Chem 163, Lecture 15

Adam E. Cohen

11/3/2022

1 Intro to polymers

Today we'll talk about polymer physics. This is one of my favorite topics in all of science. I know I say that about a lot of the things we talk about in this class, but this time I really mean it. As with most other things we discuss, this is a vast discipline, with many books, courses, and conferences devoted to the subject. Today we'll only scratch the surface, considering ideal, non-interacting chains. But before we get too tangled up in the details, I want to talk in a general way about the things we're leaving out, just so you know these things exist when you run into them in the future.

A polymer is composed of a chain of repeating elements. In a homopolymer, such as polyethylene, each of these elements is identical. Most biological polymers are heteropolymers: DNA and RNA are composed of sequences of 4 nucleotides, and proteins are composed of the 20 amino acids. What are the biophysical (as opposed to genetic) implications of the fact that DNA, RNA, and proteins are heteropolymers?

In proteins, we know that the variability in the structure is essential to the biological function. Electrostatic, van der Waals, hydrogen bonding, and hydrophobic interactions between the amino acids lead to protein folding, which ultimately determines the function of the protein. The details of calculating the folded structure from any given sequence of amino acids is a famously difficult computational problem. People have contests, use supercomputers, and do all sorts of ridiculous things to try to solve this problem. Recently, Google had good success with a program called alphaFold, but this is not a physics-based program. Rather, it is essentially a pattern-matching program trained on the large (and growing) body of crystallographic structural data. We will soon see how complicated intermolecular forces can be, and we basically don't understand the pairwise or higher order interactions.

There are two other reasons why protein folding is a difficult problem. One is that it occurs across a huge variety of timescales: the individual molecular vibrations occur on a picosecond timescale, while folding often takes between microseconds (for short, single-step folders) to minutes or hours (for molecules with many disulfide bridges). Molecular dynamicists simulate every vibration of every atom, and if you're going to do that it takes a lot of computer power to simulate 10^6 atoms for a second. So an interesting question is: do you need to simulate every vibration to describe the overall motion? We saw in the case of diffusion that the details of the microscopic steps don't matter–the central-limit theorem allows us to ignore the short-time dynamics in our description of the long-time dynamics. We used this kind of thinking in calculating the Kramers hopping rate across a potential barrier. Perhaps there are good ways to coarse-grain the dynamics over time, to simulate a protein folding trajectory. Nobody knows.

At the same time, it's not obvious that if what you want is the final equilibrium conformation, you have to simulate the trajectory that the protein actually takes to get there. After all, on a computer you can make a protein do all sorts of non-physical things (like pass through itself). And *in vivo* many proteins don't spontaneously fold into their native state; chaperones or chaperonins are required to assist this folding process, through an unknown mechanism. Perhaps there are better techniques for getting to the ground state than following the actual trajectory.

The other problem with simulating protein folding is a spatial one: there are too many spatial degrees of freedom. If we tried to sample them all computationally, it would take more than the age of the universe. In fact, even if the protein itself tried to sample all of its conformational states, going at a rate of picoseconds per conformation, an exhaustive search would take longer than the age of the universe (this is called Levinthal's paradox). People are currently working on techniques for spatial coarse-graining of the protein folding problem: using implicit water, clustering groups of atoms, or amino acids into functional bricks, and so on. The problem is choosing how to do this in an intelligent way so you don't throw out the baby with the bath water. This is a topic of much current research.

