

Today's lecture: measuring free energies along macroscopic degrees of freedom and using reweighting concepts to make predictions at other state conditions post-simulation

Generalization of reweighting to other free energies

In our previous protein-surface example, we measured the free energy along a reaction coordinate of interest, $F(z)$. It is easy to see how this could be extended to other molecular reaction coordinates, such as orientation, radius of gyration, etc. to measure a PMF. However, we can use similar approaches to determine free energies associated with macroscopic or ensemble parameters.

Macroscopic free energies

We may be concerned with the Helmholtz free energy or Gibbs free energy. We might compute changes in these as a function of their natural variables. For single-component systems:

$$A(T, V, N)$$

$$G(T, P, N)$$

For multicomponent systems,

$$A(T, V, N_1, \dots, N_M)$$

$$G(T, P, N_1, \dots, N_M)$$

Typically we are only interested in the dependence of these free energies along a single parameter, e.g.,

$$A(V), G(P), G(T), \text{ etc.}$$

for constant values of the other independent variables.

Free energies for changes in the interaction potential

It is also possible to define a free energy change associated with a change in the interaction potential. Initially the energy function is $U_0(\mathbf{r}^N)$ and we perturb it to $U_1(\mathbf{r}^N)$. If this change happens in the canonical ensemble, we are interested in the free energy associated with this perturbation:

$$\Delta A = A_1(T, V, N) - A_0(T, V, N)$$

$$= -k_B T \ln \left(\frac{\int e^{-\beta U_1(\mathbf{r}^N)} d\mathbf{r}^N}{\int e^{-\beta U_0(\mathbf{r}^N)} d\mathbf{r}^N} \right)$$

What kinds of states 1 and 0 might we use to evaluate this expression? Here is a small number of sample applications:

- **electrostatic free energy** – charging of an atom or atoms in a molecule, in which state 0 has zero partial charges and state 1 has finite values
- **dipolar free energy** – adding a point dipole to an atom between states 0 and 1
- **solvation free energy** – one can “turn on” interactions between a solvent and a solute as a way to determine the free energy of solvation
- **free energy associated with a field** – states 0 and 1 correspond to the absence and presence, respectively, of a field, such as an electrostatic field
- **restraint free energy** – turning on some kind of restraint, such as confining a molecule to have a particular conformation or location in space. Such restraints would correspond to energetic penalties for deviations from the restrained space in state 1.
- **free energies of alchemical transforms** – we convert one kind of molecule (e.g., CH₄) to another kind (e.g., CF₄). This gives the relative free energies of these two kinds of molecules in the system of interest (e.g., solvation free energies in solution).

Sampling considerations

In each of the cases above, we must sample along the coordinate of interest to compute the free energy. Unlike the case of molecular coordinates, with macroscopic coordinates we do not need to use umbrella sampling or other biasing techniques because we *choose* exactly the ensemble conditions like T, P, V and the force field $U(\mathbf{r}^N)$ in the simulation setup. Thus, to use reweighting techniques to determine free energies, we perform multiple simulations at different values of our coordinate, and post process the data using similar statistically optimal reweighting equations.

In general,

- To compute a free energy or entropy as a function of some macroscopic parameter X , the simulation must accomplish fluctuations in X . For example, $A(V)$ requires fluctuations in V and thus would necessitate use of the NPT ensemble.

- To have good statistical accuracy, a wide range of fluctuations in X must be sampled. This can be accomplished by performing multiple simulations at different values of the parameter conjugate to X . To compute $A(V)$, different pressures P would be imposed for multiple simulations at the same temperature.
- The multiple-histogram reweighting procedure will compute the relevant free energy $F(X)$ as well as the relative partition functions at the different simulation conditions. In the example for $A(V)$, we would find $A(V)$ as well as G_j , the Gibbs free energy of each NPT run.

