

Modeling site-specific amino-acid preferences deepens phylogenetic estimates of viral sequence divergence

Sarah K. Hilton^{1,2} and Jesse D. Bloom^{1,2,3,*†}

¹Basic Sciences and Computational Biology Program, Fred Hutchinson Cancer Research Center, ²Department of Genome Sciences, University of Washington, USA and ³Howard Hughes Medical Institute, Seattle, WA, USA

*Corresponding author: E-mail: jlbloom@fredhutch.org

†<http://orcid.org/0000-0003-1267-3408>

Abstract

Molecular phylogenetics is often used to estimate the time since the divergence of modern gene sequences. For highly diverged sequences, such phylogenetic techniques sometimes estimate surprisingly recent divergence times. In the case of viruses, independent evidence indicates that the estimates of deep divergence times from molecular phylogenetics are sometimes too recent. This discrepancy is caused in part by inadequate models of purifying selection leading to branch-length underestimation. Here we examine the effect on branch-length estimation of using models that incorporate experimental measurements of purifying selection. We find that models informed by experimentally measured site-specific amino-acid preferences estimate longer deep branches on phylogenies of influenza virus hemagglutinin. This lengthening of branches is due to more realistic stationary states of the models, and is mostly independent of the branch-length extension from modeling site-to-site variation in amino-acid substitution rate. The branch-length extension from experimentally informed site-specific models is similar to that achieved by other approaches that allow the stationary state to vary across sites. However, the improvements from all of these site-specific but time homogeneous and site independent models are limited by the fact that a protein's amino-acid preferences gradually shift as it evolves. Overall, our work underscores the importance of modeling site-specific amino-acid preferences when estimating deep divergence times—but also shows the inherent limitations of approaches that fail to account for how these preferences shift over time.

Key words: phylogenetics; deep mutational scanning; site-specific substitution model; influenza virus; hemagglutinin; ExpCM.

1. Introduction

Molecular phylogenetics is commonly used to estimate the historical timing of evolutionary events (Yang and Rannala 2012). This is done by estimating branch lengths based on the inferred number of substitutions, and then converting these branch lengths into units of time under the assumption of a molecular clock (Zuckermandl and Pauling 1965; Drummond et al. 2006). However, phylogenetic estimates of the divergence times of many viral lineages are clearly too recent (Duchêne, Holmes, and Ho 2014; Ho et al. 2015; Aiewsakun and Katzourakis 2016). For example, the integration of filoviruses into their host genomes indicate that Ebola

and Marburg virus diverged from their common ancestor 7 to 12 million years ago—but the estimate of this divergence time based on phylogenetic analyses of the viral sequences is only ~10,000 years ago (Carroll et al. 2013; Taylor et al. 2014). Similarly, the phylogenetic estimate of when major simian immunodeficiency virus groups diverged is almost 100 times more recent than the estimate based on the geographic isolation of their host species (Wertheim and Worobey 2009; Worobey et al. 2010). These examples, along with other similar discrepancies with measles virus (Furuse, Suzuki, and Oshitani 2010), coronavirus (Wertheim et al. 2013), and hepatitis B virus (Fares and Holmes 2002; Holmes

2003), indicate that phylogenetic methods have a systematic bias toward underestimation of deep branches.

This underestimation occurs in part because phylogenetic models do a poor job of describing the real natural selection on protein-coding genes. These genes evolve under purifying selection to maintain the structure and function of the proteins they encode. In general, these constraints are highly idiosyncratic among sites (Echave, Spielman, and Wilke 2016). However, most phylogenetic models try to account for these constraints using relatively simple approaches such as allowing the rate of substitution to vary across sites according to some statistical distribution (Yang 1994; Yang et al. 2000). These models of purifying selection are usually inadequate (Duchêne et al. 2015a; Duchêne, Di Giallonardo, and Holmes 2015b), potentially causing branch lengths to be severely underestimated (Halpern and Bruno 1998; Wertheim and Kosakovsky Pond 2011).

More recent work has used mutation-selection models to better account for purifying selection (Halpern and Bruno 1998; Yang and Nielsen 2008; Rodrigue, Philippe, and Lartillot 2010; Tamuri, dos Reis, and Goldstein 2012; McCandlish and Stoltzfus 2014). These models explicitly incorporate the fact that different protein sites prefer different amino acids, and so can improve phylogenetic estimates when there are deep branches (Philippe and Laurent 1998; Lartillot, Brinkmann, and Philippe 2007; Le, Lartillot, and Gascuel 2008; Quang, Gascuel, and Lartillot 2008; Wang et al. 2008; Susko, Lincker, and Roger 2018). However, these approaches require inferring the site-specific purifying selection from natural sequence data.

Even more recently, it has become possible to directly measure purifying selection on proteins using deep mutational scanning. This high-throughput approach involves experimentally measuring how each amino-acid mutation affects protein function in the lab (Fowler and Fields 2014). The resulting experimental measurements of which amino acids are preferred at each protein site can be used to inform phylogenetic substitution models (Bloom 2014a). These experimentally informed codon models (ExpCMs) generally exhibit much better phylogenetic fit than standard substitution models (Doud, Ashenberg, and Bloom 2015; Hilton, Doud, and Bloom 2017; Haddox et al. 2018; Lee et al. 2018).

Here we examine how ExpCMs and other models of purifying selection estimate branch lengths on a phylogenetic tree of influenza virus hemagglutinin (HA). We find that ExpCMs estimate longer deep branches, and show that this extension of branch length is mostly independent and additive with that achieved by the more conventional approach of modeling rate variation. We also show that ExpCMs estimate similar branch lengths to a mutation-selection model that infers the amino-acid preferences from the natural sequence data rather than using values obtained in experiments. However, all of these mutation-selection models are limited by their failure to account for another feature of purifying selection: the fact that a site's amino-acid preferences shift over time due to epistasis. Therefore, truly accurate analyses of deep phylogenies need to account for the fact that amino-acid preferences vary across time as well as across sites.

2. Results

2.1 Different ways substitution models account for purifying selection

Here we consider how purifying selection is handled by codon models, which are the most sophisticated of the three classes (nucleotide, codon, and amino acid) of phylogenetic

substitution models in widespread use for protein-coding genes (Arenas 2015). Standard codon models distinguish between two types of substitutions: synonymous and nonsynonymous. The relative rate of these substitutions is referred to as dN/dS or ω . In their simplest form, codon substitution models fit a single ω that represents the gene-wide average fixation rate of nonsynonymous mutations relative to synonymous ones. Here we will use such substitution models in the form proposed by Goldman and Yang (1994). When these models have a single gene-wide ω they are classified as M0 by Yang et al. (2000). We will refer to M0 Goldman-Yang models simply as GY94 models (Equation 1). The gene-wide ω is usually < 1 (Murrell et al. 2015), and crudely represents the fact that many amino-acid substitutions are under purifying selection.

A single gene-wide ω ignores the fact that purifying selection is heterogeneous across sites. The most common strategy to ameliorate this defect is to allow ω to vary among sites according to some statistical distribution (Yang 1994; Yang et al. 2000). For instance, in the M5 variant of the GY94 model (Yang et al. 2000), ω follows a gamma distribution as shown in Fig. 1a. We will denote this model as GY94+ $\Gamma\omega$. A GY94+ $\Gamma\omega$ captures the fact that the rate of nonsynonymous substitution can vary across sites. However, these models do not capture the fact that the same amino-acid mutation can have very different effects at different sites.

Mutation-selection models account for the fact that purifying selection depends idiosyncratically on the specific amino-acid mutation at each site (Halpern and Bruno 1998; Yang and Nielsen 2008; Rodrigue, Philippe, and Lartillot 2010; Tamuri, dos Reis, and Goldstein 2012; McCandlish and Stoltzfus 2014). Here we will consider mutation-selection models where the site-specific selection is assumed to act solely at the protein level (different codons for the same amino acid are treated as selectively equivalent). Such models explicitly define a different set of amino-acid preferences at each site in the protein. This more mechanistic formulation results in a site-specific stationary state (Fig. 1b). These models capture the site-to-site variation in amino-acid composition that is an obvious feature of real proteins, and usually better describe actual evolution than models with only rate variation as assessed by Bayesian or maximum-likelihood criteria (Lartillot and Philippe 2004; Le, Lartillot, and Gascuel 2008; Quang, Gascuel, and Lartillot 2008; Wang et al. 2008; Rodrigue, Philippe, and Lartillot 2010; Bloom 2014a,b; Hilton, Doud, and Bloom 2017).

However, the increased realism of mutation-selection models comes at the cost of an increased number of parameters. Codon substitution models with uniform stationary states have only a modest number of parameters that must be fit from the phylogenetic data. For instance, a GY94+ $\Gamma\omega$ model with the commonly used F3X4 stationary state has 12 parameters: two describing the shape of the gamma distribution over ω , a transition-transversion rate, and nine parameters describing the nucleotide composition of the stationary state. However, mutation-selection models must additionally specify 19 parameters defining the amino-acid preferences for each site (there are 20 amino acids whose preferences are constrained to sum to one). This corresponds to $19 \times L$ parameters for a protein of length L , or 9,500 parameters for a 500-residue protein. It is challenging to obtain values for these amino-acid preference parameters in a maximum-likelihood framework without overfitting the data (Rodrigue 2013). Here we will primarily use ExpCMs, which define the site-specific amino-acid preference parameters a priori from deep mutational scanning experiments so that they do not need to be fit from phylogenetic data (see Section 4 and Bloom 2014a,




Figure 1. Different ways codon models account for purifying selection. (a) The dN/dS parameter, ω , can be defined as one gene-wide average (orange triangle) or allowed to vary according to some statistical distribution (blue line). For computational tractability, the distribution is discretized into K bins and ω takes on the mean of each bin (blue circles) (Yang 1994; Yang et al. 2000). A gamma distribution (denoted by Γ) with $K=4$ bins is shown here. (b) A substitution model's stationary state defines the expected sequence composition after a very long evolutionary time. Most substitution models have stationary states that are uniform across sites. However, substitution models can have site-specific stationary states. In the logo plots, each column is a site in the protein and the height of each letter is the frequency of that amino acid at stationary state. (c) Substitution models can incorporate neither, one, or both of these features. Here we will use substitution models from the Goldman-Yang (GY94; Goldman and Yang 1994; Yang et al. 2000) and ExpCM (Hilton, Doud, and Bloom 2017) families with and without gamma-distributed ω to represent all possible combinations.