So much for proteins. RNA is a more interesting case. As with proteins, RNA is a heteropolymer. The old dogma was that all of the biologically relevant information in RNA was in the sequence of nucleotides. This sequence in mRNA determines the sequence of amino acids in a protein. Now we know that this is not the case. RNA folds into complex 3-D structures that mediate a variety of important functions. In 1967 Carl Woese, Francis Crick, and Leslie Orgel discovered that RNA has complex secondary structure, and in the late 1970s Thomas Cech and Sidney Altman discovered that RNA could catalyze various reactions. We now know that RNA secondary structure plays an important role in the cell, for instance as the catalytic unit of the ribosome. Usually catalytic RNAs have a complexed metal ion such as a Mg++ ion in the active site. We still don't know exactly how important RNA structure is, partially because RNA is notoriously difficult to crystallize. There has been an approach similar to alphaFold for RNA: Townshend, Raphael JL, et al. "Geometric deep learning of RNA structure." Science 373.6558 (2021): 1047-1051.

DNA is, from a mechanical perspective, apparently the simplest of the biopolymers. Bare DNA is not known to adopt any interesting secondary or tertiary structure. In a eukaryotic cell the DNA is tightly compacted, but most of this compaction seems to be done by proteins. Exactly how organized DNA is within the cell and what is the significance of this organization, is a subject of much current research. Interestingly, DNA is the only biopolymer whose folding we understand well enough to design sequences that reliably fold into arbitrary structures, like cubes or smiley faces. It is interesting that the only biopolymer that doesn't fold into interesting structures *in vivo* is the only one we know how to fold *in vitro*. It is not known whether, or to what extent, the fact that DNA is a heteropolymer is significant for its mechanical properties.

Polysaccharides are generally believed to be amorphous in solution, thought I wouldn't be surprised if some day somebody finds out that they do have functionally important secondary and tertiary structure. Short polysaccharide chains form cyclodextrins, typically with 6, 7, or 8 sugars in a ring–although much larger cyclodextrins have also been found. Cyclodextrins are good at sequestering and solubilizing hydrophobic molecules, so they are often used in drug delivery. But cyclodextrins are short enough that we shouldn't really think of them as polymers.

One important reason to understand polymers is to understand the processes by which biomolecules adapt their native states. Traditionally much of the interest in polymers stemmed from other areas: plastics, fabrics, lubricants, coatings, and other industrial applications. The questions there are how to relate the microscopic structure of the polymer to the bulk rheological, optical, electronic, and mechanical properties of the material.

Today we'll be talking about unbranched, non-self interacting, homogeneous, dilute polymers. When any of these conditions is violated, interesting things happen. Branching leads to dendrimers and star polymers, both of which have different bulk mechanical properties (for instance branched polymer solutions are far more viscous because the polymers get into tangled topologies much more easily). Self-interactions are the basis of RNA and protein folding, as well as polymeric phase transitions that, for instance, make car tires stiff in winter. When you have self-interacting heteropolymers, then things get really interesting: in addition to the well-known biological examples, there are things called block copolymers, in which different domains in the polymer have different interactions. Under the right conditions, block copolymers will self-assemble into complex 2-d or 3-d nanostructured phases. There has been a lot of interest in using block copolymers as templates for other kinds of nanostructures. When you go to concentrated solutions, you have to think about intermolecular interactions as well as intramolecular ones. Polymer go through a series of different regimes as the concentration increases: initially each polymer acts independently. Then they start to overlap, and have a few points of contact. Then they become seriously entangled, and each polymer is embedded in a sea of other polymers. A huge amount of theoretical work has gone into predicting viscosities, specific heats, and osmotic pressures of polymers under each of these conditions.

Finally, I should mention polyelectrolytes: DNA, RNA, and proteins all have many charged groups. The intra- and intermolecular interactions depend sensitively on salt concentrations, for reasons we'll discuss soon. When you apply an electric field, polymers experience a net force. This force is a complicated beast because the polymer is pulled one direction by the field, but the counterions in the Debye layer around it are pulled in the opposite direction. The separation of polymer from counterions leads to a secondary electric field that counteracts part of the initial field. The motion of the counterions also leads to hydrodynamic flows that perturb the motion of the polyeletrolyte. Furthermore, the polyelectrolyte is often immersed in a mixture of other polymers, such as agarose or acrylamide. There is a whole industry devoted to calculating the motion of polyelectrolytes in electric fields because this is a technologically important process: it is used in electrophoresis experiments for DNA sequencing, and for protein purification and identification.