Case study: T-dependent free energies at constant N, V

Theory

Let's imagine that we want to compute $\beta_2 A(T_2) - \beta_1 A(T_1)$ for a system. Statistical mechanics tells us that:

$$\begin{aligned} Z(T) &= e^{-\beta A} \\ &= \int \Omega(U) e^{-\beta U} dU \\ &= \int e^{S(U) - \beta U} dU \end{aligned}$$

where we have repressed the dependencies of quantities on N and V , which are constant. Here, Ω is the configurational density of states and S the configurational entropy. The reason that we write the free energy as this particular integral is because Ω is common between two different temperatures, i.e., it is strictly temperature independent.

This form also suggests that free energies (and differences) are related to relative energy fluctuations. Consider the energy distribution $\wp(U)$:

$$\wp(U) = \frac{\Omega(U) e^{-\beta U}}{Z(T)} = \frac{e^{S(U) - \beta U}}{Z(T)}$$

where $Z(T, V, N)$ is the configurational partition function.

If we were able to compute $S(U)$, we would be able to predict the complete energy distribution $\wp(U)$ at any temperature. For example, we can compute the average potential energy at any temperature using the expression

$$\langle U \rangle(T) = \int U \wp(U; T) dU$$

$$\begin{aligned}
&= \int \frac{U e^{S(U)-\beta U} dU}{Z(T)} \\
&= \frac{\int U e^{S(U)-\beta U} dU}{\int e^{S(U)-\beta U} dU}
\end{aligned}$$

We might extract an estimate for S by measuring $\wp(U; T)$ from a histogram in a simulation at specified temperature T . Inverting the above relationship,

$$\begin{aligned}
S(U) &= \ln \left[\frac{Z(T) \wp(U; T)}{e^{-\beta U}} \right] \\
&= \ln \wp(U; T) + \beta U + \ln Z(T) \\
&= \ln \wp(U; T) + \beta U - \beta A(T)
\end{aligned}$$

Here, $A(T)$ technically denotes the configurational part of the Helmholtz free energy, which is $k_B T \ln[N! \Lambda(T)^{3N}]$ less than the total free energy in a single component system.

We can measure $\wp(U; T)$ using a histogram along a set of discrete energies. Post-simulation, we can then take this measured distribution and use it to compute a discrete approximation to the dimensionless entropy function at the same energies,

$$S(U_k) = \ln \wp(U_k; T) + \beta U_k - \beta A(T)$$

Because the Helmholtz free energy is a temperature-dependent constant that is *independent* of potential energy U , we can compute the entropy function S to within an *additive constant* through a histogram. Fundamentally, it does not matter which temperature we use to measure $S(U)$: any temperature-dependencies on the RHS of this equation should exactly cancel to leave a T -independent function.

Reweighting a single simulation

We can use a computed a discrete approximation to the energy distribution to predict $\wp(U)$ at temperatures *other* than the original simulation temperature. This is a basic form of **reweighting**, since we use a distribution measured at one temperature to predict that at another.

Imagine $S(U)$ is computed from $\wp(U; T_1)$ in a canonical simulation at T_1 . We have

$$S(U) = \ln \wp(U; T_1) + \beta_1 U - \beta_1 A(T_1)$$

We want to predict $\wp(U; T_2)$ at another temperature T_2 . We have

$$\wp(U; T_2) = \frac{e^{S(U)-\beta_2 U}}{Z(T_2)}$$

Plugging in the above expression for S ,

$$\wp(U; T_2) = \wp(U; T_1) e^{-(\beta_2 - \beta_1)U} \frac{Z(T_1)}{Z(T_2)}$$

The last term on the RHS involving the ratio of partition functions is independent of U . We can find it using the probability normalization condition,

$$\int \wp(U; T_2) dU = 1$$

Thus,

$$\wp(U; T_2) = \frac{\wp(U; T_1) e^{-(\beta_2 - \beta_1)U}}{\int \wp(U; T_1) e^{-(\beta_2 - \beta_1)U} dU}$$

This equation states that a distribution measured at T_1 can be used to predict a distribution at a different T_2 . In principle, we would only need to perform a single simulation, measure $\wp(U)$ once, and use this expression to examine the energy distribution at any other temperature of interest. For example, we might compute the average energy as a function of temperature using the equation

$$\begin{aligned} \langle U \rangle(T_2) &= \int U \wp(U; T_2) dU \\ &= \frac{\int U \wp(U; T_1) e^{-(\beta_2 - \beta_1)U} dU}{\int \wp(U; T_1) e^{-(\beta_2 - \beta_1)U} dU} \end{aligned}$$

Similar expressions could be found for other moments of the potential energy, such as $\langle U^2 \rangle$. These could be used to compute the temperature dependence of the heat capacity, using the relationship $k_B T^2 C_V = \langle U^2 \rangle - \langle U \rangle^2$.

