

Tutorial 3, Advanced MCMC

2021/22 ICFP Master (second year)

Werner Krauth

ENS Paris

(Dated: 01 February 2023)

1. Direct sampling for hard spheres

In lecture 3 (part 3.1), we discussed a direct-sampling algorithm for four disks in a box (In week 4, we will treat a few far more successful direct-sampling algorithms for the same problem).

- (a) Prove (formally) that the *tabula-rasa* algorithm presented in the lecture samples the uniform distribution in $[0, 1]^8 \setminus \cup_{i < j} Q_{ij}$.
- (b) Prove that the alternative to the *tabula-rasa* algorithm, where you only resample the “offending” last disk, is incorrect.

2. Non-reversible hard-sphere algorithms in 1d

In lecture 3 (part 3.2), we discussed a number of reversible and non-reversible Markov chains for hard spheres in one dimension, and among them the sequential MCMC (moving disk 1 then 2, then 3... , then N , then 1...)

- (a) Precisely define the lifting that is realized in that case.
- (b) Is the sequential MCMC a reversible or a non-reversible lifting?
- (c) Prove again (as we did in the lecture) that the sequential hard-sphere MCMC satisfies global balance.
- (d) Prove that for an N -particle system, and for a single-particle-move algorithm that satisfies detailed balance, the particle-lifted sequential algorithm satisfies global balance.

3. Severely lifted non-reversible hard-sphere algorithms in 1d

In lecture 3 (part 3.2), we discussed the forward Metropolis algorithm for a one-dimensional hard-sphere system.

- (a) Precisely define the lifting that is realized in that case.

- (b) Prove again (as we did in the lecture) that the forward Metropolis algorithm satisfies global balance.
- (c) Do the same for the sequential forward algorithm. Prove that this variant of the algorithm violates the global-balance condition (sequential forward: move forward, but sequentially 1 then 2, ...).

4. Mixing times of $1d$ hard-sphere MCMC

At the end of lecture 3 (part 3.2), we discussed mixing times for non-reversible MCMC algorithms. We did not use the TVD to track the convergence but found an alternative in the “half-system-distance” distribution.

- (a) Program one or two of the non-reversible MCMC algorithms and convince yourself that the half-system distance converges to the same limit (for $t \rightarrow \infty$) in all cases.
- (b) Show that the mixing is much faster for the forward Metropolis algorithm than for the reversible Metropolis algorithm.