

Chapter 24

Markov chain Monte Carlo (MCMC)inference

24.1 Introduction

In Chapter 23, we introduced some simple Monte Carlo methods, including rejection sampling and importance sampling. The trouble with these methods is that they do not work well in high dimensional spaces. The most popular method for sampling from high-dimensional distributions is **Markov chain Monte Carlo** or **MCMC**.

The basic idea behind MCMC is to construct a Markov chain (Section 17.2) on the state space \mathcal{X} whose stationary distribution is the target density $p^*(\boldsymbol{x})$ of interest (this may be a prior or a posterior). That is, we perform a random walk on the state space, in such a way that the fraction of time we spend in each state \boldsymbol{x} is proportional to $p^*(\boldsymbol{x})$. By drawing (correlated!) samples $\boldsymbol{x}_0, \boldsymbol{x}_1, \boldsymbol{x}_2, \dots$ from the chain, we can perform Monte Carlo integration wrt p^* .

24.2 Metropolis Hastings algorithm

24.3 Gibbs sampling

24.4 Speed and accuracy of MCMC

24.5 Auxiliary variable MCMC *