

Homework 12

Exercise 1. [Gibbs sampling or heat bath dynamics]

Let $S = \{1, \dots, N\}$ and $d \geq 1$. We would like to sample from a distribution π on S^d defined as

$$\pi(x) = \frac{g(x)}{Z}, \quad x \in S^d,$$

where g is some positive function on S^d and $Z = \sum_{x \in S^d} g(x)$ is the normalization constant, which we would like to avoid computing.

A possible way to handle this problem is the following.

1. Start from a state $x \in S^d$;
2. Choose an index $u \in \{1, \dots, d\}$ uniformly at random;
3. Update the value of x_u to x'_u , which is sampled from the following conditional distribution:

$$\pi(x'_u | x_1, \dots, x_{u-1}, x_{u+1}, \dots, x_d) = \frac{\pi(x_1, \dots, x_{u-1}, x'_u, x_{u+1}, \dots, x_d)}{\sum_{y_u \in S} \pi(x_1, \dots, x_{u-1}, y_u, x_{u+1}, \dots, x_d)}$$

4. Repeat from 2.

What is the advantage of such a method? The above conditional probability can actually be rewritten as

$$\pi(x'_u | x_1, \dots, x_{u-1}, x_{u+1}, \dots, x_d) = \frac{g(x_1, \dots, x_{u-1}, x'_u, x_{u+1}, \dots, x_d)}{\sum_{y_u \in S} g(x_1, \dots, x_{u-1}, y_u, x_{u+1}, \dots, x_d)}$$

which only requires to compute one sum and not a multidimensional one, as required for computing the normalization constant Z .

Your task now is to formalize slightly the above algorithm by expressing it as a Markov chain $(X_n, n \geq 0)$ on S^d and

- a) writing down its transition probabilities $p(x, y)$, $x, y \in S^d$;
- b) showing that the detailed balance equation is satisfied, i.e. that $\pi(x)p(x, y) = \pi(y)p(y, x)$, for all $x, y \in S^d$.

Can therefore this algorithm be viewed as a Metropolis-Hastings algorithm?