3

CURVE-THINKING: UNDERSTANDING REACTION NORMS AND DEVELOPMENTAL TRAJECTORIES AS TRAITS

Joel Kingsolver[1], Sarah Diamond[2,3], and Richard Gomulkiewicz[4]

[1]*Department of Biology, University of North Carolina, Chapel Hill, NC, USA*
[2]*Department of Biology, North Carolina State University, Raleigh NC, USA*
[3]*Department of Biology, Case Western Reserve University, Cleveland, OH, USA*
[4]*School of Biological Sciences, Washington State University, Pullman,WA, USA*

INTRODUCTION

A central challenge in organismal biology is to understand how the phenotypic traits of an individual organism determine its performance and fitness in specific environmental conditions. The long traditions of research in functional morphology, environmental physiology, and animal behavior have addressed this issue for different types of traits and different environments in diverse organisms. In combination with phylogenetic information, these functional studies help us understand patterns of evolutionary adaptation in nature. More recently, many researchers have explored how phenotypic variation among individuals within populations determines variation in performance and fitness – that is, phenotypic selection – under different environmental conditions. In combination with genetic and population information, these studies help us understand the microevolutionary processes that generate adaptation, integrating individual- and population-level perspectives. Both means and variances in functionally important traits are central to analyzing patterns and processes of adaptive evolution.

Many important phenotypes are not single traits or sets of traits, but curves – traits that are functions of some continuous index variable. In general, such traits are called function-valued traits (FVTs). (A curve is a function of a single index, and a function of two or more index variables is often called a surface.) For example, phenotypic plasticity

Integrative Organismal Biology, First Edition.
Edited by Lynn B. Martin, Cameron K. Ghalambor, and H. Arthur Woods.
© 2015 John Wiley & Sons, Inc. Published 2015 by John Wiley & Sons, Inc.

involves traits whose values vary with one or more environmental variables, such as temperature, light intensity, nutrient concentration or density; in this case the index is an environmental variable. The relationship between trait value and environment for a genotype is often called the reaction norm. Developmental reaction norms indicate plastic traits in which environmental conditions during development determine trait values later in development. For example, in many animals, higher temperatures during development reduce adult body size; in many plants, greater plant density or reduced light intensity results in longer stems. In these cases, an individual organism only expresses a single fixed trait value that is determined by earlier environmental conditions. By contrast, performance curves indicate biological rates (and other aspects of performance) that vary with current or recent environmental conditions. For example locomotory rates of ectothermic animals vary with current body temperatures; photosynthetic rates of plants vary with light intensity. In these cases an individual organism may express a flexible range of different trait values depending on environmental variation. Similarly, many behavioral traits may vary flexibly with the environmental conditions experienced by the individual organism (Chapter 4, this volume). Combinations of fixed and flexible plasticity can also occur (e.g., reversible physiological acclimation). Epigenetic variation may also influence patterns of plasticity (Chapter 7, this volume).

A second important class of function-valued traits is phenotypes that are functions of the age of an individual organism. For example a growth curve or developmental trajectory represents the size of an individual as a function of its age throughout development. Similarly, most life history traits vary with age. By definition an individual organism expresses a range of sizes and other traits throughout its development. Of course, some phenotypes may vary as functions of both age and environmental variables (Chapter 2, this volume).

Developmental reaction norms, performance curves, growth curves and other FVTs can have a variety of different shapes. However, most theoretical analyses and empirical studies of phenotypic plasticity consider only two environmental levels, and quantify plasticity in terms of reaction norm slopes (or their equivalent). For example, random regression and similar statistical models may be used to estimate the variance in plasticity in terms of the variation in reaction norm slopes among individuals or genotypes (Chapter 2, this volume). While powerful, this approach assumes that reaction norms are strictly linear – that reaction norms and performance curves aren't curved. By contrast, a recent meta-analysis of continuous reaction norms suggests that variation in curvature and higher-order aspects of curve shape is greater than variation in slope for most organisms and traits (Murren *et al.* 2014).

Understanding organismal traits that are curves poses important methodological, statistical and conceptual challenges (Kirkpatrick & Heckman 1989; Kirkpatrick & Lande 1989; Kirkpatrick *et al*. 1990; Kingsolver *et al*. 2001; Stinchcombe *et al*. 2012). For example, when measuring growth rates of tadpoles at different rearing temperatures, or body mass of mice at different ages, how do we best quantify the phenotypic variation or mean differences in these curves? How can we partition genetic and environmental components of total variation in curves? If selection acts over a limited range of temperatures or ages, how does this determine evolution of the entire curve (Gomulkiewicz & Kirkpatrick 1992)? An important feature of curves is that, as the difference in environmental levels or ages decreases, the trait values at those levels or ages must necessarily become more correlated (with a perfect correlation when the difference is zero). As a result, variation and differences in the shapes of curves may be more strongly constrained than for other trait types, with implications for their evolution. How do we characterize such constraints, and assess their importance for evolutionary divergence in reaction norms and developmental trajectories (Kirkpatrick & Lofsvold 1992)?

This chapter is an exercise in curve-thinking for integrative biologists, as applied to phenotypic plasticity and developmental trajectories. The ideas and methods are quite general, but we use two main case studies to illustrate curve-thinking: thermal performance curves of insects, and growth trajectories of plants. First, we discuss various approaches to quantifying variation in curves, and their advantages and disadvantages. A key point is to distinguish between the shape of a curve, and the pattern of variation among curves. Second, we review ways of characterizing and visualizing phenotypic and genetic variation in curves using principal components analysis. We emphasize how different biological hypotheses about curve variation can generate distinctive patterns in the principal components. An important and familiar challenge is how to provide biological interpretations of PCA results for curves. Third, we will describe and illustrate a recent statistical method that identifies simple, interpretable patterns of phenotypic and genetic variation in curves. The key insight here is that we can decompose this variation in terms of different "directions" of simplicity in the curves. Finally, we briefly describe how patterns of genetic variation in curves may affect and constrain evolutionary changes in the mean curve for a population. A key idea here is that some patterns of selection on curves may lead to little or no evolutionary change, even when there is substantial genetic variation in the trait. Our overall goal is to provide conceptual and statistical tools for integrative and evolutionary biologists to develop and assess biological hypotheses about variation in curves and other function-valued traits. We conclude by discussing some challenges and promising directions for curve-thinking in integrative biology.

