7

Studying the Plasticity of Phenotypic Integration in a Model Organism

MASSIMO PIGLIUCCI

The Problem: Ecology of Phenotypic Integration

It should be clear from a survey of the heterogeneous literature on phenotypic integration that biologists still have to agree on exactly what integration is to begin with. Indeed, one of the aims of this book is to clarify the discussion on this essential aspect of the problem. Since we do not have a clear definition of integration, it may seem premature to ask other questions about it, such as how it evolves, what are its molecular underpinnings, and in what ecological context it needs to be framed. Yet, as the philosopher Ludwig Wittgenstein (1953/1973) has pointed out, most complex concepts are defined not by some essential property, but by a network of properties that intersect in different ways, depending on the viewpoint from which a given question is asked. Concepts whose exact nature eludes easy definition, often referred to as "family resemblance" concepts, may include phenotypic integration, and that may be why it is difficult to put a finger on what exactly the latter is. Moreover, scientists—by definition of what they do—cannot afford to be shy about investigating things that they cannot define in an essentialist fashion (just think of the amount of work done on speciation, despite the lack of a generally agreed-upon definition of species: Barraclough and Nee 2001; Schluter 2001). It therefore makes perfect sense to ask how to go about investigating the ecology of phenotypic integration even though we may disagree on the very definition of the concept.

This said, I still need to provide at least a general sense in which I am using both the terms "phenotypic integration" and "ecology" within the scope of this chapter. By *integration* I simply mean whatever set of evolutionary and developmental

processes result in an observable network of multivariate relationships among the phenotypic traits that define the morphology and life history of a living organism. Other chapters in this book deal with the problem of identifying the best ways to measure such relationships, but here I shall use a variety of statistical approaches that have been employed in the past to summarize the characteristics of multivariate sets of traits. In my experience, the exact choice of a metric is far less important then the choice of traits and the sample size used to study them. By *ecology* I mean the relationship between the networks of traits in question and the external environment—biotic or abiotic—in which the organisms that we study live. In other words, I am interested in what one can think of as the environmental contingency, and degree of phenotypic plasticity, of integration.

Phenotypic plasticity is a term normally used in reference to individual traits to summarize (by means of so-called reaction norm diagrams) how a given genotype (or population, or species) responds to a series of different environmental conditions by producing a more or less varied array of phenotypes (Fig. 7.1, left). It is therefore a logical extension of the concept to trace the same genotype-environment interaction in the case of complex measures of multivariate phenotypes, including whatever quantification of the amount and pattern of integration one wishes to use (Fig. 7.1, right). Consequently, most of the concepts, questions, and methodological approaches usually associated with the study of phenotypic plasticity (Pigliucci 2001) translate naturally to research on phenotypic integration (Schlichting 1989a, 1989b). This is important, because it provides us with a ready and tested toolbox from which we can choose to help build our understanding of phenotypic integration.

Now that we have set things up so that we are thinking of the ecology of phenotypic integration in a sense similar to the usual approach to genotype-environment interactions, we need to ask ourselves what it is that we want to know about the plasticity of integration. As in the case of phenotypic plasticity of individual traits, there is a risk of adopting simplistic null hypotheses (more on this below) and asking essentially trivial or irrelevant questions. For example, we

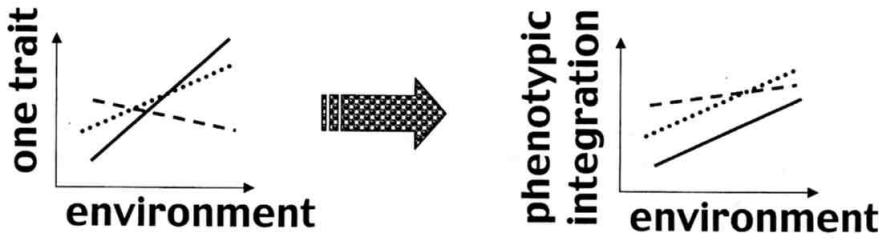

Figure 7.1 The conceptual shift from studying phenotypic plasticity of individual traits (left) to that of measures of associations among multiple characters (right) leaves unaltered the basic concepts, questions, and methodologies used in the study of phenotypic plasticity in general. The individual lines within each graph represent genotypes, populations, or species in whose plasticity one is interested. On the right, the vertical axis may represent a measure of the quantity or pattern of phenotypic integration in a group of organisms (population, species, etc.).

could wonder if there is genetic differentiation for the plasticity of integration among different populations and species. The answer ought to be yes, otherwise there would have to be a striking pattern of parallel covariances among traits in different species and environments that would be hard to miss even after a superficial glance at the natural world. Trying to be a bit more sophisticated, we could ask the same question for specific environmental factors and particular species, such as "Does the multivariate network of traits in population X of *Drosophila melanogaster* change in response to temperature?" But this would be only slightly better than the first case. The answer would probably be yes again (there are very few characters in very few species that do not show a degree of plasticity to at least some environmental factor) and it would therefore be of very little value for our understanding of the phenomenon of the modulation of integrated phenotypes in response to environmental changes. In all these cases, the null hypotheses (e.g., there is no plasticity of integration of traits X, Y, Z in response to environmental factors A, B, C) is just too simple for its (likely) rejection to be informative.

What I would rather suggest is that we need to start formulating more specific hypotheses that lead to more probing questions than we have largely been able to do so far. As any practicing organismal biologist will quickly recognize, however, this is much easier said than done. What sort of guidelines can one adopt to pursue this more informative search? How do we generate more interesting hypotheses that can make us feel that thousands of hours spent in the greenhouse are really worth the effort? I suggest that the answer lies along two distinct lines: on the one hand, we need to alter radically our ideas on hypothesis testing in general; on the other hand, we need to make intelligent use of what we already know about the ecology and genetics of certain species, especially of so-called model systems such as *Drosophila* or *Arabidopsis*. In the rest of this chapter I shall tackle first the philosophical and methodological issues related to hypothesis testing (which have a much broader application than just the study of phenotypic integration), then consider the use of a particular model system, its advantages and inevitable shortcomings, and conclude by examining in more detail two case studies in progress to illustrate some of the key concepts discussed in the chapter.

A Brief Detour on the Concept of Testing Hypotheses

So, what is wrong with standard hypothesis testing? A great deal, and evolutionary biologists (not just those interested in phenotypic plasticity or integration) have been unaware of a long discussion that has recently started to bear fruit, yielding significant changes in the *modus operandi* of several researchers in other fields (especially the social sciences). In what follows, I shall provide only a brief overview of the problem and of some of the proposed alternatives, since the topic of this book is phenotypic integration, not statistics or philosophy of science. Nevertheless, the references cited here should be very helpful to students of complex phenotypes whenever they find themselves at the point of formulating hypotheses, evaluating experimental designs, and conducting statistical analyses.

As unbelievable as it may sound, the standard statistical approach so common in the quantitative sciences is highly problematic and does not provide a very

useful tool for scientific (as opposed to statistical) inference. The main problem with the classical approach is that it involves testing the probability of the observed data given a rather arbitrarily constructed "null" hypothesis. That is, we typically set up statistical tests to tell us if, for example, the correlation coefficient between variables X and Y is zero (null hypothesis) or different from zero (any other value). We then collect the data and estimate the probability of getting those data if the null hypothesis were correct. If such probability is low, we happily "reject" the null hypothesis and proclaim the acceptance of the alternative hypothesis. Otherwise, we say that there is no evidence to reject the null, and sometimes even use statements that imply that the null has been confirmed (i.e., there is no effect of X on Y).

