Population Genetics Based Phylogenetics Under Stabilizing Selection for an Optimal Amino Acid Sequence: A Nested Modeling Approach

Jeremy M. Beaulieu,[1,2,3] Brian C. O'Meara,[2,3] Russell Zaretzki,[4] Cedric Landerer,[2,3] Juanjuan Chai,[†,3] and Michael A. Gilchrist*[,2,3]

[1]Department of Biological Sciences, University of Arkansas, Fayetteville, AR

[2]Department of Ecology & Evolutionary Biology, University of Tennessee, Knoxville, TN

[3]National Institute for Mathematical and Biological Synthesis, Knoxville, TN

[4]Department of Business Analytics & Statistics, Knoxville, TN

[†]Present address: Suite 1039, White Plains, NY

***Corresponding author:** E-mail: mikeg@utk.edu.

Associate editor: Tal Pupko

Abstract

We present a new phylogenetic approach, selection on amino acids and codons (SelAC), whose substitution rates are based on a nested model linking protein expression to population genetics. Unlike simpler codon models that assume a single substitution matrix for all sites, our model more realistically represents the evolution of protein-coding DNA under the assumption of consistent, stabilizing selection using a cost-benefit approach. This cost–benefit approach allows us to generate a set of 20 optimal amino acid-specific matrix families using just a handful of parameters and naturally links the strength of stabilizing selection to protein synthesis levels, which we can estimate. Using a yeast data set of 100 orthologs for 6 taxa, we find SelAC fits the data much better than popular models by 10^4–10^5 Akike information criterion units adjusted for small sample bias. Our results also indicated that nested, mechanistic models better predict observed data patterns highlighting the improvement in biological realism in amino acid sequence evolution that our model provides. Additional parameters estimated by SelAC indicate that a large amount of nonphylogenetic, but biologically meaningful, information can be inferred from existing data. For example, SelAC prediction of gene-specific protein synthesis rates correlates well with both empirical (r=0.33–0.48) and other theoretical predictions (r=0.45–0.64) for multiple yeast species. SelAC also provides estimates of the optimal amino acid at each site. Finally, because SelAC is a nested approach based on clearly stated biological assumptions, future modifications, such as including shifts in the optimal amino acid sequence within or across lineages, are possible.

Key words: **Wright–Fisher, stabilizing selection, allele substitution, protein function, gene expression.**

Article

Introduction

Phylogenetic analyses play a critical role in most aspects of biology, particularly in the fields of ecology, evolution, paleontology, medicine, and conservation. Although the scale and impact of phylogenetic studies have increased substantially over the past two decades, the realism of the mathematical models on which these analyses are based has changed relatively little by comparison. The most popular models of DNA substitution used in molecular phylogenetics are simple nucleotide models that date back to the early 1980s and 1990s, for example, F81, F84, HKY85, TN93, and GTR (see Yang 2014, for an overview), and are indifferent to the type of sequences they are fitted to. For example, when evaluating protein-coding sequences these models are inherently agnostic with regard to the different amino acid substitutions and their impact on gene function and, as a result, cannot describe the behavior of natural selection at the amino acid or protein level.

Two important and independent attempts to address this critical shortcoming were introduced by Goldman and Yang (1994, commonly abbreviated as GY) and Muse and Gaut (1994). These models were explicitly built for protein-coding data, assuming that differences in the physicochemical properties between amino acids, or physicochemical distances for short, could affect substitution rates. A number of researchers have used physicochemical-based models to make inferences about selection for such properties (e.g., see Hughes et al. 1990; Koshi and Goldstein 1997; Xia and Li 1998; Koshi et al. 1999; McClellan and McCracken 2001; Woolley et al. 2003; Blazej et al. 2017; and Yang 2014 for a brief summary); however, in terms of tree and/or branch length reconstruction, physicochemical-based codon models as originally introduced have rarely been used for empirical data. Instead, the often cited models of Goldman and Yang (1994) and Muse and Gaut (1994) have served as the basis for an array of simpler and, in turn, more popular ω models that, starting

with Yang and Nielsen (1998) and Nielsen and Yang (1998), typically assume an equal fixation probability for all nonsynonymous mutations. Although often attributed to GY, these later and simpler models were the first to employ the single term ω to model the differences in fixation probability between nonsynonymous and synonymous changes at all sites. Since their introduction, more complex models have been developed that allow ω to vary between sites or branches (as cited in Anisimova 2012) and include selection on different synonyms for the same amino acid (e.g., Yang and Nielsen 2008).

In Goldman and Yang (1994), Yang and Nielsen (1998), Nielsen and Yang (1998), and later studies based on their work, ω is suggested to indicate whether a given site within a protein sequence is under consistent "stabilizing" ($\omega < 1$) or "diversifying" ($\omega > 1$) selection. Contrary to popular belief, ω does not describe whether a site is evolving under a constant regime of stabilizing or diversifying selection, but instead how a very particular *selective environment* changes over time. Below we explain how the actual behavior of these models is inconsistent with how "stabilizing" and "diversifying" selection are otherwise defined and understood (e.g., see Pellmyr 2002).

For example, when $\omega < 1$ synonymous substitutions have a higher substitution rate than any possible nonsynonymous substitutions. As a result, the model behaves as if the resident amino acid *i* at a given site is favored by natural selection. Even when ω is allowed to vary between sites, symmetrical aspects of the model means that for any given site the strength of selection for the resident amino acid *i* over its 19 alternatives is equally strong regardless of their physicochemical properties. Paradoxically, natural selection for amino acid *i* persists *until* a substitution for another amino acid, *j*, occurs. As soon as amino acid *j* fixes, but not before, selection now favors amino acid *j* equally over all other amino acids, including amino acid *i*. This is now the opposite scenario from when *i* was the resident. Thus, one reasonable interpretation of ω is that it represents the rate at which the selective environment itself changes, and this change in selection at a site perfectly coincides with the fixation of a new amino acid. This change in the selective environment also differs from noninstantaneous, compensatory changes at other sites that results in reversals becoming less likely with time (Pollock et al. 2012; Shah et al. 2015).

Similarly, when $\omega > 1$, synonymous substitutions have a lower substitution rate than any possible nonsynonymous substitutions from the resident amino acid. Again due to the model's symmetrical nature, the selection *against* the resident amino acid *i* is equally strong relative to alternative amino acids. The selection against the resident amino acid *i* persists until a substitution occurs at which point selection now *favors* amino acid *i*, as well as the other amino acids, to the same degree *i* was previously disfavored. Of course, in practice it is unlikely for the nonsynonymous rate to be greater than the synonymous substitution rate in the absence of a particular process (e.g., antagonistic coevolution). Given these behaviors, ω-based models are likely to only reasonably approximate a narrow set of scenarios such as perfectly symmetrical over/underdominance or positive/negative frequency-dependent selection (Hughes and Nei 1988; Nowak 2006; Hughes 2007). Furthermore, ω-based models implicitly assume the substitution is on the same timescale as the shifts in the optimal (or pessimal) amino acid. Although ω is often viewed as a gene-wide metric, rather than site-specific one, it is not clear to what degree this averaging mitigates these issues.

New Approaches

To address these fundamental shortcomings in ω-based phylogenetic approaches, we present an approach where selection explicitly favors minimizing the cost–benefit function η of a protein whose relative performance is determined by the order and physicochemical properties of its amino acids. Our approach, which we call selection on amino acids and codons, or SelAC, is developed in the same vein as previous phylogenetic applications of the Wright–Fisher process (e.g., Muse and Gaut 1994; Koshi and Goldstein 1997; Halpern and Bruno 1998; Koshi et al. 1999; Dimmic et al. 2000; Lartillot and Philippe 2004; Rodrigue et al. 2005; Yang and Nielsen 2008; Thorne et al. 2012; Rodrigue and Lartillot 2014). Similar to Lartillot and Philippe (2004) and Rodrigue and Lartillot (2014), we assume there is a finite set of rate matrices describing the substitution process and that each position within a protein is assigned to a particular rate matrix category. Unlike that work, we assume a priori there are 20 different families of rate matrices, 1 family for when a given amino acid is favored at a site. The key parameters underlying these matrices are shared across genes with their effects scaled by each gene's expression level. By combining a cost–benefit approach with gene expression levels, SelAC produces a large set of substitution matrices using a very limited number of parameters.

Although natural selection on protein-coding regions can take many forms, one general approach to describing its effects is by relating a codon sequence to the "cost" of producing the encoded protein and the functional benefit (or potential harm) from translating its sequence. The gene-specific cost of protein synthesis can be affected by the amino acids used, the direct and indirect costs of peptide assembly by the ribosome, and the use of chaperones to aid in folding. Importantly, these costs can be computed to varying degrees of realism (e.g., Wagner 2005; Lynch and Marinov 2015). We have previously presented models of protein synthesis costs that, alternatively, take into account the cost of ribosome pausing (Shah and Gilchrist 2011) or premature termination errors (Gilchrist and Wagner 2006; Gilchrist 2007; Gilchrist et al. 2009).

Protein function or "benefit" can be affected by the amino acids at each site and their interactions. Linking amino acid sequence to protein function is a daunting task; thus for simplicity, we assume that for any given desired biological function to be carried out by a protein, that 1) the biological importance of this protein function is invariant across the tree; 2) there is a single optimal amino acid sequence that carries out this function best; and 3) the functionality of

alternative amino acid sequences declines with their physicochemical distance from the optimum on a site by site basis.

Beyond fitting the phylogenetic data better than currently available nucleotide and codon models according to model adequacy and Akaike information criterion corrected for small sample bias (AICc; Burnham and Anderson 2002, p. 66), SelAC also makes inferences about other important biological processes. By comparing these inferences to other empirical data, such as we do with protein synthesis data, we can evaluate SelAC's performance independent of the data it is fitted to. Indeed, SelAC's assumptions lead to mechanistic and, thus, testable hypothesis about the nature of and relationships between mutation, protein function, gene expression, and rates of evolution. More importantly, alternative hypotheses could be used in place of ours and, in turn, phylogenetic and other types of data could be used to evaluate the support of these alternative models. Our hope is that by moving away from the more phenomenological models we can better connect population genetics, molecular biology, and phylogenetics allowing each area to inform the others more effectively.

Results

The SelAC model requires the construction of gene and amino acid-specific substitution matrices that use a submodel nested within our substitution model. This requires only a handful of genome-wide parameters such as nucleotide-specific mutation rates μ_{ij}, which are scaled by effective population size N_e, amino acid side chain physicochemical weighting parameters α_c, α_p, and α_v, and a gamma distribution shape parameter α_G describing the distribution of site sensitivities G. In addition to these genome-wide parameters, the model requires a gene-specific functionality expression parameter that describes the average rate at which the protein's functionality is produced by the organism or a gene's "average functionality production rate" ψ. By linking transition rates $q_{i,j}$ to gene expression ψ, our approach allows use of the same model for genes under varying degrees of stabilizing selection. Specifically, we assume the strength of stabilizing selection for an optimal sequence $\vec{a}^*$ is proportional to ψ, which we can estimate for each gene.

We first evaluated the performance of our codon model by simulating data sets and estimating the bias of the inferred model parameters from these data. Overall, the simulation results indicated that our SelAC model can reasonably recover the known values of the generating model (fig. 1 and supplementary fig. S3, Supplementary Material online). This includes not only the parameters in SelAC, but also the optimal amino acids for a given sequence as well as the estimates of the branch lengths. There are, however, a few observations to note. First, the ability to accurately recover the true optimal amino acid sequence $\vec{a}^*$ will largely depend on the magnitude of the realized average protein synthesis rate of the gene ϕ, which is the target functionality rate ψ divided by the functionality of the observed amino acid sequence $\boldsymbol{B}(\vec{a})$. This is, of course, intuitive, given that ψ sets the strength of stabilizing selection towards an optimal amino acid at a site. However, the inclusion of between site variation in selection via the shape parameter α_G into SelAC generally improves the quality of our estimates of ψ and our ability to recover the optimal amino acids $\vec{a}^*$. This is true even for the gene with the lowest baseline ψ. Second, we found a strong downward bias in estimates of α_G, which actually translates to greater variation among the rate categories. The choice of a gamma distribution to represent site-specific variation in sensitivity was based on mathematical convenience and convention, rather than on biological reality. Furthermore, given the fact that the density of the gamma distribution is infinite at $G = 0$ when $\alpha_G < 1$, imputing site-specific G values will be an issue in these scenarios. Nevertheless, we suspect that this downward estimation bias of α_G is in large part due to the difficulty in determining the baseline ψ for a given gene and the value of α_G that globally satisfies the site-specific variation in sensitivity across all genes, as indicated by the slight upward bias in estimates of ψ (see supplementary fig. S5, Supplementary Material online).

