Exit frequency matrices for finite Markov chains

Andrew Beveridge

Abstract

We think of a Markov chain as a random walk on a (directed) graph. If we consider such a random walk in reverse, the result is a random walk on an associated Markov chain, called the *reverse chain*. We develop a framework that describes how this duality extends to *stopping rules*. A stopping rule is an intelligent procedure which ?looks where it is going? to sample exactly from any desired distribution. Fixing a target distribution τ, we consider a family of stopping rules, one from each possible starting node. We show that this family is dual to a family of rules on the reverse chain to an associated distribution τ^*. The key to unlocking this duality is to partition the random walks into *exit frequencies*, which are the expected number of exits at each node.