Polymers are fun because to understand them we need to draw on a wide range of ideas: random walks and diffusion, forces in solution, continuum mechanics, hydrodynamics.

2 Freely jointed chain

The simplest model of a polymer is a random walk of N steps. If each step has a length b and is independent of all the other steps, the root-mean-square end-to-end displacement in any number of dimensions is:

$$\sqrt{\langle R^2 \rangle} = \sqrt{Nb^2}. \tag{1}$$

This formula isn't as obvious as it might seem. If we were taking independent steps along the x, y, and z axes, then we'd have an rms displacement of $\sqrt{3Nb^2}$. However, if a given rod points along the x-axis, then it doesn't point along the y or z axes. This means that the steps along different axes *are not* independent. For your next homework you will calculate some other statistical properties of a random walk polymer.

There has been a resurgence of interest in polymer physics in recent years due to laser tweezers experiments in which people attach beads to opposite ends of a DNA strand and measure the elastic properties of a single strand, treating it as an entropic spring. We will sneak up on the force-extension curve of DNA by starting with a very simple picture, and then adding realism step-by-step. For those of you who have had some stat mech, you'll see that everything we do is exactly analogous to solving a 1-D Ising model. If that analogy helps you, then try to draw the parallels as we go through the derivations. If you've never heard of an Ising model, don't worry–just pay attention to the polymer physics.

2.1 1-D solution

Suppose that each strand can only point left or right and that we apply a force to the strand $F\hat{x}$. Then each rod has two states, of energy $U_{L,R} = \pm Fb$. Thus the probability of taking a step to the left or right is given by:

$$P_{L,R} = \frac{e^{\pm Fb\beta}}{e^{Fb\beta} + e^{-Fb\beta}}, \tag{2}$$

where the $+$ in the numerator is for the rod pointing to the right, and the $-$ is for the rod pointing to the left. I've substituted $\beta \equiv 1/k_B T$ as a matter of convenience (I should've done this long ago). Each rod chooses its orientation independently of its neighbors, so the total end-to-end displacement is equal to N times the expected single-rod displacement. Thus:

$$\begin{aligned} \langle x \rangle &= Nb\frac{e^{Fb\beta} - e^{-Fb\beta}}{e^{Fb\beta} + e^{-Fb\beta}} \\ &= Nb \tanh(Fb\beta). \end{aligned} \tag{3}$$

In the low force limit, the spring constant is given by $1/k = \partial x / \partial f$, so the spring constant is:

$$k = \frac{k_B T}{Nb^2}. \tag{4}$$

We could have guessed this from the Equipartition Theorem. In the absence of force we have $\langle x^2 \rangle = Nb^2$. But for a linear spring, we have $\langle x^2 \rangle = k_B T/k$. Setting these two expressions equal gives the spring constant. The same reasoning applies for the time-independent response of many systems. For instance the isothermal compressibility of a gas is related to its volume fluctuations under constant pressure.

Equivalently, the equilibrium probability distribution of x (with no applied force) is, from the central limit theorem:

$$P(x) = \frac{1}{\sqrt{2\pi N b^2}} \exp[-x^2/2Nb^2]. \tag{5}$$

If we treat this distribution as the thermal distribution arising from a potential of the form $U = \frac{1}{2}kx^2$, we get the result above, $k = k_B T/Nb^2$.

You can also get the same result in a more complex way from the fluctuation dissipation theorem. We already saw that in 1-d the equilibrium end-to-end fluctuations have mean-square magnitude $\langle x^2 \rangle = Nb^2$, and so the linear response to a potential coupling to x, $U = -kx$, is just $k_B T/\langle x^2 \rangle$.