Free energy differences from histograms

If we measure $\wp(U; T_1)$ and $\wp(U; T_2)$ from two simulations at different temperatures, we can also compute a free energy difference. We construct:


$$\begin{aligned} \ln \frac{\wp(U; T_2)}{\wp(U; T_1)} &= [S(U) - \beta_2 U + \beta_2 A(T_2)] - [S(U) - \beta_1 U + \beta_1 A(T_1)] \\ &= -(\beta_2 - \beta_1)U + \beta_2 A(T_2) - \beta_1 A(T_1) \end{aligned}$$

Rearranging,

$$\beta_2 A(T_2) - \beta_1 A(T_1) = \ln \frac{\wp(U; T_2)}{\wp(U; T_1)} + (\beta_2 - \beta_1)U$$


The RHS is dependent on U , whereas the LHS is not! In principle, any value of U could be plugged into the RHS and the same free energy would be returned.

Unfortunately, there are practical limits to the histogram reweighting procedure just described. The main problem is the measurement of $\wp(U)$ to good statistical accuracy in the tails of its distribution. Consider that a typical $\wp(U)$ is very sharply peaked:



The width of the distribution relative to the mean goes roughly as $N^{-\frac{1}{2}}$ so that, for macroscopic systems, the distribution is infinitely peaked.

The implication of this result is that it is very hard to measure the distribution at its tails, where we typically only have a few counts in each bin. If we reweight a measured distribution to a temperature where the tails change to high probability, the error can be magnified to be very large.



Here, the mean of the distribution at the new temperature is well within the tail region of the distribution at the original temperature. If we reweight $\wp(U; T_1)$ to T_2 , errors in the tails will be magnified. The error in the new distributions T can be written approximately as

$$\sigma_{\wp(U;T_2)}^2 \approx \sigma_{\wp(U;T_1)}^2 \left(\frac{\wp(U;T_2)}{\wp(U;T_1)} \right)^2$$

This formula is derived using standard error propagation rules. Notice that if $T_2 = T_1$, the error is the same at each energy as from the original measurement. Otherwise, we must weigh the error by a ratio involving the two probability distributions. In the above picture, the ratio at the mean energy

$$\frac{\wp(U^*;T_2)}{\wp(U^*;T_1)}$$

is very large due to the small probability of this energy at T_1 . Therefore, the error is greatly magnified in the reweighting procedure according to the above equation.

In general,

- Distributions are subject to statistical inaccuracies at their tails, owing to the finite number of counts in each bin.
- If the important parts of the energy distribution at T_2 correspond to the tails of a measured distribution at T_1 , a reweighting procedure will fail due to large statistical errors.
- If computed from a measured $\wp(U;T)$, the quality of an estimated $S_c(U)$ function is only good around the frequently-sampled energies, i.e., for energies where the histogram has many entries.


These techniques limit the determination of free energies and entropies using measurements from single histograms. The solution is to incorporate multiple histogram estimates of these quantities.

Multiple histogram reweighting

As we saw with umbrella sampling, we can improve our estimates of free energies using multiple histograms from multiple simulations i , each at a different temperature T_i . For each, we would get a different estimate of the underlying dimensionless configurational entropy function S :

$$S(U) = \ln \wp(U;T_i) + \beta_i U - \beta_i A(T_i)$$

We should get the same function every time. We don't know the free energies, but they are independent of U and thus can be treated as additive constants. We can therefore shift the different estimates for the entropy function up or down in value until they all match each other:



Notice that by finding the amounts to shift each curve to achieve overlap, we are determining relative free energy differences.

