Variation in Continuous Reaction Norms: Quantifying Directions of Biological Interest

Rima Izem[1],[*] and Joel G. Kingsolver[2],[†]

1. Department of Statistics, University of North Carolina, Chapel Hill, North Carolina 27599;
2. Department of Biology, University of North Carolina, Chapel Hill, North Carolina 27599

ABSTRACT: Thermal performance curves are an example of continuous reaction norm curves of common shape. Three modes of variation in these curves—vertical shift, horizontal shift, and generalist-specialist trade-offs—are of special interest to evolutionary biologists. Since two of these modes are nonlinear, traditional methods such as principal components analysis fail to decompose the variation into biological modes and to quantify the variation associated with each mode. Here we present the results of a new method, template mode of variation (TMV), that decomposes the variation into predetermined modes of variation for a particular set of thermal performance curves. We illustrate the method using data on thermal sensitivity of growth rate in *Pieris rapae* caterpillars. The TMV model explains 67% of the variation in thermal performance curves among families; generalist-specialist trade-offs account for 38% of the total between-family variation. The TMV method implemented here is applicable to both differences in mean and patterns of variation, and it can be used with either phenotypic or quantitative genetic data for thermal performance curves or other continuous reaction norms that have a template shape with a single maximum. The TMV approach may also apply to growth trajectories, age-specific life-history traits, and other function-valued traits.

Keywords: generalist-specialist trade-offs, genetic variation, *Pieris rapae*, reaction norms, thermal performance curves.

The patterns of phenotypic and genetic variation in reaction norms are important because they determine se-

[*] Present address: Department of Statistics, Harvard University, Cambridge, Massachusetts 02138; e-mail: izem@stat.harvard.edu.

[†] Corresponding author; e-mail: jgking@bio.unc.edu.

lection and evolutionary response for a population. Phenotypic and genetic variation for quantitative traits, including reaction norms, is typically characterized in terms of the phenotypic and genetic variance-covariance matrices or functions, **P** and **G** (Kirkpatrick and Heckman 1989; Phillips and Arnold 1999). The structure of **P** or **G** can be explored using principal components analysis (PCA) and similar methods to decompose into underlying components. However, interpreting principal components is often difficult, and understanding how the structure of **P** or **G** emerges from underlying biological mechanisms is a major challenge.

Thermal performance curves (TPCs) are an important class of continuous reaction norms, in which the trait value (performance) varies as a function of environment (temperature; Huey and Stevenson 1979). For many aspects of organismal performance, TPCs have a common general shape (fig. 1, *left panels*), in which performance increases with increasing temperature, reaches a maximum at some intermediate temperature, and then declines rapidly with further increases in temperature (Huey and Kingsolver 1989). Variation in TPCs can thus be considered as variation around a common shape (template) curve.

Ecological and evolutionary physiologists have identified three directions or modes of special biological interest for TPCs: vertical shift (faster-slower), horizontal shift (hotter-colder), and generalist-specialist (Huey and Kingsolver 1989). Imagine a set of phenotypes or genotypes within a population whose TPCs vary in each of these modes (fig. 1, *left panels*). Vertical shift represents variation in the height of the TPCs, or variation in overall performances across all temperatures; horizontal shift represents variation in the position of the curve along the temperature axis, or variation in the location of the thermal maximum; generalist-specialist represents variation in the width of the TPCs and the trade-off between width and maximal performance. In our analysis, each one of these directions of variation will be called a mode. Previous studies have used PCA and related methods to identify these modes of variation in TPCs in several systems, with partial success (Gilchrist 1996; Kingsolver et al. 2001).
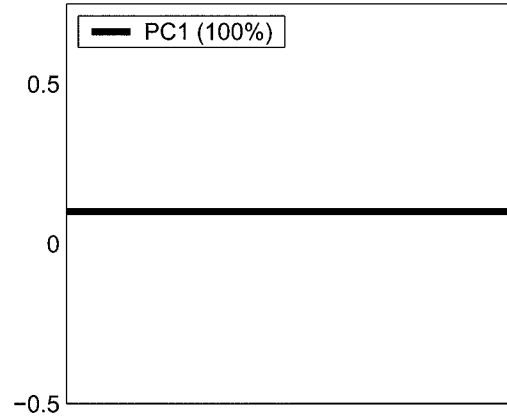
Vertical Shift: linear mode

Principal Component

PC1 (100%)

0.5

0

-0.5
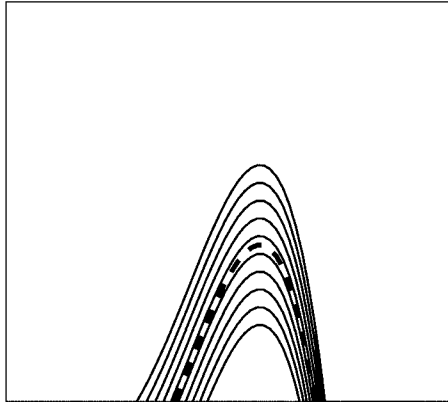
Horizontal Shift: nonlinear mode

Principal Components

PC1 (23.52 %)
PC2 (21.04 %)
PC3 (17.10 %)

0.5

0

-0.5

Performance z
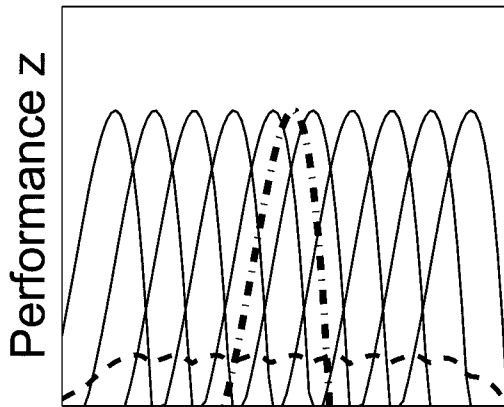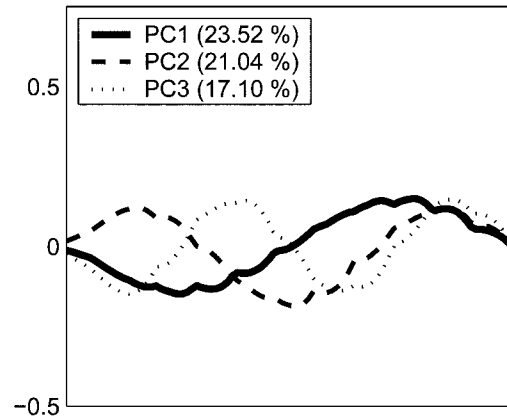
Generalist–Specialist: nonlinear mode

Principal Components

PC1 (94.45 %)
PC2 (4.52 %)
PC3 (0.69 %)

0.5

0

-0.5
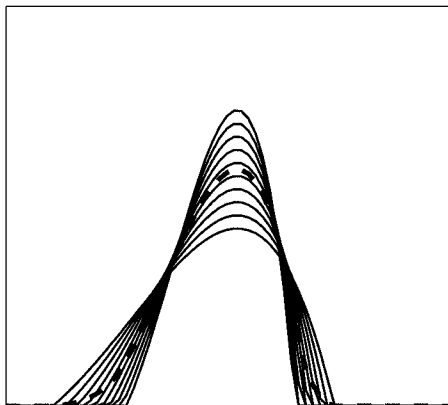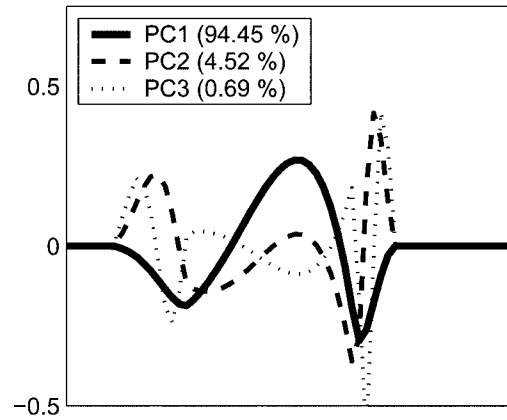
Temperature

In this article, we develop a new approach for analyzing variation in continuous reaction norms, which decomposes and quantifies the contributions of predicted modes of variation. First, we show how standard linear methods can be inadequate for estimating mean and variation and for interpreting patterns of variation in nonlinear reaction norms. Second, we describe a new method, template mode of variation (TMV), which can decompose variation in reaction norms into predetermined modes of variation. The method can be used to analyze either phenotypic or genetic variation; here we illustrate the approach using data on genetic (between-family) variation in thermal performance curves for short-term growth rate in *Pieris rapae* caterpillars (Kingsolver et al. 2004).