Proof: Assume that the autocorrelation of x decays exponentially with time (in the end it won't matter what we choose for the time constant). Then $\langle x(t)x(0) \rangle = \langle x^2 \rangle e^{-\alpha t}$. From the FDT we get the response function of the system to a perturbation:

$$\begin{aligned} G(t) &= = -\frac{1}{k_B T}\frac{d}{dt}\langle x(t)x(0)\rangle \\ &= \frac{\langle x^2 \rangle}{k_B T}\alpha e^{-\alpha t}. \end{aligned} \tag{6}$$

On the other hand, we want that at long times the displacement should obey $x = f/k$, where k is the spring constant. One way to implement this with a finite-duration response function is:

$$x = \int_{-\infty}^{t} G(t-t')f(t')dt' \tag{7}$$

with

$$G(t) = \frac{1}{k}\alpha e^{-\alpha t}. \tag{8}$$

Equating the two expressions for G yields $k = k_B T/\langle x^2 \rangle$.

2.2 1 D solution in terms of free energy and partition function

It is instructive to redo the simple calculation above, using the partition function. Even though this is redundant, the expressions we get will be very helpful when we have a more complex system. Recall that for a discrete state system, the partition function is defined as:

$$Z = \sum_i e^{-\beta U_i} \tag{9}$$

where I've introduced the notation $\beta = 1/k_B T$. This expression is super useful because you can calculate the expectation values of lots of different things by taking derivatives of $\ln Z$ with respect to different quantities. Here's how the game works.

For the two-state system, the partition function is $Z_1 = e^{-\beta Fb} + e^{\beta Fb}$, or $Z = 2\cosh \beta Fb$. Suppose I want to calculate $\langle x \rangle$. This is just

$$\langle x_1 \rangle = \frac{1}{\beta}\frac{\partial \ln Z_1}{\partial F}. \tag{10}$$

The 1 subscripts indicate that this is for a single step in the chain. We could apply the CLT and say the net extension for N steps is $\langle x_N \rangle = N\langle x_1 \rangle$, or equivalently we could say that $Z_N = Z_1^N$ (assuming each link in the chain is independent of all other links).

Recall the relation between the partition function and the free energy:

$$\begin{aligned}
\langle S \rangle &= -k_B \sum_i p_i \ln p_i \\
&= -k_B \sum_i p_i(-\epsilon_i \beta - \ln Z) \\
&= k_B(\langle E \rangle \beta + \ln Z) \\
&= \frac{\langle E \rangle}{T} + k_B \ln Z. \qquad (11)
\end{aligned}$$

In going from the first line to the second, we used the fact that

$$p_i = \frac{e^{-\beta \epsilon_i}}{Z} \tag{12}$$

The above derivation shows that

$$A = -\frac{1}{\beta}\ln Z. \tag{13}$$

where $A = \langle E \rangle - TS$ is the Helmholtz Free Energy, which for our purposes we can treat the same as the Gibbs Free Energy, G. We see that

$$\langle x \rangle = -\frac{\partial G}{\partial F}. \tag{14}$$

For a purely mechanical spring ($U = F^2/2k$) you'll see that this relation just returns Hook's law.

You can use the same trick to get the fluctuations, $\langle x^2 \rangle$. These are:

$$\langle x^2 \rangle = \frac{1}{Z\beta^2}\frac{\partial^2 Z}{\partial F^2}. \tag{15}$$