CHARACTERIZING CURVES

There are three general approaches to characterizing organismal curves. One approach is to choose a landmark value that represents a key feature of the curve: the body size at adulthood; the critical thermal maximum, indicating the maximum temperature for survival or activity; the saturation light intensity for photosynthetic rate; the LD50 for a toxin. These can provide valuable information that allows straightforward comparisons among organisms and systems, but this essentially avoids the issue by reducing a continuous curve to a single point – not curve-thinking. The second approach is to choose a specific functional form for the curve, and estimate the parameters of the function. In this case, the parameters become the traits of interest. This approach has been widely used for characterizing mean curves for both developmental trajectories (e.g., logistic and Gompertz functions for growth) and reaction norms (e.g., Michaelis-Menten model for physiological rates, modified Gaussian functions for thermal performance curves). This makes efficient use of limited data – in most organismic studies, traits are only measured at 3–6 index values – and the parameters often have simple biological interpretations. An important limitation of this approach is that for many (probably most) reaction norms and developmental trajectories, the appropriate parametric model is unknown. More generally, choosing a specific parametric model assumes that all individuals and genotypes follow the same model; differing only their parameter values. This assumption can be quite limiting when considering patterns of phenotypic and genetic variation and potential evolutionary changes in reaction norms and developmental trajectories. In addition, model parameters are estimated with some associated error, making it difficult to consider parameter values as traits of an individual organism for further analysis (but see Hadfield (2010) for one approach to this problem). As a consequence, parametric models are a powerful way to estimate mean curves for a sample of individuals, but are less useful for characterizing variation among individuals (see next page).

A third, nonparametric way of characterizing curves is by using a set of orthogonal polynomials, Fourier series, splines or other sets of functions (known as basis functions). An important advantage of this approach is that one can use the data to choose the number of basis functions (called the order) that best characterize mean and variation in the curves: the greater the order, the more complex (wiggly) the curves that can be represented. This approach can also be readily extended to quantify patterns of phenotypic and genetic variation in curves.

It is important to recognize that, as for any variable object, the mean of a sample of curves can differ greatly from the individual curves. Among other things, this implies that the mean curve need not represent the shape of the individual curves (see Figures 3.1 and 3.2 for examples). Our point here is that individual variation can affect mean curves estimated for populations or species; this may be particularly important when populations or species differ in the amount of individual variation.

VARIATION AMONG CURVES

For many reaction norms and developmental trajectories, the qualitative shape of the curve for a given trait may be similar for different individuals and genotypes within a population or species. As a result there may be important constraints on the patterns of phenotypic and genetic variation in these curves. This can lead to simple biological hypotheses about variation among curves (Huey & Kingsolver 1989; Izem & Kingsolver 2005). For example, imagine a sample of thermal performance curves (TPCs) for 5 individuals or genotypes (Figure 3.1, left panels). One pattern or "direction" of variation, called *vertical shift* (Figure 3.1, first row), describes variation in overall performance across all temperatures: relative to the mean curve, an individual with high performance at one temperature has high performance at all temperatures. *Maximum shift* (Figure 3.1, second row) describes variation in maximal performance at the optimal temperature: individuals have similar, low performance at very low and high temperatures, with greater variation among individuals at temperatures closer to the optimum. *Horizontal shift* (Figure 3.1, third row) describes variation in the location (e.g., the optimal temperature) of the curve: in this case individuals with relatively high performance at low temperatures have relatively low performance at high temperatures (and vice versa). A fourth pattern of variation involves tradeoffs between thermal breadth and maximal performance (aka *specialist-generalist* variation): here individuals with higher maximal performance at intermediate (optimal) temperatures have relatively low performance at both low and high temperatures (Figure 3.1, fourth row). These different hypotheses about variation have been widely explored by evolutionary physiologists interested in thermal adaptation (Huey and Kingsolver 1989; Angilletta 2009).

These different patterns of correlation in relative performance across different temperatures can be quantified using principal components analysis. Figure 3.1 (right panels) shows the loadings for the first principal component (explaining the greatest proportion of total variance) for each of the four hypotheses. The loadings describe the contributions from each temperature to differences among individuals in their performance across temperatures. For the *vertical shift* example, the loadings all have the same sign (i.e., the line is above zero) across all temperatures (Figure 3.1, first row): individuals have relatively high (or low) performance at all temperatures. (Recall that a principal component with all positive loadings is equivalent to one with all negative loadings – it is the change in sign across temperatures that is relevant.) For *maximum shift*, the loadings all have the same sign across

