

Mathematical Physics

Lectures by Professor Michael Aizenman
Scribes: Holden Lee and Kiran Vodrahalli

April 19, 2016

1	Introduction to statistical mechanics	7
1	Canonical Ensembles for the Lattice Gas	7
1.1	Configurations and ensembles	7
1.2	The equivalence principle	9
1.3	Generalizing Ensemble Analysis to Harder Cases	11
2	Concavity and the Legendre transform	13
2.1	Basic concavity results	13
2.2	Concave properties of the Legendre transform	15
3	Basic setup for statistical mechanics	18
3.1	Gibbs equilibrium measure	19
3.2	Introduction to the Ising model	20
3.3	Entropy, energy, and free energy	21
4	Large deviation theory	25
5	Free energy	27
5.1	Basic Properties	28
5.2	Convexity of the pressure and its implications	36
5.3	Large deviation principle for van Hove sequences	39
6	1-D Ising model	41
6.1	Transfer matrix method	44
6.2	Markov chains	44
7	2-D Ising model	46
7.1	Ihara graph zeta function	54
8	Gibbs states in the infinite volume limit	56
8.1	Conditional expectation	59
8.2	Symmetry and symmetry breaking	61
8.3	Phase transitions	64
9	Random field models	70
10	Proof of symmetry-breaking of continuous symmetries	70
10.1	The spin-wave perspective	71
10.2	Infrared bound	73
10.3	Reflection positivity	74

Introduction

Notes from Michael Aizenman's class "Mathematical Physics" at Princeton in Spring 2016.

This class is of interest to both physicists and mathematicians. Several recent Fields medals are for work related to these topics.

I plan to cover the following topics. The focus is on **Topics in Mathematical Statistical Mechanics**.

1. **The statistical mechanic perspective:** systems can be described at the microscopic level with many degrees of freedom. We observe their collective behavior and find emergent behavior.
2. **Thermodynamics principles:** Intellectually preceding statistical mechanics is thermodynamics, a field of physics which emerged through experimental and intellectual work trying to understand what is happening with the transfer of heat. A couple of principles emerged. This framework is more appropriate to the macroscopic description of physical systems. In departure from mechanics, which cares about equalities like $F = ma, E = mc^2$, a unique thing about thermodynamics is that its key principle is in inequality: entropy increases.

$$\Delta S \geq 0$$

3. **The emergence of thermodynamics from statistical mechanics via the equidistribution assumption and the "large deviation theory"**. I would like to discuss the emergence of thermodynamics from statistical mechanics. Mathematicians formalized the theory but the concepts were introduced earlier by physicists. These principles led Boltzmann to introduce these ideas and it took a while for physics to absorb the ideas.¹
4. **Phase transitions:** When you have a system, say H_2O , which you can control with temperature and pressure, you can induce changes in state. Continuous change in control parameters results in discrete jumps in the result.
5. **Critical phenomena, critical exponents, universality classes.** Phase transitions are fascinating since there are interesting critical phenomena which are characterized

¹This may have led to his premature death via suicide.

by critical exponents, which turn out (this is one of the surprising discoveries experimentally) to result in universality classes of critical phenomena. Systems are macroscopically different, but the singularities you observe are given by the same power laws.

6. **Exact solution of the 2-D Ising model.** Mathematically, 3 is the hardest dimension to comprehend.
- (a) 1-D is solvable: correlations can be described by Markov chains, and can be computed.
 - (b) In 2-D, the conformal group is very important. It gives many constraints on critical behavior, leading to a rich behavior.
 - (c) In 3-D, this does not apply except for ongoing work finding consequences in 3-D from results in 2-D.
 - (d) Anything with ≥ 4 dimensions gets simpler. High dimensions are characterized by the fact that loop effects do not play a large role. In 1-D a simple random walk is recurrent. In sufficient high dimensions, a simple random walk is not recurrent. The infinite-dimensional case reduces to models on trees; high dimensions exhibit the same behavior as infinite dimensions.

In coding theory and discrete math people study phase transitions of different graphs and are interested in the same topics.

The 2D Ising model is usually a highly specialized topic, but we will do it in a way done a bit differently from normal. We would like to do this through “graph zeta functions”. What I found fascinating about this topics is that there are a lot of connections to other topics. There are analogues of the Riemann zeta function on graphs, and there is a relationship. It is actually one of the simplest paths to tackle this proof by.

7. **Stochastic geometry behind correlation functions at criticality:** in the Ising model we have a collection of spin variables $\sigma_x = \pm 1$ for $x \in \mathbb{Z}^d$. The spins are correlated in such a way that agreement among neighbors is encouraged. Thus there are correlations which spread through the system.

It is interesting and powerful to represent this correlation between spins is via a “shadow system” in which you play the following game: decompose the collection of spins at random into connected clusters. You see spin values but you don’t see who is connected to whom. An analogy is that students form cliques in class, and then each clique chooses what to do and votes unanimously. So if you just saw the voting pattern, you would see some cliques, but the nature of correlations among the votes become transparent if you know the clusters.

The states for critical Ising models become larger, but also become fractal. Fractal geometric objects can be used to explain the structure. There is an interesting fractal geometry which tells us about correlation functions.

We will introduce all of this later in much detail from the ground up. This is related to percolation.

8. **Scaling limits:** Finally, if time permits, I would like to discuss scaling limits. You may want to know properties of a substance when you have something like a lattice that is very fine-grained, a statistical mechanical system which macroscopically is described by a myriad of variables. There are techniques of describing the relevant quantities from the macroscopic perspective, and this becomes relevant for describing the critical state. This provides a fascinating link between statistical mechanics and field theory. It's related to quantum field theory, except we are now in the Euclidean regime. There are some related results for quantum spin systems.

In a way we have a mouthful here, and one could probably give a full course on each of these topics. I will try not to be exhaustive, as we may not progress far. Previously I gave a course on random operators. Each week would cover some area of this subject, saying enough so you have a glimpse of the essence and get some comfort, and then moving on. It's now in a book format, and I would count it as a success if out of these lectures, something similar would emerge. The idea is not to be exhaustive, but to give enough key results. There is much more to be said, but I do not suppose to cover it all.

There is a broad spectrum of references, none of which I'll follow exclusively.

- Sacha Friedli and Yvan Velenik's online book project (on mathematical statistical mechanics) at <http://www.unige.ch/math/folks/velenik/smbook/index.html>. I recommend this for people totally new to the subject.
- David Ruelle formulated models and basic results mathematically (late 70's). His book helped physicists organize their thoughts. It became outdated quickly, but remains a good starting point and reference for the formalism.

Chapter 1

Introduction to statistical mechanics

2-2-16

1 Canonical Ensembles for the Lattice Gas

1.1 Configurations and ensembles

One way to start is with the axioms of statistical mechanics. Instead I'll take a simple problem, see how it works, and present results in that context. There are simple problems that teach us a lot. ("The elementary problems are the most precious, once you absorb them they are part of your makeup.") The simplest is a lattice gas.

Definition 1.1.1 (Lattice gas): A **lattice gas** is a substrate where at each lattice site there may or may not be a particle. This model has been used to describe alloys where you have a substrate which you can draw as a simple lattice (in particular, we will use \mathbb{Z}^d for simplicity). Each point may have a particle of a certain type.

The **configuration** is a function $n : \mathbb{Z}^d \rightarrow \{0, 1\}$:

$$n_x = \begin{cases} 1, & x \text{ is occupied,} \\ 0, & x \text{ is vacant.} \end{cases}$$

For now let us ignore energy conservation. Let us suppose the system is neutral, or at some infinite temperature; energy is not an issue.

I'll use L to denote the size of the box we are considering, and $\Lambda \subset \mathbb{Z}^d$ to be a region (subset of the system) and we have m particles. The **configuration space** is the space of possible n 's, $\Omega = \{0, 1\}^\Lambda$.

If you have a finite system and energy does not play a role, but there is conservation of the number of particles, if you shake the box the state may change. Essentially each configuration gets equal weight, and particles do not overlap.

Definition 1.1.2 (Equidistribution assumption): All particle configurations which have an equal number of particles have equal probability.

This gives rise to the notion of ensembles.

Definition 1.1.3 (Ensemble): An **ensemble** is a probability measure with respect to which you do averages over the configuration space Ω .

First, we define some notation.

Definition 1.1.4 (Indicator function):

$$\mathbb{1}[\text{cond}] := \begin{cases} 1, & \text{condition satisfied} \\ 0, & \text{else} \end{cases}$$

Now we define two specific ensembles, the **microcanonical ensemble** and the **grand canonical ensemble**.

Definition 1.1.5 (Microcanonical ensemble): In the microcanonical ensemble,

$$\mathbb{P}(n_\Omega) = \frac{\mathbb{1}[\sum_{x \in \Omega} n_x = N]}{Z}$$

where Z is a normalization constant.

For every function $f : \Omega \rightarrow \mathbb{R}$ assigning a real number to each configuration, define the **microcanonical ensemble average** by

$$\langle f \rangle_{N,n}^{\text{Can}} = \frac{\sum_{n \in \Omega} \mathbb{1}[n_x = N] f(n)}{\sum_{n \in \Omega} \mathbb{1}[\sum n_x = N]}.$$

The ensemble we will focus on for this course is the grand canonical ensemble.

Definition 1.1.6 (Grand canonical ensemble): The **grand canonical ensemble average** is defined as

$$\langle f \rangle_{\mu, \Lambda}^{\text{Gr.C}} = \frac{\sum_{n \in \Omega} e^{-\mu \sum_{x \in \Lambda} n_x} f(n)}{\sum_{n \in \Omega} e^{-\mu \sum_{x \in \Lambda} n_x}}$$

(Later on we will omit superscripts where it is clear.)

We can relate these two notions with the **equivalence principle**.

Definition 1.1.7 (Equivalence principle): Loosely, the equivalence principle says that for any “local function”, the microcanonical average is approximately the grand canonical ensemble average

$$\langle f \rangle_{N, \Lambda}^{\text{Can}} \approx \langle f \rangle_{\mu, \Lambda}^{\text{Gr.C}}$$

when we take $\mu = \frac{N}{|\Lambda|}$.

1.2 The equivalence principle

Consider functions which depend only on a system $\Lambda \subset \tilde{\Lambda}$ of much smaller volume, $|\Lambda| \ll |\tilde{\Lambda}|$. This is the sense in which the averages match up.

We can think of the difference as the difference between two societies. One is draconian. The first is highly centralized control economy. If you don't fit the build, you get weight 0 and you're thrown out. In the grand canonical, everything goes. Some contribute more than others, and the contributions depend on the parameter μ which is adjustable. The density of particles depends on μ . And there is a value of μ for which the density is equal to $N/|\tilde{\Lambda}|$. Here, the average of the draconian system is equal to the average of the lackadaisical system, asymptotically.

The micro-canonical ensemble is draconian: the number of particles is prescribed, all other configurations get weight 0. In the grand canonical ensemble, each configuration contributes. There is a value of μ where the density is the same; at that value the local average of the draconian system is asymptotically the same as that of the more relaxed system. This \approx becomes $=$ when you take the thermodynamic limit,

$$\begin{aligned}\tilde{\Lambda} &\rightarrow \mathbb{Z}^d \\ N &\rightarrow \infty \\ \frac{N}{|\tilde{\Lambda}|} &\rightarrow \rho.\end{aligned}$$

Now let us see where this comes from. What is the induced distribution of the micro-canonical ensemble on Λ ? If all you care about is the number of particles in the average you care about, you just care about getting a small system. If the whole system has a large number of particles, then it does not matter if you focus on a much smaller box since you can trade particles very easily, if the larger volume has a ton of particles. Under the microcanonical ensemble, the probability that $n_{|\Lambda}$ (n restricted to Λ) takes a particular value depends only on one quantity, which is the number of particles the configuration has ($\sum_{\Lambda} n_x = k$ in the big box). This expression is equal to the cardinality (number) of configurations in the rest of the big box ($n_{|\Lambda^c}$, Λ complement) such that the total number of points in Λ^c is equal to $N - k$. And that of course must be normalized. Now we come to the question of how many configurations there are in the complement of Λ .

We count the number of ways to complete the configuration in $\Lambda^c = \tilde{\Lambda} \setminus \Lambda$:

$$\frac{|\{n_{\Lambda^c} : \sum_{x \in \Lambda^c} n_x = N - k\}|}{Z}$$

where Z is a normalization constant.

The number of configurations of M particles in volume V is $\binom{V}{M} = \frac{V!}{M!(V-M)!}$. Using Stirling's approximation

$$\ln(M!) = M(\ln M - 1)(1 + o(1)),$$

After some gymnastics of an elementary nature, there is a fundamental formula. The logarithm of the number of configurations is of tremendous importance. Ludwig Boltzmann's

grave has the formula which opened people's eyes to what this mysterious **entropy** S truly is: the logarithm of the number of configurations.

Lemma 1.1.8 (Entropy). *Defining*

$$s(\rho) = -[\rho \ln \rho + (1 - \rho) \ln(1 - \rho)],$$

we have

$$\binom{V}{M} = \frac{V!}{M!(V-M)!} =: e^{S(M,V)} \approx e^{Vs(\rho)}$$

where density $\rho = m/V$ and $s(p) = -p \ln(p) - (1-p) \ln(1-p)$.

The ratio ρ varies between 0 and 1. $s(\rho)$ is concave and attains maximum of $\ln 2$ at $\frac{1}{2}$ where it has quadratic behavior.

Proof. Please do this exercise once in your life; it's good to do it once but not too often. \square

Shannon also found such a formula for entropy.

The implication is that if you slightly change the density, the number of configurations changes drastically. In physical substances V may be 10^{23} . The change would then be $e^{10^{23}\Delta s}$. In any average over configurations, only those at the peak contribute: "winner takes all."

What is the probability of observing k particles in the small box given n in the big box?

$$\mathbb{P}(n_\Lambda) \approx \frac{e^{|\Lambda^c|s\left(\frac{N-k}{|\widetilde{\Lambda}| - |\Lambda|}\right)}}{Z}$$

where Z is a normalizing constant. The probability is proportional in exponent to the volume of the complement multiplied times the entropy of the density, which we saw in the exercise above. Then we re-write the density of the small box in terms of the overall density:

$$\frac{N-k}{|\widetilde{\Lambda}| - |\Lambda|} = \rho - \underbrace{\left(\rho - \frac{N}{|\widetilde{\Lambda}| - |\Lambda|}\right)}_{\rightarrow 0 \text{ independent of } k} - \frac{k}{|\widetilde{\Lambda}| - |\Lambda|}$$

The entropy is infinitely differentiable except at the endpoints, so we can just expand. Whatever you see inside affects the number outside, but does not effect the density outside. This explains why we only have a tiny correction to ρ . Thus

$$s\left(\frac{N-k}{|\widetilde{\Lambda}| - |\Lambda|}\right) \approx s(\rho) - s'(\rho) \frac{k}{|\Lambda^c|}$$

Changing k by a little bit affects how many particles are outside but not so much the density outside: the correction term is small. Hence for $\sum_{x \in \Lambda} n_x$,

$$\mathbb{P}(n_\Lambda) \approx \frac{e^{|\Lambda^c|s(\rho)} e^{-s'(\rho)k}}{Z}.$$

$e^{|\Lambda^c|s(\rho)}$ is a huge factor but it does not vary with k so we can omit it. Writing out Z as the sum of the numerators over all n_Λ , and letting $\mu = s'(\rho)$, we get that this equals

$$= \frac{e^{-\mu k}}{\sum_{n' \in \Omega_\Lambda} e^{-\mu \sum_{x \in \Lambda} n'_x}},$$

which is the grand canonical ensemble average. So all you have to do is be sure to pick μ as the derivative of the thermodynamic function at the correct density. That is how we derive the canonical ensemble in this case. The rest of the system acts on the small system as a “particle (heat) bath.”

Here we used very explicit machinery, namely, the Stirling formula. We want a general expression that doesn't rely on the Stirling formula because in more complicated models, we will not have the luxury of using the Stirling approximation. We'll do this in the next section.

Then we will be able to consider models where there are more energy constraints. The general procedure is first make a list of energy constraint. There is a generalization of the equivalence principle where

- in the micro-canonical ensemble we average over configurations where the constraints have prescribed values. (We considered the special case where just the number of particles was prescribed.)
- in the grand canonical ensemble, we add an energy term to the exponential: $e^{-\mu N(n) - \beta \mathcal{E}(n)}$ where $N(n)$ is the number of particles, $-\mu N(n)$ is the Gibbs factor, and $\mathcal{E}(n)$ the energy.

How can we construct an alternative method without the Stirling formula?

2-4-16

1.3 Generalizing Ensemble Analysis to Harder Cases

Previously we used very specific machinery to derive the grand canonical ensemble for the lattice gas; how might we derive similar expressions for more complicated systems? We would like to avoid using the Stirling formula, which only happened to work in the simple case. To derive the entropy bypassing Stirling's approximation, you may proceed by analyzing partition functions.

Define the **partition function** for the grand canonical ensemble as the (unnormalized) sum of the likelihoods over all configurations. At each site there are 2 possibilities and they

are independent, so the sum factors into a product.¹

$$\begin{aligned}
 Z_\Lambda &= \sum_{n \in \Omega} e^{-\mu N(n)} \\
 &= \sum_{n \in \Omega} \prod_{x \in \Lambda} e^{-\mu \mathbb{1}[n_x=1]} \\
 &= \prod_{x \in \tilde{\Lambda}} (1 + e^{-\mu}) \\
 &= (1 + e^{-\mu})^{|\Lambda|} \\
 \frac{1}{|\Lambda|} \ln Z_\Lambda &= 1 + e^{-\mu} \tag{1.1}
 \end{aligned}$$

So what does this have to do with entropy? We can calculate Z_Λ a different way, as follows. Note that the only thing that matters in the summand is the number of particles in n , so let's group the summands by this. Letting $\rho = \frac{N}{|\Lambda|}$, and using the fact that the number of states with density ρ is $e^{|\Lambda|s(\rho)}$,²

$$\begin{aligned}
 Z_\Lambda &= \sum_{n \in \Omega} e^{-\mu N(n)} \\
 &\approx \sum_{\rho \in \frac{1}{|\Lambda|}\mathbb{Z}} e^{-\mu N(n)} e^{|\Lambda|s(\rho)} \\
 &= \sum_{\rho \in \frac{1}{|\Lambda|}\mathbb{Z}} e^{|\Lambda|[s(\rho) - \mu\rho]} \\
 \tag{eq:z2} &\approx e^{|\Lambda| \max_{\rho \in [0,1]} [s(\rho) - \mu\rho]}. \tag{1.2}
 \end{aligned}$$

To get the last step, note that the error we are making by focusing on the maximal point rather than counting with multiplicity the near-maximal value is at most a factor equal to the volume, since we can upper bound our error by taking all points in Λ instead of merely an ε -ball around the maximum. "In analysis, it often pays to avoid being excessively generous." Since $\frac{1}{|\Lambda|} \ln |\Lambda| \rightarrow 0$, this term is negligible. The expression $\max_{\rho \in [0,1]} [s(\rho) - \mu\rho]$ is called the **Legendre transform** of the entropy and denoted by $s^*(\mu)$.

Matching (1.1) and (1.2), we find

$$s^*(\mu) := \max_{0 \leq \rho \leq 1} [s(\rho) - \mu\rho] = \ln(1 + e^{-\mu})$$

Using this, we can derive the expression the formula for $s(\rho)$, using the inverse Legendre transform.

¹Note that Λ in the last section is the micro-canonical ensemble, but here it is the grand canonical ensemble.

²Rather than getting this with $s(\rho) = -[\rho \ln(\rho) - (1 - \rho) \ln(1 - \rho)]$ by the Stirling approximation in the last section, we can take this as the *definition* of s here, and use it to solve for s .

To check that taking the maximum in (1.2) was legit, note that³

$$\begin{aligned} \max_{\rho} [s(\rho) - \mu\rho] &\leq \frac{1}{|\Lambda|} \ln Z_{\Lambda} \leq \max_{\rho} [s(\rho) - \mu\rho] + \underbrace{\frac{1}{|\Lambda|} \ln |\Lambda|}_{\rightarrow 0} \\ \implies \lim_{|\Lambda| \rightarrow \infty} \frac{1}{|\Lambda|} \ln Z_{\Lambda} &= \sup_{0 \leq \rho \leq 1} [s(\rho) - \rho\mu]. \end{aligned}$$

We can verify that $s(\rho) = -[\rho \ln \rho + (1 - \rho) \ln(1 - \rho)]$ does indeed make (??) hold. We find the maximum (critical point) of $s(\rho) - \mu\rho$ by setting the derivative to 0.⁴ We use the trick

$$[x(\ln x - 1)]' = \ln x.$$

We can subtract 1 from each of the logs changing the expression by a constant. Thus

$$s'(\rho) = -\ln \rho + \ln(1 - \rho) - \mu \implies \frac{\rho}{1 - \rho} = e^{-\mu}.$$

We can use this to solve for ρ and find $s^*(\mu)$ (do this yourself).

In general, a micro-canonical ensemble specifies all the conserved quantities: Particles, energies, and whatever else there is. The grand canonical ensemble also generalizes by changing the factor to $e^{-\mu N} e^{\beta H}$. Such factors are referred to as **Gibbs factors** or Gibbs measures.

2 Concavity and the Legendre transform

2.1 Basic concavity results

Definition 1.2.1: A function on \mathbb{R}^k is **concave** if for any $x_0, x_1 \in \mathbb{R}^k$, $0 \leq \lambda \leq 1$,

$$F(\lambda x_1 + (1 - \lambda)x_0) \geq \lambda F(x_1) + (1 - \lambda)F(x_0).$$

For a **convex** function, the same inequality with the sign flipped holds. A negative convex function is concave. Concavity (convexity) means if you draw a chord between two points, it will lie below (above) the curve.⁵

For a strictly concave function, a maximum, whenever it exists, is unique.

Theorem 1.2.2. *For any concave function on \mathbb{R} ,*

³Mathematicians are paranoid, so we use sup instead of max. For many practical purposes they are the same.

⁴Don't differentiate in public.

⁵"As Richard Feynmann pointed out, getting the sign right is the hardest thing."

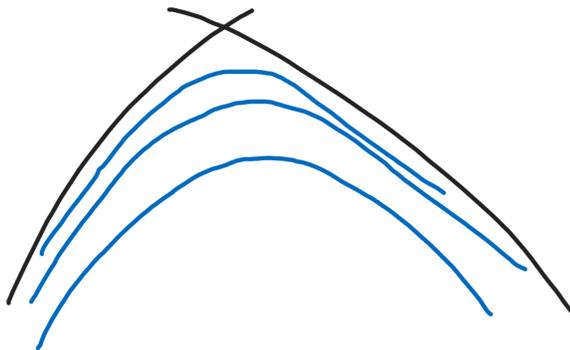
1. The directional derivatives $F'(x \pm 0)$ exist at all $x \in \mathbb{R}$. The **directional derivative** is defined as

$$F'(x+0) = \lim_{\varepsilon \rightarrow 0^+} \frac{F(x+\varepsilon) - F(x)}{\varepsilon}$$

$$F'(x-0) = \lim_{\varepsilon \rightarrow 0^-} \frac{F(x+\varepsilon) - F(x)}{\varepsilon}.$$

2. $F'(x-0) \geq F'(x+0)$ and $F'(x \pm 0)$ are decreasing.
3. For all but countably many values $x \in \mathbb{R}$, $F'(x-0) = F'(x+0)$, i.e., F is differentiable at x .
4. Let F_n be a sequence of concave functions which converge pointwise: for all x , $\lim_{N \rightarrow \infty} F_N(x) =: \tilde{F}(x)$ exists. Then
- (a) \tilde{F} is concave.
- (b) At points of differentiability of \tilde{F} , the derivatives also converge,⁶

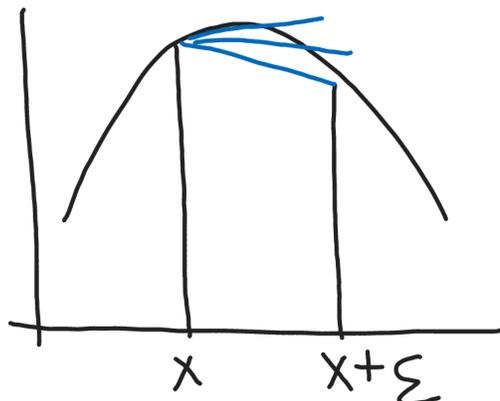
$$F'(x \pm 0) \rightarrow \tilde{F}'(x).$$



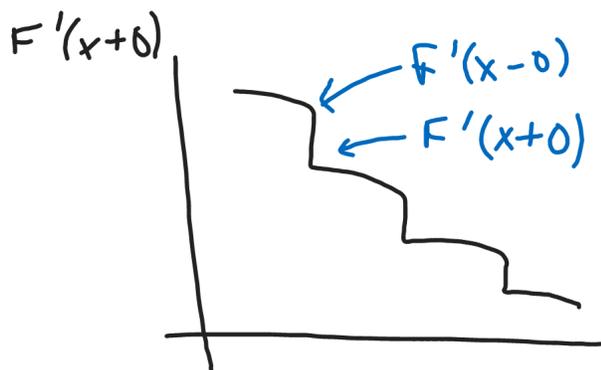
Much of this generalizes to directional derivatives in n dimensions.

