1. **A simple model of a polymer.** You're no doubt aware that the shape a protein folds up into is influenced by energetic interactions between amino acids. You may not have given as much thought to the influence of entropy on a protein's shape. To focus on that concept, we consider the simplest model for conformational fluctuations of a long chain molecule, formed from n polymer segments connected end to end. Each segment $i = 1, 2, \ldots, n$ (perhaps comprising many chemical units) has a fixed length ℓ and an orientation $\hat{b}_i$ that is parallel to one of d Cartesian axes ($\hat{x}, \hat{y}$, or $\hat{z}$ in 3 dimensions). In other words, the molecular configuration traces a random walk on a $d-$dimensional cubic lattice:

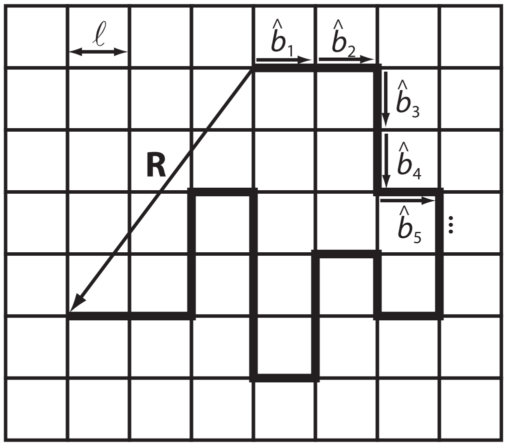

Imagine that the orientations of different segments are statistically independent, and that there is no preferred orientation, $\langle \hat{b}_i \rangle = 0$ and $\langle \hat{b}_i \cdot \hat{b}_j \rangle = \delta_{ij}$, where $\delta_{ij} = 1$ if $i = j$ and vanishes otherwise.

(i) Show that the entropy of such an ideal chain molecule has the form $S = k_{\mathrm{B}} n \ln a$. Determine the parameter a as a function of dimensionality d.

The molecule consists of n bonds, each of which can be selected from $2d$ options. Hence the total number of possible polymers (assuming we always start at the origin) is $\Omega = (2d)^n$. The entropy (excluding any translational entropy from the location of the first monomer that I'm fixing at the origin) is therefore $S = k_{\mathrm{B}} \ln \Omega = k_{\mathrm{B}} n \ln 2d$.

(ii) What is $\langle \vec{R} \rangle$, where $\vec{R} = \ell \sum_{i=1}^{n} \hat{b}_i$ is the end-to-end vector corresponding to a single polymer configuration?

By symmetry, $\langle \vec{R} \rangle = 0$. Every step is as likely to go left as to go right, up as down, *etc*.

(iii) Given your response to part (ii), suggest a reason that the root mean squared end-to-end distance $\sqrt{\langle R^2 \rangle}$, where $R^2 = \vec{R} \cdot \vec{R}$, would be a better measure of the size of the polymer than the average displacement vector $\langle \vec{R} \rangle$.

The root mean squared end-to-end distance circumvents the cancellations described in part (ii), in the sense that it gives a typical breadth of the distribution even when the distribution of displacements has zero mean.

(iv) Calculate the mean squared end-to-end distance of the chain molecule, $\langle R^2 \rangle$. Your result should indicate that the typical distance between ends of the molecule grows with chain length as $\sqrt{\langle R^2 \rangle} \propto \sqrt{n}$. How does your result depend on dimensionality d?

Notice that $\vec{R} = \ell \sum_{i=1}^{n} \hat{b}_i$, so

$$\langle \vec{R} \cdot \vec{R} \rangle = \ell^2 \left\langle \sum_{i=1}^{n} \sum_{j=1}^{n} \hat{b}_i \cdot \hat{b}_j \right\rangle = \ell^2 \sum_{i=1}^{n} \sum_{j=1}^{n} \langle \hat{b}_i \cdot \hat{b}_j \rangle = \ell^2 \sum_{i=1}^{n} \sum_{j=1}^{n} \delta_{ij} = \ell^2 \sum_{i=1}^{n} 1 = n\ell^2.$$

Thus, $\sqrt{\langle R^2 \rangle} = \ell n^{1/2}$. Remarkably the answer does not depend on the dimensionality d!

