year. We see here several important markers of phase variation. There are two minimum temperatures in most winters, the first positioned around January 15 and the January thaw that separates them typically arrives on January 25. We notice, too, the increased volatility in temperature in the two months in the dead of winter. The two horizontal lines mark temperatures of great importance to Canada's economy. The five degree Celsius threshold is the point at which cash crops in the Canadian prairies germinate, and their total growth depends on, in addition to precipitation, the total number of degrees above this threshold prior to harvest. Minus seven degrees is the threshold below which ice has enough structural integrity to support winter river crossings and year-round ice




FIG. 2. Temperature variation in Montreal, Canada, over three winters. The solid curve is a smooth of the daily min/max averages, which are shown as dots. The dashed line is a strictly periodic smooth of the data over the years 1960 to 1994. The vertical dotted lines indicate the “orbital” year boundaries separated by 365.25 days. The upper dashed horizontal line is the temperature at which growth begins for most crops on the prairies; and the lower dashed line is the temperature below which ice is structurally sound. Note strong variation from year to year.

dams around tailing ponds for the many mines in the north. Global warming is altering the dates at which these thresholds are crossed. The small plateaus in the spring and fall mark out the arrival and departure of snow, respectively. We see that winter arrived in both 1988 and 1989 particularly early, and with an intense cold snap in 1989, while the 1987 winter was typical in its timing. Summer phase variation, by contrast, seems small. Predicting phase variation is of great importance in weather prediction, crop management, and far northern transportation.

Once recognized, one sees phase variation everywhere. Parents see children reaching puberty over a wide range of ages, and perhaps wonder if there is some connection between the timing of the pubertal growth spurt and adult height. Growth implies positive change, and Figure 3 displays the growth of ten girls in the Berkeley growth study (Tuddenham and Snyder, 1954) as the positive first derivative of height in the top panel, as well as the acceleration of height or derivative of growth in the bottom panel. Musicians alter the timing of notes in subtle ways to create tension and define mood, achieving in this way their unique auditory signature as performers. Golfers and baseball players, on the other hand, tend to find phase variation in their swings to be an impediment to fine control over am-

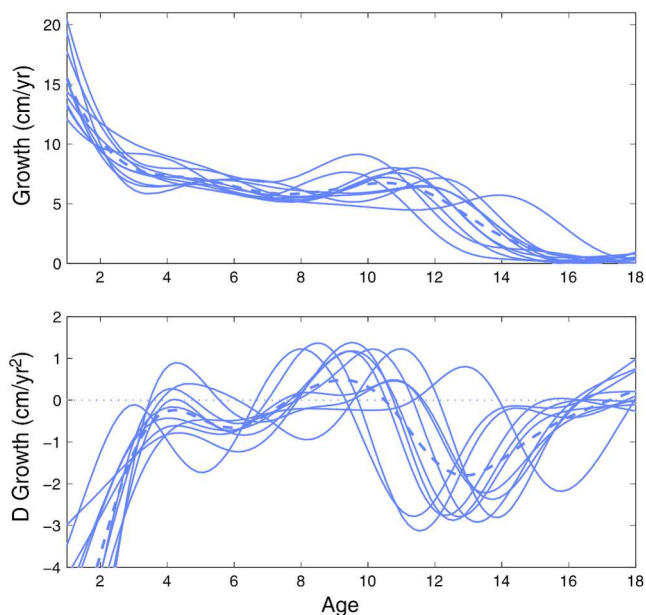


FIG. 3. The top panel plots the growth, understood as the first derivative of height, of ten girls, and the bottom panel contains the corresponding height-acceleration or growth-derivative curves. The dashed curve in both plots is the cross-sectional mean. Both these plots indicate both phase and amplitude variability.

plitude variation, and train to the point where it is nearly eliminated.

1.2 Clock Time, System Time, and the Time-Warping Function

We can articulate the concept of phase variation by distinguishing between *clock time* s and *system time* t . That is, we envisage the spectra of wines, the weather, and children as evolving over their respective continua at variable rates determined by processes that we may at least partially understand and would like to know more about. Consequently, when large-scale phase variation is compared to the clock time, defined these days in terms of the number of oscillations of the cesium atom, we envisage a functional relationship $s = h(t)$ that can vary from one wine type to another, over successive winters, and across children even within the same family. However, the system times are defined so that all girls will reach puberty at the same age.

In most cases, we can expect that the mapping h , often called the *time warping function*, will be smooth and strictly increasing, two properties captured in the term *diffeomorphism*. In other words, we require that the inverse function value $h^{-1}(s)$ exists everywhere in the support of the functional

data since we need to use $t = h^{-1}(s)$ to align a feature such as the pubertal growth spurt across multiple curves. As statisticians, we look for ways to estimate the h 's associated with different units of data distributed over the base continuum, as well as ways of using discrete and continuous covariate observations to explain and predict them.

Other conditions such as specified boundary behavior are added as makes sense for the context at hand. For example, the time taken to produce a sample of handwriting will vary from replication to replication, so that h_i may map, say, the interval $[0, T_0]$ into the interval $[0, T_i]$ where T_0 is a fixed template time. But if the observation is also supposed to reflect *when* the handwriting event took place, then simple shifts, $h_i(t) = t + \delta_i$, will provide a better model. If the process under study may reasonably be expected to have one or more derivatives, then the chain rule requires that h , too, be differentiable to the same extent. In any case, it seems unlikely that in many real-world applications the problem constraints will allow for sharp jumps in h , so that smoothness can be added to monotonicity as a property.

The following single-parameter expression for h mapping $[0, T]$ into itself serves as an illustration and is often useful: for $\beta \neq 0$,

$$(1) \quad h(t|\beta) = T \left[\frac{e^{\beta t} - 1}{e^{\beta T} - 1} \right] \quad \text{and}$$

$$h^{-1}(s|\beta) = \frac{1}{\beta} \log \left[\frac{s(e^{\beta T} - 1) + T}{T} \right].$$

The expression converges to the identity warp $h(t) = t$ as $\beta \rightarrow 0$. This model, taken from Kneip and Ramsay (2008), can also be derived from a later equation [equation (11)] by setting the function $W(t) = \beta t$.

Some warping functions corresponding to early and late growth spurts are shown in Figure 4. The warping function [of the type given in equation (1)] in each right panel maps pubertal growth spurt on the growth (system) time scale into the clock or observed time scale, as indicated by the zero crossing of the growth derivative function in the left panel, that is, the peak of the spurt, as shown by a circle. The early pubertal spurt in the top panel is modeled using an h which moves quickly through early growth phases relative to clock time (i.e., curves downward) so as to produce the early clock time of about nine years, whereas the bottom panel is better modeled

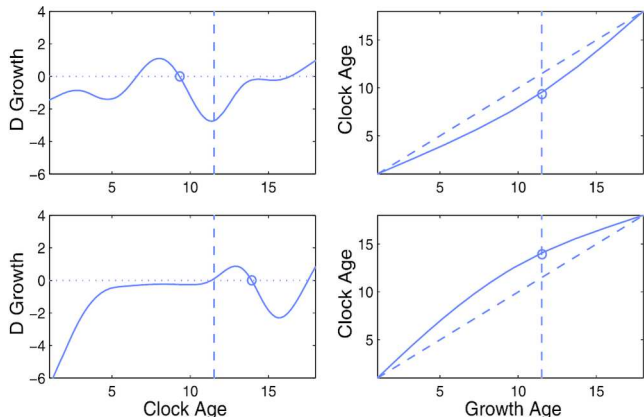


FIG. 4. The top left panel displays the derivative of growth for a girl with an early growth spurt, and the bottom left panel for a girl with a late growth spurt. The top right panel plots a warping function h that maps the growth time of the pubertal growth spurt, indicated by the circle, into the early clock time in the left panel. The bottom right panel shows the corresponding warping function for the late growth spurt. This shows how phase variation is effectively modeled by warping functions.

with an upward curving h that reflects slower transition through early growth phases to reach the clock time of the late growth spurt of about 14 years.