The Limits of Linear Methods

As an example, figure 2 presents mean thermal performance curves (TPCs) for 32 full-sib families of *Pieris rapae* caterpillars. In this study, the growth rate of each fourth instar caterpillar was measured at six temperatures between 11° and 40°C during a 48-h period within a single larval instar (for details, see Kingsolver et al. 2004). As with many TPCs, the curves have a common shape in which performance (growth rate) increases slowly, reaches a maximum, and decreases rapidly.

One standard method for analyzing variation in TPCs (and other sets of quantitative traits) is to apply PCA, treating performance at each temperature as a distinct but correlated trait (Via and Lande 1985; Gomulkiewicz and Kirkpatrick 1992; fig. 3). The loadings on each principal component (eigenvector) represent the contribution of performance at that temperature to the variation along the direction of the principal component; the eigenvalues indicate the amount of variation explained by each principal component. Principal components analysis for the TPC data in figure 1 results in a complex pattern of loadings on each principal component that cannot be easily interpreted biologically (fig. 3), a common problem with PCA.

Suppose we apply PCA to our proposed biological modes of variation: vertical shift, horizontal shift, and generalist-specialist (fig. 1). Here we have generated simulated data for TPCs that represent variation in each of these modes. For the vertical shift, the pointwise (arithmetic) mean curve is a good estimate of the common shape or template shape, and the variation is completely identified by the first principal component direction PC1 (fig. 1, *top row*). Principal components analysis successfully identifies vertical shift variation because this mode represents a simple linear translation around the mean curve (Fry 1992; Kingsolver et al. 2001).

Applying PCA to the other two modes of variation yields a very different pattern, because these modes involve nonlinear "directions." For the horizontal shift (fig. 1, *middle row*), the pointwise mean curve (*dashed line*) does not describe the common shape of any of the individual curves, and PCA results in a complex pattern of components and loadings that are difficult to interpret. Similarly, generalist-specialist variation does not simply correspond to the loadings on a particular principal component (fig. 1, *bottom row*). Both the horizontal shift and generalist-specialist modes generate complex patterns of negative and positive loadings on multiple principal components; PCA cannot readily distinguish between these two different types of variation. Moreover, when there is more than one mode of variation of interest, for example, simultaneous vertical shift and horizontal shift, PCA also fails to decompose the variation into these two modes of interest.

These considerations suggest that for nonlinear reaction norms, standard linear methods of characterizing the mean or variation may be inadequate or difficult to interpret. Izem (2004) has recently developed a new statistical method, TMV, to characterize and quantify nonlinear modes of variation. The approach introduces new measures to estimate the "average" curve and variation among curves, which take into account the nonlinear geometry of the space of variation. The TMV method allows for the decomposition and quantification of the variation in the data along linear and nonlinear modes of interest.