Proof. For (1) note that concavity implies that the slope of the line between $x, x + \varepsilon$ is increasing as $\varepsilon \rightarrow 0^+$.

⁶Finite energy functions are always smooth, but their limit can have discontinuous derivative.



To prove (3), consider the graph of the derivative $F'(x+0)$. It's a monotone decreasing function: it starts somewhere and goes down. At points where $F'(x+0)$ is continuous, we have $F'(x+0) = F'(x-0)$. The discontinuities, or “steps,” are the points where $F'(x+0) \neq F'(x-0)$. Now we use the fact that the sum of any uncountable collection of nonzero numbers is ∞ . Applying this to the steps, we find that the number of steps is countable, i.e., the total collection of discontinuity points has to be countable.^{7 8}



□

2.2 Concave properties of the Legendre transform

Definition 1.2.3: The Legendre transform of a function is defined as

$$(TG)(y) = \inf_x [y \cdot x - G(x)] = -\sup_x [G(x) - y \cdot x]$$

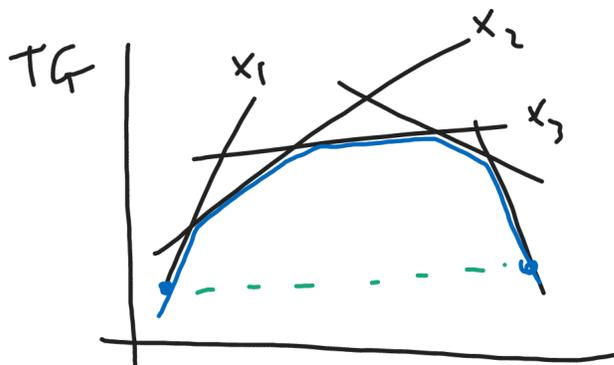
⁷More rigorously, for every nonzero interval $[F'(x+0), F'(x-0)]$, we can associate with it a rational number. Then note \mathbb{Q} is countable.

⁸The set of discontinuities can still be devilishly dense, ex. all the rational numbers. In physics it used to be thought that weird functions with Cantor-like sets of discontinuities could not occur, but there are materials whose free energy discontinuities are dense in certain areas.

Theorem 1.2.4. For any function G , TG is concave.

Proof. An efficient way to think about this is the following: for each value of the parameter x , as a function of y this is a linear function. For each x we get a linear function. Define the transform by taking the infimum over that.

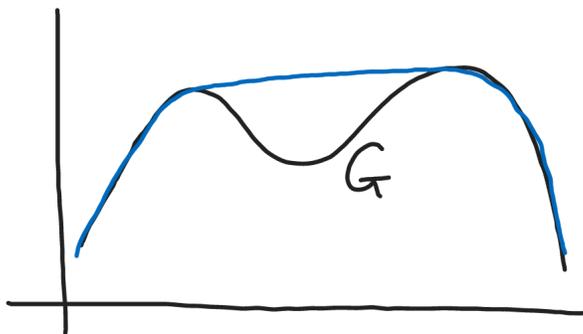
Take 2 points and draw the chord between them. For each linear function the chord lies below it.⁹



□

Definition 1.2.5: The **concave hull** \tilde{G} of G is defined as the smallest concave function that is at least the function value at every point:

$$\tilde{G}(x) = \inf \{F(x) : F \text{ concave}, \forall u, F(u) \geq G(u)\}.$$



Theorem 1.2.6. For concave G ,

$$T(TG) = G.$$

In general, $T(TG)$ is the concave hull of G .

⁹I.e., we use the following : Let \mathcal{F} be a collection of concave (e.g. linear) functions. Then $\inf_{f \in \mathcal{F}} f(x)$ is concave.

Proof for G differentiable. Use the fact that if G is differentiable, then

$$\inf_x [y \cdot x - G(x)]$$

occurs at $y = G'(x)$. □

Note that we plotted the function and the dual function on the same graph. However, they have inverse units, for example, energy and inverse temperature. You get a lot of insights into physics if you keep track of the units.

Recall $\rho = \frac{N}{|\Lambda|}$. If all $2^{|\Lambda|}$ configurations are given equal weight, the typical value or ρ , the particles per unit volume is $\frac{1}{2}$. The Law of Large Numbers says that with probability 1 the ratio tends to $\frac{1}{2}$. Is it possible that the density is $\frac{1}{3}$? Yes, but the probability of such a density is given by the entropy: it's exponentially small, $e^{|\Lambda|[s(\frac{1}{3}) - \ln 2]}$. Anything other than $\frac{1}{2}$ is a large deviation; they occur with exponentially small probability. We want to quantify the probability of large deviation events. Here is a language that people found useful.

Definition 1.2.7: def:ldp A sequence of probability measures on \mathbb{R} is said to satisfy a **large deviation principle** with speed $\{a_N\}$, and rate function $I(x)$ if for each $x \in \mathbb{R}$, $\varepsilon > 0$,

$$-\inf_{|u-x|<\varepsilon} I(u) \leq \liminf_{N \rightarrow \infty} \frac{1}{a_N} \ln \mathbb{P}_N((x-\varepsilon, x+\varepsilon]) \leq \overline{\lim}_{N \rightarrow \infty} \frac{1}{a_N} \ln \mathbb{P}_N((x-\varepsilon, x+\varepsilon]) \leq -\inf_{|x-u| \leq \varepsilon} I(u)$$

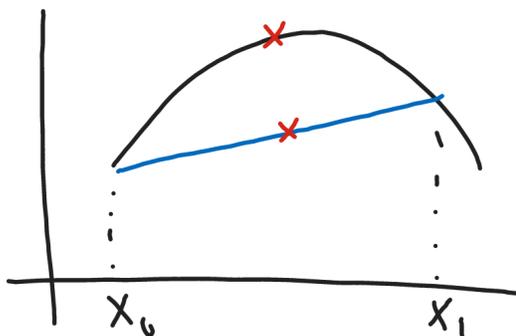
For us, $I(x) = \ln 2 - s(x)$.

2-9-16: Today we'll talk about Gibbs states: definition, variational principle, and relations to thermodynamic foundations.

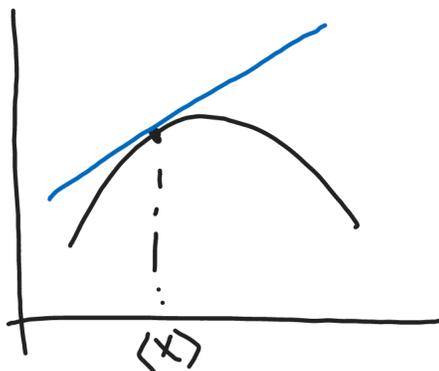
Theorem 1.2.8 (Jensen's inequality). Let $\rho(dx)$ be a probability measure on \mathbb{R} (\mathbb{R}^d) with finite expectation $\int |X| \rho(dx) < \infty$, and let $F : \mathbb{R} \rightarrow \mathbb{R}$ be a concave function. Then

$$\int F(X) \rho(dx) \leq F\left(\int X \rho(dx)\right).$$

To remember this, draw a picture. Consider the case where the measure is concentrated on two points. The interpolated value $\int F$ is less than the function value $F(\int)$. (In the case of two points, this is the definition of Jensen.)



Proof. Let $\langle X \rangle = \int X \rho(dx)$. Take any tangent to F at $\langle X \rangle$. Note that F may not be differentiable, so take the line to have slope $F'(\langle X \rangle + 0)$. The first inequality (1.3) follows from concavity.



Integrate and note that $F(\langle X \rangle)$ is constant to get (1.4).

$$\text{eq:jensen1} \quad F(X) \leq F(\langle X \rangle) + (X - \langle X \rangle)F'(\langle X \rangle + 0) \quad (1.3)$$

$$\text{eq:jensen2} \quad \int F(X) \rho(dx) \leq \underbrace{F(\langle X \rangle)}_{\int \rho(dx)=1} + \underbrace{0}_{\int [X - \langle X \rangle] \rho(dx)=0} \quad (1.4)$$

□

This theorem is elementary but very useful.

3 Basic setup for statistical mechanics

The basic setup consists of...

1. A **lattice or a homogeneous graph** like \mathbb{Z}^d . It will be important that for $\Lambda_L = [-L, L]^d$,

$$\text{eq:bdary-ratio} \quad \frac{|\partial \Lambda_L|}{|\Lambda_L|} \xrightarrow{L \rightarrow \infty} 0. \quad (1.5)$$

Here $|\cdot|$ means the size in terms of number of points (it doesn't matter much how you count—e.g. whether you count just the points on the edge, or adjacent too, etc.). (??) says when you chop space into regions, the boundary plays a small role. This is good because we should be talking about extensive quantities.

2. **Collection of local variables** like $\{n_x\}$ taking values in $0, 1$, or $\{\sigma_x\}$ taking values in ± 1 , magnetizations, etc.

3. An extensive **energy function**, defined on finite subsets. For example,

$$H_\Lambda(\sigma) = - \sum_{(x,y) \subset \Lambda, |x-y| \leq r} \sigma_x \sigma_y - h \sum_{x \in \Lambda} \sigma_x.$$

This is over a finite range; we can also consider unbounded ranges with some decay.

Here there are only pairwise interactions, but more generally, there can be interactions between more variables: 3, 4... Typically the interactions are translation invariant. The general equation is

$$H_\Lambda(\sigma) = \sum_{A \subset \mathbb{Z}^d, \text{diam}(A) \leq R} K_A(\sigma),$$

where $K_A(\sigma)$ depends on $\sigma|_A$ (σ restricted to A), and are shift covariant.

4. A **reference a-priori (probability) measure** $\rho_0(d\sigma)$ (a probability distribution with respect to which we integrate) on the configuration space $\Omega_\Lambda (= \{-1, +1\}^{\mathbb{Z}^d})$.

For example, $\rho_0(d\sigma)$ could be the product measure where $\{\sigma_x\}$ are iid variables. (Think of a system at high temperature.)¹⁰¹¹

We will also allow measures which are not probability measures—normalization may be “part of the game.”

Particle configurations are given by specifying locations and momenta. We can chop space into boxes, and specify the number of particles in each box, and their positions and momenta. The starting point is the Liouville measure, which is invariant under time evolution by the Hamiltonian.

3.1 Gibbs equilibrium measure

Definition 1.3.1: df:gibbs-eq The finite-volume **Gibbs equilibrium measure** at temperature $T = \beta^{-1}$ is

$$\text{eq:gibbs-eq} \text{Prob}(d\sigma) = \frac{e^{-\beta H_\Lambda(\sigma)}}{Z_\Lambda} \rho_0(d\sigma_\Lambda). \quad (1.6)$$

Here Z_Λ is the normalizing (partition) function

$$\text{eq:zla} Z_\Lambda(\beta) = \int e^{-\beta H_\Lambda(\sigma_\Lambda)} \rho_0(d\sigma_A). \quad (1.7)$$

¹⁰Note that it makes sense to talk about e.g. an infinite number of coin flips. There is such a thing as an infinite product of probability measures. In an infinite product space, the result of any finite collection is independent. Any single value has probability 0.

¹¹How does the Declaration of Independence go? “We hold these truths to be self-evident... inalienable rights...” The definition of person is time-dependent. But there is a reference measure, individuals are treated equally. That’s how we start in statistical mechanics. E.g. Every spin configuration gives equal value. If the spins are continuous, what would be a good starting point? Perhaps they are independently distributed on the sphere.

This is a **generating function** because taking the derivatives we can learn about the distribution of the random variable.

The Gibbs equilibrium measure is the uniform measure multiplied by the Gibbs factor. Here, β is a factor corresponding to the inverse of the **temperature**.

- When $\beta = 0$, local variables are independently distributed. We have chaos; all states are equally likely.
- When β cranked up, i.e., temperature is lowered, the probability distribution becomes more concentrated (near the “ground state”). When $\beta = \infty$, the distribution becomes concentrated on configurations which minimize the energy.

We will see that the Gibbs equilibrium measure is the distribution of a system at thermal equilibrium at temperature T . Why is that so and what is the relation to microcanonical ensemble?

We will try to understand the structure of these measures and the phase transition they manifest.

3.2 Introduction to the Ising model

In biology, they study *Drosophila*, the fruit fly. Studying this simple organism tells us a lot. The Ising model is the *Drosophila* of statistical physics. What you learn from it extrapolates to many other systems, but not everything.

Consider the **Ising model** on \mathbb{Z}^d ,

$$H_\Lambda(\sigma) = -\frac{1}{2} \sum_{|x-y|=1} J\sigma_x\sigma_y - h \sum_{x \in \Lambda} \sigma_x, \quad J \geq 0.$$

At $T = 0, \beta = \infty$, the state with all +’s and the state with all –’s are equally likely. There is no continuous way to go from all spins + to all spins –. For nonzero, analytic even after take infinite limit.

However, the analyticity would fail, and in fact you would find a line of first order transitions at $H = 0$ up to some temperature T_c . What happens is a natural extension at zero temperature and infinite β . You go from configuration that is all + to all – going through a discontinuity. The Gibbs state—the trivial distribution at this temperature—comes in 2 flavors (at least, other possibilities can hold): The state would remember whether the magnetic field H was turned to 0 from the positive or negative side. This is a beautiful example of what you observe in magnets.

The floor under the Atlantic ocean has ferromagnetic rocks. It was detected that the direction of magnetization changes. When you cool a ferromagnet (ferromagnets develops magnetic moments), which way it points is affected by the prevalent external field. According to prevailing wisdom, as the sea floor was expanding, the Earth’s magnetic moment flipped. The rocks have encoded in them the direction of magnetization when they were cooled past this critical temperature. This phenomenon is called residual magnetization.

This phenomenon is eliminated when you raise to high enough temperature; then the state becomes the analytic. We will discuss more about this phase transition, including behavior near the critical point, later.

This refers to the infinite volume limit of such measures. In finite volume, everything is analytic. At 0 temperature the configuration is all +, but at small temperature, thermal fluctuations occur. Unlike when $\beta = \infty$, when β is finite, every local configuration gets some nonzero weight. Even though there is a preference for agreement, the system exhibits fluctuations. Among the +'s there will be islands of minority spins. As temperature increase, minority fluctuations increase to a point where each spin tries to agree both with its neighbors and with magnetic field. When there is a lot of fluctuation among its neighbors, the effect of magnetic field is not so significant. That's how this discontinuity evaporates when you increase temperature.

3.3 Entropy, energy, and free energy

Definition 1.3.2: Let ρ_0 be a reference measure.

1. For any probability measure μ on Ω_Λ ¹², we can write it in terms of ρ_0 ,

$$\mu(d\sigma_\Lambda) = g(\sigma)\rho_0(d\sigma_\Lambda).$$

The function $g(\sigma)$, the “ratio” of the measures, is called the **Radon-Nikodym derivative**, and denoted by

$$g(\sigma) = \frac{\delta\mu}{\delta\rho_0}.$$

2. Define the **entropy** of μ as

$$S_\Lambda(\mu|\rho_0) = - \int g(\sigma_\Lambda) \ln g(\sigma_\Lambda) \rho_0(d\sigma_\Lambda).$$

3. Define the **energy content** of μ as

$$E(\mu) = \int H_\Lambda(\sigma_\Lambda) \underbrace{g(\sigma_\Lambda)\rho_0(d\sigma_\Lambda)}_{\mu(d\sigma_\Lambda)}.$$

Theorem 1.3.3. *thm:gibbs-eq* For each $\beta \in [0, \infty)$, the Gibbs equilibrium measure (1.6) is the unique minimizer of

$$\beta F_\Lambda(\mu) := \beta E_\Lambda(\mu) - S_\Lambda(\mu|\rho_0),$$

equivalently, the unique maximizer of $S_\Lambda(\mu|\rho_0) - \beta E_\Lambda(\mu)$.

Definition 1.3.4: The quantity $\beta F_\Lambda(\mu)$ is called the **free energy**.

¹²that is absolutely continuous with respect to ρ_0

We'll keep β positive, but we can make sense of some of the theory when β is negative. Sometimes we can even talk about β complex!

The second law of thermodynamics says that nature “maximizes entropy under constraints.” Here we're maximizing entropy minus energy. How to reconcile this? There is some reservoir where the system trades with the entropy. As it trades off energy with the reservoir, the energy of the reservoir is affected. It is the effect of the entropy of the reservoir that the system has energy E . What the system is really maximizing is the entropy of the *system and reservoir*.

The free energy $\beta F_\Lambda(\mu)$ is the energy you get from the system when in thermal contact.¹³

What is the consequence? The thermal states are the states which maximize the difference. This allows to quantify the difference. When β is small, the energy plays a minor role. The states maximize the entropy. As the temperature decreases and β increases, $\beta E_\Lambda(\mu)$ in the variational principle plays an increasing role; the energy has to be low; β controls how low it has to be.

What's coming next? We'll pay attention to the factor $\ln Z_\Lambda(\beta)$. (Recall $Z_\Lambda(\beta)$ was defined in (1.10).) We will calculate

$$\lim_{\Lambda \nearrow \mathbb{Z}^d} \frac{1}{|\Lambda_L|} \ln Z_\Lambda(\beta).$$

The relevant contribution comes only from configurations which are at an energy where the entropy is maximized.

Entropy appears in many guises. We talked about entropy over a measure, but you can also talk about the entropy function of energy; the contribution from those particles would be $e^{S(E)}e^{-\beta E}$. Now it's just a function of the energy. This function picks up the Legendre transform of the entropy.

In the Ising model, there is a discontinuity in the nature of the states. Differentiability of the free energy function breaks down across the line. There is an interplay between convexity properties of thermodynamic functions. The left and right derivative—the mean values of the magnetization—corresponds to 2 different limiting values, that are the two different magnetizations depending on which way you arrive.

Statistical mechanics translates into the nonuniqueness of Gibbs equilibrium state. Statistical mechanics provides much more information because it looks at the joint distribution of all these variables, whereas thermodynamics just fixes attention on a few relevant parameters. We will clarify this relation. We will clarify this and then discuss techniques to find the phase transitions.

2-11-16: I started a couple of lines of discussion involving entropy; I would like to tie up a few of those loose ends. Just to remind you, please do not hesitate to ask questions about notation and the like; we have a very mixed audience and it's good to be reminded of elementary questions.

¹³It's funny to talk about the energy crisis: we never run out of energy. The problem is that we have too little *free* energy, we have too much entropy.

Remark 1.3.5: Let Ω_0 be the possible states (e.g., $\{-1, +1\}$ for spin states) for a single particle, and $\Omega = \Omega_0^G$ where G is the lattice graph. Any $\sigma \in \Omega$ is in the form $\sigma = \{\sigma_x\}_{x \in G}$, where $\sigma_x \in \Omega_0$. In the discrete case (Ω_0 and G are finite), the probability of configuration σ is $\rho(\sigma) = \rho(\{\sigma\})$. The expectation in the discrete case is a sum,

$$\mathbb{E}_\rho(f) = \sum_{\sigma \in \Omega} f(\sigma)\rho(\sigma).$$

In general, the expectation is an integral,

$$\int f(\sigma)\rho(d\sigma).$$

Note that the sum is just integration respect to a discrete measure. The integral notation is more flexible because it adapts to non-discrete cases, for instance, when $\Omega_0 = S$ takes values on the sphere.

We talked about relative entropy, and entropy entered the situation different times in different ways. Entropy appears in many different areas, and they are all either related by precise relations or analogy.

Definition 1.3.6: Let ρ, μ be probability measures such that μ is absolutely continuous with respect to ρ . Let $\mu = G(\sigma)d\sigma$, i.e., $G = \frac{\delta\mu}{\delta\rho}$ is the Radon-Nikodym derivative with respect to ρ . The **relative entropy** is

$$S(\mu|\rho) = - \int G(\sigma) \ln(G(\sigma))\rho(d\sigma) = - \int \ln G(\sigma)\mu(d\sigma)$$

Before we talk about the Gibbs measure and the variational principle, let us give a useful lemma.

Lemma 1.3.7. *For any pair of probability measures,*

$$S(\mu|\rho) \leq 0$$

with equality iff $\mu = \rho$ (as measures).

Proof. Take $\psi(x) = -x \ln(x)$. Look at the tangent line $(1 - x)$ intersecting the x -axis (see photo) and see that the function is less:

$$\begin{aligned} -x \ln(x) &\leq 1 - x && (1.8) \\ S(\mu|\rho) &= \int \psi\left(\frac{\delta\mu}{\delta\rho}\right) \rho(d\sigma) \\ &= \int \left[\psi\left(\frac{\delta\mu}{\delta\rho}\right) - \left(1 - \frac{\delta\mu}{\delta\rho}\right) \right] \rho(d\sigma) \\ &\leq 0 \end{aligned}$$

where we used the fact that the derivative of ρ against the Radon-Nikodym derivative with respect to ρ is just 1,

$$\int \left(\frac{\delta \mu}{\delta \rho} \right) \rho(d\sigma) = \int \mu(d\sigma) = 1,$$

so that $1 - \frac{\delta \mu}{\delta \rho}$ integrates to 0. We added this term order to use (1.8). \square

Hence, this is a variational principle for entropy.

Remark 1.3.8: For cases where things are not absolutely continuous, we define the relative entropy to be $-\infty$, since things will blow up.

Now, we derive a useful relation for $\rho_\beta(d\sigma) = \frac{e^{-\beta H(\sigma)}}{Z} \rho_0(d\sigma)$:

$$\text{eq:2-11-2} S(\mu|\rho_0) - \beta \int H(\sigma) \mu d(\sigma) = \ln(Z) + S(\mu|\rho_\beta) \quad (1.9)$$

With a bit of political license, you may refer to this as the amount of free energy: It's β times the entropy.

Proof. We have

$$\left(S(\mu|\rho_0) - \beta \int H(\sigma) \mu d(\sigma) \right) - \ln(Z) = - \int \ln \left(\frac{\delta \mu}{\delta \rho} \cdot \frac{Z}{e^{-\beta H(s)}} \right) \mu(d\sigma)$$

Now we must show this is the same thing as $S(\mu|\rho_\beta)$. We can look at it slightly differently, as though we are modifying the Radon-Nikodym derivative with respect to the modified measure. The above equals

$$= - \int \ln \left(\frac{\delta \mu}{\delta \rho_\beta} \right) \mu d\sigma = S(\mu|\rho_\beta)$$

by definition. \square

This is rather cool: It says a thermodynamic flavored quantity is equal to the relative entropy with respect to the Gibbs-measure plus a positive constant.

So with this we have proof of what we said last time.