2. **A Gaussian polymer.** Real polymers don't live on lattices, so you would be excused for thinking the model in the first problem would have some serious limitations. Perhaps a more realistic model for a classical polymer should consist of $n+1$ monomers connected together in a chain by n harmonic springs. We use $\vec{R}_i$ to denote the position (in three-dimensional space) of the i^{th} monomer. The energy of the polymer has kinetic and potential energy contributions which depend on the positions and velocities of every monomer:

$$E(\vec{R}_0, \ldots, \vec{R}_n, \dot{\vec{R}}_0, \ldots, \dot{\vec{R}}_n) = \frac{k}{2} \sum_{i=1}^{n} (\vec{R}_i - \vec{R}_{i-1})^2 + \frac{m}{2} \sum_{i=0}^{n} \dot{\vec{R}}_i^2,$$

where m is the mass of a monomer, $\dot{\vec{R}}_i$ is the velocity of monomer i, and the strength of the spring connecting polymers is k with units of energy per length squared. The polymer is immersed in a solution which has temperature T, so the energy of the polymer can fluctuate.

(i) The potential energy in a single bond is $\frac{k}{2}(\vec{R}_i - \vec{R}_{i-1})^2$. Compute the average potential energy stored in the single bond as a function of k_{B}, T, and n. Your final answer should not depend on k or on m. (The result is closely connected with the so-called equipartition theorem.) Should the velocities matter? Why or why not? Does your answer depend on the fact that the polymer is in three-dimensional space?

The average energy stored in a single bond is given by

$$\langle E_{\text{single bond}} \rangle = \frac{k}{2} \langle R^2_{\text{single bond}} \rangle .$$

Note that $\vec{R}_{\text{single bond}}$ is a Gaussian random variable (by the Boltzmann distribution, and because the energy is a quadratic function of the displacements), so the expectation value of $R^2_{\text{single bond}}$ is simply the variance of a Gaussian distribution. Recall that in 3D,

$$\langle R^2 \rangle = \left(\frac{1}{2\pi\sigma^2} \right)^{3/2} \int \mathrm{d}\vec{R}\, R^2 e^{-R^2/(2\sigma^2)} = 3\sigma^2.$$

Comparing the integrand with the Boltzmann distribution tells us that $\sigma^2 = 1/\beta k$. Hence $\langle R^2_{\text{single bond}} \rangle = 3k_{\mathrm{B}}T/k$. Because the energy in the bond is $k/2$ times this average distance squared, we find

$$\left\langle E_{\text{single bond}}\right\rangle = \frac{3k_{\mathrm{B}}T}{2}.$$

The velocities do not matter because the energy is a function of the positions only; velocities will not appear in the integral. The dimension **does** matter: for a three-dimensional system, we have $\langle R^2\rangle = \langle X^2+Y^2+Z^2\rangle = 3\langle X^2\rangle = 3\sigma^2$ because the statistics along each dimension are the same, and this result naturally generalizes to $\langle R^2\rangle = d\sigma^2$ in d dimensions. Alternatively, you may show this by explicitly evaluating the integral (use Cartesian coordinates and argue that some terms vanish based on the symmetry of the integrand).

(ii) Let us define $\vec{R}_{0\to n}$ to be the vector from monomer 0 to monomer n. Owing to the harmonic form of the potential, the probability distribution $\rho(\vec{R}_{0\to n})$ will be a three-dimensional Gaussian:

$$\rho(\vec{R}_{0\to n}) = \left(\frac{1}{2\pi\langle\delta R^2_{0\to n}\rangle}\right)^{3/2} \exp\left(-\frac{\left(\vec{R}_{0\to n} - \langle\vec{R}_{0\to n}\rangle\right)^2}{2\langle\delta R^2_{0\to n}\rangle}\right),$$

in terms of the mean $\langle\vec{R}_{0\to n}\rangle$ and the variance $\langle\delta\vec{R}^2_{0\to n}\rangle$. As in Problem 1's polymer model, the second moment grows linearly with the number of monomers: $\langle(\vec{R}_{0\to n})^2\rangle = n\langle(\vec{R}_{0\to 1})^2\rangle$. (In Problem 1, the length of a single bond was **always** ℓ, but here it fluctuates.) Determine the normalized Gaussian distribution $\rho(\vec{R}_{0\to n})$ in terms of β, k, and n.

