

Automation and Evaluation of the SOWH Test with SOWHAT

SAMUEL H. CHURCH^{1,*}, JOSEPH F. RYAN^{2,3}, AND CASEY W. DUNN¹

¹Department of Ecology and Evolutionary Biology, Brown University, Providence, Rhode Island, USA; ²Whitney Laboratory for Marine Biosciences, St. Augustine, Florida, USA; and ³Sars International Centre For Marine Molecular Biology, Bergen, Norway

*Correspondence to be sent to: Ecology and Evolutionary Biology, Brown University, 69 Brown Street, Providence, Rhode Island, 02912, USA; E-mail: church@g.harvard.edu

Received 18 June 2014; reviews returned 23 July 2015; accepted 24 July 2015
Associate Editor: Mark Holder

Abstract.—The Swofford–Olsen–Waddell–Hillis (SOWH) test evaluates statistical support for incongruent phylogenetic topologies. It is commonly applied to determine if the maximum likelihood tree in a phylogenetic analysis is significantly different than an alternative hypothesis. The SOWH test compares the observed difference in log-likelihood between two topologies to a null distribution of differences in log-likelihood generated by parametric resampling. The test is a well-established phylogenetic method for topology testing, but it is sensitive to model misspecification, it is computationally burdensome to perform, and its implementation requires the investigator to make several decisions that each have the potential to affect the outcome of the test. We analyzed the effects of multiple factors using seven data sets to which the SOWH test was previously applied. These factors include a number of sample replicates, likelihood software, the introduction of gaps to simulated data, the use of distinct models of evolution for data simulation and likelihood inference, and a suggested test correction wherein an unresolved “zero-constrained” tree is used to simulate sequence data. To facilitate these analyses and future applications of the SOWH test, we wrote SOWHAT, a program that automates the SOWH test. We find that inadequate bootstrap sampling can change the outcome of the SOWH test. The results also show that using a zero-constrained tree for data simulation can result in a wider null distribution and higher *p*-values, but does not change the outcome of the SOWH test for most of the data sets tested here. These results will help others implement and evaluate the SOWH test and allow us to provide recommendations for future applications of the SOWH test. SOWHAT is available for download from <https://github.com/josephryan/SOWHAT>. [Phylogenetics; SOWH test; topology test]

A phylogenetic topology test evaluates whether the difference in optimality criterion score between incongruent hypotheses is significant. In some cases, the test is used to determine whether a data set provides significantly more support for one of several previously proposed phylogenetic hypotheses. In other cases, the test is used to compare a novel or unexpected phylogenetic result to a previously proposed hypothesis. In both scenarios, the observed difference in the optimality criterion score between two trees is compared to an estimated null distribution of differences in scores. Phylogenetic topology tests differ largely in how this null distribution is created. The Kishino–Hasegawa test (KH) (Kishino and Hasegawa 1989), for example, creates a null distribution by analyzing data sets created by sampling with replacement from the original data set (Goldman et al. 2000). However, this approach is subject to selection bias and is only appropriate for tests of hypotheses selected *a priori*, such as the comparison of two alternative hypotheses from the literature. In cases where hypotheses are not determined *a priori*, as when comparing a topology to an unanticipated maximum likelihood tree produced from the same data, tests such as the approximately unbiased test (AU) (Shimodaira 2001), the Shimodaira–Hasegawa test (SH) (Shimodaira and Hasegawa 1999), and the Swofford–Olsen–Waddell–Hillis test (SOWH) are more appropriate (Goldman et al. 2000). The SOWH test (Swofford et al. 1996), as a parametric bootstrap method, can provide greater statistical power than non-parametric methods such as the SH test, though this comes at the cost of an increased reliance on the model of evolution (Goldman et al.

2000). This test complements other approaches such as Bayesian topology tests (Bergsten et al. 2013).

A typical SOWH test where an *a priori* hypothesis is compared to the maximum likelihood tree (Fig. 1) begins with performing two maximum likelihood searches on a data set. One is unconstrained and results in an estimate of the most likely tree. The other is constrained by an alternative hypothesis, which will result in a less likely topology. The difference in log-likelihood scores between these two topologies (δ) is the test statistic. New data sets are then simulated using the topology and parameter estimates (base frequencies, rates, alpha value, etc.) from the constrained likelihood search. On each of these data sets, two maximum likelihood searches—one unconstrained and one constrained by the alternative hypothesis—are performed and δ values are calculated. The observed value of δ is then compared to this distribution of simulated δ values. A significant test statistic is one which falls above some percentage, such as 95%, of simulated δ values.

The SOWH test can be burdensome to implement with existing tools. Several helpful step-by-step instructions for manual implementation are available (Crawford 2009; Anderson et al. 2014), but these approaches require extensive hands-on time which makes it difficult to systematically examine the behavior of the test under different conditions. Performing the SOWH test requires an investigator to make multiple decisions which may not be informed without an evaluation of the behavior of the test. These decisions include how many bootstrap samples to generate, which likelihood software to use, and how to treat gaps during data simulation. These