Theorem (Theorem 1.3.3). *The (finite volume) **Gibbs equilibrium** measure $\rho_\beta(d\sigma)$ is the unique maximum of*

$$-\beta F(\mu) := S(\mu|\rho_0) - \beta \int H(\sigma) \mu(d\sigma)$$

Proof of Theorem 1.3.3. By (1.9), this equals $\ln Z + S(\mu|\rho_\beta)$. $S(\mu|\rho_\beta)$ is maximized for $\mu = \rho_\beta$. \square

The point was to show that this measure would be dominated by the Gibbs state. It's useful to know when things are maximum.

This kind of correction factor is typical in thermodynamics. It appears in dynamics with some kind of conservation law, where the energy changes only due to interaction with a heat path or some other large system at some fixed temperature which can exchange energy. Then the fluctuations of the system energy are affected by fluctuations in the reservoir. This second term corresponds to the state of the universe if your state is at state μ . The fact that the Gibbs state is where this is maximal is a reflection of the second law of thermodynamics. The fact is that the measure ρ_0 is not just any measure; it's a natural notion of an *a priori* measure which is what the system typically does.

You may want to think of the statistical mechanics of a regular system like a lattice. We can specify the dynamics by specifying the number of particles in each box, and let everything flow freely. What would be a natural measure ρ_0 ? We want a probability measure of the system that is invariant with respect to state. You could just take the product $\frac{1}{N!} \prod dq_i \cdots dq_j$: This is the classical measure which is stationary under any Hamiltonian evolution regardless of the state function. Of course, people working in probability theory are aware of the Bayesian approach to probability, where some notion of what a priori measure is is fundamental.

4 Large deviation theory

We've touched on **large deviation theory** (Definition 1.2.7). For instance, consider spin states:

$$\sigma_x = \begin{cases} +1, & w.p. \frac{1}{2} \\ -1 & w.p. \frac{1}{2}. \end{cases}$$

Let's take a space Λ with some finite volume. The empirical average spin is $\frac{1}{|\Lambda|} \sum_{x \in \Lambda} \sigma_x$. The Law of Large Numbers says

$$\mathbb{P} \left(\left| \frac{1}{|\Lambda|} \sum_{x \in \Lambda} \sigma_x - \langle \sigma \rangle \right| > \varepsilon \right) \rightarrow 0$$

as $|\Lambda| \rightarrow \infty$. Significant deviation from the mean $\langle \sigma \rangle$ tends to zero. So we might ask now: what is the probability that the empirical average is close to any other value m is close to 0?

$$\mathbb{P} \left(\left| \frac{1}{|\Lambda|} \sum_{x \in \Lambda} \sigma_x - m \right| \leq \varepsilon \right) \approx e^{-a_n I(m)}$$

In this situation, $a_n = |\Lambda_n|$, where we are taking a sequence of increasing volumes. I am a bit ambiguous here, since for our purposes, these constants will be proportional to the volume. A book on large deviation theory would formulate things in a more general way. It goes without saying the volume is very large in the limit we are discussing, so this is a very tiny probability. What we mean by approximate is also a little ambiguous here: it means

when you take a log of the RHS, you'll get the exponent in the probability. For us, the $I(m)$ is nothing more than the entropy corrected by a constant. Again, we use the I notation to agree with large deviation books.

You never ask for a precise value of m since the probability is typically 0, unless it's a multiple of $|\Lambda|$, instead you have some tolerance level ε . Now how to prove this? In the first lecture we used the Stirling formula, and derived it by hand. But I would now like to present a method to conclude this kind of result. These variables are not independent variables, they might be correlated like spin states.

The validity of a **large-deviation principle** can often be deduced using the following theorem:

Theorem 1.4.1 (Gärtner-Ellis). *Assume that for a sequence of random variables E , the limit*

$$Q(\lambda) = \lim_{n \rightarrow \infty} \frac{1}{a_n} \ln \mathbb{E}(e^{\lambda E})$$

exists and is finite for all λ (here, a_n is the volume). We will also use more generally X as a vector of random variables, in which case define

$$Q(\lambda) = \lim_{n \rightarrow \infty} \frac{1}{a_n} \ln \mathbb{E}(e^{\lambda \cdot X}).$$

Then for any closed set F and open set G :

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{a_n} \ln \mathbb{P}(E \in F) &\leq - \inf_{x \in F} Q^*(x) \\ \lim_{n \rightarrow \infty} \frac{1}{a_n} \ln \mathbb{P}(E \in G) &\geq - \inf_{x \in G'} Q^*(x) \end{aligned}$$

with $Q^(x) = \sup_{\lambda} [\lambda x - Q(\lambda)]$ and G' is the set of **exposed points** of Q^* .*

We would like to consider systems of large volume. Let the number of points be $a_n = |\Lambda_n|$. Let E_n be the total energy of configuration σ in the volume Λ_n ,

$$E_n = H_{\Lambda_n}(\sigma).$$

Alternatively, we can take the total energy in the box and break it down as the volume times the energy per volume, $E_n = a_n \cdot x$, where $x = \frac{E_n}{a_n}$ is the energy density. We will calculate the mean value of the total energy in the n^{th} box, and estimate how large it should be should a large deviation principle apply.

$$\mathbb{E}(e^{\lambda E_n}) = \int e^{\lambda x a_n} \mathbb{P}(E_n \in a_n dx)$$

We want to reflect the fact that this probability is ridiculously small, as we saw before in large deviation principle: $\mathbb{P}(E_n \in a_n dx) = e^{-a_n I(x)}$. Combining these factors, you get the following integral:

$$\int e^{a_n[\lambda x - I(x)]} a_n dx$$

In effect we are integrating x over the exponential of a volume times some reasonable quantity. So when you take the logarithm of that and divide by the volume,

$$\frac{1}{a_n} \ln(\mathbb{E}(e^{\lambda E_n})) \approx \max_x [\lambda x - I(x)].$$

Seeing this, you say, “Ah, I know what this is—this is the Legendre transform of the rate function!”. What you’re really doing is taking the Legendre transform which is sensitive to the convex hull of the rate function.

Suppose there is a certain interval where there are few configurations in a fallen entropy. What would contribute to the expected value? In that situation, there will be competition between situations: The volume may decompose into situations where you have a small energy density for part of the volume and large energy density for the rest of the volume. So it averages to something larger: Thus the expected value of the partition function gives you information about the Legendre transform. This theorem tells you that if you want to learn about the function, you need to do the inverse Legendre transform. The Legendre transform has the property that you can recover the points of $I(x)$ on the convex hull of $I(x)$.

Now what are “exposed points of $Q^*(x)$ ”? Exposed points are points on the intersection of the convex hull and the original function.

We can learn about the thermodynamics of the system $\frac{1}{|\Lambda_n|} \ln(Z(\beta))$. Notions of convexity carry very nicely from the real line to general affine spaces. In that case, what you want to do for spin systems like this, is to ask more detailed questions, like “What is the probability that the energy for unit volume is in some area dE , and the sum of the magnetization is in some interval dm ?”. Such things are typically governed by large deviation principles. In order to answer these kinds of questions, you can just take the partition function as $Z_n = \mathbb{E}(e^{-\lambda \cdot X})$ where X is the quantities you are interested in, and then you can just apply the multidimensional version of the theorem we just stated.

[I would also post some homework about using the calculation we derived here instead of the Stirling inequality.](#)

2-16-16

5 Free energy

We now discuss free energy.

Let us first start with a finite system. Consider a finite subset Λ of the grid \mathbb{Z}^d . In particular, consider $\Lambda_L = [-L, L]^d$, a cube of size $2L$ in each dimension. The configuration space is Ω_Λ , which is the space of $\sigma = \{\sigma_x\}_{x \in \Lambda}$ with

- $\sigma_x \in \{\pm 1\}$ for the Ising model, or
- $\sigma_x \in \{1, \dots, Q\}$ for the Potts model.

There are other possibilities as well.

We have the a priori probability measure $\rho_0(d\sigma) = \otimes_x \rho(d\sigma_x)$.

First we define free energy. Recall the partition function

$$\text{eq:zla} Z_\Lambda = \int_{\Omega_\Lambda} e^{-\beta H_\Lambda(\sigma_\Lambda)} \rho_0(d\sigma_\Lambda) \quad (1.10)$$

where the integral is over the spins in the cube with respect to the product measure, and

$$H_\Lambda(\sigma_\Lambda) = \sum_{A \subset \Lambda, \text{diam} A \leq R} J_A \Phi_A(\sigma_A)$$

where Φ is a translation invariant function and J_A is a coupling constant.

For instance, in the Ising model,

$$\text{eq:ising-h} H_\Lambda = - \sum_{x,y,|x-y|=1} J \sigma_x \sigma_y - h \sum_x \sigma_x \quad (1.11)$$

(we sum over neighbors). This is an example of a formula in the type given above. Here J is a coupling constant, and for single sites, the translation invariant coupling constant is h .

Definition 1.5.1 (Free energy): The free energy is defined as

$$F(\beta, \Phi) = -\frac{1}{\beta} \ln(Z_\Lambda)$$

where Z_Λ is defined in (1.10).

Here Φ can in general be complicated and $\beta = \frac{1}{k_B T}$, where k_B is the Boltzman constant (in the following, choose units so $k_B = 1$). For simplicity, given a specific model, we could just write (J, h) instead of Φ .

5.1 Basic Properties

When we formulate the following models, we assume that $\max_{\sigma_A} |\phi_A(\sigma_A)| < \infty$.

1. Assume the system obeys a large deviation principle with an entropy function $S(E)$ where E denotes the energy per unit volume. Using $E = \frac{H_\Lambda(\sigma_\Lambda)}{|\Lambda|}$, we calculate

$$Z_\Lambda = \int e^{-\beta H_\Lambda(\sigma)} \rho_0(d\sigma) \approx \int e^{-\beta |\Lambda| \cdot E} e^{|\Lambda| s(E)} \text{“} dE \text{”}$$

where dE is kind of notation abuse: The number of values the energy can take is essentially integer values over volume units, and this gives you essentially an integral. The integral is dominated by the energy that gives the maximal exponent, and the fact that the integral is an approximation is not too important.

$$\int e^{|\Lambda| \max_E (s(E) - \beta E)} \text{“} dE \text{”}$$

Then we can write the free energy as ($T = \frac{1}{\beta}$)

$$F(\beta) = -\frac{1}{\beta} \max_E (s(E) - \beta E) = \inf_E \{E - Ts(E)\}$$

This is why it's called the free energy: It's the energy term corrected by $Ts(E)$ and is related to the Legendre transform.¹⁴ The free energy is as much energy you can extract from the system when it's in contact with a thermal reservoir, and this is limited.

Recalling that energy is conserved, if you have a finite but huge system and let the system evolve under whatever dynamics it has, it is reasonable to assume that all states are equally likely (as in the microcanonical ensemble). If you look at a sub-system, however, then energy fluctuates. The bulk of the system serves as a reservoir for sub-systems.

At what temperature is this reservoir? Through the large deviation principle for entropy, you are led to realize that even if in this entire universe the energy is fixed, for subsystems, the energy is not constrained, and the distribution within the subsystem is given by $e^{-\beta H}$, suitably normalized.

Hence, systems serve as the reservoir for the subsystems.

2. The free energy function is a generating function. The Gibbs state average energy per volume is

$$\frac{\langle H \rangle}{|\Lambda|} = \frac{1}{|\Lambda|} \int H_\Lambda(\sigma_\Lambda) \frac{e^{-\beta H_\Lambda(\sigma_\Lambda)} \rho_0(d\sigma_\Lambda)}{Z_\Lambda} = \frac{\partial}{\partial \beta} [\beta F(\beta)]$$

We have $\beta F(\beta) = -\frac{1}{|\Lambda|} \ln(Z_\Lambda)$.

Suppose you want to know the total value of the Ising model magnetization, $\langle \frac{1}{|\Lambda|} \sum_{x \in \Lambda} \sigma_x \rangle$. Here, H_Λ is given by (1.11). To find $\langle \frac{1}{|\Lambda|} \sum_{x \in \Lambda} \sigma_x \rangle$, differentiate with respect to h ,

$$\frac{\partial}{\partial h} \ln(Z_\Lambda) = \frac{\beta}{|\Lambda|} \int \left(\sum_{x \in \Lambda} \sigma_x \right) \frac{e^{-\beta H_\Lambda(\sigma)} \rho_0(d\sigma_\Lambda)}{Z}$$

In general, derivatives of the log of the partition function generate the averages you desire. For instance, if you are somewhat sensitive to sums over triangles, then differentiating the free energy with respect to this parameter would give you the average value of that.

3. Variance of H_Λ : To find the variance of H_Λ , differentiate with respect to β again. We have

$$\text{eq:var-h} \quad -\frac{\partial^2}{\partial \beta^2} \beta F(\beta) = -\frac{\partial}{\partial \beta} \int \frac{H_\Lambda}{|\Lambda|} \frac{e^{-\beta H_\Lambda}}{Z_\Lambda(\beta)} \rho_0(d\sigma) = \left\langle \frac{H^2}{|\Lambda|} \right\rangle - \frac{\langle H \rangle^2}{|\Lambda|} = \frac{\langle (H - \langle H \rangle)^2 \rangle}{|\Lambda|}, \quad (1.12)$$

¹⁴I would love it if energy was written in terms of entropy. Somehow, mankind discovered “energy” before “entropy”, so we are continuously having to change notation between energy and entropy.

which is the fluctuation, or variance, in the total energy.

What is the difference between the bulk (average) energy and the minimum energy? Assuming the free energy is twice differentiable, how big would the fluctuation of the total energy minus its mean be? It's going to be big (we're talking about energy in a big universe).

(1.12) is of order 1, so $\langle (H - \langle H \rangle)^2 \rangle$ is on the order of $\sqrt{|\Lambda|}$ which is reminiscent of independence of random variables. The Central Limit Theorem says that if you have a sum of random variables, the fluctuation is $O(\sqrt{|\Lambda|})$. We can prove an analogue of the CLT if the function is twice differentiable, but this requires more work.

Note $F(\beta)$ is convex in β since it's a supremum of linear functions in β (?), and we just proved that $\beta F(\beta)$ is concave. Concavity immediately implies differentiability with the possible exception of some countable set of points¹⁵. In infinite systems, there is more to be said, but for most sets, the second derivative is finite. Hence for Lebesgue almost every value of β , $\beta F(\beta)$ is differentiable in β and has a finite and bounded second derivative.

Our next goal is the following two theorems.

Theorem 1.5.2. *For any translation invariant system as described previously, the following limit exists:*

$$\lim_{L \rightarrow \infty} \frac{1}{|\Lambda_L|} \ln(Z_{\Lambda}(\beta)) = -\frac{1}{\beta} F(\beta)$$

The limiting function is concave in β because the limit of concave functions is concave. The theorem says that the finite volume free energies converge (the limit is infinite dimensional).

If the system is not translation invariant, the limit need not exist: if you set coupling constant to new values as you go along, in that case, there's no consistency between different scales and the limit will not exist. However, translation invariance is a very rigid statement: "Whatever will be". This general principle can be relaxed a bit, and it requires that the system is stochastically invariant: It looks similar at different places.

Theorem 1.5.3. *If the infinite volume free energy F is differentiable at β ,*

1.

$$\langle E \rangle_{L,\beta} = \left\langle \frac{H_{\Lambda_L}}{|\Lambda_L|} \right\rangle_{\Lambda_L,\beta} \rightarrow \frac{\partial}{\partial \beta} [\beta F(\beta)]$$

2. *For any $\epsilon > 0$, then*

$$\mathbb{P}_{\Lambda_L;\beta} \left(\left| \frac{1}{|\Lambda_L|} H_{\Lambda_L}(\sigma_{\Lambda_L}) - \langle E \rangle_{\Lambda_L,\beta} \right| \geq \epsilon \right) \xrightarrow{L \rightarrow \infty} 0$$

the probability is with respect to the Gibbs measure $\rho_{\Lambda,\beta}$.

¹⁵which could even be dense, this set of points can be of extreme interest: they are first order phase transitions

To explain part 1, we know the energy density in the finite volume is given by the finite volume free energy. But what do the derivatives of finite volume functions have to do with derivatives of limiting functions? There are functions which converge pointwise but whose derivatives do not! So the first statement gives us a nontrivial fact. However, derivatives of convex functions converge pointwise as well if a function converges pointwise, so we get a bit extra.

Now, not only is the mean energy given in the limit, but we get part 2 above. Part 2 says that if we consider a growing sequence of finite boxes of size L , the empirical average of the energy per volume is predictable.

In case of independent random variables, the weak law of large numbers says exactly this. But we're in the domain of **correlated** systems, and it is still true. All you need to know is that the free energy is differentiable at a point.

Next we will prove this theorem and the full spelling out of this argument will be left to you (Hint: Use Chebyshev bounds).

2-18

Good references include David Ruelle and Friedli-Velenik (see the introduction).

What is a thermodynamic limit? We would like to discuss systems of large volumes.

Definition 1.5.4: Let $\Lambda(a) = \{x \in \mathbb{R}^d : 0 \leq x_j \leq a\}$. (When we talk about the lattice, replace \mathbb{R}^d with \mathbb{Z}^d .) Let $\Lambda_n(a) = \Lambda(a) + na$ where $n \in \mathbb{Z}^d$. (The shifts tile space by boxes of size a .)

For any $\Lambda \subset \mathbb{R}^d$, let

$$N_a^+(\Lambda) = \left| \{n \in \mathbb{Z}^d : \Lambda_n(a) \cap \Lambda \neq \emptyset\} \right|$$

$$N_a^-(\Lambda) = \left| \{n \in \mathbb{Z}^d : \Lambda_n(a) \subseteq \Lambda\} \right|.$$

They are the number of cells needed to cover Λ and the number of cells completely inside Λ , respectively.

Definition 1.5.5: A sequence Λ_k converges to \mathbb{Z}^d in the **van Hove** sense if for all $0 < a < \infty$,¹⁶

1. $\Lambda_k \rightarrow \mathbb{Z}^d$,
2. $N_a^-(\Lambda_k) \rightarrow \infty$,
3. $\frac{N_a^+(\Lambda_k)}{N_a^-(\Lambda_k)} \rightarrow 1$ as $k \rightarrow \infty$. (This is equivalent to the surface-to-volume ratio $\frac{|\partial\Lambda_k|}{|\Lambda_k|} \rightarrow 0$.)

An example of a sequence violating (3) is a sequence of boxes with many ‘‘arms.’’ The arms have volume proportional to the whole volume; in the arms, the distance to the boundary is $O(1)$. We don't want shapes whose boundaries are on the order of the volume.

¹⁶Careful: we use subscripts here in a different sense than in the previous definition.

It is important that the conditions hold for all a . We will compute the free energy in regular cubes, and see that whatever we prove for regular cubes is valid for a sequence converging in the van Hove sense.

Define the partition function in a box Λ as

$$Z_\Lambda^\#(\beta, h) = \int_{\Omega(\Lambda)} e^{-\beta H_\Lambda^\#(\sigma_\Lambda)} \rho_0(d\sigma_\Lambda)$$

where the $\#$ means we take into account the boundary conditions,

$$H_\Lambda^\#(\sigma_\Lambda) = \sum_{A \subseteq \Lambda} \phi_A(\sigma_A) + \sum_{\substack{B \cap \partial\Lambda \neq \emptyset \\ \text{periodic terms}}} \phi_B^\#(\sigma_A)$$

Here,

- we allow finite-range interactions at the boundary.
- we also allow periodic (wrap-around) boundary conditions, e.g. a term to depend on a pair of one point on the far left and a point on the far right.

We define the **pressure** in Λ as

$$\psi_\Lambda(\beta, h, \dots) = \frac{1}{|\Lambda|} \ln Z_\Lambda$$

(last time we defined this as $-\beta F$).

Theorem 1.5.6. *For any Hamiltonian with a finite range translation invariant interaction,*

$$\lim_{k \rightarrow \infty} \psi_{\Lambda_k}(\beta, h) = \psi(\beta, h)$$

exists for any van Hove sequence $\Lambda_k \rightarrow \mathbb{Z}^d$ and is independent of the boundary conditions.

Proof. Consider first $\Lambda_k = \Lambda(ka)$ where

$$H_{\Lambda(ka)} = \sum_{n, 0 \leq n_j < k} H_{\Lambda_n(a)} + R_k$$

where the sum gives the interactions within boxes, and the second term gives interactions at the boundary of the boxes. We bound the total effect of the R_k (the boundary corridors) to estimate $Z(ka)$. Here r is the radius of interactions (the width of the corridors),

$$\begin{aligned} \|R_k\|_\infty &\leq C_r |\partial\Lambda(a)| k^d \\ Z_{\Lambda(ka)} &= \int e^{-\beta H_{\Lambda(ka)}} \rho(d\sigma_{\Lambda(ka)}) \\ (Z_{\Lambda(a)})^{k^n} e^{-C|\partial\Lambda(a)|h^d} &\leq Z(ka) \leq Z_{\Lambda(a)}^{k^n} e^{C|\partial\Lambda(a)|k^d} \\ \frac{1}{\Lambda(ka)} \ln Z(ka) &= \frac{1}{|\Lambda(a)|} \frac{1}{k^d} \ln Z(ka) \\ \frac{\ln Z(a)}{\Lambda(a)} - C_r \frac{|\partial\Lambda(a)|}{|\Lambda(a)|} &\leq \frac{1}{\Lambda(ka)} \ln Z(ka) \leq \frac{\ln Z(a)}{\Lambda(a)} + C_r \frac{|\partial\Lambda(a)|}{|\Lambda(a)|} \end{aligned}$$

($C_r \propto r$, but the exact dependence is not so important.) The log of the main term is proportional to the number of boxes. Dividing by the number of boxes, we are left with the following estimate.

$$\left| \frac{1}{|\Lambda(ka)|} \ln Z(ka) - \frac{1}{|\Lambda(a)|} \ln Z(a) \right| \leq C_r \frac{|\partial\Lambda(a)|}{|\Lambda(a)|}.$$

(The partition function in the larger box, up to small error, is equal to the partition function in the smaller box.) For any $\varepsilon > 0$, choose a such that $C_r \frac{|\partial\Lambda(a)|}{|\Lambda(a)|} \leq \varepsilon$. Then for all $k_1, k_2 \in \mathbb{N}$,

$$\left| \frac{\ln Z(k_1, a)}{|\Lambda(k_1, a)|} - \frac{\ln Z(k_2, a)}{|\Lambda(k_2, a)|} \right| \leq 2\varepsilon.$$

This says that $\frac{\ln Z(k_1, a)}{|\Lambda(k_1, a)|}$ forms a Cauchy sequence for $k \nearrow \infty$.

Claim 1.5.7. For all a ,

$$\begin{aligned} \limsup_{m \rightarrow \infty} \frac{1}{|\Lambda(m)|} \ln Z(\Lambda(m)) &\leq \frac{\ln \Lambda(a)}{\Lambda(n)} + C_r \frac{|\partial\Lambda(a)|}{|\Lambda(a)|} \\ \limsup_{m \rightarrow \infty} \frac{1}{|\Lambda(m)|} \ln Z(\Lambda(m)) &\geq \frac{\ln \Lambda(a)}{\Lambda(n)} - C_r \frac{|\partial\Lambda(a)|}{|\Lambda(a)|} \end{aligned}$$

Using tiling of boxes of length a , for large cubes you the free energy is approximately what you get from the boxes of length a . The lim sup and lim inf are close up to small error.

$$\psi_{\Lambda(a)} - C_r \frac{|\partial\Lambda(a)|}{|\Lambda(a)|} \leq \liminf_{m \rightarrow \infty} \psi_{\Lambda(m)} \leq \limsup_{m \rightarrow \infty} \psi_{\Lambda(m)} \leq \psi_{\Lambda(a)} + C_r \frac{|\partial\Lambda(a)|}{|\Lambda(a)|}$$

The interval is of width $C_r \frac{|\partial\Lambda(a)|}{|\Lambda(a)|} \rightarrow 0$. We have to take the 2 limits in the right order:

$$\lim_{a \rightarrow \infty} \lim_{m \rightarrow \infty} .$$

□

It's of interest to extend this theorem (of fundamental importance in physics, free energy exists independent of boundary conditions) to (decaying) long-range interactions; what is the cutoff at which the argument breaks? The analogy is that analysts first prove for the nicest (smooth) functions, and then estimate, for what range of slowly decaying functions does this still work? [2-23: Make great reference to Prof. Aizenman's other notes. Some of this stuff will be added and corrected by Aizenman soon, for next class.](#)

Refer to the notes given by hand as a basis for the discussion today.