The key to these calculations are these shift amounts. The best ways of evaluating them considers the errors for each energy in each measured $\phi(U)$. For this approach to work, we must have good **overlap** between the histograms. That is, there must be ranges of the histograms that overlap with a reasonable number of counts in each. Otherwise, it will be challenging to determine the shift amounts to good statistical accuracy.

Statistically optimal multiple histogram reweighting

As we saw with umbrella sampling, we can use statistically optimal methods to stitch together the results from multiple simulations to compute relative free energies. In this context, we are determining $S(U)$ from J simulations at different temperatures T_j . We know that the underlying $S(U)$ for each should be the same, but that the measured histograms $c_j(U)$ will be different because the simulations are performed at different temperatures. The observed events are the energies tabulated in the histograms.

Here, the variables are:

- n – number of observations at each temperature

- i – index of observations ($i = 1, \dots, n$)
- j – index of temperature ($j = 1, \dots, J$)
- k – index of discrete energy values
- U_k – discrete energy values
- S_k – entropy at the values U_k
- U_{ij} – energy of the i th observation at temperature j
- c_{jk} – count of histogram entries for temperature j and energy bin k

With this notation, we construct the total likelihood L of making the nJ observations from the simulations in terms of a yet-unknown unknown entropy function:

$$L = \prod_{i=1}^n \prod_{j=1}^J \wp(U_{ij}; T_j)$$

$$= \prod_{i=1}^n \prod_{j=1}^J \left(\frac{e^{S(U_{ij}) - \beta_j U_{ij}}}{Z_j} \right)$$

where the Z_j are given by

$$Z_j = e^{-\beta_j A_j} = \sum_k e^{S_k - \beta_j U_k}$$

Using a similar maximum likelihood approach as before, we find that the optimal entropy S_k and free energy A_j values are found from the iterative solution of the **multiple histogram reweighting equations**:

$$S(U) = \ln c_{tot}(U) - \ln n - \ln \sum_{j=1}^J e^{-\beta_j U + \beta_j A_j} \quad \text{for all } U$$

and

$$-\beta_j A_j = \ln \sum_U e^{S(U) - \beta_j U} \quad \text{for all } j$$

Keep in mind that we must set one free energy to zero, e.g., $A_1 = 0$. This iterative procedure can be fairly computationally expensive and require a nontrivial number of iterations to converge. The pymbar library has very nice, automated numerical methods to solve these equations.

Once convergence has been achieved, energy averages at any temperature—even temperatures not included in the original dataset—can be computed with the expression

$$\wp(U; T) \propto e^{S(U) - \beta U}$$

The constant of proportionality is determined by normalizing the probability distribution.

Property averages

As before, we can compute averages of properties at any temperature by defining a **weight** associated with each configuration i at each temperature j at the reweighting temperature T :

$$w_{ij} = \frac{e^{S(U_{ij}) - \beta U_{ij}}}{c_{tot}(U_{ij})}$$

such that any property average for an observable X follows:

$$\langle X \rangle = \frac{\sum_{j=1}^J \sum_{i=1}^n w_{ij} X_{ij}}{\sum_{j=1}^J \sum_{i=1}^n w_{ij}}$$

Here, X_{ij} is the value of the property X for the i th configuration observed at temperature j .

Binless implementation

Because we are interested in free energies at different discrete states (e.g., temperatures), versus along some continuous coordinate, it's worth considering the role of the bin width δU for computing histograms and tabulating energies. It is possible to follow this derivation in the limit that $\delta U \rightarrow 0$. Then the expression for the free energies becomes:

$$-\beta_j A_j = -\ln n + \ln \left[\sum_{l=1}^J \sum_{i=1}^n e^{-\beta_j U_{il}} \left(\sum_{m=1}^J e^{-\beta_m U_{il} + \beta_m A_m} \right)^{-1} \right]$$

where l and m are also indices over temperature. As with the previous approach, this equation must be iterated to solve for the free energies A_j . Notice that it involves a quadruple loop, once over observations and three times over temperatures (the indices j, l, m). Therefore, convergence can be much slower than the histogram version. The MBAR and pymbar methods essentially solve this problem efficiently.