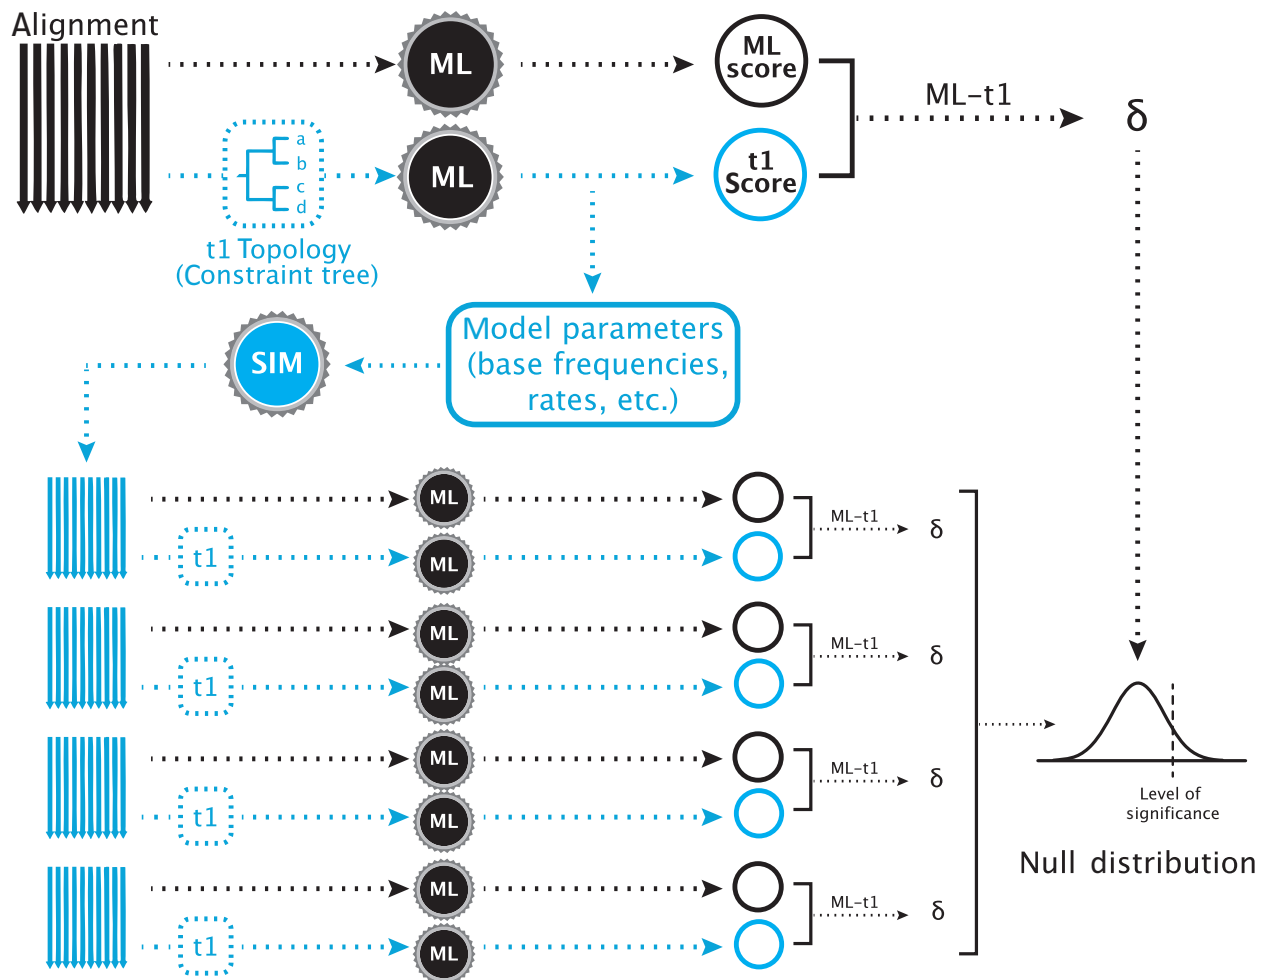


FIGURE 1. A typical SOWH test. The test begins with two maximum likelihood searches on a single alignment. One search, represented by the black arrow, is performed with no constraining topology. Another test, represented by the shaded arrow, is constrained to follow an *a priori* topology that represents a phylogenetic hypothesis incongruent with the maximum likelihood topology. The black gears represent maximum likelihood software used to score the trees (i.e., GARLI, RAxML). These two searches result in two maximum likelihood scores, the difference (δ) between which is the test statistic. From the constrained search, the optimized parameters and topology are retrieved and used to simulate new alignments with software (shaded gear) such as Seq-Gen. For each simulated alignment (shaded), two maximum likelihood searches are performed, one unconstrained (black arrow) and one constrained (shaded arrow), scores are obtained, and a δ value is calculated. The test statistic is compared to this distribution of δ values. A significantly large δ value is one which falls above some proportion of those generated by data simulation (i.e., 95%)

factors regularly vary across SOWH tests published in the literature, and before the analyses described here there has not been an evaluation of their effects on multiple data sets.

Furthermore, the test has been shown to have a high type I error rate as compared to other topology tests (Buckley 2002; Susko 2014). To correct this, two adjustments to the implementation of the SOWH test have been recommended. Before the present study, these hypothesized corrections have not been evaluated across multiple data sets. Buckley (2002) suggested that when the model of evolution for simulating data is the same model used in a maximum likelihood search on the data, it is more probable that the generating topology will be found. Because the generating topology in a SOWH test conforms to the alternative hypothesis,

both the constrained and unconstrained searches on the simulated data will find a tree very similar to the generating tree. This will result in smaller δ values, a more narrow null distribution, and a potentially higher rate of type I error. To correct for this, previous investigators recommended that the model of evolution for likelihood inference should be distinct from the model for simulating data (Buckley 2002; Pauly et al. 2004; Fig. 2).

More recently Susko (2014) observed that likelihood differences between unconstrained and constrained searches on simulated data are smaller than would be expected. Susko (2014) suggests that when data are simulated on a fully resolved tree, it is more likely that the generating topology used for simulation will be found. This will result in a more narrow distribution




FIGURE 2. A SOWH test using two models of evolution. In this test, two maximum likelihood searches are performed as described above using a model of evolution (Model 1). Instead of retrieving parameter values from the constrained search, as would be done in a typical SOWH test, an additional constrained maximum likelihood is performed using a different model of evolution (Model 2). Parameters are retrieved from this test and used to simulate new data. These simulated data sets are then scored using the same model of evolution used to score the original data sets (Model 1). This adjustment to the typical SOWH test was suggested following the assumption that a SOWH test performed with the same model of evolution for likelihood scoring and data simulation would result in a smaller δ values on simulated data, a smaller null distribution, and a more liberal test. For tests in our study which use the CAT-GTR model in PhyloBayes, both the additional constrained search and data simulation are performed using PhyloBayes—all other likelihood searches are performed using the specified likelihood software (i.e., GARLI or RAXML).

and a potentially higher type I error rate, as described above. As a correction, Susko (2014) suggests using a tree which is not fully resolved, specifically the newly proposed “zero-constrained” tree, as the topology for data simulation. This tree is generated by manipulating the most likely tree estimated in an unconstrained search so that all edges incongruent with the alternative hypothesis are reduced to zero or near zero, creating a polytomy.

Using seven published data sets previously analyzed with the SOWH test, we examined the effects on test outcome of multiple factors (sample size, likelihood software, and treatment of gaps) and proposed adjustments (model specification and generating topology). To facilitate implementation, we developed SOWHAT (as in, “The maximum likelihood tree differs from my alternative phylogenetic hypothesis,

so what?”), a program that automates the SOWH test (Fig. 1) and includes features that allow for additional complexities such as partitioned data sets. Using our results we provide recommendations for future implementation of the SOWH test. These recommendations, along with the program SOWHAT, allow for a more informed and less burdensome application of the SOWH test.

