

1. The climb up.

“[The Boltzmann distribution] is the summit of statistical mechanics, and the entire subject is either the slide-down from this summit, as the principle is applied to various cases, or the climb-up to where the fundamental law is derived.” —Richard Feynman

In class we presented a reasonable but less than rigorous derivation of the Boltzmann distribution,

$$P(\nu) \propto e^{-\beta E(\nu)},$$

for microstates ν of a system that can exchange energy with a very large bath. This result was obtained from the fundamental relationship

$$P(\nu) \propto \Omega_B(E_T - E(\nu))$$

through a Taylor expansion of $\ln \Omega_B$, where E_T denotes the total energy shared by system and bath, and $\Omega_B(E_B)$ is the number of bath microstates with energy E_B . Here you will develop this argument a little more carefully (and perhaps a little more convincingly—you may have wondered, for example, why we were Taylor expanding $\ln \Omega_B$ and not Ω_B).

We will assume that Ω_B has a large deviation form,

$$\Omega_B = [\omega_B(\epsilon_B)]^{N_B},$$

where N_B is the number of molecules in the bath, $\epsilon_B \equiv E_B/N_B$ is the corresponding energy per molecule, and $\omega_B(\epsilon_B)$ is a smooth function that does not depend on the size of the bath. As in lecture the dependence of Ω_B on E_B will be used to define a property β of the bath:

$$\beta \equiv \left(\frac{\partial \ln \Omega_B}{\partial E_B} \right)_{N_B, V_B}$$

(i) Show that β is insensitive to the extent of the bath. In particular, relate β to ω_B and derivatives of ω_B with respect to ϵ_B . Explain why this relationship indicates independence of the bath's size. (As we will see in future lectures, this β is an inverse temperature, so it should be pleasing β doesn't depend on bath size. When I say to put a test tube in a 25° C water bath, I shouldn't have to specify if the bath is 1 liter or 2 liters in volume. Assuming the volume of the bath is very big compared to the system, I should be able to specify β alone.)

(ii) Show that

$$\left(\frac{\partial \Omega_B}{\partial E_B} \right)_{N_B, V_B} = \beta \Omega_B$$

(iii) Show that

$$\left(\frac{\partial^2 \Omega_B}{\partial E_B^2} \right)_{N_B, V_B} = \beta^2 \Omega_B + c.$$

Identify the quantity c and explain why it can be neglected in the limit $N_B \rightarrow \infty$.

(iv) Calculate

$$\left(\frac{\partial^n \Omega_B}{\partial E_B^n} \right)_{N_B, V_B}$$

for arbitrary (integer) n . Do not include any terms that are negligible in the limit $N_B \rightarrow \infty$.

(v) Consider the Taylor expansion

$$\begin{aligned}\Omega_B(E_T - E) &= \Omega_B(E_T) - E \left(\frac{\partial \Omega_B}{\partial E_B} \right)_{N_B, V_B} + \frac{1}{2} E^2 \left(\frac{\partial^2 \Omega_B}{\partial E_B^2} \right)_{N_B, V_B} + \dots \\ &= \sum_{n=0}^{\infty} \frac{1}{n!} (-E)^n \left(\frac{\partial^n \Omega_B}{\partial E_B^n} \right)_{N_B, V_B},\end{aligned}$$

where all partial derivatives are implicitly evaluated at $E = 0$.


Combining your results together with this expansion, show that the relationship

$$\Omega_B(E_T - E(\nu)) \propto e^{-\beta E(\nu)}$$

is exact in the limit $N_B \rightarrow \infty$ of an infinitely large bath.

2. 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 $i = 1, 2, \dots, n$ (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 number of microstates of such an ideal chain molecule has the form $\Omega = a^n$. Determine the parameter a as a function of dimensionality d . (Assume that every chain starts at the origin.)

(ii) What is the expected value of \mathbf{R} , that is to say what is $\langle \mathbf{R} \rangle$, where $\mathbf{R} = \ell \sum_{i=1}^n \hat{b}_i$ is the end-to-end vector?