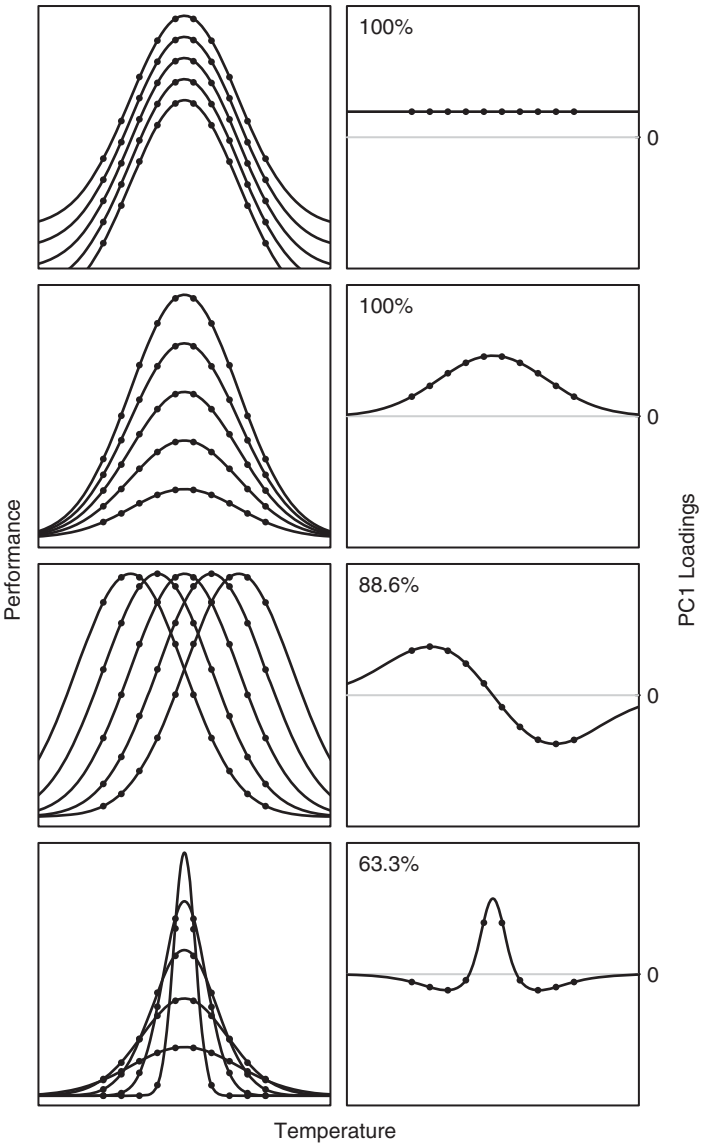

Figure 3.1. Toy thermal performance curves based on normal functions (left panels) and associated eigenfunction loadings for the first principal component (PC1) (right panels). Row 1: *vertical shift*; row 2: *maximum shift*; row 3: *horizontal shift*; row 4: *specialist-generalist tradeoff*. The points along each thermal performance curve correspond with the points along the eigenfunction. Percent variance explained by PC1 is indicated in the upper left of each panel.

all temperatures, but the magnitude of the loadings increases at intermediate temperatures (Figure 3.1, second row), which indicates that some individuals achieve higher performance than others at intermediate temperatures. For *horizontal shift*, the first principal component has positive loadings at low temperatures and negative loadings at high temperatures (or vice versa) (Figure 3.1, third row), which indicates a reversal of relative performance at

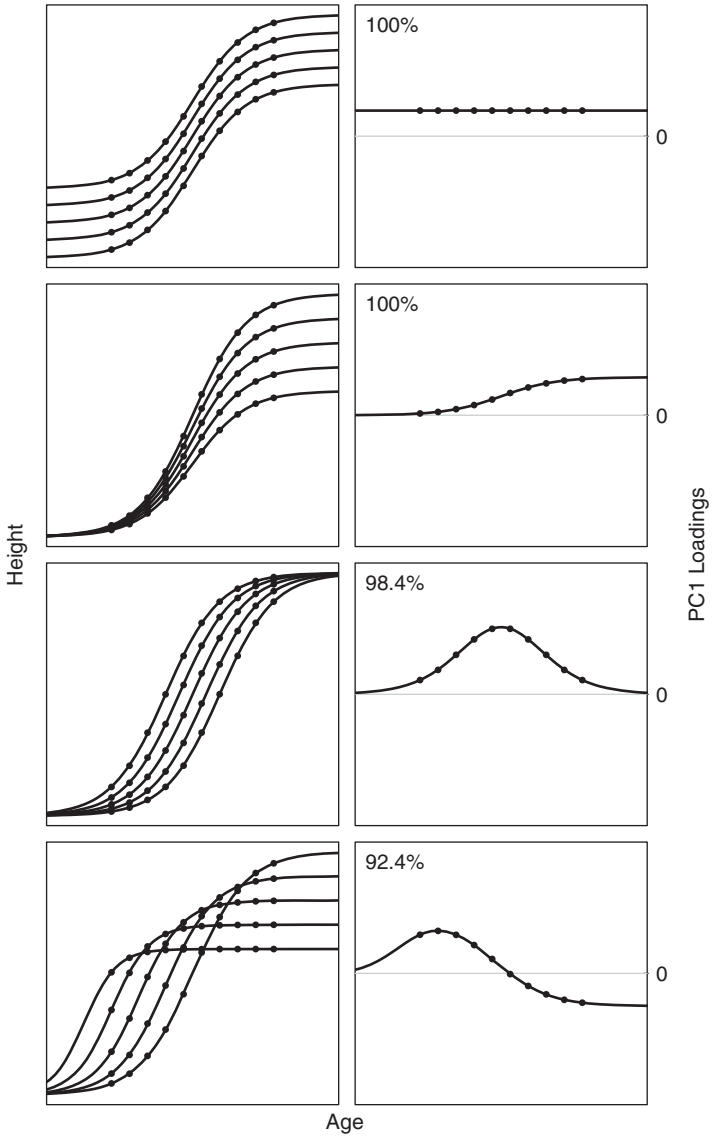

Figure 3.2. Toy growth curves based on logistic functions (left panels) and associated eigenfunction loadings for the first principal component (PC1) (right panels). Row 1: *vertical shift*; row 2: *maximum shift*; row 3: *horizontal shift*; row 4: *early growth-final size tradeoff*. The points along each growth curve correspond with the points along the eigenfunction. Percent variance explained by PC1 is indicated in the upper left of each panel.

low and high temperatures. In this sense there are "cool" and "warm" adapted individuals or genotypes. For the *specialist-generalist* case, the combination of positive loadings at intermediate temperatures with negative loadings at both low and high temperatures (or vice versa) indicates that specialist individuals with high relative performance at intermediate temperatures have low performance at low *and* high temperatures (Figure 3.1, fourth row).

The key insight here is that biological patterns of variation generate principal components with distinctive, characteristic shapes.

