

Parsimony Accelerated Maximum Likelihood Searches

Kenneth Sundberg*

Computer Science Department,
Brigham Young University

E-mail: kasundberg@gmail.com

*Corresponding author

Timothy O'Connor

University of Cambridge
E-mail: timothydoconnor@gmail.com

Hyrum Carroll

Computer Science Department,
Brigham Young University
E-mail: hdc@cs.byu.edu

Mark Clement and Quinn Snell

Computer Science Department,
Brigham Young University
E-mail: clement@cs.byu.edu, nell@cs.byu.edu

Abstract: Phylogenetic search is a key tool used in a variety of biological research endeavors. However, this search problem is known to be computationally difficult, due to the astronomically large search space, making the use of heuristic methods necessary. The performance of heuristic methods for finding maximum likelihood trees can be improved by using parsimony as an initial estimator for maximum likelihood. The time spent in performing the parsimony search to boost performance is insignificant compared to the time spent in the maximum likelihood search, leading to an overall gain in search time. These parsimony boosted maximum likelihood searches lead to topologies with scores statistically similar to the unboosted searches, but in less time.

Keywords: parsimony, maximum likelihood, phylogenetic search

Reference to this paper should be made as follows: Sundberg, K., O'Connor, T., Carroll, H., Clement, M, and Snell, Q. (2007) 'Parsimony Accelerated Maximum Likelihood Searches', Int. J. Computational Biology and Drug Design, Vol. x, Nos. xx, pp.xx-xx.

1 Introduction

Phylogenetic analysis is an important tool in many biological research programs. It is used in fields such as epidemiology [7, 28], viral transmission [8, 16], and biogeography [10] among many others. New data for inferring phylogenies is being rapidly produced. This creates a need for faster heuristic methods to analyze this data.

There are two common methods for inferring phylogenies from multiple alignment data, maximum parsimony (MP) [2] and maximum likelihood (ML) [13]. The search problem for both methods is known to be NP-Hard [6]

with parsimony known to also be NP-Complete [9], however the scoring of individual topologies during the search is quite different. The parsimony score of a topology can be computed in linear time with respect to the number of nodes in the topology, whereas the best known ML computations run in exponential time with respect to the number of nodes. As a result MP searches are much faster than ML searches.

Unfortunately parsimony has a strong bias toward long branch attraction, and can lead to positively misleading topologies [12]. Maximum likelihood also models more of

the biology, including probabilities and estimates of branch lengths, information that can be required by other methods of phylogenetic analysis. As a result of the flaws of parsimony and the additional information provided by maximum likelihood, maximum likelihood is preferred by many researchers.

2 Data Sets

For this work we used three types of alignment data sets: four taxa synthetic data, small real data sets and large data sets.

2.1 Synthetic Alignments

The four taxa synthetic data sets were used to establish areas of general concordance between ML and MP. To produce them we used the program Dawg [5] under a General Time Reversible (GTR) [31, 21, 26] model of evolution. We modeled the parameters according to the examples included with the program and explored a range of branch lengths as seen in Figure 1. The lambda value of 0.1 was used for the indel evolution rate and can be interpreted as one indel for every ten substitutions. The sequence length was increased to 2000 as this gives a reasonably sized sequence to allow for the expected value of any simulation to be seen.

We ran two types of trees, a Felsenstein topology and a Farris topology as shown in Figure 2. Dawg generated data sets for trees under both topologies where the α and β branch lengths were varied from a branch length of 0.1 to 4.05 incremented by 0.05. A branch length of one is interpreted to mean that each site is expected to have one substitution from the internal node under the GTR definition of branch length. For each permutation of α and β branch lengths we ran a hundred replicates to get a percentage of matches between the two methods. This created a total of 1,280,000 data sets. In addition, this entire procedure was repeated with the the model changed to Jukes-Cantor (JC) [19], creating an additional 1,280,000 data sets.

2.2 Small Real Alignment Data Sets

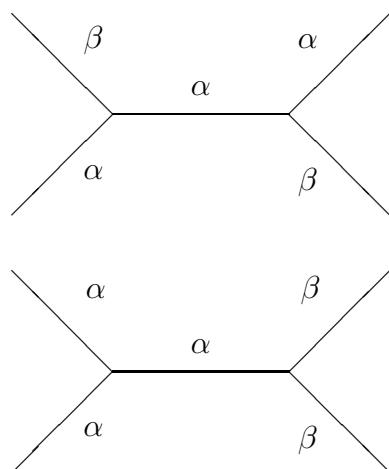
We used all the data sets with nine, ten, and eleven taxa from the `mdsa_100s` version of the benchmarks published by Carroll *et al.* [3]. These 279 data sets are protein-coding DNA alignments, derived from reference amino acid alignments [32, 24, 11, 23]. These small real alignments were used to perform exhaustive MP and ML searches.