2017; Hilton, Doud, and Bloom 2017). Because the amino-acid preference parameters in an ExpCM are obtained from experiments, the number of ExpCM free parameters is similar to a non-site-specific substitution model. An alternative strategy to account for site-specific amino-acid preferences is to formally consider them as random effects across sites, rather than parameters, and infer them using Bayesian methods (Lartillot and Philippe 2004; Rodrigue and Lartillot 2014). This strategy is discussed in Section 2.6.

Importantly, these two strategies for modeling purifying selection are not mutually exclusive. Mutation-selection models such as an ExpCM can still incorporate an ω parameter, which now represents the relative rate of nonsynonymous to synonymous substitution *after* accounting for the constraints due to the site-specific amino-acid preferences (Bloom 2017; Rodrigue and Lartillot 2017). This ω parameter for an ExpCM can be drawn from a statistical distribution (e.g., a gamma distribution) just like for GY94-style models (Rodrigue and Lartillot 2014; Haddox et al. 2018). We will denote such models as ExpCM+ $\Gamma\omega$. Figure 1c shows the full spectrum of models that incorporate all combinations of gamma-distributed ω and site-specific stationary states.

2.2 Effect of stationary state and rate variation on branch-length estimation

Given a single branch, a substitution model transforms sequence divergence into branch length. Under a molecular-clock assumption, this branch length is proportional to time. The transformation from sequence divergence to branch length is trivial when the sequence identity is high. For instance, when there has only been one substitution, then the sequence identity will simply be $(L-1)/L$ for a gene of L sites, and even a simple exponential model (Zuckermandl and Pauling 1965) will correctly infer the short branch length of $1/L$ substitutions per site. However, as substitutions accumulate it becomes progressively more likely for multiple changes to occur at the same site. In this regime, the accuracy of the substitution model becomes critical for transforming sequence divergence into branch length. Any time-homogenous substitution model predicts that after a very large number of substitutions, two related sequences will approach some asymptotic amino-acid sequence identity. For instance, if all 20 amino acids are equally

likely in the stationary state, then this asymptotic sequence identity will be $\frac{1}{20} = 0.05$. If the substitution model underestimates the asymptotic sequence identity then it will also underestimate long-branch lengths, since it will predict that sequences that have evolved for a very long time should be more diverged than is actually the case.

Figure 2 shows how different substitution models predict amino-acid sequence identity to decrease as a function of branch length using model parameters fit to a phylogeny of H1 influenza HA genes. The GY94 model predicts the same behavior for all sites, since it does not have any site-specific parameters, with an asymptotic sequence identity of 0.062. While this predicted sequence identity is higher than $\frac{1}{20} = 0.05$ due to redundant codon and nucleotide biases favoring certain amino acids, it is much lower than the pairwise identity of even the most diverged HAs in nature. While it is of course possible that the identity of HAs in nature would become even lower given more time, it seems biochemically improbable that it would ever become as low as 0.062. The reason is that like many proteins HA has a highly conserved structure and function that imposes constraints that cause many sites to sample only a small subset of the 20 amino acids among all known HA homologs (Nobusawa et al. 1991).

Accounting for site-to-site dN/dS rate variation in GY94 models affects the rate at which the asymptotic sequence identity is approached, but not the actual value of this asymptote. For instance, Fig. 2 shows that the GY94+ $\Gamma\omega$ model takes longer to reach the asymptote than GY94, but that the asymptote is identical for both models. This fact holds true even if we use experimental measurements of HA's site-specific amino-acid preferences (Doud and Bloom 2016) to calculate a different ω_r value for each site using the method of Spielman and Wilke (2015b) (see Equation 7). Specifically, this GY94+ ω_r model predicts that different sites will approach the asymptote at different rates, but the asymptote is always the same (Fig. 2). The invariance of the asymptotic sequence identity under different schemes for modeling ω is a fundamental feature of the mathematics of this type of reversible substitution model. These models are reversible stochastic matrices, which can be decomposed into stationary states and symmetric exchangeability matrices (Nielsen 2006). The stationary state is invariant with respect to multiplication of the symmetric exchangeability matrix by any nonzero number. Different schemes for modeling

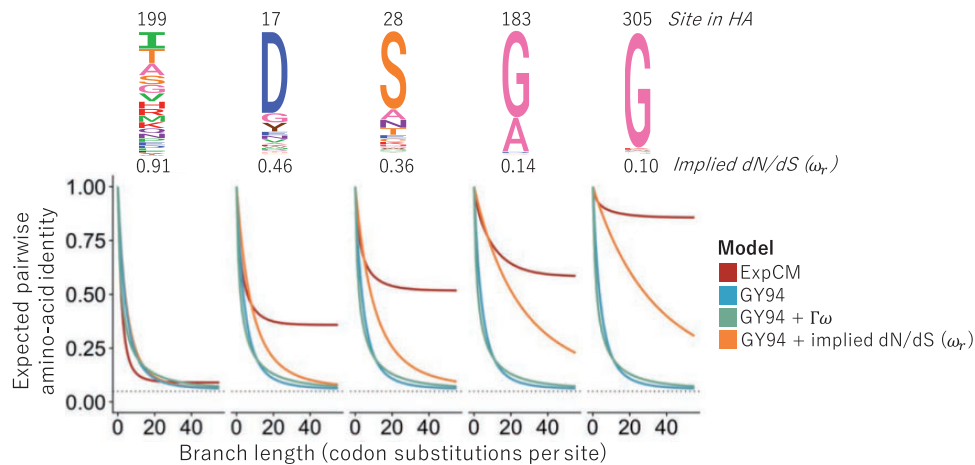


Figure 2. Effect of stationary state and $\Gamma\omega$ rate variation on predicted asymptotic sequence divergence. The logo plots at top show the amino-acid preferences for some sites in an H1 influenza HA protein as experimentally measured by Doud and Bloom (2016). The graphs show the expected amino-acid identity at that site for two sequences separated by a branch of the indicated length (Equation 9). For the GY94 model, the graphs are identical for all sites since this model does not have site-specific parameters; the same is true for GY94+ $\Gamma\omega$. The graphs do differ among sites if we calculate a different ω_r for each site r in the GY94 model using the amino-acid preferences (Equation 7; Spielman and Wilke 2015b). However, all GY94 models, including the one with site-specific ω_r values, approach the same asymptote since they all have the same stationary state. The ExpCM has different asymptotes for different sites since it accounts for how amino-acid preferences lead to site-specific stationary states.

ω only multiply elements of the symmetric exchangeability matrix. Therefore, no matter how ‘well’ a model accounts for site-to-site variation in ω , it will always have the same stationary state as a simple GY94 model.

However, mutation-selection models such as ExpCMs have site-specific stationary states. They predict that different sites will have different asymptotic sequence identities (Fig. 2)—a prediction that accords with the empirical observation that some sites are much more variable than others in alignments of highly diverged sequences. For instance, Fig. 2 shows that at sites such as 183 and 305 in the H1 HA, an ExpCM but not a GY94-style model predicts that the identity will always be relatively high. When sites with highly constrained amino-acid preferences such as these are common, an ExpCM can estimate a long-branch length at modest sequence identities that a GY94 model might attribute to a shorter branch.

2.3 Simulations demonstrate how failure to model site-specific amino-acid preferences leads to branch-length underestimation

To directly demonstrate the effect of stationary state and $\Gamma\omega$ rate variation on branch-length estimation, we tested the ability of a variety of models to accurately infer branch lengths on simulated data (Fig. 3). Specifically, we simulated alignments of sequences along the HA phylogenetic tree using an ExpCM parameterized by the amino-acid preferences of H1 HA as experimentally measured by deep mutational scanning (Doud and Bloom 2016). We then estimated the branch lengths from the simulated sequences using all the substitution models in Fig. 1c, and compared these estimates to the actual branch lengths used in the simulations. Note that these simulations closely parallel those performed by Halpern and Bruno (1998) and Wertheim and Kosakovsky Pond (2011).

The models with a uniform stationary state underestimated the lengths of long branches on the phylogenetic tree of the simulated sequences (Fig. 3). The GY94 model estimated branch lengths that are ~60 per cent of the true values for the longest branches. Accounting for site-to-site variation in ω did not fix the fundamental problem: the GY94+ $\Gamma\omega$ did slightly better, but

still substantially underestimated the longest branches. However, there was no systematic underestimation of long branches by the ExpCM and ExpCM+ $\Gamma\omega$ models. The improved performance of the ExpCMs is due to their modeling of the site-specific amino-acid preferences: if we parameterize ExpCMs by amino-acid preferences that have been averaged across HA sites (and so are no longer site-specific), then they perform no better than GY94 models (Fig. 3). Therefore, models with uniform stationary states underestimate the length of long branches in phylogenies of sequences that have evolved under strong site-specific amino-acid preferences.

2.4 Experimentally informed site-specific models estimate longer branches on real data

The foregoing section shows the superiority of ExpCMs to GY94 models for estimating long branches on phylogenies simulated with ExpCMs. But how do these models perform on real data? Real genes do evolve under functional constraint, but these constraints are almost certainly more complex than what is modeled by an ExpCM. However, if ExpCMs do a substantially better job than GY94 models of capturing the true constraints, then we might still expect them to estimate more accurate branch lengths.