METHODS AND MATERIALS

Implementation of the SOWH test with SOWHAT

We used SOWHAT, our automation of the SOWH test, for all analyses. SOWHAT is available at <https://github.com/josephryan/SOWHAT>.

This tool evaluates the significance of the difference between the unconstrained maximum likelihood tree and a maximum likelihood tree inferred under a topology constraint provided by the user (Fig. 1). At a minimum, the user specifies the model of evolution under which the maximum likelihood searches will be performed, and provides as input an alignment file (in PHYLIP format) as well as the constraint topology to be tested (in Newick format). Two maximum likelihood trees—an unconstrained tree, and a tree constrained according to the provided constraint topology—are then inferred with either RAxML (Stamatakis 2006) or GARLI (Zwickl 2006), as specified by the user.

New data sets are simulated either by Seq-Gen (Rambaut and Grass 1997) or by PhyloBayes (Lartillot et al. 2009) if the CAT-GTR model is selected for simulation. The simulated alignments are generated using the topology, branch lengths, and model parameters (i.e., state frequencies, rates, and the alpha parameter for the gamma rate heterogeneity approximation) from the constrained analysis as inferred by the likelihood software or by PhyloBayes. If the correction of Susko (2014) is applied, the data are instead simulated on the zero-constrained tree. If the original data set is partitioned, parameters are estimated separately for each partition and new alignments are then generated following the partitioning scheme. Each of the simulated alignments is then used to estimate two maximum likelihood trees, using both an unconstrained search and a constrained search. The difference in log-likelihood of these two searches (δ) is calculated for each simulated data set, and the set of these differences make up the null distribution.

The observed log-likelihood difference (δ) is then compared to the null distribution and a p -value is calculated by the equation $p = d/n$ where d is the number of δ values greater than or equal to the observed value of δ and n is the sample size.

SOWHAT can be used to evaluate a hypothesized topology given data sets of nucleotide, amino acid, or binary characters. All models for likelihood inference available in RAxML or GARLI are available in SOWHAT.

Calculation of Confidence Intervals

SOWHAT calculates confidence intervals around a p -value following the addition of each new sample to the null distribution. The confidence intervals are calculated using the Clopper–Pearson method assuming a significance level of 95%. The confidence values can be used to determine adequacy of sample size by evaluating whether both the lower and upper interval fall on the same side of the significance level as the p -value.

Examined Data sets

We used Google Scholar (<http://scholar.google.com>) to perform a literature search for papers that cited

Goldman et al. (2000) and included the terms “SOWH” or “parametric bootstrap”.

We narrowed these roughly 400 results down to 40 which had deposited data on TreeBase (<http://treebase.org/>). From these we selected seven data sets for examination that represent a range of p -values and data set sizes. They are included with SOWHAT so that users can easily validate their installation of SOWHAT, explore the tool, and reproduce our analyses. The results presented here were prepared with the version of SOWHAT (v0.20) available at <https://github.com/josephryan/sowhat/tree/3137601014e24274ebc115acabe086290b5f46e7>. We analyzed the data with GARLI version 2.01.1067, RAxML version 8.1.15, PhyloBayes version 3.3f, seq-gen version 1.3.2x, and R version 3.0.0.

Buckley (2002).—From the analysis of the SOWH test by Buckley (2002), we selected a data set of mitochondrial ribosomal RNA genes (12S) originally assembled by Sullivan et al. (1995). Likelihood analysis of this alignment produces a topology of sigmoidontine rodents, which differs from the topology recovered by analyses of morphological, chromosomal, allozyme, and other DNA data sets (Buckley 2002). Buckley (2002) performed a series of SOWH tests using PAUP* and multiple models, all of which rejected the alternative hypothesis using a significance level of 0.05. Here we repeat the analyses using the models GTR+I+ Γ and GTR+ Γ . The data set contains eight taxa and 791 characters, and is referred to here as Buckley.

Dixon et al. (2007).—This data set contains 14 individuals of five species of European *Androsace*, a perennial herbaceous plant (Dixon et al. 2007). The alternative hypothesis to the maximum likelihood topology includes a monophyletic haplotype group of chloroplast regions for two species, *A. laggeri* and *A. pyrenica*. Dixon et al. (2007) performed a SOWH test using PAUP* which rejected this hypothesis using the Hasegawa, Kishino and Yano (HKY) model model of evolution, but not when using HKY+ Γ . The data set contains 2180 nucleotide characters, and is referred to here as Dixon.

Dunn et al. (2005).—This data set is composed of mitochondrial ribosomal RNA genes (16S and 18S) from 52 siphonophores (Hydrozoa: Cnidaria), as well as four out-group taxa (Dunn et al. 2005). The authors tested several hypotheses—the one we examine here unites *Bargmannia* with Agalmatidae *sensu stricto*. Dunn et al. (2005) performed a SOWH test which rejected this hypothesis, using parsimony (PAUP*) to score the original and simulated trees and maximum likelihood to estimate parameters and simulate data. The data set contains 2748 nucleotide characters, and is referred to here as Dunn.

Edwards et al. (2005).—This data set is composed of five distinct gene regions from 38 taxa in the group Cactaceae

(Edwards et al. 2005). The alternative hypothesis to the maximum likelihood topology creates a monophyletic group, *Pereskia*. Edwards et al. (2005) performed a SOWH test which rejected this hypothesis, using PAUP* (ML) and the model GTR+I+ Γ . The data set contains 6150 nucleotide characters, and is referred to here as Edwards.

Liu et al. (2012).—This data set includes 12 DNA loci, including 3 nuclear loci, 5 chloroplast loci, and 4 mitochondrial loci from 26 species of bryophytes (Liu et al. 2012). The alternative hypothesis to the maximum likelihood topology constrains the monophyly of *Physcomitrella*. Liu et al. (2012) performed a SOWH test which rejected this hypothesis using RAxML and the model GTR+I+ Γ , and incorporated a partitioning scheme which we did not utilize in the analyses presented here. The data set contains 44 taxa and 6657 nucleotide characters, and is referred to here as Liu.