2.3 Large Data Sets

We also used several large data sets provided by Stamatakis (<http://icwww.epfl.ch/~stamatak/index-Dateien/Page443.htm>). These data sets include seven real alignments ranging from 101 to 500 taxa and 10 simulated data sets, each with 4000 taxa and

```
Tree =
((a:0.1,b:0.25):0.05,
 (c:0.25,d:0.1):0.05);
Length = 2000
Model = "GTR"
#Rates of Substitution:
#AC, AG, AT, CG, CT, GT
Params = {1.5, 3.0, 0.9, 1.2, 2.5, 1.0}
#Frequencies of Nucleotides: A,C,T,G
Freqs = {0.20, 0.30, 0.30, 0.20}
Format = "Phylip"
Lambda = 0.1
GapModel = "NB"
GapParams = {1,0.5}
```

Figure 1: Typical file given to Dawg (Felsenstein zone)



**Felsenstein
Topology**

**Farris
Topology**

Figure 2: The Felsenstein and Farris Trees

2000 base pairs. These data sets are used to evaluate our heuristic on large data sets using RAxML [29].

3 Correlations

Both parsimony and maximum likelihood methods try to extract the evolutionary signal out of a multiple sequence alignment. Though they do so in very different ways it is reasonable to expect a fair correspondence between the two methods.

3.1 Areas of General Concordance of ML and MP

There has been a large inquiry into the relationship of ML and MP from the perspective of accuracy and strength of methods [27, 17, 18, 20] but the comparison has never taken the view of how well one estimates the other. Here we have used simulated data to elucidate this relationship. To do so we focused on the standard four taxa case and compared the performance of both methods under a Felsenstein topology and a Farris topology.

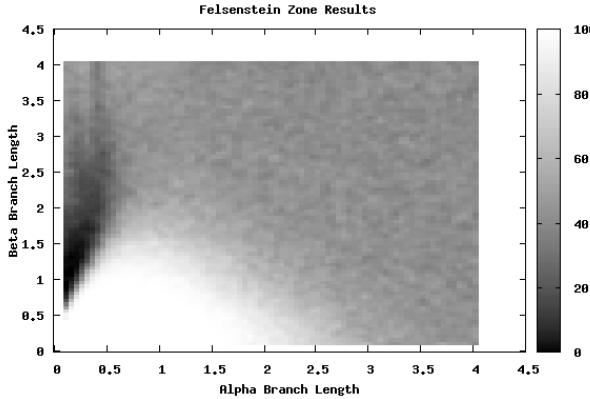


Figure 3: Chart showing the percentage of datasets of a given true Felsenstein Topology where Maximum Parsimony and Maximum Likelihood predicted the same optimal topology.

This topology as seen in Figure 2 is a four taxa tree with non-sibling branches that have branch length β and all other branches with length α . The second tree in Figure 2 shows a Farris topology where the two β length branches are siblings.

The Felsenstein zone [12] can be loosely described as the subspace of all Felsenstein topology branch lengths where parsimony will predict a tree topology that is incorrect. Later work went on to show that ML performs well on Felsenstein topologies, including those that lie in the Felsenstein zone and reasonably well on Farris topologies. Also shown was that MP failed in certain areas of branch lengths in the Felsenstein topology but excelled at the Farris topology [27, 17, 18, 20].

Parsimony, has a strong bias toward placing longer branches as siblings. This bias gives it an advantage when longer branches are ‘truly’ siblings and a disadvantage when the ‘true’ tree has non-sibling long branches.

To illustrate this we scored simulated alignments, as described in section 2.1 and compared the topological result produced from each method, the results are shown in Figures 3 and 4. These graphs show the percentage of time each method chose the same tree of the three possible trees.

One area, the Felsenstein zone, does not correlate well and forms the dark parabolic structure in Figure 3. This area of disagreement is due to the long branch attraction problem (LBA). From the perspective of parsimony as an estimator of likelihood topology, this should not be a problem because the number of long branches in a large data set are usually few and would only vary the topology by a few branches that are incorrect according to likelihood.

