[Andrieu, Doucet & Holenstein, 2010]

Introduce algorithms that use SMC proposals in MCMC

Given a target distribution $\pi(x)$ on $\mathcal{X}$ (usually high-dimensional space), assume that we can run an SMC algorithm that returns an n-sample $X^1, \ldots, X^n$ together with associated unnormalized weights $\omega^1, \ldots, \omega^n$; let $Z = 1/n \sum_{j=1}^n \omega^j$.

Particle Independent Metropolis-Hastings (PIMH) algorithm

Given the current state (X_i, Z_i),

1. Run the SMC algorithm to obtain $\bar{X}^1, \ldots, \bar{X}^n$, $\bar{\omega}^1, \ldots, \bar{\omega}^n$, $\bar{Z} = \sum_{k=1}^n \bar{\omega}^k$.
2. Draw an index $\bar{K}$ in $\{1, \ldots, n\}$ with probability $\mathrm{P}(\bar{K} = k) = \bar{\omega}^k / \sum_{j=1}^n \bar{\omega}^j$.
3. Accept $X_{i+1} = \bar{X}^{\bar{K}}$, $Z_{i+1} = \bar{Z}$ with probability $1 \wedge \bar{Z}/Z_i$ (otherwise stay in X_i, Z_i).

Proof [Assuming that $(\bar{X}^j, \bar{\omega}^j)^{j=1,\dots,n}$ is produced by direct importance sampling with instrumental pdf q (i.e. no resampling)]

1. The important idea is that the state of the chain is in fact $(X^1, \dots, X^n, K)$ and that the targeted auxiliary distribution has pdf

$$\pi^{aux}(x^1, \dots, x^n, k) = \frac{1}{n}\pi(x^k) \prod_{j \neq k} q(x^j)$$

2. Check that indeed X^K has marginal distribution π under π^{aux}
3. The proposal distribution is independent of the current state and is given by

$$q^{aux}(\bar{x}^1, \dots, \bar{x}^n, \bar{k}) = \frac{\pi(\bar{x}^{\bar{k}})/q(\bar{x}^{\bar{k}})}{\sum_l \pi(\bar{x}^l)/q(\bar{x}^l)} \prod_j q(\bar{x}^j)$$

4. The Metropolis-Hastings acceptance ratio is given by

$$\frac{\pi^{aux}(\bar{x}^1,\ldots,\bar{x}^n,\bar{k})}{\pi^{aux}(x^1,\ldots,x^N,k)}\frac{q^{aux}(x^1,\ldots,x^n,k)}{q^{aux}(\bar{x}^1,\ldots,\bar{x}^N,\bar{k})}=\frac{\bar{z}}{z}$$

as

$$\begin{aligned}\frac{\pi^{aux}(\bar{x}^1,\ldots,\bar{x}^n,\bar{k})}{q^{aux}(\bar{x}^1,\ldots,\bar{x}^n,\bar{k})} &= \frac{\frac{1}{n}\pi(\bar{x}^{\bar{k}})\prod_{j\neq\bar{k}}q(\bar{x}^j)}{\frac{\pi(\bar{x}^{\bar{k}})/q(\bar{x}^{\bar{k}})}{\sum_l \pi(\bar{x}^l)/q(\bar{x}^l)}\prod_j q(\bar{x}^j)}\\ &= \frac{1}{n}\sum_l \pi(\bar{x}^l)/q(\bar{x}^l)=\bar{z}\end{aligned}$$

5. Keeping the whole state $(X^1,\ldots,X^n,K)$ is not required and one only needs to keep track of X^K and Z

The method can defeat the *curse of dimensionality* (at a certain price...)

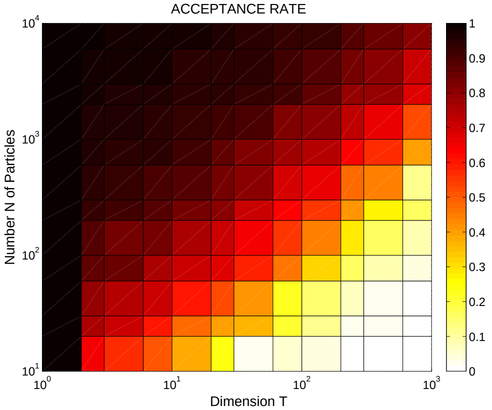

Figure: PIMH acceptance rate as a function of the dimension T of the target and the number N of particles. The target pdf is $\pi_T(x_1, \ldots, x_T) = \prod_{t=1}^{T} \pi(x_t)$, where π is the normal pdf truncated to the range $[-4, 4]$; the SMC proposal "kernel" q is an independent proposal, uniformly distributed in the range $[-4, 4]$. To assess the difficulty of the simulation task, note that for direct self-normalized importance sampling targeting π_T the Effective Sample Size (ESS) statistic, normalized by N, tends to 2.26^{-T} $(2.26 = \int_{-4}^{4} 8\pi^2(x)dx)$ as N increases, which is about 10^{-6} for $T = 17$.