To test the models on real data, we used actual sequences of influenza HA. The topology of HA phylogenetic trees makes these sequences an interesting test case for branch-length estimation. HA consists of a number of different subtypes. Sequences within a subtype have >68 per cent amino-acid identity, but sequences in different subtypes have as little as 38 per cent identity. However, HA proteins from all subtypes have a highly conserved structure that performs a highly conserved function (Ha et al. 2002; Russell et al. 2004). We used RAxML (Stamatakis 2006) with a nucleotide substitution model (GTRCAT) to infer a phylogenetic tree for 92 HA sequences drawn from 15 of the 18 subtypes (we excluded bat influenza and one other rare subtype). For the rest of this article, we fix the tree topology to this RAxML-inferred tree. Although the nucleotide model used with RAxML to infer this tree topology is probably less accurate than codon models, the modular subtype structure of the HA phylogeny (the tree is clearly divided into

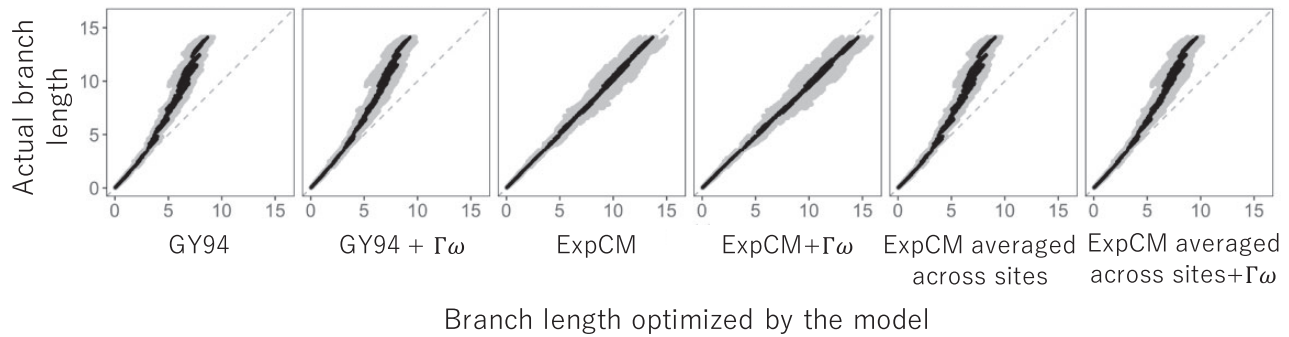


Figure 3. Branch lengths inferred on data simulated under a model with site-specific amino-acid preferences. We simulated alignments along a HA phylogenetic tree using an ExpCM parameterized by the actual site-specific amino-acid preferences for an H1 HA (Doud and Bloom 2016). We then inferred the branch lengths of this tree from the simulated alignments. The inferred branch lengths for various models are plotted on the x-axis, and the actual branch lengths used in the simulations are on the y-axis. We performed 10 simulations and inferences, and gray points show each inferred branch length from each simulation, and black points show the average of each branch length across simulations. The gray dashed line at $y = x$ represents what would be seen if the inferred branch lengths exactly matched those used in the simulations.

widely separated clades) means that most of the phylogenetic uncertainty lies in the length of the long branches separating the subtypes rather than in the tree topology itself. We note that other genes may have phylogenetic structures that are more prone to topological uncertainty. In such cases, the accuracy of the substitution model may be important for avoiding topological errors such as long-branch attraction (Felsenstein 1978; Lartillot, Brinkmann, and Philippe 2007).

Deep mutational scanning has been used to measure the amino-acid preferences of all sites in two different HAs. One scan measured the preferences of an H1 HA (Doud and Bloom 2016) and the other measured the preferences of an H3 HA (Lee et al. 2018). The amino-acid preferences measured for these two HAs are shown in Supplementary Figs S1 and S2. The H1 and H3 HAs have only ~42 per cent amino-acid identity. As described in Lee et al. (2018), the amino-acid preferences clearly differ between the H1 and H3 HA at a substantial number of sites (these differences are apparent in a simple visual comparison of Supplementary Figs S1 and S2; see site 33 as an example). Therefore, we also created a third set of amino-acid preferences by averaging the measurements for the H1 and H3 HAs, under the conjecture that these averaged preferences might better describe the ‘average’ constraint on sites across the full HA tree (Supplementary Fig S3). These three sets of HA amino-acid preferences define three different ExpCMs.

We fit the GY94 model and each of the three ExpCMs to the fixed HA tree topology estimated using RA_{xML} , and also tested a version of each model with Γ_{ω} rate variation. Table 1 shows that all ExpCMs fit the actual data much better than the GY94 models. The best fit was for the ExpCM informed by the average of the H1 and H3 deep mutational scans. For all models, incorporating Γ_{ω} rate variation improved the fit, although even ExpCMs without Γ_{ω} greatly outperformed the GY94+ Γ_{ω} model (Table 1). As mentioned in the previous section, ω is generally < 1 when a single value is fit to all sites in a gene (Murrell et al. 2015), and this is the case for all the models we tested (Table 1). However, the ExpCMs always fit an ω greater than the GY94 model, suggesting that the site-specific amino-acid preferences capture some of the purifying selection that the GY94 models can represent only via a small ω . Among the models with Γ_{ω} , the GY94+ Γ_{ω} model fits all four ω categories to values $\ll 1$, but the ExpCM+ Γ_{ω} models fit one of the ω categories to a value > 1 . This increase in ω values makes sense given the different interpretation of ω for each family of models. The ExpCM ω is the

relative rate of fixation of nonsynonymous to synonymous mutations *after* accounting for the functional constraints described by the amino-acid preferences. This more realistic null model gives ExpCMs enhanced power to detect diversifying selection for amino-acid change (Bloom 2017; Rodrigue and Lartillot 2017), which is known to occur at some sites in HA due to immune selection (Bedford et al. 2014).

Importantly, models that account for purifying selection via either Γ_{ω} rate variation or site-specific amino-acid preferences do not just exhibit better fit—they also estimate longer branches on the HA tree. Figure 4 shows the branch lengths optimized by each model on a common scale. The tree’s deepest branches are shortest when they are optimized by the GY94 model, which lacks both Γ_{ω} and site-specific amino-acid preferences. Adding either Γ_{ω} rate variation or site-specific amino-acid preferences increases the length of the deep branches. Specifically, the tree’s diameter (the distance between the two most divergent tips) for the GY94+ Γ_{ω} model is 159 per cent of the GY94 model tree diameter (Supplementary Table S1). The tree diameter is 122 per cent and 135 per cent of the GY94 model tree diameter for ExpCMs informed by H1 or H3 amino-acid preferences, respectively, and 160 per cent of the GY94 model for the ExpCM informed by the average of the H1 and H3 preferences (Supplementary Table S1).

The deepening of branch lengths that results from the Γ_{ω} and site-specific amino-acid preference approaches to modeling purifying selection are largely independent. This can be seen by examining the ExpCM+ Γ_{ω} models, which combine Γ_{ω} rate variation with site-specific amino-acid preferences. As shown in Fig. 4, these ExpCM+ Γ_{ω} models estimate longer branches than models with just Γ_{ω} rate variation (GY94+ Γ_{ω}) or just site-specific amino-acid preferences (ExpCMs). The near independence of these effects is quantified in Supplementary Table S1, which shows that 76 per cent of the tree diameter extension of ExpCM(H1+H3 avg)+ Γ_{ω} versus can be explained by simply adding the extension from incorporating Γ_{ω} (GY94+ Γ_{ω} versus GY94) to the extension from incorporating site-specific amino-acid preferences (compare ExpCM(H1+H3 avg) to GY94).

However, while adding Γ_{ω} rate variation increases the length of deep branches in a roughly uniform fashion across the tree, the branch lengthening from adding site-specific amino-acid preferences is not uniform across the tree (Figs 4 and 5). Instead, the increase in branch length is most pronounced on

Table 1. Fitting of substitution models to the HA phylogenetic tree.

Model	Δ AIC	Log likelihood	ω	Stringency parameter (β)
ExpCM (H1+H3 avg) + $\Gamma\omega$	0	-51083	0.19, 0.50, 0.91, 1.86	1.69
ExpCM (H1+H3 avg)	1,063	-51,616	0.14	1.77
ExpCM (H1) + $\Gamma\omega$	1,321	-51,744	0.12, 0.42, 0.89, 2.13	1.11
ExpCM (H3) + $\Gamma\omega$	1,777	-51,972	0.10, 0.36, 0.76, 1.84	1.28
ExpCM (H1)	2,670	-52,419	0.12	1.21
ExpCM (H3)	3,377	-52,773	0.12	1.43
GY94 + $\Gamma\omega$	4,817	-53,487	0.00, 0.03, 0.08, 0.24	-
GY94	7,892	-55,025	0.07	-

Notes: All ExpCMs describe the evolution of HA better than the GY94 models, as evaluated by the Akaike information criteria (Δ AIC, Posada and Buckley 2004). The models fit here are the same ones in Fig. 4. The ω value for each of the K = 4 bins is shown for the models with $\Gamma\omega$ rate variation. All ExpCMs fit a stringency parameter > 1 .

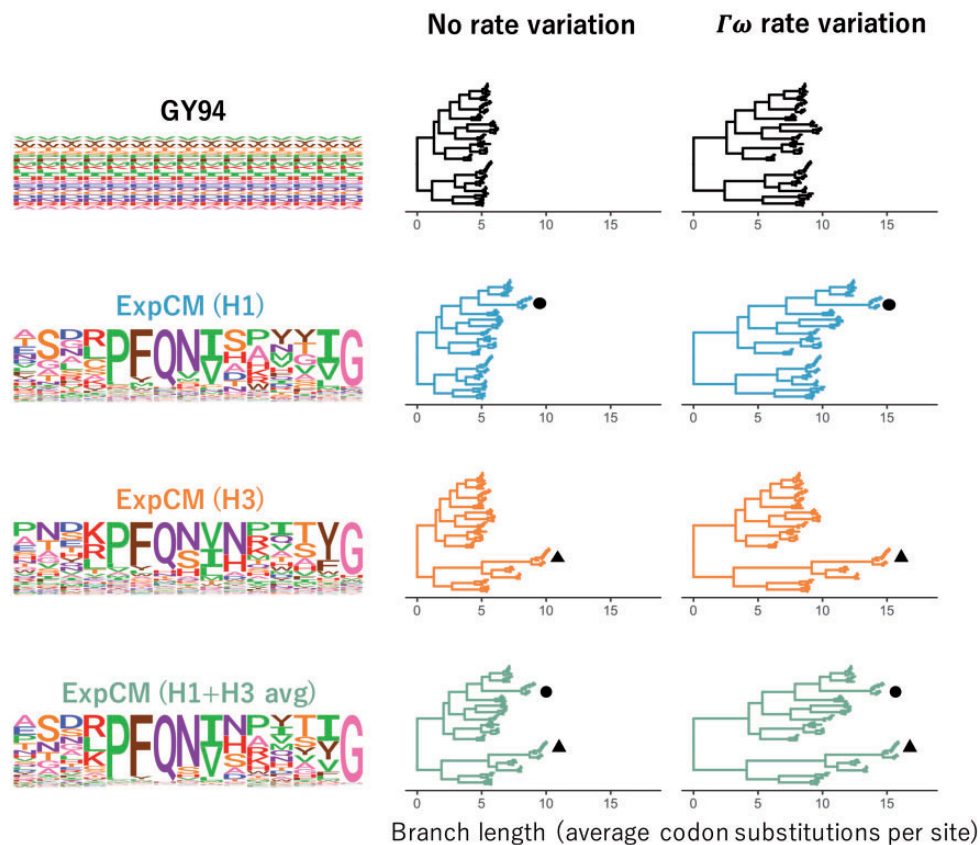


Figure 4. Effect of site-specific amino-acid preferences and $\Gamma\omega$ rate variation on HA branch length estimation. The branch lengths of the HA tree were optimized using the indicated ExpCM or GY94 model. The amino-acid preferences defining the model (ExpCM) or implied by the model (GY94) are shown as logo plots for 15 sites in HA; the full set of experimentally measured amino-acid preferences are in Supplementary Figs 1–3. The ExpCMs use amino-acid preferences measured in deep mutational scanning of an H1 HA (Doud and Bloom 2016), an H3 HA (Lee et al. 2018), or the average of the measurements for these two HAs. Circle denotes the H1 clade and triangle denotes the H3 clade. The root of each tree is placed where it would fall if the tree was midpoint rooted using the branch lengths inferred by RAxML using the GTRCAT model. This figure enables qualitative visualization of the trees; for a quantitative comparison of branch lengths optimized by different models, see Fig. 5.