By running a likelihood search based on the best parsimony tree we can save time by getting the majority of branches correct and leaving likelihood to figure out the incorrect branches. This saves on the cost of using likelihood

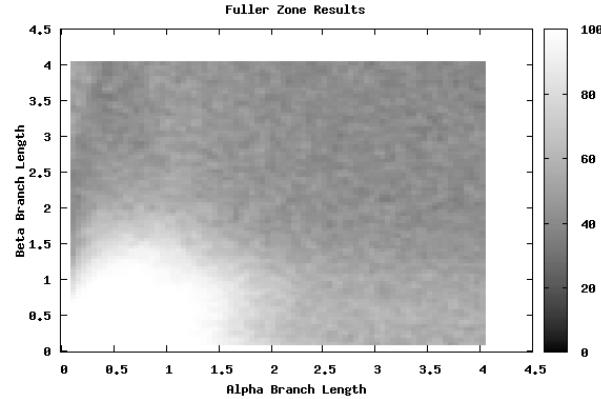


Figure 4: Chart showing the percentage of datasets of a given true Farris Topology where Maximum Parsimony and Maximum Likelihood predicted the same optimal topology.

to search the whole tree space but lets it fine tune after the majority of bad trees are ruled out by parsimony. We are not trying to use parsimony as a function to predict ML but as a method to boost the search pass the large number of trees that both procedures predict are incorrect.

Most trees are assumed to be within the area where both methods are completely consistent, the light colored areas of Figures 3 and 4. In the gray areas of the figures the methods are doing no better than randomly guessing. This is to be expected as this occurs where there is an expected rate of 2 or more substitutions per site, which causes a large loss of phylogenetic signal. As expected, the major limitation to the concordance between ML and MP are those topologies that lie within the Felsenstein zone.

Each tree was scored using the phylogenetics program PAUP* [30] under both MP and ML. We derived the results shown in Figures 3 and 4 using the defaults of PAUP* for parsimony and likelihood. The model for likelihood used was HKY [15].

This is an under-parametrization of the base model, GTR, used by Dawg to generate the data. The use of this under-parametrization, as HKY is a more simplified model compared to GTR, was used because it allowed for a universal model for all data sets without having to calculate the optimum parameters or model for each data set. To further explore the effects of this assumption we reran the analysis using a base model of JC in Dawg to see if there was a change in the results. In this case the use of HKY in the likelihood evaluation would be an over-parametrization of the the data sets model. We found no particularly different results (data not shown).

3.2 Gap State

When using parsimony as a heuristic for likelihood the gap mode used for the parsimony score is critical for good cor-

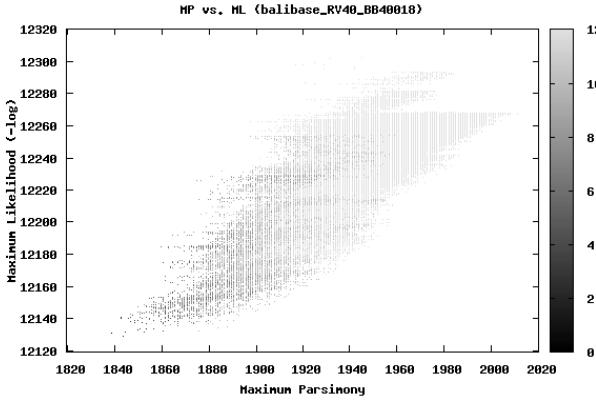


Figure 5: All possible trees for BALiBASE dataset RV40 BB40010 scored with parsimony (GapMode = missing) and maximum likelihood (HKY under PAUP* defaults). Color indicates RF distance from optimal parsimony tree. Note that the optimal tree under likelihood is a small RF distance from the optimal parsimony tree.

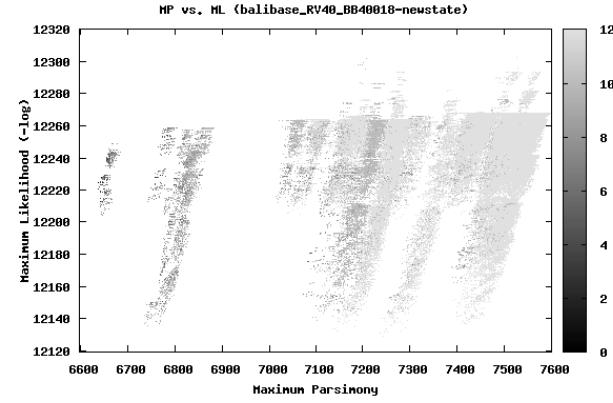


Figure 6: All possible trees for BALiBASE dataset RV40 BB40010 scored with parsimony (GapMode = new state) and maximum likelihood (HKY under PAUP* defaults). Color indicates RF distance from optimal parsimony tree. Note that the optimal tree under likelihood is a large RF distance from the optimal parsimony tree.

