Advanced topics in Markov-chain Monte Carlo

Lecture 8:
Meta algorithms, consensus sampling
Part 2/2: Consensus sampling

Werner Krauth

ICFP -Master Course Ecole Normale Supérieure, Paris, France

16 March 2022

W. Krauth Frontiers in Physics (2022)

Lifted MCMC in one dimension

Probability distribution $\pi = (1/n, \ldots, 1/n)$ (Diaconis et al. 2000)

- Transport + resampling

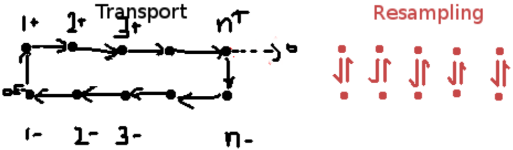

- "Lifted" Markov chain $\hat{\Omega} = \Omega \times \{-, +\}$:

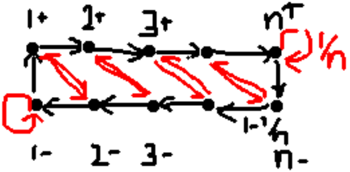

Lifted MCMC in higher dimensions (1/5)

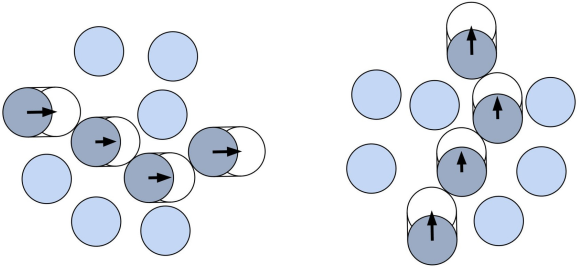

- Infinitesimal moves (avoid overlaps).
- Particle lifting + direction lifting.
- This algorithm is correct without return moves.

Lifted MCMC in higher dimensions (2/5)

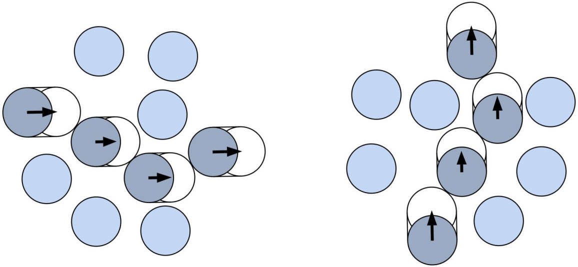

- Consensus sampling (trivial case)

Lifted MCMC in higher dimensions (3/5)

- Metropolis filter:

$$p^{\text{Met}}(a \to b) = \min \left[1, \prod_{i<j} \exp\left(-\beta \Delta U_{i,j}\right) \right]$$

- Factorized Metropolis filter:

$$p^{\text{Fact.}}(a \to b) = \prod_{i<j} \min\left[1, \exp\left(-\beta \Delta U_{i,j}\right)\right].$$

  satisfies detailed-balance condition, with a symmetric choice of a priori probabilities.

- Interpretation in terms of Boolean random variables.

$$X^{\text{Fact.}}(a \to b) = X_{1,2} \wedge X_{1,3} \wedge \cdots \wedge X_{N-1,N}$$

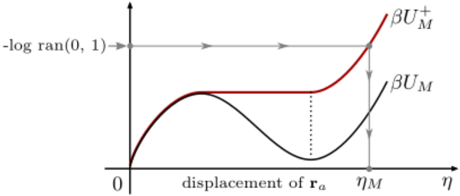

- $$p_M(m) = \underbrace{\prod_{l=1}^{m-1} e^{-\beta\,\Delta U_M^+(l)}}_{\text{accepted}} \overbrace{\left[1 - e^{-\beta\,\Delta U_M^+(m)}\right]}^{\text{move } m \text{ rejected}}, \qquad (1)$$
- Peters, de With (2012)

Lifted MCMC in higher dimensions (5/5)

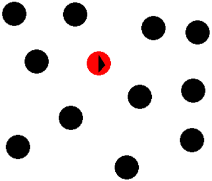

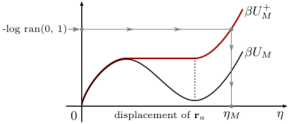

- Compute event-times for all factors.
- Select smallest one.
- Complexity $\mathcal{O}(N)$ unless short-range.

Thinning and sampling (1/1)

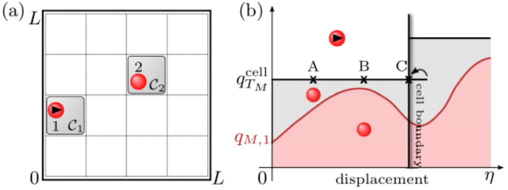

- Time-dependent Poisson process $q(x)\mathrm{d}x$

$$\underbrace{q(x)\mathrm{d}x}_{\text{variable}} = \underbrace{q^{\max}\mathrm{d}x}_{\text{constant}} \underbrace{\frac{q(x)}{q^{\max}}}_{\text{rejection}}$$

- This is called “Thinning”, many generalizations.
- ... reduces ECMC to the sampling from a vector of probabilities, one for each cell.

Sampling from a discrete distribution (1/3)

- Rejection sampling (see book)

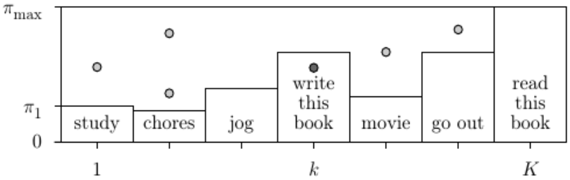

Fig. 1.28 Saturday night problem solved by Alg. 1.13 (`reject-finite`).

Sampling from a discrete distribution (2/3)

- Tower sampling (see book)

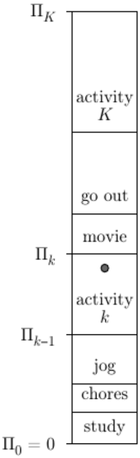

Fig. 1.29 Saturday night problem solved by tower sampling.

Sampling from a discrete distribution (3/3)

- Walker's method of aliases

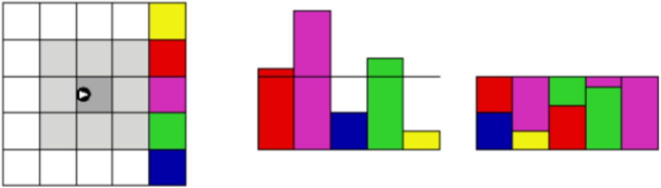

.. Cell-veto sampling using Walker's method.

- Complexity $\mathcal{O}(1)$