1.3 The Problems that Come with Ignoring Phase Variation

The presence of phase variation can play havoc with classical data analyses that are designed for data structures without phase changes. The heavy dashed line in Figure 1 is the average of the ethanol peaks across forty wine samples, of which 31 are red. The heights of the mean peaks are lower than almost all corresponding sample peaks, their widths are substantially wider, and no sample peak displays the step in the middle of the down-slope of each average peak. That is, a statistical analysis as elementary as averaging takes the data well outside of their normal modalities of variation, causing it to fail as an effective data summary. A recent review of chemometrics (Lavine and Workman, 2013) highlights the importance of aligning peaks in spectral data as a first step, and warns spectroscopists that getting this step right can be crucial to the quality of subsequent analyses. In fact, most familiar data analyses are found to fail in the presence of phase variation; variances are inflated, fits by regression models are degraded, and additional principal components are required.

This paper began as a follow-up to a workshop on curve registration at the Mathematical Biosciences

Institute at the Ohio State University in 2012 [see Marron et al. (2014) and companion papers]. An effective workshop raises many more questions than it answers, and this workshop left us with much to consider. Is there a clear distinction between amplitude and phase variation, or is there variation that can be represented either way? Can the transformation h be considered as a full data object, or does it just represent nuisance variation to be discarded once identified? When phase data objects are meaningful, how can we incorporate known covariates, such as pH in the NMR context, into the estimation? Are “features” in a curve or surface always things like peaks, points of inflection, and threshold crossings, or can models define more general properties that become invisible on the model side of the equation when phase is properly incorporated and estimated? Are traditional fitting criteria such as error sums of squares still useful, or are they only usable when there is no phase variation? What role should derivatives play? Can the warping function h be as complex as is required to align features, or is it wise to impose some regularity? When is it useful to develop data analyses that reveal aspects of the *joint* variation in phase and amplitude? We will discuss some of these questions in this paper.

Section 2 defines some possible goals for curve and surface alignment or registration, and discusses ways of understanding what is amplitude and phase variation. Section 3 considers various optimization strategies and statistical models that separate phase and amplitude variations. Section 4 provides some links for downloading relevant softwares. Section 5 considers what has been learned in working with these and other data sets, and looks forward to future research and generalizations in this fascinating area.

2. VIEWPOINTS AND GOALS

2.1 The Identification of Phase Variation

In this paper we will use y_1, y_2, \dots , to denote the observed functions with both phase and amplitude variability and x_1, x_2, \dots , to be the underling functions denoting only the amplitude variability, that is, after removing phase variability, such that $x_i(t) = y_i(h_i(t))$.

An important challenge is *identifiability* of amplitude and phase variation, since which is which is apt to depend very much on prior intuitions and knowledge about how each type of variation is caused. For example, while it may seem obvious that the peaks

after age eight in the top panel of Figure 3 exhibit phase variation, a close look at the lower panel shows that a number of the growth-derivative functions display more than one negative slope episode prior to the final crossing of zero. What we are tempted to call early spurts may only be due to the presence of a single pre-pubertal spurt, and a late spurt may be due to two or even more pre-pubertal spurts. This tends to sound more like an amplitude-oriented explanation.

A simple example of this is a data set of linear functions on \mathbb{R} , y_1, \dots, y_n , having the same slope, but differing intercepts. Using the notation $x(t) = y[h(t)]$, that mode of variation could be entirely modeled as linear shifts, $h_i(t) = a_i t + b_i$ constructed so that $x_1 = x_2 = \dots = x_n$ (i.e., all variation is in the phase variation), or it could equally well be modeled as $h_i(t) = t$, the identity warp, with all of the variation in the original data appearing in the intercepts of the y_i , or the variation could be split between these modes.

We have tied phase variation in the wine data to a known causal factor, the pH level of the wine, but, for the weather data, it seems to depend on intuition as to whether spring came late in a particular year or whether that year was simply unusually cold. Even an early velocity peak defining the pubertal growth spurt can be seen in part as a year of strong growth followed by a year of weaker growth. It is not surprising, as a consequence, that we see very little attention given to the phase variation in the evolution of statistical methodology. In particular, the distinction between phase and amplitude variation is generally not univocal, but instead depends on both the application under study and the goals of a particular analysis.

2.2 Types of Phase Variations

We have mentioned the linear shifts earlier, but there are several possibilities when choosing a class of warpings to specify phase variation. Depending on the application context, one may prefer one class over the others. We enumerate some possibilities below and illustrate an example of each in Figure 5:

- *Uniform Scaling*: Here the warping of the time domain simply rescales it by a positive constant $a \in \mathbb{R}_+$, that is, $h(t) = at$ for all $t \in \mathbb{R}_+$.
- *Uniform Shift*: In this case the time axis gets shifted by a constant $c \in \mathbb{R}$, that is, $h(t) = c + t$.

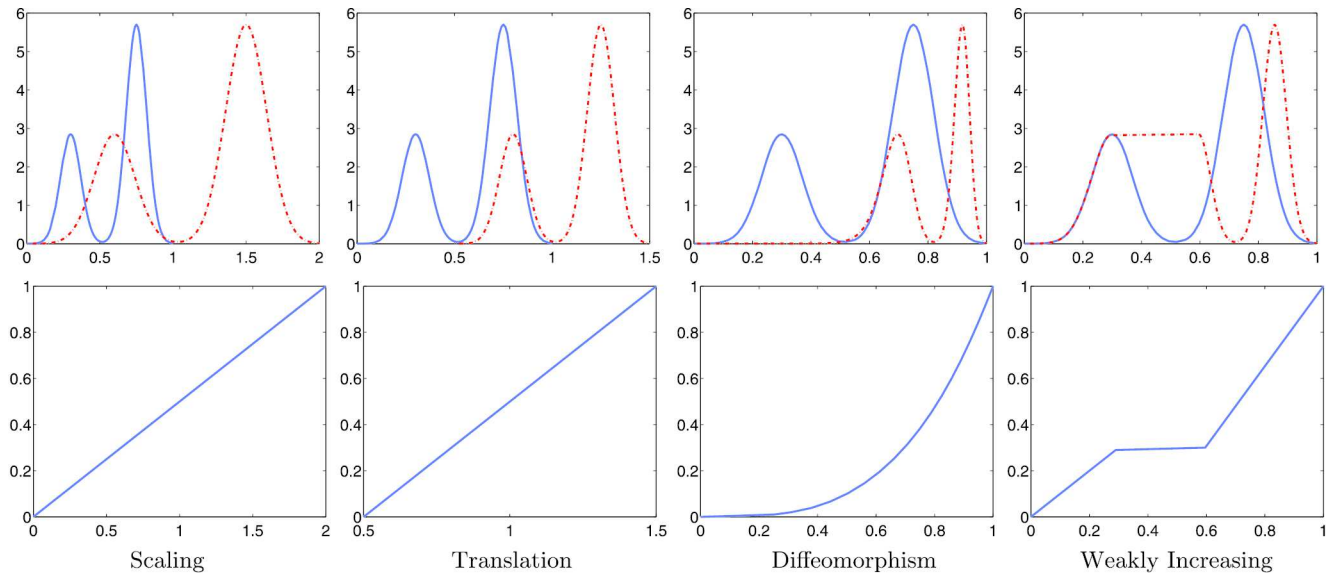


FIG. 5. Illustrations of different types of warping functions applied to the same function y . The top row shows $y(t)$ (solid line) and $y(h(t))$ (dashed line), and the bottom row shows the corresponding warping functions $h(t)$.

- *Linear or Affine Transform*: A combination of the previous two leads to a linear or affine transformation: $h(t) = c + at$, $a \in \mathbb{R}_+$ and $c \in \mathbb{R}$.
- *Diffeomorphisms*: A general class that includes domain warpings is given by the set of diffeomorphisms of the domain to itself. While it is possible to define diffeomorphisms on the full real line, practical considerations make it interesting to restrict warpings to compact intervals. The set of linear transformations is contained in the set of diffeomorphisms if the domain is defined to be the full real line.

While these are the main types of warping transformation, one can further enlarge the scope by including functions that allow for some flat regions; an example is shown in the rightmost column of Figure 5. Please refer to Srivastava et al. (2011a) for a discussion on the need for such functions and a rigorous approach to handling them.

