June 29, 2009

A Phylogenetic Mixture Model for the Evolution of Gene Expression

Research Article

Kevin H. Eng^{1,3}, Héctor Corrada Bravo², Sündüz Keleş^{1,3*}

Department of Statistics
 Department of Computer Sciences
 Department of Biostatistics and Medical Informatics

*Corresponding Author: Sündüz Keleş Department of Statistics 1300 University Ave University of Wisconsin-Madison Madison, WI 53706 608-262-2598 608-262-0032 (fax) keles@stat.wisc.edu

Running Title: Phylogenetic Mixture Model

Keywords: gene expression, likelihood ratio test, mixture model, natural selection, phylogenetic correction.

Abstract

Microarray platforms are used increasingly to make comparative inferences through genome-wide surveys of gene expression. While recent studies focus on describing the evidence for natural selection using estimates of the within and between taxa mutational variances, these methods do not explicitly or flexibly account for predicted nonindependence due to phylogenetic associations between measurements. In the interest of parsing the effects of selection, we introduce a mixture model for the comparative analysis of variation in gene expression across multiple taxa. This class of models isolates the phylogenetic signal from the non-phylogenetic and the heritable signal from the non-heritable while measuring the proper amount of correction. As a result, the mixture model resolves outstanding differences between existing models, relates different ways to estimate the across taxa variance and induces a likelihood ratio test for selection. We investigate by simulation and application the feasibility and utility of estimation of the required parameters and the power of the proposed test. We illustrate analysis under this mixture model with a gene duplication family data set.

Keywords: gene expression, likelihood ratio test, mixture model, natural selection, phylogenetic correction.

Introduction

The availability of gene expression data en masse admits a genomic resolution comparative expression experiment which measures many homologous gene transcripts simultaneously across many taxa in the interest of determining which genes are likely to undergo selective forces (Rifkin et al., 2003; Nuzhdin et al., 2004; Whitehead and Crawford, 2006). Through such an experiment the investigator may determine the relative strengths of neutral drift and natural selection forces on gene expression traits (Fay and Wittkopp, 2008) at the single gene level while isolating whole groups of genes which act together and which might have a common evolutionary history. These investigators propose the use of the variance within and between taxa to determine the strength and form of hypothesized selection forces. The expression of each gene is a single, continuously-valued trait, and, as in the usual comparative experiment, the analysis is potentially obfuscated by the evolutionary dependence common to the taxa. (Felsenstein, 1985; Harvey and Pagel, 1991; Martins and Garland, 1991; Purvis and Garland, 1993; Garland et al., 2005; Rohlf, 2006)

To account for this dependence, we may examine the structured form of the phylogenetic covariance matrix defined between taxa. The investigator typically considers the evolutionary relationship evidenced by a phylogenetic tree estimated from molecular sequence characters, but, for model based comparative analyses, we wish to translate these trees into covariance matrices. Under the assumption of a Brownian motion process underlying the evolution of the trait, we may construct a phylogenetic covariance from a known tree (Felsenstein, 1988). For general phylogenetic covariance matrices, Martins and Housworth (2002) suggested an eigenvector decomposition to identify variance with specific tree shapes. In Corrada Bravo et al. (2008), we developed a new algorithm for estimating a tree and its matching Brownian motion covariance directly from observed continuous-trait data. As opposed to methods like neighbor-joining (Saitou, 1987), this method globally optimizes a projection criterion over all possible tree topologies using proven, efficient methods for combinatorial optimization. For expression from gene duplication families, Gu (2004) and Oakley et al. (2005) both re-parameterize the mutational rates on each branch of a known tree covariance allowing it to better fit the phylogeny information. Of particular note, the addition of an error component allows these covariances to extend to a model for the entire experiment with a single covariance matrix (Ives et al., 2007; Felsenstein, 2008).

Practically, linear models model both dependence and error by implicitly assuming a covariance structure which decomposes the observed or experimental variance. Such decompositions are especially desirable since they correspond to known structures in the experiment. Lynch (1991) defines a mixed effects model across multiple traits, capturing the phylogenetic structure in a relationship matrix G and covariance between traits as a series of single parameter variance components. While adapting Lynch's model for biological replicate data, Christman et al. (1997) extend a memetic, due to Cheverud et al. (1985), where the trait value (T) is separated into a phylogenetic component (P), a specific value (S) and a random error component (E), namely T=P+S+E. This decomposition leads the authors to conclude that Lynch's model isolates heritable effects (P+S) from noise (E) but fails to separate them from one another (P from S). Housworth et al. (2004) reformulate Lynch's model to address this deficiency by emphasizing a parameter which indirectly estimates the degree of phylogenetic signal in the sample. Guo et al. (2007) fit three types of Bayesian flavored mixed effects models each parameterizing an increasing amount of phylogenetic signal, finding that modeling the degree of signal present yields better models.

The importance of determining the amount of phylogenetic signal in a sample cannot be understated. If there is a phylogenetic signal, the comparative analysis ought to find and remove the extra variation. If no signal can be detected, then corrective methods will overzealously bias the final estimates. Permutation tests at the level of tree estimation offer a way of testing for the presence of a signal or not (Blomberg et al., 2003). Pagel (1999) introduced λ as a measure of the strength of the signal and developed a likelihood ratio test for its presence. A continuous estimate carries more information than a dichotomous hypothesis test and should it indicate a strong signal, we ought to apply an appropriate phylogenetic correction.

Our goal in this paper is to integrate a framework for studying selection forces into phylogenetic, variance-decomposing models in a gene expression context. With respect to tests of selection, Rifkin et al. (2003) proposed the use of the estimated mean squares to model expected variation between and within taxa. In this framework, evidence of deviation from expectation under neutrality is evidence of the effect of natural selection. Nuzhdin et al. (2004) revised this idea using nested random effects in an ANOVA model and proposing the numerator and denominator of the standard, uncorrected F-ratio to be estimates of the between and within taxa variance. In particular, they give forms of the tests which distinguish between purifying and adaptive selection. Whitehead and Crawford (2006) continue the use of plain mean square estimates, adding a test for stabilizing selection.

