3 Disordered systems: the spherical spin glasses

Disordered systems are broadly defined as ones whose properties depend on the value of a large number of parameters that are frozen in the system's evolution but vary from sample to sample. The prototypical example are *spin glasses*, which are materials made from doping metals with magnetic impurity atoms. The positions of the impurities in the metal are disordered, meaning that they are frozen but vary from sample to sample. The interaction between the impurities is sensitively dependent on their distance apart, so much so that even the sign of the interaction (whether the spins want to align or anti-align) changes from very small changes of distance. The effect of this is very well modeled by drawing the interaction coefficient for each pair of spins at random from a fixed distribution, with a magnitude that dies off with distance.

The Ising model of spin glasses is the Edwards–Anderson model, which consists of spins on a lattice taking values ±1, and has Hamiltonian

$$H_J(\mathbf{s}) = -\frac{1}{2}\sum_{\langle ij\rangle} J_{ij}s_is_j \tag{1}$$

At first glance it is exactly the same as the Ising model, with the only difference being that the coefficients J_{ij} now depend on the bond. We have also made clear in the notation that the Hamiltonian H *depends* on the values of the set of Js. The most common choice for their distribution is to make them independent Gaussians, with

$$P(J_{ij}) = \frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}J_{ij}^2} \tag{2}$$

for each of the J_{ij}. We also take J to be symmetric: $J_{ji} = J_{ij}$ (so not every element is independent). We write the average over all the J with an overbar, like

$$\overline{\star} = \int\prod_{\langle ij\rangle} dJ_{ij}\,P(J_{ij})\star \tag{3}$$

Because the J are Gaussian, we can characterize them completely by their mean and covariance, which are

$$\overline{J_{ij}} = 0 \qquad \overline{J_{ij}J_{kl}} = \delta_{ik}\delta_{jl} + \delta_{il}\delta_{jk} \tag{4}$$

Despite its simplicity, like the Ising model the Edwards–Anderson model can be solved under few circumstances and must be studied by approximations or numeric simulation.

The Curie–Weiss model of spin glasses is the Sherrington–Kirkpatrick model, which like the Curie–Weiss model is the Edwards–Anderson on a fully-connected graph. It is defined by

$$H_J(\mathbf{s}) = -\frac{1}{2\sqrt{N}}\sum_{ij} J_{ij}s_is_j \tag{5}$$

Despite its simplicity, the solution of the Sherrington–Kirkpatrick model is extremely rich, and its details were not completely settled for decades after the model was introduced.

For various reasons of convenience, we will instead study another family of mean-field models, the so-called spherical spin glasses. Like the spherical Curie–Weiss model, these are formed by taking the Sherrington–Kirkpatrick model and moving from the configuration space of the hypercube $\pm1^N$ to that of the sphere in N dimensions. However, unlike the spherical Curie–Weiss model, using the same Hamiltonian on the sphere does not produce a model with the properties we want. On the sphere, the quadratic Hamiltonian H_J is *convex*, which means it has a simple structure. To see this, consider the fact that a symmetric matrix J can be diagonalized by $J = O^TDO$ where O is an orthogonal matrix and D is a diagonal matrix containing the eigenvalues of J. Making the change of variables $\mathbf{s} \mapsto O\mathbf{s}$ is equivalent to choosing a set of rotated axes, and results in the Hamiltonian

$$H_J(\mathbf{s}) = -\frac{1}{2\sqrt{N}}\sum_{i} \lambda_i s_i^2 \tag{6}$$

which is a Hamiltonian for non-interacting quadratic degrees of freedom. This can be solved trivially and does not behave like a spin glass. In order to have a nontrivial model, we need to tweak the model we study. One way to do this is to move from pairwise interactions to multi-spin interactions. We consider a model where every spin interactions with $p-1$ others for $p>2$, where

$$H_J(\mathbf{s}) = -\frac{1}{\sqrt{2p!N^{p-1}}}\sum_{i_1,\ldots,i_p} J_{i_1,\ldots,i_p}s_{i_1}\cdots s_{i_p} \tag{7}$$

where the prefactor constants have been chosen to make the Hamiltonian of order N. Now J is not a symmetric matrix, but a symmetric p tensor, with Gaussian entries whose mean and covariance are

$$\overline{J_{i_1,\ldots,i_p}} = 0 \qquad \overline{J_{i_1,\ldots,i_p}J_{j_1,\ldots,j_p}} = \delta_{i_1j_1}\cdots\delta_{i_pj_p} + \text{permutations of js} \tag{8}$$

For instance, for $p=3$ this Hamiltonian is

$$H_J(\mathbf{s}) = -\frac{1}{\sqrt{12}N}\sum_{ijk} J_{ijk}s_is_js_k \tag{9}$$

This may seem like a strange choice, somewhat far from the original models we were trying to study. However, there are strong motivations for studying the spherical spin glasses. As we will see later in the course, their dynamics are schematically the same as those predicted for *structural* glasses by mode-coupling theory. In addition, the form of the phase transitions in these models, though not identical to that of the Sherrington–Kirkpatrick model, are closely related, and can be made effectively identical by combining multiple orders of interaction.

