

Published in final edited form as:

Syst Biol. 2014 September ; 63(5): 838–842. doi:10.1093/sysbio/syu041.

Testing for Universal Common Ancestry

Leonardo de Oliveira Martins and David Posada

Department of Biochemistry, Genetics and Immunology, University of Vigo, Vigo 36310, Spain. Tel +34 986 818 681, Fax: +34 986 812 556

Keywords

model selection; common ancestry; tree of life; UCA; independent origins; IO; AIC; homology

A phylogenetic model selection test to quantify the evidence for the Universal Common Ancestry (UCA) of life forms was proposed recently (Theobald 2010), based on the comparison of the statistical support, using likelihoods, the Akaike Information Criterion or Bayes Factors, for two different phylogenetic models representing the UCA and the Independent Origins (IO) hypotheses (Sober and Steel 2002). In this test, the former is represented by a single phylogeny connecting all sequences, while the latter is depicted by several, independent phylogenetic trees (Figure 1). Importantly, in the original UCA test the same alignment was used to represent both hypotheses. When applied to a particular data set of 23 universally conserved proteins, the test strongly favored a UCA scenario.

While there is no question of the common ancestry of the particular set of aligned sequences analyzed, since its publication several criticisms of the test have been raised. Yonezawa and Hasegawa (2010) showed how the UCA test failed to detect that the mitochondrial genes *cytb* and *nd2* are not homologous, to which Theobald replied that when the test is applied to codon or protein models, as originally devised, then the IO hypothesis is correctly preferred (Theobald 2010b). More recently, the same authors extended their analysis and commented on the possible failure of the test for cases of convergence towards similar amino acid composition (Yonezawa and Hasegawa 2012). Koonin and Wolf (2010) simulated alignments lacking phylogenetic structure (site columns came from an independent distribution of amino acid frequencies) and showed that the test would spuriously favor UCA, probably because it was misled by column-wise similarity. In a recent reply, Theobald (2011) included the model used to simulate K&W's data, which was indeed preferred over the UCA model. In his reply, he also suggested that Koonin and Wolf's simulations "were produced by a well-known common ancestry model", which we believe is incorrect because the IO model described by Theobald (2010) corresponds mathematically, in the limit, to a tree with at least one infinite branch length (see Supplementary Text S1, doi:). We also pointed out (Martins and Posada 2012) that the original UCA analysis was affected by selection bias: the query data consisted of sequences already subjected to a similarity search (e.g. BLAST) whose putative column-wise homology status had then been optimized by an alignment algorithm (Brown et al. 2001). In addition, we showed that under the

representation of UCA and IO as one versus multiple phylogenies, we can easily distinguish sequences simulated under UCA vs IO by simply observing similarity measures, concluding that similarity should not be used to select which data sets are eligible for the UCA test.

In this point of view, we demonstrate a fundamental drawback of the original UCA test, which is the use of the same alignment to represent both the UCA and IO hypotheses. The UCA test uses the standard phylogenetic likelihood (L), which is the probability (P) of the 'aligned sequences' (D) given a phylogenetic hypothesis (H ; which is UCA or IO), $L = P(D|H)$. Phylogenetic studies usually consider alignments as raw data (D) and so there is an underlying assumption that all sites from a column are homologous. However, in reality, the unaligned sequences are the true raw data and the fixed alignment should be recognized as a point estimate of the homology relations (Kumar and Filipski 2007; Wong et al. 2008). In any case, in order to make the competing model likelihoods comparable, they have to be based on the same data, which in the original UCA test translates into using a single, fixed global sequence alignment to represent both UCA and IO, even if the global alignment is later split for the calculation of the IO model likelihood. Given the global homology assumption made by multiple alignment programs (Meng et al. 2011; Varon and Wheeler 2012), the possibility that a fixed alignment could bias the test towards UCA has been raised before (Yonezawa and Hasegawa 2010, Theobald 2011), although never demonstrated. In fact, in phylogenetics, alignment algorithms try to optimize the data to conform to a common ancestry hypothesis, and many even use a guide tree, like ClustalW (Thompson et al. 1994) which was the program utilized to align the protein sequences studied by Theobald (2010). In order to better understand the performance of the UCA test, we carried out the simulation study described next, followed by a proposed solution that might alleviate the bias.