In this article, we present a mixture model for the covariance in order to resolve predecessor models' inability to separate phylogenetic effects from non-phylogenetic ones clearly and to resolve the exclusion of clear consideration for the structured dependence in the testing frameworks of Nuzhdin et al. (2004) and Whitehead and Crawford (2006). In such a model, the necessary degree of correction is freely estimated so the investigator may draw inferences on parameters un-confounded by dependence. We discuss the convergence of existing models by demonstrating the relationships between their assumptions on the covariance; the mixture formulation covers a continuum of models set between independent contrasts and the class of phylogenetic mixed effects models. We describe the main assumptions and implications of the model from the practical analysis point of view, illustrating its effect with a simulation study and demonstrating its use in the study of gene family evolution in *Saccharomyces cerevisiae* (Oakley et al., 2005).

Methods

We are interested in modeling the variance of gene expression measurements, $Y_{gr} = (Y_{gr1}, \ldots, Y_{grT})$, which are made across the T taxa of a known phylogenetic tree for $g = 1, \ldots, G$ many genes with $r = 1, \ldots, n$ many microarray replicates. Through a Brownian motion process (Felsenstein, 1985), a rooted, bifurcating, phylogenetic tree has a well understood translation into the covariance matrix of a multivariate normal random vector.

Pagel (1999) and Freckleton et al. (2002) introduce the parameter λ as a measure of the strength of phylogenetic correlation, or the "loss of history," which induces a covariance matrix V(λ). In defining V(λ) to be a phylogenetic covariance matrix whose off-diagonals are multiplied by λ , the authors implicitly assume that opposing the phylogenetic structure V₀ is a specific, non-phylogenetic structure Λ_0 :

$$V(\lambda) = [\lambda J_T + (1 - \lambda)I_T] \circ V_0$$

= $\lambda V_0 + (1 - \lambda)\Lambda_0.$ (1)

Here, J_T is a $T \times T$ matrix of ones, I_T is the identity matrix of the same dimension and \circ is the element-wise (Hadamard) product. We define Λ_0 to be the diagonal matrix with the same main diagonal as V_0 and assume that $0 \le \lambda \le 1$.

Adding the experimental error to this variance we might model the variability of this set of measurements as

$$Var(Y_{gr}) = \tau^2 a \lambda V_0 + \tau^2 a (1 - \lambda) \Lambda_0 + \sigma^2 I_T$$
⁽²⁾

$$= \frac{\tau^2 \lambda}{\tau^2 + \sigma^2} a(\tau^2 + \sigma^2) \mathbf{V}_0 + \frac{\tau^2 (1 - \lambda)}{\tau^2 + \sigma^2} a(\tau^2 + \sigma^2) \Lambda_0 + \frac{\sigma^2}{\tau^2 + \sigma^2} (\tau^2 + \sigma^2) \mathbf{I}_T$$
(3)

$$= p_1(\kappa_1 \mathcal{V}_0) + p_2(\kappa_1 \Lambda_0) + p_3(\kappa_2 \mathcal{I}_T), \qquad (4)$$

where τ^2 is the rate of variation in the expression trait across taxa, σ^2 is the rate of variation within taxa, *a* is a nuisance scale factor accounting for the difference in units between sequence based trees (typically, the expected number of sequence substitutions) and the log ratio scale of the gene expression measurements. This point is discussed further in the following sections. The proportions p_1 , p_2 , p_3 are constrained to sum to one and κ_1 and κ_2 are measurements of total variation on the sequence scale and the expression scale.

This variance has two important interpretations. The relative rate interpretation (Equation 2) expresses the variance in terms of rates of mutation so that sequence based models and expression based models of evolution may be compared. The second, mixture model interpretation (Equation 4) is that p_1 , p_2 , and p_3 represent the proportion of the observed variance attributable to certain archetypical signals: the phylogenetic history, non-phylogenetic variation and within-taxa variation respectively. One should note that this is precisely the desired decomposition of the variance into T = P + S + E components from Christman et al. (1997).

Under this second interpretation, formalized in the following section, these parameters are estimable and evidence for selection forces can be evaluated. It may appear that a density with this mixture covariance is not identifiable because different sets of $(\mathbf{p}, \mathbf{\kappa})$ yield the same marginal covariance. In the mixture model, those densities are not the same; that is two identifiable mixtures may produce the same marginal covariance.

A Phylogenetic Mixture Model

We suppose that Y_{gr} follows a mixture distribution with variance Equation 4 and some mean vector $\boldsymbol{\mu}$. The mixture model probability density function of Y_{gr} is given by

$$Pr(Y_{gr}; \mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\kappa}, \mathbf{V}_0, \Lambda_0, \mathbf{I}_0) =$$

$$p_1 f(Y_{gr} \mid \boldsymbol{\mu}, \kappa_1 \mathbf{V}_0) + p_2 f(Y_{gr} \mid \boldsymbol{\mu}, \kappa_1 \Lambda_0) + p_3 f(Y_{gr} \mid \boldsymbol{\mu}, \kappa_2 \mathbf{I}_0).$$
(5)

where $f(y \mid \mu, \Sigma)$ is a normal density with mean μ and covariance Σ and $p_1 + p_2 + p_3 = 1$, $\kappa_1, \kappa_2 > 0$. This mixture model supposes that the observed variation has three sources, the correlated, phylogenetic signal; the uncorrelated, non-phylogenetic signal; and residual, experimental variance each represented by the three distributions.

These component distributions may be interpreted as particular archetypical scenarios. If the data show phylogenetic signal (a particular type of non-independence) then we believe that they come from the $f(Y_{gr} \mid \boldsymbol{\mu}, \kappa_1 V_0)$ component. If the data were independent but not identically distributed (each has is own specific variance) then $f(Y_{gr} \mid \boldsymbol{\mu}, \kappa_1 \Lambda_0)$ is the correct model. If the taxa were truly independent, identically distributed noise then $f(Y_{gr} \mid \boldsymbol{\mu}, \kappa_2 I_0)$ takes precedence. Mixing proportions p_1 , p_2 , p_3 represent the relative strengths or the probabilities of each component.

Since V_0 and Λ_0 are sequence based estimates they require a different scale (κ_1) than the expression log ratio based error term (κ_2) . Recall that V_0 is a tree structured covariance, Λ_0 is a diagonal matrix with the same diagonal entries (same specific variance but no covariance), and we typically assume that I_0 is the identity matrix of size T.

This model can be fit using the Expectation Maximization algorithm (Dempster et al., 1977) outlined in the appendix, where the strategy is to find the maximum likelihood estimates $(\mathbf{p}, \boldsymbol{\kappa})$ and to transform them into $(\tau^2, \sigma^2, \lambda, a)$. Some technical details (Everitt

and Hand, 1981) require that V_0 and Λ_0 are identifiable, that the off diagonal entries of V_0 are not too small or that V_0 is a reasonable tree estimate, a case readily checked by the investigator.