respondence. Treating the gaps in an alignment as a new state creates large variances in the parsimony scores, depending on the extent to which the gaps line up. One solution when using parsimony as a precursor to a maximum likelihood search to set the gap mode to missing to avoid this difficulty. Figures 5 and 6 show the difference in correspondence between the two methods. Both figures are a plot of an exhaustive search through one of the BALiBASE [32] data sets (RV40 BB40010). Every possible topology was scored by both parsimony and maximum likelihood, and then plotted. The color of each point indicates the topological distance, (measured using Robinson Foulds [25]), from the most parsimonious topology. In both cases maximum likelihood was scored using the HKY method. These results are representative of the other data sets.

3.3 Correlation Between The Most Likely and The Most Parsimonious Topology

While the correspondence between MP and ML methods exists it does not preserve a partial order between topologies. It is therefore not possible to directly substitute one method for the other. However it is possible to use parsimony to aid a maximum likelihood search. Figure 5 shows that as parsimony score decreases, the negative log maximum likelihood score generally decreases also. This trend also generalizes to the other data sets studied.

In all but two of the nearly 300 real data sets, less than 70 possible topologies had a better likelihood score than the most parsimonious topology for that dataset. It is clear that in the majority of cases, parsimony eliminates most of the incorrect topologies. Figure 7 is a histogram of data sets grouped by the percentage of trees with better likelihood scores than the best parsimony tree. Note that

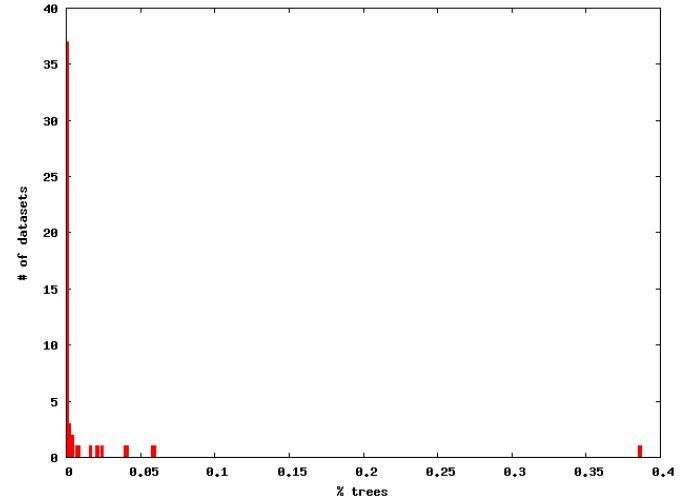


Figure 7: Percentage of trees with a better likelihood score than the most parsimonious tree

37 datasets fell into the 0% trees bucket. For the 9 taxa case, from which these data sets have been drawn 0.05% corresponds to 67 trees.

Furthermore the likelihood score of the best parsimony tree is often very comparable to the likelihood score of the best maximum likelihood tree. Figure 8 shows the relative error of the best parsimony tree, taking the best maximum likelihood score as the true value.

3.4 Correlations between likelihood methods

Model selection can have significant impact on maximum likelihood scores. Even when the same model, such as GTR+ γ is used variations in the implementation of a phy-

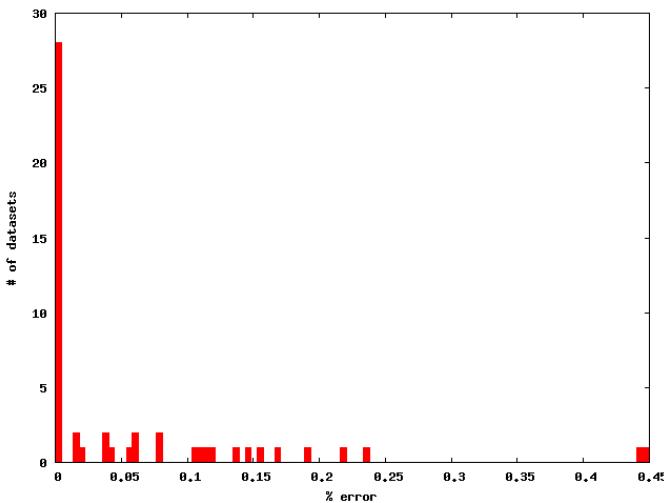


Figure 8: Relative error in ML Score of MP tree

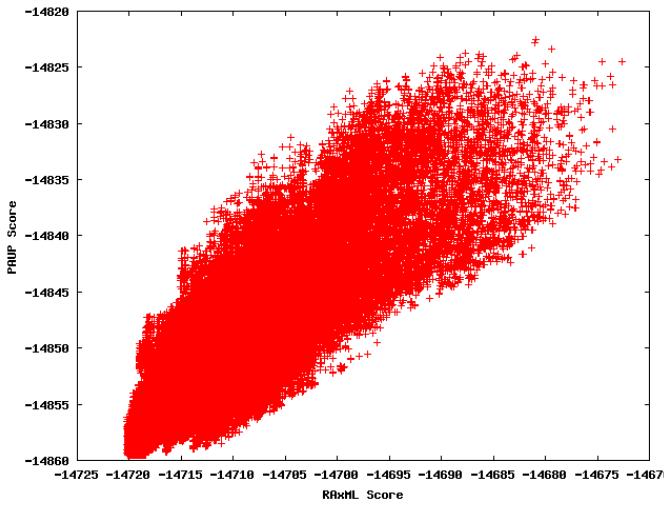


Figure 9: A comparison between ML scores computed by RAxML and PAUP using GTR+ γ

logenetic search algorithm can cause variations in the scoring of trees. One such example is the differences in the scoring between RAxML and PAUP. Figure 9 shows a comparison of the scores produced by these two programs. Both programs agree on the best tree, though not on its score. However, beyond this agreement there is at best a general concordance between the two scores. The precise relationship between these two scoring methods is not clear.

Comparing Figures 5 and 9 shows how maximum parsimony also has this same general concordance with maximum likelihood as the scores of the same model under two different software implementations. This is not unexpected as maximum parsimony can be thought of as an extremely simple model for maximum likelihood. This correlation leads us to the Parsimony Heuristic.

```
begin paup;
  set criterion = parsimony;
  hsearch
  set criterion = likelihood;
  hsearch start=current swap=TBR
    rearrlimit=x;
  lscores all scorefile=<filename>;
  quit;
end;
```

Figure 10: PAUP* block to perform heuristic search

4 Parsimony Heuristic

Not only do likelihood and parsimony generally improve together, the most likely and the most parsimonious trees are often very close topologically. The major difficulty with parsimony methods is long branch attraction [1]. This problem causes two long branches to be incorrectly made siblings in the final topology. This topological change is however correctable with a very small number of TBR swaps. A simple heuristic that can be applied to many existing phylogenetic programs is to perform a parsimony search and use the result as the seed tree for a maximum likelihood search. This maximum likelihood search can be further limited to trees within a small TBR distance of the most parsimonious tree.

The idea of using parsimony to improve likelihood searches is not entirely new. RAxML uses a step-wise maximum parsimony tree as its starting tree [29]. This is a step in the right direction but performing a complete parsimony search further improves performance of this search.

4.1 Improving PAUP*

Our heuristic is very easily implemented in PAUP*. PAUP* already has the machinery to perform both parsimony and maximum likelihood searches. The heuristic can be expressed in a PAUP* block, as shown in Figure 10. As parsimony gives us a tree that is topologically close to the desired tree, we limit the length of the maximum likelihood search, gaining further savings in time. The x used as the *rearrlimit* was set to the number of trees within one TBR swap.

We achieved identical scores on the majority of data sets using this method. Figure 11 shows a histogram of the differences in the scores of the final tree. It is not statistically significant that occasionally the heuristic outperforms the standard PAUP* search, nor is it significant that the standard PAUP* search occasionally produces a better score than the heuristic.

The use of the heuristic does improve search times significantly. As shown in Figure 12, the run time was on average improved by 25%. In a handful of the data sets, parsimony was not a good indicator of likelihood and the search took longer using the heuristic than it would have otherwise. All times are the average of five runs.