Sullivan et al. (2000).—This data set is composed of mitochondrial sequences from 26 individuals of the rodent species *Reithrodontomys sumichrasti*. The alternative hypothesis to the maximum likelihood topology constrains phylogeographic groups similar to other rodents (Sullivan et al. 2000). Sullivan et al. (2000) performed SOWH tests using both maximum likelihood (under a GTR+I+ Γ model) and parsimony optimality criterion scores in PAUP*. The hypothesis was rejected using maximum likelihood scores but not using parsimony scores. The data set contains 1130 nucleotide characters, and is referred to here as Sullivan.

Wang et al. (2008).—This data set is composed of mitochondrial sequences from 41 individuals of the species of pygmy rain frog *Pristimantis ridens* and a single out-group taxon (Wang et al. 2008). The alternative hypothesis constrains a phylogeographic group of seven individuals. Wang et al. (2008) performed a SOWH test which rejected this hypothesis using PAUP* and fixed parameter values for maximum likelihood searches of simulated data sets. The data set contains 1672 nucleotide characters, and is referred to here as Wang.

Sensitivity to Number of Sample Replicates

We selected three of the seven data sets for an analysis of sensitivity to sample size. These data sets were selected because the results reported in the original study did not strongly reject the alternative hypothesis. For each of these data sets we performed 100 SOWH tests using a sample size of 100 each, and then subsequently performed 100 SOWH tests using a sample size of 500 each. We compared the range of *p*-values returned to the average confidence intervals calculated by SOWHAT at the two sample sizes. Each of these tests was performed using RAxML and the model GTR+ Γ .

Sensitivity to Choice of Likelihood Software

For all seven data sets we performed a test using GARLI and a test using RAxML and compared the results, with the exception of Dixon because the model HKY is not an option in RAxML. We compared the results of five of these tests to the originally reported results which were performed using PAUP* (ML). Our analyses were performed using the same model, sample size, and constraint topology as reported in the literature. All other program settings were the default. For our analyses, gaps were propagated into simulated data—treatment of gaps was mentioned only in the study by Edwards et al. (2005) where gaps were not present in simulated data. It is not clear how gaps were treated in the remaining data sets. We did not directly compare the results for Dunn, because the original study used parsimony to score data sets, or Liu, because the original study used a partitioning scheme not used for these analyses. Following consultation with the authors of RAxML, we reran the RAxML analyses with the Broyden-Fletcher-Goldfarb-Shanno (BFGS) routine suppressed (using the --nobfgs flag in RAxML) and compared the results.

Sensitivity to Treatment of Gaps

By default, SOWHAT propagates the same number and position of gaps present in the original data set into each simulated data set. This is accomplished by simulating a full data set and then subsequently removing the cells in matrix positions corresponding to gaps in the original data set. We analyzed the effect of not including gaps by suppressing this feature (i.e., simulated data matrices were complete, where originals were not) and comparing the resulting *p*-values to those calculated with gaps in simulated data. Four of the seven data sets analyzed in this study had gaps in the original data (Dunn, Edwards, Liu, and Wang). We used RAxML as the likelihood engine in all of these searches, and the same sample size as reported in the original study.

Sensitivity to Model Specification

We analyzed the effects of using a different model for likelihood inference from the model used for parameter estimation and data simulation (Fig. 2). For all seven data sets, we performed a SOWH test in which the likelihood score was calculated using a model with a reduced number of parameters free to vary (JC69) in comparison to the model used for parameter estimation and data simulation (GTR+I+ Γ). We compared the resulting *p*-values to those calculated using the originally reported model of evolution for both inference and data simulation. For these tests we used GARLI as the likelihood engine and the same sample size as reported in the original study. We then ran a SOWH test in which the originally reported model of evolution was used exclusively for likelihood inference and a model with

TABLE 1. Number of sample replicates

Data set	Samples	Tests	ML soft.	P-values			Min. conf. interval	
				Average	Lowest	Highest	Lower	Upper
Buckley	100	100	RAxML	0.261	0.140	0.312	0.114	0.185
Buckley	500	100	RAxML	0.263	0.219	0.264	0.171	0.211
Sullivan	100	100	RAxML	0.411	0.290	0.510	0.256	0.327
Sullivan	500	100	RAxML	0.401	0.344	0.458	0.329	0.360
Dixon	100	100	RAxML	0.092	0.030*	0.160	0.017	0.051
Dixon	500	100	RAxML	0.095	0.065	0.132	0.056	0.073

Notes: 100 SOWH tests were performed for three data sets with a sample size of 100, and 100 tests were performed with a sample size of 500. *P*-values for the Dixon data set at a sample size of 100 vary from 0.030 to 0.169, indicating repeated SOWH tests at this sample size could result in different outcomes using a significance level of 0.05. The minimum confidence interval is 0.017–0.051, indicating that all *p*-values which fall below 0.05 are accompanied by a confidence interval which spans the significance level, therefore more sampling is required. At a sample size of 500, all *p*-values and all confidence intervals fall entirely above the confidence level (the minimum interval is 0.056–0.073), indicating a sufficient sample size. * indicates *p*-value less than 0.05.

an increased number of free parameters (the CAT-GTR model found in PhyloBayes) was used for parameter estimation and data simulation. The resulting *p*-values were compared to those calculated using the originally reported model for both likelihood and inference. For these tests we used RAxML as the likelihood engine and the same sample size as reported in the original study.

Sensitivity to Generating Topology

We analyzed the effects of simulating data under the zero-constrained tree in place of a fully resolved tree estimated under the alternative hypothesis. In this method, following estimation of the most likely tree in an unconstrained search, SOWHAT uses the R package *hutan* (<https://bitbucket.org/caseywdunn/hutan>) to identify incongruent nodes and reduce them to a value near zero (here 1×10^{-6}). All other branch lengths are preserved, resulting in a partially resolved tree referred to as the zero-constrained tree (Susko 2014). For all seven data sets we compared the results of a SOWH test using a fully resolved generating topology to a test using the zero-constrained tree. We used RAxML as the likelihood engine in all of these searches, with the exception of Dixon where GARLI was used, and the same sample size as reported in the original study.

RESULTS AND DISCUSSION

Number of Sample Replicates

When an investigator applies the SOWH test they must decide how many simulated bootstrap samples to generate. Without sampling an adequate number of replicates, stochasticity can lead to poor estimates of the null distribution and different outcomes between repeated SOWH tests. Increasing the number of replicates, however, comes with an increased computational cost. Few previously reported SOWH

tests provide justification for sample size. A sufficient sample size depends on the data in question. To explore the effects of sample size, we ran multiple identical SOWH tests and compared the resulting *p*-values.