Depending on how big of a challenge you want, you could really *prove* this Gaussian form. To do so, you'd have to average over the position of all monomers other than the two endpoints. Those integrals could be performed one-by-one in a straightforward and tedious manner. Alternatively you could do the integrals inductively by first considering a two-bead polymer and thinking about what happens when you add one more bead. Another cute method I like involves doing a "decimation and renormalization" calculation. If you're feeling more ambitious/bored this week, you can learn more about that strategy at http://gingrich.chem.northwestern.edu/teaching/348/2021/hws/ExtraChallenge.pdf.

The normalized Gaussian is

$$\rho(\vec{R}_{0\to n}) = \left(\frac{1}{2\pi n\langle R^2_{0\to 1}\rangle}\right)^{3/2} \exp\left(-\frac{R^2_{0\to n}}{2n\langle R^2_{0\to 1}\rangle}\right).$$

In the previous problem we saw $\langle R^2_{\text{single bond}}\rangle = 3k_{\mathrm{B}}T/k$. Hence

$$\rho(\vec{R}_{0\to n}) = \left(\frac{\beta k}{6\pi n}\right)^{3/2} \exp\left(-\frac{\beta k R^2_{0\to n}}{6n}\right).$$

The velocities do not matter because they fluctuate independently from the positions.

(iii) How much less likely is it to find the stretched out polymer than the compact one? To answer this, let's say the compact polymer has an end-to-end distance of zero and the stretched polymer is extended a distance L along the x-axis (meaning $\vec{R}_{0\to n} = L\hat{x}$). Express your answer in terms of L, β, and n.

The prefactor cancels when we compare $\rho(\vec{R}_{0\to n} = L\hat{x})$ to $\rho(\vec{R}_{0\to n} = 0)$:

$$\frac{\rho(L)}{\rho(0)} = \exp\left(-\frac{\beta k L^2}{6n}\right).$$

(iv) The polymer prefers to be folded up because there are fewer elongated configurations. Suppose the polymer starts out with both monomers 0 and n at the origin. If both monomers were attached to glass beads, laser tweezers could be used to hold one bead fixed at the origin while the other bead is moved a distance L away. What is the reversible work required to stretch the polymer?

By definition,

$$W_{\text{rev}} = -k_{\text{B}}T \ln \frac{\rho(L)}{\rho(0)} = -k_{\text{B}}T\left(-\frac{\beta k L^2}{6n}\right) = \frac{kL^2}{6n}.$$

3. **Reversible work simulations.**

I have written a simulation of the polymer in Problem 1 being pulled apart at a finite rate (so not necessarily reversibly). You can run the simulation on this website:

http://gingrich.chem.northwestern.edu/teaching/polymer/gaussianchain.html

Each time you run the simulation you get a trajectory of the polymer chain of harmonically bound nearest neighbors beads, with velocities of each bead assigned at the start of each trajectory from a Maxwell-Boltzmann distribution. Each bead of the polymer experiences forces from the neighboring beads, trying to pull the neighbors closer together, but it also feels random "kicks" from the environment. These kicks are meant to mimic the forces from molecules in a fixed-temperature solvent colliding with the polymer. In the simulation, the solvent is not explicitly simulated, rather at each step of the polymer's dynamics, a random Gaussian number is drawn by the computer to represent the extra force a monomer feels from the effect of the solvent. Finally, the endpoint of the polymer experiences an external force that we apply to force the polymer to spread out at a fixed stretching rate. This applied force is measured and plotted in the top right of the screen. When the applied force is integrated, we compute the total work exerted by the external force, and this accumulated work is plotted in the bottom right of the screen.

If you slow down the pulling rate, you will see that the total work to stretch out the chain matches the reversible work.