In regard to model fit in an empirical setting, our results clearly indicated that linking the strength of stabilizing selection for the optimal sequence to gene expression substantially improves our model fit. Furthermore, including the shape parameter α_G for the random effects term $G \sim \text{Gamma}\left(\text{shape} = \alpha_G,\ \text{rate} = \alpha_G\right)$ to allow for heterogeneity in this selection between sites within a gene improves the ΔAICc of SelAC $+\,\Gamma$ over the simpler SelAC models by over 22,000 AICc units. Using either ΔAICc or AIC_w as our measure of model support, the SelAC models fit extraordinarily better than GTR $+\,\Gamma$, GY, or FMutSel (table 1). This is in spite of the need for estimating the optimal amino acid at each position in each protein, which accounts for 49,881 additional model parameters. Even when compared with the next most parameter rich codon model in our model set, FMutSel (with 178 parameters), SelAC $+\,\Gamma$ model shows over 160,000 AICc unit improvement over FMutSel. SelAC models also appeared to outperform, based on likelihood and reported AIC and AICc from each program, the 161 codon models in IQtree (Nguyen *et al.* 2015). See; supplementary table S1, Supplementary Material online for results.

We note our use of AICc, as calculated in Burnham and Anderson (2002, p. 66) and as opposed to the standard AIC, in the above model comparisons. At the outset of our study it was unclear what the appropriate sample size n is when comparing models of sequence evolution. Building upon the work of Jhwueng et al. (2014), our simulations suggest that using the number of taxa times the number of sites as the sample size correction performed best as a small sample size correction for estimating Kullback–Liebler (KL) distance in phylogenetic models (Supporting Materials). This also has an intuitive appeal. In models that have at least some parameters shared across sites and some parameters shared across taxa, increasing the number of sites and/or taxa should be adding more samples for the parameters to estimate. This is consistent considering how likelihood is calculated for phylogenetic models: the likelihood for a given site is the sum of the probabilities of each observed state at each tip, which is then multiplied across sites. It is arguable that the conventional

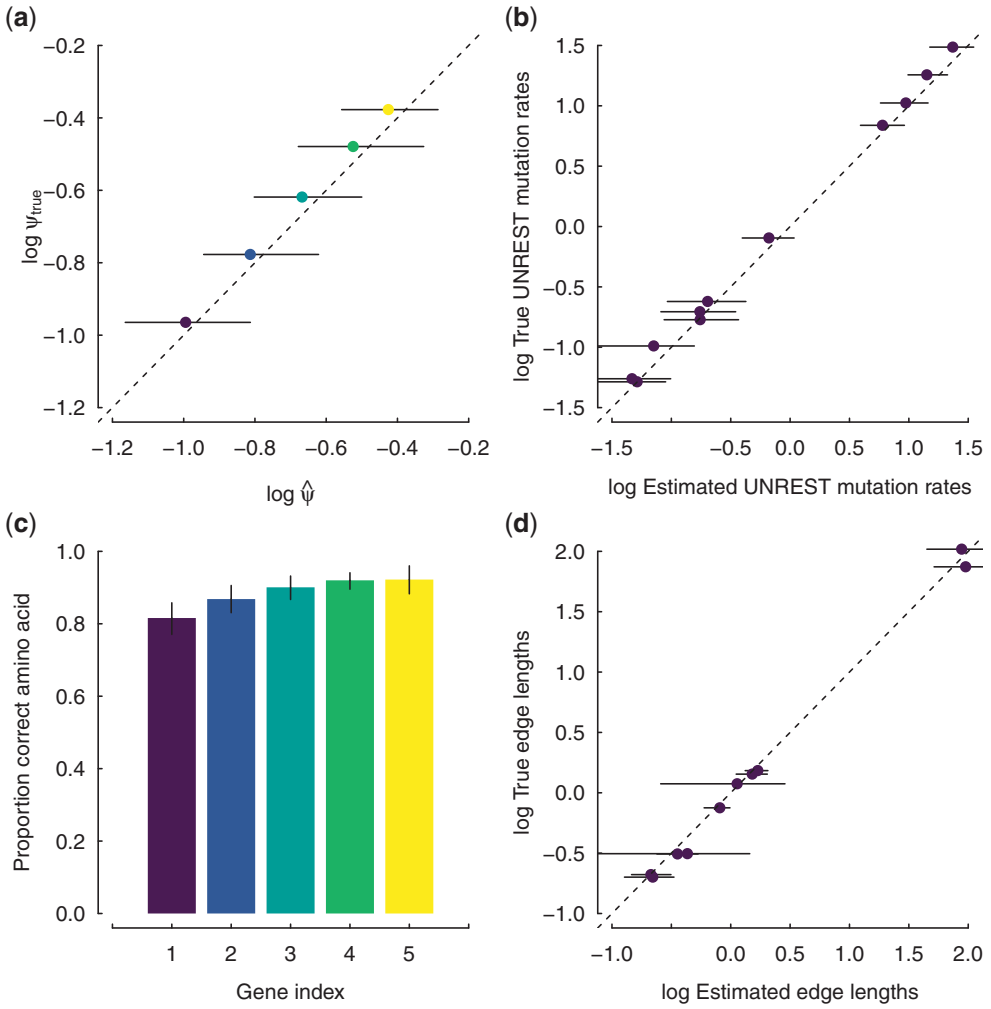

Fig. 1. Summary of a five-gene simulation for a SelAC model where we assume $\alpha_G = \infty$, and thus, no site-specific sensitivity in the generating model. The "known" parameters were based on fitting the SelAC model to the 106 gene data set and phylogeny of Rokas et al. (2003), with gene choice being based on five evenly spaced points along the rank order of the gene-specific composite parameter ψ'_g. The points and associated uncertainty in the estimates of the gene-specific average protein synthesis rate, or ψ (calculated from ψ') (*a*), nucleotide mutation rates under the UNREST model (*b*), proportion of correct optimal amino acids for a given gene (*c*), and estimates of the individual edge lengths are based the mean and 2.5% and 97.5% quantiles across all 50 simulated data sets (*d*). Gene index on the *x*-axis refers to the arbitrary number assigned to the simulated gene.

approach in comparative methods is calculating AICc in the same way. That is, if only one column of data (or "site") is examined, as remains remarkably common in comparative methods, when we refer to sample size, it is technically the number of taxa multiplied by number of sites, even though it is referred to simply as the number of taxa.

With respect to estimates of gene expression SelAC, actually has two related measures. The first ψ represents the average rate at which the gene's function is produced. The second ϕ represents the average rate at which a gene's protein is produced. These two parameters are closely linked and thus highly correlated because $\phi = \psi/\boldsymbol{B}$ where $\boldsymbol{B}$ is the relative functionality of each taxa's sequence. For simplicity, we will focus on ϕ as our measure of gene expression. SelAC's ϕ values were strongly correlated with both empirical measurements (Pearson's r=0.33–0.48) and theoretical predictions (Pearson's r=0.45–0.64) of gene expression (fig. 2 and supplementary figs. S1 and S2, Supplementary Material online, respectively). These correlations are remarkable given that they were uncovered using only codon sequences. The estimate of the α_G parameter, which controls the shape of the site-specific, gamma-distributed, variation in sensitivity of the protein's functionality, indicated a moderate level of variation in strength of selection among sites. Our estimate of α_G =1.36, produced a distribution of sensitivity terms G ranged from 0.342 to 7.32, but with more than 90% of the weight for a given site-likelihood being contributed by the 0.342 and 1.50 rate categories. In simulation, however, of all the parameters in the model, only α_G showed a consistent bias, in that the maximum likelihood estimate (MLE) were generally lower than their actual values (see Supporting Materials). Other parameters in the model, such as the Grantham weights, provide an indication as to the physicochemical distance between amino acids. Our estimates of

Table 1. Comparison of Maximum Likelihood Fits for SelAC and Commonly Used Models Based on Negative Log-Likelihood ($-\ln L$), AIC, AICc, and AIC_w from Analyses of 100 Selected Genes from Six Yeast Taxa (Salichos and Rokas 2013).

Model	$-\ln \mathcal{L}$	Parameters Estimated	AIC	AICc	$_{\Delta}$AICc	Model Weight
SelAC+Γ	453,620.8	50,005	1,007,252	1,027,314	0	>0.999
SelAC	464,114.8	50,004	1,028,238	1,048,299	20,985	<0.001
$SelAC_M$+Γ	465,106.9	50,005	1,030,224	1,050,286	22,972	<0.001
$SelAC_M$	478,302.4	50,004	1,056,613	1,076,674	49,360	<0.001
FMutSel	597,140.7	178	1,194,637	1,194,638	167,324	<0.001
GY	612,670.4	111	1,225,563	1,225,563	198,249	<0.001
GTR+Γ	655,166.4	610	1,311,553	1,311,554	284,240	<0.001

NOTE.—The subscripts *M* indicate model fits where the most common or "majority rule" amino acid was fixed as the optimal amino acid a^* for each site. As discussed in text, despite the fact that a^* for each site under *M* was not fitted by our algorithm, its value was determined by examining the data and, as a result, represent an additional parameter estimated from the data and are accounted for in our table. Sample size used in the calculation of AICc is assumed to be equal to the size of the matrix (number of taxa *x* number of sites $= 6 \times 49,881 = 299,286$). For the comparison between the different SelAC and 192 other models fitted using IQTree (Nguyen et al. 2015), see supplementary table S1, Supplementary Material online. In summary, the different SelAC models and FMutSel fitted the data better than any of the IQTree models.

these weights only strongly deviate from Grantham's (1974) original estimates in regard to composition weight, α_c, which is the ratio of noncarbon atoms in the end groups or rings to the number of carbon atoms in side chains. Our estimate of the composition weighting factor of $\alpha_c = 0.459$ is one-fourth the value estimate by Grantham suggests that the substitution process is less sensitive to this physicochemical property when shared ancestry and variation in stabilizing selection are taken into account.

It is important to note that the nonsynonymous/synonymous mutation ratio, or ω, which we estimated for each gene under the FMutSel model strongly correlated with SelAC's gene expression parameter ϕ. In fact, ω showed similar, though slightly reduced correlations, with the same empirical estimates of gene expression described above (fig. 3) This would give the impression that the same conclusions could have been gleaned using a much simpler model, both in terms of the number of parameters and the assumptions made. However, as we discussed earlier, not only is this model greatly restricted in terms of its biological feasibility, SelAC clearly performs better in terms of its fit to the data and biological realism.

For example, when we simulated the sequence for *S. cerevisiae*, starting from the ancestral sequence under both GTR + Γ and FMutSel, the functionality of the simulated sequence, defined as protein function in relation to the physicochemical distance from each amino acid to the optimal, moves away from the observed sequence. In contrast, SelAC remains near the functionality of the observed sequence (fig. 4*b*). This is somewhat unsurprising, given that both GTR + Γ and FMutSel are agnostic to the functionality of the gene, but it does highlight the improvement in biological realism in amino acid sequence evolution that SelAC provides. We do note that the adequacy of the SelAC model does vary among individual taxa, and does not always match the observed functionality. For instance, our simulations of *S. castellii* gene function is consistently higher than estimated from the data (fig. 4*c*). We suspect this is an indication that assuming a single set of optimal amino acid across all taxa is too simplistic. However, we cannot rule out violations of SelAC's other model assumptions such as, a single set of Grantham weights, a single α_G, or reductions in protein functionality **B** being solely a function of physicochemical distances *d* between sites.

Discussion

A central goal in evolutionary biology is to quantify the nature, strength, and, ultimately, shifts in the forces of natural selection relative to genetic drift and mutation. As data set size and complexity increase, so does the amount of potential information on these forces and their dynamics. As a result, there is a need for more complex and realistic models to accomplish this goal (Goldman et al. 1996; Thorne et al. 1996; Goldman et al. 1998; Halpern and Bruno 1998; Lartillot and Philippe 2004). Although extremely popular due to their elegance and computational efficiency, the utility of ω-based models in helping us reach this goal is substantially more limited than commonly recognized. Because these ω models use a single substitution matrix, they are only applicable for situations in which the substitution process and shifts in the selective environment are intrinsic to the sequence, such as with positive or negative frequency-dependent selection; these models do not describe stabilizing or diversifying selection as commonly envisioned (Endler 1986; Pellmyr 2002).