UCA test performance under simulated IO

We simulated sets of sequences evolving under the IO hypothesis, using parameter values estimated from the data in Theobald (2010), a concatenated data set of 6591 sites of 4 eukaryotic (E) and 4 bacterial (B) sequences. We used INDELible (Fletcher and Yang 2009) to simulate protein evolution (without indels) independently along the E and B trees under the best-fit amino-acid replacement models in Theobald (2010) (rtREV+GF), forming two sets of 4 sequences (quartets) per simulated data set. Both quartets were then grouped together into a single data set composed of 8 sequences.

Next, we carried out the UCA test for the simulated data sets. We aligned all sequences with MUSCLE (Edgar 2004) and estimated the AIC scores for the UCA and IO models as described by Theobald (2010). These alignments were not subjected to further processing such as removal of gapped columns or regions of low quality, and presented between 7 and 11% of gaps. In Figure 2, we show the results for 400 simulated replicates (200 under each simulation model), where $AIC = AIC(B) + AIC(E) - AIC(BE)$, such that positive values for AIC favor UCA. Clearly, we can see that the UCA hypothesis is incorrectly preferred by a large margin in all simulated datasets.

To investigate whether this bias was caused by the alignment, we implemented the UCA test without the alignment step. As explained above, in the original UCA test, the likelihoods were calculated upon the aligned sequences, so an alignment is the minimum input requirement. As expected, if the alignment operation is not performed (though indels were not simulated so total sequence length was conserved) the test ‘correctly’ favors IO (Figure 2). Obviously, nobody would (or should) carry out in practice such a phylogenetic test, without aligning the sequences, but this experiment served here to demonstrate that the fixed alignment of the UCA test biases the outcome towards UCA. For the conditions described in Theobald (2010) and replicated here, the UCA test has a false positive rate of 100% in our simulations. Our simulations showed that even if one aligns the E and B subsets independently under the IO hypothesis on one hand and the B+E sequences under UCA on the other hand, the AIC (or AICc, or BIC) values would still favor UCA (data not shown, scripts available as Supplementary material, [doi:10.5061/dryad.gn376](https://doi.org/10.5061/dryad.gn376)), although we reprove this procedure because the likelihoods compared do not correspond to the same data. This predilection of the test for UCA is due to the fact that the alignment optimization allows for the B+E sequences to have a much better AIC than their unaligned counterparts, at the cost of adding less than 11% of indels.

Moreover, our re-analysis of previously published data sets purportedly showing the original UCA test favoring IO (Theobald 2010; 2011) indicates that under proper conditions the UCA hypothesis is in fact spuriously preferred (section S2 of the Supplementary Text). Not surprisingly, given its bias towards UCA, the test always correctly favored the UCA hypothesis for alignments simulated under common ancestry.

We did find other scenarios where the UCA test ‘correctly’ favored IO (results not shown) for the wrong reason, like simulating each life domain under a different amino acid replacement model - which suggests that, in this case, the UCA test is in fact identifying the misspecification of the amino acid replacement model. This implies that whenever the UCA test favors IO we should further analyze the data before making a decision, since it may not distinguish IO from certain amino acid replacement heterogeneities -- an issue already highlighted in Theobald (2010).