We are discussing the first major result of the subject, which is the free energy. For an extensive system, a useful form for Hamiltonian is an extensive function in finite volume,

and $H_\Lambda(\sigma)$ refers to a sum over spin configurations, where $\Lambda \subset \mathbb{Z}^d$ is a subset of a lattice or more generally a homogenous graph. We can write

$$H_\Lambda(\sigma) = \sum_{A \subset \Lambda} \phi_A(\sigma) = \sum_{x \in \Lambda} \left[\sum_{x \in A} \frac{1}{|A|} \phi_A(\sigma_A) \right]$$

This is true by summation by parts. Now why is this useful? This collection includes the sum of those terms which affect the spin at x . It's convenient to define as $\|\phi\| = \sum_{0 \in A} \frac{1}{|A|} \sup_\sigma |\psi_A(\sigma)|$. What are spins in Ising models? For Ising, there are two types of interactions: Pairs of spins, so we sum over $\|\phi_{Ising}\| = \frac{1}{2} \sum_{u \neq 0} |J_{0,u}| + h$, where h comes from the external field. What is the range of interactions tolerated by the formalism? Well, the norm must be finite, so the interactions must be summable. Also note that this expression is good for translation invariant actions (see Eq. 1.12).

The basic theorem we are after is in on pg. 32 of the notes. Let us discuss boundary conditions: If you have a system, it's open to effects outside of the system, which we denote via the $\#$ symbol. Sometimes we like to put periodic boundary conditions. When I was student, it took me a while to figure out why people talk about periodic boundary conditions (we're effectively saying the boundary is neighbors). The reason is for very large systems, you see emergent system invariants, and locally, often, things look similar in different places. So by adding interactions across boundaries, you encourage this directly and it's very mathematically convenient. Thus we have equation 1.18. In fact, in the infinite volume limit, if you had a tiny bit of external field which is moved, the resulting state may or may not depend on it. The argument can tolerate this extra term, as long as it's $\|\frac{\phi_\Lambda^\#(\sigma)}{|\Lambda|}\| \rightarrow 0$ (this part is more important). (See Eq. 1.20).

Now the strategy for the more general theorem is on page 32 (see Thm 1.4.4). Assuming ϕ is translation invariant, we can just assume that $\|\phi\| < \infty$ is finite, and the extra boundary terms are $o(|\Lambda|)$, the quantity $\frac{1}{|\Lambda|} \ln(Z_\Lambda^\#) \rightarrow 4(\beta, \phi)$. For any van-Hove sequence, $\Lambda_n \rightarrow \mathbb{Z}^d$ (approaches the lattice) and asymptotically covers the graph, and the boundary over the volume tends to 0 ($\frac{|\partial N|}{|\Lambda|} \rightarrow 0$), where you measure it in terms of a discretized system. If you take the volume and pack it by large cubes, then first pick the size of the cube you're using for packing, measure the boundary by the number of cubes which overlap, and the volume by the number of cubes which fit inside, then the fraction goes to 0 no matter what size of box you used.

Now this is called the **pressure**. The boundary conditions you add do not effect the pressure. This is the physical notion of pressure. To explain this we will have to go into thermodynamics, so we will postpone that.

I would like to discuss the strategy, since it is an example of how analysts approach problems. You were trying to analyze a system with in principle, a potentially unbounded range of interactions. It seems the energy per volume is dominated by the norm of the interaction. Remember the norm of the interaction is defined in such a way as $|H_\Lambda(\sigma)| \leq \left| \sum_{x \in \Lambda} \left[\sum_{x \in A} \frac{1}{|A|} \phi_A(\sigma) \right] \right|$. So now the question to ask is how can I trim to get the system to get something similar. Now as long as the norm is finite, as long as you chop diameters with terms larger than R , you're not making much of a dent in the volume. For every $\epsilon > 0$, there

exists R_ϵ which is finite such that $|H_\Lambda(\sigma) - H_\Lambda^{R_\epsilon}(\sigma)| \leq \epsilon|\Lambda|$. So you are shifting the energy per density by a small amount. After all, we are interested in the logarithm of the partition function. We have $Z = \int e^{-\beta H(\sigma)} \rho_0(d\sigma_\Lambda)$. The effect of terms which are that small leads us to bound Z by a factor at most exponential. Then $|\frac{1}{|\Lambda|} \ln(Z) - \frac{1}{|\Lambda|} \ln Z_\Lambda^{R_\epsilon}| \leq \epsilon$ to prove it converges.

Now, let us divide up the space of size $n \times n$ into squares of size $m \times m$ with $n \gg m$. Now there may be some boundary layer around the space which includes no m inside. Now how much is the partition function affected by turning off the boundary, and terms which involve different boxes. These terms can be estimated by removing thin rectangles of size R_ϵ (call these of terms of type 1). What is effect of the energy of this on these terms on one over the volume on the partition in this box. The right question to ask is how much energy can you pack into the rectangle. We don't care about constants since this is analysis.

Each box contributes a corridor of order R_ϵ , so we have

$$\left| \frac{1}{|\Lambda_m|} \ln(Z_{\Lambda_m}^{R_\epsilon}) - \frac{1}{|\Lambda_m|} \ln(Z_{\Lambda_m}) \right| \leq d\beta \|\phi\| \left[\frac{R}{m} + \frac{m}{n} \right]$$

Here's a cool way to prove this works: For n fixed, the limit when $n \rightarrow \infty$ yields (if it exists, however, just take the lim sup if this is the case), you learn $|\limsup_{n \rightarrow \infty} \left| \frac{1}{|\Lambda_m|} \ln(Z_{\Lambda_m}^{R_\epsilon}) - \frac{1}{|\Lambda_m|} \ln(Z_{\Lambda_m}) \right| \leq d\beta \|\phi\| \left[\frac{R}{m} + \frac{m}{n} \right] \rightarrow \frac{R}{m} \|\phi\| d\beta$. So now we see that $\frac{1}{|\Lambda_m|} \ln(Z_{\Lambda_m}^{R_\epsilon})$ is a number. We learn that the finite volume partition functions are close to that number, with distance at most $1/m$. Taking $m \rightarrow \infty$, we see that the upper bound goes to 0, and therefore $\frac{1}{|\Lambda_m|} \ln(Z_{\Lambda_m}) \rightarrow \frac{1}{|\Lambda_m|} \ln(Z_{\Lambda_m}^{R_\epsilon})$. This is very typical argument in analysis. If you now have an unbounded interaction, you just use a norm estimate. That's the proof outlined in the notes. Later today I'll upload the full proof using all approximations.

Now what is this free energy/pressure used for? We have $\psi(\beta, \phi)$ or $\psi(\beta, h)$. Finite volume Gibbs equilibrium states are ρ_β are measures in the volume on spin configurations (in discrete case, this amounts to δ -measure where weights add up to 1), and the measure is defined as $\rho_{\beta, \Lambda}^\#(\sigma_\Lambda) = \frac{e^{-\beta H_\Lambda^\#(\sigma_\Lambda)} \rho_0(d\sigma_\Lambda)}{Z_\Lambda^\#}$ where $\beta = 1/T$.

In thermodynamics, you take about extensive quantities, temperature, density, etc. In statistical mechanics, you want to know the joint distribution of states in multitude of local variables in Gibbs equilibrium. How does this inform you? It tells you some information, but not everything. All of this is not complicated once you understand it, but let's talk about probability referring to ρ . Now consider $\frac{1}{|\Lambda|} \ln(Z_\Lambda^\#(\beta + \Delta\beta)) = \frac{1}{|\Lambda|} \ln(Z_\Lambda^\#(\beta))$. Then

$$\frac{1}{|\Lambda|} \ln \int e^{-(\Delta\beta)H_\Lambda(\sigma)} \rho_\beta(d\sigma_\Lambda)$$

So we write

$$Z_\Lambda(\beta + \Delta\beta) = \int e^{-(\beta + \Delta\beta)H_\Lambda} \rho_0(d\sigma) = Z_\Lambda(\beta) \mathbb{E}_\beta(e^{-(\Delta\beta)H_\Lambda})$$

We basically split up the exponent, and divided by the normalizing factor, pulling it outside the integral. The resulting integral is no more than $e^{-(\Delta\beta)H_\Lambda}$ averaged with respect

to an average measure. Thus we have

$$\mathbb{E}_{\beta,\Lambda}(e^{-(\Delta\beta)H_\Lambda}) = \frac{Z_\Lambda(\beta + \Delta\beta)}{Z_\Lambda(\beta)} \geq e^{-\Delta\beta E} \mathbb{P}_\beta(H_\Lambda < E)$$

This is exponentially small. Hence, for $\Delta\beta < 0$, we learn that $\mathbb{P}_\beta(H_\Lambda < E) \leq e^{E\Delta\beta} \frac{Z_\Lambda(\beta+\Delta\beta)}{Z_\Lambda(\beta)}$. Now divide by the volume to get

$$\frac{1}{|\Lambda|} \ln(\mathbb{P}_\beta(H_\Lambda < E)) \leq \Delta\beta \frac{E}{|\Lambda|} [\psi(\beta + \Delta\beta) - \psi(\beta)]$$

Let me just stop and give a picture and see the organized notes. Let's just explain what's going on: Convex functions have directional derivatives exists, this means that as you move along ψ (image) and β (domain), the pressure along β gives the mean value of the energy. What does this tell you about actual values? We want to extract that the energy deviates by ϵ is a probability distribution, we can get it by studying the way the function behaves, the main idea is that the slope of the function gives you mean energy, and if you want to study energy which differs from slope, look at the free energy function. If you move by slope different from tangent, you reach right above it. Then the biggest gap you can create is a gap of exponential decay which is how much probability there is between yourself and the mean.

At those temperatures where the pressure is differentiable, then the energy density is equal to the thermodynamic value and is exponentially small in the volume. Each of the Hamiltonian terms will get its own shift-invariant coefficient. Functions which are jointly convex in a finite number of parameters are differentiable almost everywhere. You can then read the expected values of the energy, the quantity related to H (magnetization), and various other local variables, can all be read off of this thermodynamic function ψ , and if you want more parameters, just add those. So you capture parameters which are existing where pressure is differentiable. People who study statistical mechanics not where everything is smooth, but where the pressure is not differentiable. What happens at such points? The Gibbs equilibrium states behave as if the slope of the function is the same at the point according to higher up, it's like there is one state. From the other side, you get a different energy density. It's very similar to phase transitions (solid \rightarrow liquid \rightarrow gas).

2-25

5.2 Convexity of the pressure and its implications

Theorem 1.5.8 (Hölder inequality). *For any pair of functions $f, g : \Omega \rightarrow \mathbb{R}$ on a (positive) measure space, for $\frac{1}{p} + \frac{1}{q} = 1$,*

$$\left| \int f(\omega)g(\omega) d\mu(\omega) \right| \leq \int \left[\int |f|^p d\mu \right]^{\frac{1}{p}} \left[\int |g|^q d\mu \right]^{\frac{1}{q}}.$$

This can be proved using Jensen's inequality.

Theorem 1.5.9. For any function $\phi : \Omega \rightarrow \mathbb{R}$,

$$P(h) = \ln \int e^{h\phi(\omega)} d\mu(\omega)$$

is convex in h .

Proof 1. Differentiate:

$$\begin{aligned} P'(h) &= \langle \phi \rangle_h \\ P''(h) &= \langle (\phi - \langle \phi \rangle_h)^2 \rangle_h \geq 0. \end{aligned}$$

□

Proof 2. Let $h_t = th_0 + (1-t)h_1$. Then

$$\begin{aligned} \int e^{h_t \phi} d\mu &= \int (e^{h_0 \phi})^t (e^{h_1 \phi})^{1-t} d\mu \\ &\leq \left(\int e^{h_0 \phi} d\mu \right)^t \left(\int e^{h_1 \phi} d\mu \right)^{1-t}. \end{aligned}$$

Hence

$$P(h_t) \leq tP(h_0) + (1-t)P(h_1).$$

□

Note the proof doesn't assume differentiability of ϕ .

We used the density for fluctuation estimates.

Our goal is to estimate

$$\mathbb{P}(H(\sigma) \leq \langle H \rangle_\beta - \varepsilon |\Lambda|) = \int_{\Omega_\Lambda} \mathbb{1}[H(\sigma) \leq \langle H \rangle_\beta - \varepsilon |\Lambda|] \frac{e^{-\beta H_\Lambda(\sigma)} \rho_0(d\sigma)}{Z_\Lambda(\beta)}.$$

We use the bound

$$\mathbb{1}[H \leq \langle H \rangle - \varepsilon |\Lambda|] \leq e^{-tH} e^{t(\langle H \rangle - \varepsilon |\Lambda|)}.$$

Then

$$\begin{aligned} \mathbb{P}(H(\sigma) \leq \langle H \rangle_\beta - \varepsilon |\Lambda|) &\leq \left[\int_{\Omega_\Lambda} e^{-tH} \rho_\beta(\sigma) \right] e^{t(\langle H \rangle - \varepsilon |\Lambda|)} \\ &= \frac{Z(\beta + t)}{Z(\beta)} e^{t(\langle H \rangle - \varepsilon |\Lambda|)}; \end{aligned}$$

we got the partition temperature at shifted temperature. The partition function is a generating function: its derivative up to sign gives energy,

$$\psi(\beta) = \frac{1}{|\Lambda|} \ln \left[\underbrace{\int e^{-\beta H_\Lambda} \rho_0(d\sigma)}_{Z_\Lambda(\beta)} \right]$$

$$\langle H \rangle_\beta = -\psi'_A(\beta) |\Lambda|.$$

Using this, we get

$$\begin{aligned} \mathbb{P}(H(\sigma) \leq \langle H \rangle_\beta - \varepsilon |\Lambda|) &\leq \frac{Z(\beta + t)}{Z(\beta)} e^{-t[\psi'(\beta) + \varepsilon]|\Lambda|} \\ &\leq \frac{e^{\psi(\beta+t)|\Lambda|}}{e^{(\psi(\beta) + [\psi'(\beta) + \varepsilon]t)|\Lambda|}}. \end{aligned}$$

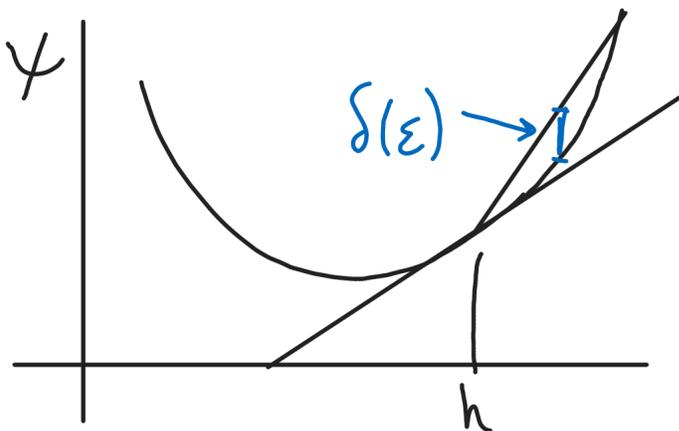
Thinking of ψ as a function of inverse temperature β , ψ is convex. The slope at a point is $\psi'(\beta)$. In the denominator is a linear function which is equal to $\psi(\beta)$ at β but increases a bit faster than the tangent line.

By convexity, for all $\varepsilon > 0$, $t > 0$ small enough,

$$\psi(\beta) + t_\varepsilon[\psi'(\beta) + \varepsilon] - \psi(\beta + t) \geq 0$$

Pick t to maximize this difference. Denoting $\delta(\varepsilon) = \sup_{t>0} \delta(t)$ we have

$$\mathbb{P}_\beta(H_\Lambda \leq \langle H_\Lambda \rangle - \varepsilon |\Lambda|) \leq e^{-\delta(\varepsilon)|\Lambda|}.$$



The probability that the energy density falls by ε below its mean decays exponentially. This is the kind of decay we were talking about with large deviation principles. Similarly we can show

$$\mathbb{P}_\beta(H \geq \langle H \rangle + \varepsilon |\Lambda|) \leq e^{-\delta_+(\varepsilon)|\Lambda|}.$$

This is left as an argument. We can repeat the argument, or apply what we proved to a system where the Hamiltonian is the negative.

The convex functions converge to a convex function. In the infinite volume limit, the function may not be differentiable, but the left and right derivatives exist. Our argument is really comparing with the slope of the function coming from above; for the other inequality, compare with the slope of the function from below.

Since the pressure converges, you can conclude that given a van Hove sequence, for Λ large enough, you can conclude the bounds (not optimally) with $\delta/2$.

5.3 Large deviation principle for van Hove sequences

Theorem 1.5.10. *Assume that the infinite volume pressure ψ is differentiable at β_0 . Then for all $\Lambda_n \uparrow \mathbb{Z}^d$ (in the van Hove sense). Then for any ε , there is $N < \infty$ such that for all $n \geq N$,*

$$\mathbb{P}_{\beta, \Lambda_n} \left(\left| \frac{1}{|\Lambda|} H_\Lambda(\sigma) + \psi'(\beta) \right| \geq \varepsilon \right) \leq e^{-\frac{\delta(\varepsilon)}{2} |\Lambda|}.$$

If the infinite volume pressure is differentiable at a particular β , then in large enough finite volume, then the fluctuations per volume is small. with high probability. Look at the infinite volume function; for every ε , ask when we move right (left) at ε greater (less) than the slope, how far we move to have the largest discrepancy. Wait for the gap between the infinite and finite volume case function to be less than half.

So far we studied the distribution of the energy. But you may want to know the distribution of other quantities such as magnetization, spin, and number of clusters of a certain class. These are covered by the same analysis except you consider the dependence of the free energy.

TODO: edit What is the physical implication for a lattice model? If you know the pressure as a function of inverse temperature. Pressure is a thermodynamic function. We are interested in sysm of deg ... tell us something... typical range of values of the energy of this configuration. At places where the energy is differentiable, the total energy per unit volume does not fluctuate much. We expect $\sqrt{\text{volume}}$ of fluctuations. This is like CLT for nonindependent variables.

What happens if the pressure is not differentiable? Without assuming differentiability of $\psi(\beta)$, we have

$$\begin{aligned} \mathbb{P}_{\beta, \Lambda} \left(\frac{1}{|\Lambda|} H \geq \psi'(\beta + 0) + \varepsilon \right) &\leq e^{-\frac{\delta(\varepsilon)}{2} |\Lambda|} \\ \mathbb{P}_{\beta, \Lambda} \left(\frac{1}{|\Lambda|} H \leq \psi'(\beta - 0) - \varepsilon \right) &\leq e^{-\frac{\delta(\varepsilon)}{2} |\Lambda|}. \end{aligned}$$

(Note $\psi'(\beta - 0) \leq \psi'(\beta + 0)$.) To extend to the densities of other local observables,

$$H(\sigma) = \sum_{A \subseteq \mathbb{Z}^d} J_A \phi_A(\sigma)$$

with $J_A = J_{S_n} A$ for all shifts S_n . (The proper way to do this is to introduce coupling constants depending on the equivalence class of terms modulo translation. For example, Ising coupling has 2 parameters corresponding to nearest neighbor coupling and external field.)

For a collection of translation-invariant interactions, we denote

$$\psi(\beta, J) = \lim_{\lambda \uparrow \mathbb{Z}^d} \frac{1}{|\Lambda|} \ln Z_\Lambda(J, \phi).$$

Theorem 1.5.11. *For any translation-invariant interaction with $\|J\| < \infty$, the limit $\psi(\beta, J)$ exists and is a convex function of J (it is jointly convex in all parameters).*

We think of pressure as a function of both the temperature and parameters such as the magnetic field. For example, for $H = -J \sum_{\{x,y\}, |x-y|=1} \sigma_x \sigma_y - h \sum_x \sigma_x$, ψ is a function $\psi(\beta, J, h)$. We have

$$\frac{\partial \psi_\Lambda}{\partial h} = \beta \left\langle \frac{1}{|\Lambda|} \sum_{x \in \Lambda} \sigma_x \right\rangle.$$

and

$$\begin{aligned} \psi_\Lambda(\beta, J, h) &= \frac{1}{|\Lambda|} \ln \int e^{\beta(J \sum_{\langle x,y \rangle} \sigma_x \sigma_y + h \sum_{x \in \Lambda} \sigma_x)} \rho_0(d\sigma) \\ \psi(\beta, J) &= \lim_{\Lambda \uparrow \mathbb{Z}^d} \frac{1}{|\Lambda|} \ln Z_\Lambda(J\phi) \end{aligned}$$

Introduce as many parameters as you think are relevant.

In the homework, you will look at the simplest system, ideal gas. There are no interactions. We introduce something akin to the magnetization. Compute the pressure for an ideal gas and use this to compute the entropy of the system. The Legendre transform allows you to read the ... from the free energy. You'll learn how to arrive at this celebrated formula

$$-\rho \ln \rho + (1 - \rho) \ln(1 - \rho).$$

It becomes a challenge to evaluate partition functions and find their singularities. Using calculated the partition function for the 1-D model; you will do the calculation too. He said it has no singularity. If you move to 2 dimensions, it does exhibit a phase transition. The 2-D model is solvable by other techniques. [3-1: First, notational adjustments](#)

Let

$$H_\Lambda(\sigma) = \sum_{A \subseteq \Lambda} J_A \phi_A(\sigma)$$

and assume ϕ is normalized so that

$$\sup_{\sigma} |\phi_A(\sigma)| = 1$$

(if it is not identically 0). For $J = \{1_A\}$, define

$$\|J\| = \sum_{A \ni 0} \frac{1}{|A|} |J_A|.$$

Definition 1.5.12: H is translation invariant if for all $A \subseteq \Lambda$,

1. (Coupling is translation covariant) $\phi_{A+u}(\sigma) = \phi_A(S_u \sigma)$, where $(S_u \sigma)_x = \sigma_{x+u}$.
2. $J_A = J_{A+u}$.

Note that ϕ is really “shift covariant” whereas H is “shift invariant”, i.e. $H(S_n\sigma) = H(\sigma)$ (formally). This is a formal expression because for an infinite system, the energy is infinite. (There are various ways to make sense of this mathematically. For example, you can look at a finite system with periodic boundary conditions, and we can talk about shift invariance there.)

For translation invariant Hamiltonians,

$$\psi(\beta, J) = \lim_{\Lambda \uparrow \mathbb{Z}^d} \frac{1}{|\Lambda|} \ln Z_\Lambda(\beta, J)$$

where $Z_\Lambda(\beta, J) = \int_{\Omega_\Lambda} e^{-\beta H_\Lambda(\beta)} \rho_0(d\sigma_\Lambda)$.

We proved 2 theorems, Theorem 1.5.11 and the following. Let h be one of the parameters of $J = \{J_\Lambda\}$ and M be the conjugate quantity

$$M_\Lambda(\sigma) = -\frac{\partial}{\partial h} H_\Lambda(\beta, J).$$

$$M_\Lambda(\sigma) = -\frac{\partial}{\partial h} H_\Lambda(\beta, J).$$

The example to keep in mind is the following.

Example 1.5.13: The Ising model is $H_\Lambda(\sigma) = \sum_{\{x,y\}} J_{x-y} \sigma_x \sigma_y - h \sum_{x \in \Lambda} \sigma_x$. Here,

$$-\frac{\partial}{\partial H_\Lambda(\sigma)} \partial h = \sigma_{x \in \Lambda} \sigma_x \equiv M.$$

We have $\frac{1}{|\Lambda|} \mathbb{E}(M_\Lambda) = \frac{1}{\beta} \frac{\partial}{\partial h} \psi$.