branches leading to the HA sequence that was used in the deep mutational scanning experiment that informed the ExpCM. For instance, the ExpCM informed by the H1 data most dramatically lengthens branches near the H1 clade of the tree, while the ExpCM informed by the H3 data has the largest effect on branches near the H3 clade. The ExpCM informed by the average of the H1 and H3 data has a more uniform effect across the tree, but still most strongly extends branches leading to either the H1 or H3 clade. Therefore, Figs 4 and 5 show that ExpCMs estimate longer branches, but that the effect is shaped by the set of amino-acid preferences used to inform the model.

2.5 Shifting amino-acid preferences limit the benefits of models with site-specific stationary states for estimating long-branch lengths

The fact that an ExpCM leads to the most profound increase in branch length leading to the sequence used in the experiment can be rationalized in terms of existing knowledge about epistasis during protein evolution. Each ExpCM is informed by a single set of experimentally measured amino-acid preferences. But in reality, the effect of a mutation at one site in a protein can depend on the amino-acid identities of other sites in the protein

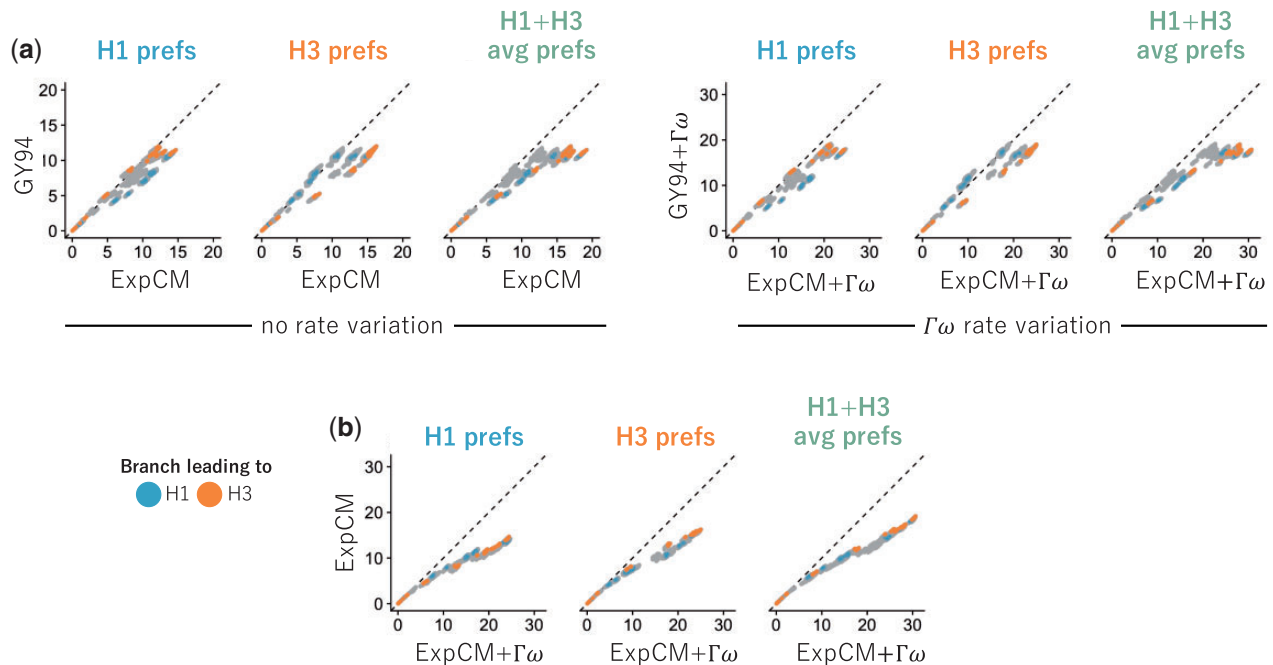


Figure 5. Modeling site-specific amino-acid preferences using deep mutational scanning experiments extends branch lengths, especially for branches leading to the HA used in the experiment. The points indicate the total length of branches separating all pairs of tips on the HA phylogenetic tree when the tree is optimized under the indicated model. Blue and orange denote branches that lead to the H1 and H3 HAs used in the deep mutational scanning. The amino-acid preference set defining the ExpCM is labeled above each plot. (a) ExpCMs defined by amino-acid preferences from any of the deep mutational scanning experiments estimate generally longer branches than the GY94 model, with the increase particularly pronounced for branches leading to the HA used in the experiment. (b) The addition of $\Gamma\omega$ rate variation further extends branch lengths, without any apparent bias toward the HAs used in the experiment. Note that this figure shows the same data as Fig. 4 in a different form.

(Ortlund et al. 2007; Gong, Suchard, and Bloom 2013; Harms and Thornton 2014; Tufts et al. 2014; Starr et al. 2018). This epistasis can lead to shifts in a protein's amino-acid preferences over evolutionary time (Pollock, Thiltgen, and Goldstein 2012; Bazykin 2015; Doud, Ashenberg, and Bloom 2015; Shah, McCandlish, and Plotkin 2015; Haddock et al. 2018). Because the deep mutational scanning experiments that inform our ExpCMs were each performed in the context of a single HA genetic background, their measurements do not account for the accumulation of epistatic shifts in amino-acid preferences as HA evolves. Therefore, an ExpCM is expected to most accurately describe the evolution of sequences closely related to the one used in the experiment.

We can observe how shifting amino-acid preferences degrade the accuracy of an ExpCM by fitting the model to trees containing increasingly diverged sequences. For both H1 and H3 HAs, we created three phylogenetic trees (Supplementary Fig. S5): a 'low' divergence tree that contains sequences with ≥ 59 per cent amino-acid identity to the HA used in the experiment, an 'intermediate' divergence tree that contains sequences with ≥ 46 per cent amino-acid identity to the HA in the experiment, and a 'high' divergence tree that contains all HAs (which have as little as 38 per cent identity to the HA in the experiment). Figure 6 shows the subtrees containing each of these sets of HA sequences. For each subtree, we examined the congruence between site-specific natural selection and the amino-acid preferences measured in the deep mutational scanning experiment using the ExpCM stringency parameter β (Bloom 2014b; Hilton, Doud, and Bloom 2017). Values of β that are >1 indicate that natural selection prefers the same amino acids as the experiments but with a greater stringency, suggesting strong

congruence between natural selection and the experimental preferences. In contrast, values of β that are <1 flatten the preferences, suggesting that they provide a relatively poor description of natural selection on the protein.

The value of β decreases as the divergence from the sequence used in the deep mutational scan increases Fig. 6. This inverse relationship between β and overall divergence is seen for the ExpCMs informed by both the H1 and H3 experiments. As β decreases, the preferences 'flatten' and so the ExpCM draws less information from the experiment. At the most extreme value of $\beta=0$, the preferences would be perfectly uniform and look similar to the GY94 preferences in Fig. 4. In reality, β never reaches a value this low, indicating the deep mutational scanning experiments remain somewhat informative about real natural selection across the entire swath of HAs. However, Fig. 6 shows that the amino-acid preferences clearly become less informative about natural selection as we move away from the experimental sequence on the tree. This shifting of amino-acid preferences helps explain why the ExpCM informed by the average of the H1 and H3 experiments performs best (Table 1, Figs 4 and 5): averaging the measurements across these two HAs is a heuristic method of accounting for shifts in preferences during HA evolution.

The fact that amino-acid preferences shift as a protein evolves leaves us with an inherent tension: models with site-specific amino-acid preferences only become important for accurate branch-length estimation as sequences become increasingly diverged, but this same divergence degrades the accuracy of extrapolating the amino-acid preferences from any given experiment across the phylogenetic tree. Crucially, this problem is more fundamental than the inability of a single deep