The principal components described above, and represented as the lines or curves in the right panels of Figure 3.1, are called eigenfunctions. There is a simple and direct relationship between eigenfunctions and more familiar eigenvectors. Imagine that we measure performance at 10 different temperatures for each individual (points in Figure 3.1, left panels). We can consider performance at each temperature as a separate, but potentially correlated, trait, so there are 10 traits for each individual: in this sense we represent the curve for each individual as a vector of 10 traits. We can estimate the variance-covariance matrix for the 10 traits, yielding the variance in performance at each temperature, and the covariance in performance across each pair of temperatures. For example, element [2,3] in this matrix indicates the covariance between performance at temperature 2 and performance at temperature 3. We can then use PCA to compute the eigenvectors and associated eigenvalues for the variance-covariance matrix. The loadings for each trait (i.e., performance at temperature 1, at temperature 2, etc.) on the first principal component (the first eigenvector) are the points along the principal component (Figure 3.1, right panels). In this sense, an eigenfunction is the natural extension of an eigenvector as we move from a discrete set of points to a continuous function. As a result we can approximate continuous curves as discrete vectors, and use the familiar method of principal components analysis (PCA) to decompose variation in the curves into a set of orthogonal vectors (the eigenvectors), based on the amount of variance explained. As illustrated above (Figure 3.1), plotting the loadings of the eigenvectors (e.g., the first PC) allows us to visualize patterns of variation in the curves.

Similar hypotheses about patterns of variation may occur for other types of reaction norms and for developmental trajectories, even if the shape of the mean curve is different. For example, for logistic growth curves we can consider *vertical shifts* that reflect variation in overall size at all ages (Figure 3.2, first row, left panel). The resulting principal component (Figure 3.2, first row, right panel) is identical to that for vertical shift of the TPCs (Figure 3.1, first row, right panel). *Maximum shift* represents variation in final size (Figure 3.2 second row), resulting in a first principal component with loadings that are near zero at early ages but that increase with age until reaching final (maximum) size. For growth curves, *horizontal shift* represents variation in the age at the inflection (midway) point (where growth rate is most rapid (Figure 3.2, third row). Loadings for the resulting principal component are near zero at initial and later ages, and are largest at intermediate ages. Note the shape of the principal component in this case is similar to that for the maximum shift case for TPCs (Figure 3.1, second row). We can also consider tradeoffs between the timing of growth (age at the inflection point) and final size (*early growth-final size tradeoff*), such that individuals who grow to larger size at early ages have relatively small final sizes (Figure 3.2, fourth row). Loadings for the resulting principal component have positive values at earlier ages and negative values at later ages, illustrating the tradeoff between early growth and final size. Note that the shape of the principal component for this case is qualitatively similar to that for the horizontal shift case in TPCs (Figure 3.1, third row). The key to this similarity in both cases is the changes in relative ranks of the curves: relative performance at low vs high temperatures, and relative height at early vs late ages.

Note that for both the TPC and growth curve examples, the *vertical shift* and *maximum shift* patterns (Figures 3.1 and 3.2, top two rows) are fully explained by a single principal component: PC1 explains 100% of the variance. This is because these patterns represent linear changes or deformations of the curves (and principal components analysis is a linear method of decomposition). In contrast the other two patterns of variation (Figures 3.1

and 3.2, bottom two rows) represent nonlinear deformations in the curves, which cannot be fully described by a single principal component: in these cases, the variance explained by PC1 is less than 100%.

There are two important insights from these toy examples. First, there is often a direct relationship between simple biological patterns of variation in curves and the shape of their associated principal component (eigenvector or eigenfunction), for reaction norms, developmental trajectories, and other biological curves. The key to understanding this association is that changes in the relative rankings of curves across (e.g.) ages or temperatures produce changes in the sign of the associated principal component with age or temperature. Second, the shape of the principal components may have different biological meanings for curves of different shapes (e.g., normal vs logistic). This is a key insight as we develop biological interpretations about important patterns of phenotypic and genetic variation in curves.

Armed with these conceptual and statistical tools, let us now move from toy examples to data for real biological curves. The methods can be applied to both phenotypic and genetic components of variance. For example, phenotypic data for a sample of individuals from a population can be used to estimate the phenotypic variance-covariance matrix, P. Quantitative genetic studies frequently decompose this total phenotypic variation into genetic and environmental components, to estimate the genetic variance-covariance matrix, G (Chapter 4, this volume). Estimation of G is useful for predicting evolutionary responses of quantitative traits, including curves, to selection (Lande & Arnold 1983; Via & Lande 1985). In our two case studies we will consider G matrices that approximate patterns of genetic variation for curves.

The first example considers thermal performance curves for short-term growth rates of *Pieris rapae* (Imported Cabbageworm) larvae, in which each 4th instar larva was measured at a series of 6 different temperatures (for details and data see (Kingsolver *et al.* 2004). These data were used to estimate the (broad-sense) genetic variance-covariance matrix, G. Principal components analysis of G shows that the first 4 principal components (eigenvectors) explain over 99% of the total genetic variance (Gaydos *et al.* 2013) (Figure 3.3). The shape of the curve for the first PC (explaining 60% of the variance) is quite complex, but is dominated by strong positive loadings at the highest temperature (40 °C) and strong negative loadings at 35 °C (and vice versa). This pattern indicates that most differences among individual curves involve higher growth rates near the optimal temperature (around 35 °C for most individuals and families) coupled with lower growth rates at high temperature (and vice versa; see Izem & Kingsolver (2005) for additional analysis and discussion). The second PC, explaining 19% of the variance, has positive loadings for all temperatures between 17 and 40 °C, and loadings near zero at 11 °C (Figure 3.3). Similar to the *vertical shift* or *maximum shift* patterns (Figure 3.1, rows 1–2), this secondary axis of variation suggests among-genotype variation that corresponds to curves with relatively high (or low) performance across a wide range of temperatures. The other two PCs are difficult to interpret.