Figure 1: Three biological modes of variation in thermal performance curves (performance as a function of temperature), using toy simulated and discretized data. *Top row*, vertical shift mode. *Left*, simulated data (*solid lines*) and pointwise mean (*dashed line*). *Right*, principal component direction. Because vertical shift variation is linear and one dimensional, the pointwise mean curve falls in the middle of the variation, and all of the variation in the toy data (100%) is explained by one principal component. *Middle row*, horizontal shift mode. *Left*, simulated data (*solid lines*), pointwise mean (*dashed line*), and Fréchet mean (*dashed and dotted line*). Because this variation is nonlinear, the pointwise mean curve differs in shape from any of the individual curves and does not lie at the center of the variation in the curves, whereas the Fréchet mean has the same shape as the individual curves and lies at the center of the variation in the curves. *Right*, principal component directions. Because horizontal shift variation is nonlinear, it is not simply explained by one principal component. The first three principal component directions explain only about 61% of the total variation. *Bottom row*, generalist-specialist mode. *Left*, simulated data (*solid lines*) and pointwise mean (*dashed line*). *Right*, first three principal component directions. This case is less extreme than the horizontal shift case. However, because this mode is nonlinear, the variation is totally explained by three, rather than one, principal component directions.
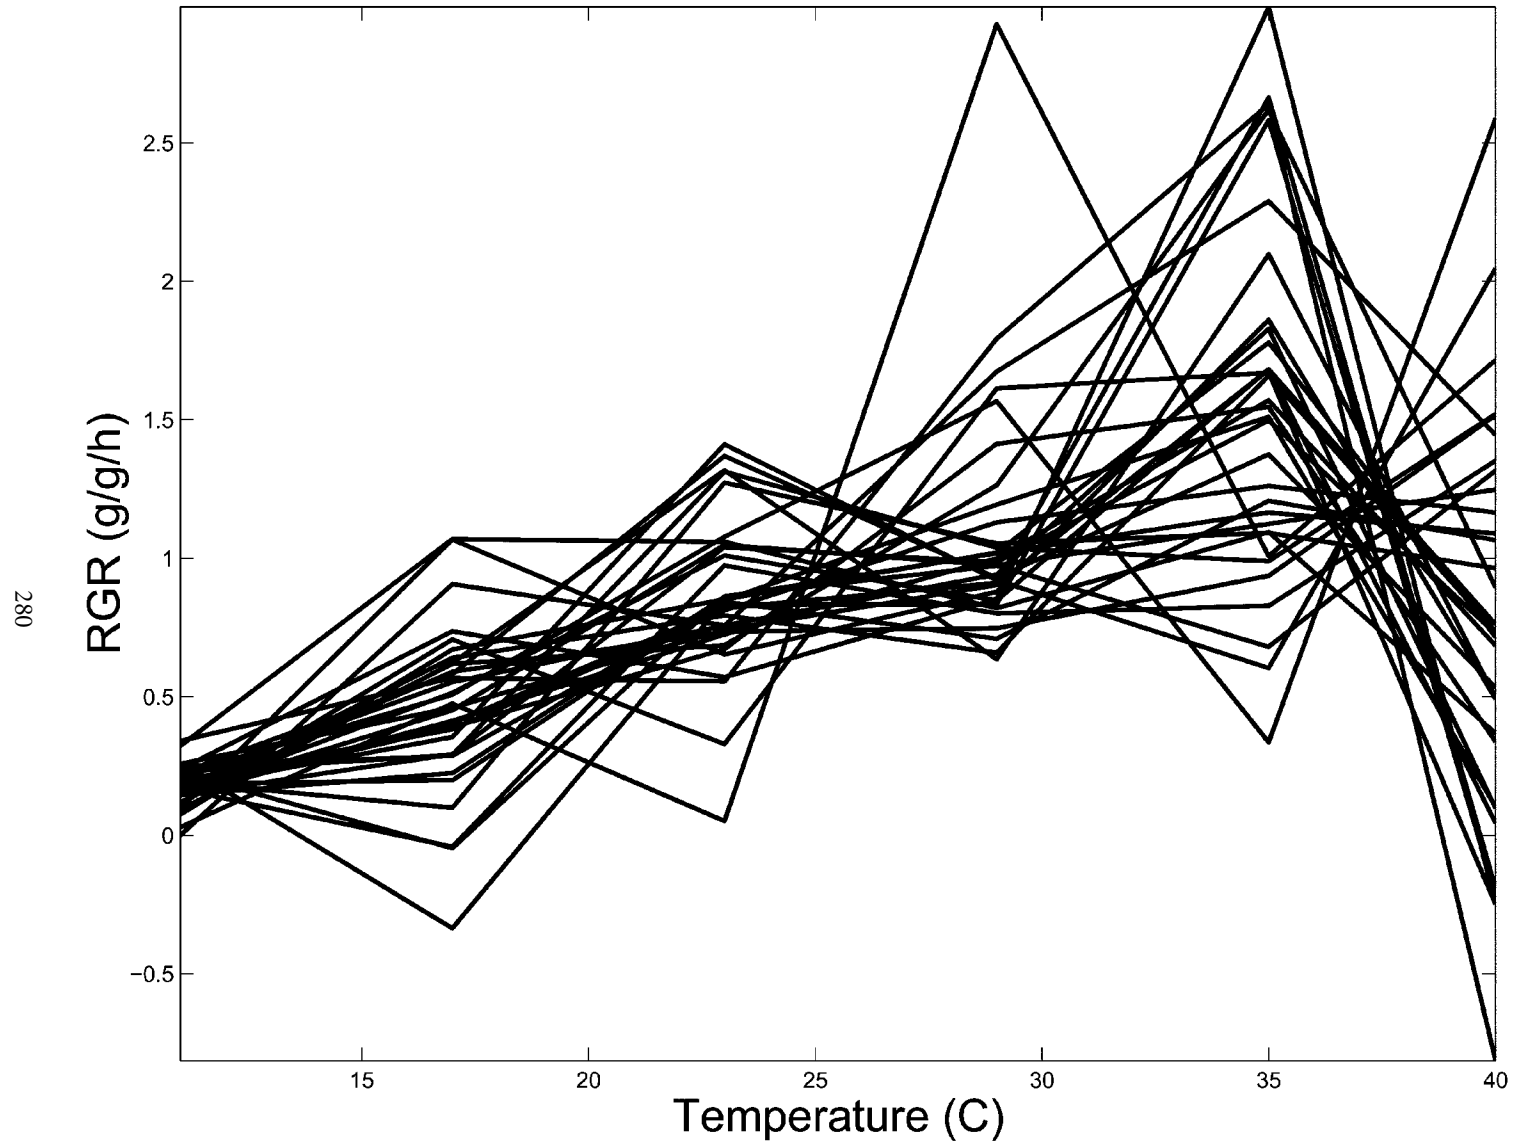
Figure 2: Mean relative growth rate (in 100 × g/g/h) of fourth instar *Pieris rapae* caterpillars as a function of temperature for 32 full-sib families; data from Kingsolver et al. (2004)
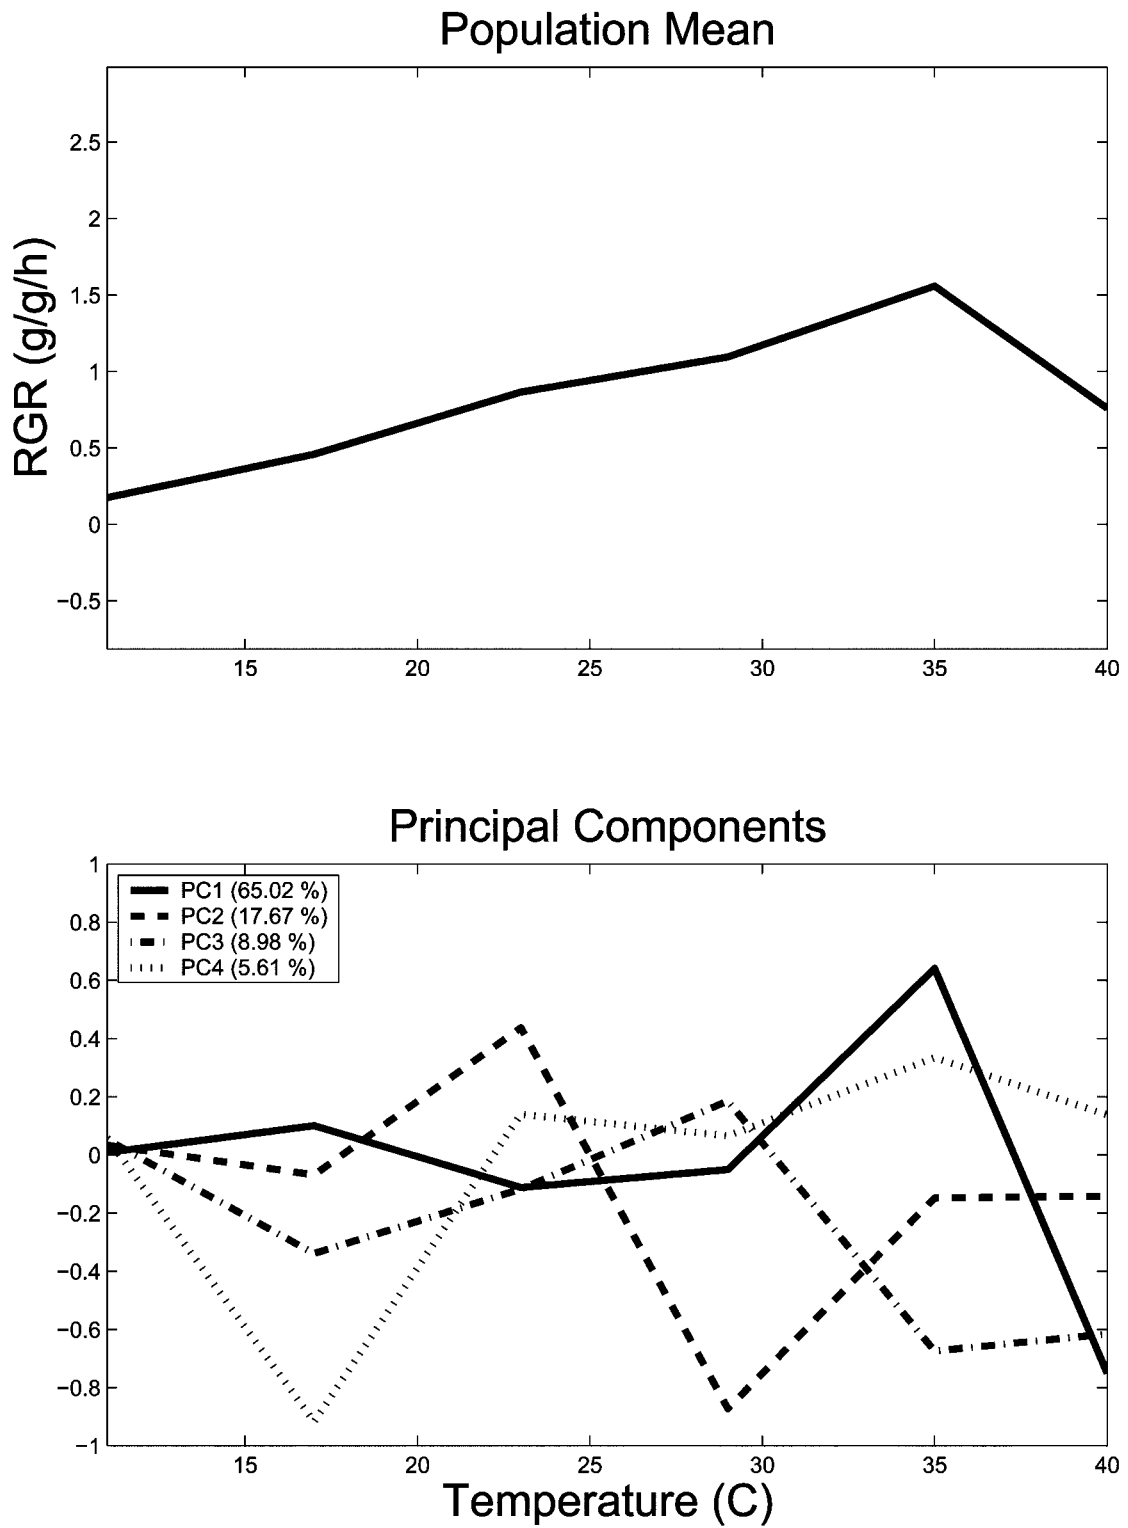
Figure 3: Principal components analysis of the caterpillar data. *Top*, weighted population mean relative growth rate. *Bottom*, first four principal component (PC) directions, with the percent variation explained by each PC.