(i) First consider the case of an infinitesimally slow pulling rate. Compute the average force (averaged over fluctuations, **not** over the length of the polymer) that must be applied on the end monomer as a function of k, L, and n. This should follow from your answer to Problem 2, Part (iv).

Note that the average force is a function of the end-to-end length l. Let us write this mean force as $f(l)$. Integrating the mean force over the distance will give the work:

$$W(L) = \int_0^L \mathrm{d}l\, f(l).$$

Taking a derivative of W with respect to L thus yields

$$f(L) = \frac{\partial W}{\partial L} = \frac{kL}{3n}.$$

(ii) Use the provided simulation, starting with $n = 20$ beads and a pulling rate of 0.1 (in reduced units where $\beta = 1, l = 1$, and bead mass $m = 1$). Then try a pulling rate of 0.01 and a pulling rate of 0.001. Observe that for very slow pulling the total work is very close to the reversible work. Notice also that the mean force fluctuates around the correct average value that you found in (i). But the fluctuations in the measured force are quite large, even at slow pulling rates. Explain why the force fluctuations do not decrease for slow pulling rates even though the work fluctuations do.

The force fluctuations are quite large even for slow pulling, while the work fluctuations are small when the pulling rate is slow. Mathematically, the work can be represented as

$$W = \int_0^L F \,\mathrm{d}x = \int_0^{L/\dot{x}} F\dot{x} \,\mathrm{d}t = \dot{x} \int_0^{L/\dot{x}} F \,\mathrm{d}t.$$

Note that $\dot{x}$ is simply the pulling rate, which is fixed and can hence be brought outside the integral. The smaller $\dot{x}$ gets, the longer the simulation time $\tau = L/\dot{x}$, and hence the more averaged and deterministic the integrated force.

The origin of the force fluctuations is that the particular structure of the polymer can radically change the forces felt at the endpoint. If, for example, the second-to-last bead ends up to the right of the final bead, it will pull with a negative force. Even when the pulling rate is very slow, these less-typical polymer configurations can be explored. The integral of these forces over distance, however, becomes very deterministic for slow pulling rates. In this case, the work being performed at a separation of L is found by averaging over the fluctuating forces times a distance element. In the case of very slow pulling, the entire spectrum of possible force measurements are made and averaged over before the endpoints are appreciably spread further. The subsequent averaging over many polymer configurations makes the work fluctuations very small.

4. **A perpetual motion machine?!** Your friend at the University of Chicago has come up with a brilliant idea. When he played with the simulation, he set $n = 15$ and pulled with a rate of 0.1 (in the reduced units). He observed that sometimes he gets negative work by stretching the polymer out. In other words, he expected to have to do work to pull the polymer apart, but sometimes the polymer did work on him. Your friend is super excited because he thinks he can build some sort of DNA/laser tweezer machine that stretches and unstretches a strand of DNA and gets work out in the process. Use the simulation posted at:

   http://gingrich.chem.northwestern.edu/teaching/polymer/cyclicworkdistribution.html

   to help you evaluate his plan.

   (i) During a single cycle, will the UChicago scientist ever extract work from the system (the measured work is negative)?

   Yes, sometimes fluctuations will yield negative values of work!

(ii) Will the average extracted work ever be negative? To answer, compute the average work $\langle W \rangle$ for pulling rates of $0.05, 0.1$, and 0.2. You will want to check "Repeat Pulling to Collect Work Statistics" then click restart to automatically generate statistics of the pulling experiments. Make sure you run the program long enough for the values of $\langle W \rangle$ to converge. What do you notice about how the distribution of measured work values depends on the pulling rate? Will there be a pulling rate so that your friend's machine will reliably extract work and make him billions of dollars?

No, it looks like the average work will never be less than zero. It seems my friend is out of luck because he's never going to find a pulling rate that will yield a negative work on average. The average work decreases as the pulling rate slows down. Additionally, the breadth of the distribution decreases with slower pulling rates.