Starting with Halpern and Bruno (1998), a number of researchers have developed methods for linking heterogenous or site-specific selection on protein sequence and phylogenetics (e.g., Yang et al. 1998; Koshi et al. 1999; Dimmic et al. 2000; Koshi and Goldstein 2000; Robinson et al. 2003; Lartillot and Philippe 2004; Thorne et al. 2012; Rodrigue and Lartillot 2014). Halpern and Bruno (1998) calculated a vector of 20 expected amino acid frequencies for each amino acid site, making it the most general and most parameter rich of these methods. This generality, however, comes at the cost of being purely descriptive; there is no explicit biological mechanism proposed to explain the site-specific amino acid frequencies estimated. By grouping together amino sites with similar evolutionary behaviors, Lartillot and Philippe (2004) and Rodrigue and Lartillot (2014) retained the descriptive

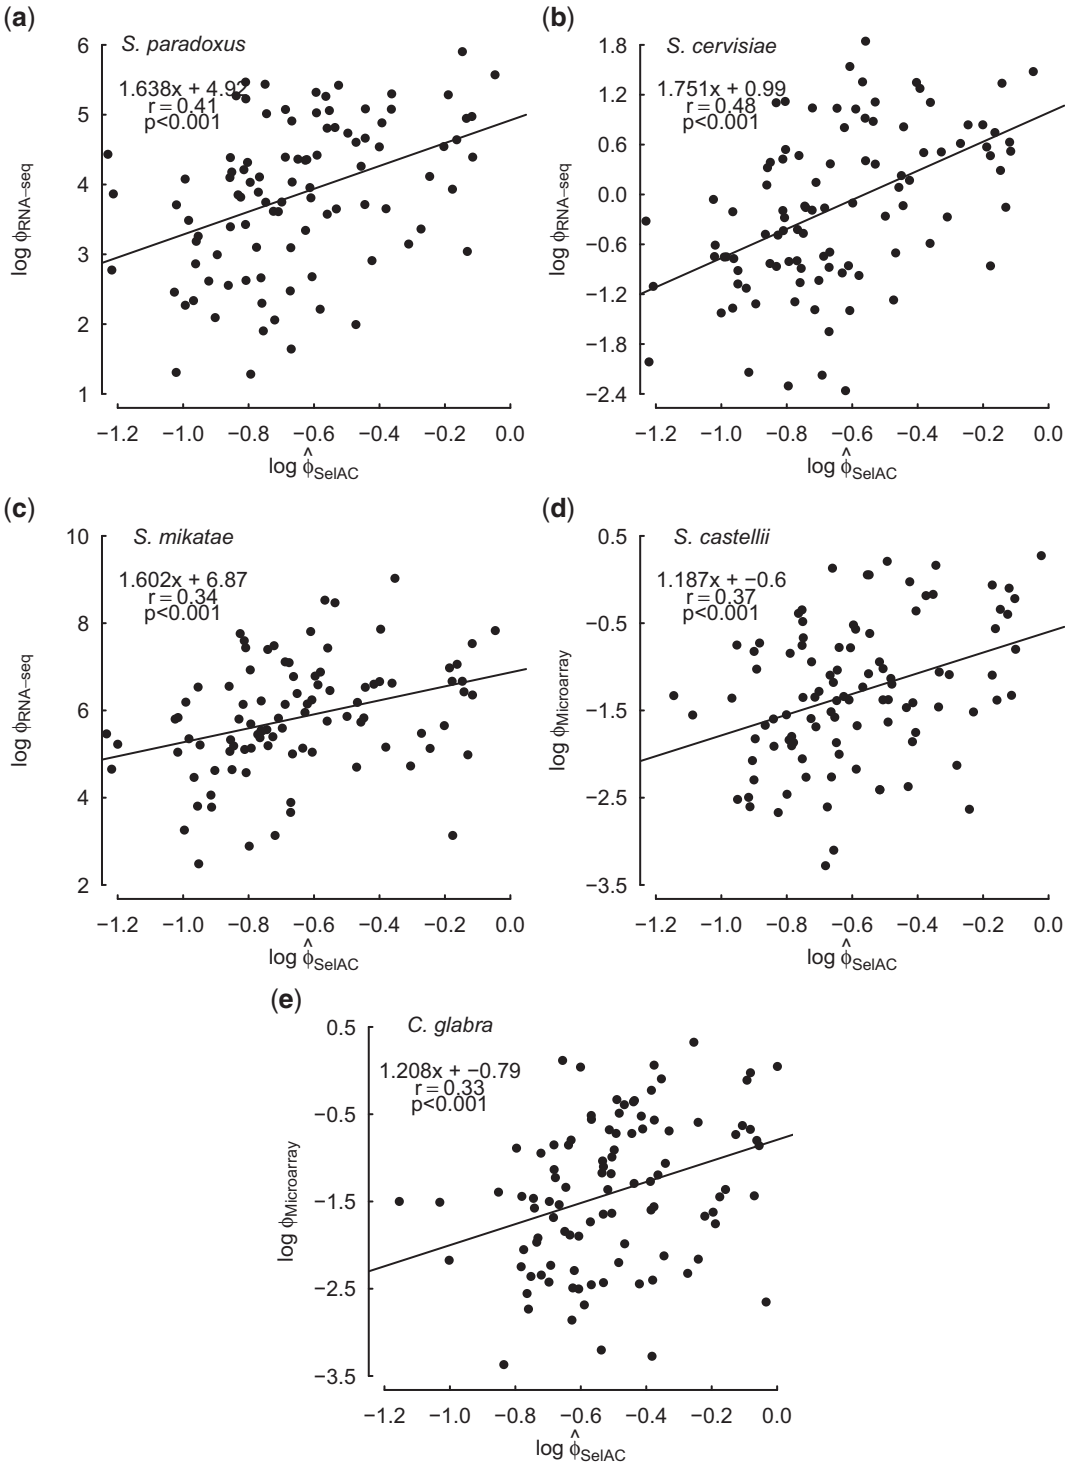

Fig. 2. Comparisons between estimates of average protein translation rate $\hat{\phi}_{\text{SelAC}}$ obtained from SelAC + Γ and direct measurements of expression for individual yeast taxa across the 100 selected genes from Salichos and Rokas (2013) measured during log-growth phase. Estimates of $\hat{\phi}_{\text{SelAC}}$ were generated by dividing the composite term ψ' by **B** ($\vec{a}_i|\vec{a}^*$). Gene expression was measured using either RNA-Seq (*a*–*c*) or microarray (*d*–*e*). The equations in the upper left-hand corner of each panel represent the regression fit and the Pearson correlation coefficient *r*.

nature of Halpern and Bruno (1998) work while greatly reduced the number of model parameters needed.

SelAC follows in this tradition of using multiple substitution matrices, but includes some key advances. First, by nesting a model of a sequence's cost–benefit function **C**/**B** within a broader model, SelAC allows us to formulate and test a hierarchical, mechanistic models of stabilizing selection. More precisely, our nested approach allows us to relax the

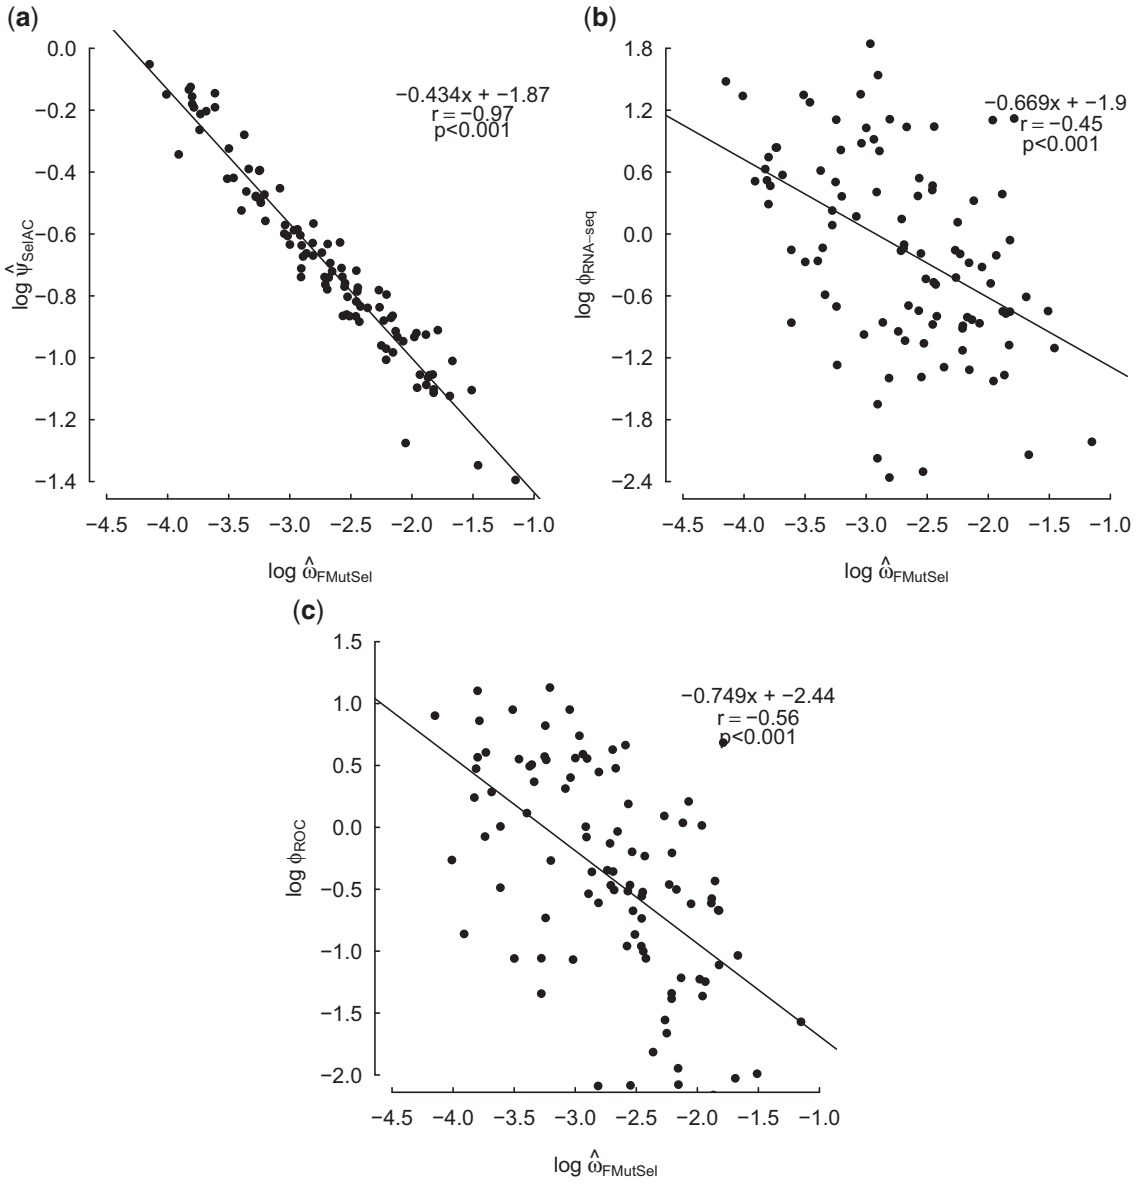

Fig. 3. Comparisons between $\omega_{FMutSel}$, which is the nonsynonymous/synonymous mutation ratio in FMutSel, SelAC + Γ estimates of protein functionality production rates ψhat (*a*), RNA-Seq-based measurements of mRNA abundance $\phi_{RNA-seq}$ (*b*), and ROC-SEMPPER's estimates of protein translation rates ϕ_{ROC}, which are based solely on *S. cerevisiae*'s patterns of codon usage bias (*c*), for *S. cerevisiae* across the 100 selected genes from Salichos and Rokas (2013). As in figure 2, the equations in the upper right-hand corner of each panel provide the regression fit and correlation coefficient.

assumption that physicochemical deviations from the optimal sequence $\vec{a}^*$ are equally disruptive at all sites within a protein. Indeed, SelAC results are consistent with the idea that the strength of stabilizing selection against physicochemical deviations from $\vec{a}^*$ varies between sites (ΔAICc =20,983; table 1). Second, because our substitution matrices are built on a formal description of a sequence's cost–benefit function **C**/**B**, we are able to efficiently parameterize 20 different matrices using a relatively small number of genome-wide parameters—for example, our physicochemical weightings, α_c, α_p, and α_v, and the shape parameter α_G for the distribution of selective strength *G* and one gene-specific expression parameter ψ. Although the **C**/**B** function on which SelAC currently rests is very simple, nevertheless, it leads to a dramatic increase in our ability to explain the sequence data we analyzed. Importantly, because SelAC uses a formal description of a sequence's **C**/**B**, replacing our assumptions with more sophisticated ones in the future is relatively straightforward. Third, our use of nested models also allows us to make biologically meaningful and testable predictions. By linking a gene's expression level to the strength of purifying selection it experiences, we are able to provide coarse estimates of gene expression. The anticorrelation between ω and gene expression indicates ω is often just a proxy for the strength, rather than nature, of natural selection on a sequence.