The second example considers growth curves (height as a function of age) for *Impatiens capensis* (Jewelweed) plants raised in full-sun conditions at high density, in which each plant was measured at 6 different ages (see Stinchcombe *et al.* (2010) for details and data). Forty nine clonal genotypes were used to estimate the G matrix. PCA of G shows that the first two principal components explain more than 99% of the total genetic variance (Stinchcombe *et al.* 2010; Gaydos *et al.* 2013) (Figure 3.4). The first PC, explaining 99.8% of the variation, has positive loadings at all ages with the magnitude increasing with age, a pattern similar to that for *maximum shift* (see Figure 3.2, second row). This indicates variation in overall growth rate among clones, such that genotypes that are tall (or short) at one age are tall (or short) at all ages, and with increasing variation with age. The second

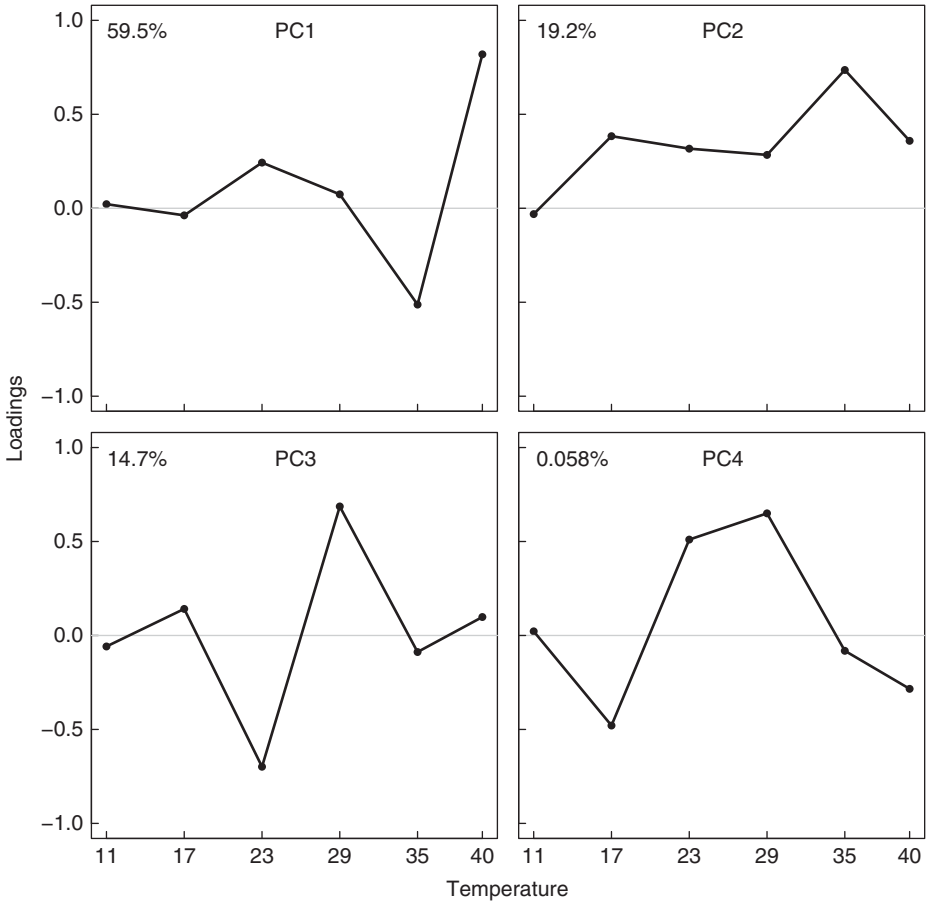

Figure 3.3. The first four principal components of the G-matrix for thermal performance curves of growth rates for Imported Cabbageworms. Percent variance explained by each component is indicated in the upper left of each panel.

PC, explaining only 0.004% of the genetic variation, is qualitatively similar to the *early growth-final size tradeoff* pattern (Figure 3.2, fourth row), with negative loadings at early ages and positive loadings at later ages (or vice versa). This secondary axis of variation suggests that (once overall variance is accounted for by PC1) genotypes that are relatively short at early ages are also relatively tall at later ages (and vice versa).

The structure of genetic covariation in these two examples is quite different. Only two PCs explain nearly all of the genetic variance in the Jewelweed growth trajectories, whereas four PCs are required for the Cabbageworm TPCs; in this sense the patterns of genetic variance are simpler (of lower dimension) in Jewelweed (Stinchcombe *et al.* 2010). In addition, much of the genetic variation in Jewelweed growth can be interpreted in terms of a single biological hypothesis (*maximum shift*: Figure 3.2, second row), but this is not the case for the Cabbageworm TPCs. Whether these differences are specific to these particular examples, or reflect more general differences between growth trajectories and TPCs, is unknown. PC analyses of P and G matrices for a variety of datasets for TPCs, developmental reaction norms, and growth trajectories would be most valuable in addressing this issue.

SIMPLICITY AND BIOLOGICAL HYPOTHESES

Our examples illustrate the important and well-known limitation of PCA that the principal components are often difficult to interpret biologically. The principal components estimated from real datasets (Figures 3.3 and 3.4) may rarely match those predicted by biological hypotheses (Figures 3.1 and 3.2). One solution is to develop statistical methods that directly evaluate specific hypotheses. For example, the "template mode of variation" (TMV) method was developed to decompose and quantify variation in thermal performance curves in terms of the easily interpreted vertical shift, horizontal shift, and generalist-specialist tradeoffs (Izem & Kingsolver 2005). The method places certain constraints on the qualitative shape of the curves (e.g., the curve is continuous, differentiable, and has a single maximum value), but does not assume a specific parametric form for the curve. This nonlinear approach has been used to analyze TPCs in a number of different study systems (Stinchcombe *et al.* 2012). However, the method has not been generalized to other types of function-valued traits, and it involves fancy mathematics like non-Euclidean manifolds and Dirichlet variances (Izem 2004).

An alternative, more general approach is to analyze variation in curves in terms of the simplicity of their eigenfunctions or eigenvectors. For current purposes we can consider simplicity to describe how wiggly or curvy a curve is: a simpler curve is less wiggly. The rationale for this approach is that, for many biological hypotheses, the qualitative shape of the associated eigenfunction is quite simple (Figures 3.1 and 3.2). For example, a flat (nonzero) eigenfunction indicates a vertical shift pattern of variation (Figures 3.1–3.2, top row). An eigenfunction that crosses zero once (e.g., is negative at low index values and positive at high index values, or vice versa) indicates a tradeoff in performance or size across temperatures (Figure 3.1, 3rd row) or ages (Figure 3.2, 3rd row). An eigenfunction that crosses zero twice indicates a specialist-generalist tradeoff: relative performance or size at intermediate temperatures or ages is inversely related to relative performance or size at low and high temperatures (Figure 3.1, 4th row).