There are many problems with this (Gregson 1997; Dixon and O'Reilly 1999). For one thing, we are actually interested in determining the probability of a series of alternative hypotheses given the data, not vice versa. In other words, scientific inference is about comparing hypotheses on the basis of data, not the other way around (which is what we do when we use the standard statistical approach). Also, the null hypotheses are often uninteresting (why did you bother starting the experiment if you really thought that there was no correlation between X and Y?), and almost certain to be rejected given a large enough sample size. The latter is because the null hypothesis is stated very precisely (e.g., $r_{xy} = 0$, the correlation between X and Y is zero). In reality, what we mean to test is whether the correlation, measure of phenotypic integration, or whatever other value we are interested in, is close enough to zero to be scientifically uninteresting. But how close is close enough? Furthermore, the standard approach does not take into account any prior knowledge one may have about the system (like the expectation that not only will there be a correlation between X and Y, but that it will be positive, for example); but scientists do not start each project with a blank slate. It should also be obvious that the so-called alternative hypothesis is in reality an infinite family of hypotheses, which means that concluding that the result does not accord with the null is not very informative. Indeed, good science proceeds by considering several alternative hypotheses (Chamberlain 1897), not just two: the more alternatives, and the better defined, the more we can design informative experiments and make progress (Platt 1964).

There are serious problems also with the companion concept of "p-values," since in the standard statistical approach they are often misinterpreted in a variety of ways (Loftus 1993; Gregson 1997). The most important thing to understand is that the p-value is neither a measure of $P(H|D)$ (the probability of a hypothesis given the data, which is what is of interest to the scientist) nor (exactly) of $P(D|H_0)$ (the probability of the data given the null hypothesis, what we allegedly test with the standard approach). In fact, it is the cumulative probability of the point-null hypothesis and of all values more extreme than the one observed. This is related to $P(D|H_0)$, but clearly is not the same. Furthermore, it is absolutely incorrect to interpret a p-value as an estimate of the strength of the effect under investigation (as in "the smaller the p-value the more 'significant' the results are": significant in what sense, statistically or scientifically?). This is because for a fixed effect size (large or small), the p-value is a function of the sample size. One can repeat the experiment with a

much larger sample size, obtain a much smaller *p*-value associated with the null hypothesis, and yet the effect size (e.g., the strength of a correlation coefficient) can be unchanged (if it were properly estimated before).

The reader will notice throughout this discussion that many of the objections and alternatives to the standard approach to hypothesis testing are compatible with a Bayesian framework for scientific and statistical inference (Howson and Urbach 1991; Jefferys and Berger 1992). However, I shall not take that route here (even though I do consider it valid) for two reasons: first, there are several interesting philosophical problems raised by the adoption of a Bayesian framework (such as the estimation of prior probabilities) that I do not have space to discuss in this chapter. Second, it is currently very difficult to utilize the Bayesian insight in practice because of the lack of general-purpose software. One has instead to write custom programs tailored to each specific problem, something that people have started doing in a variety of fields within organismal biology, from systematics (Huelsenbeck et al. 2000), to genetics (Shoemaker et al. 1999), to evolutionary biology (Rudge 1998), to conservation biology (Wade 2000). Here I shall instead briefly examine some alternatives to the standard approach that are of immediate application for any researcher who is familiar with widely available statistical packages.

One powerful alternative practice to the standard approach is simply to think in terms of multiple hypotheses, none of which gets to play default. Interesting and informative scientific papers are those that consider two (or more) plausible hypotheses and can actually conclude for an increase or decrease in likelihood of one or more of them (Chamberlain 1897). Most scientists of course do realize the distinction between statistical and scientific inference but, because of the constraints (imposed by editors and reviewers, as well as by available textbooks and software) favoring standard statistical analyses, they end up writing rather schizophrenic papers in which the Results section is presented in terms of uninformative *p*-values while the Introduction and Discussion sections are where the real theoretical action lies. There are also better ways of reporting the relationship between the actual results one obtains and the predictions of the alternative hypotheses. A simple but often-neglected approach is to plot the data and their confidence intervals. As Loftus (1993) puts it, "a picture is worth a thousand *p*-values." Loftus provides examples from the primary literature from which it is clear that a graph of means and standard errors gives us all the information that a *p*-value provides, and then some. Moreover, the really important information can be gathered only from the graph, while the rest is not that interesting scientifically. In particular, graphs of means and their associated measures of dispersion provide us with a visual estimate of effect sizes, of the power of the analysis (which is inversely proportional to the standard errors), and with an immediate understanding of the patterns identified by the data. Loftus points out that these quantities are completely invisible in a table of *p*-values, where all we are told is that certain, largely irrelevant, null hypotheses are to be rejected. Furthermore, even in those cases in which the null is not rejected, Loftus reminds us that we can immediately discern from a plot if this is because there is little power in the analysis (large standard errors when compared to the differences between means), or because the population means are really well estimated, and indeed

close to each other. This distinction is important, because in the first case we deduce that we need more data, in the latter that we really have no reason to think there is a scientifically relevant difference among means.

The next alternative to consider combines the insights of the Bayesian approach (*sans* the difficulties related to the estimation of priors) and the simplicity of already available statistical approaches (and therefore of computer packages): maximum likelihood ratios (Dixon and O'Reilly 1999). A maximum likelihood ratio is a quantity that compares the likelihood of the data based on one model with the likelihood of the data based on a second, competing, model (comparisons can also be made among multiple models, pairwise or cumulatively). The ratio therefore provides a measure of the relative match of models and data. Most sources suggest that a ratio of 10:1 is equivalent to a p-value of 0.05, that is, it should be the minimum ratio at which one should consider one model better than another. However, remember the fact that any such threshold (including the magical p-value of 0.05) is arbitrary and that there is no shortcut to using one's own judgment. The maximum likelihood ratio is obviously conceptually (and mathematically; see Dixon and O'Reilly 1999) related to the Bayesian approach, but it is also conveniently close to the standard statistics we are all familiar with, which makes it a powerful practical alternative. There are several ways to calculate likelihood ratios from standard statistics, which are explained in detail by Dixon and O'Reilly: all one needs are standard estimates of variances obtainable from ANOVA or regression output tables.

Yet another approach to the solution of the same problems concerning model selection is provided by the use of the Akaike Information Criterion (AIC). This is a measure based on information theory and derived from the concept of entropy in physics (Anderson et al. 2000). Briefly, AIC measures the fit of the data to a given model, when the model is penalized in proportion to the number of parameters it employs. This is necessary because, other things being equal, a model with more parameters will always fit the data better than one with fewer parameters, but the improved fit will not always be scientifically relevant. The Akaike Information Criterion and its derivates (which, like likelihood ratios, can be calculated from standard statistical quantities) lend themselves very easily to the comparison of several alternative models, while not being at all "tests" in the sense of requiring the calculation of associated p-values. The AIC also benefits from solid theoretical and philosophical foundations given its relationship to information theory.

At the end of this brief tour one obvious question comes to mind: should we throw away everything that has been published using the standard statistical methods? Fortunately, no: it turns out that in many cases inferences based on the classical p-value-based approach are similar to those arrived at using the alternative methods described here. However, there are some important differences that are especially useful to quantitative biologists, particularly in fields such as phenotypic integration where hypotheses are complex and sample sizes may vary dramatically from study to study. First, there are cases in which the two approaches will yield very different results. In particular, as we have seen, p-values are problematic when sample sizes are either too small or very large. In the first case, a failure to reject the null hypothesis may actually be premature, since the

problem may be the low power inherent in small sample sizes. The *p*-values do not carry any information about this problem, and one needs to plot the data and their confidence intervals to get an idea, as mentioned above (just reporting the sample sizes will not do, because that information does not provide clues to the relevant measures of dispersion of the data). In the second case, *p*-values may lead us to reject the null when the difference between effects is scientifically irrelevant, though technically nonzero. As we have seen, this is because p-values refer to a point-hypothesis, which is never of interest to scientists; any small, but systematic, unaccounted source of variation will produce deviations from the point-null which will be detectable for large enough sample sizes. In this case again, plotting the data will reveal more readily both the statistical and the scientific significance of the results.

The second reason to opt for alternatives to standard hypothesis testing is a matter of training ourselves to think along realistic, and inevitably complex, lines: typically, we are simply not interested in the hypotheses tested with the standard methods. Again, good science actually proceeds by comparing the relative fitness (using the data) of a series of reasonable alternative models (hypotheses), something that standard hypothesis testing simply cannot do.

It is now time to turn to the second source of help that students of integration can use to make some progress in this murky field: model systems.

