15. January 2013

Exercise 1: Simulate trees from several prior distributions and let sequences evolve along the trees. Estimate the trees from the data by maximum-likelihood methods and apply bootstrap with PHYLIP and/or RAxML. Repeat this many times and explore whether the bootstrap support values are good approximations of the probabilities that branches are correct.

Exercise 2: Simulate data for some example trees and explore how the type-I-errors and the type-II-errors of the KH test and the SH test (as implemented in the ape package in R) depend on the amount of available data and on the branch lengths of the trees.

Exercise 3: It is known that the balls in an urn are numbered from 1 to n, but there are two Hypotheses about n. H_0 says that $n = 100$, and H_A says that n could be any positive number. Ball number 99 is drawn from the urn.

(a) Apply hypothesis testing to decide whether H_0 can be rejected on the 5% level.

(b) Apply a Bayes factor analysis to decide between the two hypotheses. For H_A assume a uniform prior on $\{1, 2, \ldots, N\}$ for n. How does the result depend on N?

Exercise 4: The two values 1.0 and 1.1 are independently drawn from the same normal distribution $\mathcal{N}(\mu, \sigma^2)$ with unknown μ and unknown σ^2. We would like to know whether $\mu = 0$.

(a) Apply a t-test to address this question.

(b) Compute Bayes-factors to compare the following two hypotheses

$H_0 :$ $\mu = 0$ and σ^2 has a uniform prior on $[0, s]$

$H_A :$ μ has a uniform prior on $[-m, m]$ and σ^2 has a uniform prior on $[0, s]$

How does the decision for one or the other model depend on m and s?