(iii) The squared length of the end-to-end vector is $R^2 = \mathbf{R} \cdot \mathbf{R}$. 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 $\sqrt{\langle R^2 \rangle} \propto n^\nu$. Identify the exponent ν . How does your result depend on the dimensionality d ?

(iv) For a (three-dimensional) polymer in “good” solvent, experiments yield $\nu \approx 3/5$. (A “good” solvent is one that prevents the molecule from collapsing onto itself, i.e., effective interactions among different segments are repulsive.) Compare this measured value with the one you calculated, and comment on the discrepancy.

3. Exploring the lattice polymer with a computer

One way to get more insight into the previous problem is to *sample* realizations of the polymer using a computer. To illustrate the idea, let’s temporarily think about last week’s coin flip problem. With 3 coin flips, one can enumerate (that is list out) all of the possible sequences of coin flips: HHH, HHT, HTH, THH, HTT, THT, TTH, TTT. You could even write a computer program to systematically generate the possibilities, but for N coin flips there will be 2^N possibilities. As N gets bigger and bigger, it becomes less and less reasonable to use a computer to generate all of the possibilities. Rather, it is commonly useful to generate *representative samples* from the probability distribution. You might, for example, use a computer to randomly flip 100 coins in a row then use that particular 100-flip sequence as a good guess for what you would expect. Since all of the coin flip microstates are equally likely, it is kind of crazy to think that any one of your randomly generated sequences is actually a good guide for what to expect. Nevertheless, you could collect several different 100-flip sequences and take that collection of samples as a good reflection of the distribution of possibilities. In particular, you could approximate the distribution for a macrostate, like the total number of times you got heads, by constructing a histogram that counts up how many times your samples returned 0 heads, 1 head, 2 heads, etc.

For the polymer problem, we generate representative samples of these polymers by randomly picking a direction in which to extend the polymer during each step, with all of the directions on the grid being equally probable. Online, you will find a link to some code that generates such random polymers in both two- and three-dimensional space.

(i) Using the provided code (or by writing your own), generate some samples of the random lattice polymer in both two and three dimensions with $n = 1000$. Is it more common to see the polymer with bunched up regions that appear as though it is partially “folded” or more likely to see the polymer completely spread out as if it is “unfolded”? Notice that there is nothing like a Coulomb attraction pulling any of the monomers closer together. What you are seeing is entirely entropic (due to counting the number of possibilities).

(ii) Generate 100 different $n = 1000$ 3d polymers. How many of the 100 samples have the polymer start and stop at the origin? Does this observation contradict your response to 2(ii)?

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

(iv) Using sufficiently many samples of the $n = 1000$ polymer in three dimensions, plot a histogram to approximate the probability distribution for R^2 . Does your histogram agree with your answer to 2(iii)? Discuss.

4. **Fun with Gaussians.** On the last problem we had several occasions to confront the Gaussian distribution,

$$\rho(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right).$$

Some students expressed that it would be useful to build a little more familiarity with the distribution. Consider these exercises to be either a crash course or a refresher. I do not consider it important that you “discover” the standard tricks for computing the integrals, but I do think you should be aware of them. Please talk to me or to other students if you’re having trouble figuring out how to proceed!

(i) **Normalization:** Show that $\int dx \rho(x) = 1$. [*Standard trick: Convert to a two-dimensional integral over a joint distribution of identical, independent Gaussians for the x and y coordinates. Convert to polar coordinates, taking care to include the appropriate Jacobian factor.*]

(ii) **Mean:** Show that $\langle x \rangle = \int_{-\infty}^{\infty} dx x P(x) = \mu$. [*Standard trick: Substitute $u = x - \mu$ and notice that an integral cancels by symmetry.*]

(iii) **Variance:** Show that $\langle (\delta x)^2 \rangle = \int_{-\infty}^{\infty} dx (x - \mu)^2 P(x) = \sigma^2$. [*Standard trick: From (i) and symmetry find the integral of $e^{-\alpha x^2}$ from 0 to ∞ . Differentiate with respect to α .*]