The Use of Model Systems in Evolutionary Ecology

The virtues and benefits of the use of model systems in biological research have been discussed (Kellogg and Shaffer 1993), and a reasonable position to take is that while they provide certain advantages that are impossible to gather from the majority of potential experimental organisms, there are some serious limitations to the generalizations one can make from research conducted on *Drosophila*, *Caenorhabditis*, *Arabidopsis* and the few other species of choice that have emerged over the last few decades.

On the one hand, it is undeniable that remarkable progress has been made, especially in molecular, but also in organismal biology, with the use of model systems. It is simply not possible to imagine modern genetics without *Drosophila*, and we owe much of our understanding of flower development to the famous ABC model developed in *Arabidopsis* (Pidkowich et al. 1999). The very idea of a "model system" is central to that part of biology most closely resembling research in physics and chemistry: molecular biology is the search for universals analogous to the structure of the atom and the principles of thermodynamics. In this framework, it is logical to focus the attention on a few, easily manipulated systems, and attempt to generalize to the rest of the living world.

On the other hand, organismal biology (ecology and evolutionary biology) is a much messier business, dominated by the importance of historical events, where organisms differ from each other in important respects, and where the focus is on understanding the differences, not the similarities. In this context, concentrating our attention on a few (often unrepresentative) species can be seen as a fatal flaw that will lead us into the false security of very narrow understanding.

While this distinction between ahistorical and historical science is indeed of fundamental interest, and does have important consequences for our understanding of the very nature of science (Dupré 1993), I think that a useful middle ground can be found. I shall illustrate the advantages of such an intermediate position for the study of phenotypic integration in the remainder of this chapter using research on *Arabidopsis thaliana*, from which discoveries can easily be extended to its close phylogenetic relatives.

Arabidopsis has been a model system in molecular, developmental, and cell biology for quite some time now (Griffing and Scholl 1991; Pyke 1994; Anderson and Roberts 1998), but several researchers have recently also been proposing it as an interesting tool for quantitative genetics, ecology, and evolutionary biology (Pigliucci 1998; Alonso-Blanco and Koornneef 2000; Mitchell-Olds 2001). While *A. thaliana* is a highly selfing plant, several of its relatives are not, which allows us to study the evolution of mating systems and of their relationship to the ecology of a group of taxa. Furthermore, from an ecological perspective *A. thaliana* can be described as an opportunistic weed, a plant of ruderal habitats (Ratcliffe 1965; Napp-Zinn 1985). It grows along a huge latitudinal span (from northern Africa to the polar circle), and it is found in a surprising number of soil types in varying association with other vegetation. The ecology of the other species of this group is much less understood, but the clade includes alpine species, annuals, biennials, and perennials, as well as species with a different base chromosome number and even hybrids and polyploids (Mummenhoff and Hurka 1995; Lee and Chen 2001). For a long time the systematics of this group was confusing (Price et al. 1994; O'Kane and Al-Shehbaz 1997), but recent morphological and especially molecular research has significantly cleared up the relationships among many of the more or less close relatives of *A. thaliana* (Al-Shehbaz et al. 1999; Koch et al. 1999, 2000, 2001). Currently a robust phylogenetic hypothesis is available (Fig. 7.2) to provide an invaluable comparison base for research on the evolution of phenotypes in this clade using standard phylogenetic comparative methods (Harvey and Pagel 1991; Huelsenbeck et al. 2000; Martins 2000). Examples of this in *Arabidopsis* have already appeared in the literature (Pigliucci et al. 1999, 2003).

Having made my case for a judicious use of *Arabidopsis* as a model system in organismal biology, it is now time to turn to a few examples of ongoing research on phenotypic integration that illustrate several of the concepts I have been discussing so far.

Case Study 1: Bushy Phenotypes in Response to Touch

An interesting and understudied type of phenotypic plasticity affecting several aspects of the phenotype is the so-called thigmomorphogenic response, that is, the reaction that many plants display to mechanical stimulation (Biddington 1986; Jaffe and Forbes 1993; Wisniewski 1996). The ecological context of thigmomorphogenesis can be varied and complex in itself, as this plasticity can be a reaction to abiotic mechanical stimulation as imposed by wind, rain, or snow (Braam and Davis 1990; Telewski and Pruyn 1998; Cordero 1999), to contact

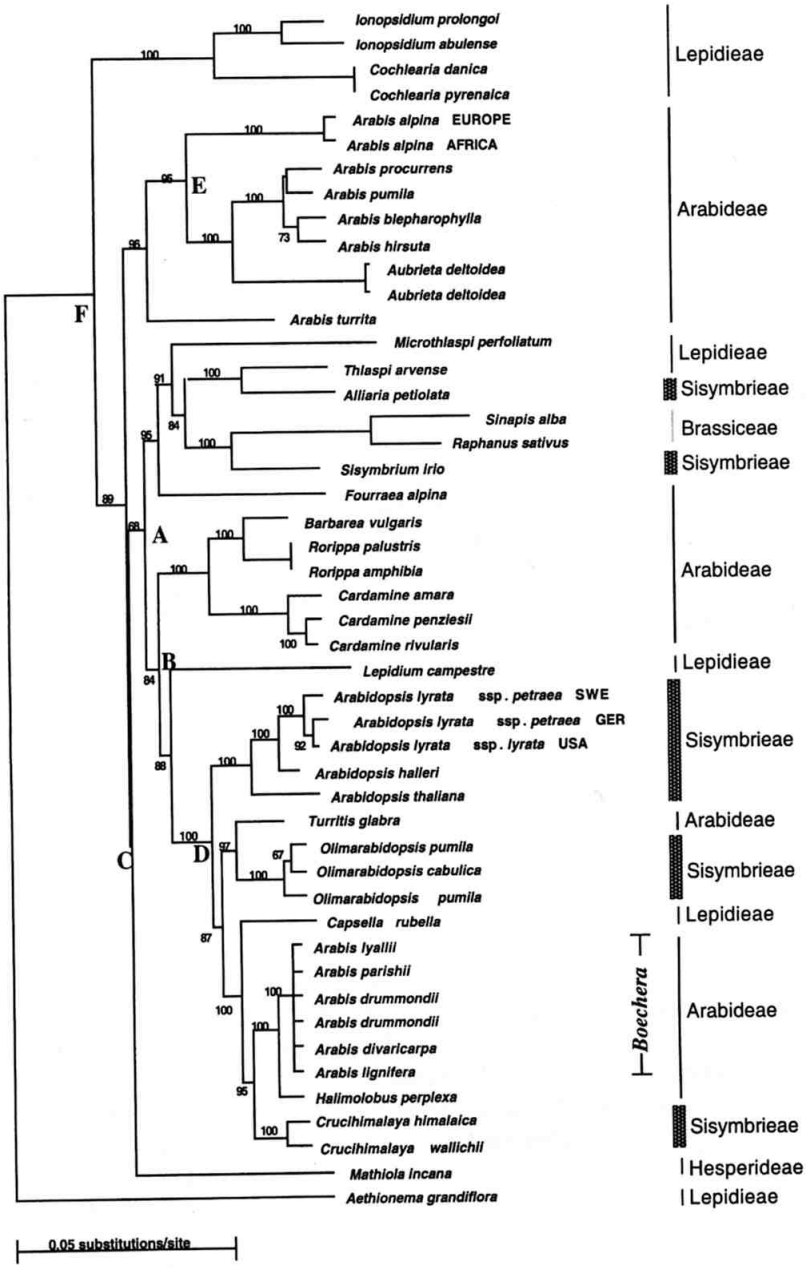

Figure 7.2 Neighbor-joining distance tree of *Arabidopsis* and some of its close relatives based on *matK* and *Chs* sequences. Bootstrap values in percent from 1000 replicates are shown above branches. Tribal assignments are given in the right margin. Approximate divergence dates (million years before present) for nodes A–F are: A, 16–21; B, 13–19; C, 19–25; D, 10–14; E, 15–17; F, 26–32. Several important taxa are not shown in this figure: *Arabidopsis arenosa* and *A. suecica* are closely related to *A. lyrata*. *Brassica* is near *Raphanus* and *Sinapis*. Finally, *Leavenworthia* is related to *Barbarea*. Reproduced with permission from Koch et al. (2001).

by insects (van Emden et al. 1990; Cipollini 1997), or to contact by other plants in crowded stands (Emery et al. 1994). It has even been suggested that touch-response plasticity may have practical applications, since in some plants it increases resistance to pests (Cipollini 1997).