Our results indicate that a sample size of 100 is not sufficient for all data sets (Table 1). For the Dixon data set, repeated SOWH tests with 100 samples returned *p*-values ranging from 0.030 to 0.160, which results in a different interpretation of outcome at a significance level of 0.05. When the sample size was increased to 500, the test consistently returned a *p*-value greater than 0.05.

The confidence intervals calculated by SOWHAT provide an explicit tool for assessing the adequacy of sampling for a given test. With each new value added to the null distribution, SOWHAT recalculates the *p*-value and confidence intervals, allowing the adequacy of the sample size to be evaluated simultaneously with the results of the test. In the case of the SOWH tests on Dixon with 100 samples, the minimum confidence interval is 0.017–0.051, indicating that all *p*-values which fall below 0.05 are accompanied by a confidence interval which spans the significance level. This signals to the investigator that there is not adequate sampling to provide a robust interpretation of the outcome. As sample size increases, the confidence interval decreases, and at a sample size of 500 all *p*-values and confidence intervals fall above the level of significance for this data set, indicating that sample size is sufficient.

Choice of Likelihood Software

Application of the SOWH test also requires deciding which software to use for likelihood estimation. Our results indicate that, all else being equal, the choice of likelihood software between GARLI and RAxML does not affect the outcome of the SOWH test as long as a minor adjustment is made to how RAxML is run. Specifically, BFGS optimization of RAxML model parameters results in suboptimal performance on simulated data, but suppressing this optimization resolves the problem.

TABLE 2. Choice of likelihood software

Data set	Source	Samples	Model	ML software	<i>P</i> -value	Conf. interval	
						Lower	Upper
Buckley	New	1000	GTR+I+Γ	GARLI	0.018*	0.011	0.028
Buckley	New	1000	GTR+I+Γ	RAxML	0.010*	0.005	0.018
Buckley	Reported	1000	GTR+I+Γ	PAUP* (ML)	0.018*	–	–
Buckley	New	1000	GTR+Γ	GARLI	0.118	0.099	0.140
Buckley	New	1000	GTR+Γ	RAxML	0.252	0.225	0.280
Buckley	Reported	1000	GTR+Γ	PAUP* (ML)	0.015*	–	–
Dixon	New	500	HKY	GARLI	0.012*	0.004	0.026
Dixon	Reported	500	HKY	PAUP* (ML)	0.258	–	–
Dixon	New	500	HKY+Γ	GARLI	0.002**	<0.005	0.011
Dixon	Reported	500	HKY+Γ	PAUP* (ML)	<0.005**	–	–
Dunn	New	100	GTR+I+Γ	GARLI	<0.01**	<0.01	0.036
Dunn	New	100	GTR+I+Γ	RAxML	<0.01**	<0.01	0.036
Edwards	New	100	GTR+I+Γ	GARLI	<0.01**	<0.01	0.036
Edwards	New	100	GTR+I+Γ	RAxML	<0.01**	<0.01	0.036
Edwards	Reported	100	GTR+I+Γ	PAUP* (ML)	<0.01**	–	–
Liu	New	100	GTR+I+Γ	GARLI	<0.01**	<0.01	0.036
Liu	New	100	GTR+I+Γ	RAxML	<0.01**	<0.01	0.036
Sullivan	New	100	GTR+I+Γ	GARLI	<0.01**	<0.01	0.036
Sullivan	New	100	GTR+I+Γ	RAxML	<0.01**	<0.01	0.036
Sullivan	Reported	100	GTR+I+Γ	PAUP* (ML)	<0.01**	–	–
Wang	New	500	GTR+Γ	GARLI	<0.005**	<0.005	0.007
Wang	New	500	GTR+Γ	RAxML	<0.005**	<0.005	0.007
Wang	Reported	500	GTR+Γ	PAUP* (ML)	<0.005**	–	–

Notes: SOWH tests were performed using GARLI and RAxML for each data set, with the exception of Dixon et al. (2007) as HKY is not an option in RAxML. Each SOWH test was performed using the model, sample size, and constraint topology specified in the original performance of the test. Buckley and Dixon were analyzed using two different models of evolution, as reported originally. The resulting *p*-values were compared to those reported in the literature for five data sets. The other two data sets were not directly compared due to known differences in implementation; the SOWH test performed by Dunn et al. (2005) used parsimony to score tree; the test by Liu et al. (2012) was performed with a partition scheme not used here. The outcome of the tests differed from the literature for two data sets, Buckley using GTR+Γ and Dixon using HKY. * indicates *p*-values less than 0.05; ** indicates less than 0.01.

We compared the results of SOWH tests performed using the same data, model, and sample size, but different likelihood software (GARLI and RAxML) for six data sets (Table 2). For the Buckley data set, we performed this comparison using two different models for each likelihood software tool. Under BFGS optimization (the RAxML default), the choice of likelihood software had an effect on the outcome for two data sets, Sullivan and Wang (see Supplementary Table 1 in online Appendix 1 available on Dryad at <http://dx.doi.org/10.5061/dryad.c80p7>). Rerunning the tests using RAxML with the optimization suppressed resulted in universally smaller ranges of null distribution and the outcomes of the tests run under RAxML were the same as those under GARLI (Table 2).

We also compared the results of the SOWH test reported in the literature to the results obtained here for five of these seven data sets (Table 2). Each of the tests compared was performed using the same hypothesis, model, and sample size as originally reported, but differed in likelihood software—all of the reported tests were previously performed using PAUP*. In addition to the previously stated difference in results of the Sullivan and Wang tests, two other tests returned

outcomes different than those previously reported. For Buckley using the model GTR+Γ, the SOWH tests performed using GARLI and RAxML both failed to reject a hypothesis reported to be rejected in the original study ($P < 0.05$). For Dixon, the SOWH test performed using GARLI and the model HKY rejected a hypothesis ($P < 0.05$) that was not rejected in the original study. As the reported SOWH tests were not performed using SOWHAT, we cannot determine whether these results are due to differences in performance of likelihood software or other aspects of test implementation. We did not directly compare the results of Dunn or Liu due to known differences in implementation beyond likelihood software—despite these difference, for both data sets the outcome of the SOWH tests performed here were consistent with those originally reported.