Reducing the false positive rate of the UCA test

If we want to reduce the bias towards UCA induced by the alignment step, we should work with the unaligned sequences as our primary data, in order to obtain likelihood values associated to the raw sequences. One way of doing this is estimating the alignment and the phylogeny at the same time (Lunter et al. 2005; Fleissner et al. 2005; Redelings and Suchard 2005, 2007; Novák et al. 2008). Under this framework, the data (D) are the (unaligned) sequences, while the alignment is one of the parameters of the model, to be treated as an unknown random variable. This type of model is implemented, for example, in the program BAli-phy (Redelings and Suchard 2005), that not only accounts for substitutions but also explicitly models indels. Therefore the likelihood values are very different from those obtained by standard phylogenetic models. In order to evaluate the performance of this approach, we simulated protein sequences of 500 amino acids under IO exactly as described before, but this time conducting the test with BAli-Phy instead of MUSCLE + ProtTest +

Phyml (under BAli-Phy, the alignment optimization program is redundant). We used BAli-Phy to jointly estimate the posterior distribution of alignments, branch lengths and of the shape parameter of the gamma distribution for rate variation among sites assuming the LG +G (Le and Gascuel 2008) model under a fixed tree topology with variable branch lengths. For each replicate we ran the software 3 times: once for each domain (E and B) independently (the product of these two analyses give us the likelihood for the IO model), and once for the 8-sequence E+B data set (which give us the likelihood for the UCA model). Although BAli-Phy can also sample from the space of phylogenies, we fixed the topologies at their true values (the ones used in the simulation) and allowed only the branch lengths to vary in the interests of straightforward computation.

We used the marginal likelihoods, calculated as the harmonic mean of the sample likelihoods (Kass and Raftery 1995), in order to estimate the Bayes factor between the UCA and IO hypotheses. Notice that for each replicate we will have an alignment distribution for B only, then one for E only and finally one for B+E, together with their respective marginal likelihoods $P(B)$, $P(E)$ and $P(B+E)$. Therefore we have $BF = \log[\text{Prob}(D/UCA)] - \log[\text{Prob}(D/IO)] = \log[P(B+E)] - \log[P(B)] - \log[P(E)]$, such that positive values support UCA. In Figure 3 we show the results from 100 replicates, where we can see that 86% of the simulations were correctly identified as supporting IO, 12% favored UCA, and two simulations were inconclusive. Figure 3a shows the histogram with BF values normalized per site -- that is, divided by the posterior median alignment length -- while Figure 3b plots the raw BF values against the posterior median total tree length. Note that there's no apparent correlation between tree length and support for UCA. Here, we must note that these Bayes factors should not be taken at face value: the harmonic mean estimator (HME) is numerically unstable and tends to favor more complex models, and although better estimators exist, they are not implemented yet in most Bayesian phylogenetic software (Lartillot and Philippe 2006; Xie et al. 2010). The HME also tends to overestimate the marginal likelihood, which will favor IO more easily (Lartillot and Philippe 2006). In any case, we believe that these results clearly suggest that considering alignment and phylogeny coestimation should reduce to a large extent the bias towards UCA evidenced by the original UCA test.

Discussion

We have shown that the UCA test described in Theobald (2010) is unable to detect the independent origins of two sets of unrelated sequences. While our simulations are not exhaustive -- we did not explore many possible combinations of trees, branch lengths, sequence sizes, and evolutionary models for instance -- they show that there are many cases not unlike real data sets where the UCA test fails. Our general impression is that the original UCA test would not reject a common origin for any but obviously unrelated set of sequences. Certainly, one can argue that for a specific, particular data set the UCA test has worked. But the high 'quality' of the original data set should not be used to justify the correctness of the method. We have previously noted (Martins and Posada 2012) that selecting the sequences based on similarity can make the alignment bias disappear due to the lower number of introduced indels, but then this selection procedure clearly introduces its own bias.

Theobald (2011) offered a few suggestions for situations when we are not very confident about the alignment. The first was to use structural alignments, which might be a promising approach in the future but depends on the ability of structurally aligning simulated or empirical independent sequences of arbitrary similarity. The second was to account for ‘alignment bias and uncertainty’, which according to our simulations is in fact a prerequisite if the UCA test is to be applied as devised. Moreover, we believe that any formal attempt to quantify the UCA hypothesis must take into account the selection and alignment of sequences into the test. The third suggestion was a permutation procedure whereby sites for certain sequences are shuffled followed by recalculation of the AICs after realignment. This would tell us by how much the original data departs from data sets whose phylogenetic structure has been partially removed. However, using AIC to compare different data sets is not a valid approach. Therefore, AIC values between distinct alignments can not be interpreted in probabilistic terms. Still, this procedure can lead to a permutation test (similar to the permutation tail probability tests of Faith and Cranston 1991; Swofford et al., 1996), in which a wide collection of test statistics can be used in place or in addition to the AIC.