Once the free energies are determined, the configurational weights are:

$$w_{ij} \propto \frac{e^{-\beta U_{ij}}}{\sum_{l=1}^J e^{\beta_l A_l - \beta_l U_{ij}}}$$

Precision issues

It is very important to keep in mind that the combined exponentiation and logarithm functions in these expressions require careful treatment in numerical evaluation in the computer in order to avoid precision inaccuracies. More on this issue is described in the document “Simulation best practices.” As a quick example, consider the computation of

$$-F = \ln \sum_{i=1}^M e^{w_i}$$

To get around this numerical precision issues, we rearrange this expression so that the terms in the exponential are better-behaved:

$$\begin{aligned} -F &= \ln \left[\left(\frac{e^{w_{\max}}}{e^{w_{\max}}} \right) \sum_{i=1}^M e^{-w_i} \right] \\ &= w_{\max} + \ln \sum_{i=1}^M \frac{e^{w_i}}{e^{w_{\max}}} \\ &= w_{\max} + \ln \sum_{i=1}^M e^{w_i - w_{\max}} \end{aligned}$$

where

$$w_{\max} = \max w_i$$

Case study: predicting phase behavior with reweighting methods

The computation of phase equilibria—saturation conditions, critical points, phase diagrams, etc.—is a major enterprise in the molecular simulation community. There are many ways that one might go about this, including the Gibbs ensemble, and they are discussed in detail in [Chipot & Pohorille, *Free Energy Calculations: Theory and Applications in Chemistry and Biology*, Springer, 2007]. Of all of these methods, however, the current gold-standard for phase equilibrium computations is one that relies on multiple histogram reweighting. This approach is often the quickest and has the best accuracy.

Here we discuss these calculations in the context of liquid-vapor equilibria. Other phase behavior—such as liquid-solid, vapor-solid, or liquid-liquid (in mixtures)—can be treated using this approach as well, although special techniques may be needed to treat structured phases.

Grand canonical perspective

The basic idea of our calculations is the following:

- For a given temperature, we want to find the conditions at which phase equilibrium occurs. In a single component system, we only have one free parameter to specify this condition in addition to T , per the Gibbs phase rule. This could be P or μ .
- We will focus on μ instead of P and perform simulations in the grand canonical ensemble. Thus, we want to find the value of μ at coexistence that generates an equilibrium between liquid and vapor phases. We could equally choose P and perform NPT simulations; however, particle insertions and deletions are generally much more efficient than volume fluctuations and thus we go with the former for computational reasons.
- If the state conditions place us at conditions of phase equilibrium (i.e., on the liquid-vapor phase line), we expect the GCMC simulation to spontaneously fluctuate between a dense liquid (low U , large N) and dilute gas (high U , small N). These fluctuations can be observed by examining the joint distribution

$$\wp(U, N)$$


- At phase equilibrium, $\wp(U, N)$ will be bimodal, and the total probability weight under the two peaks will be equal. One peak corresponds to the liquid phase, and the other to the gas phase.
- We will find the conditions of phase equilibrium (T, μ) that result in a bimodal $\wp(U, N)$. We will generate $\wp(U, N; T, \mu)$ by reweighting simulation data from multiple simulations, each not necessarily on the phase coexistence line.

GCMC implementation

We need to gather enough histogram data so that we will be able to reweight $\wp(U, N; T, \mu)$ with good statistical accuracy. We perform J simulations in the grand-canonical ensemble at different temperatures T_j and chemical potentials μ_j . For each we measure the joint distribution $\wp(U, N; T_j, \mu_j)$. The temperatures and chemical potentials should be chosen so that the energy and particle number distributions all overlap with each other and span the complete range of liquid-to-gas densities. Typically we perform the following range of simulations:

- one near the (presumed) critical point that has large fluctuations in U and N
- one in the gas phase
- several in the liquid phase, of varying temperature

The point is that we must have good histogram statistics for all the particle numbers and energies that will ultimately be relevant at phase equilibrium.