Testing Selection Hypotheses

We consider the evidence in favor of natural selection forces characterized by the variance between and within taxa (Rifkin et al., 2003) and the ratio between them, the F-ratio. Nuzhdin et al. (2004) identify genes with both variance estimates low as undergoing stabilizing selection; genes with low F-ratios may be undergoing balancing selection; and genes with large F-ratios may undergo adaptive divergence. Whitehead and Crawford (2006) add the constraint that genes undergoing adaptive divergence ought to favor a particular direction, i.e., correlate with an additional environmental covariate.

The variance estimates in these studies vary: the first article uses the mean squared error for the variance within taxa and the mutational variance scaled by time for the variance between taxa. The second uses the variance of a nesting factor (species) and the nested factor (line). The last uses the variance among the population means and the variance within populations. Using these ANOVA sums of squares implicitly assumes the following variances,

$$Var(Y_{gr}) = \sigma_{\text{error}}^2 I_T$$
 (Rifkin et al. 2003) (6)

$$= \sigma_{\text{species}}^2 \mathbf{I}_T + \sigma_{\text{line (species)}}^2 \mathbf{I}_T \qquad (\text{Nuzhdin et al. 2004}) \qquad (7)$$

$$= \sigma_{\text{pop.}}^2 I_T + \sigma_{\text{individual (pop.)}}^2 I_T \qquad \text{(Whitehead and Crawford 2006)} \qquad (8)$$

Compare these variances with Equation 2 which uses a phylogenetic tree covariance, $V(\lambda)$, in place of the uncorrelated I_T . While these studies do consider phylogenetic corrections at other points in their analysis, their ANOVA mean square estimates for the variances are uncorrected for possible phylogenetic dependence (they use diagonal I_T).

In our mixture parametrization, τ^2 is the gene specific between taxa variance (numerator of the F-ratio) and σ^2 is the gene specific within taxa variance (denominator of the F-ratio). Since σ^2 is interpreted as the rate of mutation in the expression trait, the relative sizes of τ^2 and σ^2 imply the following different evolutionary scenarios. When $\tau^2 = \sigma^2$, the signal is consistent with a Brownian motion process evolving along the given tree, representing the neutral drift null hypothesis. If $\tau^2 < \sigma^2$, there is less expression divergence than predicted by sequence divergence suggesting that the gene may be undergoing balancing selection. Inversely, $\tau^2 > \sigma^2$ favors directional selection since the observed divergence is larger than expectation. We relax the requirement that the residuals must also show correlation with environmental covariates, that is that they show a particular direction as well. If τ^2 and σ^2 are both "small" then we conclude that there is evidence of purifying or stabilizing selection. Because it is not clear what constitutes an unusually small variance, we do not consider testing stabilizing selection hypotheses at this time (see also the discussion at the end of the article).

For relative rate type models, a neutral model variance supposes that the divergence given by a sequence based tree directly predicts the divergence in the expression trait up to a mutation rate constant, that is it assumes that $\sigma^2 = \tau^2$ or that

$$Var(Y_{qr}) = 2\sigma^2 \left(p_1 a \mathcal{V}_0 + p_2 a \Lambda_0 + p_3 \mathcal{I}_T \right), \tag{9}$$

where $p_1 = \lambda/2$, $p_2 = (1 - \lambda)/2$, $p_3 = 1/2$. This variance may be used as the null model for a likelihood ratio test (LRT) which compares the log likelihoods of the model fit under the general mixture variance (Equation 4) and the model under the neutral variance (Hulsenbeck and Rannala, 1997). Since the within and between taxa variance estimates can take several possible values, we can test for evidence of each of these types of selection using a single omnibus test. If the test is not significant we cannot reject the neutral drift model, but if it is significant we must look at the estimates of τ^2 and σ^2 to determine the type of selection evidenced.

To conduct the test, compare the LRT statistic versus its asymptotic distribution, where $ll(\sigma^2, \tau^2, \lambda)$ is the log likelihood of the unrestricted model and $ll_0(\sigma^2, \lambda)$ is the log likelihood of the model when $\sigma^2 = \tau^2$:

$$LR = -2\left\{ ll(\hat{\sigma}^2, \hat{\tau}^2, \hat{\lambda}) - ll_0(\tilde{\sigma}^2, \tilde{\lambda}) \right\} \sim \chi_1^2, \tag{10}$$

for unrestricted maximum likelihood estimates $(\hat{\sigma}^2, \hat{\tau}^2, \hat{\lambda})$ and estimates $(\tilde{\sigma}^2, \tilde{\lambda})$ under the null model. An algorithm for computing both likelihoods is given in the appendix.

Results

Simulation: Need for corrections

While it is well accepted that phylogenetic corrections are necessary in comparative studies, we construct the following simulation study to illustrate the cost of failing to correct a phylogenetic signal on the statistical evolutionary hypotheses posited above. Suppose that V_0 is the following tree structured matrix with corresponding tree in Figure 1, that is the main diagonal entries are the specific variances for each taxa (total branch length) and the off diagonals are the covariances between taxa (shared branch lengths). For example, with reference to the branch lengths in Figure 1, $Var(Taxa A) = 1 + 2 + 1 + 1 = 5 = [V_0]_{11}$; $Var(Taxa B) = 1 + 2 + 1 + 3 = 7 = [V_0]_{22}; and Cov(Taxa A, Taxa B) = 1 + 2 + 1 = 4 = [V_0]_{12}.$

$$V_0 = \begin{pmatrix} 5 & 4 & 3 & 1 & 0 \\ 4 & 7 & 3 & 1 & 0 \\ 3 & 3 & 7 & 1 & 0 \\ 1 & 1 & 1 & 5 & 0 \\ 0 & 0 & 0 & 0 & 8 \end{pmatrix}$$

[Figure 1 about here.]