The full BAli-Phy analysis on each of the 500 sites replicate took more than one week on a single thread, even assuming a fixed topology, restricting right now these type of analyses to small data sets. In any case, any data set must be aligned to be amenable to the original UCA test, and here we have demonstrated that by doing so the test will often favor UCA. We want to emphasize again that we are not denying the common ancestry of the data set analyzed in Theobald (2010). What we and others have been pointing out are shortcomings of the UCA test itself.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank Mateus Patrício and Ramon Fallon for helping us with the large scale analyses and discussing the manuscript, and Douglas Theobald for a fruitful email exchange since December 2010. We also appreciate the careful reading and guidance by the anonymous reviewers.

Funding

This work was supported by the European Research Council (ERC-2007-Stg 203161-PHYGENOM to D.P.).

Bibliography

- Brown JR, Douady CJ, Italia MJ, Marshall WE, Stanhope MJ. Universal trees based on large combined protein sequence data sets. *Nature genetics*. 2001; 28(3):281–285. http://www.nature.com/ng/journal/v28/n3/abs/ng0701_281.html. [PubMed: 11431701]
- Darriba, Diego; Taboada, Guillermo L.; Doallo, Ramón; Posada, David. ProtTest 3: fast selection of best-fit models of protein evolution. *Bioinformatics* (Oxford, England). 2011 Apr 15; 27(8):1164–5. DOI: 10.1093/bioinformatics/btr088
- Edgar, Robert C. MUSCLE: multiple sequence alignment with high accuracy and high throughput. *Nucleic acids research*. 2004 Jan; 32(5):1792–7. DOI: 10.1093/nar/gkh340 [PubMed: 15034147]