The molecular basis of touch-response plasticity has been investigated, especially in *Arabidopsis*, and a series of relevant genes has been identified and is currently under study (Braam and Davis 1990; Sistrunk et al. 1994; Xu et al. 1995, 1996; Mizoguchi et al. 1996; Antosiewicz et al. 1997; Braam et al. 1997). Braam and collaborators (1997) found that the *TCH* (touch) gene products in *Arabidopsis* include a xyloglucan endotransglycosylase and several calmodulin and calmodulin-related proteins. They found that the functions of these proteins comprise cell and tissue expansion under mechanical strain, and hypothesize that a subset of *TCH* genes may be involved in cell wall biogenesis, all of which makes sense in light of their relationship to phenotypic plasticity triggered by mechanical stimulation.

I studied thigmomorphogenesis in response to winds sustained for different periods of time in 11 populations of *A. thaliana* to determine the degree of genetic differentiation among populations in plasticity to mechanical stimulation, as well as to test the hypothesis that recognizable "wind-specialized" populations exist and are characterized by an identifiable pattern of phenotypic integration (Pigliucci 2002b). More specifically, I hypothesized that when plants are exposed to sustained winds, populations can be grouped by their multivariate response into "bushy" and "nonbushy." If true, the clustering of populations could then be used to predict which populations are actually found in high-wind environments in the field (a piece of information that is still currently missing from the puzzle).

I did find that 8 of the 11 populations responded to sustained winds, especially by increasing their branching intensity, which made them look more "bushy." The remaining 3 did not respond at all to the environmental stimulus (Fig. 7.3). More interestingly, the multivariate pattern of phenotypic integration—measured by the correlation matrices among all traits—clearly separated bushy from nonbushy phenotypes (Fig. 7.4). This result confirmed the existence of a recognizable phenotypic syndrome related to mechanical stimulation under sustained winds, but not under wind-free conditions. Since I had previous experience with *A. thaliana*, I had also set up a series of three alternative models predicting the patterns of phenotypic integration based on my knowledge of the biology of this species and on basic considerations of plant physiology. The three hypotheses were set up as matrices of character correlations which where then compared with the observed matrices using standard matrix correlation tests (Cheverud et al. 1989). The hypothetical matrices assumed one of three scenarios: (1) integration within trait classes, where characters within vegetative, architectural, and reproductive classes are hypothesized to be highly positively correlated, but independent of characters in the other classes; (2) developmental integration with no tradeoffs, in which traits expressed contiguously (in time) during the life cycle are assumed to be positively correlated with each other, but the correlation decays with time (i.e., between traits expressed very early and very late); and (3) developmental integration with tradeoffs, similar to scenario (2) but where tradeoffs (manifested as negative correlations) are predicted between vegetative and reproductive traits, as well as between different reproductive traits (such as main versus

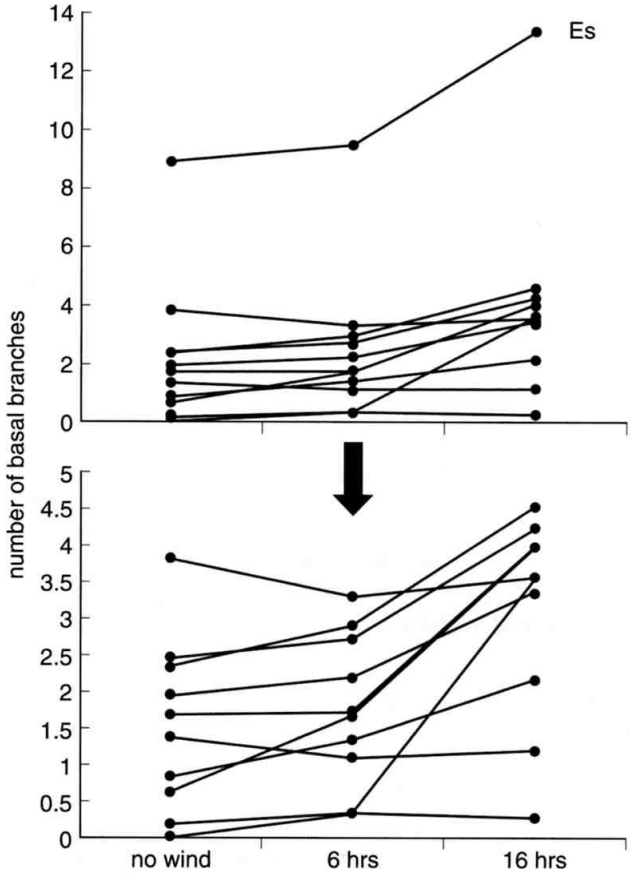

Figure 7.3 Reaction norms of individual populations of *Arabidopsis thaliana* graphing the response of branching to exposure to wind for increasing amounts of time. Each line represents a different population. The lower graph is a zoom into the lower level of the upper one with the exclusion of the late-flowering population Es (Espoo, from Finland), which produced many more branches than the others regardless of environmental conditions. (From Pigliucci, 2002b.)

basal branching). Models were compared by means of likelihood ratios, one of the alternative statistical approaches mentioned above. I found that under all three environmental conditions the third model was the best predictor of the observed patterns of integration, with amounts of explained variance ranging from 63% under no wind to 74% for 6-hour wind and 52% for 16-hour wind.

This study pointed out a series of interesting things about the plasticity of phenotypic integration when the response of *Arabidopsis* to mechanical stimulation is concerned. First, it is possible to predict the general pattern of phenotypic integration as reflected in character correlations, if one has a minimal knowledge of the life cycle and physiological characteristics of the organism in question. (This has been done in other systems for which prior knowledge of the developmental trajectory was available: Cheverud 1996; Mezey et al. 2000). Second,

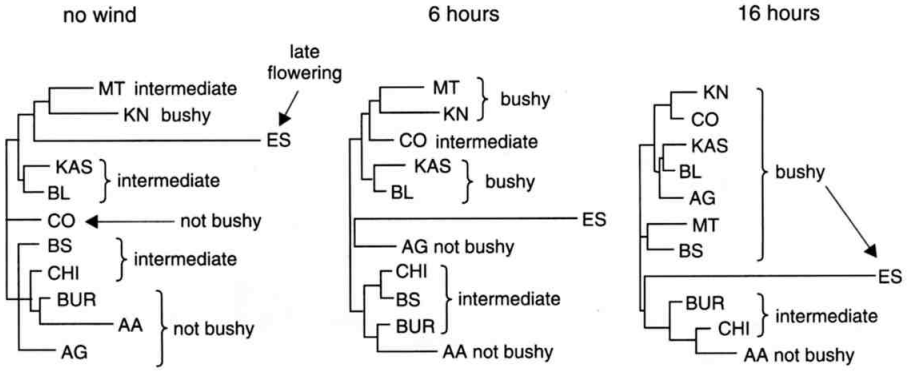

Figure 7.4 Multivariate clustering of 11 populations of *Arabidopsis thaliana* according to their patterns of phenotypic integration (measured as matrices of character correlations) within each of three environmental treatments. Notice how the populations characterized by the plastic response that alters the degree of their branching are scattered throughout the dendrogram on the left (no wind) but become increasingly separated as a group with more and more sustained winds (especially under 16 hours of exposure, right). The clustering was obtained by computing additive trees based on Euclidean distances calculated on phenotypic data. (From Pigliucci, 2002b.)

phenotypic integration can nevertheless be substantially altered by environmental changes. Third, if the environmental stimuli can be connected even roughly with the field ecology of the species under investigation, one can in fact predict the existence of specific patterns of integration, predictions that can later be used to test for the existence of a link between phenotypic syndromes expressed under controlled conditions and prevalent environmental regimes in the field. To show that making sense of phenotypic integration in this way can be applied to other circumstances, I shall now turn to our second case study, in which the same species (but different populations) is examined in response to two more environmental factors relevant to its ecology.