Under the mixture model proposed above, we define the selection hypotheses in the table below. We construct an artificial array of 350 genes (50 genes undergo each type of selection) and an artificial experiment where each gene is measured 500 times (5 taxa in V₀ above, 100 individuals). λ is generated by sampling a Uniform(0,1) random variable once for each of the 350 genes. Each gene has replicates drawn from one of three sources with probabilities $p_1 = \frac{\tau^2 \lambda}{\tau^2 + \sigma^2}, p_2 = \frac{\tau^2(1-\lambda)}{\tau^2 + \sigma^2}, p_3 = \frac{\sigma^2}{\tau^2 + \sigma^2}$. The details of the data generation are given below.

$$X_{gr} \sim \text{Multinomial}\left(p_1, p_2, p_3\right) \tag{11}$$

$$Y_{gr} \mid \{X_{gr} = 1\} \sim \mathcal{N}\left(0, (\tau^2 + \sigma^2) \mathbf{V}_0\right)$$
 (12)

$$Y_{gr} \mid \{X_{gr} = 2\} \sim \mathcal{N}\left(0, (\tau^2 + \sigma^2)\Lambda_0\right) \tag{13}$$

$$Y_{gr} \mid \{X_{gr} = 3\} \sim \mathcal{N}\left(0, (\tau^2 + \sigma^2)I_0\right)$$
 (14)

[Table 1 about here.]

Note that we choose a large number of individuals (arrays) to illustrate this problem clearly; if the problem exists for a large number of individuals then it ought to exist for a small number of individuals. Also, all the methods presented operate gene by gene so that these conclusions scale to any sized experiment.

We are most interested in considering the adequacy of methods based on ANOVA sums of squares. For each gene g, the mean square approach generates estimates from the usual one-way ANOVA table and considers significant effects using an F-ratio test. Since there are T = 5 rows in V₀, and we select n = 100 replicates, the proper reference for this test is the F-distribution with 4 and 495 degrees of freedom.

The resulting simulated data is analyzed in Figure 2 which plots the logged values of τ^2 and σ^2 under the scenarios tabled above. The seven versions of the variance based hypotheses are color coded: the black points represent a neutral drift null scenario, the two shades of blue are genes undergoing strong and weak directional selection; two shades of red, balancing selection and two shades of green, stabilizing selection. Two grey lines indicate the level 0.05 two-sided thresholds for the F-ratio test. Points above the upper threshold show evidence of directional selection. Points below the lower threshold show evidence of balancing selection. We do not implement the corresponding stabilizing selection tests.

[Figure 2 about here.]

The top two panels illustrate the same data when V₀ captures the true correlation between taxa. The true values of τ^2 and σ^2 are the same for every gene in the same group, so the spread of points represents sampling variability (and to some extent the effect of λ). The plot on the left (2a) shows the ANOVA estimates and on the right (2b) shows the estimates from the mixture model. Intuitively, both variance estimation procedures partition the total observed variance into within and between taxa parts. Since we assume the mixture model is the true generating model, we can see that the ANOVA estimates tend to over estimate σ^2 and make up for the excess by increasing the variance in the estimate of τ^2 . In a joint bias-variance tradeoff, the ANOVA estimate trades low variance in the σ^2 estimate for bias and higher overall error in the τ^2 estimate.

The left hand plots (Panels 2a and 2c) employ the mean square estimates for the between and within taxa variances. In Panel 2a, since all the groups of genes in each class of hypotheses are centered about the identity line, it is clear that choosing genes using their F-ratio is not specific for the directional alternative or the balancing alternative; some genes from each group fall above or below the corresponding threshold. This pattern holds even in the lower left plot (2c) where we assume the data really are independent and identically distributed.

In addition to problems with the hypothesis tests, the estimates of σ^2 and τ^2 appear over-estimated when we do not account for the dependence structure. In Figure 2a and 2c, the neutral drift cluster and the stabilizing selection clusters (which are and should be centered on the identity line), appear biased much farther up the identity line than they should.

Contrast these observations with the estimates from with the mixture model (Panel 2b). The plot shows what we would ideally like to see: all the genes clearly separate based on the true values of their parameters. The effects are clearly separated implying that there are a sufficient number of replicates to identify all the effects. Note that this scenario represents artificially ideal conditions: a large number of observations, good separation, each gene class has the same true parameter. The point is that the mean square estimates do not behave as expected under this optimal setting and we would not expect them to do so under more realistic experimental conditions. In practice, we might expect each gene to have a different set of parameters (τ^2 and σ^2) and the groups to overlap significantly. Furthermore, the proportion of genes undergoing natural selection may alter the plot significantly, the plot will depend on the proportion of genes under each type of selection and their relative strengths.

Simulation: Calibration

As we discussed in the methods section, the mixture model relies on a V₀ matrix that captures the phylogenetic relationship between the taxa. Additionally, tests of selection hypotheses require a model that preserves the relationship between τ^2 and σ^2 . Since estimates of the phylogenetic tree are typically obtained from sequence information (or similar independent sources), there is no reason to believe that it is of the appropriate scale for expression level data. If the given covariance structure is scaled too small then estimates of τ^2 will be artificially large; likewise if the given covariance is too large, τ^2 will be too small. The mixture model variance includes the parameter *a* to account for this scale (Equation 2).

We need to emphasize that the mixture assumption, made in the methods section, is necessary to obtain an identifiable estimate of a. Had we assumed a marginal mixed effects model (Lynch, 1991; Martins and Hansen, 1997; Guo et al., 2007) with the same variance (Equation 2), the scale parameter and the variance would only be estimable as $a\tau^2$. Practically, the investigator would have to assume some value of a in order to conduct selection tests, but this would create an uncorrectable bias in the testing framework.

Figure 3 summarizes the scaling problem and this correction using the same set of seven hypotheses from the simulation section. For these plots, we assume that we know V_0 , the same as in the last section, but that data are generated from scaled versions of V_0 . In the left panels (3a and 3d), the true $V_{1/100} = V_0/100$ is 100 times smaller than the given V_0 . In the middle (3b and 3e), the scale is correct. In the right panels (3c and 3f), the true $V_{100} = 100V_0$ is 100 times larger than the given V_0 .

The top row of Figure 3 demonstrates the effect of the wrong sized covariance by fitting the mixture with V_0 given. Notice that estimates are drawn uniformly downwards in Panel 3a but pushed upwards in panel 3c. For reference panel 3b is the same plot from Figure 2b. Fewer points appear in the latter plot since estimates may be unobtainable when this scaling is too far off. The bottom row of Figure 3 shows the effect of estimating nuisance scale a for large and small true values.

[Figure 3 about here.]

It makes sense that the procedure fails for a very small, since this case corresponds to the scenario where the heritable component is weak, i.e., there is very little signal. At present, this case can be identified by observing an unusually large proportion of genes for which $\hat{\lambda} = 0$ since very small a forces λ to shrink even if the signal is present.

