Estimating phylogenetic trees with phangorn (Version 1.3-1)

Klaus P. Schliep*

February 14, 2011

1   Introduction

These notes should enable the user to estimate phylogenetic trees from alignment data with different methods using the *phangorn* package [7]. For more background on all the methods see e.g. [2, 9]. This document illustrates some of the *phangorn* features to estimate phylogenetic trees using different reconstruction methods. Small adaptations to the scripts in section 6 should enable the user to perform phylogenetic analyses.

2   Getting started

The first thing we have to do is to read in an alignment. Unfortunately there exists many different file formats that alignments can be stored in. The function `read.phyDat` is used to read in an alignment. There are several functions to read in alignments depending on the format of the dataset (nexus, phylip, fasta) and the kind of data (amino acid or nucleotides) in the *ape* package [4] and *phangorn*. The function `read.phyDat` calls these other functions. For the specific parameter settings available look in the help files of the function `read.dna` (for phylip, fasta, clustal format), `read.nexus.data` for nexus files. For amino acid data additional `read.aa` is called. We start our analysis loading the *phangorn* package and then reading in an alignment.

```
> library(phangorn)
> primates = read.phyDat("primates.dna", format = "phylip",
+       type = "DNA")
```

---

*mailto:kschliep@snv.jussieu.fr

3   Distance based methods

After reading in the alignment we can build a first tree with distance based methods. The function dist.dna from the ape package computes distances for many DNA substitution models. To use the function dist.dna we have to transform the data to class DNAbin. For amino acids the function dist.ml offers common substitution models ("WAG", "Dayhoff", "JTT" and "LG"). After constructing a distance matrix we reconstruct a rooted tree with UPGMA and alternatively an unrooted tree using Neighbor Joining [6, 8].

```
> dm = dist.dna(as.DNAbin(primates))
> treeUPGMA = upgma(dm)
> treeNJ = NJ(dm)
```

We can plot the trees treeUPGMA and treeNJ (figure 1) with the commands:

```
> par(mfrow = c(1, 2), mar = c(1, 1, 4, 1))
> plot(treeUPGMA, main = "UPGMA")
> plot(treeNJ, "unrooted", main = "NJ")
```

Distance based methods are very fast and we will use the UPGMA and NJ tree as starting trees for the maximum parsimony and maximum likelihood analyses.

4   Parsimony

The function parsimony returns the parsimony score, that is the number of changes which are at least necessary to describe the data for a given tree. We can compare the parsimony score or the two trees we computed so far:

```
> parsimony(treeUPGMA, primates)
[1] 751
> parsimony(treeNJ, primates)
[1] 746
```

The function optim.parsimony performs tree rearrangements to find trees with a lower parsimony score. So far the only tree rearrangement implemented is nearest-neighbor interchanges (NNI). However is also a version of the parsimony ratchet [3] implemented, which is likely to find better trees than just doing NNI rearrangements.

```
> treePars = optim.parsimony(treeUPGMA, primates)
Final p-score 746 after  1 nni operations
> treeRatchet = pratchet(primates, trace = 0)
> parsimony(c(treePars, treeRatchet), primates)
[1] 746 746
```

Figure 1: Rooted UPGMA tree and unrooted NJ tree

5 Maximum likelihood

The last method we will describe in this vignette is Maximum Likelihood (ML) as introduced by Felsenstein [1]. We can easily compute the likelihood for a tree given the data

```
> fit = pml(treeNJ, data = primates)
> fit
 loglikelihood: -3077.846

unconstrained loglikelihood: -1230.335

Rate matrix:
  a c g t
a 0 1 1 1
c 1 0 1 1
g 1 1 0 1
t 1 1 1 0

Base frequencies:
0.25 0.25 0.25 0.25
```

The function pml returns an object of class pml. This object contains the data, the tree and many different parameters of the model like the likelihood etc. There are many generic functions for the class pml available, which allow the handling of these objects.

```
> methods(class = "pml")
[1] anova.pml*  logLik.pml* plot.pml*   print.pml*  update.pml*
[6] vcov.pml*

    Non-visible functions are asterisked
```