- Faith, Daniel P.; Cranston, Peter S. Could a Cladogram This Short Have Arisen By Chance Alone?: on Permutation Tests for Cladistic Structure. *Cladistics*. 1991 Mar; 7(1):1–28. DOI: 10.1111/j.1096-0031.1991.tb00020.x
- Fleissner, Roland; Metzler, Dirk; von Haeseler, Arndt. Simultaneous statistical multiple alignment and phylogeny reconstruction. *Systematic biology*. 2005 Aug; 54(4):548–61. DOI: 10.1080/10635150590950371 [PubMed: 16085574]
- Fletcher, William; Yang, Ziheng. INDELible: a flexible simulator of biological sequence evolution. *Molecular biology and evolution*. 2009 Aug; 26(8):1879–88. DOI: 10.1093/molbev/msp098 [PubMed: 19423664]
- Guindon, Stéphane; Dufayard, Jean-François; Lefort, Vincent; Anisimova, Maria; Hordijk, Wim; Gascuel, Olivier. New algorithms and methods to estimate maximum-likelihood phylogenies: assessing the performance of PhyML 3.0. *Systematic biology*. 2010 May 1; 59(3):307–21. DOI: 10.1093/sysbio/syq010 [PubMed: 20525638]
- Kass, Robert E.; Raftery, Adrian E. Bayes Factors. *Journal of the American Statistical Association*. 1995 Jun; 90(430):773. doi: 10.2307/2291091
- Koonin, Eugene V.; Wolf, Yuri I. The common ancestry of life. *Biology direct*. 2010 Jan; 5(1):64. doi: 10.1186/1745-6150-5-64 [PubMed: 21087490]
- Kumar, Sudhir; Filipinski, Alan. Multiple sequence alignment: In pursuit of homologous DNA positions. *Genome Research*. 2007; 17(2):127–135. DOI: 10.1101/gr.5232407 [PubMed: 17272647]
- Lartillot N, Philippe H. Computing Bayes factors using thermodynamic integration. *Systematic biology*. 2006; 55(2):195–207. DOI: 10.1080/10635150500433722 [PubMed: 16522570]
- Le, Si Quang; Gascuel, Olivier. An improved general amino acid replacement matrix. *Molecular biology and evolution*. 2008 Jul; 25(7):1307–20. DOI: 10.1093/molbev/msn067 [PubMed: 18367465]
- Lunter, Gerton; Miklós, István; Drummond, Alexei; Jensen, Jens Ledet; Hein, Jotun. Bayesian coestimation of phylogeny and sequence alignment. *BMC bioinformatics*. 2005 Jan; 6:83. doi: 10.1186/1471-2105-6-83 [PubMed: 15804354]
- Martins, Leonardo de Oliveira; Posada, David. Proving universal common ancestry with similar sequences. *Trends in Evolutionary Biology*. 2012; 4:e5. doi: 10.4081/eb.2012.e5 [PubMed: 23814665]
- Meng, Lu; Sun, Fengzhu; Zhang, Xuegong; Waterman, Michael S. Sequence alignment as hypothesis testing. *Journal of computational biology*. 2011 May; 18(5):677–91. DOI: 10.1089/cmb.2010.0328 [PubMed: 21554016]
- Novák, Adám; Miklós, István; Lyngsø, Rune; Hein, Jotun. StatAlign: an extendable software package for joint Bayesian estimation of alignments and evolutionary trees. *Bioinformatics (Oxford, England)*. 2008 Oct 15; 24(20):2403–4. DOI: 10.1093/bioinformatics/btn457
- Redelings, Benjamin D.; Suchard, Marc A. Joint Bayesian estimation of alignment and phylogeny. *Systematic biology*. 2005 Jun; 54(3):401–18. DOI: 10.1080/10635150590947041 [PubMed: 16012107]
- Redelings, Benjamin D.; Suchard, Marc a. Incorporating indel information into phylogeny estimation for rapidly emerging pathogens. *BMC evolutionary biology*. 2007 Jan; 7:40. doi: 10.1186/1471-2148-7-40 [PubMed: 17359539]
- Sober E, Michael Steel. Testing the Hypothesis of Common Ancestry. *Journal of Theoretical Biology*. 2002; 218:395–408. DOI: 10.1006/yjtbi.3086 [PubMed: 12384044]
- Swofford, David L.; Thorne, Jeffrey L.; Felsenstein, Joseph; Wiegmann, Brian M. The Topology-Dependent Permutation Test for Monophyly Does Not Test for Monophyly. *Systematic Biology*. 1996 Dec; 45(4):575. doi: 10.2307/2413533
- Theobald, Douglas L. A formal test of the theory of universal common ancestry. *Nature*. 2010 May 13; 465(7295):219–22. DOI: 10.1038/nature09014 [PubMed: 20463738]
- Theobald, Douglas L. Theobald reply. *Nature*. 2010b Dec 16; 468(7326):E10–E10. DOI: 10.1038/nature09483
- Theobald, Douglas L. On universal common ancestry, sequence similarity, and phylogenetic structure: The sins of P-values and the virtues of Bayesian evidence. *Biology direct*. 2011 Nov 24; 6(1):60. doi: 10.1186/1745-6150-6-60 [PubMed: 22114984]

- Thompson JD, Higgins DG, Gibson TJ. CLUSTAL W: improving the sensitivity of progressive multiple sequence alignment through sequence weighting, position-specific gap penalties and weight matrix choice. *Nucleic Acids Research*. 1994; 22(22):4673–80. [PubMed: 7984417]
- Varon A, Wheeler WC. The tree alignment problem. *BMC Bioinformatics*. 2012; 13(1):293.doi: 10.1186/1471-2105-13-293 [PubMed: 23140486]
- Wong KM, Suchard Ma, Huelsenbeck JP. Alignment uncertainty and genomic analysis. *Science*. 2008; 319(5862):473–6. DOI: 10.1126/science.1151532 [PubMed: 18218900]
- Xie W, Lewis PO, Fan Y, Kuo L, Chen M-H. Improving marginal likelihood estimation for Bayesian phylogenetic model selection. *Systematic biology*. 2011; 60(2):150–60. DOI: 10.1093/sysbio/syq085 [PubMed: 21187451]
- Yonezawa, Takahiro; Hasegawa, Masami. Was the universal common ancestry proved? *Nature*. 2010 Dec 16.468(7326):E9. discussion E10. doi: 10.1038/nature09482 [PubMed: 21164432]
- Yonezawa, Takahiro; Hasegawa, Masami. Some Problems in Proving the Existence of the Universal Common Ancestor of Life on Earth. *The Scientific World Journal*. 2012; :479824.doi: 10.1100/2012/479824 [PubMed: 22666131]




Figure 1.

Diagram showing how the UCA and IO hypotheses can be represented by phylogenies, according to Theobald (2010). While the UCA assumes that all sequences are connected by one single phylogeny, the IO posits that there is no branch (represented in black) connecting the two domains. It is mathematically equivalent to an infinite length for this branch (see Supplementary Text S1).

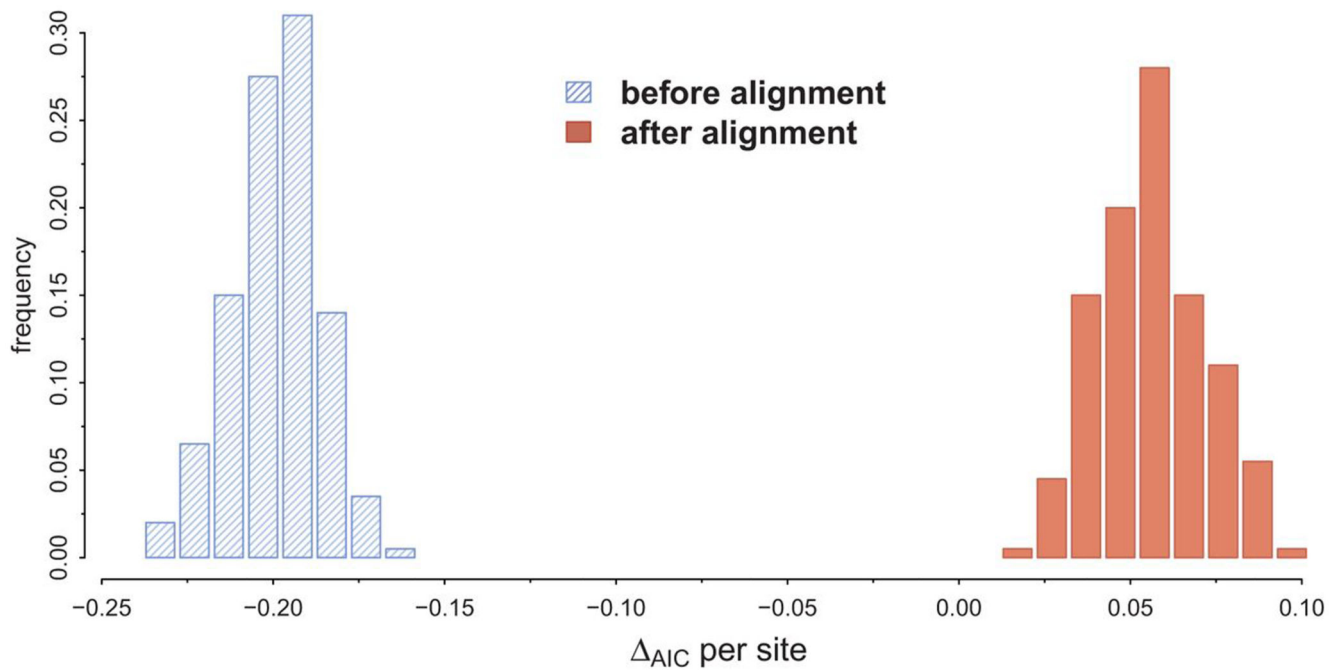


Figure 2.

Data set simulations under IO before and after optimizing the alignment, where positive values for ΔAIC suggest a UCA. It shows ΔAIC per site for data sets simulated under the best model and parameters according to the original study (rtREV+GF). The simulated data sets have 6591 sites before optimizing the alignment, and for each parameter set we simulated 200 replicates. All replicates favor IO before aligning the sequences, but then spuriously favor UCA after the alignment step.

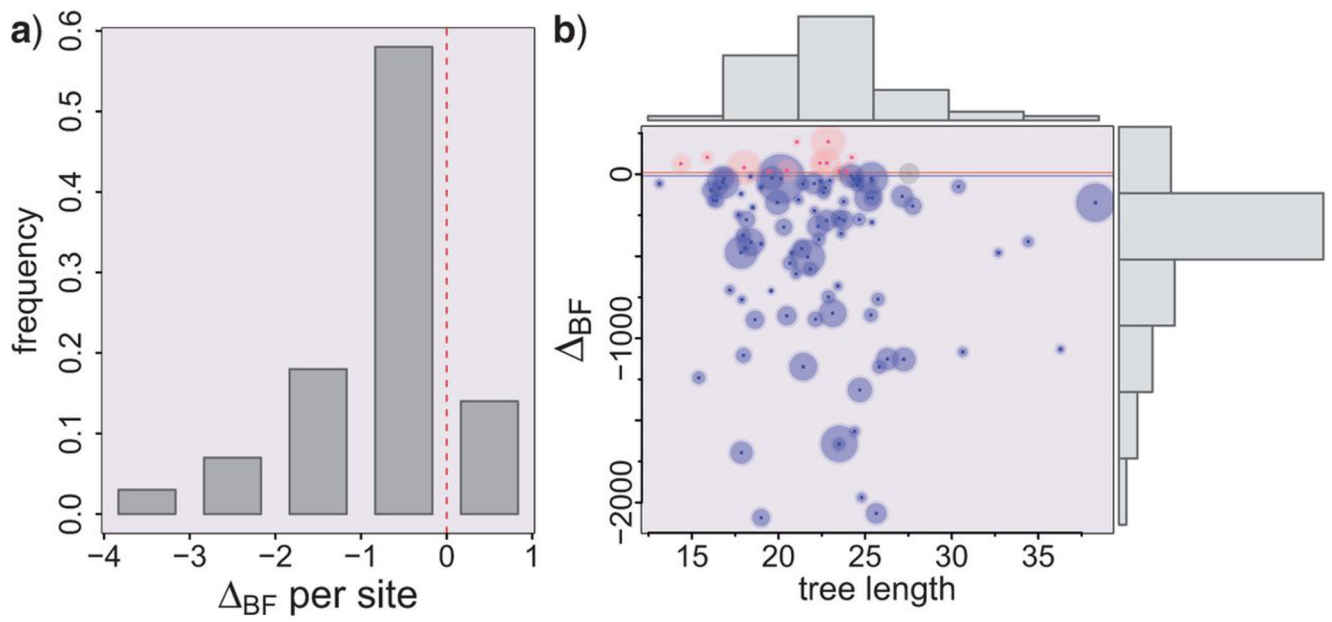


Figure 3.

Bali-phy results for IO simulated data sets. a) histogram of the log Bayes factor values per site as calculated by Δ_{BF} divided by the posterior median of the alignment length for the BE data set. b) unscaled Δ_{BF} against posterior median estimate of tree length under the UCA hypothesis, for 100 replicates. The circle diameter represents the posterior median alignment length for the BE data set, going from 506 to 868 sites. The 12 data sets shown in red wrongly support UCA, while the gray circles are two inconclusive simulations, assuming that more than 10 BF units between the hypotheses corresponds to strong evidence. The correctly identified IO data sets are shown in blue.