Once these histograms are taken, **multiple histogram reweighting** (Ferrenberg-Swendsen like) iterations can be used to stitch together all of the data and to compute the free energy at each state point. The relevant equations for GCMC simulations are:

$$S(U, N) = \ln c_{tot}(U, N) - \ln n - \ln \sum_{j=1}^J e^{-\beta_j U + \beta_j \mu_j N + \beta_j F_j} \quad \text{for all } U, N$$

$$-\beta F_j = \ln \sum_U \sum_N e^{S(U, N) - \beta_j U + \beta_j \mu_j N} \quad \text{for all } j$$

This procedure is performed until values for the discretized entropy $S(U, N)$ and free energies $F_j = -k_B T_j \ln \Xi_j$ converge. Keep in mind that one value of F must be set equal to zero.

Once these values are determined, the joint distribution $\wp(U, N)$ can be computed for an arbitrary reweighting T, μ using


$$\wp(U, N; T, \mu) \propto e^{S(U, N) - \beta U + \beta \mu N}$$

The constant of proportionality is given by the normalization condition.


Using the reweighting equation, one finds phase equilibrium by adjusting μ given a value of T until a bimodal distribution appears and the integral under each peak (the total probability of the

phase) is the same. Notice that this is a very fast operation since the Ferrenberg-Swendsen reweighting only needs to be performed once, at the beginning, to determine $S(U, N)$.

The following figure shows what this distribution might look like, taken from a review article by Panagiotopoulos [Panagiotopoulos, J Phys: Condens Matter 12, R25 (2000)].



By repeating this procedure at different reweighting temperatures, one can map out the phase diagram. In the T - ρ plane, this might look like:



The average density of each phase can be determined using:

$$\langle \rho \rangle_L = \frac{1}{V} \frac{\sum_U \sum_{N > N^*} N \wp(U, N)}{\sum_U \sum_{N > N^*} \wp(U, N)} \quad \langle \rho \rangle_G = \frac{1}{V} \frac{\sum_U \sum_{N < N^*} N \wp(U, N)}{\sum_U \sum_{N < N^*} \wp(U, N)}$$

Here, N^* is the value of N at the minimum between the two peaks in the distribution.

Pressures can be found using the fact that

$$\beta PV = \ln \Xi(\mu, T, V)$$

We cannot compute absolute pressures using this equation, because we cannot compute Ξ absolutely (one F_j must be set to zero). However, we can compare the pressures of two different state points:

$$\begin{aligned} \beta_2 P_2 V - \beta_1 P_1 V &= \ln \frac{\Xi(\mu_2, T_2, V)}{\Xi(\mu_1, T_1, V)} \\ &= \ln \frac{\sum_U \sum_N e^{S(U, N) - \beta_2 U + \beta_2 \mu_2 N}}{\sum_U \sum_N e^{S(U, N) - \beta_1 U + \beta_1 \mu_1 N}} \end{aligned}$$

By letting one state point correspond to a very dilute, high temperature gas phase, we can compute its absolute pressure using the ideal gas law. This equation then can be used to relate the pressure at other state points back to the ideal gas reference.

Critical point reweighting and finite-size scaling

How can one predict the location of the critical point? It turns out that systems have a kind of universal behavior at this point in the phase diagram. For many fluids, their behavior at the critical point follows the same trends as the three-dimensional Ising model (or, equivalently, the lattice gas); that is, they fall in the **Ising universality class** or the **Ising criticality class**. In particular, the probability distribution $\wp(U, N)$ assumes a universal form.

We can locate the true critical point by finding reweighting values of T_c, μ_c so that $\wp(U, N)$ maps onto the universal Ising distribution, which has been computed to high accuracy by a number of researchers.

Near the critical point, fluctuations become macroscopic in size. That means that computed properties become sensitive to the size of the system being studied. In particular, the computed T_c, μ_c can have a substantial dependence on the system size V . Fortunately, these parameters have a scaling law behavior with V . To locate the infinite system size T_c, μ_c one performs multiple studies for different values of V , each time computing the critical point using the procedure above. Then, one extrapolates to $V \rightarrow \infty$. An example of this taken from [Panagiotopoulos, 2000] is:

