

Growing phylogenetic trees with TreeLine

Erik S. Wright

April 30, 2024

Contents

1	Introduction	1
2	Performance Considerations	1
3	Growing a Phylogenetic Tree	2
4	Plotting Branch Support Values	5
5	Ancestral State Reconstruction	11
6	Exporting the Tree	13
7	Session Information	13

1 Introduction

This document describes how to grow phylogenetic trees using the `TreeLine` function in the DECIPHER package. `TreeLine` takes as input a set of aligned nucleotide or amino acid sequences and returns a phylogenetic tree (i.e., *dendrogram* object) as output. This vignette focuses on building maximum likelihood (ML) and maximum parsimony (MP) phylogenetic trees starting from sequences, but `TreeLine` can also be used to build additive trees from a distance matrix.

Why is the function called `TreeLine`? The goal of `TreeLine` is to find the most likely/parsimonious tree for a given sequence alignment. There are often many trees with nearly maximal likelihood/parsimony. Therefore, `TreeLine` seeks to find a tree as close as possible to the treeline, analogous to how no trees can grow above the treeline on a mountain.

Why use `TreeLine` versus other programs? The `TreeLine` function is designed to return an excellent phylogenetic tree with minimal user intervention. Many tree building programs have a large set of complex options for niche applications. In contrast, `TreeLine` simply builds a great tree when relying on its defaults. This vignette is intended to get you started and introduce additional options/functions that might be useful.

`TreeLine` uses multi-start optimization followed by hill-climbing to find the highest trees on the likelihood or parsimony landscapes. Since `TreeLine` is a stochastic optimizer, it optimizes many trees to prevent luck from influencing the final result. The main difference from most other approaches to tree optimization is that `TreeLine` heavily uses past trees to generate new trees as it optimizes. With any luck it'll find the treeline!

2 Performance Considerations

Finding a tree with very high likelihood/parsimony is no easy feat. `TreeLine` systematically optimizes hundreds to thousands of candidate trees before returning the best one. This takes time, but there are things you can do to make it

go faster.

- Only use the sequences you need: *TreeLine* scales approximately quadratically with the number of sequences. Hence, limiting the number of sequences is a worthwhile consideration. In particular, always eliminate redundant sequences, as shown below, and remove any sequences that are not necessary. This concern is shared for all tree building programs, and *TreeLine* is no exception.
- Choose a model: Automatic model selection is a useful feature, but frequently this time-consuming step can be skipped. For most modestly large sets of nucleotide sequences, the "GTR+G4" model will be automatically selected. Typical amino acid sequences will tend to pick the "LG+G4" or "WAG+G4" models, unless the sequences are from a particular origin (e.g., mitochondria). Pre-selecting a subset of the available `MODELS` and supplying this as the *model* argument can save considerable time.
- Set a timeout: The `maxTime` argument specifies the (approximate) maximum number of hours you are willing to let *TreeLine* run. If you are concerned about the code running too long then simply specify this argument.
- Compile with OpenMP support: Significant speed-ups can be achieved with multi-threading using OpenMP. See the "Getting Started DECIPHERing" vignette for how to do this on your computer platform. Then you only need to set the argument `processors=NULL` and *TreeLine* will use all available processors.
- Compile for SIMD support: *TreeLine* is configured to make use of SIMD operations, which are available on some processors. The easiest way to enable SIMD is to add "-O3 -march=native" to the end of `PKG_CFLAGS` in the "DECIPHER/src/MAKEVARS" text file. Then, after recompiling, there may be an automatic speed-up on systems with SIMD support. Note that enabling SIMD makes the compiled code non-portable, so the code always needs to be compiled on the hardware being used.

3 Growing a Phylogenetic Tree

TreeLine takes as input a multiple sequence alignment when constructing a maximum likelihood or maximum parsimony phylogenetic tree. Multiple sequence alignments can be constructed from a set of (unaligned) sequences using `AlignSeqs` or related functions. *TreeLine* will optimize trees for amino acid (i.e., `AAStringSet`) or nucleotide (i.e., `DNAStringSet` or `RNAStringSet`) sequences. Here, we are going to use a set of sequences that is included with DECIPHER. These sequences are from the internal transcribed spacer (ITS) between the 16S and 23S ribosomal RNA genes in several *Streptomyces* species.

```
> library(DECIPHER)
> # specify the path to your sequence file:
> fas <- "<<path to FASTA file>>"
> # OR find the example sequence file used in this tutorial:
> fas <- system.file("extdata", "Streptomyces_ITS_aligned.fas", package="DECIPHER")
> seqs <- readDNAStringSet(fas) # use readAAStringSet for amino acid sequences
> seqs # the aligned sequences
DNAStringSet object of length 88:
      width seq
[1] 627 TGTACACACCGCCCGTCA-CGTC...GGGGTTTCCGAATGGGGAAACC supercont3.1 of S...
[2] 627 NNNNCACACCGCCCGTCA-CGTC...GGGGTTTCCGAATGGGGAAACC supercont3.1 of S...
[3] 627 TGTACACACCGCCCGTCA-CGTC...GGGGTTTCCGAATGGGGAAACC supercont1.1 of S...
[4] 627 CGTACACACCGCCCGTCA-CGTC...GGGGTTTCCGAATGGGGAAACC supercont1.1 of S...
[5] 627 TGTACACACCGCCCGTCA-CGTC...GGGGTTTCCGAATGGGGAAACC supercont1.1 of S...
...
[84] 627 TGTACACACCGCCCGTCA-CGTC...GGGGTTTCCGAATGGGGAAACC gi|297189896|ref|...
[85] 627 TGTACACACCGCCCGTCA-CGTC...GGGGTGTCCGAATGGGGAAACC gi|224581106|ref|...
```

```
[86] 627 TGTACACACCGCCCGTCA-CGTC...GGGGTGTCCGAATGGGGAAACC gi|224581106|ref|...
[87] 627 TGTACACACCGCCCGTCA-CGTC...GGGGTGTCCGAATGGGGAAACC gi|224581106|ref|...
[88] 627 TGTACACACCGCCCGTCA-CGTC...GGGGTTTCCGAATGGGGAAACC gi|224581108|ref|...
```

Many of these sequences are redundant or from the same genome. We can de-replicate the sequences to accelerate tree building:

```
> seqs <- unique(seqs) # remove duplicated sequences
> ns <- gsub("^.*Streptomyces( subsp\\. | sp\\. | | sp_) ([^ ]+).*$", "\\2", names(seqs))
> names(seqs) <- ns # name by species (or any other preferred names)
> seqs <- seqs[!duplicated(ns)] # remove redundant sequences from the same species
> seqs
DNAStrngSet object of length 19:
      width seq                                     names
[1]    627 TGTACACACCGCCCGTCA-CGTC...GGGGTTTCCGAATGGGGAAACC  albus
[2]    627 TGTACACACCGCCCGTCA-CGTC...GGGGTTTCCGAATGGGGAAACC  clavuligerus
[3]    627 TGTACACACCGCCCGTCA-CGTC...GGGGTGTCCGAATGGGGAAACC  ghanaensis
[4]    627 TGTACACACCGCCCGTCA-CGTC...GGGGTTTCCGAATGGGGAAACC  griseoflavus
[5]    627 TGTACACACCGCCCGTCA-CGTC...GGGGTGTCCGAATGGGGAAACC  lividans
...    ...    ...
[15]   627 TGTACACACCGCCCGTCA-CGTC...GGGGTGTCCGAATGGGGAAACC  cattleya
[16]   627 TGTACACACCGCCCGTCA-CGTC...GGGGTTTCCGAATGGGGAAACC  bingchenggensis
[17]   627 TGTACACACCGCCCGTCA-CGTC...GGGGTTTCCGAATGGGGAAACC  avermitilis
[18]   627 TGTACACACCGCCCGTCA-CGTC...GGGGTGTCCGAATGGGGAAACC  C
[19]   627 TGTACACACCGCCCGTCA-CGTC...GGGGTGTCCGAATGGGGAAACC  Tu6071
```

Now, it's time to try our luck at finding the most likely tree. Here, we will set a stringent time limit (0.01 hours) to make this example faster, although longer time limits (e.g., 24 hours) are advised because setting very short time limits leaves the result partly up to luck.

Note that `TreeLine` automatically selects a substitution model based on Akaike information criterion (by default). It is possible to specify specific model(s) (e.g., `model="GTR+G4"`) to limit the possible selections and test your luck with fewer models.

Also, since `TreeLine` is a stochastic optimizer, it is critical to always set the random number seed for reproducibility. You can pick any lucky number, and if you ever wonder how much you pushed your luck, you can try running again from a different random number seed to see how much the result came down to luck of the draw.

Note that setting a time limit, as done below with `maxTime`, negates the purpose of setting a seed – never set a time limit if reproducibility is desired or you'll have no such luck.

```
> set.seed(123) # set the random number seed
> tree <- TreeLine(seqs,
  method="ML",
  model="GTR+G4",
  reconstruct=TRUE,
  maxTime=0.01)
```

Optimizing model parameters:
GTR+G4 -ln(L) = 4368, AICc = 8832, BIC = 9017

PHASE 1 OF 3: INITIAL TREES

```
1/3. Optimizing initial tree #1 of 10 to 100:
-ln(L) = 4366.5 (-0.044%), 1 Climbs
1/3. Optimizing initial tree #2 of 10 to 100:
-ln(L) = 4365.3 (-0.027%), 7 Climbs
1/3. Optimizing initial tree #3 of 11 to 100:
-ln(L) = 4363.8 (-0.035%), 6 Climbs
```

PHASE 2 OF 3: REGROW GENERATION 1 OF 10 TO 20

```
2/3. Optimizing regrown tree #1 of 10 to 100:
-ln(L) = 4363.7 (-0.002%), 2 Climbs
2/3. Optimizing regrown tree #2 of 10 to 100:
-ln(L) = 4377.0 (+0.305%), 2 Climbs
2/3. Optimizing regrown tree #3 of 10 to 100:
-ln(L) = 4363.8 (+0.002%), 2 Climbs
2/3. Optimizing regrown tree #4 of 10 to 100:
-ln(L) = 4365.7 (+0.045%), 2 Climbs
2/3. Optimizing regrown tree #5 of 10 to 100:
-ln(L) = 4363.7 (+0.002%), 2 Climbs
2/3. Optimizing regrown tree #6 of 10 to 100:
-ln(L) = 4363.8 (+0.002%), 2 Climbs
2/3. Optimizing regrown tree #7 of 10 to 100:
-ln(L) = 4363.8 (+0.002%), 2 Climbs
2/3. Optimizing regrown tree #8 of 10 to 100:
-ln(L) = 4367.4 (+0.086%), 4 Climbs
2/3. Optimizing regrown tree #9 of 10 to 100:
-ln(L) = 4366.6 (+0.066%), 4 Climbs
```

PHASE 3 OF 3: SHAKEN TREES

Grafting 1 tree to the best tree:
-ln(L) = 4363.7 (0.000%), 0 Grafts of 0

```
3/3. Optimizing shaken tree #1 of 3 to 1000:
-ln(L) = 4364.6 (+0.022%), 7 Climbs
3/3. Optimizing shaken tree #2 of 3 to 1000:
-ln(L) = 4373.9 (+0.235%), 4 Climbs
3/3. Optimizing shaken tree #3 of 3 to 1000:
-ln(L) = 4484.3 (+2.689%), 3 Climbs
```

Grafting 3 trees to the best tree:
-ln(L) = 4363.7 (0.000%), 0 Grafts of 3

Model parameters:
Frequency(A) = 0.175
Frequency(C) = 0.245
Frequency(G) = 0.344

4 Plotting Branch Support Values

Maybe it was just beginner's luck, but we already have a reasonable looking starting tree! TreeLine automatically returns a variety of information about the tree that can be accessed with the `attributes` and `attr` functions:

```
> #attributes(tree) # view all attributes
> attr(tree, "members") # number of leaves below this (root) node
[1] 19
> attr(tree, "height") # height of the node (in this case, the midpoint root)
[1] 2.237168
> attr(tree, "state") # ancestral state reconstruction (if reconstruct=TRUE)
[1] "-----CACCGCCCGTCA-CGTCACGAAAGTCGGTAACACCCGAAGCCGGTGGCCCAACCCCTTG-GGGAGGGAGCTGTCGAA
> attr(tree, "siteLnLs") # LnL for every alignment column (site)
  [1] -2.021099 -1.584987 -2.021099 -2.359708 -1.938356 -2.359708
  [7] -2.080347 -2.508487 -2.080347 -2.080347 -1.715876 -2.080347
 [13] -2.080347 -2.080347 -1.715876 -2.167185 -2.080347 -2.508487
 [19] 0.000000 -2.080347 -1.715876 -2.167185 -2.080347 -2.508487
 [25] -4.595481 -1.715876 -2.508487 -2.508487 -2.508487 -1.715876
 [31] -2.167185 -2.080347 -1.715876 -1.715876 -2.167185 -2.508487
 [37] -5.917834 -4.595481 -5.917834 -2.080347 -2.080347 -2.080347
 [43] -1.715876 -2.508487 -2.508487 -1.715876 -2.080347 -2.080347
 [49] -5.092905 -4.465175 -2.167185 -1.715876 -1.715876 -2.080347
 [55] -2.080347 -2.080347 -2.508487 -2.508487 -2.080347 -2.080347
 [61] -2.080347 -12.805520 -8.061224 -16.153763 -13.592698 -6.686965
 [67] -1.715876 -1.715876 -1.715876 -2.508487 -1.715876 -1.715876
 [73] -1.715876 -2.508487 -1.715876 -8.719578 -12.015230 -1.715876
 [79] -2.167185 -2.080347 -1.715876 -2.508487 -2.508487 -1.715876
 [85] -1.715876 -2.167185 -1.715876 -1.715876 -1.715876 -2.508487
 [91] -2.080347 -13.719782 -12.562330 -1.715876 -2.080347 -1.715876
 [97] -2.508487 -2.167185 -2.167185 -1.715876 -1.715876 -1.715876
[103] -2.508487 -2.080347 -1.715876 -2.508487 -2.508487 -1.715876
[109] -2.167185 -2.080347 -1.715876 -2.167185 -2.508487 -2.508487
[115] -2.080347 -2.508487 -2.508487 -1.715876 -1.715876 -2.167185
[121] -2.508487 -1.715876 -2.080347 -2.080347 -1.715876 -2.167185
[127] -2.508487 -2.080347 -2.080347 -1.715876 -1.715876 -2.508487
[133] -2.508487 -1.715876 -1.715876 -2.167185 -1.715876 -2.080347
[139] -1.715876 -1.715876 -2.080347 -2.167185 -1.715876 -1.715876
[145] -2.508487 -2.167185 -2.080347 -2.508487 -2.080347 -2.080347
[151] -2.167185 -2.080347 -2.080347 -2.167185 -2.167185 -2.167185
[157] -2.080347 -2.167185 -2.508487 -2.508487 -1.715876 -1.715876
[163] -2.508487 -1.715876 -2.080347 -2.508487 -14.427612 -19.584501
[169] -5.552565 -12.101565 -12.321482 -13.494799 -17.278981 -13.710475
[175] -11.286164 -17.472163 -7.626318 -9.024096 -8.372031 -15.859536
[181] -22.391255 -19.848818 -7.653704 -12.269797 -11.500951 -10.734978
[187] -8.777308 -12.886197 -9.783944 -11.406940 -14.185190 -11.113153
[193] -6.850939 -7.277518 -6.055092 -10.108391 -9.352544 -8.527166
[199] -14.288102 -11.053019 -13.897823 -11.957607 -8.864576 -2.062260
[205] -7.749086 -6.349919 -17.580597 -15.590796 -16.453472 -9.340225
[211] -15.602586 -10.269438 -13.840275 -14.573967 -6.498563 -2.080347
[217] -15.404740 -10.520493 -13.364853 -16.888469 -1.715876 -15.297375
```

[223]	-17.557815	-15.697284	-16.661305	-14.727623	-18.288867	-12.007639
[229]	-15.735110	-9.744544	-1.715876	-4.892642	-2.167185	-15.561990
[235]	-2.426825	-4.465175	-2.080347	-5.738164	-2.080347	-6.109724
[241]	-11.413849	-1.715876	-5.092905	-4.465175	-4.738262	-1.715876
[247]	-5.092905	-4.810437	-4.810437	-4.595481	-1.715876	-2.167185
[253]	-2.167185	-1.715876	-6.109724	-16.760204	-5.110062	-4.810437
[259]	-16.039684	-4.738262	-7.561318	-4.892642	-1.715876	-7.138108
[265]	-11.118977	-15.342746	-18.822633	-19.610132	-13.362261	-21.596448
[271]	-20.274354	-22.038136	-24.480470	-23.424857	-20.963407	-18.773643
[277]	-22.782456	-25.297190	-23.492093	-26.872406	-21.112862	-24.229300
[283]	-17.846469	-11.575613	-18.263239	-13.073302	-18.978337	-21.696945
[289]	-4.810437	-1.715876	-2.167185	-5.917834	-5.752411	-2.167185
[295]	-7.004105	-5.153175	-11.673362	-14.318881	-13.778521	-6.249467
[301]	-5.362389	-7.225713	-10.604252	-10.456408	-8.238050	-1.715876
[307]	-13.581309	-6.900542	-9.735685	-1.715876	-10.646801	-9.227094
[313]	-8.983879	-25.408828	-17.356342	-15.587094	-1.404772	-1.446734
[319]	-1.446734	-7.388083	-21.394655	-23.918267	-22.174177	-22.127758
[325]	-26.511420	-24.369232	-23.226896	-21.676464	-22.908750	-19.801430
[331]	-12.655255	-3.755978	-19.011252	-23.588347	-24.010499	-18.348843
[337]	-21.038007	-23.428279	-17.481979	-13.109775	-17.427216	-9.153633
[343]	-16.560149	-5.153175	-8.847971	-1.715876	-1.715876	-5.752411
[349]	-8.459037	-2.080347	-7.589702	-12.578330	-4.738262	-1.715876
[355]	-5.738164	-5.110062	-1.715876	-1.404772	-1.715876	-1.715876
[361]	-5.738164	-13.586797	-12.143921	-2.080347	-2.167185	-1.715876
[367]	-4.810437	-10.566812	-1.715876	-1.715876	-24.059782	-12.391613
[373]	-7.315316	-1.715876	-10.978917	-17.232579	-4.469938	-16.697932
[379]	-17.781011	-4.498594	-1.670524	-1.446734	-21.616535	-21.084619
[385]	-23.144063	-8.441417	-20.109347	-12.416362	-22.990176	-14.618806
[391]	-20.440186	-16.250487	-15.173105	-17.196065	-10.030076	-17.455080
[397]	-23.910734	-16.268053	-1.924175	-1.404772	-1.446734	-1.446734
[403]	-9.965890	-19.491949	-12.518061	-7.961238	-4.595481	-5.153175
[409]	-1.715876	-12.680365	-8.947220	-2.080347	-6.781348	-4.595481
[415]	-18.009432	-8.781208	-7.725505	-17.192164	-33.366038	-2.508487
[421]	0.000000	-9.420021	-9.156338	-13.969619	-21.409739	-29.017181
[427]	-21.351683	-19.888591	-22.017168	-17.352209	-18.196029	-18.119674
[433]	-21.328525	-30.392294	-15.002718	-25.144990	-23.248891	-24.201599
[439]	-22.609674	-13.744040	-19.290657	-16.296700	-1.715876	-7.845890
[445]	-7.639532	-15.889483	-19.060879	-4.892642	-1.715876	-1.715876
[451]	-7.845890	-1.715876	-5.633110	-16.050976	-6.621564	-1.584987
[457]	-1.584987	-2.021099	-4.595481	-1.715876	-5.110062	-2.167185
[463]	-1.715876	-13.339430	-2.167185	-2.167185	-1.715876	-2.508487
[469]	-1.715876	-2.508487	-2.508487	-2.080347	-9.532919	-10.579471
[475]	-4.595481	-2.508487	-12.039484	-2.508487	-1.715876	-2.167185
[481]	-1.715876	-1.715876	-2.508487	-4.595481	-1.715876	-2.080347
[487]	-8.968796	-2.508487	-1.715876	-2.080347	-2.508487	-2.167185
[493]	-2.080347	-2.167185	-1.446734	-1.446734	-1.715876	-2.167185
[499]	-1.715876	-1.715876	-4.595481	-2.080347	-2.508487	-2.508487
[505]	-1.715876	-2.167185	-2.167185	-5.110062	-2.167185	-2.167185
[511]	-2.508487	-2.508487	-1.715876	-1.715876	-1.715876	-2.080347
[517]	-4.465175	-2.080347	-2.508487	-4.595481	-1.715876	-1.715876
[523]	-2.167185	-1.715876	-1.715876	-2.508487	-2.167185	-1.715876

```

[529] -4.595481 -2.080347 -2.167185 -2.167185 -1.715876 -1.715876
[535] -10.874286 -5.917834 -4.595481 -2.080347 -2.508487 -1.715876
[541] -1.715876 -2.508487 -4.810437 -2.080347 -2.080347 -1.715876
[547] -2.508487 -2.167185 -1.715876 -2.508487 -2.508487 -1.715876
[553] -1.715876 -2.508487 -2.080347 -1.715876 -2.167185 -1.715876
[559] -6.498563 -1.715876 -2.508487 -1.715876 -1.715876 -2.080347
[565] -4.595481 -10.300653 -2.080347 -1.715876 -2.508487 -2.167185
[571] -2.508487 -4.892642 -17.926594 -2.080347 -2.080347 -4.595481
[577] -2.080347 -1.715876 -1.715876 -1.715876 -1.715876 -2.508487
[583] -1.715876 -17.768105 -11.001933 -1.715876 -7.026184 -2.080347
[589] -2.508487 -2.508487 -2.080347 -6.779694 -10.300653 -11.385137
[595] -1.715876 -2.080347 -2.167185 -11.339433 -2.167185 -1.715876
[601] -2.508487 -2.167185 -2.080347 -2.080347 -1.715876 -4.465175
[607] -1.715876 -1.715876 -4.465175 -2.167185 -18.305001 -2.167185
[613] -2.080347 -2.080347 -1.715876 -2.508487 -2.508487 -2.167185
[619] -1.715876 -1.715876 -1.715876 -1.715876 -2.508487 -2.508487
[625] -2.508487 -2.080347 -2.080347
> attr(tree, "score") # best score (in this case, the -LnL)
[1] 4363.674
> attr(tree, "model") # either the specified or automatically select transition model
[1] "GTR+G4"
> attr(tree, "parameters") # the free model parameters (or NA if unoptimized)
      FreqA      FreqC      FreqG      FreqT      FreqI      A/G      C/T      A/C
0.1749531 0.2454231 0.3442860          NA          NA 3.1520671 2.8618858 0.7178655
      A/T      C/G      Indels      alpha
1.0982960 0.5949656          NA 0.1919559
> attr(tree, "midpoint") # center of the edge (for plotting)
[1] 9.82373

```

The tree is (virtually) rooted at its midpoint by default. For maximum likelihood trees, all internal nodes include aBayes branch support values [1]. These are given as probabilities that can be used in plotting on top of each edge. We can also italicize the leaf labels (species names).

```

> plot(dendrapply(tree,
  function(x) {
    s <- attr(x, "probability") # choose "probability" (aBayes) or "support"
    if (!is.null(s) && !is.na(s)) {
      s <- formatC(as.numeric(s), digits=2, format="f")
      attr(x, "edgetext") <- paste(s, "\n")
    }
    attr(x, "edgePar") <- list(p.col=NA, p.lwd=1e-5, t.col="#CC55AA", t.cex=
    if (is.leaf(x))
      attr(x, "nodePar") <- list(lab.font=3, pch=NA)
    x
  })),
  horiz=TRUE,
  yaxt='n')
> # add a scale bar (placed manually)
> arrows(0, 0, 0.4, 0, code=3, angle=90, len=0.05, xpd=TRUE)
> text(0.2, 0, "0.4 subs./site", pos=3, xpd=TRUE)

```




Figure 2: Tree with (aBayes) support probabilities at each internal node.

We lucked out because maximum likelihood and maximum parsimony trees both provide branch supports in the form of the fraction of optimized trees that contained a given partition (branch). These are accessible from the “support” attribute. As luck would have it, support values and (aBayes) probabilities are correlated, but support tends to be more conservative.

```

> getSupports <- function(x) {
  if (is.leaf(x)) {
    NULL
  } else if (is.null(attr(x, "support"))) {
    rbind(getSupports(x[[1]]), getSupports(x[[2]]))
  } else {
    rbind(cbind(attr(x, "support"),
                attr(x, "probability"),
                attr(x, "members")),
          getSupports(x[[1]]), getSupports(x[[2]]))
  }
}
> support <- getSupports(tree)
> plot(support[, 1],
      support[, 2],
      cex=log10(support[, 3]),
      xlab="Support", ylab="aBayes probability", asp=1)
> abline(a=0, b=1, lty=2) # line of identity (y=x)

```




Figure 3: Comparison of aBayes probabilities and branch support values.

5 Ancestral State Reconstruction

We're in luck —one of the advantages of maximum likelihood and maximum parsimony tree building methods is that they automatically predict states at each internal node on the tree [2]. This feature is enabled when *reconstruct* is set to `TRUE`. These character states can be used by the function `MapCharacters` to determine state transitions along each edge of the tree. This information enables us to plot the total number of substitutions occurring along each edge. The state transitions can be accessed along each edge by querying a new “change” attribute.

```

> new_tree <- MapCharacters(tree, labelEdges=TRUE)
> plot(new_tree, edgePar=list(p.col=NA, p.lwd=1e-5, t.col="#55CC99", t.cex=0.7))
> attr(new_tree[[1]], "change") # state changes on first branch left of (virtual) root
 [1] "G168A" "G171T" "G176C" "G182C" "G184T" "G185T" "G199A" "G208C" "G214A"
[10] "G224T" "G225C" "G227C" "G229C" "C266T" "G272T" "G274C" "G276T" "G277T"
[19] "G279T" "G280T" "G283T" "G287C" "G288T" "G302C" "G314C" "G316T" "G321T"
[28] "G324T" "G325C" "G328T" "G330A" "G338T" "G371C" "G379A" "G385T" "G387C"
[37] "G389C" "G391C" "G394C" "G396T" "G397C" "G419A" "G440T" "G447T" "G477T"
[46] "G584C"

```




Figure 4: Edges labeled with the number of state transitions.

6 Exporting the Tree

We've had a run of good luck with this tree, so we'd better save it before our luck runs out! The functions `ReadDendrogram` and `WriteDendrogram` will import and export trees in Newick file format. If we leave the *file* argument blank then it will print the output to the console for our viewing:

```
> WriteDendrogram(tree, file="")
(("cattleya":0.2262257, ("bingchenggensis":0.09951155, ("clavuligerus":0.09547068, "pristi
```

To keep up our lucky streak, we should probably include any model parameters in the output along with the tree. Luckily, Newick format supports square brackets (i.e., “[]”) for comments, which we can append to the end of the file for good luck:

```
> params <- attr(tree, "parameters")
> cat("[", paste(names(params), params, sep="=", collapse=","), "]", sep=" ", append=TRUE
[FreqA=0.174953146897847, FreqC=0.245423105888693, FreqG=0.34428600457599, FreqT=NA, FreqI=N
```

7 Session Information

All of the output in this vignette was produced under the following conditions:

- R version 4.4.0 beta (2024-04-15 r86425), x86_64-pc-linux-gnu
- Running under: Ubuntu 22.04.4 LTS
- Matrix products: default
- BLAS: /home/biocbuild/bbs-3.19-bioc/R/lib/libRblas.so
- LAPACK: /usr/lib/x86_64-linux-gnu/lapack/liblapack.so.3.10.0
- Base packages: base, datasets, graphics, grDevices, methods, stats, stats4, utils
- Other packages: BiocGenerics 0.50.0, Biostrings 2.72.0, DECIPHER 3.0.0, GenomeInfoDb 1.40.0, IRanges 2.38.0, S4Vectors 0.42.0, XVector 0.44.0
- Loaded via a namespace (and not attached): compiler 4.4.0, crayon 1.5.2, DBI 1.2.2, GenomeInfoDbData 1.2.12, httr 1.4.7, jsonlite 1.8.8, KernSmooth 2.23-22, R6 2.5.1, tools 4.4.0, UCSC.utils 1.0.0, zlibbioc 1.50.0

References

- [1] Anisimova, M., Gil, M., Dufayard, J., Dessimoz, C., & Gascuel, O. Survey of branch support methods demonstrates accuracy, power, and robustness of fast likelihood-based approximation schemes. *Syst Biol.*, 60(5), 685-699.
- [2] Joy, J., Liang, R., McCloskey, R., Nguyen, T., & Poon, A. Ancestral Reconstruction. *PLoS Comp. Biol.*, 12(7), e1004763.