(iii) Based upon the time-reversibility of dynamics (i.e., Newton's laws look the same forward as they do backwards), it can be shown that the probability of observing a value of work, W, over one of the cycles is related to the probability of measuring $-W$ by

$$\frac{\rho(W)}{\rho(-W)} = e^{\beta W}, \tag{1}$$

where $\beta = 1/k_{\mathrm{B}}T$ is the inverse temperature. (This formula is a special case of the Crooks fluctuation theorem, which we will discuss further in the next problem.) At a single pulling rate, run the simulation long enough to generate a smooth work probability distribution. Once the distribution has converged use the raw data provided below the plot to confirm that the probability of positive and negative values of work are related as predicted by Equation (1). (Agreement will be good but not perfect only because of finite statistics and because histogram bin sizes analyzed in the applet are fairly wide.) $\beta = 1$ in the reduced units of the simulations.

Here is a plot of the work distribution, collected by pulling a polymer with 15 beads at a rate of 0.1 (in the reduced units of the simulation).

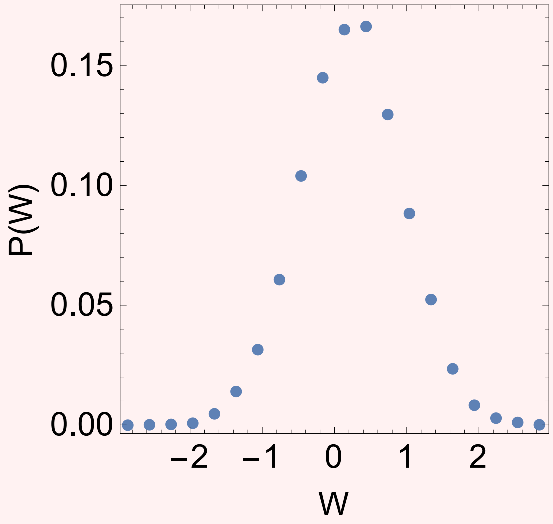

I have also plotted $\ln \rho(W) - \ln \rho(-W)$ for all positive W and compared that with the line βW (remember $\beta = 1$), to see that Equation (1) appears to be satisfied (at least for the values of W that are well-sampled. There are not many samples with large negative values of W, so those points should have big error bars.

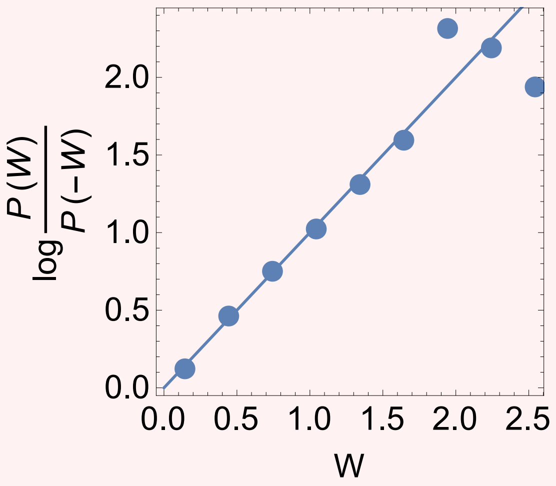

(iv) The average work from Part (ii) can be considered to be an average over the work distribution function, $\rho(W)$, as

$$\langle W\rangle = \int_{-\infty}^{\infty} \mathrm{d}W \, W \rho(W).$$

Show that Equation (1) implies that your answer to Part (ii) must be positive.

$$\langle W\rangle = \int_{-\infty}^{0} \mathrm{d}W \, W\rho(W) + \int_{0}^{\infty} \mathrm{d}W \, W\rho(W) = \int_{0}^{\infty} \mathrm{d}W \, W\rho(W)(1-e^{-\beta W}) \geq 0.$$

The first equality simply splits up the integral over W into the two intervals $(-\infty, 0)$ and $(0, \infty)$. The second equality follows by the substitution $W \to -W$ in the first integral and then combining it with the second. The final inequality follows from noticing that W is positive over $(0, \infty)$, $\rho(W)$ is a probability density and hence positive by construction, and the term in parentheses also must be positive for $W > 0$.

(v) We could find the average value of any observable in the same way we did in Part (iii). For example, the average value of some function of work, $f(W)$ is given by

$$\langle f(W)\rangle = \int_{-\infty}^{\infty} \mathrm{d}W \, f(W)\rho(W).$$

With this formula, show that

$$1 = \left\langle e^{-\beta W} \right\rangle.$$

The average in this equation is over all possible values of work, which means averaging over all of the possible non-equilibrium pulling experiments.