We can examine some statistical properties of this model. At any point $\mathbf{s}$ the average energy is zero, since

$$\overline{H_J(\mathbf{s})} = -\frac{1}{\sqrt{2p!N^{p-1}}}\sum_{i_1,\ldots,i_p}\overline{J_{i_1,\ldots,i_p}}s_{i_1}\cdots s_{i_p} = 0 \tag{10}$$

since $\overline{J_{i_1,\ldots,i_p}} = 0$ for each of the elements of the tensor J. Since the Hamiltonian is a linear combination of Gaussian random variables, it is itself a Gaussian random variable. Therefore, its statistical properties are completely described by its mean and covariance. We have seen the mean is zero, but the covariance at two different points $\mathbf{s}$ and $\mathbf{s}'$ on the sphere is

$$\begin{aligned}\overline{H_J(\mathbf{s})H_J(\mathbf{s}')} &= \frac{1}{2p!N^{p-1}}\sum_{i_1,\ldots,i_p}\sum_{j_1,\ldots,j_p}\overline{J_{i_1,\ldots,i_p}J_{j_1,\ldots,j_p}}s_{i_1}\cdots s_{i_p}s'_{j_1}\cdots s'_{j_p}\\ &= \frac{1}{2p!N^{p-1}}\sum_{i_1,\ldots,i_p}\sum_{j_1,\ldots,j_p}(\delta_{i_1j_1}\cdots\delta_{i_pj_p} + \text{permutations of js})s_{i_1}\cdots s_{i_p}s'_{j_1}\cdots s'_{j_p}\\ &= \frac{1}{2N^{p-1}}\sum_{i_1,\ldots,i_p}s_{i_1}s'_{i_1}\cdots s_{i_p}s'_{i_p}\\ &= \frac{1}{2N^{p-1}}\left(\sum_i s_is'_i\right)^p = N\frac{1}{2}\left(\frac{\mathbf{s}\cdot\mathbf{s}'}{N}\right)^p\end{aligned}$$

Here, we see that the covariance of the energy at two different points depends only on the *overlap* $q = \frac{\mathbf{s}\cdot\mathbf{s}'}{N}$ between those configurations, or

$$\overline{H_J(\mathbf{s})H_J(\mathbf{s}')} = Nf\left(\frac{\mathbf{s}\cdot\mathbf{s}'}{N}\right) \tag{11}$$

with $f(q) = \frac{1}{2}q^p$ for the p-spin models. The fact that the covariance depends only on the overlap between the two configurations is reflects the fact that the spherical models are statistically isotropic: statistical properties of the energy depend only on the distance between points where it is measured. In particular, we see that configurations that are orthogonal to each other on the sphere have energies that are statistically independent.

How do we study the statistical mechanics properties of this model? We can define the partition function of the model for a given set of Js in the usual way by writing

$$Z_J(\beta) = \int_\circ d\mathbf{s}\, e^{-\beta H_J(\mathbf{s})} \tag{12}$$

This clearly depends on the specific set of Js. As a first attempt, we can try to calculate the partition function averaged over the Js to try to understand the model's properties. We would have

$$\overline{Z_J(\beta)} = \int_\circ ds\, \overline{e^{-\beta H_J(s)}} \tag{13}$$

Using properties of Gaussian random variables, we can write

$$\overline{e^{-\beta H_J(s)}} = e^{\frac{1}{2}\beta^2 \overline{H_J(s) H_J(s)}} = e^{\frac{1}{2}N\beta^2 f(\frac{s\cdot s}{N})} = e^{\frac{1}{2}N\beta^2 f(1)} \tag{14}$$

for the covariance function $f(q) = \frac{1}{2}q^p$. We therefore find

$$\overline{Z_J(\beta)} = \int_\circ ds\, e^{\frac{1}{2}N\beta^2 f(1)} = e^{\frac{1}{2}N\beta^2 f(1)} \int_\circ ds \tag{15}$$

completely independent of s! The integral over the sphere just gives the volume of the sphere, which to principal order in N gives what we found for the Curie–Weiss model, or

$$\overline{Z_J(\beta)} \propto e^{\frac{1}{2}N\beta^2 f(1) + \frac{1}{2}N(1+\log(2\pi))} \tag{16}$$

and the model appears to already be solved! Since $f(1) = \frac{1}{2}$ for all values of p, we have a 'free energy' per site of

$$F = -\frac{1}{\beta N}\log \overline{Z_J(\beta)} = -\frac{1}{4}\beta - \frac{1}{2\beta}(1+\log(2\pi)) \tag{17}$$

