

Exercise 1:

- (a) Implement Felsenstein's pruning algorithm in R as a function that reads a tree in newick format, a sequence alignment and parameters of the F84 model and returns the log-likelihood of the tree. You can use the R script `pruning.R` as a starting point.
- (b) Discuss how your implementation can be made more efficient.

Exercise 2: Based on the fossil record for lizards it is claimed that the most recent common ancestor (MRCA) of *Sphenodon punctatus* and *Cnemidophorus tigris* is at least 228.0 Mya old, the MRCA of *Cnemidophorus tigris* and *Rhineura floridana* is at least 113 Mya old, the MRCA of *Rhineura floridana* and *Gallotia galloti* at least 64.2 Mya, and the MRCA of *Timon lepidus* and *Dalmatolacerta oxycephala* at least 5.3 Mya.

- (a) Use this information to estimate the phylogeny of lizards with time calibration.
- (b) How does this result depend on the choice of model for the relaxed clock?
- (c) How does the result change if you use only two or three of the given calibration points?
- (d) How well are the given calibration points supported by fossils documented in the following references?

Science (1990) 249:1020-1023,

J Vertebr Paleontol (2006) 26:795-800,

J Vertebr Paleontol (2002) 22:286-298,

J Paleontol (1985) 59:1481-1485

Exercise 3: Analyze the primates data data with `beast` and with `RAXML` with bootstrapping and compare the results, that is the tree topologies, the branch lengths, and the posterior probabilities vs. bootstrap values. Do the same also for the lizards data. Explore the role of priors in the Bayesian analyses.

Exercise 4: Assume that 10000 trees haven been MCMC-sampled from the posterior distribution for a given data set. To summarize these trees we would like to show a single tree that has all branches that appear in more than 5000 of the sampled trees (as splits of the taxa sets). Proof that this is always possible or describe a counter-example.

Exercise 5: We flip a coin 1000 times to test whether it is fair. If we apply Bayesian inference with a uniform prior on $[0, 1]$ for the probability p that the coin shows "head", how probable is it that we estimate the probability of $\{p > 0.5\}$ to be higher than 90% or lower than 10%?