2.3 Some Goals for an Amplitude/Phase Analysis

We can distinguish three motivations for a model that allows for phase variation. First, amplitude variation could be the main focus, with phase variation being a nuisance to be removed and then cast aside. The wine NMR spectra in Figure 1 illustrate this nicely, in part because the goal of the analysis is specifically to model the relative heights of the clearly visible peaks, the widths of which tend to be

proportional to their height. Prairie crop scientists tend to focus on the total heat and precipitation available to plants in the growing season as predictors of crop yield, leaving the issue of when the season starts and finishes to the producers to wrestle with. Auxologists, who study human growth, may be preoccupied by the variation in the shape characteristics of growth curves such as the variation in their amplitudes, and see the variation in the timings of the pubertal growth spurt as a nuisance to be eliminated by lining up the corresponding peaks.

On the other hand, phase variation could instead contain all of the interesting information, in contexts where issues such as timing are more important than relative peak heights, such as the locations of bursts in neuronal spike train data. Crop producers know that they have little control over heat and precipitation budgets, but they can look for indicators of when they can sow their seeds and when certain variants will mature. In this situation, the time warp functions are the center of attention.

Finally, both amplitude and phase variation, and in fact the *joint variation* between these, can be central issues in the analysis. It turns out, for example, that there is a simple relation between the strength of a pubertal growth spurt and its timing, namely, that early spurts are stronger and later ones are weaker, resulting in adult final heights that do not depend much on either factor. That is, it appears that each child has a wired-in capacity for growth,

but that the distribution of the expenditure of the growth energy over time can vary over children with similar growth capacities.

2.4 The Role of the Model in the Amplitude/Phase Partition

Assuming the relevance of phase variation, it will be clear that both its nature and estimation strategies will depend critically on the model being proposed for the data. The cross-sectional mean is often the model of choice in feature alignment strategies; peaks and threshold crossings are considered aligned when the mean curve is centrally located within the registered curves at all points over the interval of observation. More generally, the mean can be taken as one of many *template* or gold-standard curves to be approached as closely as possible in some sense by the application of phase transformations. Alternatively, as described in the next section, one can compute the mean under a different metric and use that as a model for alignment. Finally, functional linear equations, low-dimensional principal component representations, differential equations, and many other mathematical structures may provide model spaces for amplitude variation that, simultaneously, identify what is meant by phase variation. That is, if a diffeomorphic transformation of the substrate of the data, possibly within some predefined class, can improve the fit of the model to the data, we define it as phase. Models, of course, are usually chosen to represent a conjecture or hypothesis about what generates the data, and in this sense the identification of the amplitude/phase dichotomy is very much centered on the science underlying the application.

2.5 Amplitude/Phase Separation via Equivalence Classes

One way to study amplitude and phase variation is through equivalence classes. The use of equivalence classes is not new to statistics. In fact, they form the core idea in statistical shape analysis (Dryden and Mardia, 1998) and in Grenander’s work on pattern theory (Grenander, 1993), including its applications to computational anatomy (Grenander and Miller, 1998). In Kendall’s shape analysis the experimental units are configurations of (landmark) points in an appropriate space, usually two- or three-dimensional Euclidean space. To focus the analysis on the *shape* variation in the data, nonshape aspects, such as location, rotation,

and perhaps scaling, are incorporated into equivalence classes, where point configurations are identified with each other (i.e., called equivalent) when they can be translated, rotated, and scaled into each other. Then, one compares shapes of objects by comparing their equivalence classes. While the past shape approaches were restricted to point sets and simple transformations (rigid motions and global scales), the more recent literature has studied continuous curves with transformations that include time warpings (more precisely, reparameterizations) [see Younes et al. (2008) and Srivastava et al. (2011b), among others].

In an entirely parallel fashion, one can define amplitude and phase variability in functional data using equivalence classes. As laid out in Srivastava et al. (2011a), the main idea is to understand amplitude variation through a quantity that incorporates all aspects of phase variation inside it. This is done by defining an equivalence relation, where curves are identified or deemed equivalent when they can be time warped into each other. Figure 6 shows some elements of an equivalence class—a set of warps of a single three-peak curve. The equivalence class is actually much bigger, including all diffeomorphic time warps of this curve, only some of which are shown here. These equivalence classes are now taken as representing *amplitudes* because they model the essence of vertical variation in a simple and natural way. The phase variation is incorporated *within equivalence classes*, while the amplitude variation appears *across equivalence classes*. Further motivation for how equivalence classes provide clear definitions for separation of amplitude and phase is given in Section 3.4. (See also Vantini (2012).)

While the origins of these ideas lie in shape theory, an understanding of these concepts can also be ob-




FIG. 6. Different time warps of a function (left) form an equivalence class from the perspective of defining its amplitude.

tained using the terminology of object-oriented data analysis (OODA), as defined in Wang and Marron (2007) and more recently discussed in Marron and Alonso (2014). An important special case of OODA is FDA, where functions are the data objects. A natural approach to the decomposition of amplitude and phase variation is to model each with appropriate data objects, with specific goals as laid out in Section 2.3. In some situations, such as the wine NMR data in Figure 1, the phase variation can be viewed as a nuisance, so the data objects of interest are registered curves, that is, time warped to match their peaks. In other situations, for example, the temperature data shown in Figure 2 and for human growth curve data in Figure 4, interesting data objects can be any of the registered amplitude curves, or the transformations used to achieve registration (reflecting phase variation), or else the concatenation of both, for situations where joint amplitude–phase variation is key. In the same spirit, the data objects in an equivalence-class approach are the equivalence classes themselves.

3. SOME CURRENT CURVE REGISTRATION METHODS

In this section we look at a few curve registration techniques for estimating warping functions h . In the first two sections, the focus is on using a template function x_0 as a target, so that $y(s) \approx x_0[h(t)]$ and, inversely, $x_0(t) \approx y[h^{-1}(s)]$. We will see that the sense in which the approximation is defined requires considerable care, with least squares approximations computed in the usual way not being a viable candidate. The template x_0 is often defined using an objective function whose solution is iterative, starting with the cross-sectional mean and alternating between a registration and a recalculation of the cross-sectional mean of the registered functions. Typically this process, often referred to as *Procrustes iterations*, converges in only a few steps.

3.1 Dynamic Time Warping (DTW)

Before examining current registration methods, it is worthwhile mentioning dynamic time warping (DTW), an early registration method applied to discrete sequences of phonemes (a basic unit of language). Sakoe and Chiba (1978) devised an insertion/deletion algorithm that is rather like that of isotonic regression (Barlow et al., 1972). The underlying algorithm, which is a dynamic programming

algorithm, is an optimization technique where one partitions the graph space using a finite grid and the warping h is restricted to be a piecewise-linear function passing through the nodes of this grid. Depending on the context, one may allow it to have vertical jumps or be horizontal for multiple time-steps. In the classical DTW, the dynamic programming algorithm is applied to minimize the least-squares cost function given in equation (2). DTW can be effective as a feature alignment method, as it provides a globally optimal solution, albeit on the restricted search space (piecewise-linear h on a fixed grid). But the classical DTW has the conceptual problem that it may not provide smooth differentiable time warps that many applications require. Also, the computational algorithm can be greedy, in the sense of warping regions where no alignment seems called for. These problems, in general, can be handled by adding a regularization term to the cost function.

3.2 Landmark Registration

In terms of functional data alignment, we begin with the easiest situation in which each curve $y_i(s)$ has clearly-defined features, the timings of which can be used to estimate h_i at a series of points $(t_\ell, h_{i\ell})$. This requires, in turn, a consideration of what we might mean by “feature.”

In the case of the wine data, there seems to be little confusion. In most types of spectra, the presence of a chemical compound is marked by a single peak, the location of which is the desired landmark, and automatic methods for peak detection are relatively easy to devise. For multi-peak structures such as the NMR peaks in Figure 1, the average of the peak locations would serve the purpose. Alternatively, a template can be set up for a peak shape, and a peak detector can be devised by computing correlations with moving windows of the curve shape with the template pattern.