Theorem 1.5.14. For any (β, J) , for $\varepsilon > 0$, there exists $\delta(\varepsilon) > 0$ such that for any van Hove sequence,

$$\lim_{\Lambda \uparrow \mathbb{Z}^d} \frac{1}{|\Lambda|} \ln \mathbb{P}_\Lambda \left(\begin{array}{l} \frac{1}{|\Lambda|} M_\Lambda(\sigma) \geq \frac{1}{\beta} \frac{\partial}{\partial h} \psi \Big|_{+0} + \varepsilon \text{ or} \\ \frac{1}{|\Lambda|} M_\Lambda(\sigma) \leq \frac{1}{\beta} \frac{\partial}{\partial h} \psi \Big|_{-0} - \varepsilon \end{array} \right) \leq -\delta(\varepsilon)$$

This says that

$$\mathbb{P}_\Lambda \left(\frac{1}{|\Lambda|} M_\Lambda \geq \frac{\partial}{\partial h} \psi + \varepsilon \right) \approx e^{-\delta(\varepsilon)|\Lambda|}.$$

6 1-D Ising model

Consider the 1-D Ising model. The spins are $\sigma_n = \pm 1$; the energy is

$$H(\sigma) = -J \sum_n \sigma_n \sigma_{n+1} - h \sum_n \sigma_n.$$

The first term encourages neighbors to agree, and the second term encourages the spins to agree with the applied field.

(Here we just have nearest-neighbor interactions. In general we can have interactions between longer distances, which complicates the calculations. When there are long-range interactions (following a power law), there can be a phase transition.)

Theorem 1.6.1. *Let $\hat{h} = \beta h$. Then*

$$\psi(\beta, h) := \lim_{N \rightarrow \infty} \frac{1}{N} \ln Z_N = \ln[e^{\beta J} \cosh \hat{h} + (e^{2\beta J} \sinh^2 \hat{h} + e^{-2\beta})^{\frac{1}{2}}]$$

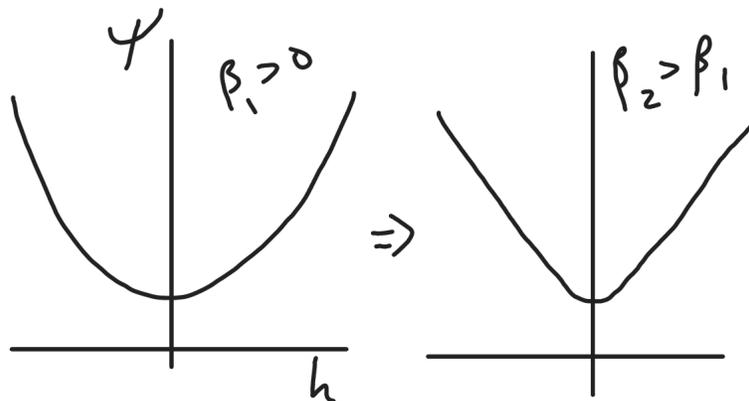
The particular form of the formula is not so important; what's important is that it is explicit. As a corollary, ψ is differentiable at all $\beta \geq 0$, $h \in \mathbb{R}$ with

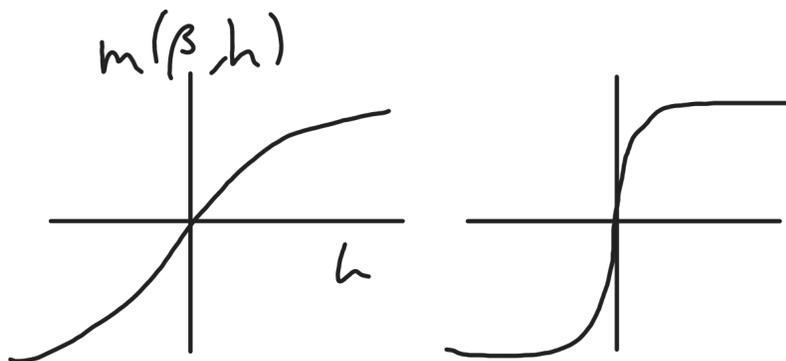
$$m(\beta, h) := \frac{\partial \psi}{\partial h} = \frac{e^{\beta} \sinh \hat{h}}{[e^{2\beta} \sinh^2 \hat{h} + e^{-2\beta}]^{\frac{1}{2}}}$$

In particular,

$$\lim_{\beta \rightarrow \infty} m(\beta, h) = \begin{cases} 1, & h > 0 \\ 0, & h = 0 \\ -1, & h < 0. \end{cases}$$

Note that we do not see the signature of a first-order phase transition, that the energy is discontinuous. The derivative is discontinuous, but the energy is not. We observe something like this only at infinite β . This is a general feature of 1-D statistical mechanical models: there is no symmetry breaking unless the model has very long-range interactions.





Proof. Take an Ising model with periodic boundary conditions $\sigma_L = \sigma_0$. Then (it helps to symmetrize the sum)

$$\begin{aligned} Z_{[0,L]}^{\text{per}} &= \sum_{\substack{\sigma_j = \pm 1 \\ \sigma_0 = \sigma_L}} e^{\beta \sum_{m=1}^L \sigma_m \sigma_{m+1} + \hat{h} \sum_{n=1}^L \frac{\sigma_n + \sigma_{n+1}}{2}} \\ &= \text{tr}(A^L) \end{aligned}$$

where

$$A = \begin{pmatrix} A_{++} & A_{+-} \\ A_{-+} & A_{--} \end{pmatrix} = \begin{pmatrix} e^{\beta + \hat{h}} & e^{-\beta} \\ e^{-\beta} & e^{\beta - \hat{h}} \end{pmatrix}.$$

Recall that

$$\text{tr}(A^L) = \sum_{\sigma_0 = \sigma_L \in \pm 1} (A^L)_{\sigma_0 \sigma_L} = \sum_{\substack{\sigma_i = \pm 1 \\ \sigma_0 = \sigma_L}} A_{\sigma_0, \sigma_1} A_{\sigma_1, \sigma_2} \cdots A_{\sigma_{L-1}, \sigma_0}.$$

Since A is self-adjoint, A is diagonalizable. The eigenvalues have different modulus, $|\lambda_1| > |\lambda_2|$, and

$$\begin{aligned} \text{Tr}(A^L) &= \lambda_1^L + \lambda_2^L = \lambda_1^L \left[1 + \left(\frac{\lambda_2}{\lambda_1} \right)^L \right] \\ \frac{1}{L} \ln Z_L &= \ln \lambda_1 + \underbrace{\frac{1}{L} \ln \left[1 + \left(\frac{\lambda_2}{\lambda_1} \right)^L \right]}_{\approx \frac{1}{L} \left(\frac{\lambda_2}{\lambda_1} \right)^L} \quad \text{using } \ln(1 + \varepsilon) \approx \varepsilon. \end{aligned}$$

$$\lim_{L \rightarrow \infty} \frac{1}{L} \ln Z_L = \ln \lambda_1$$

The matrix A is a **transfer matrix** (as in a Markov chain). Fixing the boundary conditions σ_0, σ_L , we have

$$Z_L^{b.c.} = \sum A_{\sigma_0 \sigma_1} A_{\sigma_1 \sigma_2} \cdots A_{\sigma_{L-1} \sigma_L} = (A^L)_{\sigma_0 \sigma_L}$$

(Exercise: write this in terms of the eigenvalues.) □

Consider $Z_L^{++}, Z_L^{+-}, Z_L^{-+}, Z_L^{--}$. Assuming the magnetic field is positive ($h > 0$), the maximum is Z_L^{++} ; if $h < 0$ then the maximum is Z_L^{--} . We have

$$\left| \frac{Z^{++} - Z^{--}}{Z^{++}} \right| \leq \left| \frac{\lambda_2}{\lambda_1} \right|^L$$

With open-ended boundary conditions, the spins are an array of variables with the Markov property. If you specify a spin, then the distribution of the future does not depend on the past beyond the value of the previous spin. Thus the values form a Markov chain, and the transfer matrix is familiar from the theory of Markov chains.

What happens beyond 1 dimension? Then the Markov property re-emerges in an interesting way. When you talk about the distribution of a finite system, it's sufficient to give the spins at the boundary. There is a multidimensional Markov property, specified by boundary spins. We will arrive at the theory of Gibbs states. [3-3-16](#)

6.1 Transfer matrix method

This method applies to any finite model in 1 dimension. Suppose the individual states are $\sigma = \pm 1$, $\{\sigma_i\}_{i=0}^L$. The probability of observing a particular configuration is

$$\mathbb{P}_\beta^{b.c.}(\{\sigma_{\Lambda_L}\}) = \frac{e^{-\beta H_L^{bc}(\sigma)}}{Z_L^{br}(\beta)} \mathbb{1}[\sigma_0, \sigma_L \text{ satisfy boundary conditions}] \quad (1.13)$$

$$\mathbb{P}_\beta(\sigma \text{ has boundary conditions } \sigma_0, \sigma_L) = (T^L)_{\sigma_0, \sigma_L} \quad (1.14)$$

$$= \sum_{\sigma=\{-1,1\}^{[1,L-1]}} T_{\sigma_0, \sigma_1} \cdots T_{\sigma_{L-1}, \sigma_L}. \quad (1.15)$$

Here,

$$T_{\sigma, \sigma'} = e^{-[\beta J \sigma \sigma' + \hat{h}(\frac{\sigma + \sigma'}{2})]} \quad (1.16)$$

$$= e^{-\beta H(\sigma)} e^{\hat{h}(\frac{\sigma_0 + \sigma_L}{2})}. \quad (1.17)$$

The product of these is

$$e^{-\beta H(\sigma)} = e^{-\beta \sum_{j=0}^{L-1} \sigma_j \sigma_{j+1} - \hat{h} \sum_{j=0}^L \sigma_j}.$$

Another natural boundary condition is the periodic boundary conditions: we only allow states where $\sigma_0 = \sigma_L$, so

$$Z_L^{\text{per}} = Z_L^{+,+} + Z_L^{-,-} = \text{Tr}(T^L).$$

6.2 Markov chains

We summarize some basic facts about Markov chains.

Definition 1.6.2: A Q -state **Markov chain** is a collection of random variables $\{\omega_n\}$ with

values in $[Q] = \{1, \dots, Q\}$. Let $\Psi_n = \begin{pmatrix} P_n(1) \\ \vdots \\ P_n(Q) \end{pmatrix}$ where $P_n(k) = \mathbb{P}(\omega_n = k)$.

Ψ_1 is in a given initial state¹⁷, i.e., the probability distribution of ω_1 , and

$$\Psi_{n+1} = A\Psi_n$$

where $A_{ij} = \mathbb{P}(\omega_{n+1} = j | \omega_n = i)$ is **doubly stochastic** $\sum_j A_{ij} = 1 = \sum_i A_{ij}$, $A_{ij} \geq 0$.¹⁸

Think of n as the discrete time. Taking powers, we have

$$\Psi_{n+1} = A^n \Psi_1.$$

Markov chains have finite memory; the effect of the initial value decays exponentially. Nice Markov chains have a “destiny” (equilibrium distribution).¹⁹

A 1-D finite system looks like a Markov chain; boundary conditions are equivalent to specifying initial and final states.

Theorem 1.6.3 (Perron-Frobenius). *Let $A = (a_{ij})$ be a $Q \times Q$ positive matrix $a_{ij} \geq 0$.*

1. *There is a positive eigenvalue r (the Perron-Frobenius eigenvalue) such that the spectral radius of A is $\rho(A) = r$ and A has no other eigenvalue with $|\alpha| = r$.*
2. *This eigenvalue is simple, and the corresponding eigenvector is positive.*
3. *Letting $|v\rangle, \langle w|$ be the left and right eigenvectors of A corresponding to the largest eigenvalue, $A|v\rangle = r|v\rangle, \langle w|A = r\langle w|$,²⁰*

$$\lim_{L \rightarrow \infty} \frac{1}{r^L} A^L = |v\rangle \langle w|$$

In other words, there is S such that

$$S^{-1}AS \approx r^L \begin{pmatrix} 1 & 0 & \cdots \\ 0 & 0 & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix}.$$

¹⁷Probabilists use “state” to mean the value of the random variable; physicists use “states” to correspond to probability measure. (In quantum physics, not all variables have concrete values.)

¹⁸The LHS equality says that the Markov chain is reversible.

¹⁹The basic laws of physics seem reversible; how to explain the arrow of time? Consider colliding molecules; the impact parameters seem to be coming without precise aim in an uncorrelated way (Boltzmann’s hypothesis), which isn’t true if it is played in reverse. Entropy keeps increasing. Is there a moment when the entropy starts decreasing; does the arrow of time reverse?

²⁰We use Dirac’s notation $|v\rangle = v^T, \langle w| = w$.

Applying this to the 1-D Ising model,

$$Z^{\text{per}} = \sum_{j=1}^Q \lambda_j^L \quad (1.18)$$

$$\approx \lambda_{\max}^L \left(1 + \left| \frac{\lambda_2}{\lambda_1} \right|^L + \dots \right) \quad (1.19)$$

where

$$1 > \left| \frac{\lambda_2}{\lambda_1} \right| = e^{-\delta}$$

where $\delta\lambda > 0$ is the spectral gap, the gap between λ_{\max} and the next smaller eigenvalue (in absolute value). Using this, it is an exercise to show that in the Ising model, in the infinite volume limit the state exhibits exponential decay of correlations,

$$\lim_{L \rightarrow \infty} \left| \langle \sigma_x, \sigma_y \rangle_{[-L, L]} \right| \leq A e^{-\delta|x-y|},$$

where $\langle A, B \rangle = \langle AB \rangle - \langle A \rangle \langle B \rangle$ is the covariance. In the 1-D Ising model,

$$\Psi(\beta, h) = \lim_{L \rightarrow \infty} \frac{1}{L} \ln Z_L = \ln \lambda_{\max},$$

where λ_{\max} is the Perron-Frobenius eigenvalue of $T = \begin{pmatrix} e^{\beta+\hat{h}} & e^{-\beta} \\ e^{-\beta} & e^{\beta-\hat{h}} \end{pmatrix}$.

The matrix elements are analytic. Can we conclude from this the analyticity of eigenvalues? They are when they don't bump into each other. When they do, funny things can happen. But here, Perron-Frobenius tells us that eigenvalues will not be equal (all entries in T are positive). Thus the eigenvalues will be differentiable in β, h and there is no phase transition.

7 2-D Ising model

Studying the transfer matrix going up, certain models are integrable; you can get a formula for the partition function. We'll present a different method which is culturally very important. There is no closed-form solution in 3-D.

In contrast to the 1-dimensional case, in 2 dimensions, the free energy is not differentiable.

Next we will show that for the 2D Ising model, for β large enough ($\beta > \beta_c$) there exists $m(\beta) > 0$ such that

$$\mathbb{P}_{\beta, k=0}^+ \left(\frac{1}{|\Lambda_L|} \sum_{x \in \Lambda_L} \sigma_x m \geq \frac{m}{2} \right) \rightarrow 1 \quad (1.20)$$

$$\mathbb{P}_{\beta, k=0}^- \left(\frac{1}{|\Lambda_L|} \sum_{x \in \Lambda_L} \sigma_x m \leq \frac{m}{2} \right) \rightarrow 1 \quad (1.21)$$

as $L \rightarrow \infty$. This indicates a phase transition. The signature of a phase transition is that when you take $\beta > \beta_c$, $h = 0$ (so the Hamiltonian is invariant under spin-flip).

The Hamiltonian is $H = -\sum_{x,y \text{ neighbors}} \sigma_x \sigma_y$. With $+$ boundary conditions, the spins are overwhelmingly positive. With $-$ boundary conditions, the spins are overwhelmingly negative. The system deep inside remembers the boundary conditions. Locally, there is symmetry; globally we have symmetry breaking.

What if we take symmetric boundary conditions, with $+$'s and $-$'s? Since it is invariant under a reflection, the probability of observing $\geq \frac{m}{2} + \varepsilon$ vanishes exponentially. The system organizes itself as follows: close to the $+/-$ boundary there is mostly $+/-$ with islands of $-/+$; in between there is an interface.

We will see Peiech's argument (?).

3-8

Peiech showed that these models can exhibit symmetry breaking.

In the 2-D Ising model,

$$H = -J \sum_{\{x,y\}} \sigma_x \sigma_y - h \sum_x \sigma_x.$$

Theorem 1.7.1 (Peiech). *For the 2D nonnegative Ising model, there exists $\beta_0 < \infty$ such that for $\beta > \beta_0$, at $h = 0$,*

$$\begin{aligned} \langle \sigma_x \rangle_{\Lambda, \beta, 0}^+ &\geq m(\beta) \\ \langle \sigma_x \rangle_{\Lambda, \beta, 0}^- &\leq -m(\beta). \end{aligned}$$

with some $m(\beta) > 0$.

Why is this “symmetry breaking”? If you take the model at external field $h = 0$ (so that the measure is invariant under global spin flip), the expected value of any anti-symmetric quantity like spin is 0. If you put $+$ boundary conditions all around, then you expect there to be more $+$'s. More surprisingly, if you fix *any* spin, even if the spins are far away, the bias persists.

The theorem implies by linearity

$$\left\langle \frac{1}{|\Lambda|} \sum_{x \in \Lambda} \sigma_x \right\rangle^+ \geq m(\beta).$$

Taking the van Hove limit,

$$\begin{aligned} \frac{\partial}{\partial h} \Psi(\beta, 0+) &= \lim_{\Lambda \uparrow \mathbb{Z}^2} \left\langle \frac{1}{|\Lambda|} \sum_{x \in \Lambda} \sigma_x \right\rangle^+ \geq m(\beta) \\ \frac{\partial}{\partial h} \Psi(\beta, 0-) &= \lim_{\Lambda \uparrow \mathbb{Z}^2} \left\langle \frac{1}{|\Lambda|} \sum_{x \in \Lambda} \sigma_x \right\rangle^- \leq -m(\beta). \end{aligned}$$

Thus the left and right derivatives of Ψ are unequal. There is a cusp at 0; there is a phase transition.

The theorem shows there is an asymmetry in the infinite volume limit of magnetization at a site.

For each $\sigma_\Lambda \in \Omega_\Lambda$ draw a green line separating $+/-$ spins. Using

$$\sigma_x \sigma_y = 2\mathbb{1}[\sigma_x = \sigma_y] - 1$$

(this takes ± 1 values). We have

$$e^{-\beta H(0)} = \prod_{\{x,y\}} e^{\beta \sigma_x \sigma_y} = e^{\beta |\xi|} \prod_{\{x,y\}} e^{-2\beta \mathbb{1}[\sigma_x \neq \sigma_y]} = K e^{-2\beta |C|}$$

where K is a constant and C is the set of contour lines.

If the boundary conditions are $+$, then for there to be a $-$, it must be separated from the boundary by a green loop. We want to show that the probability for a given site to be separated from the boundary by a loop is strictly less than $\frac{1}{2}$. Then the probability it agrees with the boundary is $> \frac{1}{2}$.

Let Ω_γ be the set of configurations which has γ as a contour.

Lemma 1.7.2 (Peiech). *The probability that a given loop γ is among the contours lines $C(\sigma)$ satisfies*

$$\mathbb{P}_{\Lambda,\beta,0}^+(C(\sigma) \ni \gamma) = \langle \mathbb{1}_{\Omega_\gamma} \rangle_{\Lambda,\beta,0}^+ \leq e^{-2\beta |\gamma|}.$$

Proof. First we do an energy estimate.

The idea is that each configuration can be associated with another configuration by a mapping T_γ which flips the spin of all sites inside γ . The disagreements of γ would be erased. Because all spins inside are flipped, no other disagreement is affected. Note T_γ is an involution—it is its own inverse.

For each $\sigma \in \Omega_\gamma$, its flip has higher probability because it has less disagreement, so

$$\mathbb{P}^+(\{\sigma\}) = e^{-2\beta |\gamma|} \mathbb{P}(\{T_\gamma \sigma\})$$

Summing over $\sigma \in \Omega_\gamma$,

$$\mathbb{P}^+(\Omega_\gamma) = e^{-2\beta |\gamma|} \mathbb{P}(T_\gamma \Omega_\gamma) \leq e^{-2\beta |\gamma|}.$$

Next we need an entropy estimate: we need to estimate how many loops there are so we can do a union bound over them.

Lemma 1.7.3 (Entropy estimate). *The number of non-repeating loops of length l starting from a specified point is $N_l \leq \frac{4}{3} 3^l$.*

Proof. There are 4 possibilities for the first step; for each of the next step there are ≤ 3 choices for the other $l - 1$ steps because it can't backtrack. \square

Now we put the energy and entropy bound together. For each each loop around x , we can start it at an node between 0 and $\frac{l}{2}$ steps to the right of x where the path goes up. Then

$$\begin{aligned} \mathbb{P}^+(\sigma_x = -1) &\leq \mathbb{P}^+(x \text{ is separated from } \partial\Lambda) \\ &\leq \sum_{\gamma \text{ loop separating } x \text{ from } \partial\Lambda} e^{-2\beta|\gamma|} \\ &\leq \sum_{l=4}^{\infty} \frac{l}{2} 3^{l-1} e^{-2\beta l} \\ &= \frac{1}{6} \sum_{l=4}^{\infty} l(3e^{-2\beta})^l. \end{aligned}$$

This is $\leq \frac{1}{2} - \delta(\beta)$ for $\beta > \beta_0$ where β_0 is defined by the condition

$$\frac{1}{6} \sum_{l=4}^{\infty} l(3e^{-\beta_0})^l = \frac{1}{2}.$$

□

The condition is ugly. With some extra finesse, we can show $3e^{-\beta} < 1$ is sufficient for there to be symmetry breaking.

This implies that for $\beta > \beta_0$

$$\langle \sigma_0 \rangle_{\mathbb{Z}^2}^+ = \lim_{\Lambda \uparrow \mathbb{Z}^2} \langle \sigma_0 \rangle_{\Lambda, \beta, 0}^+ \geq m(\beta) > 0.$$

Here we define an infinite volume probability distribution by taking the limit of finite probability distributions with + boundary conditions. Similarly define $\langle \sigma_0 \rangle_{\mathbb{Z}^2}^-$. Symmetry breaking refers to the fact that

$$\langle \sigma_0 \rangle_{\mathbb{Z}^2}^+ \neq \langle \sigma_0 \rangle_{\mathbb{Z}^2}^-.$$

Consider a function $G(m) = G(-m)$. If G has the symmetry of the variational problem, the minimum is at 0. If the minimum does not have the symmetry of the variational problem, there may be multiple minima. We refer to this as symmetry breaking.

Remark 1.7.4: One can deduce symmetry breaking already for $\beta > \beta_1$ with β defined by $3e^{-2\beta_1} = 1$.

Actually we're counting self-avoiding walks; the rate that they grow is smaller than 3^l . So we can replace 3 by the actual rate; this is more complicated. It's harder to come up with exact values.

The part of the estimate which is not relevant is short loops. Let's look at a larger scale and only count loops which surround a larger box. If the probability of there being a large loop around x is $< \frac{1}{2}$, then we can already conclude symmetry breaking. For $3e^{-2\beta} < 1$, we can choose l large enough so that this holds.

If the box of size n is not surrounded by a loop, we can conclude there exists some loop of + spins outside the box which is connected to the boundary. This event is not compatible

with the flip of the event. Thus the probability that there is a loop of $-$'s of size $> n$ connected to the boundary (with $-$) is $< \frac{1}{2}$.

In the infinite volume case, we cannot both have a loop connected to ∞ with $-$'s and a loop connected to ∞ with $+$'s. The $+$ event is more likely.

If you start with $+$ boundary conditions, the $+$ case will occur with large probability.

There is percolation of $+$ sites. If there is no symmetry breaking, there is percolation of $-$ sites. You have to work to show that you can't have infinite line of $+$'s and infinite line of $-$'s with nonzero probability Harris developed the theory of this. We take a shortcut and show that a stronger condition (with the ring of $+/-$'s) is incompatible.

When you take this approach, there are 2 special values of the parameters: a β for which we get no information (the sum diverges), and the β for which the sum is less than a certain threshold ($\frac{1}{2}$). Looking at a renormalized version which eliminates the noise from the low cases of the sum, and you can improve the β threshold.