Simulation: Tests

Based on the data presented in the next section, we choose more realistic strong and weak presentations of balancing and directional selection forces for studying the likelihood ratio test. As we mentioned in the methods section, each hypothesis depends on the ratio τ^2/σ^2 assuming that $\tau^2 = \sigma^2$ represents the null hypothesis. We set n = 5,10 for small sample microarray studies and n = 15, 30, 50 to check that the asymptotic distribution is correctly chosen. For each hypothesis, we draw a simulation data set by selecting three nuisance parameters (σ^2, a, λ),

$$\sigma^2 \sim \chi_1^2,\tag{15}$$

$$a = a'|a' > 1, \quad a' \sim \chi_1^2$$
 (16)

$$\lambda \sim \text{Uniform}(0, 1), \tag{17}$$

and computing the LRT. We repeat this procedure 10,000 times, tabulating the proportion of tests with p-values less than 0.05 in Table 2. Repeating the procedure 50 times allows us to compute the simulation error (in all cases < 0.006). Note that we consider a > 1 since the calibration simulation indicates that when a < 1 the available signal is hard to estimate.

[Table 2 about here.]

While the likelihood ratio test has reasonable power only at moderate sample sizes (n = 15), it should be noted that it is an omnibus test in the sense that it tests for any deviation from $\sigma^2 = \tau^2$. It is likely that a more powerful test can be built when it is of interest to test for either balancing or directional selection. One could also conduct these tests separately (as in Whitehead and Crawford (2006)), but that procedure would raise concerns about multiple testing.

Housworth et al. (2004) observed that the small number of replicates available in comparative experiments may not reach the statistical power necessary to make strong inferences. For gene expression data, we observe elsewhere (Eng et al., 2008) that this sort of per gene analysis may also suffer from low power, but propose that clustering together genes with similar a covariance structure may generate additional power. That is, if we believe that several genes evolved in concert and are willing to draw selection inferences on a whole group (i.e., a each member of a group undergoes the same selection force versus a single gene under a unique force) then we may employ genes as identical replicates in order to increase the power of the test. For example, if we have n = 5 replicates, then the powers tabulated are quite poor. If we are able to suppose that g = 2 genes (or experiments in the example in the next section) have the same covariance structure, then there are effectively $n = 5 \times 2 = 10$ available replicates so the power nearly triples for the balancing and stabilizing selection hypotheses.

Gene Family Data Example

The application of phylogenetic techniques to gene duplication families in the yeast *Saccharomyces cerevisiae* supposes that the members of these families have sequences that are linked to a common ancestor sequence, the target of a duplication event, and that the expression level of these descendent sequences is itself a trait subject to evolutionary forces (Thornton and DeSalle, 2000; Gu, 2004; Oakley et al., 2005). In such an analysis, the members of these families constitute the taxa of interest and the models developed in Gu (2004) and Oakley et al. (2005) test how well a sequence derived covariance matrix matches the predicted history of the expression trait. Since there is good reason to expect a phylogenetic structure between the genes, we will re-analyze the data set presented in Oakley et al. (2005) to illustrate the mixture model.

Using Gu's procedure (2004) for searching the proteome to identify 10 large gene families (between 7 and 18 genes each), Oakley et al. (2005) process expression arrays from 19 experiments from the Stanford Microarray Database (http://genome-www5.stanford.edu) and compute maximum likelihood phylogenetic trees for each family. Each experiment represents a different experimental condition, so we may draw inferences about the evidence of selection under particular conditions. There are 19 experiments each of which contains some of the 10 gene families for a total of 169 family specific measurements. Each experiment is a separate dataset where g corresponds to a gene family, t a single transcript in the gene family and r an array in the experiment. The data are analyzed are available from the supplementary materials from Oakley et al. (2005) (http://www.lifesci.ucsb.edu/eemb/ labs/oakley/pubs/MBE2005data/).

First consider the application of an ANOVA model which assumes that the residuals from its fit will be independent and identically distributed (iid). For each gene family in Table 3, the maximum residual correlation between all pairs of taxa over all replicates in all experiments demonstrates that the residuals are frequently not independent (8 out of 10 families have correlation greater than 0.50 in at least one pair of taxa) and Levene's test for the homogeneity of variances rejects the identically distributed assumption for 6 of the 10 families. These observations reinforce the need for an adjustment to account for the violation of the iid assumptions.

[Table 3 about here.]

The same 10 gene families appear in Table 4 which lists the number of experiments in which the gene family was measured, the number of these experiments which show some evidence of phylogenetic signal ($\lambda > .5$) and the number of experiments which may had significant selection likelihood ratio tests at level 0.05 split into balancing and directional selection. Tables of the experiments with significant tests and their p-values are available in the supplementary materials.

[Table 4 about here.]

The plot in Figure 4 shows the ANOVA estimates and the mixture model corrected estimates plotted on log scale (τ^2 is between taxa and σ^2 is within taxa). We have enlarged and colored points with a significant LRT at level 0.05 to illustrate their positions on these plots. As in the simulation plots, points about the identity line favor neutral expectations and points significantly distant from the line favor selection hypotheses. Not every extreme point is consistently highlighted because each test has a different sample size (number of genes in the family).

The ANOVA estimates appear to have a strong trend where τ^2 is smaller than expected, reflecting the tendency of the ANOVA estimate to favor σ^2 at the cost of shrinking τ^2 to zero if necessary (we saw this same pattern in Figure 2). The mixture estimates are more in line with neutral expectations.

The third panel (Figure 4c) subtracts the estimated common variance estimate under the neutral model; this is the scale that the LRT considers for significant deviations from the common variance. We see that most of the points with non-significant tests appear near the origin (the unrestricted maximum likelihood estimates are close to the neutral variance estimate) and points further from the origin have significant test statistics.

[Figure 4 about here.]

Concordant with the finding in Oakley et al. (2005) that most families have a "nonphylogenetic" model in different experimental conditions (117 of 152), a large proportion of experiments corrected with the mixture model show weak phylogenetic signal, $\lambda < .5$ (115 of 169). This raises some questions about how to interpret the results since Whitehead and Crawford (2006) only defined selection scenarios for τ^2 and σ^2 supposing that $\lambda = 1$. We do find, however that the Hexose Transport gene family appears to show strong phylogenetic signal in 8 of 14 experiments versus 12 of 14 in Oakley et al.'s analysis (2005). This family is also strongly represented in the balancing selection list (5 of 7 significant experiments).

As an illustration for combining experiments, the family of heat shock proteins has ten genes and the data sets examined in Oakley et al. (2005) contain two separate heat shock related experiments with n = 2 and n = 7 arrays each. When analyzed separately, neither experiment has a significant LRT. But, if we suppose that the experiments may be treated as replicates, there are n = 9 arrays available and the LRT is significant at level 0.05 (p=0.037) with ratio 0.286 indicating balancing selection.

Discussion

We have investigated the estimation of the variance of expression traits within and between the tips of a given phylogenetic tree, demonstrating a problem with current analyses and proposing a solution. The model we developed anticipates the use of microarray platforms to make general inferences about the strength of evidence for natural selection forces in expression traits. As described previously, mutational variability/mean square type selection inference relies on estimating τ^2 and σ^2 . We have paid less attention to the frameworks put forward in Gu (2004) and Oakley et al. (2005) which define natural selection on the basis of a likely history parameterized in the form of V₀. This model is also different from the Bayesian type Guo et al. (2007) model which allows one of some, all, or no phylogenetic signal; we emphasize the transformative role of λ , τ^2 and σ^2 in the sense that they also find a tree-structured covariance matrix consistent with observed data. Further, likelihood ratio based hypothesis testing is straightforward with these parameters.

Since this model only considers likelihood based decompositions of the variance, we can augment it with the application of standard statistical linear model theory to accommodate much more complicated experiments. In time course expression experiments, this form of linear model may account for the correlation over time (Guo et al., 2007) and also approximate gene associations by clustering genes with similar mean and variance effects together (Eng et al., 2008). For comparison, Oakley et al. (2005) corrected for correlated adjacent time points by using the first order differences, while Gu (2004) found the effect negligible. It is not unbelievable that more complex factors like expression under various conditions/treatments across taxa will be of interest and the model we have presented may serve as a useful component in that analysis.

Supplementary Material

Supplementary file 1 including tables of significant gene families from the example analysis are available at *Molecular Biology and Evolution* online (http://mbe.oxfordjournals.org/).

Acknowledgements

We thank Carol Lee (UW-Madison) and Anthony Ives (UW-Madison) for their feedback on earlier drafts of this article and Andrew Whitehead (Louisiana State University) for discussion on his experiments. This work was supported by NSF grant DMS 0804597 (SK); a PhRMA Foundation Research Starter Grant (SK); NIH grant HG03747 (SK).

References

- Blomberg, S. P., T. G. Garland, and A. R. Ives. 2003. Testing for Phylogenetic Signal in Comparative Data: Behavioral Traits are More Labile. Evolution 57:717–745.
- Cheverud, J. M., M. M. Dow, and W. Leutenegger. 1985. The Quantitative assessment of phylogenetic constraints in comparative analyses: sexual dimorphism in body weight among primates. Evolution **39**:1335–1351.
- Christman, M. C., R. W. Jernigan, and D. Culver. 1997. A comparison of two models for estimating phylogenetic effect on trait variation. Evolution **51**:262–266.
- Corrada Bravo, H., K. H. Eng, S. Keleş, G. Wahba, and S. Wright. 2008. Estimating tree-structured covariance matrices via mixed integer programming with an application to phylogenetic analysis of gene expression. Technical Report 1142, University Of Wisconsin-Madison, Department of Statistics.
- Dempster, A. P., N. M. Laird, and D. B. Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. Journal of the Royal Statistical Society 39:185–197.
- Eng, K. H., S. Keleş, and G. Wahba. 2008. A Linear Mixed Effects Clustering Model for Multi-Species Time Course Gene Expression Data. Technical Report 1143, University Of Wisconsin-Madison, Department of Statistics.
- Everitt, B. S., and D. J. Hand. 1981. Finite mixture distributions. Chapman and Hall.
- Fay, J. C., and P. J. Wittkopp. 2008. Evaluating the role of natural selection in the evolution of gene regulation. Heredity 100:191–199.
- Felsenstein, J. 1985. Phylogenies and the comparative method. The American Naturalist 125:1–15.

- ———. 1988. Phylogenies and the Quantitative Characters. Annual Review of Ecological Systems **19**:447–471.
- ———. 2008. Comparative Methods with Sampling Error and Within-Species Variation: Contrasts Revisited and Revised. The American Naturalist **171**:713–725.
- Freckleton, R. P., P. H. Harvey, and M. Pagel. 2002. Phylogenetic analysis and comparative data: a test and review of evidence. The American Naturalist 160:712–726.
- Garland, T., Jr., A. F. Bennett, and E. L. Rezende. 2005. Phylogenetic approaches in comparative physiology. Journal of Experimental Biology 208:3015–3035.
- Gu, X. 2004. Statistical framework for phylogenomic analysis of gene family expression profiles. Genetics 167:531–542.
- Guo, H., R. E. Weiss, X. Gu, and M. Suchard. 2007. Time squared: repeated measures on phylogenies. Molecular Biology and Evolution 24:352–362.
- Harvey, P. H., and M. D. Pagel. 1991. The Comparative Method in Evolutionary Biology. Oxford University Press.
- Housworth, E. A., E. P. Martins, and M. Lynch. 2004. The Phylogenetic Mixed Model. The American Naturalist 163:84–96.
- Hulsenbeck, J. P., and B. Rannala. 1997. Phylogenetic methods come of age: testing hypotheses in an evolutionary context. Science **276**:227–232.
- Ives, A. R., P. E. Midford, and T. Garland, Jr. 2007. Within-Species Variation and Measurment Error in Phylogenetic Comparative Methods. Systematic Biology 56:252–270.
- Lynch, M. 1991. Methods for the analysis of comparative data in evolutionary biology. Evolution **45**:1065–1080.

- Martins, E. P., and T. Garland, Jr. 1991. Phylogenetic Analysies of the Correlated Evolution of Continous Characters: A Simulation Study. Evolution 45:534–557.
- Martins, E. P., and T. F. Hansen. 1997. Phylogenies and the comparative method: a general approach to incorporating phylogenetic information into the analysis of interspecific data. The American Naturalist **149**:646–667.
- Martins, E. P., and E. A. Housworth. 2002. Phylogeny shape and the phylogenetic comparative method. Systematic Biology **51**:873–880.
- Nuzhdin, S. V., M. L. Wayne, K. Harmon, and L. M. McIntyre. 2004. Common pattern of evolution of gene expression level and protein sequence in *Drosophila*. Molecular Biology and Evolution 21:1308–1317.
- Oakley, T. H., Z. Gu, E. Abouheif, N. H. Patel, and W. H. Li. 2005. Comparative methods for the analysis of gene-expression evolution: an example using yeast functional genomic data. Molecular Biology and Evolution **22**:40-50. http://www.lifesci.ucsb.edu/eemb/labs/oakley/pubs/MBE2005data/.
- Pagel, M. 1999. Inferring the historical patterns of biological evolution. Nature 401:877–884.
- Purvis, A., and T. Garland, Jr. 1993. Polytomies in Comparative Analyses of Continuous Characters. Systematic Biology 42:596–575.
- Rifkin, S. A., J. Kim, and K. P. White. 2003. Evolution of gene expression in the Drosophila melanogaster subgroup. Nature Genetics 33:138–144.
- Rohlf, F. J. 2006. A Comment on Phylogenetic Correction. Evolution 60:1509–1515.
- Saitou, N. 1987. The neighbor-joining method: a new method for reconstructing phylogenetic trees. Molecular Biology and Evolution 4:406–425.

- Thornton, J. W., and R. DeSalle. 2000. Gene Family Evolution and Homology: Genomics Meets Phylogenetics. Annual Review of Genomics and Human Genetics 1:41–73.
- Whitehead, A., and D. L. Crawford. 2006. Neutral and adaptive variation in gene expression. Proceedings of the National Academy of Sciences 103:5425–5430.

Appendix

EM Algorithm for \mathbf{V}_0 and Λ_0 known up to scale *a* for a single gene. For all replicates $r = 1, \ldots, R$ of a single gene, let *C* be the class variable taking values 1,2 and 3 for components $\mathcal{N}(\mu, \kappa_1 \mathbf{V}_0)$, $\mathcal{N}(\mu, \kappa_1 \Lambda_0)$, and $\mathcal{N}(\mu, \kappa_2 \mathbf{I}_T)$ respectively. Suppose *T* is the rank of \mathbf{V}_0 (the number of taxa). The algorithm stops when the observed data log-likelihood (Equation 5) increases by less than 0.001. Since the estimate μ does not depend on *C*, its estimate is $\hat{\mu} = \frac{1}{R} \sum_{r=1}^{R} (Z'Z)^{-1} Z' Y_r$ for design matrix *Z* (e.g., for balanced replicates these are the group means).

1. E-Step

$$\hat{C}_{ri} = E(C_r = i | Y_r, \hat{\mu}, \hat{\kappa}_1^{(t)}, \hat{\kappa}_2^{(t)}, \hat{p}_i^{(t)}),$$

$$= \frac{P\left(Y_r - \hat{\mu} | C_r = i, \hat{\kappa}_1^{(t)} \hat{\kappa}_2^{(t)}\right) \hat{p}_i^{(t)}}{\sum_{i'=1}^3 P\left(Y_r - \hat{\mu} | C_r = i', \hat{\kappa}_1^{(t)}, \hat{\kappa}_2^{(t)}\right) \hat{p}_{i'}^{(t)}}$$

2. M-step

$$\hat{p}_{i}^{(t+1)} = \frac{\sum_{r} \hat{C}_{ri}}{R},$$

$$\hat{\kappa}_{1}^{(t+1)} = \frac{\frac{1}{T} \sum_{r=1}^{R} \left[\hat{C}_{r1} (Y_{r} - \hat{\mu})' \mathcal{V}_{0}^{-1} (Y_{r} - \hat{\mu}) + \hat{C}_{r2} (Y_{r} - \hat{\mu})' \Lambda_{0}^{-1} (Y_{r} - \hat{\mu}) \right]}{\sum_{r=1}^{R} \hat{C}_{r1} + \hat{C}_{r2}},$$

$$\hat{\kappa}_{2}^{(t+1)} = \frac{\frac{1}{T} \sum_{r=1}^{R} \left[\hat{C}_{r3} (Y_{r} - \hat{\mu})' (Y_{r} - \hat{\mu}) \right]}{\sum_{r=1}^{R} \hat{C}_{r3}}.$$

The required estimates from this algorithm are converted back to the original parametrization:

$$\hat{\tau}^2 = \hat{\kappa_2}(1 - \hat{p}_3),$$
$$\hat{\sigma}^2 = \hat{\kappa_2}(\hat{p}_3),$$
$$\hat{\lambda} = \frac{\hat{p}_1}{1 - \hat{p}_3},$$
$$\hat{a} = \frac{\hat{\kappa}_1}{\hat{\kappa}_2}.$$

Null Model EM Algorithm for Likelihood Ratio Test. The previous EM estimates $(\tau^2, \sigma^2, \lambda, a)$ under the unrestricted, full model, to compute the LRT we compare the log-likelihood of that model with the log-likelihood of a restricted model (neutral drift). Supposing that $\tau^2 = \sigma^2$, the mixing proportions become $(p_1, p_2, p_3) = (\frac{\lambda}{2}, \frac{1-\lambda}{2}, \frac{1}{2})$. The component densities are components $\mathcal{N}(\mu, \kappa_1 V_0)$, $\mathcal{N}(\mu, \kappa_1 \Lambda_0)$, and $\mathcal{N}(\mu, \kappa_2 I_T)$ where $\tau^2 = \sigma^2 = \kappa_2/2$ and $a = \kappa_1/\kappa_2$.

1. E-Step

$$\hat{C}_{ri} = E(C_r = i | Y_r, \hat{\mu}, \hat{\kappa_1}^{(t)}, \hat{\kappa_2}^{(t)}, \hat{p}_i^{(t)}),$$

$$= \frac{P\left(Y_r - \hat{\mu} | C_r = i, \hat{\kappa_1}^{(t)} \hat{\kappa_2}^{(t)}\right) \hat{p}_i^{(t)}}{\sum_{i'=1}^3 P\left(Y_r - \hat{\mu} | C_r = i', \hat{\kappa_1}^{(t)}, \hat{\kappa_2}^{(t)}\right) \hat{p}_{i'}^{(t)}}.$$