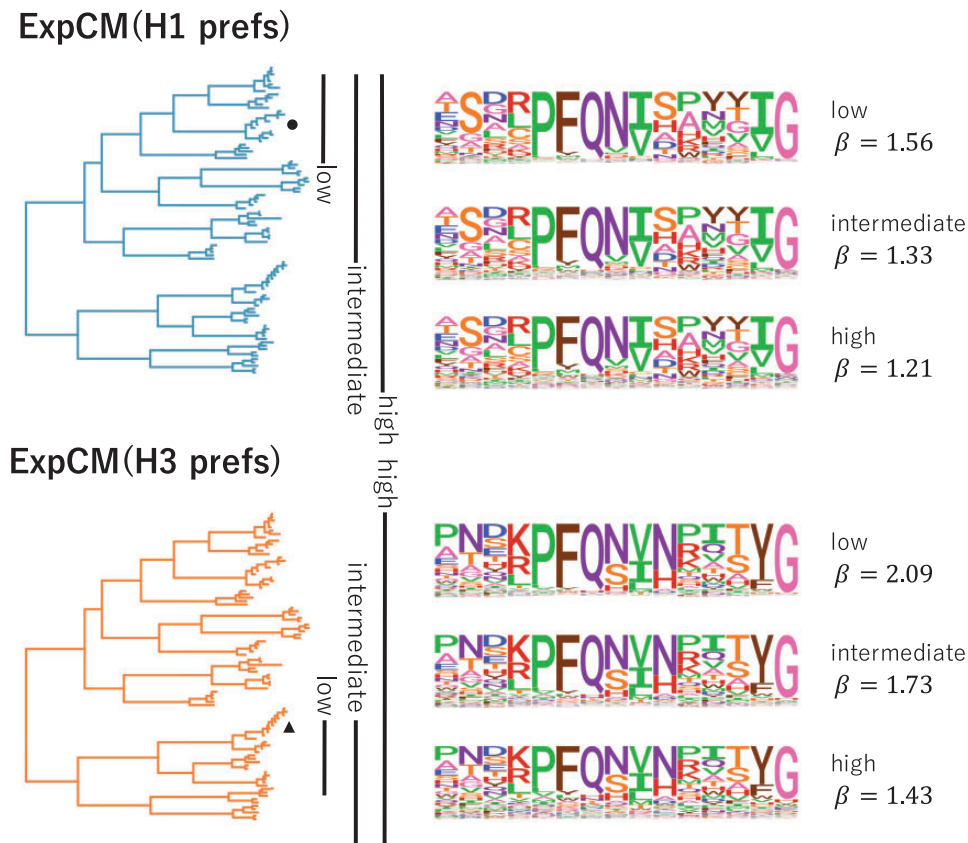


Figure 6. The congruence between natural selection and the deep mutational scanning measurements decreases with sequence divergence. We fit an ExpCM informed by the H1 or H3 deep mutational scanning experiments to trees spanning sequences with low, intermediate, and high divergence from the sequence used in the experiment. The ExpCM stringency parameter (β) is a measure of the congruence between natural selection and the experimental measurements (Bloom 2014b; Hilton, Doud, and Bloom 2017). Larger values of β indicate that natural selection prefers the same amino acids as the experiments but with greater stringency. As divergence increases between the HA used in the experiment and the other sequences in the tree, the β value decreases and the amino-acid preference ‘flatten’. Therefore, the preferences measured in each experiment are progressively less congruent with natural selection as we include increasingly diverged sequences.

mutational scanning experiment to measure amino-acid preferences in more than one genetic background. If amino-acid preferences shift during evolution, there simply will not be any single set of time-homogeneous site-specific preferences that accurately describes evolution along the entirety of a phylogenetic tree that covers a wide span of sequences.

2.6 A model with amino-acid preferences estimated from natural sequences gives similar results to an ExpCM

The previous sections used ExpCMs, which are mutation-selection models that use site-specific amino-acid preferences that have been measured by experiments. However, there are other mathematically similar implementations of mutation-selection models that infer the amino-acid preferences directly from the natural sequence data. When these models are designed for use in phylogenetic inference, they are generally implemented in a Bayesian framework, which avoids the over-fitting problems associated with trying to make maximum-likelihood estimates of the thousands of amino-acid preference parameters (Lartillot 2014). (Note that the maximum-likelihood implementations of Tamuri, dos Reis, and Goldstein (2012) and Tamuri, Goldman, and dos Reis (2014) are designed for estimating the amino-acid preferences, not for phylogenetic inference.) The model most comparable to our ExpCMs is the codon

mutation-selection model implemented in PhyloBayes-MPI, which we will refer to as pbMutSel (Rodrigue and Lartillot 2014). In the pbMutSel model, the amino-acid preferences are modeled using Dirichlet processes rather than derived from experiments. However, like an ExpCM, a pbMutSel model still assumes a single set of time-homogeneous site-specific amino-acid preferences for the entire tree.

Comparing ExpCM and pbMutSel models can help determine the ultimate limits of mutation-selection models that assign each site a single set of amino-acid preferences. If the limitations of ExpCMs described above arise simply because the deep mutational scanning experiments do not correctly measure the ‘true’ amino-acid preferences across the entirety of a highly diverged phylogenetic tree, then we would expect the pbMutSel models (which infer these preferences from the entire tree) to perform better. On the other hand, if the major limitation is that no single set of time-homogeneous amino-acid preferences can fully describe evolution over the entire tree, then we would expect ExpCM and pbMutSel models to perform similarly.

We fit a pbMutSel model to the entire HA phylogenetic tree, and compared the results to those from analyzing the same tree with the best ExpCM, which is the ExpCM(H1+H3 avg)+ $\Gamma\omega$ variant. This is a direct apple-to-apples comparison, since the pbMutSel model also draws ω from a gamma-distribution (Rodrigue and Lartillot 2014). First, we compared the amino-acid preferences inferred by the pbMutSel model to the preferences




Figure 7. Models inferred from natural sequences have similar stationary states to models defined by experimental preferences and estimate similar branch lengths. We fit an $\text{ExpCM}(\text{H1}+\text{H3 avg})+\Gamma\omega$ and a pbMutSel to the full HA tree in Fig. 4. The pbMutSel amino-acid preferences are inferred from the natural HA sequences, while the ExpCM amino-acid preferences are experimentally measured and then rescaled by the stringency parameter in Table 1. (a) The pbMutSel preferences are more correlated with the re-scaled average of the H1 and H3 deep mutational scanning preferences than the individual re-scaled H1 and H3 deep mutational scanning preferences are to each other (Pearson's r : 0.74 versus 0.52). (b) The $\text{ExpCM}(\text{H1}+\text{H3 avg})+\Gamma\omega$ and pbMutSel models estimated similar branch lengths when fit to the entire HA tree. Points denote branch lengths between all pairs of tips on the tree. Blue and orange denote branches that lead to the H1 and H3 deep mutational scanning reference sequences respectively. The `phydms` program implementing ExpCM s and the `PhyloBayes-MPI` program implementing pbMutSel models give branch lengths in different units, so to facilitate direct comparison between the models, we have normalized all branch lengths returned by each program by the length of the branches separating the earliest (A/South Carolina/1918) and latest (A/Solomon Islands/2006) seasonal human H1 sequences on the tree.

measured in the experiments. Figure 7a shows that the preferences inferred by pbMutSel are quite similar to the (H1 + H3 avg) obtained by averaging the deep mutational scanning measurements for the H1 and H3 HAs. Notably, the amino-acid preferences from the pbMutSel model are more correlated with the (H1+H3 avg) than the H1 and H3 measurements are with each other (Fig. 7a). This strong correlation indicates that the $\text{ExpCM}(\text{H1}+\text{H3 avg})+\Gamma\omega$ is unlikely to be much different than a pbMutSel model that is parameterized only using the natural sequence data.

We next compared the branch lengths estimated by using the $\text{ExpCM}(\text{H1}+\text{H3 avg})+\Gamma\omega$ and pbMutSel models. As shown in Fig. 7b, these two models estimated similar branch lengths across the entire HA phylogenetic tree. However, the estimates are not identical, and the tension between local and global accuracy of the amino-acid preferences is still apparent. Specifically, the branches leading to the H1 or H3 sequences used in the experiments were estimated to be slightly longer by the ExpCM , while some other branches were estimated to be slightly longer by the pbMutSel model. The relatively longer branches leading to the experimental sequences when using the $\text{ExpCM}(\text{H1}+\text{H3 avg})+\Gamma\omega$ suggests that the ‘tree average’ amino-acid preferences inferred by the pbMutSel model are not as accurate as the preferences from the deep mutational scanning for sequences close to those used in the experiments. However, for sequences distant from those used in the experiments, the ‘tree average’ preferences inferred by the pbMutSel model appear to be slightly better than the experimental values. Therefore, while the ExpCM and pbMutSel models differ slightly in the extent to which they lengthen different branches, neither model can avoid the tension between the local and global accuracy of amino-acid preferences.

3. Discussion

We examined how estimates of deep branch lengths on phylogenetic trees are affected by accounting for the fact that proteins prefer specific amino acids at specific sites. We did this by comparing inferences from models informed by experimental measurements of site-specific amino-acid preferences with more conventional codon substitution models, as well as with models that infer the amino-acid preferences from the natural

sequences. We found that models that account for site-specific amino-acid preferences estimated deeper long branches, regardless of whether these preferences are measured experimentally or inferred from the sequence alignment. Additionally, we showed that the extension in branch length from site-specific amino-acid preferences is mostly independent of the extension that results from simply modeling rate variation.

Overall, our results underscore the importance of modeling purifying selection in a way that is more nuanced than simply allowing the substitution rate to vary across sites. Protein sites do not simply differ in their rates of substitution—different sites also prefer different amino acids. There are now two ways to account for this fact: using models informed by deep mutational scanning experiments, or using models that infer site-specific amino-acid preferences from the natural sequence alignment. Combining either type of model with rate variation increases the inferred length of deep branches relative to models that only incorporate rate variation. We expect that further improvements could be achieved by also incorporating other factors such as host-specific substitution rates (Worobey, Han, and Rambaut 2014) that are known to be important for modeling the evolution of viral genes such as HA.

However, assuming a single set of site-specific amino-acid preferences is still an imperfect way to model evolution over a highly diverged phylogenetic tree. In the case of the experimentally informed models, it is fairly obvious why this is true: the amino-acid preferences are measured in just one genetic background, and therefore provide only a single snapshot of preferences that shift over evolutionary time due to epistasis (Pollock, Thiltgen, and Goldstein 2012; Bazykin 2015; Doud, Ashenberg, and Bloom 2015; Shah, McCandlish, and Plotkin 2015; Haddox et al. 2018; Starr et al. 2018). As a result, experimentally measured amino-acid preferences are most accurate for sequences similar to the one used in the experiment, and so cause the largest increases in branch length in that region of the phylogenetic tree. However, this limitation is not unique to experimentally informed models, but is a general limitation of describing purifying selection using a single set of site-independent and time-homogenous amino-acid preferences. For instance, we showed that averaging experimental measurements on two protein homologs does a somewhat better job of capturing the ‘average’ constraint across the tree, and performs similarly to approaches

that infer the ‘average’ preferences from natural sequence data (Rodrigue, Philippe, and Lartillot 2010; Rodrigue and Lartillot 2014). But even these ‘average’ preferences exhibit a trade-off between local and global accuracy for the inference of deep branch lengths.

So while modeling site-specific amino-acid preferences is a clear improvement over most conventional models, the next step toward greater accuracy will require relaxing the assumption that these preferences are time homogeneous and site independent. Of course, many authors have pointed out the shortcomings of models that fail to account for the full site-interdependent complexity of purifying selection (Rodrigue et al. 2005; Choi et al. 2007; Pollock, Thiltgen, and Goldstein 2012; Goldstein and Pollock 2017). However, the challenge is to overcome these shortcomings with models that are tractable for real phylogenetic questions. There are two main issues: first, the Felsenstein pruning algorithm (Felsenstein 1981) that is typically used to evaluate phylogenetic likelihoods breaks down when sites are no longer treated independently. Some alternative algorithms have been proposed (Choi et al. 2007; Rodrigue et al. 2009, 2005; Bordner and Mittelman 2014), but they are still in their infancy. Second, site-interdependent models require a realistic ‘fitness function’ that describes the interactions among sites. It appears that typical structural modeling programs are insufficient for this purpose (Rodrigue et al. 2009). But hope comes from experimental progress in measuring actual site-interdependent constraints on proteins (Olson, Wu, and Sun 2014; Li et al. 2016; Steinberg and Ostermeier 2016; Wu et al. 2016), combined with new methods for using these measurements to parameterize fitness functions (Sailer and Harms 2017; Otwinowski, McCandlish, and Plotkin 2018). Perhaps some day such truly realistic models might be useful for phylogenetic inference. Until that day, our work shows that modeling a single set of time homogenous amino-acid preferences provides at least some improvement.

4. Methods

4.1 Substitution models

All of the substitution models used in this paper have been described previously. However, here we briefly recap their exact mathematical implementations.

4.1.1 GY94 model

The GY94 model is M0 variant of the Goldman-Yang model described by Yang et al. (2000). Specifically, the substitution rate P_{xy} from codon x to codon y is

$$P_{xy} = \begin{cases} 0 & \text{if } x \text{ and } y \text{ differ by more than one nucleotide,} \\ \Phi_y & \text{if } \mathcal{A}(x) = \mathcal{A}(y) \text{ and } x \text{ is converted to } y \text{ by a single - nucleotide transversion,} \\ \omega \Phi_y & \text{if } \mathcal{A}(x) \neq \mathcal{A}(y) \text{ and } x \text{ is converted to } y \text{ by a single - nucleotide transversion,} \\ \kappa \Phi_y & \text{if } \mathcal{A}(x) = \mathcal{A}(y) \text{ and } x \text{ is converted to } y \text{ by a single - nucleotide transition,} \\ \omega \kappa \Phi_y & \text{if } \mathcal{A}(x) \neq \mathcal{A}(y) \text{ and } x \text{ is converted to } y \text{ by a single - nucleotide transition,} \\ -\sum_{z \neq x} P_{xz} & \text{if } x = y, \end{cases} \quad (1)$$

where $\mathcal{A}(x)$ is the amino-acid encoded by codon x , κ is the transition-transversion rate, Φ_y is the equilibrium frequency of codon y , and ω is the relative rate of nonsynonymous and synonymous substitutions. We define the codon frequency parameters, Φ_y , using the ‘corrected F3X4’ method from Pond et al. (2010). There are nine parameters describing the nucleotide frequencies at each codon site (the nucleotides are constrained to sum to one at each codon position), and these parameter values are calculated from the empirical alignment frequencies. The

‘corrected F3X4’ method calculates the Φ_y values from these nucleotide frequencies but corrects for the exclusion of sequences with premature stop codons from the analysis.

The frequency p_x of codon x in the stationary state of a GY94 model is simply

$$p_x = \Phi_x. \quad (2)$$

Overall, a GY94 model has eleven free parameters: κ , ω , and the nine nucleotide frequency parameters used to define Φ_y .

4.1.2 Experimentally informed codon model (ExpCM)

The ExpCM models used in this paper are the ones described in Bloom (2017). Briefly, the rate of substitution $P_{r,xy}$ of site r from codon x to y is

$$P_{r,xy} = Q_{xy} \times F_{r,xy} \quad (3)$$

where Q_{xy} is proportional to the rate of mutation from x to y , $F_{r,xy}$ is proportional to the probability that this mutation fixes, and the diagonal elements P_{xx} are set by $P_{xx} = -\sum_{z \neq x} P_{xz}$.

The rate of mutation Q_{xy} is assumed to be uniform across sites, and takes an HKY85-like (Hasegawa, Kishino, and Yano 1985) form as

$$Q_{xy} = \begin{cases} 0 & \text{if } x \text{ and } y \text{ differ by more than one nucleotide,} \\ \varphi_w & \text{if } x \text{ can be converted to } y \text{ by a transversion of a nucleotide to } w, \\ \kappa \times \varphi_w & \text{if } x \text{ can be converted to } y \text{ by a transition of a nucleotide to } w \end{cases} \quad (4)$$

where φ_w is the nucleotide frequency of nucleotide w and κ is the transition-transversion rate.

The deep mutational scanning amino-acid preferences are incorporated into the ExpCM via the $F_{r,xy}$ terms. The experiments measure the preference $\pi_{r,a}$ of every site r for every amino acid a . $F_{r,xy}$ is defined in terms of these experimentally measured amino-acid preferences as

$$F_{r,xy} = \begin{cases} 1 & \text{if } \mathcal{A}(x) = \mathcal{A}(y), \\ \omega \times \frac{\ln \left[\left(\frac{\pi_{r,\mathcal{A}(y)}}{\pi_{r,\mathcal{A}(x)}} \right)^\beta \right]}{1 - \left(\frac{\pi_{r,\mathcal{A}(x)}}{\pi_{r,\mathcal{A}(y)}} \right)^\beta} & \text{if } \mathcal{A}(x) \neq \mathcal{A}(y), \end{cases} \quad (5)$$

where β is the stringency parameter (Bloom 2014b; Hilton, Doud, and Bloom 2017) and ω is the relative rate of nonsynonymous to synonymous substitutions after accounting for the amino-acid preferences.

The stationary state of an ExpCM is

$$p_{r,x} = \frac{\varphi_{x_1} \varphi_{x_2} \varphi_{x_3} \left(\pi_{r,\mathcal{A}(x)} \right)^\beta}{\sum_z \varphi_{z_1} \varphi_{z_2} \varphi_{z_3} \left(\pi_{r,\mathcal{A}(z)} \right)^\beta} \quad (6)$$

where φ_{x_1} , φ_{x_2} , and φ_{x_3} are the nucleotides at position 1, 2, and 3 of codon x .

An ExpCM has five free parameters: κ , ω , and the three independent φ_x values. The amino-acid preferences $\pi_{r,a}$ are not free parameters since they are determined a priori by an experiment independent of the sequence alignment being analyzed.

4.1.3 $\Gamma\omega$ rate variation

The GY94+ $\Gamma\omega$ is equivalent to the M5 model in Yang et al. (2000) with ω drawn from $K=4$ categories. The ExpCM+ $\Gamma\omega$ similarly

draws ω from a Γ distribution discretized into $K=4$ bins. Each bin is equally weighted and ω takes on the mean value of the bin. Because the Γ distribution is defined by two parameters, adding $\Gamma\omega$ to a model with a single ω adds one free parameter. Therefore, the GY94+ $\Gamma\omega$ model has twelve free parameters, and the ExpCM+ $\Gamma\omega$ model has six free parameters.

4.1.4 GY94 with ω_r

In Fig. 2, we describe GY94 models where each site r has its own ω_r value that is calculated from the amino-acid preferences using the relationship described by Spielman and Wilke (2015b). This relationship defines the expected rate of nonsynonymous to synonymous substitutions given the amino-acid preferences. We first fit an ExpCM to the ‘low divergence’ H1 subtree (parameter values in Supplementary Table S2), which allows us to calculate $P_{r,xy}$ (Equation 3), Q_{xy} (Equation 4), and $p_{r,x}$ (Equation 6). We then calculated ω_r using the equation of Spielman and Wilke (2015b), normalizing by the gene-wide ω fit by the ExpCM:

$$\omega_r = \frac{\sum_x \sum_{y \in N_x} p_{r,x} \times \frac{P_{r,xy}}{\omega}}{\sum_x \sum_{y \in N_x} p_{r,x} \times Q_{xy}}, \quad (7)$$

where N_x is the set of codons that are nonsynonymous to codon x and differ from codon x by only one nucleotide.

4.2 HA amino-acid preferences from deep mutational scanning experiments

We used amino-acid preferences measured in deep mutational scans of the A/WSN/1933 H1 HA (Doud and Bloom 2016) and the A/Perth/2009 H3 HA (Lee et al. 2018) to define the amino-acid preferences that inform the ExpCMs. We only used sites that can be unambiguously aligned in these H1 and H3 HAs. These alignable sites and their mapping to sequential numbering of the HA sequences used in the deep mutational scanning experiments are in Supplementary file 1. The experimentally measured amino-acid preferences masked to just include these alignable sites are in Supplementary files 2 and 3. For the average preference set, we took the pairwise average of the H1 and H3 preferences. The preference for every amino acid a at every site r in the average preference set is

$$\pi_{r,a,(H1+H3 \text{ avg})} = \frac{\pi_{r,a,H1} + \pi_{r,a,H3}}{2} \quad (8)$$

4.3 HA sequences and tree topology

We downloaded all full-length, coding sequences for 15 of the 18 influenza A virus HA subtypes from the Influenza Virus Resource Database (Bao et al. 2008) in June 2017. We excluded rare subtypes 15, 17, and 18, which have few sequences in the database. We filtered and aligned the sequences using `phydms_s_prepalignment` (Hilton, Doud, and Bloom 2017). Specifically, we used `phydms_s_prepalignment` with the flag `-minidentity 0.3` to remove sequences with ambiguous nucleotides, premature stops, or frameshift mutations as well as redundant sequences. We also removed all codon sites which are not alignable between the H1 HA and H3 HA used in the deep mutational scanning experiments (these alignable sites are listed in Supplementary file 1). We subsampled the remaining sequences to five per subtype with ≤ 1 sequence per year per subtype. We also included a small number of sequences from the major

human and equine influenza lineages to ensure representation of these well-studied lineages. The resulting alignment contains 92 sequences, and is provided in Supplementary file 4.

We created four subalignments with ‘low’ and ‘intermediate’ divergence from either the H1 or the H3 deep mutational scanning reference sequence for the analysis in Fig. 6. The ‘low divergence’ alignments had $\geq 59\%$ amino-acid identity to the sequence used in the deep mutational scanning, and the ‘intermediate divergence’ alignments had $\geq 46\%$ identity from the reference sequence (Supplementary Fig. S5).

We inferred the tree topology of each alignment using `RAxML` (Stamatakis 2006) and the GTRCAT model. We estimated the branch lengths of this fixed topology using each ExpCM and GY94 models with `phydms_comprehensive` (Hilton, Doud, and Bloom 2017).

4.4 Asymptotic amino-acid sequence identity

For the analysis in Fig. 2, we fit models to the ‘low divergence’ H1 subtree. This gave the parameter values in Supplementary Table S2.

For each model, we calculated the expected amino-acid sequence identity for two sequences separated by a branch length of t as

$$\sum_a \sum_{x \in a} p_{r,x} \sum_{y \in a} [e^{tP_r}]_{xy} \quad (9)$$

where a ranges over all twenty amino acids, $x \in a$ indicates that x ranges over all codons that encode amino-acid a , $p_{r,x}$ is the stationary state of the model at site r and codon x (given by Equation 2 for GY94-family models, and Equation 6 for ExpCM-family models), and $[e^{tP_r}]_{xy}$ is the value in row x and column y of the matrix obtained by exponentiating the product of t and the substitution matrix P_r for site r (defined by Equation 1 for GY94-family models and Equation 3 for ExpCM-family models).

4.5 Simulations

For Fig. 3, we simulated sequences using `pyvolve` (Spielman and Wilke 2015a) along the full HA tree using an ExpCM defined by parameters fit to the ‘low divergence’ H1 subtree (Supplementary Table S2). We performed 10 replicate simulations and estimated the branch lengths for each replicate using `phydms_comprehensive` (Hilton, Doud, and Bloom 2017).

4.6 pbMutSel inference with PhyloBayes-MPI

For Fig. 7, we fit a pbMutSel model to the full HA tree. We ran one chain for 5,500 steps, saved every sample, and discarded the first 550 samples as a burn-in. We used `PhyloBayes-MPI` program `readpb_mpi` to compute the majority-rule consensus tree and the posterior average site-specific amino-acid preferences. Convergence was assessed visually using Tracer (Rambaut et al. 2018) and by the correlation of amino-acid preferences inferred by two independent chains ($r=0.996$).

In order to make the branch lengths in Fig. 7 comparable between the pbMutSel tree returned by `PhyloBayes-MPI` and the other trees returned by `phydms`, we normalized the branch lengths on the pbMutSel consensus tree and the ExpCM(H1+H3 avg)+ $\Gamma\omega$ by dividing each branch by the length from A/South Carolina/1/1918 and A/Solomon Islands/3/2006. These two H1 sequences are early and late representatives of the longest known human influenza lineage, and are of sufficiently high identity that different ExpCM and GY94

substitution models all estimate nearly identical branch lengths separating them.

5. Software versions and computer codes

All codes used for the analyzes in this article are available at https://github.com/jbloomlab/divergence_timing_manuscript.

The external computer programs that we used were

- `phydms` (Hilton, Doud, and Bloom 2017) version 2.2.2 (available at https://github.com/jbloomlab/divergence_timing_manuscript) to fit the ExpCM and GY94 models.
- `pyvolve` (Spielman and Wilke 2015a) version 0.8.7 (available at <https://github.com/sjspielman/pyvolve>) to simulate the sequences.
- `PhyloBayes-MPI` (Rodrigue and Lartillot 2014) version 1.8 (available at <https://github.com/bayesiancook/pbmpi>) to fit the pb MutSel model.
- `RAxML` (Stamatakis 2006) version 8.2.11 (available at <https://github.com/stamatak/standard-RAxML>) to infer tree topology.
- We used `ggplot2` (Wickham 2016), `ggtree` (Yu et al. 2017), and `ggseqlogo` (Wagih 2017) for visualization of the results.
- `snakemake` (Köster and Rahmann 2012) version 3.11.2 (available at <https://snakemake.readthedocs.io/en/stable/>) to run the pipelines.

Supplementary data

Supplementary data are available at Virus Evolution online.

Acknowledgements

We thank Erick Matsen and Trevor Bedford for helpful comments about the project and manuscript. S.K.H. is supported in part by training grant T32AI083203 from the NIAID of the National Institutes of Health. This work was supported by the NIAID and NIGMS of the NIH under grant numbers R01AI127893 and R35GM126911. J.D.B. is an Investigator of the Howard Hughes Medical Institute. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. All code and data used for the analyses in this article are available at https://github.com/jbloomlab/divergence_timing_manuscript.

Conflict of interest: None declared.

References

Aiweesakun, P., and Katzourakis, A. (2016) 'Time-Dependent Rate Phenomenon in Viruses', *Journal of Virology*, 90, 7184–95.

Arenas, M. (2015) 'Trends in Substitution Models of Molecular Evolution', *Frontiers in Genetics*, 6, 319.

Bao, Y. et al. (2008) 'The Influenza Virus Resource at the National Center for Biotechnology Information', *Journal of Virology*, 82, 596–601.

Bazykin, G. A. (2015) 'Changing Preferences: Deformation of Single Position Amino Acid Fitness Landscapes and Evolution of Proteins', *Biology Letters*, 11, 20150315.

Bedford, T. et al. (2014) 'Integrating Influenza Antigenic Dynamics with Molecular Evolution', *eLife*, 3, e01914.

Bloom, J. D. (2014a) 'An Experimentally Determined Evolutionary Model Dramatically Improves Phylogenetic Fit', *Molecular Biology and Evolution*, 31, 1956–78.

——— (2014b) 'An Experimentally Informed Evolutionary Model Improves Phylogenetic Fit to Divergent Lactamase Homologs', *Molecular Biology and Evolution*, 31, 2753–69.

——— (2017) 'Identification of Positive Selection in Genes Is Greatly Improved by Using Experimentally Informed Site-Specific Models', *Biology Direct*, 12, 1.

Bordner, A. J., and Mittelman, H. D. (2014) 'A New Formulation of Protein Evolutionary Models That account for Structural Constraints', *Molecular Biology and Evolution*, 31, 736–49.

Carroll, S. A. et al. (2013) 'Molecular Evolution of Viruses of the Family Filoviridae Based on 97 Whole-Genome Sequences', *Journal of Virology*, 87, 2608–16.

Choi, S. C. et al. (2007) 'Quantifying the Impact of Protein Tertiary Structure on Molecular Evolution', *Molecular Biology and Evolution*, 24, 1769–82.

Doud, M. B., Ashenberg, O., and Bloom, J. D. (2015) 'Site-Specific Amino Acid Preferences Are Mostly Conserved in Two Closely Related Protein Homologs', *Molecular Biology and Evolution*, 32, 2944–60.

——— and Bloom, J. D. (2016) 'Accurate Measurement of the Effects of All Amino-Acid Mutations to Influenza Hemagglutinin', *Viruses*, 8, 155.

Drummond, A. J. et al. (2006) 'Relaxed Phylogenetics and Dating with Confidence', *PLoS Biology*, 4, e88.

Duchêne, D. A. et al. (2015a) 'Evaluating the Adequacy of Molecular Clock Models Using Posterior Predictive Simulations', *Molecular Biology and Evolution*, 32, 2986–95.

Duchêne, S., Di Giallonardo, F., and Holmes, E. C. (2015b) 'Substitution Model Adequacy and Assessing the Reliability of Estimates of Virus Evolutionary Rates and Time Scales', *Molecular Biology and Evolution*, 33, 255–67.

———, Holmes, E. C., and Ho, S. Y. (2014) 'Analyses of Evolutionary Dynamics in Viruses Are Hindered by a Time-Dependent Bias in Rate Estimates', *Proceedings of the Royal Society of London B*, 281, 20140732.

Echave, J., Spielman, S. J., and Wilke, C. O. 2016. 'Causes of Evolutionary Rate Variation among Protein Sites', *Nature Reviews Genetics*, 17, 109–121.

Fares, M. A., and Holmes, E. C. (2002) 'A Revised Evolutionary History of Hepatitis B Virus (HBV)', *Journal of Molecular Evolution*, 54, 807–14.

Felsenstein, J. (1978) 'Cases in Which Parsimony or Compatibility Methods Will Be Positively Misleading', *Systematic Biology*, 27, 401–10.

——— (1981) 'Evolutionary Trees from DNA Sequences: A Maximum Likelihood Approach', *Journal of Molecular Evolution*, 17, 368–76.

Fowler, D. M., and Fields, S. (2014) 'Deep Mutational Scanning: A New Style of Protein Science', *Nature Methods*, 11, 801–7.

Furuse, Y., Suzuki, A., and Oshitani, H. (2010) 'Origin of Measles Virus: Divergence from Rinderpest Virus between the 11th and 12th Centuries', *Virology Journal*, 7, 52.

Goldman, N., and Yang, Z. (1994) 'A Codon-Based Model of Nucleotide Substitution for Protein-Coding DNA Sequences', *Molecular Biology and Evolution*, 11, 725–36.

Goldstein, R. A., and Pollock, D. D. (2017) 'Sequence Entropy of Folding and the Absolute Rate of Amino Acid Substitutions', *Nature Ecology & Evolution*, 1, 1923.

Gong, L. I., Suchard, M. A., and Bloom, J. D. (2013) 'Stability-Mediated Epistasis Constrains the Evolution of an Influenza Protein', *eLife*, 2, e00631.

Ha, Y. et al. (2002) 'H5 Avian and H9 Swine Influenza Virus Haemagglutinin Structures: Possible Origin of Influenza Subtypes', *The EMBO Journal*, 21, 865–75.

- Haddox, H. K. et al. (2018) 'Mapping Mutational Effects along the Evolutionary Landscape of HIV Envelope', *eLife*, 7, e34420.
- Halpern, A. L., and Bruno, W. J. (1998) 'Evolutionary Distances for Protein-Coding Sequences: Modeling Site-Specific Residue Frequencies', *Molecular Biology and Evolution*, 15, 910–7.
- Harms, M. J., and Thornton, J. W. (2014) 'Historical Contingency and Its Biophysical Basis in Glucocorticoid Receptor Evolution', *Nature*, 512, 203–7.
- Hasegawa, M., Kishino, H., and Yano, T. (1985) 'Dating of the Human-Ape Splitting by a Molecular Clock of Mitochondrial DNA', *Journal of Molecular Evolution*, 22, 160–74.
- Hilton, S. K., Doud, M. B., and Bloom, J. D. (2017) 'Phydm: Software for Phylogenetic Analyses Informed by Deep Mutational Scanning', *PeerJ*, 5, e3657.
- Ho, S. Y. et al. (2015) 'Time-Dependent Estimates of Molecular Evolutionary Rates: Evidence and Causes', *Molecular Ecology*, 24, 6007–12.
- Holmes, E. C. (2003) 'Molecular Clocks and the Puzzle of RNA Virus Origins', *Journal of Virology*, 77, 3893–7.
- Köster, J., and Rahmann, S. (2012) 'Snakemake – A Scalable Bioinformatics Workflow Engine', *Bioinformatics*, 28, 2520–2.
- Lartillot, N. (2014) 'The Bayesian Kitchen: Overcoming the Fear of Over-Parameterization'. <http://bayesiancook.blogspot.com/2014/01/the-myth-of-over-parameterization.html>. Last accessed 12 Mar 2018.
- , Brinkmann, H., and Philippe, H. (2007) 'Suppression of Long-Branch Attraction Artefacts in the Animal Phylogeny Using a Site-Heterogeneous Model', *BMC Evolutionary Biology*, 7, S4.
- and Philippe, H. (2004) 'A Bayesian Mixture Model for across-Site Heterogeneities in the Amino-Acid Replacement Process', *Molecular Biology and Evolution*, 21, 1095–109.
- Le, S. Q., Lartillot, N., and Gascuel, O. (2008) 'Phylogenetic Mixture Models for Proteins', *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 363, 3965–76.
- Lee, J. M. et al. (2018). 'Deep Mutational Scanning of Hemagglutinin Helps Predict Evolutionary Fates of Human H3N2 Influenza Variants', *Proceedings of the National Academy of Sciences of the United States of America*. 115, E8276–85.
- Li, C. et al. (2016) 'The Fitness Landscape of a Trna Gene', *Science*, 352, 837–40.
- McCandlish, D. M., and Stoltzfus, A. (2014) 'Modeling Evolution Using the Probability of Fixation: History and Implications', *The Quarterly Review of Biology*, 89, 225–52.
- Murrell, B. et al. (2015) 'Gene-Wide Identification of Episodic Selection', *Molecular Biology and Evolution*, 32, 1365–71.
- Nielsen, R. (2006) *Statistical Methods in Molecular Evolution*. Springer.
- Nobusawa, E. et al. (1991) 'Comparison of Complete Amino Acid Sequences and Receptor-Binding Properties among 13 Serotypes of Hemagglutinins of Influenza A Viruses', *Virology*, 182, 475–85.
- Olson, C. A., Wu, N. C., and Sun, R. (2014) 'A Comprehensive Biophysical Description of Pairwise Epistasis throughout an Entire Protein Domain', *Current Biology*, 24, 2643–51.
- Ortlund, E. A. et al. (2007) 'Crystal Structure of an Ancient Protein: Evolution by Conformational Epistasis', *Science*, 317, 1544–8.
- Otwinowski, J., McCandlish, D. M., and Plotkin, J. B. (2018). 'Inferring the Shape of Global Epistasis', *Proceedings of the National Academy of Sciences of the United States of America*, doi: 10.1073/pnas.1804015115.
- Philippe, H., and Laurent, J. (1998) 'How Good Are Deep Phylogenetic Trees?', *Current Opinion in Genetics & Development*, 8, 616–23.
- Pollock, D. D., Thiltgen, G., and Goldstein, R. A. (2012) 'Amino Acid Coevolution Induces an Evolutionary Stokes Shift', *Proceedings of the National Academy of Sciences of the United States of America*, 109, E1352–9.
- Pond, S. K. et al. (2010) 'Correcting the Bias of Empirical Frequency Parameter Estimators in Codon Models', *PLoS One*, 5, e11230.
- Posada, D., and Buckley, T. R. (2004) 'Model Selection and Model Averaging in Phylogenetics: Advantages of Akaike Information Criterion and Bayesian Approaches over Likelihood Ratio Tests', *Systematic Biology*, 53, 793–808.
- Quang, S. L., Gascuel, O., and Lartillot, N. (2008) 'Empirical Profile Mixture Models for Phylogenetic Reconstruction', *Bioinformatics*, 24, 2317–23.
- Rambaut, A. et al. (2018) Posterior Summarisation in Bayesian Phylogenetics Using Tracer 1.7, *Systematic Biology*, 67, 901–4.
- Rodrigue, N. (2013) 'On the Statistical Interpretation of Site-Specific Variables in Phylogeny-Based Substitution Models', *Genetics*, 193, 557–64.
- and Lartillot, N. (2014) 'Site-Heterogeneous Mutation-Selection Models within the PhyloBayes-MPI Package', *Bioinformatics*, 30, 1020–1.
- and —— (2017) 'Detecting Adaptation in Protein-Coding Genes Using a Bayesian Site-Heterogeneous Mutation-Selection Codon Substitution Model', *Molecular Biology and Evolution*, 34, 204–14.
- et al. (2005) 'Site Interdependence Attributed to Tertiary Structure in Amino Acid Sequence Evolution', *Gene*, 347, 207–17.
- et al. (2009) 'Computational Methods for Evaluating Phylogenetic Models of Coding Sequence Evolution with Dependence between Codons', *Molecular Biology and Evolution*, 26, 1663–76.
- , Philippe, H., and Lartillot, N. (2010) 'Mutation-Selection Models of Coding Sequence Evolution with Site-Heterogeneous Amino Acid Fitness Profiles', *Proceedings of the National Academy of Sciences of the United States of America*, 107, 4629–34.
- Russell, R. et al. (2004) 'H1 and H7 Influenza Haemagglutinin Structures Extend a Structural Classification of Haemagglutinin Subtypes', *Virology*, 325, 287–96.
- Sailer, Z. R., and Harms, M. J. (2017) 'Detecting High-Order Epistasis in Nonlinear Genotype-Phenotype Maps', *Genetics*, 205, 1079–88.
- Shah, P., McCandlish, D. M., and Plotkin, J. B. (2015) 'Contingency and Entrenchment in Protein Evolution under Purifying Selection', *Proceedings of the National Academy of Sciences of the United States of America*, 112, E3226–35.
- Spielman, S. J., and Wilke, C. O. (2015a) 'Pyvolve: A Flexible Python Module for Simulating Sequences along Phylogenies', *PLoS One*, 10, e0139047.
- and —— (2015b) 'The Relationship between dN/dS and Scaled Selection Coefficients', *Molecular Biology and Evolution*, 32, 1097–108.
- Stamatakis, A. (2006) 'RAxML-VI-HPC: Maximum Likelihood-Based Phylogenetic Analyses with Thousands of Taxa and Mixed Models', *Bioinformatics*, 22, 2688–90.
- Starr, T. N. et al. (2018) 'Pervasive Contingency and Entrenchment in a Billion Years of Hsp90 Evolution', *Proceedings of the National Academy of Sciences of the United States of America*. doi: 10.1073/pnas.1718133115.

- Steinberg, B., and Ostermeier, M. (2016) 'Shifting Fitness and Epistatic Landscapes Reflect Trade-Offs along an Evolutionary Pathway', *Journal of Molecular Biology*, 428, 2730–43.
- Susko, E., Lincker, L., and Roger, A. J. (2018) 'Accelerated Estimation of Frequency Classes in Site-Heterogeneous Profile Mixture Models', *Molecular Biology and Evolution*, 35, 1266–83.
- Tamuri, A. U., dos Reis, M., and Goldstein, R. A. (2012) 'Estimating the Distribution of Selection Coefficients from Phylogenetic Data Using Sitewise Mutation-Selection Models', *Genetics*, 190, 1101–15.
- , Goldman, N., and dos Reis, M. (2014) 'A Penalized Likelihood Method for Estimating the Distribution of Selection Coefficients from Phylogenetic Data', *Genetics*, 114.
- Taylor, D. J. et al. (2014) 'Evidence That Ebolaviruses and Cuevaviruses Have Been Diverging from Marburgviruses since the Miocene', *PeerJ*, 2, e556.
- Tufts, D. M. et al. (2014) 'Epistasis Constrains Mutational Pathways of Hemoglobin Adaptation in High-Altitude Pikas', *Molecular Biology and Evolution*, 32, 287–98.
- Wagih, O. (2017) 'ggseqlogo: A Versatile R Package for Drawing Sequence Logos', *Bioinformatics*, 33, 3645–7.
- Wang, H. C. et al. (2008) 'A Class Frequency Mixture Model That Adjusts for Site-Specific Amino Acid Frequencies and Improves Inference of Protein Phylogeny', *BMC Evolutionary Biology*, 8, 331.
- Wertheim, J. O., and Kosakovsky Pond, S. L. (2011) 'Purifying Selection Can Obscure the Ancient Age of Viral Lineages', *Molecular Biology and Evolution*, 28, 3355–65.
- and Worobey, M. (2009) 'Dating the Age of the SIV Lineages That Gave Rise to HIV-1 and HIV-2', *PLoS Computational Biology*, 5, e1000377.
- et al. (2013) 'A Case for the Ancient Origin of Coronaviruses', *Journal of Virology*, 87, 7039–45.
- Wickham, H. 2016. *ggplot2: Elegant Graphics for Data Analysis*. Springer.
- Worobey, M., Han, G. Z., and Rambaut, A. (2014) 'A Synchronized Global Sweep of the Internal Genes of Modern Avian Influenza Virus', *Nature*, 508, 254.
- et al. (2010) 'Island Biogeography Reveals the Deep History of SIV', *Science*, 329, 1487. (11 co-authors).
- Wu, N. C. et al. (2016) 'Adaptation in Protein Fitness Landscapes Is Facilitated by Indirect Paths', *eLife*, 5, e16965.
- Yang, Z. (1994) 'Maximum Likelihood Phylogenetic Estimation from DNA Sequences with Variable Rates over Sites: Approximate Methods', *Journal of Molecular Evolution*, 39, 306–14.
- and Nielsen, R. (2008) 'Mutation-Selection Models of Codon Substitution and Their Use to Estimate Selective Strengths on Codon Usage', *Molecular Biology and Evolution*, 25, 568–79.
- and Rannala, B. (2012) 'Molecular Phylogenetics: Principles and Practice', *Nature Reviews. Genetics*, 13, 303.
- et al. (2000) 'Codon-Substitution Models for Heterogeneous Selection Pressure at Amino Acid Sites', *Genetics*, 155, 431–49.
- Yu, G. et al. (2017) 'ggtree: An R Package for Visualization and Annotation of Phylogenetic Trees with Their Covariates and Other Associated Data', *Methods in Ecology and Evolution*, 8, 28–36.
- Zuckermandl, E., and Pauling, L. (1965) 'Evolutionary Divergence and Convergence in Proteins', in *Evolving Genes and Proteins*, pp. 97–166. New York: Academic Press.