Let us suppose that there is a gold standard template spectrum x_0 with L peaks occurring at times $t_\ell, \ell = 0, \dots, L + 1$, where times t_0 and t_{L+1} are the endpoints of the observation interval. Then, for the i th spectrum with peak locations at s_ℓ , we can estimate h_i by interpolating in some suitable way the pairs (s_ℓ, t_ℓ) . Polygonal lines might serve, or it may be important to use a smoother interpolant having, perhaps, a specified number of derivatives. Figure 4 offers an elementary example of landmark registration, where the timing of a girl’s pubertal growth spurt is the single landmark t_i , shown as a circle in

Figure 4, and the intervals $(3, t_i)$ and $(t_i, 18)$ for the i th girl are interpolated by the warping functions [formed using the expression in equation (1)].

Peak and valley locations can be translated into crossings of zero in the curve’s first derivative with negative and positive slopes, respectively. Other types of crossings may also be important. For example, the heavy-duty winter in Figure 2 can be defined as the average of the first crossing time with negative slope for -7 deg C and the second crossing time with positive slope. Prairie farmers would prefer the crossing of germination threshold of 5 deg C with positive slope, and, in fact, do just that with daily soil temperature readings in May.

The problem with landmarks, of course, is that they are not always visible or one may be faced with other types of feature time ambiguity such as two or more closely spaced -7 deg C crossings in the temperature data. Moreover, recording landmarks by hand is tedious, and fail-safe automatic detectors are sometimes hard to set up. The choice of landmark can itself be open to the kind of debate that scientists would prefer to avoid. Finally, landmark registration is only discrete evidence concerning the intrinsically continuous function h_i , and as such ignores what happens in between landmarks, where there may reside additional information about h .

3.3 Registration Using \mathbb{L}^2 Distance and Correlational Criteria

Now we look at a classical approach to functional registration that does not require the use of landmarks. Let h_i denote the time warping associated with the i th data item y_i ; this h_i can be restricted to be an element of a parametric family, defined by the value of one or more parameters, or can be fully nonparametric as in a diffeomorphism. The one-parameter warps [equation (1)], along with simple shifts, scale changes, and linear functions of t , are examples of simple parametric warping families, and we will propose more flexible representations in Section 4. It seems natural to specify a loss function L that optimizes the congruence of a set of clock-time functions y_i to corresponding warped versions of a template x_0 , that is, $y_i \approx x_0 \circ h_i$, where $(x_0 \circ h_i)(t) = x_0(h_i(t))$.

The choice, however, of standard options such as

$$(2) \quad \begin{aligned} L(h; y_i, x_0) &= \|y_i - x_0 \circ h_i\|^2 \\ &= \int [y_i(t) - x_0(h_i(t))]^2 dt \end{aligned}$$




FIG. 7. The upper panels show a Gaussian density function x_0 and its scaled version y , as dot-dashed and dotted curves, respectively. The solid curve in the upper left panel results from minimizing the squared error criterion $\int [y(t) - (x_0 \circ h)(t)]^2 dt$ with the optimal warping function h shown in the lower left panel. The solid curve in the upper right panel results from minimizing the squared error criterion $\int [(y \circ h)(s) - x_0(s)]^2 ds$ with the optimal warping function h shown in the lower right panel.

will quickly prove disappointing if combined with a flexible class of warping functions, as Figure 7 demonstrates. In the left case, the minimization of the \mathbb{L}^2 norm results in a reduction from 0.500 to 0.024, using a piecewise-linear warping and a spike that nearly eliminates the area under the registered curve corresponding to intervals where the y has larger amplitude than x_0 . In the registration process the amplitude characteristics of y have been significantly distorted.

This pinching effect can be mitigated by using warping functions that are constrained to be smooth, either by the use of a regularization strategy or by the use of a small number of basis functions. The registration procedures proposed in Ramsay and Silverman (2005), for instance, incorporate a penalization term that forces the choice of the warping functions toward functions that do not differ significantly from the identity (corresponding to the case of no registration) or from constant functions. Concerning instead the use of simple parametric families for the class of warping functions, the \mathbb{L}^2 distance will work just fine for the one-parameter shift-warp family, $h(t) = t + \delta$. Such a registration procedure performs perfectly for the example in Figure 7, where the identity warp is returned since the two peaked curves are already registered.

It thus appears fundamental to appropriately relate the definitions of amplitude variation and of

phase variation, that are jointly described by the loss function to be optimized and the class of warping functions. This motivates the simultaneous definition of phase and amplitude to avoid issues such as the one highlighted in Figure 7. For instance, the loss function L to be optimized and the class of warping functions h may be chosen so that for any two functions x_1, x_2 , and any warping function h , L satisfies the relation

$$(3) \quad L(x_1, x_2) = L(x_1 \circ h, x_2 \circ h).$$

This invariance property guarantees that it is not possible to obtain a fictitious increment of the similarity between two functional data by simply warping them simultaneously with the same warping function, and has been clarified in the context of different types of warpings in different papers. For example, Sangalli et al. (2009, 2010) and Vantini (2012) study this invariance and then specify it in the context of linear or affine transformations of the domain, while Srivastava et al. (2011a) study it for diffeomorphisms.

Moreover, as already highlighted, the concepts of amplitude variation and of phase variation are problem-specific and depend on the application goals. For instance, if two functional data x_1 and x_2 may be considered aligned when they are proportional, that is, when $x_1 = \alpha x_2$, then it is natural to use the loss function associated to the semi-norm

$$(4) \quad \left\| \frac{x_1}{\|x_1\|} - \frac{x_2}{\|x_2\|} \right\|,$$

and the corresponding correlation measure

$$(5) \quad \rho(x_1, x_2) = \frac{\langle x_1, x_2 \rangle}{\sqrt{\langle x_1, x_2 \rangle \langle x_2, x_2 \rangle}}.$$

The class of linear warping functions $h(t) = \delta + \gamma t$ is compatible, in the sense of equation (3), with the loss function associated to (4) and (5), if these are computed over the full real line. This definition of amplitude/phase variation seems, for instance, well suited for the wine data, where the amplitude variation is well described by the relative heights of the peaks, rather than by their absolute heights, and where linear transformations of the abscissa allows for a good alignment of these peaks. This also holds for growth curve data, where the emphasis is on growth velocities, rather than on the height curves per se, and the children's biological clocks, with their pubertal spurts, are aligned by aiming at

proportional growth velocities. Ramsay and Silverman (2005) used the size of the minimum eigenvalue of the order two cross-product matrix

$$(6) \quad L(h; y, x_0) = \begin{bmatrix} \int \{x_0(t)\}^2 dt & \int x_0(t)y[h(t)] dt \\ \int x_0(t)y[h(t)] dt & \int \{y[h(t)]\}^2 dt \end{bmatrix}.$$

The minimum eigenvalue criterion essentially measures the linearity of the relationship between x_0 and $y \circ h^{-1}$, and is the same thing as maximizing the correlation [equation (5)] between the two functions or, correspondingly, minimizing [equation (4)]. In other contexts, two functional data x_1 and x_2 may be considered aligned when their first derivatives are proportional, that is, $Dx_1 = \alpha Dx_2$ and, equivalently, $x_1 = \alpha x_2 + \beta$. Then it is natural to use the loss function in equations (4) and (5), but applied to the first derivative instead. Also, in this case, if the loss function is computed over the full real line, then it is compatible in the sense of equation (3) with the class of linear warping functions $h(t) = \delta + \gamma t$. And the same can of course be said for the \mathbb{L}^2 distance (2) with the shift-warp family $h(t) = \delta + t$. Sangalli, Secchi and Vantini (2014) report other examples of loss-functions/class of warping functions, that define concepts of amplitude and phase variations that are appropriate in different applications. In practice, the functional data are only available on bounded intervals, that possibly differ from curve to curve. These loss functions can then be computed over the intersection of the domains of the two functional data. In the case of the \mathbb{L}^2 distance, normalizing the distance by the length of the domain intersection helps avoiding fictitious decrements of the distance as the intersection becomes smaller.