2.3 3-D solution

Monomers in a polymer aren't constrained to lie along the x-axis; they can point any which way. Let's consider this next. Let $\mathbf{b}$ be the vector orientation of one rod (because we still assume the rods are independent, we don't need to specify which rod). Then the energy of that rod subject to a force $\mathbf{F} = F\hat{x}$ is:

$$U = -\mathbf{b} \cdot \mathbf{F}. \tag{16}$$

The partition function for the rod is:

$$\begin{aligned} Z_1 &= \int_\Omega d\Omega e^{\beta \mathbf{b}\cdot\mathbf{F}} \\ &= \int_0^{2\pi} d\phi \int_0^{\pi} d\theta \sin\theta e^{\beta b F \cos\theta}, \end{aligned} \tag{17}$$

where θ is the angle between the rod and the force. We evaluate this integral by making the substitution $W = \beta bF\cos\theta$, so that $\sin\theta d\theta = -dW/\beta bF$. Then:

$$\begin{aligned} Z_1 &= \frac{2\pi}{\beta bF}\int_{-\beta bF}^{\beta bF} e^W dW \\ &= \frac{4\pi}{\beta bF}\sinh \beta bF. \end{aligned} \tag{18}$$

Now the net extension is given by $\langle x_1\rangle = \langle b\cos\theta\rangle$. By looking at the definition of Z_1, we see that this quantity can be expressed as:

$$\begin{aligned} \langle x_1\rangle &= \frac{\int_\Omega d\Omega b\cos\theta e^{\beta bF\cos\theta}}{\int_\Omega d\Omega e^{\beta bF\cos\theta}} \\ &= \frac{1}{\beta}\frac{\partial}{\partial F}\ln Z_1 \\ &= b\coth(\beta bF) - \frac{1}{\beta F}. \end{aligned} \tag{19}$$

The total displacement is just N times this result. Our new formula looks entirely different from the formula we had with the two-state model. In particular, the zero-extension spring constant is different. Evaluating $1/k = \partial\langle x\rangle/\partial F|_{F=0}$ gives

$$k = \frac{3k_BT}{Nb^2}. \tag{20}$$

So the freely jointed chain (FJC) is 3 times stiffer in 3-D than in 1-D. I'll let you guess what its stiffness is in 2-D. The origin of the increased stiffness is that the sphere has a fat waist. There are a lot of ways for the rod to point orthogonal to the applied force, but only few ways for it to point along the applied force. So aligning the rod costs extra entropy.

[MENTION USE OF POLYMERS AS STABILIZERS OF COLLOIDAL DISPERSIONS (E.G. GUAR GUM AND GUM ARABIC). MENTION STABILIZATION OF CARBON NANOTUBES]

3 wormlike chain

The FJC model accurately describes the small-force properties of a polymer, but fails for larger forces. Why? You can think a random walk as a sum of sinusoids, each with its own amplitude and frequency. The free energy of the long-wavelength modes is entirely dominated by entropy (because they have low curvature), but the short wavelength modes have a contribution from enthalpy too: to bend a polymer in a sharp kink costs elastic energy. As you stretch out a polymer, the long wavelength modes become increasingly suppressed, and the global behavior becomes increasingly dominated by the short wavelength modes. It is these modes that we have not handled correctly in the preceding discussion.

We need to switch from describing the polymer in terms of a chain of non-interacting rods, to a language where we talk about global collective modes, which are decoupled from each other. In the language of the Ising model, adjacent spins are coupled, so now single spins are no longer the correct eigenstates of the system, but spin waves are. This approach that we're about to use is common in many-body theories of interacting systems: the individual particles are no longer good states for describing the system; instead you can decouple the motion by choosing appropriate collective motions.

Here we go. How much energy does it take to bend a rod? We will discuss this in detail when we cover continuum mechanics. For now I will simply state that the elastic energy per unit length is quadratic in the curvature. This should be plausible. The total elastic energy of a rod is:

$$H = \int_0^L \frac{1}{2}\alpha \left(\frac{\partial \theta}{\partial s}\right)^2, \tag{21}$$

where α is a parameter called the bending modulus of the rod, which describes how stiff it is, and $\theta(l)$ is the angle of the unit tangent relative to some arbitrary axis along the rod. This is called the Kratky-Porod model. In the absence of force, the unit-tangent to the wormlike chain undergoes orientational diffusion in a manner exactly analogous to the rotational diffusion of a rodlike particle.