Treatment of Gaps

Multiple sequence alignments include gaps due to both missing data and the accommodation of insertions and deletions. These gaps are often not included in the simulated bootstrap samples. The purpose of simulating

TABLE 3. Treatment of gaps

Data set	% Gaps in data set	Simulate gaps	P-value	Conf. interval	
				Lower	Upper
Dunn	14.759	yes	<0.01**	<0.01	0.036
Dunn	14.759	no	<0.01**	<0.01	0.036
Edwards	14.268	yes	<0.01**	<0.01	0.036
Edwards	14.268	no	<0.01**	<0.01	0.036
Liu	9.149	yes	<0.01**	<0.01	0.036
Liu	9.149	no	<0.01**	<0.01	0.036
Wang	13.092	yes	0.026*	0.014	0.044
Wang	13.092	no	0.022*	0.011	0.039

Notes: SOWHAT by default propagates the exact number and position of gaps present in the original data set into all simulated data sets. Suppressing this feature, thereby excluding gaps from subsequent analyses, did not change the outcome of any SOWH tests examined here. * indicates *p*-values less than 0.05; ** indicates less than 0.01.

data sets under parametric conditions is to recreate the situation under which the real-world data was generated—ideally this would recreate the presence of gaps as well. SOWH tests performed where gaps are present only in observed data sets and not in simulated data sets, have an increased amount of information in simulated data sets, which has the potential to affect likelihood analyses on these data sets and therefore affect the outcome of the test.

SOWHAT simulates the data without gaps and subsequently replaces determined sites with gaps based on the number and positions of gaps in the original data set. While this does not explicitly model the processes that led to the original gaps, this method ensures that taxa which were originally poorly sampled and which may be in question will be poorly sampled in all simulated data sets as well.

We tested the effects of gaps by suppressing this feature (i.e., all simulated data sets were completely full). Our results indicate that, for the four data sets tested here, gaps had no major effect on the outcome of the SOWH test (Table 3). Excluding these gaps from simulated data did not change the outcome of the test for any of these data sets.

Model Specification

The choice of model used for likelihood inference and for simulating data sets has previously been shown to have an impact on the outcome of the SOWH test (Goldman et al. 2000; Buckley 2002). Buckley (2002) suggest that when the likelihood search is being performed with the exact same model used to generate the data it will be prone to recover the generating topology, therefore δ values will be smaller and the null distribution will be more narrow. To minimize type I error, they suggest that the model used for parameter estimation should include a greater number of free

TABLE 4. Model specification: JC69 analysis

Data set	Models		P-value	Conf. interval	
	ML score	Parameters		Lower	Upper
Buckley	GTR+G	GTR+G	0.241	0.215	0.269
Buckley	JC69	GTR+G+I	<0.001**	<0.001	0.004
Dixon	HKY+G	HKY+G	0.012*	0.004	0.026
Dixon	JC69	GTR+G+I	<0.005**	<0.005	0.007
Dunn	GTR+G+I	GTR+G+I	<0.01**	<0.01	0.036
Dunn	JC69	GTR+G+I	<0.01**	<0.01	0.036
Edwards	GTR+G+I	GTR+G+I	<0.01**	<0.01	0.036
Edwards	JC69	GTR+G+I	<0.01**	<0.01	0.036
Liu	GTR+G+I	GTR+G+I	<0.01**	<0.01	0.036
Liu	JC69	GTR+G+I	<0.01**	<0.01	0.036
Sullivan	GTR+G+I	GTR+G+I	0.290	0.204	0.389
Sullivan	JC69	GTR+G+I	<0.01**	<0.01	0.036
Wang	GTR+G+I	GTR+G+I	0.026*	0.014	0.044
Wang	JC69	GTR+G+I	<0.005**	<0.005	0.007

Notes: Model 1 represents the model used for likelihood inference (i.e., searching and scoring both the original and simulated data sets). Model 2 was used to estimate parameter values and simulate data sets. Separating the models for scoring and simulation has been suggested as a correction for a liberal bias present in the SOWH test. Using a model for scoring with fewer parameters free to vary, such as JC69, here resulted in a more liberal test. All hypotheses were rejected when JC69 was used as Model 1. * indicates *p*-values less than 0.05; ** indicates less than 0.01.

parameters than the model used for likelihood inference (Fig. 2).

We evaluated the effect of this adjustment using two different methods. First, for seven data sets we compared the results of SOWH tests performed using the JC69 model for likelihood inference and GTR+I+ Γ for data simulation (Table 4) to tests using the same model for likelihood inference and data simulation. Our results show that, contrary to what was previously suspected, using distinct models resulted in lower δ values and a potentially higher type I error rate. For every test performed where JC69 was used for likelihood inference the *p*-value returned was <0.01 and the hypothesis was rejected, including those tests that did not reject the hypothesis when the models used for likelihood inference and data simulation were the same.

Second, for six data sets we also compared the results of SOWH tests when parameters were optimized and data simulated under the PhyloBayes CAT-GTR model and a simpler model used for inference (Table 5). In this analysis, parameters were estimated using the posterior predictive framework under the CAT-GTR model and simulated using these parameters using PhyloBayes. The CAT-GTR model introduces the additional free parameter of site-specific heterogeneity. PhyloBayes was used strictly for parameter estimation and data simulation; data sets were scored using RAXML and likelihood scores were used for the statistical test.

Our results show that this method can change the outcome of the test in both directions, depending on the data in question. For Buckley and Sullivan, the SOWH

TABLE 5. Model specification: CAT analysis

Data set	Models		<i>P</i> -value	Conf. interval	
	ML score	Parameters		Lower	Upper
Buckley	GTR+ Γ	GTR+ Γ	0.241	0.215	0.269
Buckley	GTR+ Γ	CAT	0.013*	0.007	0.022
Dunn	GTR+I+ Γ	GTR+I+ Γ	<0.01**	<0.01	0.036
Dunn	GTR+I+ Γ	CAT	0.080	0.035	0.152
Edwards	GTR+I+ Γ	GTR+I+ Γ	<0.01**	<0.01	0.036
Edwards	GTR+I+ Γ	CAT	<0.01**	<0.01	0.036
Sullivan	GTR+I+ Γ	GTR+I+ Γ	0.290	0.204	0.389
Sullivan	GTR+I+ Γ	CAT	0.030*	0.006	0.085
Liu	GTR+I+ Γ	GTR+I+ Γ	<0.01**	<0.01	0.036
Liu	GTR+I+ Γ	CAT	<0.01**	<0.01	0.036
Wang	GTR+G	GTR+G	0.026*	0.014	0.044
Wang	GTR+G	CAT	0.020*	0.010	0.036