It is also possible to consider much more flexible representations of phase variation and still define loss functions and class of warping functions satisfying the property (3). Section 3.4 is devoted to the case where the phase variation is described by arbitrary diffeomorphic transformations.

3.4 The Square-Root-Velocity Function and the Fisher–Rao Metric

Standard fitting criteria such as least squares may also be applied to transformations of the functional objects, most commonly first and second derivatives or their combinations. However, one can go even

further by choosing newer metrics that are compatible with the notion of equivalence classes mentioned earlier in Section 2.4. Application of the concept of equivalence classes as data objects in FDA needs some rethinking of important concepts. First off, the classical notion of metrics on curves needs to be extended to metrics on equivalence classes. Some consideration of this point highlights the challenges faced by classical approaches in analyzing vertical and horizontal curve variation. For example, as mentioned in the previous section, a common approach to quantifying the *vertical distance* between curves y_1 and y_2 is through \mathbb{L}^2 norm between y_1 and warped y_2 , that is, $\inf_h \|y_1 - y_2 \circ h\|_2$. From a theoretical perspective this quantity has several problems: it is not symmetric and does not satisfy the triangle inequality. Moreover, from a conceptual perspective, there are problems with this formulation, as shown in Figure 8 [constructed by Lu and Marron (2013)]. The top left panel of Figure 8 shows a toy example, using two single step functions as y_1 and y_2 . One naive approach to aligning these curves is to register y_2 to y_1 using the simple piecewise-linear warp h_2 shown in the top right panel. The result of this is a reasonable alignment shown in the top center panel. But an equally good approach to aligning these curves is to warp y_1 into y_2 , using the alternate piecewise-linear warp h_1 shown in the bottom center panel. As shown in the bottom left, this also gives a high quality of alignment. The challenge in classical approaches is what should be taken as the vertical distance between y_2 between y_1 ? The (appropriately squared, etc.) region between the aligned

curves (representing the \mathbb{L}^2 norm) in the top center panel is clearly very different from that in the bottom left panel. Now if we allow warping of both y_1 and y_2 , then many other appealing registrations could be found, for example, that in the bottom right panel, all of which are quite reasonable. A big payoff of the idea of equivalence classes as data objects is that it allows a very simple and natural metric, which essentially includes all of these reasonable alignments in its formulation.

The core idea is to choose a metric that helps compare equivalence classes, and not just individual functions, since these classes provide an identifiable representation of amplitude variability in this setting. This is done in a straightforward way, by starting with a curve metric that is invariant to identical warping of its two arguments, as in equation (3). That is, it should satisfy

$$(7) \quad d(x_1, x_2) = d(x_1 \circ h, x_2 \circ h),$$

for all warpings h . This is a particularization of equation (3) where a general loss function is replaced by a distance function. Srivastava et al. (2011a) used the nonparametric form of the Fisher–Rao metric [see Srivastava, Jermyn and Joshi (2007) for a short introduction to this metric] for this purpose. In fact, since the original Fisher–Rao metric was defined only for positive probability densities, they extended this notion to include a larger class of functions. The actual expression for this metric is complicated and thus is not discussed in detail here, except we note that the resulting Fisher–Rao distance, denoted by d_{FR} , satisfies the property stated in equation (7).

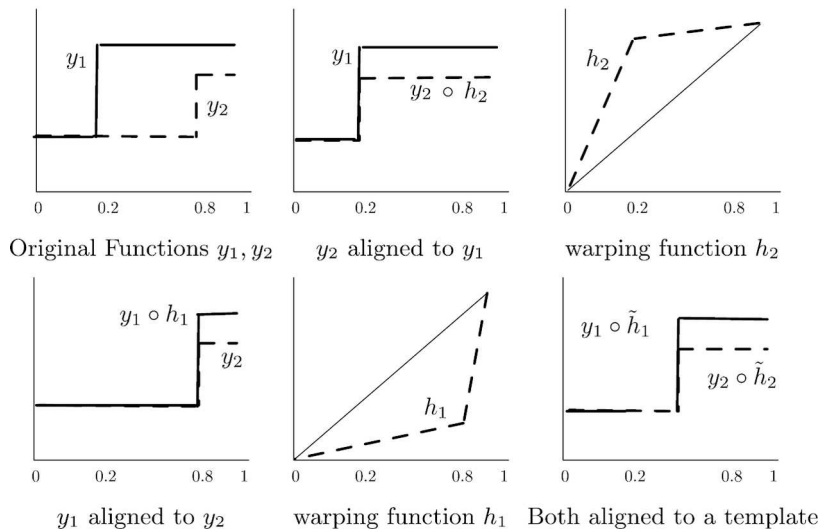


FIG. 8. Toy example showing asymmetry of the \mathbb{L}^2 norm naively applied to curve registration.

The key step in this formulation is to define a *square root velocity function* (SRVF) transform,

$$(8) \quad \text{SRVF}(x) = \text{sgn}(Dx) \sqrt{|Dx|},$$

where $\text{sgn}(u) = +1$ if $u \geq 0$ and -1 if $u < 0$ and Dx is the first derivative of x . It should be noted that SRVF is a one-to-one map up to a translation. That is, if $x(0)$ is known, then one can calculate x back uniquely from its SRVF. The SRVF transforms of the ten growth curves in Figure 3 are shown in Figure 9. In this particular case, since the x is defined to be the derivative of growth, SRVF(x) refers to the acceleration curves shown in the bottom panel of Figure 3. Consequently, the SRVF curves cross the zero axis at the same locations, but now with very steep slope.

The main reason for introducing SRVF is that the Fisher–Rao distance between any two functions is given by the \mathbb{L}^2 distance between their SRVFs, that is,

$$(9) \quad d_{\text{FR}}(x_1, x_2) = \|\text{SRVF}(x_1) - \text{SRVF}(x_2)\|.$$

We refer the reader to Srivastava et al. (2011a) for the details, but mention in passing that the proof hinges on the fact that $\text{SRVF}(x \circ h) = (q \circ h) \sqrt{Dh}$, where $q = \text{SRVF}(x)$.

This nice mathematical structure leads to *formal* definitions of amplitude and phase in functional data. For any two functions, x_1 and x_2 , the actual registration problem is given by

$$(10) \quad \begin{aligned} & \inf_h \|\text{SRVF}(x_1) - \text{SRVF}(x_2 \circ h)\| \\ &= \inf_h \|\text{SRVF}(x_1 \circ h) - \text{SRVF}(x_2)\|. \end{aligned}$$




FIG. 9. The signed root velocity transforms of the ten female growth curves displayed in Figure 3.

This formulation avoids the issue discussed in the example associated with Figure 8. It is important to note that $\text{SRVF}(x \circ h) \neq (\text{SRVF}(x) \circ h)$ and, therefore, this alignment is NOT simply a least-square alignment of SRVFs. The infimum value in equation (10) represents a comparison of the amplitudes of x_1 and x_2 and is actually a distance between the equivalence classes discussed in Section 2.5. If the optimal h on the left-hand side is invertible, then its inverse is also the optimal for the right-hand side of that equation. This has been called *inverse consistency* in the image-processing literature. The optimal h denotes the (relative) phase between x_1 and x_2 . The actual optimization over h in equation (10) can be performed in many ways, depending on the problem. If h takes a nonparametric form, a diffeomorphism of the domain, then the dynamic programming algorithm mentioned earlier is applicable. If some application calls for smooth phases, then some common smoothing idea—either restrict to a parametric family or apply a regularization penalty—can be applied, both at a loss of some mathematical structure. We emphasize that while some applications favor smooth solutions for warpings, some others, such as activity recognition in computer vision, naturally favor warping functions that are close to being vertical or horizontal over subdomains.