Let's make the *ansatz* that the normal modes of the system are sinusoids. We assume that the polymer is pinned at its two ends, so the modes are just those of a violin string:

$$\theta_n(s) = A_n \cos(n\pi s/L). \tag{22}$$

Now we need to find the energy U_n. From the first term we get

$$\begin{aligned} U_n &= \int_0^L \frac{1}{2}\alpha \left(\frac{\partial \theta}{\partial s}\right)^2 ds \\ &= \frac{1}{2}\alpha \left(\frac{A_n n\pi}{L}\right)^2 \int_0^L \sin^2(n\pi s/L) \\ &= \frac{1}{4}\alpha \frac{(A_n n\pi)^2}{L}. \end{aligned} \tag{23}$$

Note that the energy is quadratic in both the amplitude and the wavenumber.

How much does a sinusoidal oscillation shorten the end-to-end length of the rod? To calculate this, let's make the approximation that θ is small. This is a reasonable approximation, since we are dealing with the high force regime in which the polymer is already mostly stretched out. We get:

$$\begin{aligned} x_{\text{end}} &= \int_0^L \cos(\theta) ds \\ &\approx \int_0^L [1 - \frac{1}{2} A_n^2 \cos^2(n\pi s/L)] ds \\ &= L - \frac{L A_n^2}{4}. \end{aligned} \tag{24}$$

If we have an extensional force F, then the polymer must do work to shorten against the applied force. This work is just given by $-F\Delta x_{\text{end}}$, or

$$U_{\text{mech}} = \frac{F L A_n^2}{4}. \tag{25}$$

Thus the total energy of a sinusoidal deformation in mode n of amplitude A_n is

$$U[\theta_n(s)] = \frac{A_n^2}{4} \left[\alpha \frac{(n\pi)^2}{L} + FL \right]. \tag{26}$$

This is great, because we've found a solution where the energy in mode n depends quadratically on the amplitude in mode n only. (We didn't actually show that this solution *decoupled* different modes, you can either take my word for that, or prove it yourself by plugging into the Hamiltonian $\theta(s) = \sum_n A_n \theta_n(s)$). Due to the quadratic dependence of U on A we can apply equipartition to immediately get the amplitude in mode n:

$$\langle A_n^2 \rangle = \frac{2 k_B T}{\alpha \frac{(n\pi)^2}{L} + FL}. \tag{27}$$

The total extension is

$$\begin{aligned} \langle x \rangle &= L \left(1 - \frac{1}{4} \sum_n 2 \langle A_n^2 \rangle \right) \\ &= L \left(1 - \sum_{n=1}^{\infty} \frac{k_B T}{\alpha \frac{(n\pi)^2}{L} + FL} \right). \end{aligned} \tag{28}$$

That extra factor of 2 in the sum is important. It came because there are actually *two* modes at each wavevector n: one mode along the y-axis and one along the z-axis. Rather than evaluating the sum exactly, we'll look at its scaling behavior.

[INSERT DRAWING OF SUMMAND VS. n]

For $n \ll \frac{L}{\pi}\sqrt{\frac{F}{\alpha}}$, the summand is roughly constant and equal to $k_B T/(FL)$. The summand then rapidly drops to 0 around $n \approx \frac{L}{\pi}\sqrt{\frac{F}{\alpha}}$. The sum scales as:

$$\sum_{n=1}^{\infty} \frac{k_B T}{\alpha \frac{(n\pi)^2}{L} + FL} \approx \frac{k_B T}{\sqrt{F\alpha}}. \tag{29}$$

A more detailed calculation shows that there is a factor of 1/4 in the sum. Plugging this into our formula for the length gives:

$$\langle x \rangle \approx L \left(1 - \frac{k_B T}{4\sqrt{F\alpha}} \right). \tag{30}$$

This result has very different scaling from the FJC model. In the FJC the extension approached L as $1/F$. Here we approach L as $1/\sqrt{F}$. Laser tweezers experiments are accurate enough to distinguish clearly between these two cases. I've posted the Marko-Siggia paper with this derivation on the website for your amusement. It is a really neat paper, where they develop a path-integral formalism for polymer physics.