In higher dimensions, you have to enumerate surfaces. The energy bound is easy. The energy (enumerating surfaces) is messier. The entropy estimate gives worse bounds. You have to be clever to bring other techniques.

Another way is to take the restriction to 2 dimensions (a slice of the cube). If you decouple the layers you just get the 2D model. Having other spins around only increases cooperations. This brings us to the use of inequalities. For the 3D model, the magnetization is greater than the 2D model for the same temperature.

With appropriate inequalities one can show

$$\langle \sigma_0 \rangle_{[-L,L]^3}^+ \geq \langle \sigma_0 \rangle_{[-L,L]}^+.$$

This implies symmetry breaking in higher dimensions at the same temperature, so the critical temperature is smaller,

$$\beta_c^{3D} \leq \beta_c^{2D}.$$

Here

$$\beta_c = \inf \{ \beta : \langle \sigma_0 \rangle_{\beta'} > 0 \forall \beta' > \beta \}.$$

For 2 dimensions, β_c is calculable; for 3 dimensions we don't have an exact value. The 2-D Ising model is exactly solvable by a nontrivial argument. We'll present Feynman's solution.

How singular is the magnetization? That's not expected to change with local details, but with dimension. We do not know β_c for range 5 Ising model, but we expect that the power of $m(\beta)$ is the same (it grows like a power function from the critical temperature).

3-10: We continue the 2D Ising model. Much of what we say will hold true for other planar graphs.

The energy is

$$\sum_{b \in \mathcal{E}} J_b (\sigma \sigma)_b - h \sum_{x \in G} \sigma_x$$

where the first sum is over edges (bonds) and if $b = b_1 b_2$, then $(\sigma \sum_{n=0}^{\infty})_b$ denotes $\sigma_{b_1} \sigma_{b_2}$.

The state changes discontinuously across the line $h = 0$ until some critical temperature T_c . Along the symmetry line the model is solvable.

Theorem 1.7.5 (Onsager). *For the nn \mathbb{Z}^2 at $h = 0$, $J_b = J$,*

$$\frac{1}{|\Lambda|} \ln Z_\Lambda^{per}(\beta, 0) = \ln \cosh(\beta J) + \frac{1}{2} \int_{(-\pi, \pi]^2} \frac{dx_1 dx_2}{(2\pi)^2} \ln \{ [Y(\beta) + Y^{-1}(\beta) - 2] + \mathcal{E}(k) \}$$

where $\mathcal{E}(k) = 2 \left[\sin^2\left(\frac{k_1}{2}\right) + \sin^2\left(\frac{k_2}{2}\right) \right]$.

By the integral we really mean the discrete sum over $k \in (-\pi, \pi]^2 \cap \frac{2\pi}{L} \mathbb{Z}^2$. (In the infinite volume limit it is an integral.)

For square summable function in a box, we can look at a basis of delta functions, or the Fourier basis $e^{ik \cdot x}$ for $k \in (-\pi, \pi]^2 \cap \frac{2\pi}{L} \mathbb{Z}^2$. I.e., k is sampled from the dual lattice; $[-\pi, \pi]$ is subdivided to have the right number of points.

Note \sinh is an unbounded monotone increasing function.

For $h = 0$, $b \in \mathcal{E}(\mathcal{G})$,

$$\begin{aligned} Z_\Lambda &= \frac{1}{2^{|\Lambda|}} \sum_{\sigma \in \Omega_\Lambda} e^{-\beta H_\Lambda(\sigma)} \\ &= e^{-\beta H(\sigma)} = e^{\sum_b (\sigma\sigma)_b \beta J_b} = \prod_b e^{\beta J_b (\sigma\sigma)_b} && \text{(factorizing over bonds)} \\ &= \frac{1}{2^{|\Lambda|}} \sum_{\sigma \in \Omega_\Lambda} \prod_b [\cosh(\beta J_b) + \sinh(\beta J_b) (\sigma\sigma)_b] && e^x = \cosh x + \sinh x \\ &= \left[\prod_b \cosh(\beta J_b) \right] \frac{1}{2^{|\Lambda|}} \sum_{\sigma \in \Omega_\Lambda} \prod_b [1 + K_b (\sigma\sigma)_b] && \text{where } K_b = \tanh(\beta J_b) \\ &= \left[\prod_b \cosh(\beta J_b) \right] \mathbb{E}_\sigma \sum_{\Gamma \subseteq \mathcal{E}(\mathcal{G})} \prod_{b \in \Gamma} K_b \left(\prod_{b=\{x,y\} \text{ in } \Gamma} (\sigma\sigma)_b \right) \\ &= \left[\prod_b \cosh(\beta J_b) \right] \sum_{\Gamma} \prod_{b \in \Gamma} K_b \mathbb{E}_\sigma \left(\prod_{b=\{x,y\} \text{ in } \Gamma} (\sigma\sigma)_b \right) && \text{linearity of expectation} \end{aligned}$$

We expressed the sum in terms of the symmetric and anti-symmetric parts and pulled out a factor \cosh . The average is easy to compute: if there is a spin for which the power is odd, the average is 0; if the power is even for all spins, the average is 1.

$$\mathbb{E}_\sigma \left(\prod_{b \in \Gamma} (\sigma\sigma)_b \right) = \mathbb{E}_\sigma \prod_x \sigma_x^{\sum_{b \in \Gamma} \mathbb{1}[b \ni x]} = \mathbb{1}[\forall x, (-1)^{\sum_{b \in \Gamma} \mathbb{1}[b \ni x]} = 1]$$

Let $(\partial\Gamma)_2 = \{x : x \text{ belongs to an odd number of edges}\}$. The condition for the summand to not vanish is $(\partial\Gamma)_2 = \emptyset$.

We summarize this calculation as a lemma.

Lemma 1.7.6. *For $h = 0$,*

$$Z_\Lambda = \left[\prod_b \cosh(\beta J_b) \right] \sum_{\Gamma, (\partial\Gamma)_2 = \emptyset} \prod_{b \in \Gamma} K_b.$$

Theorem 1.7.7 (Feynman-Sherman). Let $K(\Lambda) = \prod_{b \in \Gamma} K_b$.

$$\sum_{\Gamma: (\partial\Gamma)_2 = \emptyset} K(\Gamma) = \prod_{\text{loops}} [1 + (-1)^{n(l)} K(l)]$$

where $n(l)$ is the number of self-intersections of l .

Each graph can be covered by a collection of loops (which can cross and repeat edges). Note the RHS is an infinite sum. By a lot of identities this infinite sum reduces to a finite sum.

Kac-Ward introduced the Ihara zeta function.

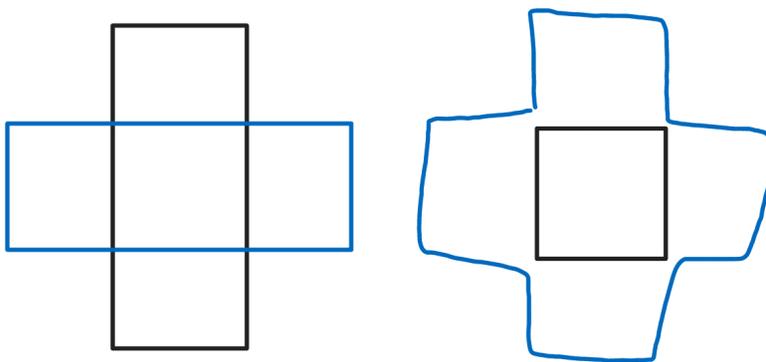
The RHS is

$$\prod_{\text{loops}} [1 + (-1)^{n(l)} K(l)] = \sum_{m: \mathcal{E}(G) \rightarrow \mathbb{Z}_+} \prod_b K_b^{m_b} W(m)$$

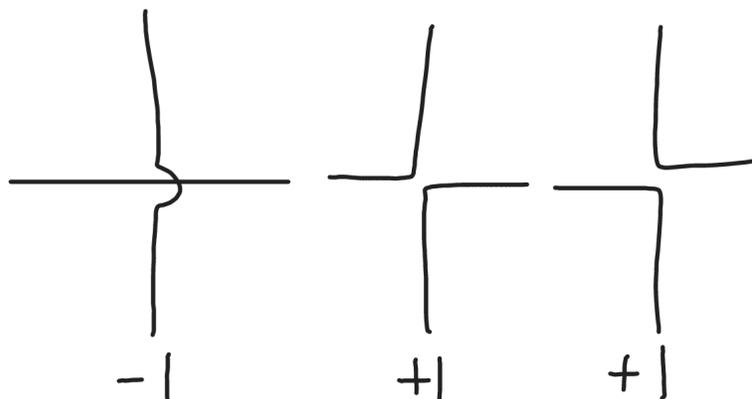
where $W(m) = \sum_{\ell_1, \dots, j} \prod_j (-1)^{n_{\ell_j}} \mathbb{1}_m[\{\ell_1, \dots, \ell_k\}]$.

Lemma 1.7.8. If $\|m\|_\infty := \max_b(m_b) = 1$ then $W(m) = 1$.

Sum over the multisets of edges. We need to enumerate how many ways the multisets can be realized as a union of loops. The Katz-Ward idea is enumerate all the different ways to resolve the traffic jams at each vertex with an intersection: sum over all possible choices with the right signs. The other vertices don't present any choices for the prospective loop.



The three possibilities are:



The parity is certain.

Lemma 1.7.9. *If $\|m\|_\infty = 2$ then $W = 0$.*

Proof. This means there is one edge traversed twice. Take one of these edges. How are the loops arranged in the complement? There are 2 possibilities for how the loops hook up through the bond. Summing you get 0. □

What if it's > 2 ? This is in fact true for any value > 1 .

Why are the higher powers 0?

Theorem 1.7.10 (Sherman).

$$\prod_{\text{loops}} [1 + (-1)^{n(\ell)} K(\ell)]^2 = \det(\mathbb{1} - \mathcal{L})$$

where \mathcal{L} is a $2|\mathcal{E}| \times 2|\mathcal{E}|$ matrix indexed by the oriented edges. Here

$$\mathcal{L}_{e,e'} = \delta_{O(e'),t(e)} e^{\frac{i\angle(e,e')}{2}} K_e.$$

(The Kronecker delta function says that e has to flow into e' , i.e., the terminal of e is the same as the origin of e' . Here $\angle(e, e')$ means the angle between e, e' . We get a phase number that's half the sum of the turns. For a simple loop, we get -1 .)

$\text{Tr } \mathcal{L}^n = \sum_{e_1, \dots, e_n} \mathcal{L}_{e_1, e_n} \cdots \mathcal{L}_{e_2, e_1}$. This is associated to an oriented loop.

Buried in this is an infinite collection of relations. In principle we have contributions in the RHS of the theorem with arbitrarily many loops. But Sherman's theorem says that it is multilinear in the weights of the edges. Higher powers do not contribute.

You can decompose the Hilbert space into Fourier modes. For each value of the momentum you have a finite-dimensional space. Edges are indexed by position and direction. If you change to momenta and direction, the different momenta are not mixed by shifting. The determinant is a block determinant where different values of the momenta are not mixed. We get 4×4 blocks; the determinants can be computed and we get the solution.

3-22

We can use the transfer matrix method for 1D systems with finite range correlations (of distance R), not just nearest-neighbor interactions. Let the states be $\tau_x = \{(\sigma_x, \sigma_{x+1}, \dots, \sigma_{x+R})\}$. Switching to this space, we can now express the energy in terms of nearest neighbor interactions,

$$H = \sum_n \Phi(\tau_n, \tau_{n+1}).$$

7.1 Ihara graph zeta function

Definition 1.7.11: A **path** in a graph \mathcal{G} is a set of edges $(x_1, x_2), \dots, (x_n, x_{n+1}) \in \mathcal{E}$. For an **oriented path**, the direction matters. In a **closed path**, $x_{n+1} = x_1$. An **oriented loop** is a closed oriented path where we don't care about the starting point, i.e., an equivalence class of paths modulo cyclic permutation. Two paths are equivalent here if $y_i = x_{i+h}$, with indices taken modulo n .

Edges are allowed to repeat.

For a path p , denote by $|p|$ the number of edges in p .

We would like to eliminate paths which are powers of more elementary paths, complete repetitions of a primitive path. A path is primitive if it is not a power of a shorter path.

Definition 1.7.12: A path is **primitive** if it cannot be written as q^k , $k > 1$, where q is a closed path.

This may remind you of prime numbers. Any number can be written as a product of numbers which cannot be decomposed. People applied number-theoretic ideas to graphs.

Theorem 1.7.13. *Let M be a matrix indexed by sites of a graph \mathcal{G} . For $\|M\|$ small enough,*

$$\det(\mathbb{1} - M) = \prod_{p^*} [1 - \chi_p(p)]$$

where the product is over equivalence classes p^* of primitive oriented loops.

Here

$$\chi_p = \prod_{n=1}^{|p|} M_{x_n, x_{n+1}}.$$

On the RHS, when we open the brackets, we can worry about convergence. For matrices of small norm, it is convergent. To define the function for all λ , note the theorem establishes that for small λ ,

$$\det(\mathbb{1} - \lambda M) = \prod_{p^*} [1 - \lambda \chi_p(p)].$$

Now we can take the analytic continuation.

Note the LHS is a finite calculation (and a multilinear function in entries of M), while the RHS is an infinite product/sum.

It follows from the proof that the RHS is a convergent product (for $\|M\|$ small enough).

Proof. First assume $\|M\| < 1$. (This is the operator norm, $\|M\| = \max_v \frac{\|Mv\|}{\|v\|}$.) Then

$$\ln \det(\mathbb{1} - M) = \operatorname{tr} \ln(\mathbb{1} - M).$$

For Hermitian matrices, prove this by diagonalization: the LHS is $\ln \prod_j (1 - \lambda_j) = \sum_j \ln(1 - \lambda_j)$.

For the general case, we can still use this method. Write $M = S^{-1}TS$ where T is upper triangular. If the diagonal entries are $\lambda_1, \dots, \lambda_n$, the determinant is still $\prod_j (1 - \lambda_j)$.

Expanding gives

$$\begin{aligned} \ln \det(\mathbb{1} - M) &= - \sum_{n=1}^{\infty} \frac{1}{n} \operatorname{tr}(M^n) \\ &= \sum_{x=(x_0, \dots, x_n)} \frac{1}{n} \left(\prod_{j=1}^n M_{x_{j+1}, x_j} \right) \\ &= - \sum_{k=1}^{\infty} \sum_{p^*} \frac{1}{k|p^*|} \chi_M(p^*)^k \\ &= - \sum_{p^*} \ln(1 - \chi_M(p^*)). \end{aligned}$$

k comes from properly weighing the volume of the path. □

We want to avoid backtracking.

Theorem (Theorem 1.7.10). *Let \mathcal{E} be the graph of oriented edges of a planar graph. Let*

$$\mathcal{L}_{e,e'} = \mathbb{1}[O(e') = t(e)] K_e e^{\frac{i\angle(e',e)}{2}}$$

For any oriented closed path,

$$\prod_j \mathcal{L}_{e_j, e_{j+1}} = \prod_i K_e (-1)^{w(p)}$$

with

$$(-1)^{w(p)n} = -(-1)^{n(p)}.$$

where $w(p)$ is the winding number and $n(p)$ is the number of crossings.

We have

$$\begin{aligned} \prod_{p^*} [1 - \chi_{\mathcal{L}}(p^*)] &= \sum_{\mathcal{P}=(p_1^*, \dots, p_n^*)} \prod_j (-1)^{n(p_j)} k[p] \\ &= \sum_m \pi_e K_e^{m_e} W_m \end{aligned}$$

where W_m is the number of paths with a given frequency of edges traversed.

Lemma 1.7.14 (Kac-Ward). *We have*

$$\sum_{P: \max_e(m(P))=1} (-1)^{n(P)} K(P) = \sum_{\Gamma: \partial\Gamma=\phi} \prod_{e \in \Gamma} K_e = Z(\beta)$$

for $K_e = \tanh(\beta J_e)$.

Each vertex has 3 possibilities.

3/24/16

Introduce $t = \frac{\beta - \beta_c}{\beta_c}$. What momenta contribute to the rise in the specific heat? That which E is on the order of $Y + Y^{-1} - 2$. Do some dimensional analysis.

The specific heat is

$$C = \frac{\partial}{\partial T} \langle H \rangle_\beta = \beta^2 \frac{\partial^2}{\partial \beta^2} \Psi = \dots \ln |\beta - \beta_c| + \dots$$

Certain fermionic structures pop up. We start from a classical system and out of that the excitations take the form of quantum objects, like quantum spinors (B Kaufman and Lee Scholz). This employed the 2D transfer matrix. When you try to solve it, you transfer from 1 layer to another; the dimension of the matrix is linearly proportional to n . In 2D, for reasons to do with integrable systems in 1D, the transfer matrix can be diagonalized.

The loops we generate can have 1 edge multiple times. We have to disentangle the system of graphs which correspond to configurations of loops.

In a high dimension, loops typically miss each other. Perhaps Onsager's solutions may emerge in higher dimensions. (Dimension 3 is the hardest.)

Weight each loop by the product we want.

3-29-16

8 Gibbs states in the infinite volume limit

Let S be the set of spin states. S is typically a compact space. For example, in the Ising/Potts spin models, $S = \{1, \dots, Q\} \subset \mathbb{N}$. In the $O(n)$ spin models, $S = \{\sigma \in \mathbb{R}^n : |\sigma| = 1\}$.²¹ Let \mathcal{G} be a transitive graph, $\Lambda \subseteq \mathcal{G}$. (Often, $|\Lambda| < \infty$.)

The configuration space is $\Omega_\Lambda = S^\Lambda$. Each $\omega \in \Omega_\Lambda$ is a map $\omega : \Lambda \rightarrow S$. Denote it by $\omega = \{\sigma_x\}_{x \in \Lambda}$.

The configuration space of the infinite system is

$$\Omega = S^{\mathcal{G}},$$

i.e., $\omega \in \Omega$ is a map $\mathcal{G} \rightarrow S$. Denote it by $\omega = \{\sigma_x\}_{x \in \mathcal{G}}$.

²¹Note that physicists say that the sphere in \mathbb{R}^n is n -dimensional, while mathematicians say it is $(n - 1)$ -dimensional.

Recall the simplest non-finite probability space $[0, 1]$. Describing a point by its dyadic representation

$$x = \sum_{j=1}^{\infty} \frac{1}{2^j} \sigma_j.$$

Represent it by

$$x = \{\sigma_j\}_{j=1}^{\infty},$$

an infinite sequence of binary variables. Similarly, points in the unit square $[0, 1]^2$ correspond to doubly infinite sequences of binary variables.

Since $[0, 1]$ is a measure space, we can use the mapping above to make the space of infinite binary sequences a measure space. How does the notion of convergence translate? We have $x_n \rightarrow x$ as $n \rightarrow \infty$ iff for all $k < \infty$, the sequence $\{\sigma_j\}_{j=1}^k$ stabilizes.

Theorem 1.8.1 (Tychonoff (Tikhonov)). *The product of any collection of compact topological spaces is compact with respect to the product topology.*

We shall typically deal with countable products.

For countable products, the topology is characterized by sequential convergence. (Otherwise, we have use nets.) Here, $\omega_n \rightarrow \omega$ in $\Omega = S^{\mathcal{G}}$ iff $\omega_n|_{\Lambda} \rightarrow \omega|_{\Lambda}$ for any finite $\Lambda \subset \mathcal{G}$ (each finite projection converges).

Our spaces are also metrizable. Note a countable product of metrizable spaces is metrizable. We can define the distance by

$$d(\sigma, \sigma') = \sum_x \frac{1}{2^n} d(\sigma_x, \sigma'_x)$$

if the distance is bounded. If it is not bounded, we normalize the distance first,

$$d(\sigma, \sigma') = \sum_x \frac{1}{2^n} \frac{d(\sigma_x, \sigma'_x)}{1 + d(\sigma_x, \sigma'_x)}.$$

Theorem 1.8.2 (Riesz-Markov (Riesz representation)). *Let Ω be a compact metric space, $C(\Omega, \mathbb{C})$ be the space of linear functionals $\rho : C(\Omega) \rightarrow \mathbb{C}$ such that*

1. $|\rho(f)| \leq \|f\|_{\infty}$,
2. for all $f \geq 0$, $\rho(f) > 0$,
3. $\rho(1) = 1$.

Then there exists a (Borel) probability measure μ on (Ω, Σ) such that $\rho(f) = \int_{\Omega} f(\omega) d\mu(\omega)$.

Definition 1.8.3: A **state** of an infinite system is a (Borel) probability measure on its configuration space Ω .

The state of a system is not a configuration on it, but a probability measure. (In quantum mechanics we can't always specify configurations. We have observables and expected values of operators; that's as much as we can ask for in describing the state of a system. We now use this not just in quantum mechanics.)

The theorem says we can capture a probability measure by giving the probabilities of various events, or by giving the expected value of functions. It's an analogy of von Neumann's Theorem, characterizing quantum states in terms of functionals.

See Barry Simon's book.

In finite volumes, equilibrium states are measures of the form

$$\rho_{\Lambda,\beta}(d\sigma) = \frac{e^{-\beta H_{\Lambda}^{(b.c.)}(\sigma_{\Lambda})} \rho_0(d\sigma_{\Lambda})}{Z^{(b.c.)}}.$$

Why can't we simply take the infinite-volume limit? The energy of the system is the sum of an infinite number of terms. The typical order of magnitude is exponential in the size of the system. When you take the system to ∞ , we get $\frac{0}{0}$ or $\frac{\infty}{\infty}$.

What survives out of this formula for conditional distributions. It is characterized by saying how the probability of a configuration changes under local changes. If Λ is a finite subset, conditioned on $G_{\Lambda \setminus \Lambda_0}$, we have a probability measure on G_{Λ} . This makes sense even in the infinite-volume limit.

This formula is due to Dobrushin, Lonford, and Ruelle (DLR).

$$\rho(\sigma_{\Lambda_0} | \sigma_{\Lambda \setminus \Lambda_0}) = \frac{e^{-\beta H_{\Lambda_0}(\sigma_{\Lambda_0} | \sigma_{\Lambda \setminus \Lambda_0})} \rho_0(d\sigma_{\Lambda_0})}{Z_{\Lambda_0}^{\sigma_{\Lambda \setminus \Lambda_0}}} \quad (1.22)$$

where $H_{\Lambda}(\sigma_{\Lambda_0} | \sigma_{\Lambda \setminus \Lambda_0}) = \sum_{A \subset \Lambda, A \cap \Lambda_1 \neq \emptyset} J_A \phi_A(\sigma_A)$. The outside provides boundary conditions, such that the conditional distribution is given by the Gibbs equilibrium formula.

Definition 1.8.4: A **Gibbs equilibrium state** of an infinite system is a probability measure on (Ω, Σ) for which the finite volume conditional expectations are given by the DLR formula (1.22).

Definition 1.8.5: For each finite Λ , a function $f : \Omega \rightarrow \mathbb{C}$ is **measurable in Λ** if $f(\sigma) = F_{\Lambda}(\sigma_{\Lambda})$ with $F_{\Lambda} : \Omega_{\Lambda} \rightarrow \mathbb{C}$ measurable.

Let \mathcal{B}_{Λ} denote the space of such functions,

$$\Sigma_{\Lambda} = \{\alpha \in \Omega : \mathbb{1}_{\alpha} \in \mathcal{B}_{\Lambda}\}$$

and let

$$\Sigma = \overline{\bigcup_{\Lambda \subset \mathcal{G}, |\Lambda| < \infty} \Sigma_{\Lambda}}.$$

Define the space of functions **measurable at ∞** to be

$$\mathcal{B}_{\infty} = \bigcap_{|\Lambda| < \infty} \mathcal{B}_{\mathcal{G} \setminus \Lambda}.$$

Example 1.8.6: For $x = \{\sigma_j\}_{j=1}^\infty$, $\overline{\lim}_{L \rightarrow \infty} \frac{1}{L} \sum_{n=1}^L \sigma_n$ is a function that is measurable at ∞ .