3.5 Representations of the Warping Function

The nature of warping functions leads to some interesting representations. The evolution of time, whether clock or system, is fundamentally a growth process, and as such, like height, has a positive first derivative. Two transformations of h play a number of useful roles in the representation and study of phase variation. Using the notation Dh for the derivative of h , the *log-derivative transformation* and its inverse

$$(11) \quad \begin{aligned} h(t) &= C_0 + C_1 \int_0^t \exp W(v) dv, \quad C_1 > 0 \quad \text{and} \\ (\log D)h(t) - \log C_1 &= W(t) \end{aligned}$$

allow us to represent any diffeomorphism h in terms of the unconstrained log-derivative function W . A natural and effective method of computing h is to use numerical differential equation solver methods to approximate the solution of the linear forced differential equation $Ds = \exp[W(t)]$ using the initial value $h_0 = C_0$. Moreover, from the equation

$h^{-1}[h(t)] = t$ the solution of the complementary non-linear unforced equation $Dt = \exp[-W(t)]$ defines the inverse of the warping function.

Since the log-derivative W is unconstrained and defined over a closed interval, it is natural to use a basis function expansion, with the B-spline basis being the likely choice. In particular, the one-parameter model [equation (1)] corresponds to $W(t) = \beta t$. The overall smoothness of h can be controlled either by the number of basis functions used or by appending a roughness penalty to a fitting criterion. It is essential that any representation be expandable to include contributions from one or more covariates z_j known or conjectured to modulate phase. For example, it is well known in climate modeling that proximity to oceans retards the seasons by two to three weeks, so that a model for phase variation across weather stations would include this factor. Because of global warming, long-term time itself is an important modifier of climate variables such as seasonal temperature and precipitation. Covariates can be easily incorporated by extending W to be a function of a covariate such as $W(t + \alpha z)$ for $W(t, z)$.

Another mathematical representation for warping functions comes from the SRVF idea. Since Dh is assumed to be positive, one can also use the positive square root $\psi(t) = \sqrt{Dh(t)}$ as a representation of h . Just like W earlier, one can use a basis expansion to express ψ if h is not constrained any further. However, if h represents a time warping of a fixed interval, for instance, $[0, 1]$, to itself, then that imposes an additional constraint on h . In order to obtain the boundary conditions $h(0) = 0$ and $h(1) = 1$, we require that $\int_0^1 \psi(t)^2 dt = 1$ or the \mathbb{L}^2 norm of ψ is one. This is an interesting geometric structure—the space of allowed ψ functions is a unit sphere and its geometry can be exploited in the ensuing analysis. The spherical geometry of this space of ψ functions has been used to perform estimation and alignment of curves in several places, including Veeraraghavan et al. (2009). This geometry has also been helpful in developing PCA of warping functions [Tucker, Wu and Srivastava (2013)] and in alternatives to PCA in the form of principal nested spheres [Jung, Dryden and Marron (2012)].

3.6 Registering Curves to Models

So far we have focused on pairwise registration of functions, but the alignment of multiple functions is often more of concern in analyzing real data. While

some methods for multiple alignment are simple extensions of the binary case, the others take a completely fresh approach and derive models tailored to such function data objects. The former approach is generally based on constructing a *template* of some kind and then registering individual functions to this template. This template may be constructed in an iterative fashion, as recursive improvements in alignments improve the resulting template, and vice-versa.

A simple idea for constructing a template is the cross-sectional mean, as mentioned earlier. At each iteration, one takes the currently aligned functions $\{x_i \circ h_i\}$ and computes their cross-sectional mean to update the template $x_0 = \frac{1}{n} \sum x_i \circ h_i$. (The cross-sectional mean is, of course, the mean of functional objects under the \mathbb{L}^2 metric.) Then, one by one, the given functions are aligned to this template to update h_i 's: $h_i = \operatorname{argmin}_h L(h; x_i, x_0)$. Depending on the nature of data, the results of this process may be sensitive to the initial conditions.

The same idea can be generalized to situations where a metric different from the \mathbb{L}^2 metric is used. In the case where equivalence classes of functions are data objects, one can compute the average of the corresponding equivalence classes $[x_1], \dots, [x_n]$, using the notion of a Karcher or Fréchet mean. This can be done under the Fisher–Rao distance mentioned in the previous section. The template is then taken to be the *center* of the Karcher mean equivalence class, chosen so that the average of the phases of x_1, \dots, x_n , with respect to this center, is the identity h_{id} . For further details of this construction and an algorithm for computing the center of an orbit, please refer to Srivastava et al. (2011a).

Shown in Figure 10 is an example of alignment using the SRVF framework applied to the wine NMR spectra shown earlier. The top row shows the original spectra, the aligned spectra, and the phase functions obtained during the alignment. The bottom row of Figure 10 shows the same data aligned using simple shifts and minimizing the loss in (4), using as a template one of the curves in the sample, the medoid curve, as detailed, for example, in Sangalli, Secchi and Vantini (2014). For these data the amplitude variation is in fact well described by the relative heights of the peaks and the shifts-warp family is able to capture very well the phase variation in the part of the spectra here considered, as also highlighted by the SRVF framework. The associated shifts display a clear clustering in the phase

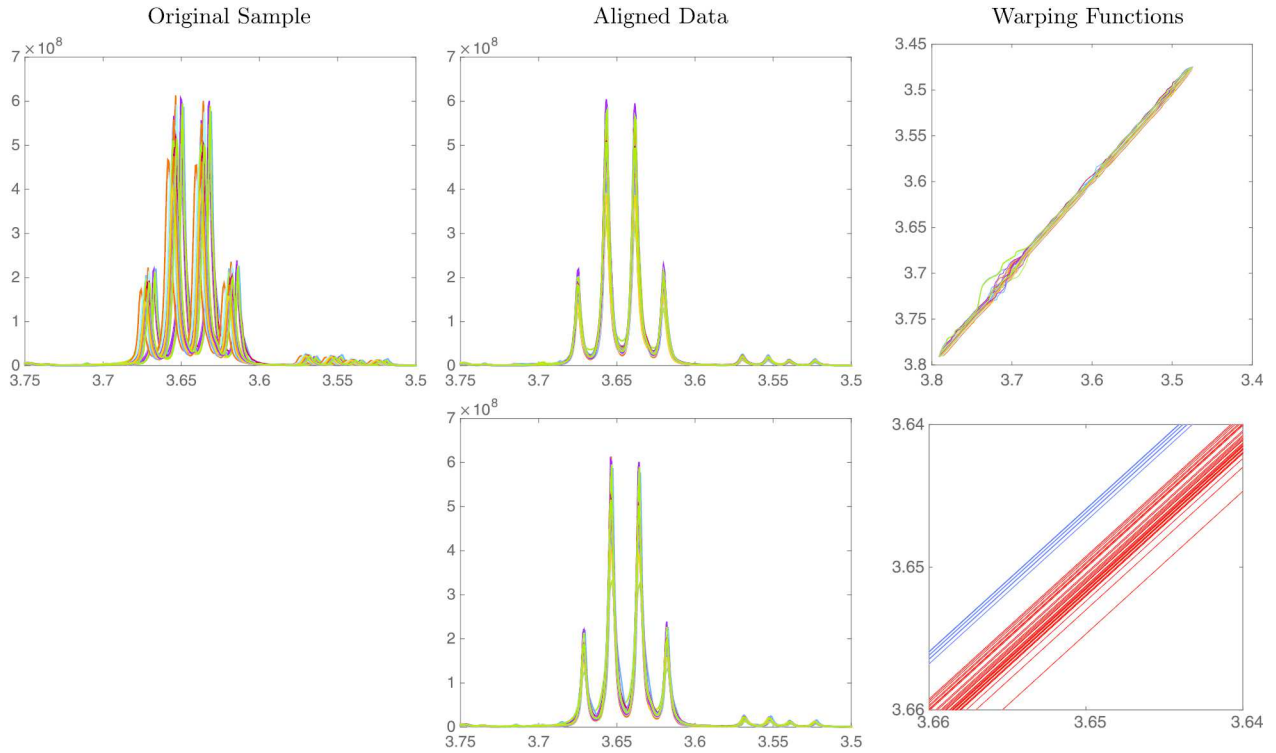


FIG. 10. Alignment of a part of 40 wine NMR spectra shown earlier. Top row: using the SRVF framework. Bottom row: using shifts and minimizing the loss in (4). A zoom of the warping functions, displayed on the bottom right panel, shows a neat separation in the phase between the red wines (warping functions colored in red) and the white and rosé wines (warping functions colored in blue).

of the red wines vs the white and rosé wines. Figure 11 shows the alignment of the growth velocities of the 54 girls in the Berkeley growth study, in the time interval 3 to 18 years—the top row displays the results obtained via the SRVF framework, while the bottom row displays the results obtained on additionally smoothed data using linear warpings and minimizing the loss in (4). The nonlinear warping in the SRVF framework allows for a visibly better alignment of the growth curves, showing that in many applicative contexts nonlinear warping is indeed necessary. The linear warping, obtained after additional smoothing of data, is nevertheless able to unveil some interesting features of the data. For instance, Sangalli et al. (2010) carry out linear warping of the growth curves of both the girls and boys in the study, highlighting a neat separation of boys and girls in the phase space and other interesting aspects of the growth dynamics of the two groups.