The object fit just estimated the likelihood for the tree it got supplied, but the branch length are not optimized for the Jukes-Cantor model yet, which can be done with the function optim.pml.

```
> fitJC = optim.pml(fit, TRUE)
> logLik(fitJC)
```

With the default values pml will estimate a Jukes-Cantor model. The function update.pml allows to change parameters. We will change the model to the GTR + $\Gamma(4)$ + I model and then optimize all the parameters.

```
> fitGTR = update(fit, k = 4, inv = 0.2)
> fitGTR = optim.pml(fitGTR, TRUE, TRUE, TRUE, TRUE, TRUE,
+       control = pml.control(trace = 0))
> fitGTR
 loglikelihood: -2609.581

unconstrained loglikelihood: -1230.335
Proportion of invariant sites: 0.006050102
Discrete gamma model
Number of rate categories: 4
Shape parameter: 3.172213

Rate matrix:
            a             c           g            t
a  0.0000000  0.646515270 33.645440903  0.4051026
c  0.6465153  0.000000000  0.008314716 14.3798327
g 33.6454409  0.008314716  0.000000000  1.0000000
t  0.4051026 14.379832726  1.000000000  0.0000000

Base frequencies:
0.3917253 0.379635 0.04027305 0.1883667
```

We can compare the objects for the JC and GTR + $\Gamma(4)$ + I model using likelihood ratio statistic

```
> anova(fitJC, fitGTR)
Likelihood Ratio Test Table
  Log lik. Df Df change Diff log lik. Pr(>|Chi|)
1  -3068.3 25
2  -2609.6 35        10         917.43  < 2.2e-16 ***
---
Signif. codes:  0 âĂŸ***âĂŹ 0.001 âĂŸ**âĂŹ 0.01 âĂŸ*âĂŹ 0.05 âĂŸ.âĂŹ 0.1 âĂŸ âĂŹ 1
```

with the AIC

```
> AIC(fitGTR)
[1] 5289.162
> AIC(fitJC)
[1] 6186.59
```

or the Shimodaira-Hasegawa test.

```
> SH.test(fitGTR, fitJC)
```
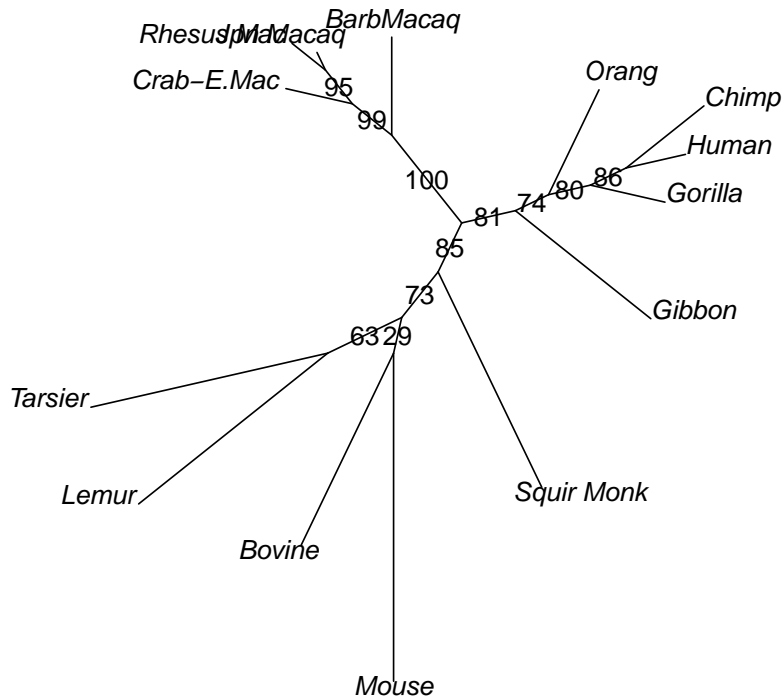
Figure 2: Unrooted tree with bootstrap support values

```
        Trees       ln L Diff ln L p-value
   [1,]     1 -2609.581    0.0000  0.5001
   [2,]     2 -3068.295  458.7142  0.0000
```

At last we may want to apply bootstrap to test which how well the edges of the tree are supported:

```
> bs = bootstrap.pml(fitJC, bs = 100, optNni = TRUE, control = pml.control(trace =
```

At last we plot the tree with the bootstrap support values on the edges

```
> par(mar = c(0.1, 0.1, 0.1, 0.1))
> plotBS(fitJC$tree, bs)
```

The bootstrap analysis can be computationally demanding, but on UNIX systems the bootstrap functions will distributed the computations using the *multicore* package.

6 Appendix: Standard scripts for nucleotide or amino acid analysis

Here we provide two standard scripts which can be adapted for the most common tasks. Most likely the arguments for `read.phyDat` have to be adapted to accommodate your file format.

```
library(multicore)
library(phangorn)
file = "myfile"
dat = read.phyDat(file)
dm = dist.ml(dat)
tree = NJ(dm)
fitNJ = pml(tree, dat, k = 4, inv = 0.2)
fit = optim.pml(fitNJ, TRUE, TRUE, TRUE, TRUE, TRUE)
fit
bs = bootstrap.pml(fit, bs = 100, optNni = TRUE)
```

You can specify different several models build in which you can specify "WAG", "JTT", "Dayhoff", "LG". Optimising the rate matrix for amino acids is possible, but would take a long, a very long time. So make sure to set optBf=FALSE and optQ=FALSE in the function `optim.pml`, which is also the default.

```
library(multicore)
library(phangorn)
file = "myfile"
dat = read.phyDat(file, type = "AA")
dm = dist.ml(dat, model = "JTT")
tree = NJ(dm)
fitNJ = pml(tree, dat, model = "JTT", k = 4, inv = 0.2)
fit = optim.pml(fitNJ, optNni = TRUE, optInv = TRUE,
    optGamma = TRUE)
fit
bs = bootstrap.pml(fit, bs = 100, optNni = TRUE)
```

References

[1] Joseph Felsenstein. Evolutionary trees from dna sequences: a maxumum likelihood approach. *Journal of Molecular Evolution*, 17:368–376, 1981.

[2] Joseph Felsenstein. *Inferring Phylogenies*. Sinauer Associates, Sunderland, 2004.

[3] K.~Nixon. The parsimony ratchet, a new method for rapid rarsimony analysis. *Cladistics*, 15:407–414, 1999.

[4] E.~Paradis, J.~Claude, and K.~Strimmer. Ape: Analyses of phylogenetics and evolution in r language. *Bioinformatics*, 20(2):289–290, 2004.

[5] Emmanuel Paradis. *Analysis of Phylogenetics and Evolution with R.* Springer, New York, 2006.

[6] N.~Saitou and M.~Nei. The neighbor-joining method - a new method for reconstructing phylogenetic trees. *Molecular Biology and Evolution*, 4(4):406–425, 1987.

[7] Klaus~Peter Schliep. phangorn: Phylogenetic analysis in R. *Bioinformatics*, 27(4):592–593, 2011.

[8] J.~A. Studier and K.~J. Keppler. A note on the neighbor-joining algorithm of saitou and nei. *Molecular Biology and Evolution*, 5(6):729–731, 1988.

[9] Ziheng Yang. *Computational Molecular evolution.* Oxford University Press, Oxford, 2006.

7 Session Information

The version number of R and packages loaded for generating the vignette were:

- R version 2.12.1 (2010-12-16), `x86_64-pc-linux-gnu`

- Locale: `LC_CTYPE=en_US.utf8`, `LC_NUMERIC=C`, `LC_TIME=en_US.utf8`, `LC_COLLATE=C`, `LC_MONETARY=C`, `LC_MESSAGES=en_US.utf8`, `LC_PAPER=en_US.utf8`, `LC_NAME=C`, `LC_ADDRESS=C`, `LC_TELEPHONE=C`, `LC_MEASUREMENT=en_US.utf8`, `LC_IDENTIFICATION=C`

- Base packages: base, datasets, grDevices, graphics, grid, methods, stats, utils

- Other packages: ape~2.6-2, phangorn~1.3-1, quadprog~1.5-3, seqLogo~1.16.0

- Loaded via a namespace (and not attached): gee~4.13-16, lattice~0.19-17, nlme~3.1-98, tools~2.12.1