[SHOW DATA ON FORCE EXTENSION CURVES OF DNA]

What's the zero-force radius of gyration of our elastic rod model? A characteristic length is related to the distance over which thermal fluctuations can bend the rod. This length is given by $l_p = \alpha/k_B T$ and is called the persistence length. It can be shown that then rms end-to-end extension is given by:

$$\sqrt{\langle R^2 \rangle} = \sqrt{2Ll_p}. \tag{31}$$

By comparison with the FJC model, we see that these models give the same zero-force distribution for $b = 2l_p$.

4 Self-interaction

Two distinct elements of a polymer chain cannot occupy the same space at the same time. This "excluded volume" interaction subtly alters the thermodynamics of the polymer and its ultimate shape distribution. Let's talk about this effect in a general way. We'll use the Flory mean-field approach.

Consider a polymer of N steps of length b in d dimensions. Suppose there is an energy cost ϵ to have two monomers overlap. We want to estimate the number of overlaps. The chain occupies a volume with linear dimension $r \sim bN^\nu$. So far we've had $\nu = \frac{1}{2}$, but we'll see that self-exclusion changes this. The occupied volume is $V \sim r^d$, and the volume fraction occupied by monomers is $\rho = aN/V \sim aNr^{-d}$ where a is the d-dimensional volume of one monomer. Now we make the grossly simplifying mean-field approximation that the monomers are uniformly and independently distributed throughout this volume (i.e. we ignore the fact that it's a polymer). Then the probability that any given monomer lands on a spot already occupied by another monomer equals ρ. Each monomer has this probability, so the total self-exclusion energy is $U_{\text{self}} = \epsilon a N^2 r^{-d}$.

Recall that we calculated an effective entropic spring constant that tries to bring the ends of the polymer together. This spring constant was $k = dk_B T/Nb^2$. The swelling caused by self-repulsion is counteracted by the attraction from the entropic spring. The total energy is

$$U_{\text{tot}} = \frac{d}{2} \frac{k_B T}{Nb^2} r^2 + \epsilon a N^2 r^{-d}. \tag{32}$$

At equilibrium $\partial U/\partial r = 0$, or

$$\frac{dk_B T}{Nb^2} r = d\epsilon a N^2 r^{-d-1}. \tag{33}$$

Solving for r in terms of N yields:

$$r = \left(\frac{\epsilon a b^2}{k_B T}\right)^{\frac{1}{d+2}} N^{\frac{3}{d+2}}. \tag{34}$$

This result is very interesting: it predicts that in 1-D $r \sim N$, in 2-D $r \sim N^{3/4}$, in 3-D $r \sim N^{3/5}$, and in 4-D $r \sim N^{1/2}$. When we ignored self-interaction we got $r \sim N^{1/2}$ in all dimensions. The effect of excluded volume is increasingly severe as we go to lower number of dimensions. This is obvious considering a 1-D self-avoiding random walk: after the first step you have to keep going in the same direction or you run into yourself. Interestingly in 4 or more dimensions the mean-field result is exact, and self-interactions are irrelevant. A huge amount of effort has been expended on developing more sophisticated models of the relation between r and N. Currently we believe that in 3-D $r \sim N^{0.588\ldots}$. This result is supported by the recent experiments of Roberts and Smith. The reason for all this interest in the scaling of polymer solutions is that the math is the same as for many other physical problems of interest, such as magnetism and other phase transitions. It turns out that mean-field models are exact in 4 or more dimensions, and very clever exact solutions have been found in 1 and 2 dimensions, but nobody has been able to make an analytical model for phase transitions in 3 dimensions. Ironic.