Note that

$$\langle e^{-\beta W}\rangle = \int_{-\infty}^{\infty} dW \rho(W) e^{-\beta W} = \int_{-\infty}^{\infty} dW \, \rho(-W) \xrightarrow{W \to -W} \int_{-\infty}^{\infty} dW \, \rho(W) = 1.$$

(vi) This average of the exponential is computed in the simulation and reported on the screen. Note that it converges to one, but the convergence is slow. Comment on why convergence is so slow. [Hint: Think about how probable it is to observe measurements that contribute the most to the average.]

The average $\left\langle e^{-\beta W}\right\rangle$ is a weighted average. You collect a bunch of samples of work, let's call them $W_1, W_2, \ldots W_n$. You then average over these measurements, giving each measurement a weight of $e^{-\beta W_i}$. That means a very negative value of work carries a huge weight (in fact an exponentially big weight) while the positive values of work carry exponentially small weights. But the negative values of work are also much more rare than the positive ones. So the running average will be strongly influenced by the rare occurances when a very negative value of work is measured. The running average cannot really be converged until the negative work tail of the distribution becomes well estimated.

5. (OPTIONAL) **Equilibrium free energies from nonequilibrium measurements.** In the previous problem, you considered a cyclic process in which the polymer endpoints returned to their original positions at the end of each run. If the reversible work between the starting and ending configurations is not zero, the general form of the Crooks fluctuation theorem applies:

$$\rho_{\mathrm{F}}(W_{\mathrm{d}}) = \rho_{\mathrm{R}}(-W_{\mathrm{d}}) e^{\beta W_{\mathrm{d}}} \,. \tag{2}$$

Here, $W_{\mathrm{d}} = W - W_{\mathrm{rev}}$ is the extra work we do on top of the reversible work. $\rho_{\mathrm{F}}(W_{\mathrm{d}})$ refers to the probability density of measuring a value W_{d} when pulling the polymer. $\rho_{\mathrm{R}}(-W_{\mathrm{d}})$ is the probability density of measuring a value of $-W_{\mathrm{d}}$ when *pushing* the endpoints of the polymer back to where they started. (F stands for forward process, R for reversed.) In experiments with DNA folding it is important to know the reversible work for unfolding a piece of DNA, but we cannot pull the DNA slowly enough to unfold the polymer reversibly. Averaging the work performed during a laser tweezing experiment is a bad estimate for W_{rev} because $\langle W\rangle \neq W_{\mathrm{rev}}$. Notice that the Crooks fluctuation theorem establishes that the difference in statistics between forward and backward processes—the nature of time's arrow—is determined entirely by value of the dissipated work, W_{d}.

Repeat the logic of Problem 4(v) to write down an expression for W_{rev} in terms of the average of some quantity you could measure in the pulling experiments.

You can check that this gives a good estimate for W_{rev} with using the simulation posted at:

http://gingrich.chem.northwestern.edu/teaching/polymer/workdistribution.html.

With the Crooks fluctuation theorem, therefore, you have a way to extract measurements that tell us about equilibrium systems (W_{rev}) from experiments that are not even close to equilibrium!

By inspection,

$$\langle e^{-\beta W_d}\rangle_F = \int_{-\infty}^{\infty} \mathrm{d}W_d \, e^{-\beta W_d} \rho_F(W_d) = \int_{-\infty}^{\infty} \mathrm{d}W_d \, \rho_R(-W_d) = \int_{-\infty}^{\infty} \mathrm{d}W_d \, \rho_R(W_d) = 1.$$

Notice that W_{rev} is a deterministic quantity and can be brought out of the expectation. Hence

$$1 = \langle e^{-\beta W_d}\rangle_F = e^{\beta W_{\mathrm{rev}}} \langle e^{-\beta W}\rangle_F \quad \Longleftrightarrow \quad e^{-\beta W_{\mathrm{rev}}} = \langle e^{-\beta W}\rangle_F .$$