Thus, we believe our cost–benefit approach to be a substantial advance of the more simplistic ω models, is

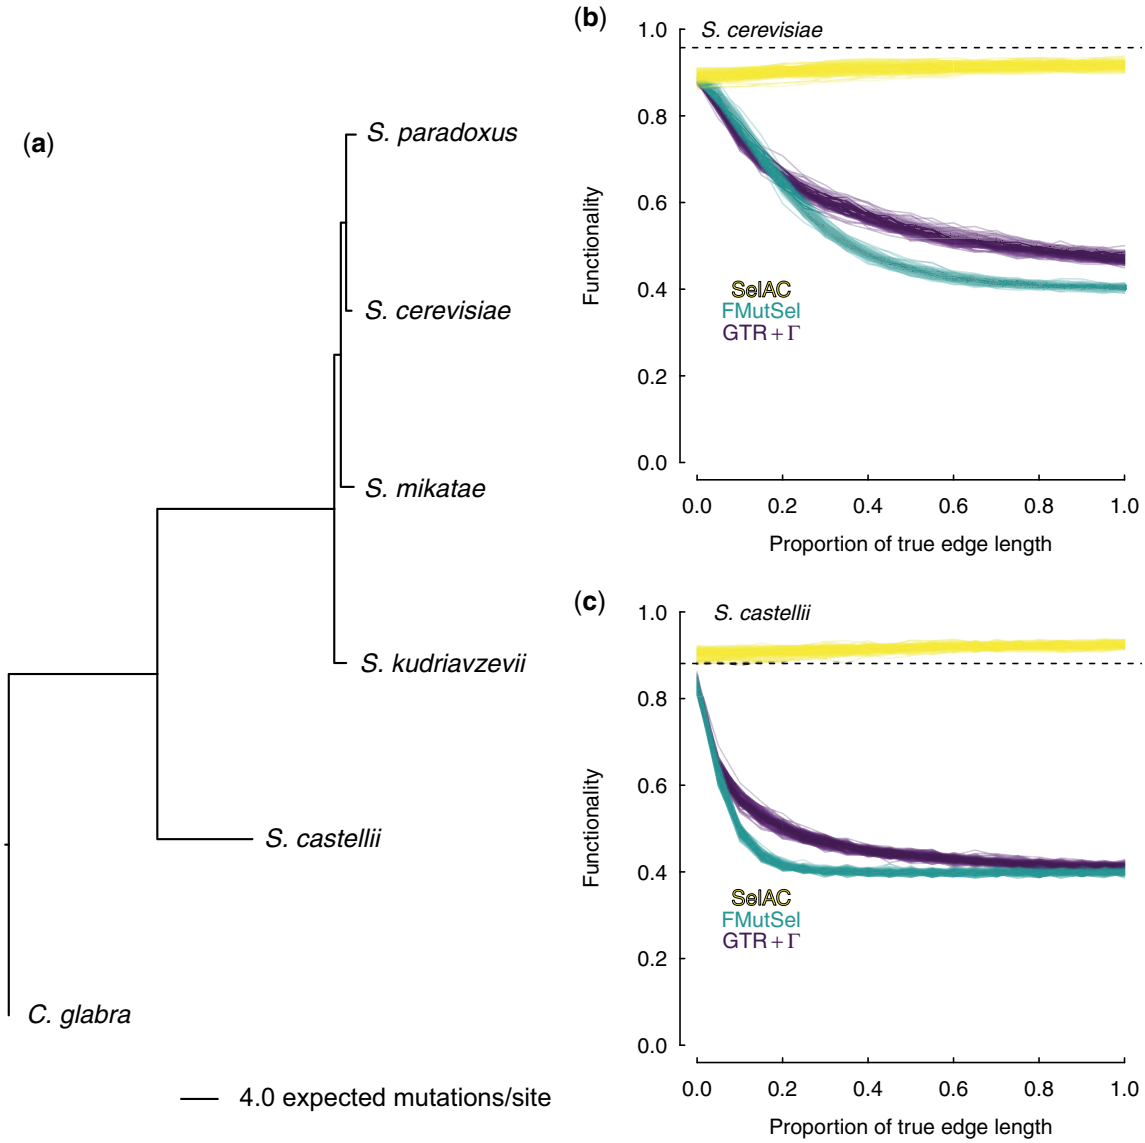

FIG. 4. (*a*) Maximum likelihood estimates of branch lengths under SelAC + Γ for 100 selected genes from Salichos and Rokas (2013). Tests of model adequacy for *S. cerevisiae* (*b*) and *S. castellii* (*c*) indicated that, when these taxa are removed from the tree, and their sequences are simulated, the parameters of SelAC + Γ exhibit functionality $\mathbf{B}(\vec{a}_{\text{obs}}|\vec{a}*)$ that is far closer to the observed (dashed black line) than data sets produced from parameters of either FMutSel or GTR + Γ.

complementary to the work of others in the field (e.g., Thorne et al. 2012; Rodrigue and Lartillot 2014), and, in turn, lays the foundation for more realistic work in the future. For instance, by assuming there is an optimal amino acid for each site, SelAC naturally leads to a nonsymmetrical and, thus, more cogent model of protein sequence evolution. Because the strength of selection depends on an additive function of amino acid physicochemical properties, an amino acid more similar to the optimum has a higher probability of replacing a more dissimilar amino acid than the converse situation. Furthermore, SelAC does not assume the system is always at the optimum or pessimum point of the fitness landscape, as occurs when $\omega < 1$ or >1, respectively.

Importantly, the cost–benefit approach underlying SelAC allows us to link the strength of selection on a protein sequence to its gene's expression level. Despite its well-recognized importance in determining the rate of protein evolution (e.g., Drummond et al. 2005, 2006), phylogenetic models have ignored the fact that expression levels vary between genes. In order to link gene expression and the strength of stabilizing selection on protein sequences, we simply assume that the strength of selection on a gene is proportional to the average protein synthesis rate of the gene.

One possible mechanism with some theoretical and empirical support which generates a linear relationship between the strength of selection and gene expression is the assumption of compensatory gene expression (Brown and Elliot 1997; Allison 2012; Lerman et al. 2012; Thiele et al. 2012; Zanger and Schwab 2013; King et al. 2015; Allison and Goulden 2017). That is, the assumption that any reduction in protein function is compensated for by an increase in the protein's production rate and, in turn, abundance. For example, a mutation that reduces the functionality of the protein to 90% of the optimal protein, would require $1/0.9 = 1.11$ of

these suboptimal proteins to be produced relative to the optimal protein in order to maintain the same amount of that protein's functionality in the cell. Because the energetic cost of an 11% increase in a protein's synthesis rate is proportional to its target synthesis rate, our assumptions naturally link changes in protein functionality and changes in gene expression and its associated costs. Such a response has been shown when error rates during protein production increases, however, the generality of such a response has to be determined (Goldsmith and Tawfik 2009). Furthermore, our model does not consider additional costs such as those imposed by misfolded proteins (Drummond and Wilke 2008). Nevertheless, the fact that our method allows us to explain 13–23% of the variation in gene expression measured using RNA-Seq, suggests that our assumption of cell compensation via increased expression is a reasonable starting point.

Furthermore, by linking expression and selection, SelAC provides a natural framework for combining information from protein-coding genes with very different rates of evolution; from low expression genes providing information on shallow branches to high-expression genes providing information on deep branches. This is in contrast to a more traditional approach of concatenating gene sequences together, which is equivalent to assuming the same average functionality production rate ψ for all of the genes, or more recent approaches where different models are fitted independently to different genes. Our results indicate that including a gene-specific ψ value vastly improves SelAC fits (table 1). Perhaps more convincingly, we find that the target functionally production rate ψ and the realized average protein synthesis rate $\phi = \psi/\boldsymbol{B}$ are reasonably well correlated with laboratory measurements and theoretical predictions of gene expression (Pearson's r=0.34–0.64; fig. 2, supplementary figs. S1 and S2, Supplementary Material online). The idea that quantitative information on gene expression is embedded within intragenomic patterns of synonymous codon usage is well accepted; our work shows that this information can also be extracted from comparative data at the amino acid level.

Of course, given the general nature of SelAC and the complexity of biological systems, other biological forces besides selection for reducing energy flux likely contribute to intergenic variation in the magnitude of stabilizing selection. Similarly, other physicochemical properties besides composition, volume, and charge likely contribute to site-specific patterns of amino acid substitution. For example, Blazej et al. (2017) have developed substitution matrices using various physicochemical properties. Thus, a larger and more informative set of physicochemical weights might improve our model fit and reduce the noise in our estimates of realized protein synthesis rates ϕ. Even if other physicochemical properties are considered, the idea of a consistent, genome-wide physicochemical weighting of these terms seems highly unlikely. Since the importance of an amino acid's physicochemical properties likely changes with its position in a folded protein, one way to incorporate such effects is to test whether the data supports multiple sets of physicochemical weights for either subsets of genes or regions within genes, rather than a single set.

Both of these points highlight the advantage of the detailed, mechanistic modeling approach underlying SelAC. Because there is a clear link between protein expression, synthesis cost, and functionality, SelAC can be extended by increasing the realism of the mapping between these terms and the coding sequences being analyzed. For example, SelAC currently assumes the optimal amino acid for any site is fixed along all branches. This assumption can be relaxed by allowing the optimal amino acid to change during the course of evolution along a branch. From a computational standpoint, the additive nature of selection between sites is desirable because it allows us to analyze sites within a gene largely independently of each other. From a biological standpoint, this additivity between sites ignores any nonlinear interactions between sites, such as epistasis, or between alleles, such as dominance. Thus, our work can be considered a first step to modeling these more complex scenarios.

For example, our current implementation ignores any selection on synonymous codon usage bias (CUB) (cf. Yang and Nielsen 2008; Pouyet et al. 2016). Including such selection is tricky because introducing the site-specific cost effects of CUB, which is consistent with the hypothesis that codon usage affects the efficiency of protein assembly or $\boldsymbol{C}$, into a model where amino acids affect protein function or $\boldsymbol{B}$, results in a cost–benefit ratio $\boldsymbol{C}/\boldsymbol{B}$ with epistatic interactions between all sites. These epistatic effects can likely be ignored under certain conditions or reasonably approximated based on an expectation of codon-specific costs (e.g., Kubatko et al. 2016). Nevertheless, it is difficult to see how one could identify such conditions without modeling the way in which codon and amino acid usage affects $\boldsymbol{C}/\boldsymbol{B}$.

This work also points out the potential importance of further investigation into model choice in phylogenetics. For likelihood models, use of AICc has become standard. However, how one determines the appropriate number of data points in a model is more complicated than generally recognized. Common sense suggests that data set size is increased by adding taxa and/or sites. In other words, a data set of 1,000 taxa and 100 sites must have more information on substitution models than a data set of 4 taxa and 100 sites. Our simple analyses support the hypothesis that the number of observations in a data set (number of sites $\times$ number of taxa) should be taken as the sample size for AICc, but this conclusion likely only applies when there is sufficient independence between taxa. For instance, one could imagine a phylogeny where one taxon is sister to a polytomy of 99 taxa that have zero length terminal branches. Absent measurement error or other intraspecific variation, one would have 100 species but only two unique trait values, and the only information about the process of evolution comes from what happens on the path connecting the lone taxon to the polytomy. Although this is a rather extreme example, it seems prudent for researchers to use a simulation-based approach similar to the one we take here to determine the appropriate means for calculating the effective number of data points in their data.

There are still significant shortcomings in the approach outlined here. Most worrisome are biological

oversimplifications in SelAC. For example, at its heart, SelAC assumes that suboptimal proteins can be compensated for, at a cost, simply by producing more of them. However, this is likely only true for proteins reasonably close to the optimal sequence. Different enough proteins will fail to function entirely: the active site will not sufficiently match its substrates, a protein will not properly pass through a membrane, and so forth. Yet, in our model, even random sequences still permit survival via an increased rate of protein production. Another deficiency is the assumption of site independence and unchanging optimal amino acids through time, neither of which matches reality (Pollock et al. 2012; Shah et al. 2015). We assume single nucleotide changes only, despite evidence from Kosiol et al. (2007) and Whelan and Goldman (2004) that better fits can be accomplished by allowing instantaneous mutations of multiple nucleotides simultaneously. Like the other oversimplifications previously discussed, these assumptions can be relaxed through further extension of this model.

A deeper potential issue comes from the nature of model fitting itself. Because SelAC, like all models, is built on very particular assumptions, it is likely these assumptions are violated in many, if not most, cases. Although all models are incomplete descriptions of reality and, thus, wrong (Box 1976), our aim with SelAC is to provide a model that is "less wrong" (Asimov 1989). Because SelAC's assumptions result in a model that is more consistent with the data, our analyses supports this idea. Whereas "less wrong," caution should be taken when interpreting SelAC's model parameters since the model itself is admittedly incomplete. In addition, other biological mechanisms could result in similar, if not identical, mathematical representations (Rabosky and Goldberg 2015; Beaulieu and O'Meara 2016) which, in turn, illustrates one limit to comparative sequence analysis.

There are also deficiencies in our implementation. Though reasonable to use for a given topology with 10s of species, it is currently too slow for practical use for large tree searches. Our work serves as a proof of concept, or of utility for targeted questions where a more realistic model may be of use (e.g., placement of particular taxa). Future work will encode SelAC models into a variety of mature, popular tree-search programs. SelAC also represents a challenging optimization problem: the nested models reduce parameter complexity vastly, but there are still numerous parameters to optimize, including the discrete parameter of the optimal amino acid at each site, which requires evaluating the fit of each of the 20 canonical amino acids. One way to avoid the use of discrete parameters at the expense of more of them would be to have SelAC estimate the optimum physicochemical values on a per site basis rather than a specific amino acid. Biologically such a model would be more realistic (as it is the properties that selection "sees," not the identity of the amino acid itself). Whereas estimating the optimal physicochemical properties for each site would increase the number of parameters estimated, it could speed up model fitting if estimating these values requires less than 20 evaluations per site. This scenario seems more likely when the number of physicochemical properties considered is small.

In spite of these difficulties, SelAC represents an important step in uniting phylogenetic and population genetic models (Higgs 2008). Most work in this area use a given tree to make inferences about the nature or distribution of selection coefficients (e.g., Hughes et al. 1990; Xia and Li 1998; McClellan and McCracken 2001; Woolley et al. 2003; Blazej et al. 2017). Whereas the work of Koshi et al. (1999), Dimmic et al. (2000), Koshi and Goldstein (2000), Robinson et al. (2003), Lartillot and Philippe (2004), Thorne et al. (2012), and Rodrigue and Lartillot (2014) do involve tree inferences, are all models of constant, stabilizing selection, SelAC can be generalized further to include diversifying selection. Specifically, by letting SelAC's sensitivity term G, which we now assume is ≥ 0, to take on negative values, SelAC will behave as if there is a pessimal, rather than optimal, amino acid for the given site. In this diversifying selection scenario, amino acids with physicochemical qualities more dissimilar to the pessimal amino acid are increasingly favored, potentially resulting in multiple fitness peaks.

The ability to extend our model and, in turn, sharpen our thinking about the nature of natural selection on amino acid sequences illustrates the value of moving from descriptive to more mechanistic models in general and phylogenetics in particular. How frequently diversifying selection of this nature occurs is an open, but addressable, question. Regardless of the frequency at which diversifying selection occurs, another question of interest to evolutionary biologists is, "How often does the optimal/pessimal amino sequence change along any given branch?" Due to its mechanistic nature, SelAC can also be extended to include changes in the optimal/pessimal sequence over a phylogeny using a hidden Markov modeling approach (Tuffley and Steel 1998; Penny et al. 2001; Whelan 2008; for an explicit shift the optimal sequence, see Tamuri et al. 2009). Extending SelAC in these ways, will allow researchers to explicitly model shifts in selection on protein sequences and, in turn, quantify their frequency and magnitude thus deepening our understanding of biological evolution. Such approaches would be challenging using nonmechanistic and parameter rich models which infer up to 19 different parameters per site category (Halpern and Bruno 1998; Tamuri et al. 2012, 2014; Rodrigue and Lartillot 2014).

In summary, SelAC allows biologically relevant population genetic parameters to be estimated from phylogenetic information while also dramatically improving fit and accuracy of phylogenetic models. By explicitly modeling the optimal/pessimal sequence of a gene, SelAC can be extended to include shifts in the optimal/pessimal sequence over evolutionary time. Moreover, it demonstrates that there remains substantially more information in the coding sequences used for phylogenetic analysis than other methods can access. Given the enormous amount of efforts expended to generate sequence data sets, it makes sense for researchers to continue developing more realistic models of sequence evolution in order to extract the biological information embedded in these data sets. The cost–benefit model we develop here is just one of many possible paths of mechanistic model development.

Materials and Methods

Overview

We model the substitution process as a classic Wright–Fisher process that includes the forces of mutation, selection, and drift (Fisher 1930; Kimura 1962; Wright 1969; Iwasa 1988; Berg and Lässig 2003; Sella and Hirsh 2005; McCandlish and Stoltzfus 2014). For simplicity, we ignore linkage effects and, as a result of this and other assumptions, sequences evolve in a site independent manner.

Because SelAC requires 20 families of 61×61 matrices, the number of parameters needed to implement SelAC would, without further assumptions, be extremely large (i.e., on the order of 74,000 parameters). To reduce the number of parameters needed, whereas still maintaining a high degree of biological realism, we construct our gene and amino acid-specific substitution matrices using a submodel nested within our substitution model, similar to approaches in Gilchrist (2007), Shah and Gilchrist (2011), and Gilchrist et al. (2015).

One advantage of a nested modeling framework is that it requires only a handful of genome-wide parameters such as nucleotide-specific mutation rates (scaled by effective population size N_e), amino acid side chain physicochemical weighting parameters, and a shape parameter describing the distribution of site sensitivities. In addition to these genome-wide parameters, SelAC requires a gene-specific functionality expression parameter ψ which describes the average rate at which the protein's functionality is produced by the organism or a gene's "average functionality production rate" for short. Currently, ψ is fixed across the phylogeny, though relaxing this assumption is a goal of future work. The gene-specific parameter ψ is multiplied by additional model terms to make a composite term ψ', which scales the strength and efficacy of selection for the optimal amino acid sequence relative to drift (see Implementation). In terms of the functionality of the protein encoded, we assume that for any given gene there exists an optimal amino acid sequence $\vec{a}^*$ and that, by definition, a complete, error free peptide consisting of $\vec{a}^*$ provides one unit of the gene's functionality. We also assume that natural selection favors genotypes that are able to synthesize their proteome more efficiently than their competitors and that each savings of an high energy phosphate bond per unit time leads to a constant proportional gain in fitness A_0 (which was q in our previous work; Gilchrist 2007). SelAC also requires the specification (as part of parameter optimization) of an optimal amino acid a^* at each position within a coding sequence. This requirement of one a^* per site makes our $\vec{a}^*$ the largest category of parameters SelAC estimates. Despite the need to specify a^* for each site, because we use a submodel to derive our substitution matrices, SelAC estimates a relatively small number of the parameters when compared with more general approaches where the fitness of each amino acid is allowed to vary freely of any physicochemical properties (Halpern and Bruno 1998; Lartillot and Philippe 2004; Rodrigue and Lartillot 2014).

As with other phylogenetic methods, SelAC generates estimates of branch lengths and nucleotide-specific mutation rates. In addition, the method can also be used to make quantitative inferences on the optimal amino acid sequence of a given protein as well as the realized average synthesis rate of each protein used in the analysis. The mechanistic basis of SelAC also means it can be easily extended to include more biological realism and test more explicit hypotheses about sequence evolution.

Mutation Rate Matrix μ

We begin with a 4×4 nucleotide mutation matrix that describes mutation rates between different bases and, in turn, different codons. For our purposes, we rely on the general unrestricted model (UNREST from Yang 1994) because it imposes no constraints on the instantaneous rate of change between any pair of nucleotides. More constrained models, such as the Jukes–Cantor, Hasegawa–Kishino–Yano, or the GTR model, could also be used.

The 12 parameter UNREST model defines the relative rates of change between a pair of nucleotides. Thus, we arbitrarily set the G $\rightarrow$ T mutation rate to 1, resulting in 11 free mutation rate parameters in the 4×4 mutation nucleotide mutation matrix. The nucleotide mutation matrix is also scaled by a diagonal matrix π whose entries, $\pi_{i,i}$, correspond to the equilibrium frequencies of each base. These equilibrium nucleotide frequencies are determined by analytically solving $\pi \times \boldsymbol{Q} = 0$. We use this $\boldsymbol{Q}$ to populate a 61×61 codon mutation matrix $\boldsymbol{\mu}$, whose entries $\mu_{i,j}$ when $i \neq j$ describes the mutation rate from codon i to j and $\mu_{i,i} = -\sum_j \mu_{i,j}$. We generate this matrix using a "weak mutation" assumption, such that evolution is mutation limited, codon substitutions only occur one nucleotide at a time. As a result, the rate of change between any pair of codons that differ by more than one nucleotide is zero.

Although the overall model does not assume equilibrium, we still need to scale our mutation matrices $\boldsymbol{\mu}$ by a scaling factor S. As traditionally done, we rescale our time units such that at equilibrium, one unit of branch length represents one expected mutation per site (which equals the substitution rate under neutrality). More explicitly, $S = -\left(\sum_{i \in codons} \mu_{i,i} \pi_{i,i}\right)$ where the final mutation rate matrix is the original mutation rate matrix multiplied by $1/S$.

Protein Synthesis Cost–Benefit Function η

SelAC links fitness to the product of the cost–benefit function of a gene η and the organism's average target synthesis rate of the functionality provided by gene ψ. As a result, the average flux energy an organism spends to meet its target functionality provided by the gene is $\eta \times \psi$. Compensatory changes that allow an organism to maintain functionality even with loss of one or both copies of a gene are widespread. There is evidence of compensation for protein function. Metabolism with gene expression models (ME-models) link those factors to successfully make predictions about response to perturbations in a cell (Lerman et al. 2012; King et al. 2015). For example, an ME-model for *E. coli* successfully predicted gene expression levels in vivo (Thiele et al. 2012). Here we assume that for finer scale problems than entire loss (e.g., a 10% loss of functionality) the compensation is more production of the

protein. The particular type of dosage compensation assumed by SelAC in response to stress (e.g., reduced functionality) is commonly assumed in microbial ecology (Allison 2012; Allison and Goulden 2017). Our assumption is also consistent with the Michaelis–Menten enzyme kinetics. Moreover, there is evidence that mutations can influence expression level, though this does not always match our expression compensation assumption (Brown and Elliot 1997; Zanger and Schwab 2013). In order to link genotype to our cost–benefit function $\eta = \mathbf{C}/\mathbf{B}$, we begin by defining our benefit function $\mathbf{B}$.

Benefit

Our benefit function $\mathbf{B}$ measures the functionality of the amino acid sequence $\vec{a}_i$ encoded by a set of codons $\vec{c}_i$, that is, $a(\vec{c}_i) = \vec{a}_i$ relative to that of an optimal sequence $\vec{a}^*$. By definition, $\boldsymbol{B}(\vec{a}^*|\vec{a}^*) = 1$ and $\boldsymbol{B}(\vec{a}_i|\vec{a}^*) < 1$ for all other sequences. We assume all amino acids within the sequence contribute to protein function and that this contribution declines as an inverse function of physicochemical distance from each amino acid to the optimal one. Formally, we assume that

$$\boldsymbol{B}(\vec{a}|\vec{a}^*) = \left(\frac{1}{n}\sum_{p=1}^{n}(1 + G_p \mathrm{d}(a_p, a_p^*))\right)^{-1}, \qquad (1)$$

where n is the length of the protein, $\mathrm{d}(a_p, a_p^*)$ is a weighted physicochemical distance between the amino acid encoded at a given position p and and the optimal amino acid for that position a_p^*. There are many possible measures for physiochemical distance; we use Grantham (1974) distances by default, though others may be chosen. For simplicity, we assume all nonsense mutations are lethal by defining the physicochemical distance between a stop codon and a sense codon as ∞. The term G_p describes the sensitivity of the protein's function to physicochemical deviation from the optimum at site position p. We assume that $G_p \sim Gamma\,(shape = \alpha_G, rate = \alpha_G)$ in order to ensure $\mathbb{E}(G_p) = 1$. Given the definition of the Gamma distribution, the variance in G_p is equal to $\mathrm{shape}/\mathrm{rate}^2 = 1/\alpha_G$. We note that at the limit of $\alpha_G \to \infty$, the model becomes equivalent to assuming uniform site sensitivity where $G_p = 1$ for all positions p. Furthermore, $\mathbf{B}\,(\vec{a}_i|\vec{a}^*)$ is inversely proportional to the average physicochemical deviation of an amino acid sequence $\vec{a}_i$ from the optimal sequence $\vec{a}^*$ weighted by each site's sensitivity to this deviation. $\mathbf{B}\,(\vec{a}_i|\vec{a}^*)$ can be generalized to include second and higher order terms of the distance measure d.

Cost

Protein synthesis involves both direct and indirect assembly costs. Direct costs consist of the high energy phosphate bonds $\sim P$ of ATPs or GTPs used to assemble the ribosome on the mRNA, charge tRNA's for elongation, move the ribosome forward along the transcript, and terminate protein synthesis. As a result, direct protein assembly costs are the same for all proteins of the same length. Indirect costs of protein assembly are potentially numerous and could include the cost of amino acid synthesis as well the cost and efficiency with which the protein assembly infrastructure such as ribosomes, aminoacyl-tRNA synthetases, tRNAs, and mRNAs are used. When these indirect costs are combined with sequence-specific benefits, the probability of a mutant allele fixing is no longer independent of the rest of the sequence (Gilchrist et al. 2015) and, as a result, model fitting becomes substantially more complex. Thus for simplicity, in this study we ignore indirect costs of protein assembly that vary between genotypes and define

$$\begin{aligned}\mathbf{C}(\vec{c}_i) &= \text{Direct energetic cost of protein synthesis,}\\ &= A_1 + A_2 n\end{aligned}$$

where A_1 and A_2 represent the direct cost, in high energy phosphate bonds, of ribosome initiation and peptide elongation, respectively, where $A_1 = A_2 = 4 \sim P$.

Defining Physicochemical Distances

Assuming that functionality declines with an amino acid a_i's physicochemical distance from the optimum amino acid a^* at each site provides a biologically defensible way of mapping genotype-to-protein function that requires relatively few free parameters. In addition, SelAC naturally lends itself to model selection since one could compare the quality of SelAC fits using different mixtures of physicochemical properties. Following Grantham (1974), we focus on using composition c, polarity p, and molecular volume v of each amino acid's side chain residue to define our distance function, but the model and its implementation can flexibly handle a variety of properties. We use the Euclidian distance between residue properties where each property c, p, and v has its own weighting term, α_c, α_p, α_v, respectively, which we refer to as "Grantham weights." Because physicochemical distance is ultimately weighted by a gene's specific average protein synthesis rate ψ, another parameter we estimate, there is a problem with parameter identifiability. The scale of gene expression is affected by how we measure physicochemical distances which, in turn, is determined by our choice of Grantham weights. As a result, by default we set $\alpha_v = 3.990 \times 10^{-4}$, the value originally estimated by Grantham, and recognize that our estimates of α_c and α_p and ψ are scaled relative to this choice for α_v. More specifically,

$$\begin{aligned}&\mathrm{d}(a_i,a^*)\\ &= \sqrt{\alpha_c[c(a_i)-c(a^*)]^2+\alpha_p[p(a_i)-p(a^*)]^2+\alpha_v[v(a_i)-v(a^*)]^2}.\end{aligned}$$

Linking Protein Synthesis to Allele Substitution

Next, we link the protein synthesis cost–benefit function η of an allele with its fixation probability. First, we assume that each protein encoded within a genome provides some beneficial function and that the organism needs that functionality to be produced at a target average rate ψ. Again, by definition, the optimal amino acid sequence for a given gene, $\vec{a}^*$, produces one unit of functionality, that is,

$\boldsymbol{B}(\vec{a}^*) = 1$. Second, we assume that the actual average rate a protein is synthesized ϕ is regulated by the organism to ensure that functionality is produced at rate ψ. As a result, it follows that $\phi = \psi/\boldsymbol{B}(\vec{a}|\vec{a}^*)$ and the energetic burden of a suboptimal amino acid increases the more it decreases the protein's functionality, **B**. In other words, the average production rate of a protein $\vec{a}$ with relative functionality $\boldsymbol{B}(\vec{a}) < 1$ must be $1/\boldsymbol{B}(\vec{a}|\vec{a}^*)$ times higher than the production rate needed if the optimal amino acid sequence $\vec{a}^*$ was encoded since $\boldsymbol{B}(\vec{a}^*|\vec{a}^*) = 1$. For example, a cell with an allele $\vec{a}$ where $\boldsymbol{B}(\vec{a}|\vec{a}^*) = 9/10$ would have to produce the protein at rate $\phi = 10/9 \times \psi = 1.11\psi$. Similarly, a cell with an allele $\vec{a}$ where $\boldsymbol{B}(\vec{a}|\vec{a}^*) = 1/2$ will have to produce the protein at $\phi = 2\psi$. In contrast, a cell with the optimal allele $\vec{a}^*$ would have to produce the protein at rate $\phi = \psi$.

Third, we assume that every additional high energy phosphate bond, $\sim P$, spent per unit time to meet the organism's target function synthesis rate ψ leads to a slight and proportional decrease in fitness W. This assumption, in turn, implies

$$W_i(\vec{c}) \propto \exp[-A_0\eta(\vec{c}_i)\psi].$$

where A_0 describes the proportional decline in fitness with every $\sim P$ wasted per unit time. Because A_0 shares the same time units as ψ and ϕ and only occurs in SelAC in conjunction with ψ, we do not need to explicitly identify our time units. Instead, we recognize that our estimates of ψ share an unknown scaling term.

Correspondingly, the ratio of fitness between two genotypes is,

$$\frac{W_i}{W_j} = \frac{\exp[-A_0\eta(\vec{c}_i)\psi]}{\exp[-A_0\eta(\vec{c}_j)\psi]} = \exp\left[-A_0\left(\eta(\vec{c}_i) - \eta(\vec{c}_j)\right)\psi\right].$$

Given our formulations of **C** and **B**, the fitness effects between sites are multiplicative and, therefore, the substitution of an amino acid at one site can be modeled independently of the amino acids at the other sites within the coding sequence. As a result, the fitness ratio for two genotypes differing at multiple sites simplifies to

$$\frac{W_i}{W_j} = \exp\left[-\left(\frac{A_0(A_1 + A_2 n)}{n}\right)\sum_{p\in\mathbb{P}}\left[\mathrm{d}\left(a_{i,p}, a_p^*\right) - \mathrm{d}\left(a_{j,p}, a_p^*\right)\right]G_p\psi\right],$$

where $\mathbb{P}$ represents the codon positions in which $\vec{c}_i$ and $\vec{c}_j$ differ. Fourth, we make a weak mutation assumption, such that alleles can differ at only one position at any given time, that is, $|\mathbb{P}| = 1$, and that the population is evolving according to a Wright–Fisher process. As a result, the probability a new mutant, j, introduced via mutation into a resident population i with effective size N_e will go to fixation is,

$$u_{i,j} = \frac{1 - \left(W_i/W_j\right)^b}{1 - \left(W_i/W_j\right)^{2N_e}} = \frac{1 - \exp\{\mathrm{C}\left[\mathrm{d}(a_i, a^*) - \mathrm{d}\left(a_j, a^*\right)\right]G_p\psi\mathrm{b}\}}{1 - \exp\{\mathrm{C}\left[d(a_i, a^*) - d\left(a_j, a^*\right)\right]G_p\psi 2N_e\}},$$

where $\mathrm{C} = -(\mathrm{A}_0/n)(\mathrm{A}_1 + \mathrm{A}_2 n)$, and b = 1 for a diploid population and 2 for a haploid population (Kimura 1962; Wright 1969; Iwasa 1988; Berg and Lässig 2003; Sella and Hirsh 2005). Finally, assuming a constant mutation rate between alleles i and j, $\mu_{i,j}$ when i the substitution rate from allele i to j can be modeled as,

$$q_{i,j} = \frac{2}{b}\mu_{i,j}N_e u_{i,j},$$

where given the substitution model's weak mutation assumption, $N_e\mu \ll 1$. In the end, each optimal amino acid has a separate 61×61 substitution rate matrix $\boldsymbol{Q}_a$, which incorporates selection for the amino acid (and the fixation rate matrix this creates) as well as the common mutation parameters across optimal amino acids. This results in the creation of 20 **Q** matrices, one for each amino acid and each with 3,721 entries that are based on a relatively small number of model parameters (1–11 mutation rates, 2 free Grantham weights, the cost of protein assembly, A_1 and A_2, the gene-specific target functionality synthesis rate ψ, and optimal amino acid at each position p, a_p*). These model parameters can either be specified a priori and/or estimated from the data.

Given our assumption of independent evolution among sites, it follows that the probability of the whole data set is the product of the probabilities of observing the data at each individual site. Thus, the likelihood $\mathcal{L}$ of amino acid a being optimal at a given site position p is calculated as

$$\mathcal{L}\left(\boldsymbol{Q}_a|\boldsymbol{D}_p, \boldsymbol{T}\right) \propto \boldsymbol{P}\left(\boldsymbol{D}_p|\boldsymbol{Q}_a, \boldsymbol{T}\right). \quad (2)$$

In this case, the data, $\boldsymbol{D}_p$, are the observed codon states at position p for the tips of the phylogenetic tree with topology $\boldsymbol{T}$. For our purposes we take $\boldsymbol{T}$ as given, but it could be estimated as well. The pruning algorithm of Felsenstein (1981) is used to calculate $\mathcal{L}\left(\boldsymbol{Q}_a|\boldsymbol{D}_p, \boldsymbol{T}\right)$. The log of the likelihood is maximized by estimating the genome scale parameters that consist of 11 mutation parameters, which are implicitly scaled by $2N_e/b$, and 2 Grantham distance parameters, α_c and α_p, and the sensitivity distribution parameter α_G. Because A_0 and ψ always co-occur and are scaled by N_e, for each gene we estimate a composite term $\psi' = A_0\psi N_e b$ and the optimal amino acid for each position a_p* of the protein. When estimating α_G, the likelihood then becomes the average likelihood which we calculate using the generalized Laguerre quadrature with $k = 4$ points (Felsenstein 2001).

Finally, we note that because we infer the ancestral state of the system, our approach does not rely on any assumptions of model stationarity. Nevertheless, as our branch lengths grow the probability of observing a particular amino acid a at a given site approaches a stationary value proportional to

$W(a)^{2N_e-b}$ and any effects of mutation bias (Sella and Hirsh 2005).

Implementation

All methods described above are implemented in the new R package, `selac` available through GitHub (https://github.com/bomeara/selac) and has been posted on CRAN. Our package requires as input a set of fasta files that each contain an alignment of coding sequence for a set of taxa, and the phylogeny depicting the hypothesized relationships among them. In addition to the SelAC models, we implemented the GY codon model of Goldman and Yang (1994), the FMutSel mutation-selection model of Yang and Nielsen (2008), and the standard GTR nucleotide model that allows for Γ-distributed rates across sites. These likelihood-based models represent a sample of the types of popular models often fit to codon data.

For the SelAC models, the starting guess for the optimal amino acid at a site comes from "majority" rule, where the initial optimum is the most frequently observed amino acid at a given site (ties resolved randomly). Our optimization routine utilizes a four stage hill climbing approach. More specifically, within each stage a block of parameters are optimized, whereas the remaining parameters are held constant. The first stage optimizes the block of branch length parameters. The second stage optimizes the block of gene-specific composite parameters $\psi' = A_0\psi N_e b$. The third stage optimizes SelAC's parameters shared across the genome α_c and α_p, and the sensitivity distribution parameter α_G. The fourth stage estimates the optimal amino acid at each site a^*. This entire four stage cycle is repeated up to six more times, using the estimates from the previous cycle as the initial conditions for the new one. The search is terminated when the improvement in the log-likelihood between cycles is less than 10^{-8} at which point we consider the ML solution found and the search is terminated. For optimization of a given set of parameters, we rely on a bounded subplex routine (Rowan 1990) in the package NLoptR (Johnson 2012) to maximize the log-likelihood function. To ensure the robustness of our results, we perform a set of independent analyses with different sets of naive starting points with respect to the gene-specific composite ψ' parameters, α_c, and α_p and were able to repeatedly reach the same log-likelihood (ln L) peak in our parameter space. Confidence in the parameter estimates can be generated by an "adaptive search" procedure that we implemented to provide an estimate of the parameter space that is some predefined likelihood distance (e.g., 2 ln $\mathcal{L}$ units) from the MLE, which follows Beaulieu and O'Meara (2016) and Edwards (1984).

We note that our current implementation of SelAC is slow, and is best suited for data sets with relatively small number of taxa (i.e., 10s not 100s). This limitation is largely due to the size and quantity of matrices we create and manipulate to calculate the log-likelihood of an individual site. Ongoing work will address the need for speed, with the eventual goal of implementing SelAC in popular phylogenetic inference toolkits, such as RevBayes (Hoehna et al. 2016), PAML (Yang 2007), and RAxML (Stamatakis 2006).

Simulations

We evaluated the performance of our codon model by simulating data sets and estimating the bias of the inferred model parameters from these data. Our "known" parameters under a given generating model were based on fitting SelAC to the 106 gene data set and phylogeny of Rokas et al. (2003). The tree used in these analyses is outdated with respect to the current hypothesis of relationships within *Saccharomyces*, but we rely on it simply as a training set that is separate from our empirical analyses (see Analysis of Yeast Genomes and Tests of Model Adequacy). Bias in the model parameters were assessed under two generating models: one where we assumed a model of SelAC with uniform sensitivity across sites (i.e., $G_p = 1$ for all sites, i.e., $\alpha_G = \infty$), and one where we used the Gamma distribution joint shape and rate parameter α_G estimated from the empirical data. Under each of these two scenarios, we used parameter estimates from the corresponding empirical analysis and simulated 50 five-gene data sets. For the gene-specific composite parameter ψ' the "known" values used for the simulation were five evenly spaced points along the rank order of the estimates across the 106 genes. The MLE estimate for a given replicate were taken as the fit with the highest log-likelihood after running five independent analyses with different sets of naive starting points with respect to the composite ψ'_g parameter, α_c, and α_p. All analyses were carried out in our `selac` R package.

Analysis of Yeast Genomes and Tests of Model Adequacy

We focus our empirical analyses on the large yeast data set and phylogeny of Salichos and Rokas (2013). As a model system, the yeast genome is an ideal system to examine our phylogenetic estimates of gene expression and its connection to real world measurements of these data within individual taxa. The complete data set of Salichos and Rokas (2013) contains 1,070 orthologs, where we selected 100 at random for our analyses. We also focus our analyses on *Saccharomyces sensu stricto* and their sister taxon *Candida glabrata*, and we used the phylogeny depicted in figure 1 of Salichos and Rokas (2013) for our fixed tree. We fit the two SelAC models described above (i.e., SelAC and SelAC + Γ), as well as two codon models, GY and FMutSel, and a standard GTR + Γ nucleotide model. The FMutSel model assumes that the amino acid frequencies are determined by functional requirements of the protein, whereas the other models make no assumptions about amino acid frequencies. In all cases, we assumed that the model was partitioned by gene, but with branch lengths linked across genes.

We also compared SelAC models with 195 codon models in IQtree (Nguyen et al. 2015). This is popular software implementing various (e.g., Muse and Gaut 1994; Goldman and Yang 1994; Schneider et al. 2005; Kosiol et al. 2007) models of evolution, including codon models. Most analyses within a set of software tools focus on differences in log-likelihood values between models. As a result, some software tools sometimes fail to include solely data-dependent terms, which function as constants, in their calculations. Failure to include these terms can make model comparison between software packages

problematic. For simplicity, no constants are dropped in the log-likelihood calculations for SelAC. Furthermore, whereas there are no identical models in SelAC and IQtree, similar models have similar likelihoods, suggesting the log-likelihood values between SelAC and IQtree are comparable. We note, however, that minor differences in model implementation can lead to small differences in log-likelihood (for example, how missing data are handled, as SelAC treats it as an absence at that site for that taxon, whereas some other software integrates across all possible states).

For SelAC, we compared our estimates of $\phi' = \psi'/\boldsymbol{B}$, which represents the average protein synthesis rate of a gene, to estimates of gene expression from empirical data. Specifically, we examined gene expression data for five of the six species measured during log-growth phase. Gene expression in this context corresponds to mRNA abundances, which were measured using either microarrays (*C. glabrata* and *S. castellii*, or RNA-Seq (*S. paradoxus*, *S. mikatae*, and *S. cerevisiae*). We obtained expression data for the remaining species, *S. kudriavzevii*, which was measured at the beginning of the stationary phase from the Gene Expression Omnibus. *Saccharomyces*, however, only enter the stationary growth phase in response to severe stress, such as starvation. In addition, only 56% of the genes examined with SelAC had expression measurements available. For these reasons, we excluded *S. kudriavzevii* from our comparisons of empirical gene expression.

For further comparison, we also predicted the average protein synthesis rate for each gene ϕ by analyzing gene and genome-wide patterns of synonymous codon usage using ROC-SEMPPR (Gilchrist et al. 2015) for each individual genome. Although, like SelAC, ROC-SEMPPR uses codon level information, it does not rely on any interspecific comparisons and, unlike SelAC, uses only the intra and intergenic frequencies of synonymous codon usage as its data. Nevertheless, ROC-SEMPPR predictions of gene expression ϕ correlates strongly (Pearson's r=0.53–0.74) with a wide range of laboratory measurements of gene expression (Gilchrist et al. 2015).

Although one of our main objectives was to determine the improvement of fit that SelAC has with respect to other standard phylogenetic models, we also evaluated the adequacy of SelAC. Model fit, measured with assessments such as the AIC, can tell which model is least bad as an approximation for the data, but it does not reveal whether a model is actually doing a good job of representing the data. An adequate model does the latter, one measure of which is that data generated under the model resemble real data (Goldman 1993). For example, Beaulieu et al. (2013) assessed whether parsimony scores and the size of monomorphic clades of empirical data were within the distributions of simulated data under a new model and the best standard model; if the empirical summaries were outside the range for each, it would have suggested that neither model was adequately modeling this part of the biology.

In order to test adequacy for a given gene we first remove a particular taxon from the data set and the phylogeny. A marginal reconstruction of the likeliest sequence across all remaining nodes is conducted under the model, including the node where the pruned taxon attached to the tree. The marginal probabilities of each site are used to sample and assemble the starting coding sequence. This sequence is then evolved along the branch, periodically being sampled and its current functionality assessed. We repeat this process 100 times and compare the distribution of trajectories against the observed functionality calculated for the gene. For comparison, we also conducted the same test, by simulating the sequence under the standard GTR + Γ nucleotide model, which is often used on these data but does not account for the fact that the sequences are protein coding, and under FMutSel, which includes selection on codons but in a fundamentally different way as our model.

The Appropriate Estimator of Bias for AIC

As part of the model set described above, we also included a reduced form of each of the two SelAC models, SelAC and SelAC + Γ. Specifically, rather than optimizing the amino acid at any given site, we assume the most frequently observed amino acid at each site is the optimal amino acid a^*. We refer to these "majority rule" models as SelAC$_M$ and SelAC$_M$ + Γ and note that these majority rule formulations greatly accelerate model fitting.

Since these majority rule models assume that the optimal amino acids are known prior to fitting of our model, it is tempting to reduce the count of estimated parameters in the model by the number of parameters estimated using majority rule. Whereas using majority rule does not necessarily provide the most likely parameter estimate, it nevertheless uses the data to generate the estimate and represents a parameter estimated from the data. Thus, despite having become standard behavior in the field of phylogenetics, this reduction is potentially statistically inappropriate. Because the difference in the number of parameters K when counting or not counting the number of nucleotide sites drops out when comparing nucleotide models with AIC, this statistical issue does not apply to nucleotide models. It does, however, matter for AICc, where K and the sample size n combine in the penalty term. This also matters in our case, where the number of estimated parameters for the majority rule estimation differs based on whether one is looking at codons or single nucleotides.

In phylogenetics two variants of AICc are used. In comparative methods (e.g., Butler and King 2004; O'Meara et al. 2006; Beaulieu et al. 2013) the number of data points, n, is taken as the number of taxa. More taxa allow the fitting of more complex models, given more data. However, in DNA evolution, which is effectively the same as a discrete character model used in comparative methods, the n is taken as the number of sites. Obviously, both cannot be correct. This uncertainty was highlighted by Posada and Buckley (2004): they chose to use number of sites, but mentioned in their discussion that sample size also depends on the number of taxa. Sullivan and Joyce (2005) also mention that although the number of sites is often taken as sample size, whether that is appropriate in phylogenetics is not entirely clear. One approach incorporating both number of taxa and sites in

calculating AICc is the program SURFACE implemented by Ingram and Mahler (2013), which uses multiple characters and taxa. Although its default is to use AIC to compare models, if one chooses to use AICc, the number of samples is taken as the product of number of sites and number of taxa.

Recently, Jhwueng et al. (2014) performed an analysis that investigated what variant of AIC and AICc worked best as an estimator, but the results were inconclusive. Here, we have adopted and extended the simulation approach of Jhwueng et al. (2014) in order to examine a large set of different penalty functions and how well they approximate the remaining portion of the KL divergence between two models after accounting for the deviance (i.e., $-2 \ln \mathcal{L}$) (see Supplementary material online for more details).

Supplementary Material

Supplementary data are available at *Molecular Biology and Evolution* online.

Acknowledgments

This work was supported in part by NSF Awards MCB-1120370 (M.A.G. and R.Z.), MCB-1546402 (A. Von Arnim and M.A.G.) and DEB-1355033 (B.C.O., M.A.G., and R.Z.) with additional support from The University of Tennessee Knoxville and University of Arkansas (J.M.B.). J.C. and J.M.B. received support as Postdoctoral Fellows and C.L. received support as a Graduate Student Fellow at the National Institute for Mathematical and Biological Synthesis, an Institute sponsored by the National Science Foundation through NSF Award DBI-1300426, with additional support from UTK. The authors would like to thank Premal Shah, Todd Oakley, and four anonymous reviewers for their helpful criticisms and suggestions for this work.

References

Allison S. 2012. A trait-based approach for modelling microbial litter decomposition. *Ecol Lett*. 15(9): 1058–1070.

Allison S, Goulden M. 2017. Consequences of drought tolerance traits for microbial decomposition in the dement model. *Soil Biol Biochem*. 107:104–113.

Anisimova M. 2012. Parametric models of codon evolution. In: Cannarozzi GM and Schneider A, editors. Codon evolution: mechanisms and models. Oxford: Oxford University Press. p. 12–33.

Asimov I. 1989. The relativity of wrong. *Skeptical Inquirer* 14(1): 35–44.

Beaulieu JM, O'Meara BC. 2016. Detecting hidden diversification shifts in models of trait-dependent speciation and extinction. *Syst Biol*. 65(4): 583–601.

Beaulieu JM, O'Meara BC, Donoghue MJ. 2013. Identifying hidden rate changes in the evolution of a binary morphological character: the evolution of plant habit in campanulid angiosperms. *Syst Biol*. 62(5): 725–737.

Berg J, Lässig M. 2003. Stochastic evolution and transcription factor binding sites. *Biophysics* 48(S1): S36–S44.

Blazej P, Mackiewicz D, Grabinska M, Wnetrzak M, Mackiewicz P. 2017. Optimization of amino acid replacement costs by mutational pressure in bacterial genomes. *Sci Rep*. 7(1): 1061.

Box GEP. 1976. Science and statistics. *J Am Stat Assoc*. 71(356): 791–799.

Brown L, Elliot T. 1997. Mutations that increase expression of the rpos gene and decrease its dependence on hfq function in *Salmonella typhimurium*. *J Bacteriol*. 179(3): 656–662.

Butler MA, King AA. 2004. Phylogenetic comparative analysis: a modeling approach for adaptive evolution. *Am Nat*. 164(6): 683–695.

Burnham KP, Anderson DR. 2002. Model selection and multimodel inference: a practical information-theoretic approach. New York, NY:Springer.

Dimmic MW, Mindell DP, Goldstein RA. 2000. Modeling evolution at the protein level using an adjustable amino acid fitness model. *Pac Symp Biocomput*. 5:18–29.

Drummond DA, Wilke CO. 2008. Mistranslation-induced protein misfolding as a dominant constraint on coding-sequence evolution. *Cell* 134(2): 341–352.

Drummond DA, Bloom JD, Adami C, Wilke CO, Arnold FH. 2005. Why highly expressed proteins evolve slowly. *Proc Natl Acad Sci U S A*. 102(40): 14338–14343.

Drummond DA, Raval A, Wilke CO. 2006. A single determinant dominates the rate of yeast protein evolution. *Mol Biol Evol*. 23(2): 327–337.

Edwards A. 1984. Likelihood. Cambridge: Cambridge Science Classics – Cambridge University Press.

Endler JA. 1986. Natural selection in the wild. Princeton (NJ): Princeton University Press. p. 16–17. Number 21 in Monographs in Population Biology. Reference for definition of diversifying selection.

Felsenstein J. 1981. Evolutionary trees from DNA-sequences—a maximum-likelihood approach. *J Mol Evol*. 17(6): 368–376.

Felsenstein J. 2001. Taking variation of evolutionary rates between sites into account in inferring phylogenies. *J Mol Evol*. 53(4-5): 447–455.

Fisher RAS. 1930. The genetical theory of natural selection. Oxford: Oxford University Press.

Gilchrist M, Shah P, Zaretzki R. 2009. Measuring and detecting molecular adaptation in codon usage against nonsense errors during protein translation. *Genetics* 183(4): 1493–1505.

Gilchrist MA. 2007. Combining models of protein translation and population genetics to predict protein production rates from codon usage patterns. *Mol Biol Evol*. 24(11): 2362–2373.

Gilchrist MA, Wagner A. 2006. A model of protein translation including codon bias, nonsense errors, and ribosome recycling. *J Theor Biol*. 239(4): 417–434.

Gilchrist MA, Chen W-C, Shah P, Landerer CL, Zaretzki R. 2015. Estimating gene expression and codon-specific translational efficiencies, mutation biases, and selection coefficients from genomic data alone. *Genome Biol Evol*. 7(6): 1559–1579.

Goldman N. 1993. Statistical tests of models of DNA substitution. *J Mol Evol*. 36(2): 182–198.

Goldman N, Yang ZH. 1994. Codon-based model of nucleotide substitution for protein-coding DNA-sequences. *Mol Biol Evol*. 11: 725–736.

Goldman N, Thorne JL, Jones DT. 1996. Using evolutionary trees in protein secondary structure prediction and other comparative sequence analyses. *J Mol Biol*. 263(2): 196–208.

Goldman N, Thorne JL, Jones DT. 1998. Assessing the impact of secondary structure and solvent accessibility on protein evolution. *Genetics* 149(1): 445–458.

Goldsmith M, Tawfik DS. 2009. Potential role of phenotypic mutations in the evolution of protein expression and stability. *Proc Natl Acad Sci*. 106(15): 6197–6202.

Grantham R. 1974. Amino acid difference formula to help explain protein evolution. *Science* 185(4154): 862–864.

Halpern AL, Bruno WJ. 1998. Evolutionary distances for protein-coding sequences: modeling site-specific residue frequencies. *Mol Biol Evol*. 15(7): 910–917.

Higgs PG. 2008. Linking population genetics to phylogenetics. *Banach Center Publ*. 80(1): 145–166.

Hoehna S, Landis MJ, Heath TA, Boussau B, Lartillot N, Moore BR, Huelsenbeck JP, Ronquist F. 2016. Revbayes: Bayesian phylogenetic inference using graphical models and an interactive model-specification language. *Syst Biol*. 65(4): 726.

Hughes AL. 2007. Looking for Darwin in all the wrong places: the misguided quest for positive selection at the nucleotide sequence level. *Heredity* 99(4): 364–373.

Hughes AL, Nei M. 1988. Pattern of nucleotide substitution at major histocompatibility complex class-i loci reveals overdominant selection. *Nature* 335(6186): 167–170.
Hughes AL, Ota T, Nei M. 1990. Positive Darwinian selection promotes charge profile diversity in the antigen-binding cleft of class-i major-histocompatibility-complex molecules. *Mol Biol Evol*. 7:515–524.
Ingram T, Mahler DL. 2013. Surface: detecting convergent evolution from data by fitting Ornstein-Uhlenbeck models with stepwise Akaike information criterion. *Methods Ecol Evol*. 4(5): 416–425.
Iwasa Y. 1988. Free fitness that always increases in evolution. *J Theor Biol*. 135(3): 265–281.
Jhwueng D-C, Snehalata H, O'Meara BC, Liu L. 2014. Investigating the performance of AIC in selecting phylogenetic models. *Stat Appl Genet Mol Biol*. 13(4): 459–475.
Johnson SG. 2012. The NLopt nonlinear-optimization package. http://ab-initio.mit.edu/nlopt.
Kimura M. 1962. on the probability of fixation of mutant genes in a population. *Genetics* 47(6): 713–719.
King ZA, Lloyd CJ, Feist AM, Palsson BO. 2015. Next-generation genome-scale models for metabolic engineering. *Curr Opin Biotechnol*. 35:23–29.
Koshi JM, Goldstein RA. 1997. Mutation matrices and physical-chemical properties: correlations and implications. *Proteins* 27(3): 336–344.
Koshi JM, Goldstein RA. 2000. Analyzing site heterogeneity during protein evolution. In: Russ B Altman, A Keith Dunker, Lawrence Hunker, Kevin Lauderdale. and Teri E Klein, eds. Biocomputing 2001. Singapore:World Scientific. p. 191–202. https://www.worldscientific.com/worldscibooks/10.1142/4604
Koshi JM, Mindell DP, Goldstein RA. 1999. Using physical-chemistry-based substitution models in phylogenetic analyses of HIV-1 subtypes. *Mol Biol Evol*. 16(2): 173–179.
Kosiol C, Holmes I, Goldman N. 2007. An empirical codon model for protein sequence evolution. *Mol Biol Evol*. 24(7): 1464–1479.
Kubatko L, Shah P, Herbei R, Gilchrist MA. 2016. A codon model of nucleotide substitution with selection on synonymous codon usage. *Mol Phylogenet Evol*. 94(Part A):290–297.
Lartillot N, Philippe H. 2004. A Bayesian mixture model for across-site heterogeneities in the amino-acid replacement process. *Mol Biol Evol*. 21(6): 1095–1109.
Lerman JA, Hyduke DR, Latif H, Portnoy VA, Lewis NE, Orth JD, Schrimpe-Rutledge AC, Smith RD, Adkins JN, Zengler K, et al. 2012. In silico method for modelling metabolism and gene product expression at genome scale. *Nat Commun*. 3:929.
Lynch M, Marinov GK. 2015. The bioenergetic costs of a gene. *Proc Natl Acad Sci U S A*. 112(51): 15690–15695.
McCandlish DM, Stoltzfus A. 2014. Modeling evolution using the probability of fixation: history and implications. *Q Rev Biol*. 89(3): 225–252.
McClellan DA, McCracken KG. 2001. Estimating the influence of selection on the variable amino acid sites of the cytochrome b protein functional domains. *Mol Biol Evol*. 18(6): 917–925.
Muse SV, Gaut BS. 1994. A likelihood approach for comparing synonymous and nonsynonymous nucleotide substitution rates, with application to the chloroplast genome. *Mol Biol Evol*. 11(5): 715–724.
Nguyen L-T, Schmidt HA, von Haeseler A, Minh BQ. 2015. Iq-tree: a fast and effective stochastic algorithm for estimating maximum-likelihood phylogenies. *Mol Biol Evol*. 32(1): 268–274.
Nielsen R, Yang ZH. 1998. Likelihood models for detecting positively selected amino acid sites and applications to the hiv-1 envelope gene. *Genetics* 148:929–936.
Nowak MA. 2006. Evolutionary dynamics: exploring the equations of life. Cambridge: Belknap of Harvard University Press.
O'Meara BC, Ane C, Sanderson MJ, Wainwright PC. 2006. Testing for different rates of continuous trait evolution using likelihood. *Evolution* 60(5): 922–933.
Pellmyr O. 2002. Microevolution. In: Pagel M, editor, Encyclopedia of evolution. Vol. 2. Oxford: Oxford University Press. p. 731–732.
Penny D, McComish BJ, Charleston MA, Hendy MD. 2001. Mathematical elegance with biochemical realism: the covarion model of molecular evolution. *J Mol Evol*. 53(6): 711–723.
Pollock DD, Thiltgen G, Goldstein RA. 2012. Amino acid coevolution induces an evolutionary stokes shift. *Proc Natl Acad Sci U S A*. 109(21): E1352–E1359.
Posada D, Buckley TR. 2004. Model selection and model averaging in phylogenetics: advantages of Akaike information criterion and Bayesian approaches over likelihood ratio tests. *Syst Biol*. 53(5): 793–808.
Pouyet F, Bailly-Bechet M, Mouchiroud D, Guéguen L. 2016. Senca: a multilayered codon model to study the origins and dynamics of codon usage. *Genome Biol Evol*. 8(8): 2427–2441.
Rabosky DL, Goldberg EE. 2015. Model inadequacy and mistaken inferences of trait-dependent speciation. *Syst Biol*. 64(2): 340–355.
Robinson DM, Jones DT, Kishino H, Goldman N, Thorne J. L 2003. Protein evolution with dependence among codons due to tertiary structure. *Mol Biol Evol*. 20(10): 1692–1704.
Rodrigue N, Lartillot N. 2014. Site-heterogeneous mutation-selection models within the phylobayes-mpi package. *Bioinformatics* 30(7): 1020–1021.
Rodrigue N, Lartillot N, Bryant D, Philippe H. 2005. Site interdependence attributed to tertiary structure in amino acid sequence evolution. *Gene* 347(2): 207–217.
Rokas A, Williams BL, King N, Carroll SB. 2003. Genome-scale approaches to resolving incongruence in molecular phylogenies. *Nature* 425(6960): 798–804.
Rowan T. 1990. Functional stability analysis of numerical algorithms [PhD thesis]. [Austin (TX)]: University of Texas.
Salichos L, Rokas A. 2013. Inferring ancient divergences requires genes with strong phylogenetic signals. *Nature* 497(7449): 327–331.
Schneider E, Moore M, Castro KG. 2005. Epidemiology of tuberculosis in the united states. *Clin Chest Med*. 26(2): 183.
Sella G, Hirsh AE. 2005. The application of statistical physics to evolutionary biology. *Proc Natl Acad Sci U S A*. 102(27): 9541–9546.
Shah P, Gilchrist MA. 2011. Explaining complex codon usage patterns with selection for translational efficiency, mutation bias, and genetic drift. *Proc Natl Acad Sci U S A*. 108(25): 10231–10236.
Shah P, McCandlish DM, Plotkin JB. 2015. Contingency and entrenchment in protein evolution under purifying selection. *Proc Natl Acad Sci*. 112(25): E3226–E3235.
Stamatakis A. 2006. RAxML-VI-HPC: maximum likelihood-based phylogenetic analyses with thousands of taxa and mixed models. *Bioinformatics* 22(21): 2688–2690.
Sullivan J, Joyce P. 2005. Model selection in phylogenetics. *Annu Rev Ecol Evol Syst*. 36(1): 445–466.
Tamuri AU, dos Reis M, Hay AJ, Goldstein RA. 2009. Identifying changes in selective constraints: host shifts in influenza. *PLoS Comput Biol*. 5:e1000564.
Tamuri AU, dos Reis M, Goldstein RA. 2012. Estimating the distribution of selection coefficients from phylogenetic data using sitewise mutation-selection models. *Genetics* 190(3): 1101–1115.
Tamuri AU, Goldman N, dos Reis M. 2014. A penalized-likelihood method to estimate the distribution of selection coefficients from phylogenetic data. *Genetics* 197(1): 257–271.
Thiele I, Fleming RMT, Que R, Bordbar A, Diep D, Palsson BO. 2012. Multiscale modeling of metabolism and macromolecular synthesis in e. coli and its application to the evolution of codon usage. *PLoS ONE* 7(9): 1–18.
Thorne JL, Goldman N, Jones DT. 1996. Combining protein evolution and secondary structure. *Mol Biol Evol*. 13(5): 666–673.
Thorne JL, Lartillot N, Rodrigue N, Choi SC. 2012. Codon models as a vehicle for reconciling population genetics with inter-specific sequence data. In: Gina M. Cannarozzi and Adrian Schneider, eds. Codon evolution: mechanisms and models. p. 97–110, Oxford: Oxford University Press.
Tuffley C, Steel M. 1998. Modeling the covarion hypothesis of nucleotide substitution. *Math Biosci*. 147(1): 63–91.

Wagner A. 2005. Energy constraints on the evolution of gene expression. *Mol Biol Evol*. 22(6): 1365–1374.
Whelan S. 2008. Spatial and temporal heterogeneity in nucleotide sequence evolution. *Mol Biol Evol*. 25(8): 1683–1694.
Whelan S, Goldman N. 2004. Estimating the frequency of events that cause multiple-nucleotide changes. *Genetics* 167(4): 2027–2043.
Woolley S, Johnson J, Smith MJ, Crandall KA, McClellan DA. 2003. Treesaap: selection on amino acid properties using phylogenetic trees. *Bioinformatics* 19(5): 671–672.
Wright S. 1969. Evolution and the genetics of populations. Vol. 2. The theory of gene frequencies. Chicago:University of Chicago Press.
Xia XH, Li WH. 1998. What amino acid properties affect protein evolution? *J Mol Evol*. 47(5): 557–564.
Yang Z. 2014. Molecular evolution: a statistical approach. New York: Oxford University Press.
Yang Z, Nielsen R, Hasegawa M. 1998. Models of amino acid substitution and applications to mitochondrial protein evolution. *Mol Biol Evol*. 15(12): 1600–1611.
Yang ZH. 1994. Maximum-likelihood phylogenetic estimation from DNA-sequences with variable rates over sites—approximate methods. *J Mol Evol*. 39(3): 306–314.
Yang ZH. 2007. Paml 4: phylogenetic analysis by maximum likelihood. *Mol Biol Evol*. 24(8): 1586–1591.
Yang ZH, Nielsen R. 1998. Synonymous and nonsynonymous rate variation in nuclear genes of mammals. *J Mol Evol*. 46(4): 409–418.
Yang ZH, Nielsen R. 2008. Mutation-selection models of codon substitution and their use to estimate selective strengths on codon usage. *Mol Biol Evol*. 25(3): 568–579.
Zanger U, Schwab M. 2013. Cytochrome p450 enzymes in drug metabolism: regulation of gene expression, enzyme activities, and impact of genetic variation. *Pharmacol Ther*. 138:103–141.