Notes: Model 1 and Model 2 are the same as described in Table 4. Using a model for simulation with a greater number of parameters free to vary, such as the CAT model of PhyloBayes, did not result in universally larger δ values and therefore a more conservative test, though this was true for one data set, Dunn. The outcome of two other tests also differed, for Buckley and Sullivan, but the result was a more liberal test. * indicates *p*-values less than 0.05; ** indicates less than 0.01.

test performed using the CAT–GTR model rejected a hypothesis which was not rejected using the same model for inference and simulation ($P > 0.05$). Conversely, for Dunn the SOWH test performed using the CAT–GTR model failed to reject a hypothesis which was rejected using the same model for inference and simulation ($P < 0.05$). These results do not support the hypothesis that using distinct models will universally result in higher δ values and a more conservative test. We find no consistent evidence to support the suggestion that the model for inference should be distinct from the model for simulation.

It has also been suggested that, to reduce the rate of type I error, parsimony might be used in place of likelihood inference to score the data sets (Buckley 2002). Parsimony has often been used to reduce the computational burden of running multiple likelihood searches (Dunn et al. 2005). However, using parsimony in the context of the SOWH test is problematic. Often the purpose of the test is to determine whether the most likely hypothesis is significantly more likely than some *a priori* hypothesis. It is quite possible that the *a priori* hypothesis may not be significantly less parsimonious and still be significantly less likely, as seems to be the case in the two SOWH tests reported in Sullivan et al. (2000). Given the improved speed in likelihood software, parsimony should no longer be used in place of likelihood in the context of the SOWH test.

Generating Topology

The SOWH test, as presented in Goldman et al (2000), generates new data sets using the parameter values and tree topology estimated under the alternative constraint

TABLE 6. Generating topology

Data set	Model	Generating tree	<i>P</i> -value	Conf. interval	
				Lower	Upper
Buckley	GTR+I+ Γ	Fully Resolved	0.025*	0.016	0.037
Buckley	GTR+I+ Γ	Zero-constrained	0.123	0.103	0.145
Buckley	GTR+ Γ	Fully Resolved	0.241	0.215	0.269
Buckley	GTR+ Γ	Zero-constrained	0.439	0.408	0.470
Dixon	HKY+ Γ	Fully Resolved	0.012*	0.004	0.026
Dixon	HKY+ Γ	Zero-constrained	0.032*	0.018	0.051
Dixon	HKY	Fully Resolved	0.002**	<0.005	0.011
Dixon	HKY	Zero-constrained	<0.005**	<0.005	0.007
Dunn	GTR+I+ Γ	Fully Resolved	<0.01**	<0.01	0.036
Dunn	GTR+I+ Γ	Zero-constrained	<0.01**	<0.01	0.036
Edwards	GTR+I+ Γ	Fully Resolved	<0.01**	<0.01	0.036
Edwards	GTR+I+ Γ	Zero-constrained	<0.01**	<0.01	0.036
Liu	GTR+I+ Γ	Fully Resolved	<0.01**	<0.01	0.036
Liu	GTR+I+ Γ	Zero-constrained	<0.01**	<0.01	0.036
Sullivan	GTR+I+ Γ	Fully Resolved	0.290	0.204	0.389
Sullivan	GTR+I+ Γ	Zero-constrained	0.240	0.160	0.336
Wang	GTR+ Γ	Fully Resolved	0.026*	0.014	0.044
Wang	GTR+ Γ	Zero-constrained	0.034*	0.020	0.054

Notes: We compared SOWH tests performed using a fully resolved generating topology to tests performed using the zero-constrained tree, as suggested by Susko (2014). The zero-constrained tree is created by manipulating the most likely unconstrained tree so that edges incongruent with the alternative hypothesis are reduced to nearly zero. Using this method changed the outcome of only one test, Buckley, using the model GTR+I+ Γ . * indicates *p*-values less than 0.05; ** indicates less than 0.01.

hypothesis. Susko (2014) suggest that using a fully resolved tree as the generating topology will result in a test that is too conservative and that returns lower *p*-values than it should. They suggest instead using a third tree, referred to as the zero-constrained tree, as the generating topology.

We evaluated this adjustment by comparing the results of SOWH tests performed using a fully resolved tree to the results of tests performed using the zero-constrained tree for all seven data sets included in this study (Table 6). Our results show that using the zero-constrained tree results in very similar *p*-values for nearly all of the tests in question. The outcome of the test was different for one data set, Buckley when using the model GTR+I+ Γ , where the test using the zero-constrained tree did not reject the alternative hypothesis which was rejected using a fully resolved generating tree ($P > 0.05$).

These results suggest that using the zero-constrained tree can result in higher *p*-values, but that the impact is minimal and for many data sets the outcome of the test is unchanged.

CONCLUSION

The SOWH test is of great value to the phylogenetic community. However, it can be burdensome to perform

manually and requires making multiple decisions that have the potential to affect the outcome of the test. SOWHAT eliminates the burden of manual implementation of the SOWH test. Using SOWHAT to evaluate the impact of these decisions, we provide the following recommendations for implementation of the SOWH test. Where possible, these recommendations have been adopted as the default behavior of SOWHAT.

The number of sample replicates in a SOWH test should be explicitly justified. Our results indicate that a sample size of 100 is not adequate for all data sets. SOWHAT calculates confidence intervals around the p -values, and our results indicate that these confidence values capture a reasonable amount of the observed variation in repeated tests. Following a minimum number of samples—we recommend 100—an investigator should evaluate the adequacy of sample size by determining whether the confidence intervals fall entirely on one side of the significance level.

We found that the choice of likelihood software did not depend on whether GARLI or RAxML were used for likelihood analyses, as long as BFGS optimization was suppressed in RAxML. Given that RAxML performs the likelihood searches much faster than GARLI, RAxML with the BFGS routine silenced is the default likelihood engine in SOWHAT. GARLI can be used if additional models are required.

Our results indicate that excluding gaps from simulated data is unlikely to change the overall outcome of a SOWH test, however this effect may depend on the data in question. Since there is little computational cost and strong biological justification, we recommend simulating data with gaps.

Model misspecification is a problem universal to parametric bootstrap tests, which rely on the model to both simulate and evaluate the data. It has been suggested that using a model for likelihood inference that is distinct from the model used for data simulation may result in lower type I error rate (Buckley 2002). Our results are not consistent with this hypothesis—we found that using a simpler model for inference compared to the model used for simulation does not result in universally higher δ values. We therefore suggest that the same model is used for inference and simulation.

It has also been suggested that simulating new data under a tree which is not fully resolved, such as the zero-constrained tree will result in lower type I error (Susko 2014). We find that simulating data on the zero-constrained tree can impact results, but that for many data sets using the zero-constrained tree will not change the outcome of the test. Given the well-established concern of a high type I error rate along with the low computation cost in implementation, we recommend simulating data under the zero-constrained tree.

In all statistical tests such as the SOWH test, the recommended approach is to perform all reasonable permutations of the test and to report the highest p -value. SOWHAT greatly facilitates running multiple SOWH tests with a range of parameters and comparing the effects on the final result.

SUPPLEMENTARY MATERIAL

Data available from the Dryad Digital Repository: <http://dx.doi.org/10.5061/dryad.c80p7>.

FUNDING

S.H.C. was funded by the LINK Award through Brown University and the Marine Biological Laboratory at Woods Hole, as well as the Sars International Centre for Marine Molecular Biology. J.F.R. was funded by The Whitney Laboratory for Marine Bioscience and Sars International Centre for Marine Molecular Biology. Analyses were conducted with computational resources and services at the Center for Computation and Visualization at Brown University, supported in part by the National Science Foundation EPSCoR EPS-1004057 and the State of Rhode Island.

ACKNOWLEDGMENTS

We thank Nicholas Sinnott-Armstrong, Felipe Zapata, Mark Howison, Peisi Yan, Stefan Siebert, Alexis Stamatakis, and Andreas Hejnol for comments, ideas, and support. We thank the editors, Frank Anderson and Mark Holder, as well as the anonymous reviewers for the most thorough, helpful, and longest reviews we have ever seen. Those acknowledged do not necessarily endorse the findings in this paper.

REFERENCES

- Anderson J., Goldman N., Rodrigo A. 2014. Guidelines for performing the SOWH test. Available from URL: <http://www.ebi.ac.uk/goldman/tests/sowhinstr.html>.
- Bergsten J., Nilsson A.N., Ronquist F. 2013. Bayesian tests of topology hypotheses with an example from diving beetles. *Syst. Biol.* 62: 660–673.
- Buckley T.R. 2002. Model misspecification and probabilistic tests of topology: evidence from empirical data sets. *Syst. Biol.* 51:509–523.
- Crawford A.J. 2009. horribly detailed instructions to running a parametric bootstrap test. Available from URL: <http://dna.ac/genetics.html>.
- Dixon C.J., Schoenwetter P., Schneeweiss G.M. 2007. Traces of ancient range shifts in a mountain plant group (androsace halleri complex, primulaceae). *Mol. Ecol.* 16:3890–3901.
- Dunn C.W., Pugh P.R., Haddock S.H. 2005. Molecular phylogenetics of the siphonophora (cnidaria), with implications for the evolution of functional specialization. *Syst. Biol.* 54:916–935.
- Edwards E.J., Nyffeler R., Donoghue M.J. 2005. Basal cactus phylogeny: implications of pereskia (cactaceae) paraphyly for the transition to the cactus life form. *Am. J. Bot.* 92:1177–1188.
- Goldman N., Anderson J.P., Rodrigo A.G. 2000. Likelihood-based tests of topologies in phylogenetics. *Syst. Biol.* 49:652–670.
- Kishino H., Hasegawa M. 1989. Evaluation of the maximum likelihood estimate of the evolutionary tree topologies from dna sequence data, and the branching order in hominoidea. *J. Mol. Evol.* 29: 170–179.
- Lartillot N., Lepage T., Blanquart S. 2009. Phylobayes 3: a bayesian software package for phylogenetic reconstruction and molecular dating. *Bioinformatics* 25:2286–2288.
- Liu Y., Budke S., Goffinet B. 2012. Phylogenetic inference rejects sporophyte based classification of the funariaceae (bryophyta): rapid

- radiation suggests rampant homoplasy in sporophyte evolution. *Mol. Phylogenet. Evol.* 62:130–145.
- Pauly G.B., Hillis D.M., Cannatella D.C. 2004. The history of a nearctic colonization: molecular phylogenetics and biogeography of the nearctic toads (*bufo*). *Evolution* 58:2517–2535.
- Rambaut A., Grass N.C. 1997. Seq-gen: an application for the Monte Carlo simulation of DNA sequence evolution along phylogenetic trees. *Comput. Appl. Biosci.* 13:235–238.
- Shimodaira H. 2001. Multiple comparisons of log-likelihoods and combining nonnested models with applications to phylogenetic tree selection. *Commun. Stat.-Theor. M.* 30:1751–1772.
- Shimodaira H., Hasegawa M. 1999. Multiple comparisons of log-likelihoods with applications to phylogenetic inference. *Mol. Biol. Evol.* 16:1114–1116.
- Stamatakis A. 2006. Raxml-vi-hpc: maximum likelihood-based phylogenetic analyses with thousands of taxa and mixed models. *Bioinformatics* 22:2688–2690.
- Sullivan J., Arellano E., Rogers D.S. 2000. Comparative phylogeography of mesoamerican highland rodents: concerted versus independent response to past climatic fluctuations. *Am. Nat.* 155:755–768.
- Sullivan J., Holsinger K.E., Simon C. 1995. Among-site rate variation and phylogenetic analysis of 12s rRNA in sigmodontine rodents. *Mol. Biol. Evol.* 12:988–1001.
- Susko E. 2014. Tests for two trees using likelihood methods. *Mol. Biol. Evol.* 31:1029–1039.
- Swofford D.L., Olsen G.J., Waddell P.J., Hillis D.M. 1996. Phylogenetic inference. In: Hillis D.M., Moritz C., Mable B.K., editors. *Molecular systematics*. 2nd ed. Sunderland, MA: Sinauer. p. 407–514.
- Wang I.J., Crawford A.J., Bermingham E. 2008. Phylogeography of the pygmy rain frog (*Pristimantis ridens*) across the lowland wet forests of isthmian Central America. *Mol. Phylogenet. Evol.* 47:992–1004.
- Zwickl D. 2006. Garli: genetic algorithm for rapid likelihood inference. Available from URL: <http://www.bio.utexas.edu/faculty/antisense/garli/Garli.html>.