2. M-step

$$\begin{split} \hat{p}_{1}^{(t+1)} &= \frac{\sum_{r} \hat{C}_{r1}}{2\sum_{r} \hat{C}_{r1} + \hat{C}_{r2}}, \\ \hat{p}_{2}^{(t+1)} &= \frac{\sum_{r} \hat{C}_{r2}}{2\sum_{r} \hat{C}_{r1} + \hat{C}_{r2}}, \\ \hat{\kappa}_{1}^{(t+1)} &= \frac{\frac{1}{T} \sum_{r=1}^{R} \left[\hat{C}_{r1} (Y_{r} - \hat{\mu})' \mathcal{V}_{0}^{-1} (Y_{r} - \hat{\mu}) + \hat{C}_{r2} (Y_{r} - \hat{\mu})' \Lambda_{0}^{-1} (Y_{r} - \hat{\mu}) \right]}{\sum_{r=1}^{R} \hat{C}_{r1} + \hat{C}_{r2}}, \\ \hat{\kappa}_{2}^{(t+1)} &= \frac{\frac{1}{T} \sum_{r=1}^{R} \left[\hat{C}_{r3} (Y_{r} - \hat{\mu})' (Y_{r} - \hat{\mu}) \right]}{\sum_{r=1}^{R} \hat{C}_{r3}}. \end{split}$$

Again, we obtain the required estimates by converting back to the original parametrization:

$$\hat{\sigma}^2 = \hat{\kappa}_2/2,$$
$$\hat{\lambda} = 2\hat{p}_1,$$
$$\hat{a} = \frac{\hat{\kappa}_1}{\hat{\kappa}_2}.$$



Figure 1: Example tree for simulation study.



Figure 2: Simulation Example. Simulation under ideal settings for selection hypotheses defined in the text with a tree-structured covariance (V_0) and a non-phylogenetic covariance (I_0) shows that the ANOVA estimators do not discriminate between the hypotheses. The grey lines identify tests of selection: the corresponding two-sided F-test thresholds at $\alpha = 0.05$ for $F_{4,95}$. Panel (a) shows the ANOVA estimates of data generated under V₀; Panel (b) shows the mixture model estimates under V₀ and Panel (c) illustrates that the ANOVA estimates are inflated even under independent identically distributed characters (the primary ANOVA assumption).



Figure 3: Calibration Problem/Solution. Estimation of σ^2 and τ^2 is sensitive to misspecifying the scale of the phylogenetic covariance matrix, a. When a is not accounted for (Panels a, b, c), estimates are shrunk for a small (Panel a). When a is big, estimates are too big (Panel c). The ideal pattern appears in the top center (Panel b). Simultaneously estimating a fixes the problem (Panels e, f) for all but the smallest case (Panel d).



Figure 4: ANOVA and mixture model estimates for data from Oakley et al. (2005). Uncorrected ANOVA estimates (a) show a marked trend towards small between taxa variances (τ^2) while corrected estimates (b) fit more neutral expectation. The effect of the common variance in the neutral model is removed in panel (c). The ANOVA estimates show the same low variance pattern in the σ^2 estimate as in Figure 2. Ten extreme points are omitted from panels (b,c) to make the scales comparable. Points with significant tests at level 0.05 are enlarged and colored red for balancing and blue for directional selection evidence.

Hypothesis	$ au^2$	σ^2	Number	Plot Color
			of Genes	
Neutral Drift	1.00	1.00	50	black
Balancing Selection, Weak	1.00	5.00	50	light red
Balancing Selection, Strong	1.00	10.00	50	dark red
Directional Selection, Weak	5.00	1.00	50	light blue
Directional Selection, Strong	10.00	1.00	50	dark blue
Stabilizing Selection, Weak	0.10	0.10	50	light green
Stabilizing Selection, Strong	0.05	0.05	50	dark green
			350 gen	nes total
Table 1: S	imulatio	on Desig	gn.	

<u> </u>

Table 1: Simulation Des	sign	l
-------------------------	------	---

Table 2: Simulated Power for Likelihood Ratio Test. The Neutral Drift null hypothesis row corresponds to the level of the test. The four selection hypotheses rows are estimated statistical power at level 0.05. All powers were estimated with 10,000 replicates and 50 simulation replicates (simulation errors < 0.006).

Selection	Estimated Power					
Hypothesis	τ^2/σ^2	n = 5	n = 10	n = 15	n = 30	n = 50
Neutral Drift (Null)	$^{1}/_{1}$	0.017	0.034	0.038	0.049	0.050
Balancing, Weak	$^{1}/_{2}$	0.006	0.079	0.178	0.374	0.573
Balancing, Strong	$^{1}/_{5}$	0.002	0.223	0.517	0.884	0.981
Directional, Weak	$^{2}/_{1}$	0.034	0.091	0.121	0.336	0.549
Directional, Strong	$\frac{5}{1}$	0.043	0.167	0.278	0.648	0.845

Gene Family	Max residual	Levene's Test	
	correlation	(p-value)	
ABC Transporters	0.40	0.0093	
ADP Ribosylation	0.56	0.0751	
Alpha Glucosidases	0.71	0.1139	
DUP	0.84	0.3998	
GTP Binding	0.51	< 0.0001	
HSP DnaK	0.78	< 0.0001	
Hexose Transport	0.93	< 0.0001	
Kinases	0.43	0.0001	
Permeases	0.60	< 0.0001	
Putative Helicases	0.75	0.0626	

Table 3: *Diagnostics for ANOVA residuals.* The ANOVA method for estimating the mutational variances assumes that the residuals will be independent and identically distributed. The maximum residual correlation between pairs of taxa over all replicates in all experiments demonstrates that the residuals are frequently not independent (8 out of 10 have correlation greater than 0.50) and Levene's test for the homogeneity of variances shows that the identically distributed assumption holds for only 4 of the 10 families.

Gene Family	Number of	$\lambda > .5$	Tests at level 0.05	
(Number of Taxa)	Experiments		Directional	Balancing
ABC Transporters (8)	17	4	1	0
ADP Ribosylation (7)	17	7	0	0
Alpha Glucosidases (6)	19	4	1	0
DUP (10)	13	8	0	0
GTP Binding (11)	17	7	0	1
HSP DnaK (10)	16	1	0	1
Hexose Transport (18)	14	8	0	7
Kinases (7)	16	8	2	1
Permeases (17)	12	5	0	1
Putative Helicases (11)	11	2	0	4

Table 4: Yeast Gene Family Data. Gene family data analyzed under the mixture model. A small number of experiments show strong phylogenetic signal ($\lambda > .5$), while the number of experiments with level 0.05 significant ratios τ^2/σ^2 large (directional) or small (balancing) are tabulated above.