Template Modes of Variation

The Model

An important feature of TPCs, and of continuous reaction norms in general, is continuity: while in practice performance is measured at a finite number of temperatures J, the underlying relationship between the performance and temperature is considered continuous. We model this mathematically as follows. Let $z_{i,j}$ be the mean performance for individuals of family (or genotype) i measured at temperature j. Since the reaction norm is continuous, the growth rate of a family i could be represented by a continuous function, z_i, such that

$$z_{i,j} = z_i(t_j) + \varepsilon_{i,j}; \; j \in \{1, \dots, J\}, i \in \{1, \dots, I\}, \quad (1)$$

where t_j is the jth temperature t, I is the number of families or genotypes, and $\varepsilon_{i,j}$ either are errors from the measurements or reflect environmental factors other than temperature.

A second assumption is the existence of a common shape for the TPCs of all genotypes in the population or sample of interest. The common shape can be modeled using a shape invariant model (Lawton et al. 1972). Since there are three modes of variations of interest—vertical shift, horizontal shift, and generalist-specialist—around a template shape function, we model the simultaneous three modes of variation with the following three-parameter shape invariant model:

$$z_i(t) = \frac{1}{w_i} z\left[\frac{1}{w_i}(t - m_i)\right] + h_i; \; i \in \{1, \dots, I\}. \quad (2)$$

In this model, z represents the common template shape of the curves, which is the same for all families. The curve for each family i is described by three parameters (height h_i, location of the maximum m_i, and width w_i) that characterize the variation of the curve of family i from the common shape z along the modes of interest. The height h_i parameterizes the vertical shift mode of variation, the location m_i parameterizes the horizontal shift mode of variation, and the width w_i parameterizes the generalist-specialist mode of variation. Note that h_i has the same dimensions as performance, m_i has dimensions of temperature, and width w_i is dimensionless. Since the common shape in reaction norm curves is not, in general, a straight line, z in equation (2) is not a straight line. As a result, the vertical shift is a linear mode, but the horizontal shift and generalist specialist are nonlinear modes; that is, z_i varies linearly with h_i but does not vary linearly with m_i and w_i.

To model the trade-off between generalists and spe-cialists, we assume that for m_i and h_i fixed, the positive area under the TPC curve is constant for change in the width parameter w_i. Thus, the width parameter incorporates both variation in width and maximum performance at the optimal temperature. Note that height h_i as defined here indicates the overall or "average" performance across all temperatures, not just the maximum performance at the optimal temperature.

In addition to continuity and a common shape for $z(t)$, there are three other assumptions of the model. First, the function $z(t)$ has a positive maximum at $t = 0$; that is, $\max(z) = z(0)$ and $z(0) > 0$. This ensures that the parameter m_i is the maximum for each curve z_i and that there is a positive area under the curve. Second, $z(t) = 0$ for at least two values in the neighborhood of $t = 0$. This condition ensures that the positive area under the curve is finite. Third, $\sum_{i=1}^{I} h_i = 0$. This ensures that there is only one possible function fit z with parameters h_i.

Note that the maximum performance at the optimal temperature for family i, $z_{\max, i}$, can also be obtained from equation (2):

$$z_{\max, i} = \frac{z(0)}{w_i} + h_i. \quad (3a)$$

By the assumptions above, $z(0)$ represents the maximum value for the template function $z(t)$. Rearranging this relationship, we see that

$$w_i = \frac{z(0)}{z_{\max, i} - h_i}. \quad (3b)$$

This provides a useful interpretation of the width parameter w_i. The denominator in equation (3b) represents the difference between maximum performance (at the optimum temperature) and the average performance (over all temperatures) for family i; thus, for any given h_i, decreasing maximum performance increases width w_i. Because this difference is standardized relative to the maximum value of z for the template function ($z(0)$), w_i is dimensionless.

Defining Mean and Variation

In linear methods such as PCA, the key to decomposing the variation in the data is, first, to define the center of the variation as the mean in the data; second, to quantify the spread in a direction by the sums of squares of Euclidean distances from the data to the center; and finally, to decompose the variation by using a Pythagorean theorem in the space of variation. When some of the variations of interest are nonlinear, the arithmetic mean is no longer at the center of the variation, and the Euclidean

distance is not appropriate (Izem et al. 2003; Izem 2004). For example, note that in figure 1 (*middle left panel*), the arithmetic mean curve (*dashed line*) for horizontal shift variation does not reflect the shape of any individual curve. As described in Izem (2004), we do not use the pointwise (arithmetic) mean curve as a measure of center of variation, but rather we use a curve that falls in the middle of the variation, defined as the Fréchet mean (cf. fig. 1, *middle left panel, dashed and dotted line*; Izem 2004). Similarly, we do not use the sums of squares of Euclidean distances to quantify the spread along each mode, but rather we construct a metric in the space of variation that depends on arc distances along each mode. We define a new ratio of sum of squares $\tilde{R\tilde{S}S}$, which generalizes the linear ratio of sum of squares that quantifies linear modes, to the quantification of nonlinear modes. This new ratio takes into account the curved geometry of the space of variation (Izem et al. 2003; Izem 2004; details available at http://www.fas.harvard.edu/~rizem).

Also note that since the three modes of variation are not all linear, they are not orthogonal in the usual sense. For instance, the space of variation is a manifold. This manifold is spanned by curves of variation rather than linear principal components. Decomposing the variation in this manifold was possible by linearizing this space (i.e., transforming the space into a linear space). So, the distance we constructed in the manifold relies on Euclidean distances in the transformed spaces. Since the Pythagorean theorem holds in a linear space with Euclidean metric, by equivalence, it holds also on the manifold of variation with the new metric (Izem 2004; details available at http://www.fas.harvard.edu/~rizem).