Case Study 2: No Stress, One Stress, Two Stresses

Two of the most common sources of environmental variation in most plants are heterogeneities in the levels of water and light. *Arabidopsis thaliana* occasionally grows under a relatively dense canopy of short grass, and more often has to deal with competition for light from shorter neighbors of the same or others species (such as *Trifolium*, pers. obs.). As far as water is concerned, while certainly not likely to experience prolonged conditions under flood, *A. thaliana* nevertheless lives in habitats such as ditches along roads where water can accumulate to the point of saturating the soil and causing anoxia in the roots and possibly the leaves.

As in the case of mechanical stimulation, there is a large literature on the molecular basis of responses to light intensity (Yanovsky et al. 1997; Whitelam and Devlin 1998; Lasceve et al. 1999) and water (Wang et al. 1995; Ishitani et al.

1997; Zhang et al. 2000) in *Arabidopsis*. As far as responses to light are concerned, *Arabidopsis* is equipped with a battery of photoreceptors, some sensitive to red and far-red light (phytochromes), some to the blue and UV components of the spectrum (so-called cryptochromes and UV-receptors). Phytochromes, especially phytochrome A and B (there are five in *Arabidopsis*), are the entry point for a series of distinctive responses known as the very-low-fluence, the low-fluence, and the high-irradiance responses (Yanovsky et al. 1997). The blue-light-activated signal transduction pathways themselves are of at least four different kinds, each regulated by a distinct photoreceptor as the environmental signal is received by the plant (Lasceve et al. 1999). Furthermore, we are beginning to unravel the complexity of the transduction pathways connected to these two large families of photoreceptors, and the emerging picture is one that includes a battery of repressors (the so-called *COP/DET/FUS* genes) acting downstream of multiple photoreceptors, which in turn regulate the cellular levels of hormones such as cytokinins and brassinosteroids (Whitelam and Devlin 1998). Concerning the molecular basis of responses to levels of water, Zhang et al. (2000) have demonstrated with the use of transgenic *Arabidopsis* that endogenously produced cytokinin increases flood tolerance in this species, probably in part by extending the length of the part of the life cycle between flowering and senescence, thereby allowing plants to mature more seeds under stressful conditions. Ishitani and coworkers (1997) have identified a number of mutants that affect osmosis, and therefore the reaction to water stress (drought, in this case) in *Arabidopsis*, and have argued for the existence of extensive cross-talk between pathways that are dependent of abscisic acid and those that are independent of this hormone, contrary to what was previously thought.

In the context of phenotypic integration, Ania Kolodynska and I set out to test several hypotheses about how patterns of phenotypic integration in *A. thaliana* could be altered by simultaneously varying light and water levels. In particular, we were interested in the multivariate phenotype expressed under normal levels of both environmental factors, as opposed to situations in which one or the other stress (low light or high water) was experienced, and then to a scenario in which both stresses were applied simultaneously (Pigliucci and Kolodynska, in preparation). We predicted, in agreement with other suggestions published in the integration literature (Schlichting 1986, 1989a; Chapin 1991) that the highest degrees of stress would result in a more "tight" phenotype, that is, in the expression of a higher number of correlations of large magnitude between different traits and reproductive fitness. We also made a number of specific predictions about the behavior of individual traits in response to the four imposed combinations of light and water levels. While the latter were largely confirmed, I shall focus here only on the results pertinent to phenotypic integration and not on those concerning individual traits. We obtained very clear support for our hypothesis of a positive relationship between degree of stress (as perceived by the plant and manifested in decreased fruit production) and phenotypic integration as gauged by the number of interactions between phenotypic traits and reproductive fitness, itself tested by fitting different models to the data and choosing based on the corresponding likelihood ratios (Fig. 7.5). While under nonstressful conditions the only variable that seemed to affect fitness was shoot biomass, under moderate stress (either low light

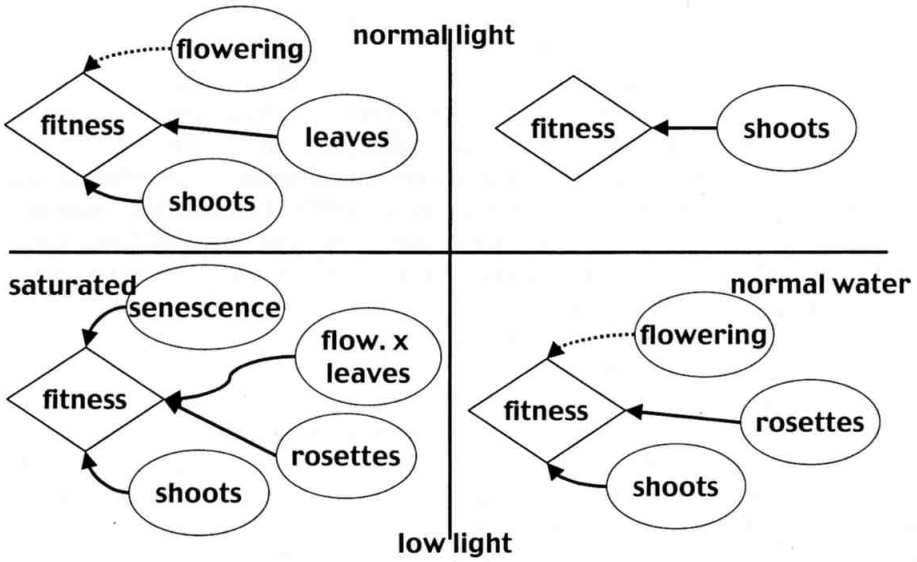

Figure 7.5 Patterns of phenotypic integration in 22 populations of *Arabidopsis thaliana* in response to four combinations of water (horizontal axis) and light (vertical axis) levels. Each quadrant shows the best model that fits the data according to a maximum likelihood ratio analysis of several alternative models. Notice how the number of traits affecting reproductive fitness (number of fruits produced) increases from no stress (upper right) to moderate stress (upper left and lower right) to high stress (lower left). Broken lines indicate a negative relationship between a given trait and fitness. "Flow. × leaves" indicates an interaction between leaf production and flowering time.

or water-saturated soil) flowering time, and either leaf number or size, were important. When the stress was highest, new traits came into play (such as the time to senescence), as well as higher-order interactions among individual characters (e.g., the particular combination of flowering time and leaf number became important).

While we do not have a good enough understanding of the ecophysiology of *A. thaliana* to allow us to predict in advance which combinations of characters will be more or less relevant under certain conditions, some of the changes in the patterns of integration are readily explainable. More importantly, such explanations can be submitted to further experimental tests. For example, if the stress was induced by water, with light at normal levels, leaf number was more important than leaf size; however, under low light it was the size, not the number, of leaves that was important. This is understandable if one thinks that under low light an increase in photosynthetic area is a more decisive parameter, because it compensates directly for reduced photosynthetic activity without costing the commitment of additional meristems to leaf production. It is the combination of such first principles of physiology and knowledge of the ecology of the particular species that it is more likely to generate detailed and testable hypotheses concerning phenotypic integration. The road ahead of us, however, is still fraught with many difficulties, as I shall discuss in the concluding section of this chapter.

What We Need to Do

For a long time molecular and organismal biology have proceeded largely independently of each other, and have managed to make significant progress (Pigliucci in press). In recent years, however, a significant amount of interdisciplinary research has been carried out, largely centering on model systems such as *Drosophila*, *Caenorhabditis*, and *Arabidopsis* (Schlichting and Pigliucci 1998, chap. 8). This multi-pronged approach is likely to be useful also within the field of phenotypic integration, providing us with valuable insights into the molecular basis of complex phenotypes. Examples of these approaches include QTL mapping and the study of the pleiotropic effects of candidate genes for important regulatory functions (Juenger et al. 2000; Mezey et al. 2000), and robust phylogenetic hypotheses to be used in comparative studies of the evolution of phenotypic integration (Losos and Miles 1994; Ackerly and Donoghue 1998). For example, the phylogenetic hypothesis of the *Arabidopsis* and closely related clades resulting from extensive work by Mitchell-Olds's laboratory (Koch et al. 1999, 2000, 2001; see Fig. 7.2) will allow countless detailed comparisons of evolutionary patterns in this group to be carried out in a fashion that is informed by the likely evolutionary history of these species.