For example, if σ_j are independent Bernoulli($\frac{1}{2}$) variables, this function is a.s. $\frac{1}{2}$ by the Law of Large Numbers. If σ_j are independent Bernoulli(p) variables

$$\begin{aligned}\mu_p(\{\sigma_n = 1\}) &= p \\ \mu_p(\{\sigma_n = 0\}) &= 1 - p,\end{aligned}$$

it is a.s. p . These measures are orthogonal (mutually singular): there is an event with probability 1 in one measure and 0 in the other (look at the value of the limsup).

In the infinite volume limit, we can ask whether what used to be the boundary values is still detectable deep inside.

We are naturally led to questions of uniqueness of Gibbs measures. A discontinuity in the derivative translates to nonuniqueness of Gibbs measures.

We already saw one proof of phase transitions. It's a simple exercise to use the Piesch argument to prove that the ferromagnetic Gibbs measure at low temperature is not unique.

Measurability at ∞ is a nice concept. Using results at probability theory, we can show in 1 dimension there is no phase transition. We will show that a sufficient condition for uniqueness of Gibbs state is that the total interaction to the left and right is uniformly bounded. We will show a 0-1 law.

3/31/16: We are graduating from free energy to states of infinite systems.

8.1 Conditional expectation

A point $\sigma \in \Omega$ is given by $\sigma = \{\sigma_x\}_{x \in \mathcal{G}}$.

The conditional probability, conditioned on an event B is $\mathbb{P}(A|B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}$. We can define probability conditioned on a σ -algebra,

$$\mathbb{P}(A|\Sigma_0)(\sigma).$$

Thinking of Σ_0 as a partition (e.g. when it's finite), it's the probability of A given the element of the partition Σ_0 that σ is in. We now define this more generally and formally.

Proposition 1.8.7 (Regular conditional expectation value): Let (Ω, Σ, μ) be a probability space and $\Sigma_0 \subseteq \Sigma$ (a sub σ -algebra). Then there is a unique map

$$\mathbb{E}_{\Sigma_0} : L^\infty(\Omega, \Sigma) \rightarrow L^\infty(\Omega, \Sigma_0)$$

such that for all $f \in L^\infty(\Omega, \Sigma)$ and $g \in L^\infty(\Omega, \Sigma_0)$,

$$\text{eq:cond-exp} \int fg d\mu(\sigma) = \int g(\sigma) \mathbb{E}_{\Sigma_0}(f) \mu(d\sigma). \quad (1.23)$$

Think of sub- σ -algebras as generalizations of partitions. Suppose for the moment that Σ_0 is a partition (the σ -algebra breaks up into atoms). What do I know about a point

given partial information about it, namely, in which set of the partition it is in? I can get a conditional probability distribution by restricting to that set. To say that $g \in \mathcal{B}_{\Sigma_0}$ means that g is constant on elements of the partition. The function $\mathbb{E}(f|\Sigma_0)(\sigma)$ is the mean value of f restricted to the partition Σ_0 . If $g \in \mathcal{B}_{\Sigma_0}$, to find the expectation of fg , (1.23) says that you can either integrate it directly, or take the average of f on each set of the partition first and then integrate.

Furthermore, \mathbb{E}_{Σ_0} acts as an orthogonal projection onto $L^2(\Omega, \Sigma_0)$.

(In general things are more complicated. The σ -algebra corresponding to specifying x is a vertical line. There is an uncountable number of “atoms” each of zero measure. Here measurable means only depending on x . The discussion on partitions is for primarily for intuition.)

Recall that a Gibbs equilibrium state in a finite volume Λ with boundary conditions is a probability measure on Ω_Λ of the form $\frac{e^{-\beta H_\Lambda^{(b.c.)}(\sigma_\Lambda)}}{Z_\Lambda^{(b.c.)}} \rho_0(d\sigma_\Lambda)$.

Let $\Lambda_0 \subseteq \Lambda$ and assume that the boundary condition term in H does not involve $\sigma_{\Lambda_0} = \sigma|_{\Lambda_0}$. Denote

$$\mathbb{E}(f|\Sigma_{\Lambda \setminus \Lambda_0})(\sigma) = \mathbb{E}(f|\sigma_{\Lambda_0^c}).$$

Then (the DLR formula)

$$\text{eq:dlr2} \mathbb{E}(f|\sigma_{\Lambda_0^c}) = \int f(\sigma_{\Lambda_0}, \sigma_{\Lambda_0^c}) \frac{e^{-\beta H_\Lambda(\sigma_{\Lambda_0}|\sigma_{\Lambda_0^c})}}{Z} \rho_0(d\sigma_{\Lambda_0}). \quad (1.24)$$

This formula captures that for some equilibrium state in a huge volume, the local behavior is given by a canonical expression. In the expression the configuration in Λ_0^c is frozen; we integrate over the spins inside. This formula makes sense even if Λ is infinite. This allows us to define Gibbs states for the infinite volume limit.

Definition 1.8.8: A probability measure μ on (Ω, Σ) is a **Gibbs state** for a Hamiltonian H iff for any finite region Λ_0 , $\mathbb{E}(f|\sigma_{\Lambda_0^c})$ is given by the DLR formula (1.24).

The Metropolis-Hastings algorithms gives a way to compute expectations by running dynamical system. Pick a site at random and redraw it with right conditional probability distribution. This is one way to numerically compute conditional expectation values. The Gibbs state is stationary under these dynamics.

Theorem 1.8.9. Let $\Lambda_n \nearrow \mathcal{G}$, (e.g. $\mathcal{G} = \mathbb{Z}^d$) and μ_{Λ_n} be a sequence of finite volume

$$\mu_{\Lambda_n} = \frac{e^{-\beta H_{\Lambda_n}^{(b.c.)}(\sigma_{\Lambda_n})}}{Z} \rho_0(d\sigma_{\Lambda_n}).$$

If the weak limit

$$w \lim_{n \rightarrow \infty} \mu_{\Lambda_n} = \mu$$

exists then the limiting measure is a Gibbs state.

Definition 1.8.10: We say $w \lim_{n \rightarrow \infty} \mu_{\Lambda_n} = \mu$ if

$$\lim_{n \rightarrow \infty} \int f d\mu_n(d\sigma)$$

exists for all local families, and this equals $\int f \mu(d\sigma)$.

For finite systems you can ask questions like “what is the boundary”? In the infinite limit, it doesn’t make sense to ask such questions.

The theorem follows from the DLR formula.

Lemma 1.8.11. *The collection of Gibbs states is closed under convex combinations. I.e., if μ_0 are Gibbs states for a common Hamiltonian, then so is*

1. $\mu = (1 - t)\mu_0 + t\mu_1$, $0 \leq t \leq 1$.
2. When ν is a probability measure ($\int \nu(d\alpha) = 1$) $\mu(-) = \int \mu_\alpha(-) \nu(d\alpha)$.

Definition 1.8.12: A Gibbs state is said to be a **pure** Gibbs measure if it does not admit a decomposition as a convex combination of distinct Gibbs states.

For example, consider the Ising model at $h = 0$, $\beta > \beta_c$. We can take + boundary conditions, – boundary conditions, or free/periodic boundary conditions. We have that the Gibbs state with free/periodic boundary conditions is not a pure state:

$$\mu = \lim_{\Lambda_n \nearrow \mathbb{Z}^d} \mu_{\Lambda_n}^{(\text{per})} = \frac{1}{2}\mu^{(+)} + \frac{1}{2}\mu^{(-)}.$$

What if we mix +’s and –’s? The state of Gibbs states is large.

8.2 Symmetry and symmetry breaking

Definition 1.8.13: A **symmetry** of a statistical mechanical systems is an invertible (measurable) mapping $T : \Omega \rightarrow \Omega$ which preserves ρ_0 and H in the following sense:

1. for all $A \in \Sigma$, $\rho_0(T^{-1}A) = \rho_0(A)$.
2. $H(T\sigma) = H(\sigma)$, $\phi_\Lambda(T\sigma) = \phi_\Lambda(\sigma)$.

Definition 1.8.14: A Gibbs state is said to exhibit **simple symmetry breaking** if μ is not invariant under one of the system’s symmetries.

Example 1.8.15: The Ising model at $h = 0$ is symmetric with respect to $T(\sigma) := -\sigma$ because $\rho_n(d\sigma)$ and $H = -\sum J_{x,y} \sigma_x \sigma_y$ are both invariant under $T(\sigma) := -\sigma$.

All spins get equal weight makes sense for a finite system. What's the meaning for an infinite collection, when all configurations have 0 probability? We mean that for all f local (depending on a finite number of spins),

$$\int f(\sigma)\rho_0(d\sigma) = \int f(-\sigma)\rho_0(d\sigma).$$

We can talk about expectation values of functions, rather than probabilities of individual configurations.

Consider taking the limit

$$\rho_0(d\sigma) \mapsto \lim_{\Lambda \nearrow \mathbb{Z}^d} \frac{e^{-\beta H(\sigma)}}{Z} \rho_0(d\sigma).$$

Before the limit, if we have free boundary conditions, this would be symmetric. If you put boundary conditions, this would not be symmetric. Does the asymmetry persist in the limit? It can't disappear; the limit is characterized by the local DLR equation.

4-5-16

We define symmetry breaking slightly differently.

Definition 1.8.16: A **symmetry** is a measurable map $T : \Omega \rightarrow \Omega$ that satisfies the following.

1. It is invertible and preserves $\rho_0(d\sigma)$.
2. $H(T\sigma) = H(\sigma)$ in the sense that $\phi_A(T\sigma) = \phi_A(\sigma)$ for all $A \subseteq \mathcal{G}$.

Definition 1.8.17: A Gibbs state μ is said to exhibit symmetry breaking if μ is decomposable into a linear combination of Gibbs state of which a positive fraction fail to be invariant under one of the system's symmetries.

Consider the 2D Ising model at low temperature with $h = 0$ and free boundary conditions. In this model, $H = -\sum J\sigma_x\sigma_y$, and the symmetry is global spin flip, $(T\sigma)_\Lambda = -\sigma_\Lambda$.

Theorem 1.8.18 (Symmetry breaking). *The 2D Ising model with $h = 0$ with free boundary conditions exhibits symmetry breaking for all β such that $3e^{-2\beta} < 1$.*

Recall we saw that $\mathbb{P}^+(\sigma_0 = -1) \leq \frac{4}{3} \sum_{l \geq 4} (3e^{-2\beta})^l$. For each closed path γ ,

$$\mathbb{P}(\gamma \text{ occurs among the contours of } 0) \leq e^{-2\beta|\gamma|}.$$

The probability there exists a contour within the annular region of outer radius R and inner radius r is $\leq \sum_{4r \leq l} (3e^{-2\beta})^l$. This is the tail of a convergent series which can be made as small as desired. By picking r large enough, we do not find any contour enclosing that region. There must be an infinite cluster of the same spins.

There are 2 types of configurations in the infinite volume limit: there is an infinite cluster of $+$, or an infinite cluster of $-$.

Let $C_+(R) = \{x \in \mathbb{Z}^2 : \dots\}$. (to be filled in)

At low temperatures the Gibbs states show agreement among neighboring spins. The predominant behavior is a sponge that percolates to infinity (an infinite connected component) where all spins have the same sign; there may be islands within the cluster. There are many paths (exponentially many). If the energy penalty beats the entropy, then the probability of seeing a lone line goes to 0. Any finite configuration will be realized somewhere, so you have to localize the bounds.

Taking the Gibbs state with + boundary conditions, defined as

$$\mu_\beta^+ = \lim_{\Lambda \nearrow \mathbb{Z}^d} \mu_\Lambda^{(+)}$$

We can prove $\mu^+ \neq \mu^-$ by showing $\mu^{(+)} \neq \mu^{(-)}$ for a well-chosen event A .

Theorem 1.8.19. *A sufficient condition for H to have a unique Gibbs state of given β is that for any pair of Gibbs equilibrium measures, there exists $B < \infty$ such that $\mu_1(f) \leq B\mu_2(f)$ for all positive functions $f \in C(\Omega, \mathbb{R})$.*

Theorem 1.8.20. *For 1D systems, a sufficient condition for uniqueness of the Gibbs state is that for all x ,*

$$\sum_{\substack{A \subset \mathbb{Z}, A \cap [x, \infty) \neq \emptyset \\ A \cap (-\infty, x] \neq \emptyset}} J_A \|\phi_A\| \leq B < \infty.$$

For example, for the Ising model this is satisfied when $\sum_{\substack{x \geq 0 \\ y < 0}} |J_{x-y}| < \infty$ or equivalently, $\sum_{u \geq 0} |J_u| < \infty$.

Here is a continuous version of the theorem.

Theorem 1.8.21 (Merrin-Wagner). *Translation invariant systems in $D \leq 2$ dimensions with finite-range interactions and continuous (compact) symmetry do not exhibit symmetry breaking.*

For example consider $O(N)$, with $H = -\sum J_{x-y} \langle \sigma_x, \sigma_y \rangle$, $S = \{\sigma \in \mathbb{R}^n : |\sigma| = 1\}$. Now the system is invariant under global rotation. This symmetry is not broken if you only have finite-range interactions. We use the fact that the symmetry group here is compact.

We have

$$\mu(-) = \int \mu(-|\mathcal{B}_\infty)(\eta) \mu(d\eta).$$

Note that for each $\eta \in \Omega$,

$$\mu(-|\mathcal{B}_\infty)(\eta)$$

is a Gibbs state.

Recall DLR: For $f \in \mathcal{B}_{\Lambda_0}$,

$$\mu(f) = \int \mu(f|\mathcal{B}_{\Lambda_0^c})(\eta) \mu(d\eta)$$

where $\mu(f|\mathcal{B}_{\Lambda_0}^c)$ are given by DLR. (We have $\mu(f) = \int \mu(f|\mathcal{B})(\sigma) \mu(d\eta)$ for any σ -algebra \mathcal{B} .)

The formula says that the functional (measure) μ can be decomposed as an integral of functionals (measure).

In the limit $\mathcal{B}_\infty = \bigcap_{|\Lambda_0| < \infty} \mathcal{B}_{\Lambda_0}^c$. “Does the configuration have a limit for the magnetization and is it nonzero?” is a valid question.

The conditional expectation is a Gibbs measure in any finite volume so is a Gibbs measure.

Any Gibbs measure can be decomposed into a collection of extremal Gibbs measures so that the σ -algebra at ∞ is trivial with respect to each of those. The Gibbs measures are mutually singular.

Claim 1.8.22. For almost every pair $\eta_1, \eta_2 \in \Omega$,

$$\mu(-|\mathcal{B}_0)(\eta_1), \quad \mu(-|\mathcal{B}_0)(\eta_2)$$

are either equal or mutually singular.

Mutually singular means that their supports are disjoint.

Proof. For all $g \in \mathcal{B}_\infty$,

$$\mu(g|\mathcal{B}_\infty)(\eta) = g(\eta).$$

$\mu(g|\mathcal{B}_\infty)$ is the conditional mean of g when only this information about η ? If g is measurable with respect to \mathcal{B}_∞ , it's just $g(\eta)$.

“Is there a $+/-$ measure” is measurable with respect to \mathcal{B}_∞ . □

Take the $O(N)$ model with boundary spins oriented the same. Does it exhibit symmetry breaking?

Symmetry breaking occurs even when you have a state that is itself symmetric. For example, $\frac{1}{2}(\mu^{(+)} + \mu^{(-)})$.

Consider a functional symmetric under rotation.

Symmetric breaking is associated with the picture of the Mexican hat, with minimum attained on a ring. These points are not symmetric.

What if we rotate all boundary spins to another value? It is ferromagnetic so there is a bias towards the condition prescribed by the boundary. We ask how these compare, $\mu^{(\nearrow)}(f(\sigma))$ and $\mu^{(\nearrow)}(f(R\sigma))$ where R is a rotation. The second is equal to $\mu^{(R\searrow)}(f(\sigma))$.

What if you gradually untwist the rotation, doing it gradually throughout the system? Do soft rotations.

8.3 Phase transitions

4-7-16

What may cause a phase transitions? There are other reasons besides symmetry breaking.

Consider the 1D Ising model $H = -\sum J_{x,y} \sigma_x \sigma_y$.

Theorem 1.8.23. For the 1D Ising model with $\sum_{\substack{x \geq u \\ y < u}} |J_{x,y}| \leq B < \infty$ for all u , the Gibbs state is unique for all $\beta < \infty$.

For example, if $J_{x,y} = \frac{1}{|x-y|^2}$, there is no first order phase transition if $\tau > 2$.

To see this, index the sum by the difference $v = x - y$ and note each difference appears v times.

$$\sum_{\substack{x \geq 0 \\ y < 0}} |J_{x-y}| = \sum_{v \geq 0} v |J_v|.$$

Proof. It suffices to show that there exists C such that for all μ_1, μ_2 and $f \in \mathcal{B}_0$ with $f \geq 0$,

$$\mu_1(f) \leq D\mu_2(f).$$

The basic idea is to use conditional expectation. By DLR,

$$\mu(f) = \int \mathbb{E}_{[u,v]^c}(f|\sigma) \mu(d\sigma) \quad (1.25)$$

$$\mathbb{E}_{\Lambda^c}(f|\sigma) = \int f(\eta_\Lambda, \sigma_{\Lambda^c}) \frac{e^{-\beta H_\Lambda(\eta_\Lambda, \sigma_{\Lambda^c})}}{Z_\Lambda^{(\sigma_{\Lambda^c})}} \rho_0(d\eta_\Lambda). \quad (1.26)$$

Included in $e^{-\beta H_\Lambda(\eta_\Lambda, \sigma_{\Lambda^c})}$ include interactions between sites inside the finite region and between sites in the finite region and the outside. Not included is interactions between outside sites.

Now

$$|H_\Lambda^{\text{b.c.}}(\eta_\Lambda, \sigma_{\Lambda^c}) - H_\Lambda^{\circ}(\eta_\Lambda)| \leq 2B.$$

Let f.b.c. denote “free boundary conditions.” We conclude that

$$\mathbb{E}_{\Lambda^c}(f|\sigma) \leq \mathbb{E}_{\Lambda^c}^{\text{f.b.c.}}(f) e^{4B} \quad (1.27)$$

$$\mathbb{E}_{\Lambda^c}(f|\sigma) \geq \mathbb{E}_{\Lambda^c}^{\text{f.b.c.}}(f) e^{-4B} \quad (1.28)$$

(Note we use the fact that f is positive.)

Hence

$$\mu_j(f) = \int \mathbb{E}_{\Lambda^c} \mathbb{E}_{\Lambda^c}(f|\sigma) \leq \mu_j(f) \quad (1.29)$$

$$e^{-4B} \mathbb{E}_{\Lambda^c}^{\text{f.b.c.}}(f) \leq \mu_j(f) \leq e^{4B} \mathbb{E}_N^{\text{f.b.c.}}(f). \quad (1.30)$$

This proves that there cannot be 2 mutually singular Gibbs states. \square

For the $O(N)$ model in 2D, we show there is also a unique Gibbs state by showing that any extremal state is rotation invariant. Letting R_θ denote rotation by θ and defining its action on functions and measures as

$$(Rf)(\sigma) = f(R\sigma) \quad (1.31)$$

$$(R\mu)(f) = \mu(Rf). \quad (1.32)$$

We will show that that in the 2D model with nearest neighbor interactions, for all $f \in B$, $\mu(Rf) \leq C\mu(f)$.

Every extremal Gibbs state is actually equal to its rotation, so it is rotationally invariant. Thus there is no symmetry breaking.

In 1D we disconnected the coupling. We can't do this in 2D. The reason is that. Removing couplings across the boundary costs $L^{d-1} = L$, which is not uniformly bounded. However, for rotational symmetry, consider the effect of gradual rotation. If we rotate spins uniformly in the volume, we still pay the price of discontinuity. But if we rotate at a linear rate across space, the effect on the energy is smaller.

$$\sum_{\substack{x, y \in \Lambda_L \setminus \Lambda_{L/2} \\ |x-y|=1}} ((R_{\theta(x)}\underline{\sigma}_x) \cdot R_{\theta(x)}(\underline{\sigma}_y)) - (\underline{\sigma}_x, \underline{\sigma}_y).$$

Here is the heuristic idea. The quadratic term introduces a correction term which behaves as $\frac{1}{L^2}$. The quadratic term is under control. Next time I'll consider the linear term.

4/12

It suffices to prove in 2D there exists D such that for all Gibbs states μ ,

$$\mu(R_\alpha f) \leq D\mu(f)$$

for all positive $f \in \mathcal{B}_0$. Note that the measure $\tilde{\mu}(f) = \mu(R_\alpha f)$ is also a Gibbs state.

To see this, consider a finite volume Gibbs state

$$\mu_\Lambda(f) = \int f(\sigma_\Lambda) \frac{e^{-\beta H_\Lambda^{\text{b.c.}}(\sigma_\Lambda)}}{Z_\Lambda^{\text{b.c.}}} \rho_0(d\sigma_\Lambda).$$

Applying a rotation to this gives

$$\tilde{\mu}(f) := \mu(R_\alpha f) = \int f(\underbrace{R_\alpha \sigma_\Lambda}_\eta) \frac{e^{-\beta H_\Lambda^{\text{b.c.}}(\sigma_\Lambda)}}{Z_\Lambda^{\text{b.c.}}} \rho_0(d\sigma_\Lambda) \quad (1.33)$$

$$= \int f(\eta) \frac{e^{-\beta H_\Lambda^{R_\alpha^{-1}(\text{b.c.})}(R_\alpha^{-1}\eta)}}{Z_\Lambda^{\text{b.c.}}} \rho_0(dR_\alpha^{-1}\eta) \quad (1.34)$$

$$= \mu^{R_\alpha(\text{b.c.})}(f). \quad (1.35)$$

(The boundary conditions are rotated by $-\alpha$.) The DLR condition is invariant under rotations (exercise).

See lemma 11.5 in notes. Change D to C and $<$ to \geq in Lemma 11.5.

Proof. Let

$$\theta(x) = \pi \begin{cases} 1, & \|x\| \leq L_0 \\ \frac{2L - \|x\|_\infty}{L}, & L_0 \leq \|x\| \leq 2L_0 \\ 0, & \|x\|_2 \geq 2L_0. \end{cases}$$

(We rotate the configuration in such a way so that the rotation does not extend to ∞ .)