Fitting Criterion

In model equation (1), the common template shape z is unknown, and the I triplets (h_i, m_i, w_i) are unknowns. For simplicity, for TPCs we assume that z was a polynomial of degree 4. We fit the model (eq. [1]) by optimizing for the five coefficients of this polynomial as well as the parameters (h_i, m_i, w_i) for each curve. Because families may not all have the same sample size, it is useful to take the sample size within each family into account in our fitting. The criterion for optimization is to minimize the weighted sum of squared errors (SSE; weighted by sample size within each family) such that

$$\text{SSE} = I \times \sum_{i=1}^{I} \sum_{j=1}^{J} \frac{n_i}{N} \|z_{i,j} - z_i(t)\|^2, \quad (4)$$

where n_i is the sample size of family i and N is the total number of individuals in the study. Note that equation (4) represents a weighted sum of the squared deviations of the data $(z_{i,j})$ from the model $z_i(t)$ as given by equation (2).

Applying the Model

Fitted Curves

Here we apply the model to data on TPCs for growth rate of fourth instar *Pieris rapae* caterpillars summarized (see fig. 2). There are 32 full-sib families, with two to 22 individuals measured per family (median sample size = 17/family), and relative growth rate (mass increase) was measured at six different temperatures (11°, 17°, 23°, 29°, 35°, and 40°C) during a 48-h period with a single larval instar for each individual (Kingsolver et al. 2004). Here relative growth rate is the proportional mass increase of an individual per unit time and has dimensions of time^{-1}. The result of the fit of the common template curve z for these data is

$$z(t) = \frac{256 + 5.45t - 5.77t^2 - 0.61t^3 - 0.02t^4}{100}. \quad (5)$$

Figure 4 (*top*) shows the results of fitted polynomials for each family; figure 4 (*bottom*) presents both the data (*solid lines*) and the fitted curves (*dashed line*) that have been discretized at the six measurement temperatures. Since the plots in figure 4 are on the same scale, we see that the fits are close to the data and that our model explains most of the variation in the data. Fits for each individual family curve (figs. A1, A2 in the online edition of the *American Naturalist*) indicate that the model fits very closely for most families. One family with maximum growth rate at 29°C (see fig. 4, *top*) is poorly fit by the model; this family has a sample size of 2 and thus does not contribute strongly to the model fit or decomposition. More generally, those families with rapid increases and/or decreases in growth rate at adjacent measurement temperatures had poorer fits, as expected for a polynomial model.

The bottom of figure 4 gives an important visual assessment of the residuals in the model and how well the template shape fits the data. In this figure, each data curve i was rescaled with respect to the three estimates of the parameters (height h_i, location of the maximum m_i, and width w_i), and all the rescaled discrete data are compared with the continuous common shape z; the remaining variation thus represents the residual variation not explained by the model. The fitted common shape is a good approximation of the common shape of the curves, and there is no apparent systematic variation that remains after fitting the model.
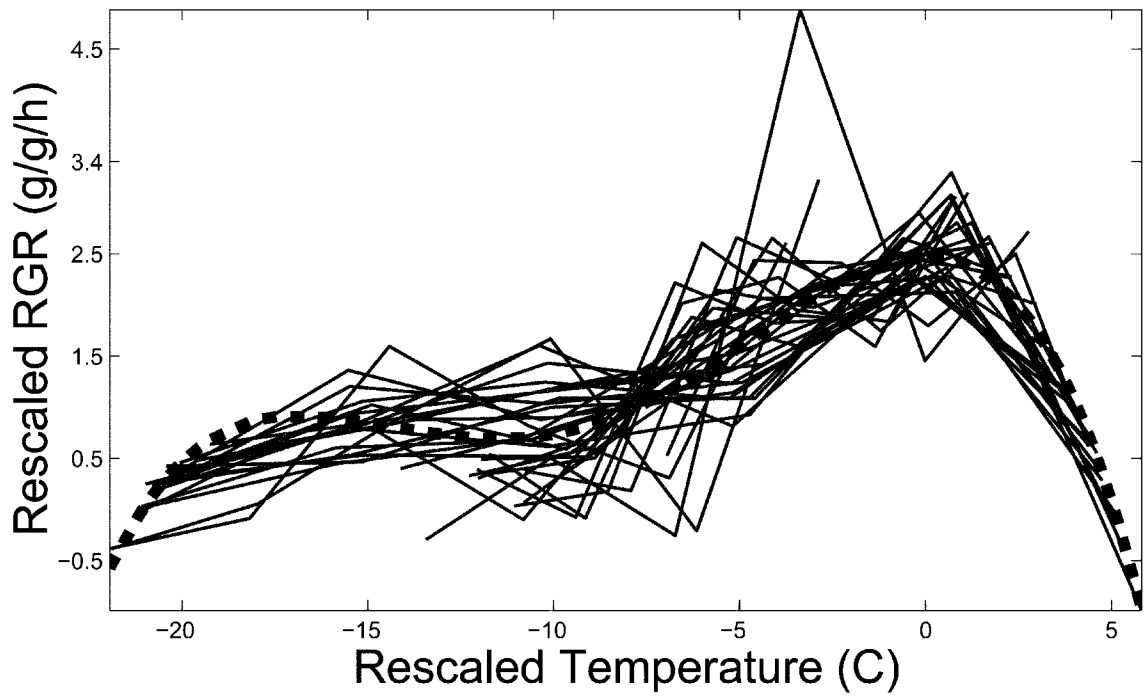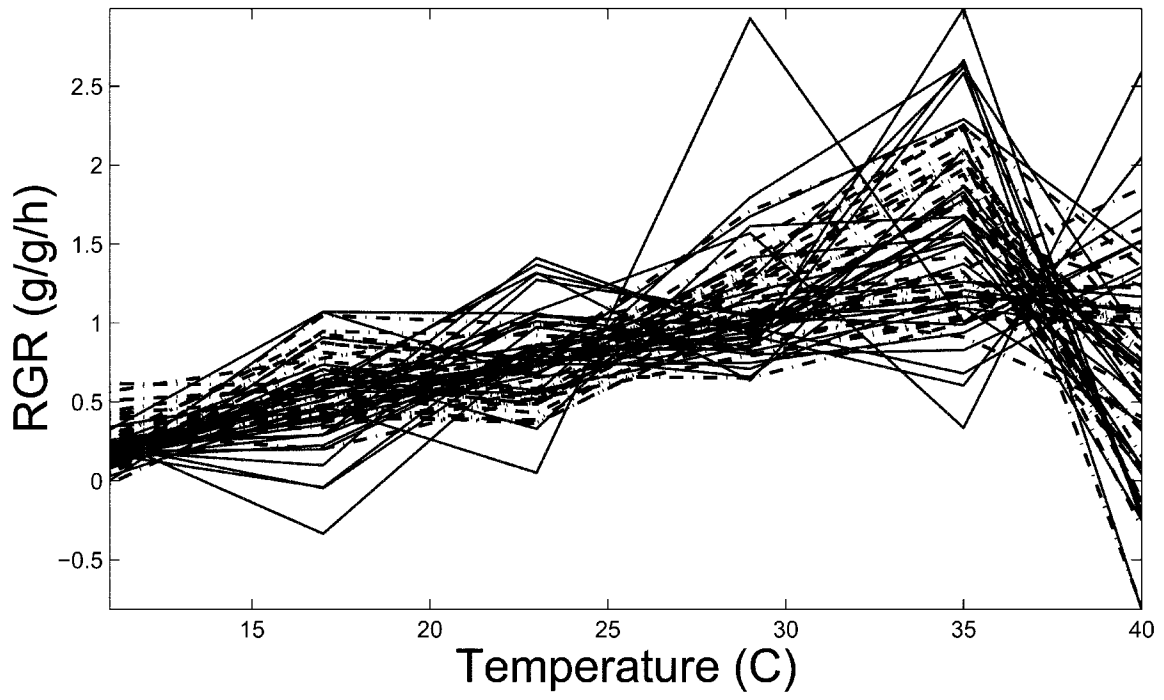
Decomposition of the Variation

The decomposition of the variation into the modes of interest by our TMV method (table 1) confirms our visual assessment of the fit. The model explains more than two-thirds (66.9%) of the total variation in the data. This high percentage shows the importance of these biologically interpretable modes of variation. A major biological result is that the generalist-specialist mode accounts for the largest component of variation (38.0%). Horizontal shift (16.2%) and vertical shift (12.6%) explain smaller fractions of the total variation in the data. Importantly, the model clearly distinguishes horizontal shift from generalist-specialist variation, both of which can generate complex patterns of positive and negative covariances across temperatures (see fig. 1 and "Discussion").