Instead of using just one template, it is often beneficial to divide data into smaller sets and use different templates for alignment in these subsets. An instance of this idea is when clustering and alignment are performed together. For instance, Sangalli

et al. (2010) propose a k -mean alignment procedure that jointly performs alignment and (unsupervised) clustering of functional data. Other proposals in this context are given by Tang and Müller (2009), Liu and Yang (2009), Boudaoud, Rix and Meste (2010). Another set of papers (Tang and Müller (2008), Liu and Müller (2004), Gervini and Gasser (2004)) takes the approach where some data points serve as templates for others, and the individual warping functions are averaged to find ultimate warpings.

Kneip and Ramsay (2008) perform registration of functional observations to the fits provided by a K -dimensional principal components analysis. In other words, the template is constructed individually for each function using an orthonormal basis. As an illustration, consider the 15 sections of mean-centered log-transformed mass-spectrometry intensities in the top panel of Figure 12. The large peaks on the right are fairly well registered by a preliminary landmark registration of the whole sequence, but we see substantial phase variation in the rest of these spectrum sections that obscures important amplitude variation. Three principal components were computed from these data combined

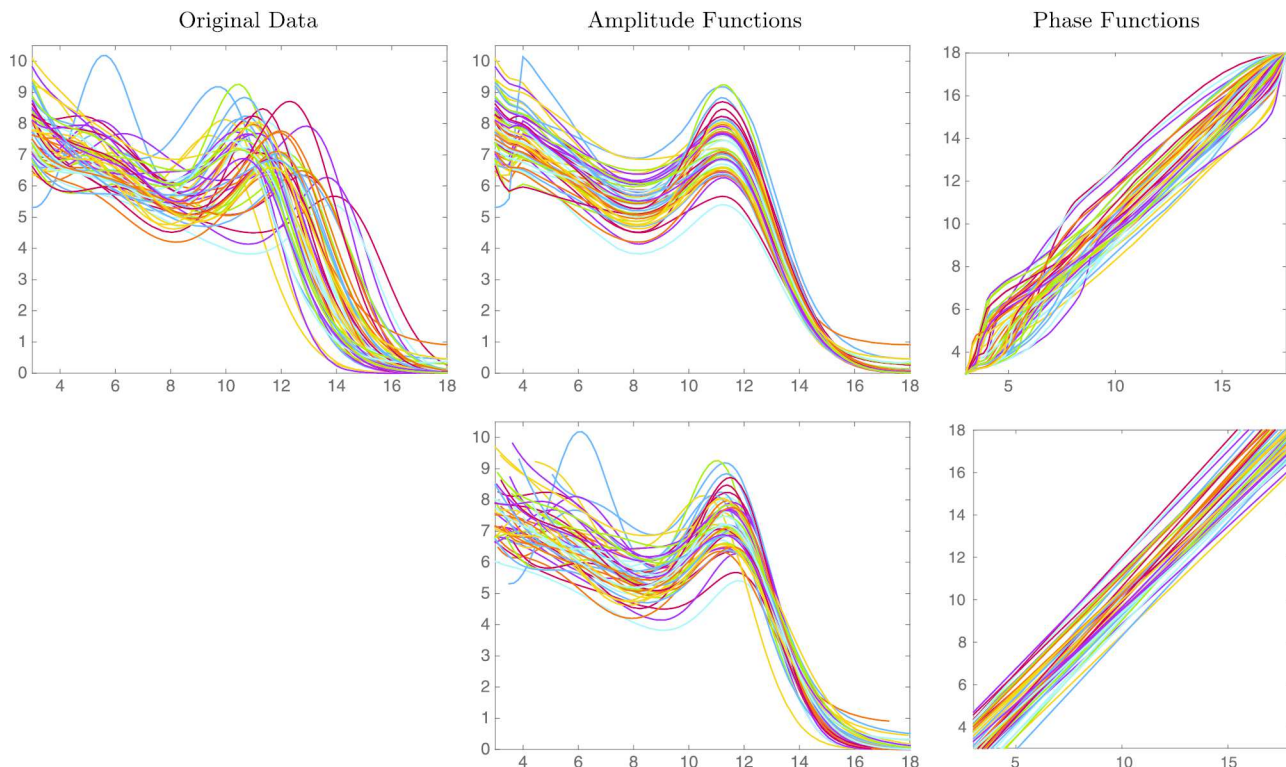


FIG. 11. Alignment of the growth velocities of the 54 girls in the Berkeley growth study using the SRVF framework (top) with the original data and the linear alignment (bottom).

with a registration of each section y_i to its fit \hat{y}_i using a method currently under development, as well as a principal components analysis without registration. The mean squared residuals for unregistered and registered PCA's were 0.0052 and 0.0038, respectively, corresponding to a squared multiple correlation 0.47. That means nearly half of the variation around the unregistered fit can be accommodated by modeling phase variation. The bottom part of Figure 12 displays the fits for y_1 and y_{15} in its left panels, along with the deformations $d_i(t) = h_i(t) - t$ associated with the registration in the right panels. The PCA is able to nicely accommodate the amplitude variation, and its fits after time warping are well aligned with all of the peaks. Choice of the number of components has an important impact on this type of analysis. Combining registration with model estimation or using multiple templates further blurs the distinction between amplitude and phase variation, suggesting that a successful analysis may depend heavily on prior choices guided by knowledge and intuitions about which type of variation is the primary focus.

In some contexts it also makes sense to combine the registration problem with other inferences,

such as a regression problem, for a more comprehensive solution. For instance, Hadjipantelis et al. (2015, 2014) study the problem of regression using phase and amplitude components of the given functions.

4. AVAILABLE SOFTWARE

Software implementations of many of the methods illustrated here are available publicly. R and Matlab code for implementation of the minimum eigenvalue method of Ramsay and Silverman can be found at <http://www.psych.mcgill.ca/misc/fda/downloads/FDAfuns/>. Matlab software for the extended Fisher-Rao SRVF approach of Srivastava et al. (2011a) is available at <http://ssang.stat.fsu.edu/software> and the R package is available from CRAN under *fdasrvf*. The R package *fdakma* (Parodi et al., 2014) implementing the k-mean alignment procedure described in Sangalli et al. (2010) is available from CRAN.

5. DISCUSSION AND CONCLUSIONS

In this paper we highlight the concept of phase variability that is present in functional data and the pitfalls of ignoring it in statistical analysis. After




FIG. 12. *Top: Fifteen mean-centered \log_{10} -transformations of sections of mass spectrometry analyses of blood samples. Bottom left: The fits (dotted curve) to the data (solid curve) for two observations, y_1 and y_{15} , produced by three principal components and registration. Bottom right: The corresponding warping functions display using $d_i(t) = h_i(t) - t$.*

motivating the importance of phase–amplitude separation, or alignment of functional data, in statistical analyses we proceed to summarize different ideas present in the literature for accomplishing this task. Specifically, we describe the problem of pinching associated with the classical L^2 -norm-based matching, and present several solutions to avoid this problem. These solutions involve either restricting the amount of warping or using an alternative metric to perform matching.