This optimism notwithstanding, there are some cautionary considerations to be made. First, it should be obvious to every evolutionary biologist that our discipline in general (not just the study of phenotypic integration) is a strange mixture of historical and experimental science (Cleland 2001). We can carry out experiments on phenotypic integration under controlled conditions (in the laboratory, greenhouse, or field), but the results will be highly dependent on which particular species, populations, and genotypes we have chosen. The quest to make evolutionary biology into a "hard" science on the model of physics or chemistry (or even, to some extent, molecular biology) is simply a misguided one (Pigliucci 2002a).This means that not only the study of model systems, but the study of *any* particular group of organisms, will not likely yield results that can be generalized beyond fairly narrow taxonomic or life-history categories. As evolutionary biologists we should embrace this as a natural outcome of the diversity of life, not as a limitation on our ability to conduct science.

Similarly, perhaps talk of "phenotypic integration" is itself misleading, as is talk about "phenotypic plasticity," without reference to specific combinations of traits and environments. Bradshaw (1965) pointed out decades ago that plasticity is not a meaningful property of a genotype, but it can be different for different traits, and even for the same trait when expressed under different environmental circumstances. Similarly, phenotypic integration may not intelligibly be attributed to the totality of an organism's traits, but make sense only for certain groups of characters (modules) that have been subjected to joint selection, or that have recently ceased to be coupled functionally and are now evolving independently. There is no reason to think that natural selection works on the full phenotypic variance-covariance matrix of an organism any more than we are well grounded in thinking that it operates simultaneously to coordinate the plasticities of all traits (and of course, there is the additional problem

of what exactly a trait is, a discussion that recently has deserved a whole book in itself: Wagner 2001). Rather, we should focus our attention on groups of characters about which we have reasonable grounds to think that they are functionally integrated (Klingenberg and Zaklan 2000; Mezey et al. 2000; Wagner and Schwenk 2000; Klingenberg et al. 2001; Magwene 2001), or that may have evolved in concert as a result of the particular genetic architecture of the organism under study. Generalized measures of integration that are calculated across the whole organism are likely to be uninformative because they result from averaging the effects of distinct processes. It would be as if we calculated an "average plasticity" for a given genotype.

This problem has already emerged clearly within that part of phenotypic integration research that deals with the study of genetic variance-covariance matrices (Jernigan et al. 1994; Roff 1996; Lynch 1999). Authors write about the evolution of "**G**-matrices" as if these were a homogeneous biological class (what philosophers of science would call a natural category). But they are clearly not, because **G**-matrices can be measured in phylogenetically highly divergent taxa, and different authors concentrate on distinct subsets of traits, which can in no biologically meaningful sense be considered "samples" of the whole **G**-matrix. The result is that comparative studies of **G**-matrices among disparate biological entities tend to be highly uninformative (Roff 1996; Roff and Mousseau 1999). Rather, one should concentrate on closely related taxa and on groups of characters that are chosen because of the specific biology of those taxa (e.g., Kassen and Bell 2000; Begin and Roff 2001; Ferguson and Fairbairn 2001). Once again, misguided quests for ahistorical generalizations are a waste of time at best in evolutionary biology research.

This leads us to a reconsideration of the phylogenetic comparative method already discussed. Unfortunately, things are not easy for students of phenotypic integration even *if* a robust phylogenetic hypothesis is available for their group of interest. This is because of a simple reason pointed out by Westoby et al. (1995) and discussed in detail by Ackerly (2000). Most phylogenetic comparative analyses include some sort of "correction" (usually in the form of independent contrasts: Felsenstein 1985; Abouheif 1999) intended to "control" for historical events (which causes statistical nonindependence of the data points represented by the terminal taxa in a phylogeny), and uncover instead ecological patterns. The problem can readily be understood by glancing at Fig. 7.6: it is possible (indeed, likely) that phylogenetic niche conservatism (i.e., the tendency of daughter taxa to retain the ecology typical of their parent species) is present in the group being studied. This causes an overlap between the components of variance explainable by phylogeny and by ecology. Thinking about it in terms of standard analysis of variance, this would be a "phylogeny-by-ecology" interaction, analogous to the genotype-by-environment interaction in plasticity studies. If we factor out any variance related to the phylogenetic effect, this may throw away valuable information on the ecology and, in extreme cases, leave us with apparently no ecological signal when in fact there is one.

There is only one solution to this conundrum, and that is to gather as detailed an ecological database as one has on the phylogenetic relatedness

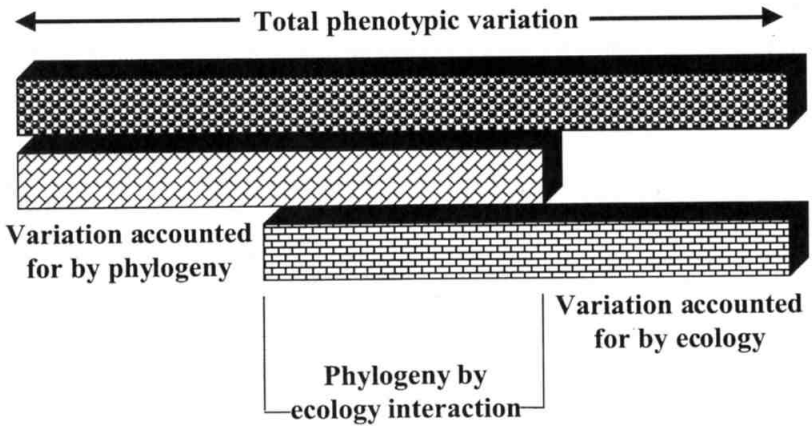

Figure 7.6 The total phenotypic variance observed in a group of phylogenetically related taxa (top) is usually thought of as the result of two components, one accounted for by historical commonality (middle), the second by the specifics of the ecology of each taxon (bottom). Standard "phylogenetic corrections" tend to eliminate the historical effect with the aim of uncovering the ecological specificities. However, if there is phylogenetic niche conservatism (i.e., daughter taxa maintain the ecological characteristics of their parent species), there may be a more or less large amount of phylogeny-by-ecology interaction that risks being thrown away by phylogenetic corrections. The only way around this problem is the difficult and tedious gathering of actual ecological information to complement the now relatively easily obtained phylogenetic one. (Based on a concept by Westoby et al. 1995).

and on the phenotypic variation. The problem identified by Westoby and colleagues stems from the fact that we normally use only the latter two types of data and infer ecological effects by difference. This clearly will not do because of the possibility of tantalizing interactions between causal agents such as selection and genetic constraints on one hand and drift and migration on the other. Unfortunately, it has been historically rather easy to gather copious (albeit not always informative) phenotypic data (Gould 1966), and recently it has been increasingly cheaper to obtain molecular data allowing us to reconstruct phylogenetic relationships. Conversely, to gather satisfactory ecological data is tedious and extremely time-consuming, and it is the sort of basic research that most funding agencies do not find sexy enough to support adequately. As a result, we may find that one of the major stumbling blocks to our understanding of phenotypic integration in the near future is the simple lack of solid field ecology for the taxa being studied.

Acknowledgments I wish to thank Ania Kolodynska for collaborating in some of the actual research summarized here and to the National Science Foundation (grant DEB-0089493) for partial funding. I should also like to thank Carl Schlichting and Gunter Wagner for many discussions on the concept of phenotypic integration and Katherine Preston for her insightful comments on this chapter.

Literature Cited

Abouheif, E. 1999. A method for testing the assumption of phylogenetic independence in comparative data. Evolutionary Ecology Research 1:895–909.

Ackerly, D. D. 2000. Taxon sampling, correlated evolution, and independent contrasts. Evolution 54:1480–1492.

Ackerly, D. D., and M. J. Donoghue. 1998. Leaf size, sapling allometry, and Corner's rules: phylogeny and correlated evolution in maples (*Acer*). American Naturalist 152:767–791.