In equation (2), each parameter represents a mode of variation; that is, h_i represents the vertical shift, m_i represents the horizontal shift, and w_i represents the change in width and the specialist-generalist trade-off. The distribution of the parameters and the relationship between the parameters of variation are indicators of how much variation there is along each mode and how they co-vary (fig. 5). Recall that for our example data, height h_i represents the relative growth rate (g/g/h) averaged across all temperatures for family i (note that mean h_i over all families for the sample is by definition 0), m_i represents the optimal temperature (°C) at which relative growth rate is maximum for family i, and w_i represents the (dimensionless) width and generalist-specialist trade-off for family i. From these estimated parameters, we can also compute the maximal performance $z_{\max, i}$, the performance (here relative growth rate) at the estimated optimal temperature for family i. Several interesting patterns emerge from inspection of the parameter estimates (fig. 5). Estimated optimal temperature (m) varies between 29° and 46°C among families, although most families cluster between 32° and 38°C, with a second cluster of three families at 45°–46°C (see "Discussion"). Height (h) varies from −0.3 to 0.5 g/g/h, with most families between −0.2 and 0.2 g/g/h. In contrast, maximum performance ($z_{\max}$) varies over a much wider range among families (1.0–2.4 g/g/h). This pattern supports the conclusion that vertical shift (variation in h)

Table 1: Decomposition of the total between-family variation in relative growth rate for fourth instar *Pieris rapae* caterpillars, using template mode of variation

Mode of variation	Ratio of sum of squares (%)
Generalist specialist	38.04
Horizontal shift	16.23
Vertical shift	12.63
Model total	66.90

makes only minor contributions to overall between-family variation. There is no clear correlation between optimal temperature m and either maximum performance $z_{\max}$ or average height h.

Discussion

Biological Insights

In this article, we have described some of the limitations of applying standard linear methods to the analysis of nonlinear reaction norms and proposed a new nonlinear method that allows the decomposition and quantification of biological modes of variation for a set of continuous reaction norms. It is useful to compare our present results with more traditional quantitative genetic analyses of these same data. Kingsolver et al. (2004) estimated broad-sense genetic variances and covariances for growth rate across temperatures and applied PCA to the estimated **G** matrix and **G** function. Those analyses revealed significant genetic variation in growth rate at each temperature and that genetic variance increased with increasing temperature. They also demonstrated a significant negative genetic covariance between growth rate at 35° and at 40°C, suggesting a genetic trade-off in growth rate between these temperatures. However, neither **G** nor the principal components could be readily interpreted in terms of the biological hypotheses discussed here (fig. 2). Moreover, those analyses could not resolve whether the negative genetic covariance between 35° and 40°C represented a generalist-specialist trade-off

Figure 4: *Top*, caterpillar TPCs and fitted curves on the same scale. *Solid lines*, caterpillar's TPCs. *Dashed and dotted lines*, polynomial fits discretized at six temperatures. We can see that the range of variation in the polynomial fits closely matches the range of variation in the data, and the model explains most of the variability in the data. *Bottom*, rescaled caterpillar TPCs compared with the fitted common template shape $z(t)$. *Solid lines*, rescaled TPCs. Each TPC was standardized with respect to the estimates of height, location, and width parameters from the fit to model (eq. [2]). *Dashed line*, fitted common shape, polynomial of degree 4 (eq. [5]). This figure gives first an important visual assessment of the convergence of our fitting algorithm, because if our algorithm had converged to a local solution, the polynomial fit would not be in the middle of the data or the data would not be in the same scale in this representation. This figure also gives an important visual assessment of the residuals in the model and how well the template shape fits the data. We see in this figure that the fitted common shape is a good approximation of the common shape of the curves and that not very much of the variation remains after fitting the model.
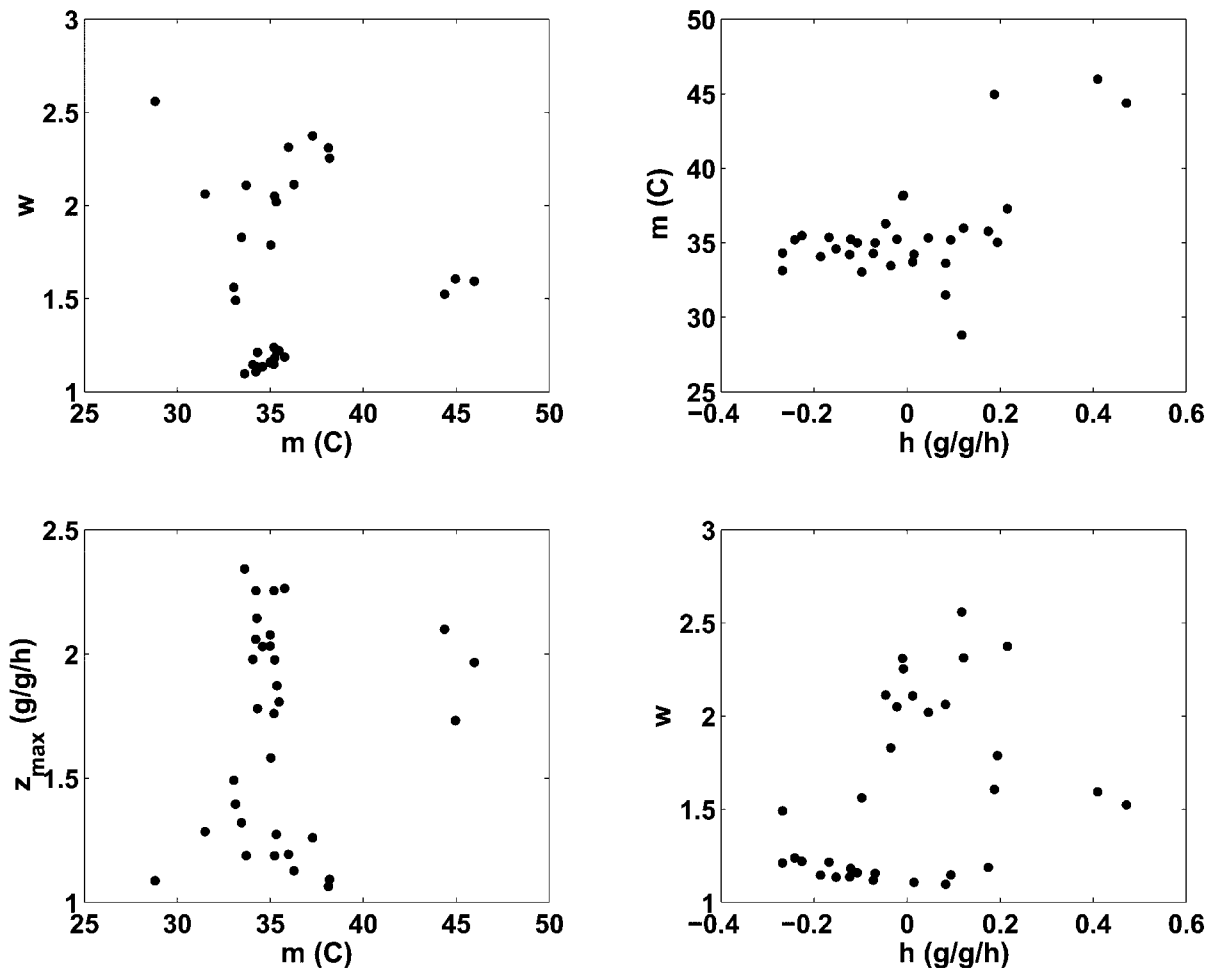
Figure 5: Values and relationships between the fitted parameters from the model. *Top left,* width parameter w versus location of the maximum parameter m (in °C). *Top right,* location parameter m (in °C) versus height parameter h (in g/g/h). *Bottom left,* fitted maximum growth rate z_{max} (in g/g/h) versus location m (in °C). *Bottom right,* width parameter w versus height h (in g/g/h).