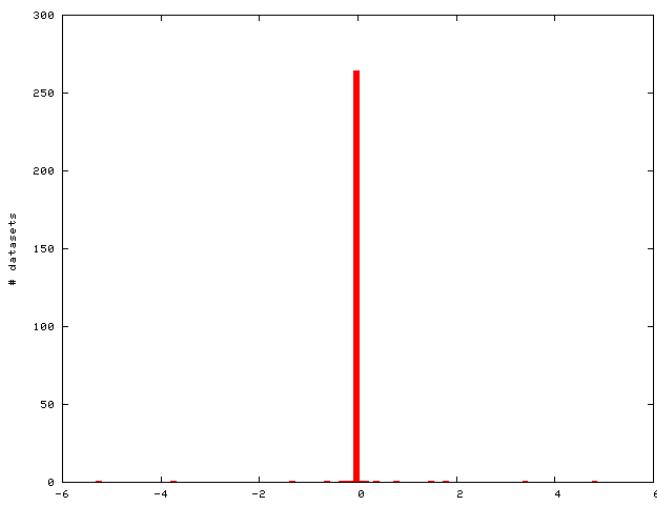


Figure 11: Score differences between PAUP* with and without the Parsimony Heuristic

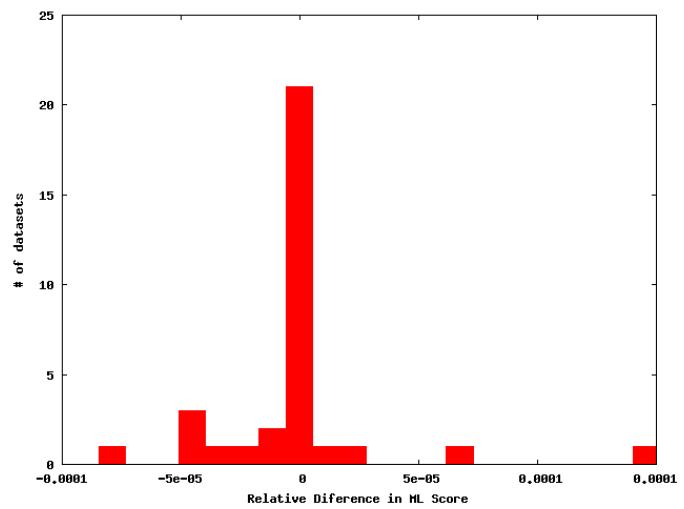


Figure 13: Score differences between RAxML with and without using TNT

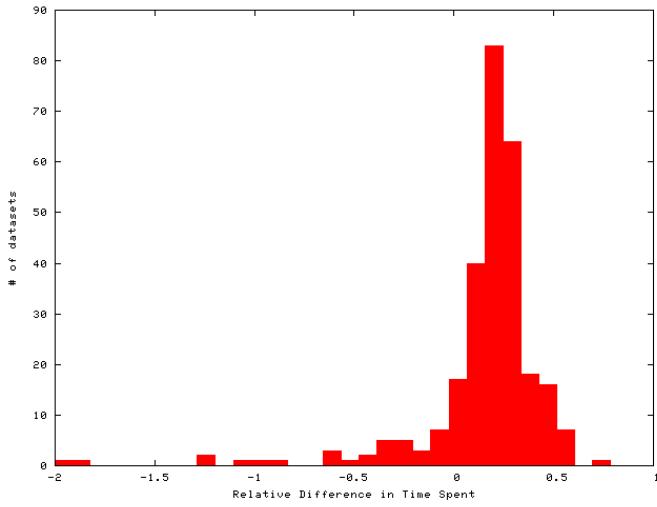


Figure 12: Running time improvement from starting with a parsimony search in PAUP*