Alonso-Blanco, C., and M. Koornneef. 2000. Naturally occurring variation in *Arabidopsis*: an underexploited resource for plant genetics. Trends in Plant Science 5:22–29.

Al-Shehbaz, I. A., S. L. O'Kane, Jr., and R. A. Price. 1999. Generic placement of species excluded from *Arabidopsis* (Brassicaceae). Novon 9:296–307.

Anderson, D. R., K. P. Burnham, and W. L. Thompson. 2000. Null hypothesis testing: problems, prevalence, and an alternative. Journal of Wildlife Management 64:912–923.

Anderson, M., and J. A. Roberts. 1998. *Arabidopsis*. Sheffield Academic Press, Sheffield, UK.

Antosiewicz, D. M., M. M. Purugganan, D. H. Polisenksy, and J. Braam. 1997. Cellular localization of *Arabidopsis* xyloglucan endotransglycosylase-related proteins during development and after wind stimulation. Plant Physiology 115:1319–1328.

Barraclough, T. G., and S. Nee. 2001. Phylogenetics and speciation. Trends in Ecology and Evolution 16:391–399.

Begin, M., and D. A. Roff. 2001. An analysis of G matrix variation in two closely related cricket species, *Gryllus firmus* and *G. pennsylvanicus*. Journal of Evolutionary Biology 14:1–13.

Biddington, N. L. 1986. The effects of mechanically-induced stress in plants—a review. Plant Growth Regulation 4:103–123.

Braam, J., and R. W. Davis. 1990. Rain-, wind-, and touch-induced expression of calmodulin and calmodulin-related genes in *Arabidopsis*. Cell 60:357–364.

Braam, J., M. L. Sistrunk, D. H. Polisensky, W. Xu, M. M. Purugganan, D. M. Antosiewicz, P. Campbell, and K. A. Johnson. 1997. Plant responses to environmental stress: regulation and functions of the *Arabidopsis TCH* genes. Planta 203:S35-S41.

Bradshaw, A. D. 1965. Evolutionary significance of phenotypic plasticity in plants. Advances in Genetics 13:115–155.

Chamberlain, T. C. 1897. The method of multiple working hypotheses. Science 15:92–96.

Chapin, F. S., III. 1991. Integrated responses of plants to stress. BioScience 41:29–36.

Cheverud, J. M. 1996. Quantitative genetic analysis of cranial morphology in the cotton-top (*Sanguinus oedipus*) and saddle-back (*S. fuscicolis*) tamarins. Journal of Evolutionary Biology 9:5–42.

Cheverud, J. M., G. P. Wagner, and M. M. Dow. 1989. Methods for the comparative analysis of variation patterns. Systematic Zoology 38:201–213.

Cipollini, D. F., Jr. 1997. Wind-induced mechanical stimulation increases pest resistance in common bean. Oecologia 111:84–90.

Cleland, C. E. 2001. Historical science, experimental science, and the scientific method. Geology 29:987–990.

Cordero, R. A. 1999. Ecophysiology of *Cecropia schreberiana* saplings in two wind regimes in an elfin cloud forest: growth, gas exchange, architecture and stem biomechanics. Tree Physiology 19:153–163.

Dixon, P., and T. O'Reilly. 1999. Scientific versus statistical inference. Canadian Journal of Experimental Psychology 53:133–149.

Dupré, J. 1993. The Disorder of Things: Metaphysical Foundations of the Disunity of Science. Harvard University Press, Cambridge, MA.

Emery, R. J. N., D. M. Reid, and C. C. Chinnappa. 1994. Phenotypic plasticity of stem elongation in two ecotypes of *Stellaria longipes*: the role of ethylene and response to wind. Plant, Cell and Environment 17:691–700.

Felsenstein, J. 1985. Phylogenies and the comparative method. American Naturalist 125:1–15.

Ferguson, I. M., and D. J. Fairbairn. 2001. Estimating genetic correlations from measurements of field-caught waterstriders. Evolution 55:2126–2130.

Gould, S. J. 1966. Allometry and size in ontogeny and phylogeny. Biological Review 41:587–640.

Gregson, R. A. M. 1997. Signs of obsolescence in psychological statistics: significance versus contemporary theory. Australian Journal of Psychology 49:59–63.

Griffing, B., and R. L. Scholl. 1991. Qualitative and quantitative genetic studies of *Arabidopsis thaliana*. Genetics 129:605–609.

Harvey, P. H., and M. D. Pagel. 1991. The Comparative Method in Evolutionary Biology. Oxford University Press, Oxford.

Howson, C., and P. Urbach. 1991. Scientific Reasoning: the Bayesian Approach. Open Court, La Salle, IL.

Huelsenbeck, J. P., B. Rannala, and J. P. Masly. 2000. Accomodating phylogenetic uncertainty in evolutionary studies. Science 288:2349–2350.

Ishitani, M., L. Xiong, B. Stevenson, and J.-K. Zhu. 1997. Genetic analysis of osmotic and cold stress signal transduction in *Arabidopsis*: interactions and convergence of abscisic acid-dependent and abscisic acid-independent pathways. Plant Cell 9:1935–1949.

Jaffe, M. J., and S. Forbes. 1993. Thigmomorphogenesis: the effect of mechanical perturbation on plants. Plant Growth Regulation 12:313–324.

Jefferys, W. H., and J. O. Berger. 1992. Sharpening Ockam's razor on a Bayesian strop. American Scientist 80:64–72.

Jernigan, R. W., D. C. Culver, and D. W. Fong. 1994. The dual role of selection and evolutionary history as reflected in genetic correlations. Evolution 48:587–596.

Juenger, T., M. Purugganan, and T. F. C. Mackay. 2000. Quantitative trait loci for floral morphology in *Arabidopsis thaliana*. Genetics 156:1379–1392.

Kassen, R., and G. Bell. 2000. The ecology and genetics of fitness in *Chlamydomonas*. X. The relationship between genetic correlation and genetic distance. Evolution 54:425–432.

Kellogg, E. A., and H. B. Shaffer. 1993. Model organisms in evolutionary studies. Systematic Biology 42:409–414.

Klingenberg, C. P., and S. D. Zaklan. 2000. Morphological integration between developmental compartments in the *Drosophila* wing. Evolution 54:1273–1285.

Klingenberg, C. P., A. V. Badyaev, S. M. Sowry, and N. J. Beckwith. 2001. Inferring developmental modularity from morphological integration: analysis of individual variation and asymmetry in bumblebee wings. American Naturalist 157:11–23.

Koch, M., J. Bishop, and T. Mitchell-Olds. 1999. Molecular systematics and evolution of *Arabidopsis* and *Arabis*. Plant Biology 1:529–537.

Koch, M., B. Haubold, and T. Mitchell-Olds. 2000. Comparative evolutionary analysis of chalcone synthase and alcohol dehydrogenase loci in *Arabidopsis*, *Arabis*, and related genera (Brassicaceae). Molecular Biology and Evolution 17:1483–1498.

Koch, M., B. Haubold, and T. Mitchell-Olds. 2001. Molecular systematics of the Brassicaceae: evidence from coding plastidic *matK* and nuclear *Chs* sequences. American Journal of Botany 88:534–544.

Lasceve, G., J. Leymarie, M. A. Olney, E. Liscum, J. M. Christie, A. Vavasseur, and W. R. Briggs. 1999. *Arabidopsis* contains at least four independent blue-light-activated signal transduction pathways. Plant Physiology 120:605–614.

Lee, H.-S., and Z. J. Chen. 2001. Protein-coding genes are epigenetically regulated in *Arabidopsis* polyploids. Proceedings of the National Academy of Sciences USA 98:6753–6758.

Loftus, G. R. 1993. A picture is worth a thousand *p* values: on the irrelevance of hypothesis testing in the microcomputer age. Behavior Research Methods, Instruments and Computers 25:250–256.