or variation in the location of the thermal maximum, since both modes can generate negative covariances (fig. 1).

The new TMV method described here resolves these issues. The analyses decompose the variation into the three modes of interest and indicate that generalist-specialist variation makes the single largest contribution (38.0%) to the overall variation in TPCs for these data (table 1). Several previous studies have documented generalist-specialist trade-offs in performance across temperatures (Gilchrist 1996; Herron 1996); to our knowledge, this study is the first to quantify the contributions of this trade-off to the total genetic variation in TPCs or other continuous reaction norms.

The results indicate that horizontal shift—variation in the optimal temperature at which performance is maximized, m_i—contributes modestly (16.2%) to variation in

TPCs in this population (table 1). Estimated values of m_i for the families ranged from 29° to 46°C, although the majority of values fall between 32° and 38°C (fig. 5). Three families had estimated m_i of 44°–45°C, outside the range of measurement temperatures, because mean relative growth rate (RGR) for these families was greatest at the highest measurement temperature (40°C). The estimates of m_i for these families are necessarily imprecise because of this extrapolation. However, the alternative of using 40°C as an estimate of m_i in this situation is also problematic.

One surprising result of the analyses is the modest contribution (12.6%) of vertical shift—variation in overall performance across temperatures—to the total genetic variation in TPCs (table 1; fig. 5). Variation in overall "vigor" or performance has been shown to be a major component
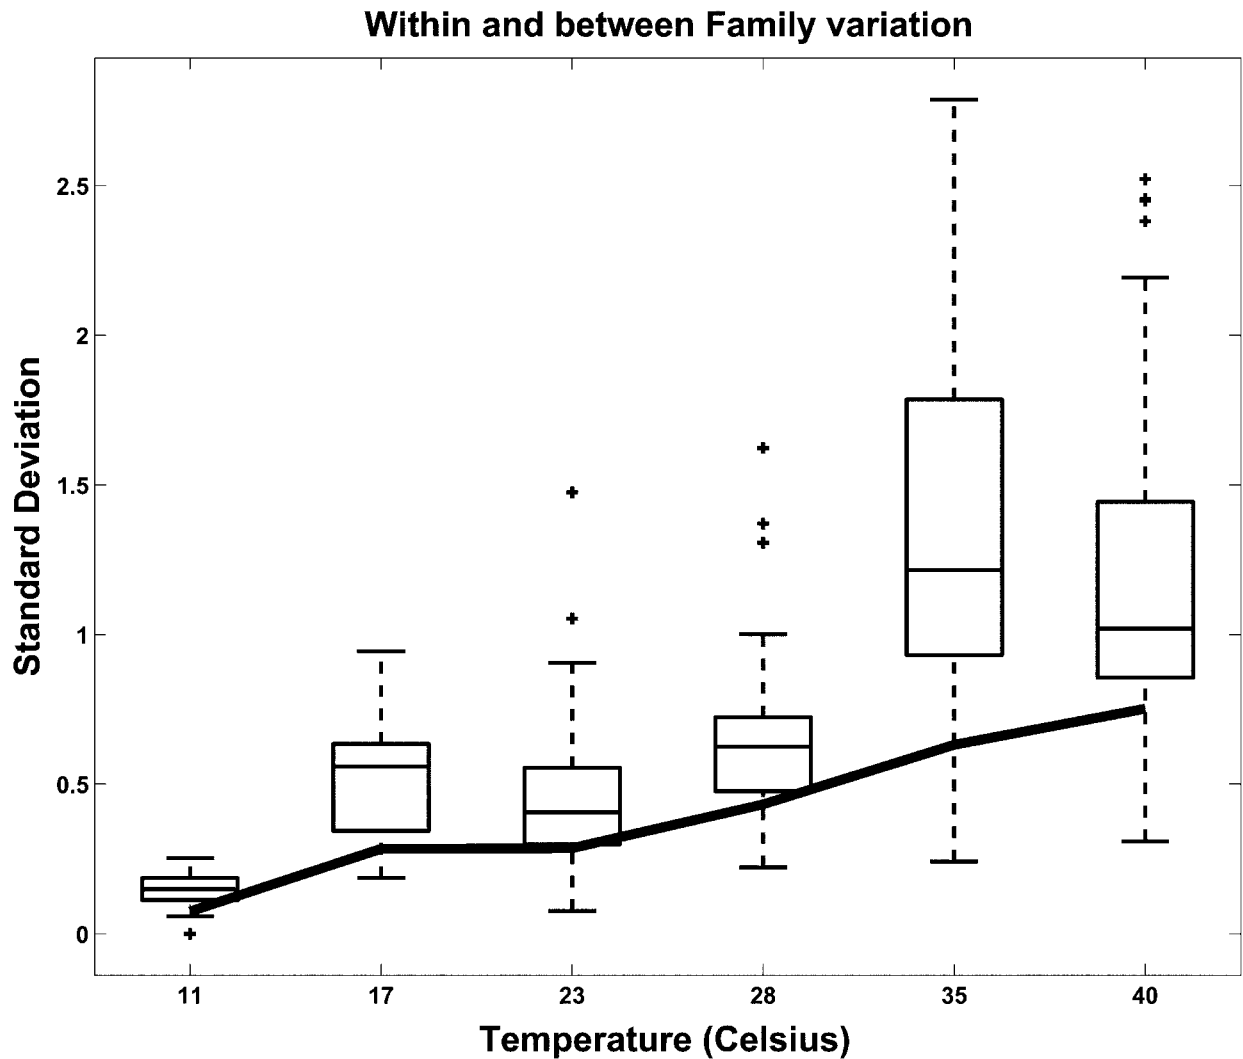
Within and between Family variation

Figure 6: Within-family variation compared with between-family variation. At each temperature, the box plot represents the variation in the 32 values of the within-family standard deviation at that temperature. The solid line shows the between-family standard deviation at each temperature. The between-family variation increases with the temperature, and so does the within-family variation. At all temperatures, most of the within-family variation exceeds the between-family variation.

of total genetic variation in many study systems and can obscure the contribution of genetic trade-offs to variation (Fry 1992). One interpretation of our results is that past directional selection for rapid growth rate at all temperatures in *Pieris rapae* has eliminated most of the genetic variation in vertical shift (Kingsolver and Gomulkiewicz 2003). As in many agricultural insect pests, *P. rapae* has the potential to grow very rapidly, completing larval growth (a mass increase of nearly 10,000-fold) in 3–4 weeks in some conditions (Kingsolver 2000).