Thus, eigenfunctions or eigenvectors with simple shapes can be readily interpreted in biological terms. We have seen the principal components analyses typically do not result in

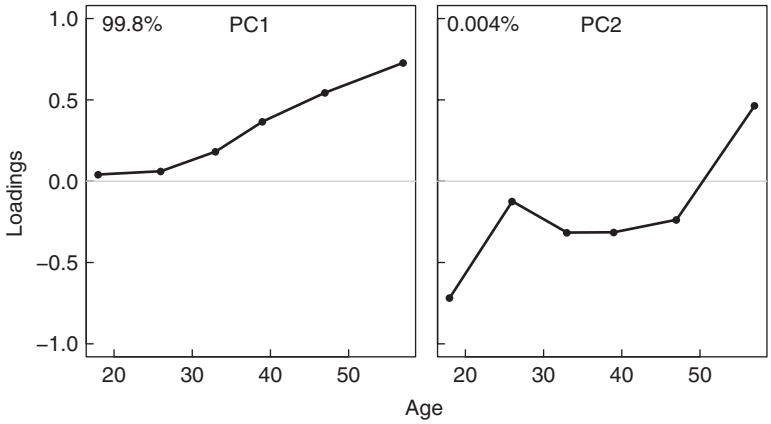

Figure 3.4. The first two principal components of the G-matrix for Jewelweed growth trajectories (height as a function of age). Percent variance explained by each component is indicated in the upper left of each panel.

simple eigenfunctions or eigenvectors. Recall that PCA decomposes a variance-covariance matrix into a set of orthogonal components (the eigenvectors) based on the amount of variance explained by each component. However, other sets of orthogonal vectors (called the basis vectors for the space of variation) can be used to quantify and visualize variation. Gaydos and colleagues (Gaydos *et al.* 2013) have recently developed a method that defines a set of orthogonal basis vectors based on a quadratic metric of simplicity, and quantifies the variation associated with each basis vector.

The Simplicity metric S takes on values between 0 (least simple, most wiggly) and 4 (most simple, a flat line). These basis vectors are ordered from highest to lowest simplicity (Figure 3.5, left panel). For the sake of discussion, suppose performance is measured at 6 temperatures. The loadings for the simplest basis vector (vector 1) give a flat line of constant (nonzero) magnitude (Figure 3.5), indicating *vertical shift*, with variation in overall performance across all temperatures (see Figure 3.1, first row). The Simplicity value for this vector is S = 4. The loadings for the 2nd basis vector (S = 3.7) give a monotonic line with negative loadings at low temperatures and positive loadings at high temperatures (or vice versa), indicating a tradeoff between high performance at low temperatures and low performance at high temperatures (and vice versa). The loadings for the 3rd basis vector (S = 3) represent a *specialist–generalist tradeoff*, in which relatively high maximal performance at intermediate temperatures is correlated with relatively low performance at both low and high temperatures. The 4th (S = 2), 5th (S = 1) and 6th (S = 0.27) basis vectors indicate increasingly complex (wiggly) patterns of variation across temperature, that are more difficult to interpret biologically.

In addition, we can compute the fraction of total variance associated with each Simplicity basis vector (Gaydos *et al.* 2013). This allows us to quantify the variation in a sample of reaction norms, developmental trajectories or other biological curves in terms of simple, biologically interpretable directions. An R library (*prinsimp*) was recently developed to implement this method for analyzing variance-covariance matrices (Cubranic *et al.* 2013).

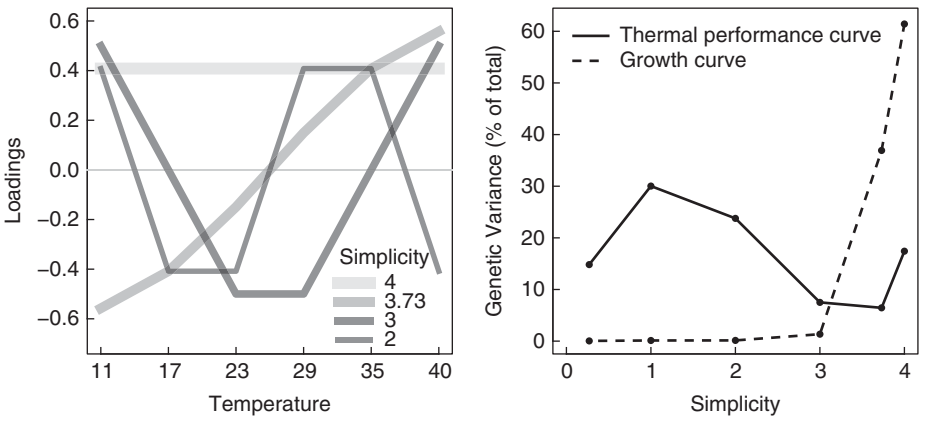

Figure 3.5. Left panel: The first four simplest basis vectors for thermal performance curves measured at 6 different temperatures. Simplicity values are indicated with different shading (4 = light grey to 2 = dark grey) and line thicknesses (higher values of simplicity have thicker line widths). Right panel: Percent variance explained as a function of simplicity based on thermal performance curves of Imported Cabbageworm growth rates (solid line) and Jewelweed growth trajectories (dashed line).

We can apply this approach to our G matrices for TPCs of Imported Cabbageworms and for growth curves of Jewelweed, by representing the matrices in terms of the Simplicity basis, and quantifying the percent variance as a function of the Simplicity value for each vector (Figure 3.5, right panel). For the TPCs, the simpler vectors (e.g., S = 3–4) explain a relatively small fraction of the total genetic variance. For example, the three simplest vectors (see Figure 3.5, left panel) account for only 32% of the variance, with the remainder accounted for by more complex patterns of variation. Returning to the PC analysis for these data (Figure 3.3), this complexity may be the result of the large loadings of opposite sign between performance at 35 °C and 40 °C – a very wiggly feature of the variation.

