1 Umbrella sampling

We have already shown how to use several runs at $\beta_1 < \ldots < \beta_R$ to compute averages for any β in the interval $[\beta_1, \beta_R]$ and to compute free energy differences. The umbrella sampling (US) method was introduced by Torrie and Valleau in 1977 to perform the same tasks by means of a single simulation. The idea consists in performing MC simulations with a non-Boltzmann-Gibbs distribution function of the form

$$\pi(x) = \frac{1}{Z_\pi} \sum_{i=1}^{R} \alpha_i e^{-\beta_i H(x)} \,, \tag{1}$$

where i runs over the R different temperatures, α_i are positive constants that should be carefully chosen as described below, and Z_π is the normalizing factor. Since $\sum_x \pi(x) = 1$, we have

$$Z_\pi = \sum_i \alpha_i Z_i.$$

By sampling the distribution (1) one aims at sampling in a single run the configurations that are typical for all β_i's and, as a consequence, all configuration space which is relevant for the computation of $\langle A \rangle_\beta$ with $\beta_1 \leq \beta \leq \beta_R$. In order for the method to work two requirements should be satisfied:

- The temperatures should be finely spaced, so that typical configurations at inverse temperature β_i overlap with those at $\beta_{i\pm1}$. If this does not occur, the system is unable to move in configuration space and does not visit the typical configuration domain of all β_i's. This condition is the same that occurs in the application of the data reweighting method.

- A second important condition fixes the coefficients α_i or, more precisely, their ratios. We require that the typical configuration domains at each β_i have approximately the same probability under π. For this purpose we require

$$\frac{1}{Z_\pi}\alpha_i Z_i \approx \frac{1}{Z_\pi}\alpha_j Z_j \qquad \Rightarrow \qquad \frac{\alpha_i}{\alpha_j} \approx \frac{Z_j}{Z_i} = e^{\beta_i F(\beta_i) - \beta_j F(\beta_j)} \,. \tag{2}$$

  The algorithm is correct, though not optimal, for any choice of the α_i's, so that it is enough to have a very rough estimate of the free-energy differences to run a US simulation.

Note that we only fix the ratios of the α_i's: this is not a limitation since one can always set, say, $\alpha_1 = 1$, by redefining Z_π.
Of course, the method is useful only if we are able to obtain canonical estimates at a given β. Let us define

$$f_\pi(x) = \sum_i \alpha_i e^{-\beta_i H(x)}$$

so that $\pi(x) = f_\pi(x)/Z_\pi$. Then, we have

$$\langle A \rangle_\beta = \frac{\sum_x A(x) e^{-\beta H(x)}}{\sum_x e^{-\beta H(x)}} = \frac{\frac{1}{Z_\pi}\sum_x f_\pi(x)[A(x)e^{-\beta H(x)}/f_\pi(x)]}{\frac{1}{Z_\pi}\sum_x f_\pi(x)[e^{-\beta H(x)}/f_\pi(x)]}$$

We obtain therefore

$$\langle A \rangle_\beta = \frac{\langle A\, e^{-\beta H}/f_\pi \rangle_\pi}{\langle e^{-\beta H}/f_\pi \rangle_\pi} \,. \tag{3}$$

The umbrella approach is very general and can be applied to any system. In general we consider R different probability distributions π_i defined on the state space and consider

$$\pi(x) = \sum_i a_i \pi_i(x)$$

with $a_i > 0$ and $\sum_i a_i = 1$. In the previous example

$$\pi_i = \frac{1}{Z_i} e^{-\beta_i H} \qquad a_i = \frac{Z_i}{Z_\pi} \alpha_i.$$

The optimality condition (2) translates in

$$\frac{\alpha_i}{\alpha_j} \approx \frac{Z_j}{Z_i} \Rightarrow \frac{Z_\pi a_i}{Z_i} \frac{Z_j}{Z_\pi a_j} \approx \frac{Z_j}{Z_i} \Rightarrow \frac{a_i}{a_j} \approx 1.$$

The weights a_i should be approximately the same to have an optimal distribution.

2 Monte Carlo implementation

The umbrella sampling method can be implemented straightforwardly. Let us consider a system of monoatomic molecules. Suppose that particles are in $\{\mathbf{r}_1, \ldots, \mathbf{r}_N\}$. A Metropolis update of one particle is performed as in the canonical case. We choose a molecule i and propose a shift in a cubic box of size Δ^3. The move is accepted with the Metropolis acceptance probability min(1,R), where

$$R = \frac{\sum_i \alpha_i e^{-\beta_i E_{\rm new}}}{\sum_i \alpha_i e^{-\beta_i E_{\rm old}}}$$

and $E_{\rm new}$ and $E_{\rm old}$ are the energies of the new and of the old configuration. This means that we accept the move if $E_{\rm new} < E_{\rm old}$ (check that, if this condition holds, $R > 1$); otherwise, we accept the move with probability R.