There is much greater between-family variation in RGR at the optimal temperature (z_{max}) than in overall RGR

(height h) in this population (fig. 5). This suggests that there may be greater scope for evolutionary response to selection for rapid growth at optimal temperatures than for overall growth. Interestingly, neither z_{max} nor h is strongly correlated with optimal temperature m (fig. A2). This result does not support the "hotter is better" hypothesis, which posits a positive association between optimal temperature and maximum performance (Huey and Kingsolver 1989).

The analysis presented in this article focuses on the between-family variation in TPCs for short-term RGR. As illustrated in figure 6, the within-family variation is higher

than the between-family variation for most families. This is because the estimated heritabilities for RGR in this study were low, ranging from 16% at 35°C to 43% at 23°C (Kingsolver et al. 2004). Low heritabilities are not uncommon for short-term measurements of organismal performance, which are influenced by environmental variability, behavioral motivation, and measurement error (Gilchrist 1996; Roff 1997; Kingsolver et al. 2004).

Limitations and Applications of the Model

A key assumption of the TMV method is the existence of a common or template shape for all curves in the sample or population (Izem 2004). The choice of template is flexible, and the template shape is estimated from the data in this method. However, if the template does not adequately describe the curves for some families or genotypes, the estimated decomposition may be misleading. Inspection of the fit of the model for each family (figs. 4, A1, A2) is important for evaluating the adequacy of the template. For the *P. rapae* data, one family of very small sample size ($n_i = 2$) was poorly fit by the model (fig. 4); because the model fit is weighted by sample size, this had little impact on the estimated template or decomposition.

In the *P. rapae* data, we had six measurements for each curve, so we chose a polynomial of degree 4 to fit the data. The main advantage of using a polynomial template shape is its simplicity, which results in a quickly convergent fitting algorithm and a unique solution and avoids overfitting. However, polynomials of a lower degree do not fit rapid changes in the curves very well. Since the template shape is estimated from the data, the larger the number of temperature measurements, the better our estimate of the common shape can be. Since a polynomial of higher degree always fits the data better (i.e., has a lower SSE), we can base our template shape selection and decomposition on minimizing SSE and penalizing for the total number of parameters in the fit. If the number of measurements for each curve is higher, it is possible to fit the model with a template shape from a larger family than polynomials by using one of several nonparametric methods developed to estimate the common shape $z(t)$ in the model (Kneip and Gasser 1988; Wang and Gasser 1997, 1999; Kneip et al. 2000).

The TMV method described here can be used to analyze both phenotypic and genetic patterns of variation for a sample or population. It may also be used to estimate the model parameters (h, w, m) for different samples, populations, or experimental treatments. The current implementation of the method is directly applicable to performance curves and other continuous reaction norms for which there is a single, intermediate environmental value (e.g., temperature, nutrient concentration) at which the

phenotypic trait value (e.g., performance, size) is maximized (Izem 2004). We have developed an algorithm in Matlab to implement the model described here. The data required consist simply of the phenotypic trait values $z_{i,j}$ (e.g., performance) at a series of index values j (e.g., temperature) for each individual, family, or clone i. The output of this code includes the coefficients of the polynomial template shape; the fitted values of the three parameters of variation (w, m, h) for each individual, family, or clone; results of the decomposition of the variation along the three modes using TMV; and several graphics illustrating the fit. This algorithm has been successfully tested with several other TPC data sets, with encouraging results. The code used in the analyses presented here is available in the appendix in the online edition of the *American Naturalist*; the most recent implementation of the model and additional information about applications are available at http://www.fas.harvard.edu/~rizem.

With some modification, the general method may be extended to other types and shapes of function-valued traits (Kingsolver et al. 2001). For example, the horizontal and vertical shift modes may be appropriate for a wide variety of function-valued traits, including physiological and developmental reaction norms, age-specific life-history traits (e.g., mortality and fecundity curves), and growth trajectories. Such extensions would require identifying an appropriate template shape and defining and parameterizing modes of variation that are of biological interest.

Acknowledgments

Literature Cited

Fry, J. D. 1992. The mixed-model analysis of variance applied to quantitative genetics: biological meaning of the parameters. Evolution 46:540–550.

Gilchrist, G. W. 1996. A quantitative genetic analysis of thermal sensitivity in the locomotor performance curve of *Aphidius ervi*. Evolution 50:1560–1572.

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

Herron, J. C. 1996. Genetic variation, thermal sensitivity, and thermal acclimation in *Volvox aureus* and *Volvox globator*. PhD diss. University of Washington, Seattle.

Huey, R. B., and J. G. Kingsolver. 1989. Evolution of thermal sensitivity of ectotherm performance. Trends in Ecology & Evolution 4:131–135.

Huey, R. B., and R. D. Stevenson. 1979. Integrating thermal physi-

ology and ecology of ectotherms: a discussion of approaches. American Zoologist 19:357–366.

Izem, R. 2004. Analyzing variation in curves of common shape. PhD diss. University of North Carolina, Chapel Hill.

Izem, R., J. S. Marron, and J. G. Kingsolver. 2003. Quantifying nonlinear variation of modes of variation of curves. Proceedings of the International Statistical Institute 54.

Kingsolver, J. G. 2000. Feeding, growth and the thermal environment of cabbage white caterpillars, *Pieris rapae* L. Physiological and Biochemical Zoology 73:621–628.

Kingsolver, J. G., and R. Gomulkiewicz. 2003. Environmental variation and selection on performance curves. Integrative and Comparative Biology 43:470–477.

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

Kingsolver, J. G., G. J. Ragland, and J. G. Shlichta. 2004. Quantitative genetics of continuous reaction norms: thermal sensitivity of caterpillar growth rates. Evolution 58:1521–1529.

Kirkpatrick, M., and N. Heckman. 1989. A quantitative genetic model for growth, shape, reaction norms, and other infinite-dimensional characters. Journal of Mathematical Biology 27:429–450.

Kneip, A., and T. Gasser. 1988. Convergence and consistency results of self-modeling nonlinear regression. Annals of Statistics 16:82–112.

Kneip, A., X. Li, K. B. MacGibbon, and J. O. Ramsay. 2000. Curve registration by local regression. Canadian Journal of Statistics 28:19–29.

Lawton, W. H., E. A. Sylvestre, and M. S. Maggio. 1972. Self modeling regression. Technometrics 14:513–532.

Phillips, P. C., and S. J. Arnold. 1999. Hierarchical comparison of genetic variance-covariance matrices. I. Using the Flury hierarchy. Evolution 53:1506–1515.

Roff, D. A. 1997. Evolutionary quantitative genetics. Chapman & Hall, New York.

Via, S., and R. Lande. 1985. Genotype-environment interaction and the evolution of phenotypic plasticity. Evolution 39:505–522.

Wang, K., and T. Gasser. 1997. Alignment of curves by dynamic time warping. Annals of Statistics 25:1251–1276.

———. 1999. Synchronizing sample curves nonparametrically. Annals of Statistics 27:439–460.

Associate Editor: Edmund D. Brodie III
Editor: Jonathan B. Losos