The Jewelweed growth curves show a very different pattern (Figure 3.5, right panel). The simplest vector (S = 4), representing overall variation in size across all ages (*vertical shift*), explains more than 61% of the total genetic variance. The 2nd vector (3.7), representing a tradeoff between relatively small size at early ages and large size at later ages (and vice versa), explains an additional 37% of the variance. Together these two simple and interpretable patterns explain over 98% of the total genetic variance in these curves. Conversely, the other 4, more complex vectors explain less than 2% of the variance (see below).

These results for Simplicity have some important implications for the evolution of these curves, because genetic variation directly determines the evolutionary responses to selection on function-valued traits (Gomulkiewicz & Kirkpatrick 1992; Kirkpatrick & Lofsvold 1992; Stinchcombe *et al.* 2012). Imagine there is directional selection that favors increased performance at all temperatures, or increased height at all ages. We would expect a strong evolutionary response to such selection in Jewelweed growth curves, because more than 61% of the genetic variance is in this "vertical shift" direction. In contrast we would expect a weaker evolutionary response to selection in Cabbageworm TPCs, as only 17% of the genetic variance is in the vertical shift direction. Whether these are general differences between growth curves and TPCs, or specific to these particular study systems, is unknown.

These results also provide some interesting insights about genetic constraints (Kirkpatrick & Lofsvold 1992). In quantitative genetics, genetic constraints are present if strong phenotypic selection results in little or no evolutionary response due to the lack of appropriate genetic variance. In practice, PCA can be used to detect and describe genetic constraints on adaptive evolution of curves: constraints are suggested by the principal components that explain the least amount of overall variation. Genetic constraints are visualized as the loadings corresponding to these PCs; the corresponding constraints can be interpreted in terms of these loadings since selection in the direction of these PCs will produce little or no evolutionary response. Gaydos and colleagues (Gaydos *et al.* 2013) provide a detailed analysis of genetic constraints on curves and other function-valued traits, and developed the simplicity basis approach as a means of visualizing these genetic constraints using these same case studies. For the Jewelweed example, consider the basis vector with Simplicity S = 3, which has positive loadings at early and late ages with negative loadings at intermediate ages (or vice versa) (Figure 3.5, left panel). This indicates individuals that are relatively large at early and late ages, but relatively small at intermediate ages: this pattern could emerge if relatively large seedlings delay their period of rapid growth, but end up at a larger final size (not shown). The simplicity analysis (Figure 3.5, right panel) shows that this pattern of variation represents less than 1% of the total genetic variance in the Jewelweed growth curves (Gaydos *et al.* 2013). If there was phenotypic selection favoring this pattern of variation (larger size at early and late ages, and smaller size at intermediate ages), there would be little or no evolutionary response to selection due to the lack of genetic variation in this "direction" – even though there is genetic variation in height at all ages. In this way, simplicity analyses can help quantify and visualize patterns of both genetic variation and genetic constraints in curves and other function-valued traits (Gaydos *et al.* 2013).

SUMMARY AND FUTURE DIRECTIONS

The goal of this chapter has been to explore curve-thinking: the conceptual and statistical tools for characterizing, quantifying and visualizing variation in biological curves such as reaction norms and developmental trajectories. Integrative biologists are familiar with the "parameters as traits" approach, in which a specific parametric model is chosen to describe the curve of interest, and the model parameters are estimated for a sample of individuals. This approach can be useful for describing mean curves, but is too restrictive for characterizing phenotypic and genetic variation for most biological curves. Here we have emphasized a familiar method, principal components, as a tool for visualizing and quantifying variation in curves.

A major theme of our discussion is developing biological hypotheses about curve variation among individuals or genotypes. Using toy examples we illustrated how different patterns of variation generate principal components with different shapes that can be readily visualized and interpreted. Our comparison between normal TPCs and logistic growth curves shows that the mapping of a hypothesis to PC shape can sometimes depend on the shape of the curve itself. This approach also provides a natural way to visualize tradeoffs: a tradeoff involves changes in the relative rankings of traits values across ages or environmental levels, and is reflected in a change in the sign of the loadings for the associated PC across ages or environmental levels. It would be useful to explore additional hypotheses about variation in thermal performance curves and growth curves beyond those considered here, and to explore hypotheses for other types and shapes of biological curves.

These hypotheses can be used to interpret the results of principal components analyses on real biological curves. We considered two case studies of genetic variation, in thermal performance curves of Imported Cabbageworms and growth curves of Jewelweeds. One striking result is that most of the genetic variance in Jewelweed growth curves is consistent with a single biological hypothesis (variation in *maximum shift*), whereas none of the hypotheses we considered explained most of the genetic variance in Cabbageworm TPCs. It would be valuable to have similar analyses of TPCs, growth curves and developmental reaction norms for other traits and study systems to explore possible generalities about these patterns. We emphasize that these methods apply to both phenotypic and genetic variation. At a minimum, trait measurements at 4 or more index levels (ages or environments) for at least 15–20 individuals or genotypes (clones, sibs, etc.) will likely be required for the analyses to be informative (Griswold *et al.* 2008).

Unfortunately, the genetic and physiological bases for natural variation in performance curves, reaction norms and growth trajectories are poorly understood. Genetic analyses with bacteriophage show that individual base substitutions (point mutations) can affect multiple aspects of performance curve shape such as optimal temperature and thermal breadth (Knies *et al.* 2006). This suggests that the different "directions" of variation described by our hypotheses may not be independent at the genetic level (Knies *et al.* 2009). The physiological bases of thermal reaction norms for adult body size have been identified recently for some holometabolous insects (Chapter 13, this volume). Despite this, how variation in the underlying physiological parameters determines variation in reaction norm shape has not been explored. Connecting physiological mechanisms to patterns of genetic and phenotypic variation in biological curves remains an important challenge for integrative and evolutionary physiologists.