This implies an energy per site of

$$E = \frac{\partial}{\partial \beta}(\beta F) = -\frac{1}{2}\beta \tag{18}$$

with apparently no phase transition. If this were true, it would be a pretty poor model of a spin glass, which are know to have a novel low-temperature phase! However, other aspects of this 'solution' should give us pause. As the temperature is lowered and β is increased, the energy density decreases without bound towards negative infinity. But the energy H_J is just a polynomial with random coefficients in a bounded configuration space, and should be finite everywhere! We can really see the inconsistency in the exact solution of the $p = 2$ model. We know from its diagonalized form that the ground state energy will correspond with a configuration s corresponding to the eigenvector of J with the smallest eigenvalue, and therefore that the energy will have the value $\frac{1}{\sqrt{2N}}\lambda_{\text{min}}N = \sqrt{\frac{N}{2}}\lambda_{\text{min}}$. In the field of random matrix theory, symmetric matrices of random Gaussian numbers are very well studied and their average spectrum is well understood. For large system size N, the spectrum of such a matrix approaches a distribution called the Wigner semicircle, which is an ellipse with support between $\pm 2\sqrt{N}$. Therefore, the minimum eigenvalue of J will approach $-2\sqrt{N}$ at large N, and therefore the ground state energy density will be -1, certainly not $-\infty$!

So what has gone wrong here? We have a hint if we write the average partition function more explicitly, like

$$\begin{aligned}\overline{Z_J(\beta)} &= \int \prod_{i_1,\ldots,i_p} dJ_{i_1,\ldots,i_p}\, P(J_{i_1,\ldots,i_p}) \int_{\circ} ds\, e^{-\beta H_J(s)} \\ &= \int \prod_{i_1,\ldots,i_p} \frac{dJ_{i_1,\ldots,i_p}}{\sqrt{2\pi}} \int_{\circ} ds\, e^{-\beta H_J(s) - \frac{1}{2}\sum_{i_1,\ldots,i_p} J^2_{i_1,\ldots,i_p}} \\ &= \int dJ\, ds\, e^{-\beta H_{\text{eff}}(s,J)}\end{aligned} \tag{19}$$

The average partition function has the form of the partition function for a model where the s and the tensor J are on equal footing as configuration space parameters and have a quadratic confinement. What we are really looking at when we look at the average partition function is the statistical mechanics of a model where the Js are not fixed but fluctuate as easily as the spins do. When the system is cooled, such a model can find arbitrarily low-energy states by exhibiting combinations of the Js that would be unlikely to be drawn as independent Gaussian random variables. The kind of average we have taken is called the *annealed* average, a reference to a procedure in the preparation of metals and glasses where the material is cooled from very high temperature very slowly, so that the internal degrees of freedom (like the position of magnetic impurities and the resulting interactions between spins) have a lot of time to explore configurations and find a low-energy state. It is in essence ignoring the disordered aspect of the problem and treating the whole system, spins and couplings, as one big statistical mechanics problem.

Mathematically, how can this happen? If we *average* over a partition function, we might have the expectation that the result is indicative of its typical value. However, this is not necessarily the case, and is often not the case in situations where the object being averaged is the large exponential of a random variable, like we have. The reason for this is that such averages can be dominated by rare configurations whose sheer magnitude outweighs the influence of the contributions from more common but smaller values. Consider the log-normal distribution, which is defined by taking a Gaussian random variable X with mean $N\mu$ and variance $N\sigma^2$ and studying e^X. We know that the most common value sampled will be around $e^{N\mu}$, the most common value of X. In fact, as N grows we have said the standard deviation should even shrink relative to the mean! However, the mean of the distribution is

$$\overline{e^X} = e^{N\mu + \frac{1}{2}\overline{(X-\mu)^2})} = e^{N(\mu + \frac{1}{2}\sigma^2)} \tag{20}$$

In traditional statistical mechanics, quantities like the energy have an average that scales with N and a standard deviation that scales like $\sqrt{N}$. In that case, the ratio between the most common value and the average value scales like

$$\overline{e^{X}}/e^{N\mu} = e^{\frac{1}{2}N\sigma^{2}} \tag{21}$$

or exponential in N! This would likewise result in an extensive difference in the free energy per site.

However, this is not what we want! We want to understand what happens to a *typical* sample of our system given a tensor of interactions J that is actually composed of independent Gaussian random numbers, not a rare outlier. How do we accomplish this? There is a hint of an answer in the log-normal distribution, where the typical value is restored by taking the average of the logarithm, or $\overline{\log e^{X}} = \overline{X} = N\mu$. This is a bit trivial, but it hints to us that taking the average of the free energy rather than the partition function should give an answer that is true for typical samples. We therefore want to compute

$$F = \overline{F_{J}(\beta)} = -\frac{1}{\beta}\overline{\log Z_{J}(\beta)} \tag{22}$$

This average is known as the *quenched average*, again after techniques in metallurgy and glassmaking, where quenching a material means to submerse it suddenly in cold water so as to cool it extremely quickly and given the material no time at all to find a low-energy configuration as it cools. The quenched free energy has the properties that we actually are seeking to study: it is the value of the free energy for typical samples. We say that the free energy, unlike the partition function, is a *self-averaging* quantity as a result. A consequence of self-averaging is that single large samples of a material will exhibit the same free energy density as the average free energy density, even though no average over disorder has been taken when looking at a single large sample (hence self averaging).