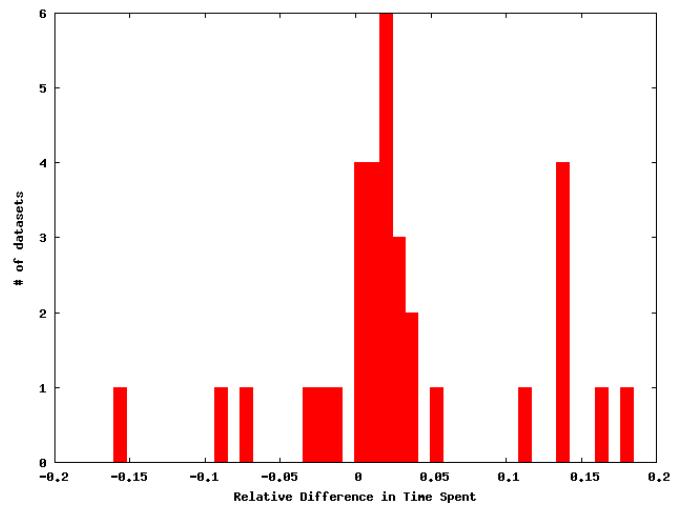


Figure 14: Execute time improvements using RAxML and TNT versus just RAxML

4.2 Improving RAxML

Due to the size of the large data sets, exhaustive searching is not feasible. RAxML is a phylogenetic search program that can perform maximum likelihood searches on large data sets [29]. One of the heuristics it uses is to start the ML searching with a stepwise maximum parsimony tree. We took a slightly different approach to applying our heuristic to RAxML. First, we heuristically searched under MP, and then started the ML search in RAxML with the most parsimonious tree found. To search with MP, we used TNT, Tree Analysis Using New Technology [14, 22]. In other analysis that we've performed, TNT out performs PAUP* in terms of time and parsimony score [4]. We used two TNT search strategies per each date set: 1) sectorial-search and xmult and 2) TBR. The most parsimonious tree found with TNT was used to start a GTR-GAMMA search with RAxML. The relative difference in scores is shown in

Figure 13. Again the scores are not statistically different, but there is more variability than with the PAUP* scores.

The execution time improvements are shown in Figure 14. There are two circumstances which are probably causing this reduction in the performance of the heuristic. First, RAxML by default starts its search with a step-wise maximum parsimony tree. This tree already has a reasonable parsimony score, so there is not as much benefit to be gained by finding a more parsimonious tree. Another difference is that RAxML does not allow us to limit the time spent in the likelihood search, which does not allow us to take full advantage of the small topological differences between the maximum parsimony and maximum likelihood trees.

5 Conclusions

Parsimony is correlated with likelihood in phylogenetic inference. In many data sets the predicted topologies are very close, if not identical. Maximum likelihood searches can be boosted using our parsimony heuristic in a manner that avoids the problems associated with the Felsenstein Zone and parsimony while maintaining some of the speed of a parsimony search.

The Parsimony Heuristic, using a simple model to eliminate clearly undesirable trees, applies to more than just parsimony. This idea can likely also be extended to the Felsenstein Zone case, especially when a score from a more complex model is desired. Not only does parsimony reflect likelihood, different models of evolution also show a correspondence of scores. This means that in a data set where long branch attraction is known to be a problem, a simple maximum likelihood model could still be used to boost the search through much of phylogenetic tree space before switching over to the more complex model desired.

REFERENCES

- [1] J. Bergsten. A review of long-branch attraction. *Cladistics*, 21(2):163–193, 2005.
- [2] J.H. Camin and R.R. Sokal. A Method for Deducing Branching Sequences in Phylogeny. *Evolution*, 19(3):311–326, 1965.
- [3] Hyrum Carroll, Wesley Beckstead, Timothy O'Connor, Mark Ebbert, Mark Clement, Quinn Snell, and David McClellan. DNA Reference Alignment Benchmarks Based on Tertiary Structure of Encoded Proteins. In press with Bioinformatics.
- [4] Hyrum Carroll, Mark Clement, Quinn Snell, and Keith Crandall. Phylogenetic analysis of large sequence data sets. In *Proceedings of the Second Biotechnology and Bioinformatics Symposium*, pages 20–24, 2005.
- [5] R.A. Cartwright. DNA assembly with gaps (Dawg): simulating sequence evolution. *Bioinformatics*, 21(3):31–38, 2005.
- [6] B. Chor and T. Tuller. Maximum likelihood of evolutionary trees is hard. *Proceedings of the 9th Annual International Conference on Research in Computational Molecular Biology (RECOMB 2005*, 3500:296–310, 2005.
- [7] A.G. Clark, K.M. Weiss, D.A. Nickerson, S.L. Taylor, A. Buchanan, J. Stengard, V. Salomaa, E. Vartiainen, M. Perola, E. Boerwinkle, and C.F. Sing. Haplotype structure and population genetic inferences from nucleotide-sequence variation in human lipoprotein lipase. *American Journal of Human Genetics*, 63:595–612, 1998.
- [8] K.A. Crandall. Multiple interspecies transmissions of human and simian t-cell leukemia/lymphoma virus type i sequences. *Molecular Biology and Evolution*, 13:115–131, 1996.
- [9] W.H.E. Day, D.S. Johnson, and D. Sankoff. The computational complexity of inferring rooted phylogenies by parsimony. *Mathematical Biosciences*, 81(33–42):299, 1986.
- [10] R. DeSalle. Molecular approaches to biogeographic analysis of Hawaiian Drosophilidae. *Hawaiian Biogeography* (ed. by WL Wagner and VA Funk), pp. 72–89, 1995.
- [11] Robert C. Edgar. MUSCLE: multiple sequence alignment with high accuracy and high throughput. *Nucleic Acids Research*, 32(5):1792–1797, mar 2004.
- [12] J. Felsenstein. Cases in which Parsimony or Compatibility Methods Will be Positively Misleading. *Systematic Zoology*, 27(4):401–410, 1978.
- [13] Joseph Felsenstein. Evolutionary Trees from DNA Sequences: A Maximum Likelihood Approach. *Journal of Molecular Evolution*, 17(6):368–376, 1981.
- [14] Pablo Goloboff, Steve Farris, and Kevin Nixon. TNT: Tree analysis using new technology. <http://www.cladistics.com/webtnt.html>, 2001.
- [15] M. Hasegawa, H. Kishino, and T. Yano. Dating of the human-ape splitting by a molecular clock of mitochondrial DNA. *Journal of Molecular Evolution*, 22(2):160–174, 1985.
- [16] BL Herring, F. Bernardin, S. Caglioti, S. Stramer, L. Tobler, W. Andrews, L. Cheng, S. Rampersad, C. Cameron, J. Saldanha, et al. Phylogenetic analysis of WNV in North American blood donors during the 2003–2004 epidemic seasons. *Virology*, 2007.
- [17] J.P. Huelsenbeck. Performance of Phylogenetic Methods in Simulation. *Systematic Biology*, 44(1):17–48, 1995.
- [18] J.P. Huelsenbeck and D.M. Hillis. Success of Phylogenetic Methods in the Four-Taxon Case. *Systematic Biology*, 42(3):247–264, 1993.
- [19] T.H. Jukes and C.R. Cantor. Evolution of protein molecules. *Mammalian Protein Metabolism*, 3:21–132, 1969.
- [20] MK Kuhner and J. Felsenstein. A simulation comparison of phylogeny algorithms under equal and unequal evolutionary rates [published erratum appears in Mol Biol Evol 1995 May; 12 (3): 525]. *Mol Biol Evol*, 11(3):459–468, 1994.
- [21] C. Lanave, G. Preparata, C. Sacone, and G. Serio. A new method for calculating evolutionary substitution rates. *Journal of Molecular Evolution*, 20(1):86–93, 1984.