Losos, J. B., and D. B. Miles. 1994. Adaptation, constraint, and the comparative method: phylogenetic issues and methods. Pp. 60–98 in P. C. Wainwright and S. M. Reilly

(eds.), Ecological Morphology: Integrative Organismal Biology. University of Chicago Press, Chicago.
Lynch, M. 1999. Estimating genetic correlations in natural populations. Genetical Research 74:255–264.
Magwene, P. M. 2001. New tools for studying integration and modularity. Evolution 55:1734–1745.
Martins, E. P. 2000. Adaptation and the comparative method. Trends in Ecology and Evolution 15:296–299.
Mezey, J. G., J. M. Cheverud, and G. P. Wagner. 2000. Is the genotype-phenotype map modular? A statistical approach using mouse Quantitative Trait Loci data. Genetics 156:305–311.
Mitchell-Olds, T. 2001. *Arabidopsis thaliana* and its wild relatives: a model system for ecology and evolution. Trends in Ecology and Evolution 16:693–699.
Mizoguchi, T., K. Irie, T. Hirayama, N. Hayashida, K. Yamaguchi-Shinozaki, K. Matsumoto, and K. Shinozaki. 1996. A gene encoding a mitogen-activated protein kinase kinase kinase is induced simultaneously with genes for a mitogen-activated protein kinase and an S6 ribosomal protein kinase by touch, cold, and water stress in *Arabidopsis thaliana*. Proceedings of the National Academy of Sciences USA 93:765–769.
Mummenhoff, K., and H. Hurka. 1995. Allopolyploid origin of *Arabidopsis suecica* (Fries) Norrlin: evidence from chloroplast and nuclear genome markers. Botanica Acta 108:449–456.
Napp-Zinn, K. 1985. *Arabidopsis thaliana*. Pp. 492–503 in A. H. Halevy (ed.), CRC Handbook of Flowering. CRC, Boca Raton, FL.
O'Kane, S. L., Jr., and I. A. Al-Shehbaz. 1997. A synopsis of *Arabidopsis* (Brassicaceae). Novon 7:323–327.
Pidkowich, M. S., J. E. Klenz, and G. W. Haughn. 1999. The making of a flower: control of floral meristem identity in *Arabidopsis*. Trends in Plant Science 4:64–70.
Pigliucci, M. 1998. Ecological and evolutionary genetics of *Arabidopsis*. Trends in Plant Science 3:485–489.
Pigliucci, M. 2001. Phenotypic Plasticity: Beyond Nature and Nurture. Johns Hopkins University Press, Baltimore, MD.
Pigliucci, M. 2002a. Are ecology and evolutionary biology "soft" sciences? Annales Zoologici Finnici 39:87–98.
Pigliucci, M. 2002b. Touchy and bushy: phenotypic plasticity and integration in response to wind stimulation in *Arabidopsis thaliana*. International Journal of Plant Science 163:399–408.
Pigliucci, M. 2003. From molecules to phenotypes? The promise and limits of integrative biology. Basic and Applied Ecology 4:297–306.
Pigliucci, M., K. Cammell, and J. Schmitt. 1999. Evolution of phenotypic plasticity: a comparative approach in the phylogenetic neighborhood of *Arabidopsis thaliana*. Journal of Evolutionary Biology 12:779–791.
Pigliucci, M., H. Pollard, and M. Cruzan. 2003. Comparative studies of evolutionary responses to light environments in *Arabidopsis*. American Naturalist 161:68–82.
Platt, J. R. 1964. Strong inference. Science 146:347–353.
Price, R. A., J. D. Palmer, and I. A. Al-Shehbaz. 1994. Systematic relationships of *Arabidopsis*: a molecular and morphological perspective. Pp. 7–20 in E. M. Meyerowitz and C. R. Somerville (eds.), *Arabidopsis*. Cold Spring Harbor Laboratory Press, Cold Spring Harbor, NY.
Pyke, K. 1994. *Arabidopsis*—its use in the genetic and molecular analysis of plant morphogenesis. New Phytologist 128:19–37.
Ratcliffe, D. 1965. The geographical and ecological distribution of *Arabidopsis* and comments on physiological variation. Web address: www.arabidopsis.org.
Roff, D. A. 1996. The evolution of genetic correlations: an analysis of patterns. Evolution 50:1392–1403.

Roff, D. A., and T. A. Mousseau. 1999. Does natural selection alter genetic architecture? An evaluation of quantitative genetic variation among populations of *Allenomobius socius* and *A. fasciatus*. Journal of Evolutionary Biology 12:361–369.

Rudge, D. W. 1998. A Bayesian analysis of strategies in evolutionary biology. Perspectives on Science 6:341–360.

Schlichting, C. D. 1986. The evolution of phenotypic plasticity in plants. Annual Review of Ecology and Systematics 17:667–693.

Schlichting, C. D. 1989a. Phenotypic integration and environmental change. BioScience 39:460–464.

Schlichting, C. D. 1989b. Phenotypic plasticity in *Phlox*. II. Plasticity of character correlations. Oecologia 78:496–501.

Schlichting, C. D., and M. Pigliucci. 1998. Phenotypic Evolution: A Reaction Norm Perspective. Sinauer, Sunderland, MA.

Schluter, D. 2001. Ecology and the origin of species. Trends in Ecology and Evolution 16:372–380.

Shoemaker, J. S., I. S. Painter, and B. S. Weir. 1999. Bayesian statistics in genetics: a guide for the uninitiated. Trends in Genetics 15:354–358.

Sistrunk, M. L., D. M. Antosiewicz, M. M. Purugganan, and J. Braam. 1994. *Arabidopsis TCH3* encodes a novel Ca^{2+} binding protein and shows environmentally induced and tissue-specific regulation. Plant Cell 6:1553–1565.

Telewski, F. W., and M. L. Pruyn. 1998. Thigmomorphogenesis: a dose response to flexing in *Ulmus americana* seedlings. Tree Physiology 18:65–68.

van Emden, H. F., R. J. Macklin, and S. Staunton-Lambert. 1990. Stroking plants to reduce aphid populations. Entomologist 109:184–188.

Wade, P. R. 2000. Bayesian methods in conservation biology. Conservation Biology 14:1308–1316.

Wagner, G. 2001. The Character Concept in Evolutionary Biology. Academic Press, San Diego.

Wagner, G. P., and K. Schwenk. 2000. Evolutionarily stable configurations: functional integration and the evolution of phenotypic stability. Evolutionary Biology 31:155–217.

Wang, H., R. Datla, F. Georges, M. Loewen, and A. J. Cutler. 1995. Promoters from *kin1* and *cor6.6*, two homologous *Arabidopsis thaliana* genes: transcriptional regulation and gene expression induced by low temperature, ABA, osmoticum and dehydration. Plant Molecular Biology 28:605–617.

Westoby, M., M. R. Leishman, and J. M. Lord. 1995. On misinterpreting the "phylogenetic correction." Journal of Ecology 83:531–534.

Whitelam, G. C., and P. F. Devlin. 1998. Light signalling in *Arabidopsis*. Plant Physiology and Biochemistry 36:125–133.

Wisniewski, M. 1996. Recent advances in plant responses to stress. HortScience 31:30–57.

Wittgenstein, L. 1953/1973. Philosophical Investigations. Macmillan, New York.

Xu, W., M. M. Purugganan, D. H. Polisensky, D. M. Antosiewicz, S. C. Fry, and J. Braam. 1995. *Arabidopsis TCHA*, regulated by hormones and the environment, encodes a xyloglucan endotransglycosylase. Plant Cell 7:1555–1567.

Xu, W., P. Campbell, A. K. Vargheese, and J. Braam. 1996. The *Arabidopsis XET*-related gene family: environmental and hormonal regulation of expression. Plant Journal 9:879–889.

Yanovsky, M. J., J. J. Casal, and J. P. Luppi. 1997. The *VLF* loci, polymorphic between ecotypes Landsberg *erecta* and Columbia, dissect two branches of phytochrome A signal transduction that correspond to very-low-fluence and high-irradiance responses. Plant Journal 12:659–667.

Zhang, J., T. Van-Toai, L. Huynh, and J. Preiszner. 2000. Development of flooding-tolerant *Arabidopsis thaliana* by autoregulated cytokinin production. Molecular Breeding 6:135–144.