We have emphasized the use of discrete, matrix-based methods for computing eigenvectors and eigenvalues using PCA because of its familiarity to most integrative and population biologists. However, this discrete approach to curves has important limitations. For example, it requires that all individuals be measured at the same index values;

it does not take into account the order or spacing of index values; it is less efficient and less powerful in hypothesis testing (Griswold *et al.* 2008). True curve-thinking requires consideration of eigenfunctions and variance-covariance functions, rather than their more familiar discrete counterparts. These can be estimated using random regression and related models (see Stinchcombe *et al.* (2012) for a recent overview). However for cases where traits are consistently measured at a fixed and limited set of index values (e.g., 3–8), the discrete approximations can be very useful.

Another general theme we have explored is Simplicity. PCA decomposes total variation onto different orthogonal axes (principal components) based on the amount of variance explained, but the principal components are often difficult to interpret. We describe an alternative approach that decomposes variance onto orthogonal axes based on the simplicity of the basis vectors (curves) that may have simple biological interpretations. One interesting result is that more than 98% of the variance in Jewelweed growth curves is explained by two, simple vectors that can be interpreted in terms of two hypotheses (*vertical shift* and *early growth-final size tradeoff*). Applying this new approach to other traits, study systems and types of curves will be interesting to see whether there are general patterns of simplicity in the variation of biological curves.

ACKNOWLEDGMENTS

We thank Travis Gaydos, Nancy Heckman and Steve Marron for useful discussion; Davor Cubranic and Nancy Heckman for help with *prinsimp*; and Cameron Ghalambor, Marty Martin, and Art Woods for suggestions on an earlier version of the paper. Work supported in part by NSF IOS-1120500 to JGK and NSF DEB-1354264 to RG.

REFERENCES

Angilletta, M.J. (2009) *Thermal Adaptation: A Theoretical and Empirical Synthesis*, Oxford University Press, Oxford.

Cubranic, D., Zhang, J., Heckman, N.E., Gaydos, T.L. & Marron, J.S. (2013) prinsimp: Finding and plotting simple basis vectors for multivariate data. Comprehensive R Archive Network, URL: http://CRAN.R-project.org/package=prinsimp.

Gaydos, T.L., Heckman, N.E., Kirkpatrick, M. *et al.* (2013) Visualizing genetic constraints. *Annals of Applied Statistics*, **7**, 860–82.

Gomulkiewicz, R. and Kirkpatrick, M. (1992) Quantitative genetics and the evolution of reaction norms. *Evolution*, **46**, 390–411.

Griswold, C., Gomulkiewicz, R. and Heckman, N. (2008) Hypothesis testing in comparative and experimental studies of function-valued traits. *Evolution*, **62**, 1229–42.

Hadfield, J. (2010) MCMC methods for multi-response generalized linear mixed models: the MCMCglmm R package. *Journal of Statistical Software*, **33**, 1–22.

Huey, R.B. and Kingsolver, J.G. (1989) Evolution of thermal sensitivity of ectotherm performance. *Trends in Ecology and Evolution*, **4**, 131–5.

Izem, R. (2004) Analyzing variation in curves of common shape. PhD, University of North Carolina.

Izem, R. and Kingsolver, J.G. (2005) Variation in continuous reaction norms: Quantifying directions of biological interest. *The American Naturalist*, **166**, 277–89.

Kingsolver, J.G., Gomulkiewicz, R. and Carter, P.A. (2001) Variation, selection and evolution of function-valued traits. *Genetica*, **112/113**, 87–104.

Kingsolver, J.G., Ragland, G.J. and Shlichta, J.G. (2004) Quantitative genetics of continuous reaction norms: Thermal sensitivity of caterpillar growth rates. *Evolution*, **58**, 1521–9.

Kirkpatrick, M. and Heckman, N. (1989) A quantitative genetic model for growth, shape, reaction norms, and other infinite-dimensional characters. *Journal of Mathematical Biology*, **27**, 429–50.

Kirkpatrick, M. and Lande, R. (1989) The evolution of maternal characters. *Evolution*, **43**, 485–503.

Kirkpatrick, M. and Lofsvold, D. (1992) Measuring selection and constraint in the evolution of growth. *Evolution*, **46**, 954–71.

Kirkpatrick, M., Lofsvold, D. and Bulmer, M. (1990) Analysis of the inheritance, selection and evolution of growth trajectories. *Genetics*, **124**, 979–93.

Knies, J.L., Izem, R., Supler, K.L. *et al.* (2006) The genetic basis of thermal reaction norm evolution in lab and natural phage populations. *PLoS Biology*, **4**, 1–8.

Knies, J.L., Kingsolver, J.G. and Burch, C.L. (2009) Hotter Is better and broader: Thermal sensitivity of fitness in a population of bacteriophages. *The American Naturalist*, **173**, 419–30.

Lande, R. and Arnold, S.J. (1983) The measurement of selection on correlated characters. *Evolution*, **37**, 1210–26.

Murren, C.J., MacLean, H.J., Diamond, S.E., Steiner, U.K., Keskel, M.A., Handelsman, C.A., Ghalambor, C.K., Auld, J.R., Callahan, H.S., Pfennig, D.W., Relyea, R.A., Schlichting, C.D. & Kingsolver, J.G. (2014) Evolutionary change in continuous reaction norms. *American Naturalist*. In press.

Stinchcombe, J.R., Group, F.-V.T.W. and Kirkpatrick, M. (2012) Genetics and evolution of function-valued traits: understanding environmentally responsive phenotypes. *Trends in Ecology and Evolution*, **27**, 637–47.

Stinchcombe, J.R., Izem, R., Herschel, M.S. *et al.* (2010) Across-environment genetic correlations and the frequency of selective environments shape the evolutionary dynamics of growth rate in *Impatiens capensis*. *Evolution*, **64**, 2887–2903.

Via, S. and Lande, R. (1985) Genotype-environment interaction and the evolution of phenotypic plasticity. *Evolution*, **39**, 505–22.