We note that while several methods exist for phase–amplitude separation, this is not a completely solved problem and forms an active area of research. A major challenge comes from the lack of a single mathematical definition or algorithm that can work in all, or even most, applications and contexts. For instance, one can argue that the goals of warping in weather data will be different from that in wine spectra. Similarly, while in some cases a simple translation and scaling may be sufficient for alignment

of curves, the other cases require genuine nonlinear warpings for proper alignment. In some cases effective data analysis is done by seeking the best possible peak/valley alignment, for example, in spectral data. In those cases the Fisher–Rao method is the most effective that we have seen so far. However, in other cases too much peak alignment can be a distraction, for example, the growth curve data. Therefore, it seems more natural to tailor objective functions and algorithms to the problem area.

Although we have focused on phase–amplitude separation of real-valued functions in this paper, this problem is prevalent in several other data object contexts. For instance, the problem of registration of images is considered a central issue in medical image registration. See Sotiras, Davatzikos and Paragios (2013) for a recent survey of warping-based techniques in this problem area. The ideas presented in this paper can be extended to included higher-dimensional signals such as images.

ACKNOWLEDGMENTS

The authors thank the Mathematical Biosciences Institute at the Ohio State University, Columbus, OH, for its support in organizing a workshop which led to the writing of this paper.

REFERENCES

- BARLOW, R. E., BARTHOLOMEW, D. J., BREMNER, J. M. and BRUNK, H. D. (1972). *Statistical Inference Under Order Restrictions. The Theory and Application of Isotonic Regression*. Wiley, New York. [MR0326887](#)
- BOUDAUD, S., RIX, H. and MESTE, O. (2010). Core shape modelling of a set of curves. *Comput. Statist. Data Anal.* **54** 308–325. [MR2756428](#)
- DRYDEN, I. L. and MARDIA, K. V. (1998). *Statistical Shape Analysis*. Wiley, Chichester. [MR1646114](#)
- GERVINI, D. and GASSER, T. (2004). Self-modelling warping functions. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **66** 959–971. [MR2102475](#)
- GRENANDER, U. (1993). *General Pattern Theory*. Oxford Univ. Press, New York. [MR1270904](#)
- GRENANDER, U. and MILLER, M. I. (1998). Computational anatomy: An emerging discipline. *Quart. Appl. Math.* **56** 617–694. [MR1668732](#)
- HADJIPANTELOS, P. Z., ASTON, J. A. D., MÜLLER, H. G. and EVANS, J. P. (2015). Unifying amplitude and phase analysis: A compositional data approach to functional multivariate mixed-effects modeling of Mandarin Chinese. *J. Amer. Statist. Assoc.* **110** 545–559. [MR3367246](#)
- HADJIPANTELOS, P. Z., ASTON, J. A. D., MÜLLER, H.-G. and MORIARTY, J. (2014). Analysis of spike train data: A multivariate mixed effects model for phase and amplitude. *Electron. J. Stat.* **8** 1797–1807. [MR3273597](#)
- JUNG, S., DRYDEN, I. L. and MARRON, J. S. (2012). Analysis of principal nested spheres. *Biometrika* **99** 551–568. [MR2966769](#)
- KNEIP, A. and RAMSAY, J. O. (2008). Combining registration and fitting for functional models. *J. Amer. Statist. Assoc.* **103** 1155–1165. [MR2528838](#)
- LAVINE, B. K. and WORKMAN, J. J. (2013). Chemometrics. *Anal. Chem.* **85** 705–714.
- LIU, X. and MÜLLER, H.-G. (2004). Functional convex averaging and synchronization for time-warped random curves. *J. Amer. Statist. Assoc.* **99** 687–699. [MR2090903](#)
- LIU, X. and YANG, M. C. K. (2009). Simultaneous curve registration and clustering for functional data. *Comput. Statist. Data Anal.* **53** 1361–1376. [MR2657097](#)
- LU, X. and MARRON, J. S. (2013). Principal nested spheres for time warped functional data analysis. Preprint. Available at [arXiv:1304.6789](#).
- MARRON, J. S. and ALONSO, A. M. (2014). Overview of object oriented data analysis. *Biom. J.* **56** 732–753. [MR3258083](#)
- MARRON, J. S., RAMSAY, J. O., SANGALLI, L. M. and SRIVASTAVA, A. (2014). Statistics of time warpings and phase variations. *Electron. J. Stat.* **8** 1697–1702. [MR3273584](#)
- PARODI, A., PATRIARCA, M., SANGALLI, L., SECCHI, P., VANTINI, S. and VITELLI, V. (2014). fdakma: Functional data analysis: K-mean alignment. R package version 1.1.1.
- RAMSAY, J. O. and SILVERMAN, B. W. (2005). *Functional Data Analysis*, 2nd ed. Springer, New York. [MR2168993](#)
- SAKOE, H. and CHIBA, S. (1978). Dynamic programming algorithm optimization for spoken word recognition. *IEEE Trans. Acoust. Speech Signal Process.* **26** 43–49.
- SANGALLI, L. M., SECCHI, P. and VANTINI, S. (2014). Analysis of AneuRisk65 data: k -mean alignment. *Electron. J. Stat.* **8** 1891–1904. [MR3273609](#)
- SANGALLI, L. M., SECCHI, P., VANTINI, S. and VENEZIANI, A. (2009). A case study in exploratory functional data analysis: Geometrical features of the internal carotid artery. *J. Amer. Statist. Assoc.* **104** 37–48. [MR2663032](#)
- SANGALLI, L. M., SECCHI, P., VANTINI, S. and VITELLI, V. (2010). k -mean alignment for curve clustering. *Comput. Statist. Data Anal.* **54** 1219–1233. [MR2600827](#)
- SOTIRAS, A., DAVATZIKOS, C. and PARAGIOS, N. (2013). Deformable medical image registration: A survey. *IEEE Trans. Med. Imag.* **32** 1153–1190.
- SRIVASTAVA, A., JERMYN, I. and JOSHI, S. H. (2007). Riemannian analysis of probability density functions with applications in vision. In *IEEE Conference on Computer Vision and Pattern Recognition, 2007. CVPR '07* 1–8. Minneapolis, MN, USA.
- SRIVASTAVA, A., WU, W., KURTEK, S., KLASSEN, E. and MARRON, J. S. (2011a). Registration of functional data using Fisher–Rao metric. Preprint. Available at [arXiv:1103.3817v2](#).
- SRIVASTAVA, A., KLASSEN, E., JOSHI, S. H. and JERMYN, I. H. (2011b). Shape analysis of elastic curves in Euclidean spaces. *IEEE Trans. Pattern Anal. Mach. Intell.* **33** 1415–1428.

- TANG, R. and MÜLLER, H.-G. (2008). Pairwise curve synchronization for functional data. *Biometrika* **95** 875–889. [MR2461217](#)
- TANG, R. and MÜLLER, H.-G. (2009). Time-synchronized clustering of gene expression trajectories. *Biostatistics* **10** 32–45.
- TUCKER, J. D., WU, W. and SRIVASTAVA, A. (2013). Generative models for functional data using phase and amplitude separation. *Comput. Statist. Data Anal.* **61** 50–66. [MR3063000](#)
- TUDDENHAM, R. D. and SNYDER, M. M. (1954). Physical growth of California boys and girls from birth to eighteen years. *University of California Publication in Child Development* **1** 183–364.
- VANTINI, S. (2012). On the definition of phase and amplitude variability in functional data analysis. *TEST* **21** 676–696. [MR2992088](#)
- VEERARAGHAVAN, A., SRIVASTAVA, A., ROY-CHOWDHURY, A. K. and CHELLAPPA, R. (2009). Rate-invariant recognition of humans and their activities. *IEEE Trans. Image Process.* **18** 1326–1339. [MR2742162](#)
- WANG, H. and MARRON, J. S. (2007). Object oriented data analysis: Sets of trees. *Ann. Statist.* **35** 1849–1873. [MR2363955](#)
- YOUNES, L., MICHOR, P. W., SHAH, J. and MUMFORD, D. (2008). A metric on shape space with explicit geodesics. *Atti Accad. Naz. Lincei Cl. Sci. Fis. Mat. Natur. Rend. Lincei (9) Mat. Appl.* **19** 25–57. [MR2383560](#)