

20. November 2013

Exercise 1: Assume that 10000 trees have 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 2: 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%?

Exercise 3: Write a computer program that reads a tree in newick format and simulates the evolution of quantitative traits along the branches of the tree. The output should be the values of the traits at the tips of the tree.

Exercise 4: The files QuantTraitsA.csv, QuantTraitsB.csv, and QuantTraitsC.csv contain datasets of quantitative traits for species whose phylogeny is given in the file QuantTraits_Tree.txt. Do you find evidence for selection and/or for correlated evolution of some of the traits?

Exercise 5: Assume that 200 different traits have been measured for 50 species of known phylogeny. For only a few pairs of traits there is some evidence for correlated evolution. How would you measure the significance of this evidence?

Exercise 6: 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 $[0, s]$

$H_A : \mu$ 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 ?