- [22] R. Meier and F.B. Ali. Software Review. The newest kid on the parsimony block: TNT (Tree analysis using new technology) . *Systematic Entomology*, 30(1):179, 2005.
- [23] CP Ponting, J. Schultz, F. Milpetz, and P. Bork. SMART: identification and annotation of domains from signalling and extracellular protein sequences. *Nucleic Acids Research*, 27(1):229–232, 1999.
- [24] G. P. S. Raghava, Stephen M. J. Searle, Patrick C. Audley, Jonathan D. Barber, and Geoffrey J. Barton. OXBench: a benchmark for evaluation of protein multiple sequence alignment accuracy. *BMC Bioinformatics*, 4:47, 2003.
- [25] DF Robinson and LR Foulds. Comparison of Phylogenetic Trees. *Mathematical Biosciences*, 53(1):131–147, 1981.
- [26] F. Rodriguez, JL Oliver, A. Marin, and JR Medina. The general stochastic model of nucleotide substitution. *J Theor Biol*, 142(4):485–501, 1990.
- [27] M. Siddall. Success of Parsimony in the Four-Taxon Case: Long-Branch Repulsion by Likelihood in the Farris Zone. *Cladistics*, 14(3):209–220, 1998.
- [28] CF Sing, MB Haviland, KE Zerba, and AR Templeton. Application of cladistics to the analysis of genotype-phenotype relationships. *European Journal of Epidemiology*, 8:3–9, 1992.
- [29] A. Stamatakis. RAxML-VI-HPC: maximum likelihood-based phylogenetic analyses with thousands of taxa and mixed models. *Bioinformatics*, 22(21):2688, 2006.
- [30] David L. Swofford. *PAUP*: Phylogenetic Analysis Using Parsimony (* and Other Methods). Version 4*. Sinauer Associates, Sunderland, Massachusetts, 2003.
- [31] S. Tavare. Some probabilistic and statistical problems on the analysis of DNA sequences. *Lect. Math. Life Sci*, 17:57–86, 1986.
- [32] Julie D. Thompson, Patrice Koehl, Raymond Ripp, and Olivier Poch. BALiBASE 3.0 latest developments of the multiple sequence alignmnet benchmark. *Proteins: Structure, Function, and Bioinformatics*, 61(1):127–136, 2005.