We will consider the rotation $\widehat{R}_{\frac{\alpha}{\pi}\theta}$ (we use a hat to emphasize that the rotation depends on x). In the box $\|x\| \leq L_0$, $\widehat{R}_{\frac{\alpha}{\pi}\theta}$ acts as uniform rotation:

$$\mu(R_\alpha f) = \mu(f(\widehat{R}_{\frac{\alpha}{\pi}\theta}\sigma)) \quad (1.36)$$

$$= \int \left[\int_{\Omega_\Lambda(2L_0)} f(\underbrace{\widehat{R}_{\frac{\alpha}{\pi}\theta}\sigma_\Lambda}_{\sigma'_\Lambda}) \frac{e^{-\beta H_\Lambda(\sigma'_\Lambda|\sigma_{\Lambda^c})}}{Z_\Lambda(\sigma_{\Lambda^c})} \rho_0(d\sigma'_\Lambda) \right] \rho(d\sigma_{\Lambda^c}) \quad (1.37)$$

$$= \int \left[\int_{\Omega_\Lambda(2L_0)} f(\sigma'_\Lambda) \frac{e^{-\beta H_\Lambda(\widehat{R}_{\frac{\alpha}{\pi}\theta}\sigma'_\Lambda|\sigma_{\Lambda^c})}}{Z_\Lambda(\sigma_{\Lambda^c})} \rho_0(d\sigma'_\Lambda) \right] \rho(d\sigma_{\Lambda^c}) \quad (1.38)$$

When you carry out a gradual rotation, the energy does not change by a lot. We have

$$\mu(R_\alpha f) = \mu(\widehat{R}_{\frac{\alpha}{\pi}\theta} f) \quad (1.39)$$

$$= \mu(f e^{-\beta[H(\widehat{R}_{\frac{\alpha}{\pi}\theta}\sigma) - H(\sigma)]}). \quad (1.40)$$

Formally, we rewrote

$$f(0) e^{-\beta H(\widehat{R}^{-1}\sigma)} \frac{\rho_0(d\sigma)}{Z} = e^{-\beta[R^{-1}H - H]} e^{-\beta H(\sigma)} \frac{\rho_0(d\sigma)}{Z}$$

Note $R^{-1}H - H$ is finite; what is left is to estimate it. We estimate it in the case $H = -\sum_{|x-y|=1} \underline{\sigma}_x \cdot \underline{\sigma}_y$. We calculate

$$H(R^{-1}\sigma) - H(\sigma) = - \sum_{\text{nearest neighbor}} [(R_{\theta(x)}^{-1}\sigma_x) \cdot (R_{\theta(y)}^{-1}\sigma_y) - \sigma_x \cdot \sigma_y] \quad (1.41)$$

$$= \sum_{\text{nearest neighbor}} [\sigma_x \cdot \sigma_y - \sigma_x \cdot R_{\theta(x)-\theta(y)}\sigma_y] \quad (1.42)$$

$$= \sum_{\text{nearest neighbor}} [\alpha(\sigma_x \cdot \mathcal{L}\sigma_y)] \quad (1.43)$$

$$= \sum_{\text{nearest neighbor}} \left[\frac{\alpha}{\pi} \delta_1(H) + \delta_2(H) \right], \quad (1.44)$$

where δ_i is the i th order correction, from the Tylor expansion. Then

$$\mu(R_\alpha f) = \mu(\widehat{R}_{\frac{\alpha}{\pi}\theta} f) \quad (1.45)$$

$$= \mu(f e^{-\beta[H(\widehat{R}_{\frac{\alpha}{\pi}\theta}^{-1}\sigma) - H(\sigma)]}) \quad (1.46)$$

$$= \mu(f e^{[\frac{\alpha}{\pi}\delta_1(H) + \delta_2(H)]}) \quad (1.47)$$

with $|\delta_2(H)| \leq C \sum |\theta(x) - \theta(y)|^2 \leq CL^{D-2}$.

If the dimension is $D \leq 2$ this is bounded. The soft rotation dropped this quantity from L^{D-1} to L^{D-2} .

However, $\delta_1(H)$ might be bigger; a priori we only have $|\delta_1(H)| \leq C \sum |\theta(x) - \theta(y)| = O(L^{D-1})$. The idea is that we can rotate “from the other side.” To deal with the first order term,

$$\mu(f) = \mu(fe^{-\frac{(2\pi-\alpha)}{\pi}\delta_1(H)+\tilde{\delta}_2(H)}) \quad (1.48)$$

The second-order term is controllable, the first-order term is hard to control. We have two expressions where the coefficient of the first-order terms has different signs. Let convexity work for us.

Consider

$$Q(t) = \mu(fe^{(1-t)[\frac{\alpha}{\pi}\delta_1(H)+\delta_2(H)]}e^{t[-\frac{(2\pi-\alpha)}{\pi}\delta_1(H)+\tilde{\delta}_2(H)]}).$$

The function $\ln Q(t)$ is convex in t with $Q(0) = Q(1) = \mu(R_\alpha f)$. Thus by convexity, or Hölder’s inequality, for all $0 \leq t \leq 1$,

$$Q(t) \leq Q(0)^{1-t}Q(1)^t = \mu(Rf).$$

$(1-t)\alpha - t(2\pi - \alpha) = 0$. At $t = t_\alpha$,

$$Q(t) = \mu(f)e^{(1-t)\delta_2+t\tilde{\delta}_2} \geq e^{-CL^{D-2}}\mu(f)$$

giving us $\mu(R_\alpha f) \geq e^{-CL^{D-2}}\mu(f)$. □

What gave physicists the intuition to look at second order term is that at low temperature, when the spins line up, the first order term vanishes. The second-order term is $\frac{L^d}{L^2}$. But there is chaos and term-by-term it’s not true that the first-order term can be ignored. We sidestep this in the analysis.

If $D > 2$, the bound says nothing.

The theorem implies that the pressure $\Psi(\beta, \underline{h})$ is differentiable in \underline{h} at $\underline{h} = \underline{0}$. We know it’s symmetric, but it could have a kink or be smooth with zero derivative. The theorem tells us it is smooth.

The argument works equally well for quantum systems and for systems with disorder. For a while people had two predictions based on 2 theories. Disordered systems in D dimensions behaved as ordered systems in $D - 2$ dimensions. Imry and Ma predicted that for $D \leq 4$ there is no symmetry breaking. it uses this technology with bootstrapping.

Rigorous proofs arrived before physicists had a consensus.

Symmetry breaking occurs in dimensions $D \geq 3$.

[4-14-16](#)

Remarks on DLR condition: For any measure μ on (Ω, \mathcal{B}) , we have a map

$$P_{\Lambda^c}(f) = \mathbb{E}(f|\sigma_{\Lambda^c}) = \mathbb{E}_{\mathcal{B}_{\Lambda^c}}(f|\sigma),$$

giving the expected value of a function which depends on everything, conditioned on the spins outside Λ .

The tower property is the following: For $\Lambda_1 \subset \Lambda_2 \subset \dots$, for all $m > n$,

$$(P_{\Lambda_m^c} P_{\Lambda_n^c})f = P_{\Lambda_m^c} f.$$

The measures $\mathcal{B}_{\Lambda_n^c}$ are decreasing.

\mathcal{B}_0 are “local functions”: each only depends on variables in a finite region. An infinite sum of these may not be local, for example, $\sum_{|x| \in \mathbb{Z}^d} e^{-\alpha|x|} \sigma_x \notin \mathcal{B}_0$. However, up to $\varepsilon > 0$ it can be approximated uniformly by local functions.

\mathcal{B} is the smallest collection of functions containing the local functions and that is closed under the usual operations.

For all finite Λ , P_{Λ^c} is an orthonormal projection onto $\{f \in L^2(d\mu) : f \in \mathcal{B}_{\Lambda^c}\}$.

Theorem 1.8.24 (Based on martingale convergence theorem). *For all bounded f ,*

$$\mathbb{E}_{\mathcal{B}_\infty}(f|\sigma) = \lim_{n \nearrow \infty} P_{\Lambda_n^c} f(\sigma)$$

exists for μ -almost every σ .

The averages stay the same but have diminishing fluctuations.

There are σ for which the limit does not converge, e.g., spins organized in rings of + and -.

Gibbs measures are characterized by the fact that the conditional expectations are given by nothing else other than the DLR conditions.

Any two configurations that differ by a finite number of spins are the same at ∞ .

Suppose you have a function which is measurable and bounded. Suppose we only know integrals of continuous functions. We can take pointwise limits. Use Monotone Convergence Theorems.

Given a measure on the unit interval, how do we read the integral of measurable functions which are not continuous?

Remarks on the Mermin-Wagner Theorem: In 2 dimensions there is no symmetry breaking. The idea was to look at the energy penalty for gradual rotation. To estimate $H(\widehat{R}_\theta \sigma) - H(\sigma)$, we bounded the second-order term $\delta_2(H) \leq c \sum_{\{x,y\}, |x-y|} |\theta_x - \theta_y|^2$.

Using discrete notation, we saw that (for $r = L, R = 2L$)

$$\theta_x = \begin{cases} 0, & |x| = R \\ 1, & |x| = r. \end{cases} \quad \min_x \sum |\nabla \theta|^2 \leq cL^{d-2}$$

We could do better:

$$\min \left\{ \int_{r < |x| < R} |\nabla \theta|^2 d^2x : \theta(x) = \begin{cases} 0, & |x| = R \\ 1, & |x| = r. \end{cases} \right\}.$$

The solution is a harmonic function, satisfying $\Delta \theta = 0$:

$$\theta(x) = \frac{\ln\left(\frac{1}{R}\right)|x|}{\ln\frac{R}{r}}, \quad \ln|x| = \Re \ln x.$$

We have $\nabla\theta = \frac{1}{\ln\frac{R}{r}} \frac{1}{|x|}$, and

$$\int_{r < |x| < R} |\nabla\theta|^2 d^2x = \frac{1}{\ln\left(\frac{R}{r}\right)}.$$

In the energy estimate, the second-order term can be made as small as desired.

9 Random field models

Let

$$H_0 = - \sum J_{x-y} (\underline{\sigma}_x \cdot \underline{\sigma}_y) = \sum_A J_A \Phi(\sigma_A);$$

this is rotationally invariant.

Consider now

$$H = H_0 - \sum (\underline{h} + \varepsilon \underline{\eta}_x, \underline{\sigma}_x);$$

we couple with an external field \underline{h} that encourages the spins to line up with it. However, there is disorder; $\underline{\eta}_x$ are iid with $|\eta_x| = 1$. What are the statistical mechanics?

Would the system remember boundary conditions or would the local disorder $\underline{\eta}_x$ dominate?

Theorem 1.9.1 (Imry-Ma (Aizenmann-Wehr)). *For such systems, in $D \leq 4$, for such a system with continuous symmetry,*

$$\psi(\beta, \varepsilon, \underline{h})$$

is differentiable at $\underline{h} = 0$.

Why is the pressure of such a system defined in the infinite-volume limit? We first generate a random environment. It is defined because the way we prove existence of thermodynamic limit was to partition it into noncommunicating regions. If you add randomness, the law of large numbers for independent variables implies that the $\psi(\beta, \varepsilon, \underline{h})$ is independent of η almost surely (independent of disorder).

The pressure and free energy exist. With frozen disorder, If there is initially some bias, then we maintain it. The disorder itself may create some symmetry breaking. In low dimensions the disorder plays a decisive role.

10 Proof of symmetry-breaking of continuous symmetries

Let's return to homogeneous systems. We are coming to research-level questions. There is a dearth of arguments to prove symmetry breaking for continuous spaces.

10.1 The spin-wave perspective

The key results were by Fröhlich-Simon-Spencer, Fröhlich-Israel-Lieb-Simon.

We can write

$$\text{eq:csb-h1} H(\sigma) = \frac{1}{4} \sum_{x,y} J_{x,y} |\underline{\sigma}_x - \underline{\sigma}_y|^2. \quad (1.49)$$

(Note that $|\underline{\sigma}_x|^2 + |\underline{\sigma}_y|^2$ is constant, while the cross-term is what we had originally.) Define

$$\text{eq:csb-f} \hat{\sigma}(p) = \frac{1}{\sqrt{|\Lambda|}} \sum_{x \in \Lambda} e^{ip \cdot x} \underline{\sigma}_x. \quad (1.50)$$

Being invariant under translations, we can express H as a sum of Fourier modes,

$$\text{eq:csb-h} H(\sigma) = \frac{1}{4} \sum_{\underline{k}} \mathcal{E}(\underline{k}) |\hat{\sigma}(\underline{k})|^2 \quad (1.51)$$

where $\underline{k} \in \frac{2\pi}{L} \mathbb{Z}^d \cap [-\pi, \pi]^d$, the momenta for which the function is periodic.

You may be tempted to say the problem is now trivial. There is energy for each mode and no coupling. Had this been the end of the story, we have a bunch of independent random variables. However, they are not independent. There are bizarre correlations between the Fourier modes. If we ignored them we would have a quadratic interaction and the equipartition law says that the expected energy per mode is on the order of

$$\mathcal{E}(\underline{k}) |\hat{\sigma}(\underline{k})|^2 \approx \frac{1}{2\beta}.$$

Under certain conditions (reflection positivity), we can prove

$$\mathcal{E}(\underline{k}) |\hat{\sigma}(\underline{k})|^2 \leq \frac{1}{2\beta}.$$

We'll show how to derive this and how to conclude the theorem from it. It's a beautiful argument but it's a crystal vase. When you have reflection positivity, you use this physically motivated argument; when you don't, the vase breaks.

[4-19-16: We continue talking about continuous symmetry breaking.](#)

We will cover the infrared bound, reflection positivity, and the chessboard inequality.

Now the symmetry of the lattice will play a role.

Here $H = -\sum_{\{x,y\}} J_{x-y} \underline{\sigma}_x \cdot \underline{\sigma}_y$. We first cover the nearest-neighbor case. Our argument is not very robust; it only covers the nearest-neighbor case and several other special cases. It would be an achievement to find a generalization.

We pass to the Fourier transform. A basis for the space of periodic functions is

$$\frac{1}{\sqrt{|\Lambda|}} e^{ik \cdot x}, \quad \underline{k} = \frac{2\pi}{L} (n_1, \dots, n_d).$$

We scale to take the $(0, L]^d$ grid into the grid $(-\pi, \pi]^d$ with intervals $|\Delta k_i| = \frac{2\pi}{L}$.

Let $\hat{\sigma}(p)$ be as in (1.50). The Fourier expansion of σ is

$$\sigma_x = \frac{1}{\sqrt{|\Lambda|}} \sum_k^* e^{-ikx} \hat{\sigma}(k). \quad (1.52)$$

Letting H be as in (1.49), we find that the momentum representation of H is given by (1.51). Think of (1.49) as the discrete analogue of the integral of a gradient. Similarly, (1.51) is the discretized version of an integral.

We show that

$$\hat{\sigma}(k) = 2 \sum_{j=1}^d \sin^2 \left(\frac{k_j}{2} \right) \approx \frac{1}{2} |k|^2.$$

We calculate

$$H = -\frac{1}{2} \left(\frac{1}{\sqrt{|\Lambda|}} \right)^2 \sum_{x,y} \sum_{\underline{k}, \underline{k}'} \hat{\sigma}(\underline{k}) \hat{\sigma}(\underline{k}') e^{i\underline{k} \cdot \underline{x}} e^{-i\underline{k}' \cdot \underline{y}} J_{x-y} \quad (1.53)$$

$$= -\frac{1}{2} \frac{1}{|\Lambda|} \sum_{x,y} \sum_{\underline{k}, \underline{k}'} \hat{\sigma}(\underline{k}) \hat{\sigma}(\underline{k}') e^{i\underline{k} \cdot (\underline{x}-\underline{y})} e^{-i(\underline{k}-\underline{k}') \cdot \underline{y}} J_{x-y} \quad (1.54)$$

$$= -\frac{1}{2} \frac{1}{|\Lambda|} \sum_{u,y} \sum_{\underline{k}, \underline{k}'} \hat{\sigma}(\underline{k}) \hat{\sigma}(\underline{k}') e^{i\underline{k} \cdot \underline{u}} e^{-i(\underline{k}-\underline{k}') \cdot \underline{y}} J_u \quad (1.55)$$

$$\text{using } \sum_{\underline{k}, \underline{k}'} e^{i(\underline{k}-\underline{k}') \cdot \underline{y}} = \frac{1}{|\Lambda|} \delta_{\underline{k}, \underline{k}'} \quad (1.56)$$

$$= -\frac{1}{2} \sum_{\underline{k}} |\hat{\sigma}(\underline{k})|^2 \hat{J}(\underline{k}) \quad (1.57)$$

$$\hat{J}(\underline{k}) = \sum_u e^{i\underline{k} \cdot \underline{u}} J_u \quad (1.58)$$

$$= \sum_{j=1}^d (e^{ik_j} + e^{-ik_j}) \quad (1.59)$$

$$= -\frac{1}{2} \sum_{j=1}^d \cos(k_j) \quad (1.60)$$

$$= -\sum_{j=1}^d \left(1 - 2 \sin^2 \left(\frac{k_j}{2} \right) \right) \quad (1.61)$$

Looking at this intuitively, how big can we expect to be $\sum_{x \in \Lambda} e^{ikx} \sigma_x$? The Central Limit Theorem suggest that the order of magnitude should be $\sqrt{|\Lambda|}$, making $\hat{\sigma}(k) = O(1)$ (in high temperature). The energy H is given by a sum of terms of order 1, hence is of order $O(|\Lambda|)$.

This leads us to the right ‘‘infrared behavior.’’ $\mathcal{E}(\underline{k}) = -\hat{J}(\underline{k}) + \hat{J}(0)$. (?)

Plancharel's identity tells us

$$\frac{1}{|\Lambda|} \sum_x |\sigma_x|^2 = \frac{1}{|\Lambda|} \sum_{\underline{k}}^* |\hat{\sigma}(k)|^2 \quad (1.62)$$

$$1 = \frac{1}{|\Lambda|} |\hat{\sigma}(0)|^2 + \frac{1}{|\Lambda|} \sum_{\underline{k}=\frac{2\pi}{L}\underline{n} \in (-\pi, \pi]^d} |\hat{\sigma}(k)|^2 \quad (1.63)$$

Think of this as a Riemann sum approximation to the integral $\frac{1}{(2\pi)^d} \int_{(-\pi, \pi]^d} du$. (Each cube has volume $(\frac{2\pi}{L})^d$.)

10.2 Infrared bound

Theorem 1.10.1 (Infrared bound). *For the nearest neighbor (and other reflection positive models),*

$$\langle |\hat{\underline{\sigma}}(k)|^2 \rangle_{\beta, L}^{p, b, c} \leq \frac{1}{2\beta \mathcal{E}(k)}.$$

Equivalently,

$$\text{eq:csb-fs} \mathcal{E}(k) \langle |\hat{\underline{\sigma}}(k)|^2 \rangle \leq \frac{1}{2\beta}. \quad (1.64)$$

Near 0, $\mathcal{E}(k) \approx \frac{1}{2}|k|^2$. Thus the infrared bound diverges at 0.

This is a very physical statement. The physical interpretation is that $\mathcal{E}(k) \langle |\hat{\underline{\sigma}}(k)|^2 \rangle$ is the mean energy in the k th mode.

Take the mean value of (1.64) to obtain

$$1 = \left\langle \frac{1}{|\Lambda|} |\hat{\sigma}(0)|^2 \right\rangle + \frac{1}{|\Lambda|} \sum_{\underline{k}=\frac{2\pi}{L}\underline{n} \in (-\pi, \pi]^d} \underbrace{\langle |\hat{\sigma}(k)|^2 \rangle}_{\frac{1}{2\beta} \frac{1}{\mathcal{E}(k)}}.$$

If $d > 2$, although the bound diverges at 0, it is integrable.

We have

$$1 \leq \frac{1}{|\Lambda|} \langle |\hat{\sigma}(0)|^2 \rangle + \left[\int \frac{1}{(2\pi)^d} \int_{(-\pi, \pi), k \neq 0} \frac{1}{\mathcal{E}(k)} \right] \frac{1}{2\beta} \quad (1.65)$$

$$\lim_{L \rightarrow \infty} \frac{1}{L^d} \langle |\hat{\sigma}(0)|^2 \rangle_{L, \beta} \geq 1 - \frac{1}{2\beta} \underbrace{\frac{1}{(2\pi)^d} \int_{(-\pi, \pi]^d} du^d \frac{1}{\mathcal{E}(k)}}_{=: C_d < \infty \text{ when } d > 2} \quad (1.66)$$

This says that “zero momentum requires positive mass.” Compare this to Bose-Einstein condensation, when there is a point mass at a particular mode.

This is more informative if β is large. For $\beta > 2C_D$,

$$\left\langle \left| \frac{1}{|\Lambda|} \sum \sigma_x \right|^2 \right\rangle = \frac{1}{|\Lambda|} \langle |\hat{\sigma}(0)|^2 \rangle_{L, \beta} \geq \left(1 - \frac{C_D}{2\beta} \right) > 0.$$

Also,

$$\langle \underline{\sigma}_0 \underline{\sigma}_x \rangle = \left\langle \left(\frac{1}{\sqrt{|\Lambda|}} \sum_{k'}^* \hat{\sigma}(k') \right) \left(\frac{1}{\sqrt{|\Lambda|}} \sum_k^* e^{-ik \cdot x} \hat{\sigma}(k) \right) \right\rangle \quad (1.67)$$

$$\langle \underline{\sigma}_0 \underline{\sigma}_x \rangle = \frac{1}{|\Lambda|} \sum_k \langle |\hat{\sigma}(k)|^2 \rangle e^{-ik \cdot x} \quad (1.68)$$

$$\frac{1}{|\Lambda|} \sum_x |\langle \underline{\sigma}_0 \underline{\sigma}_x \rangle|^2 = \frac{1}{|\Lambda|^2} \langle |\hat{\sigma}(0)| \rangle^2 + \frac{1}{2\pi} \int_{(-\pi, \pi]^d} du'' \langle |\hat{\sigma}(k)|^2 \rangle e^{ikx}. \quad (1.69)$$

The function has an integrable singularity at the origin. A Riemann-Lebesgue argument shows this tends to 0 as $x \rightarrow \infty$ with $|x| \leq \frac{L}{2}$. (I.e., take $\lim_{|x| \rightarrow \infty} \lim_{L \rightarrow \infty}$.)

We get that

$$\langle \underline{\sigma}_0, \underline{\sigma}_x \rangle_{L, \beta} = \Omega(1)$$

for all L (long-range order).

To saturate the sum rule, we have to get massive density/volume-type occupation at 0.

The IR bound gives

$$Z = \int e^{\frac{\beta}{2} \sum_{x,y} J_{x,y} |\underline{\sigma}_x - \underline{\sigma}_y|^2} \rho_0(d\sigma) \quad (1.70)$$

$$Z(\underline{h}) = \int e^{\frac{\beta}{2} \sum_{x,y} J_{x,y} |\underline{\sigma}_x + \underline{h}_x - (\underline{\sigma}_y + \underline{h}_y)|^2} \rho_0(d\sigma). \quad (1.71)$$

We yank different spins out of ± 1 by shifting. What is the optimal shift? The best you can do is do nothing.

We claim $Z(\underline{h}) \leq Z(\underline{0})$.

Then $\frac{\partial^2}{\partial x^2} Z(\lambda \underline{h})|_{\lambda=0} \leq 0$, giving the IR bound.

10.3 Reflection positivity

Recall the Cauchy-Schwarz inequality: If $\langle \cdot, \cdot \rangle$ is a bilinear positive form ($\langle \psi, \psi \rangle \geq 0$), then

$$|\langle \psi, \varphi \rangle| \leq |\langle \psi | \psi \rangle|^{\frac{1}{2}} |\langle \varphi | \varphi \rangle|^{\frac{1}{2}}.$$

Definition 1.10.2: A system is **reflection positive** with respect to a plane if letting R be the reflection around the plane, letting L and R be the left and right side of the plane, for every $f \in \mathcal{B}_L$,

$$(Rf)(0) = f(R\sigma)$$

and

$$\langle \overline{F}(\sigma) F(R\sigma) \rangle \geq 0.$$

The Schwarz inequality and reflection positivity say that for all $F, G \in \mathcal{B}_L$,

$$|\langle F(\sigma) G(R\sigma) \rangle| \leq \langle \overline{F} R F \rangle^{\frac{1}{2}} \dots \langle \overline{G} R G \rangle^{\frac{1}{2}}.$$

The idea is to take reflections over many different planes. Consider a product over boxes α ,

$$\langle \pi_\alpha F_\alpha(\sigma) \rangle \leq \prod_\alpha \left\langle \prod_{\alpha'} F_\alpha^\#(\sigma_{\alpha'}) \right\rangle^{\frac{1}{|\Lambda|}} \geq 0$$

where $F_\alpha^\#$ is F_α extended to all squares by reflection, the “chessboard” extension.

Index

- concave, 13
- concave hull, 16
- configuration, 7
- convex, 13

- directional derivative, 14
- doubly stochastic, 45

- energy content, 21
- energy function, 19
- entropy, 10
- Equidistribution assumption, 8
- equivalence principle, 8
- exposed points, 26

- generating function, 20
- Gibbs equilibrium measure, 19
- Gibbs equilibrium state, 58
- Gibbs factors, 13
- grand canonical ensemble, 8
- grand canonical ensemble average, 8

- Ising model, 20

- Jensen's inequality, 17

- large deviation principle, 17
- lattice gas, 7
- Legendre transform, 12, 15

- Markov chain, 45
- microcanonical ensemble, 8
- microcanonical ensemble average, 8

- partition function, 11
- pressure, 32

- Radon-Nikodym derivative, 21

- temperature, 20
- transfer matrix, 43

- van Hove, 31