Chapter 1

Lattice spin models: Crash course

1.1 Basic setup

Here we will discuss the basic setup of the models to which we will direct our attention throughout this course. The basic ingredients are as follows:

- *Lattice*: We will take the d-dimensional hypercubic lattice $\mathbb{Z}^d$ as our underlying graph. This is a graph with vertices at all points in $\mathbb{R}^d$ with integer coordinates and edges between any *nearest neighbor* pair of vertices, which are those at Euclidean distance one. We will use $\langle x, y\rangle$ to denote a nearest-neighbor pair.

- *Spins*: At each $x \in \mathbb{Z}^d$ we will consider a spin S_x, by which we will mean a random variable taking values in a closed subset Ω of $\mathbb{R}^\nu$, for some $\nu \geq 1$. We will use $S_x \cdot S_y$ to denote a scalar product between S_x and S_y (Euclidean or otherwise).

- *Spin configurations*: For $\Lambda \subset \mathbb{Z}^d$, we will refer to $S_\Lambda = (S_x)_{x\in\Lambda}$ as the spin configuration in Λ. We will be generically interested in describing the statistical properties of such spin configurations with respect to certain (canonical) measures.

- *Boundary conditions*: To describe the law of S_Λ, we will not be able to ignore that there are (in general) also the spins outside Λ. We will refer to S_{Λ^c} as the boundary condition. The latter will usually be fixed and may often even be considered a parameter of the game. When both S_Λ and S_{Λ^c} are known, we will write

  $$S = (S_\Lambda, S_{\Lambda^c}) \tag{1.1}$$

  to denote their concatenation on all of $\mathbb{Z}^d$.

The above setting incorporates rather varied physical interpretations. The spins may be thought of as describing magnetic moments of atoms in a crystal, displacement of atoms from their equilibrium position or even orientation of grains in nearly-crystalline granular materials.

To define the dynamics of such spin systems, we will need to specify the energetics. This is conveniently done by prescribing the *Hamiltonian* which is a function on the spin-configuration space $\Omega^{\mathbb{Z}^d}$ that tells us how much energy each spin configuration has. Of course, to have all quantities well defined we need to fix a *finite* volume $\Lambda \subset \mathbb{Z}^d$ and compute only the energy in Λ. The most general formula we need is

$$H_\Lambda(S) = \sum_{\substack{A\subset\mathbb{Z}^d \text{ finite} \\ A\cap\Lambda\neq\emptyset}} \Phi_A(S) \tag{1.2}$$

where Φ_A is a function that depends only on S_A. To make everything well defined, we require e.g. that Φ_A is translation invariant and that $\sum_{A\ni 0} \|\Phi_A\|_\infty < \infty$. (The infinity norm may be replaced by some other norm, should the need to talk about unbounded spin systems arise.) It is often more convenient—and is invariably done by physicist—to write the above as a formal sum

$$H(S) = \sum_A \Phi_A(S) \tag{1.3}$$

with the above specific understanding in the situation where rigorous definition is required.

The energy is not sufficient on its own to determine the spin system; we also need to specify the *a priori measure* on the spins. This will be done by prescribing a Borel measure μ_0 on Ω (which may or may not be finite); the spin configurations (in finite volume) will be "distributed" according to the product measure, e.g., the *a priori* law of S_Λ is $\bigotimes_{x\in\Lambda} \mu_0(\mathrm{d}S_x)$.

1.2 Examples

Here are a few examples of spin systems:

(1) *O(n)-model*: Here $\Omega = \mathbb{S}^{n-1} = \{z \in \mathbb{R}^n : |z|_2 = 1\}$ with μ_0 = surface measure. The Hamiltonian is

$$H(S) = -J \sum_{\langle x,y\rangle} S_x \cdot S_y \tag{1.4}$$

where the dot denotes the usual (Euclidean) dot-product in $\mathbb{R}^n$ and $J \geq 0$. (The sign of J can be reversed by reversing the spins on the odd sublattice of $\mathbb{Z}^d$.)

Note that if $A \in O(n)$—i.e., A is an n-dimensional orthogonal matrix—then

$$AS_x \cdot AS_y = S_x \cdot S_y \tag{1.5}$$

and so $H(AS) = H(S)$. Since also $\mu_0 \circ A^{-1} = \mu_0$, the model possesses a *global rotation invariance* (with respect to simultaneous rotation of all spins).

Two instances of this model are known by other names: $n = 2$ is the *rotor model* while $n = 3$ is the (classical) *Heisenberg ferromagnet*.

(2) *Ising model*: Formally, this is $O(1)$-model. Explicitly, the spin variables σ_x take values in $\Omega = \{-1, +1\}$ with uniform *a priori* measure; the Hamiltonian is

$$H(\sigma) = -J \sum_{\langle x,y\rangle} \sigma_x\sigma_y \tag{1.6}$$

Note that the energy is smaller when the spins at nearest neighbors align and higher when they antialign. A similar statement holds, of course, for all $O(n)$ models. This is due to the choice of the sign $J \geq 0$ which makes these models *ferromagnets*.

(3) Potts model: This is a generalization of the Ising model to more spin states. Explicitly, we fix $q \in \mathbb{N}$ and let σ_x take values in $\{1, \dots, q\}$ (with uniform *a priori* measure). The Hamiltonian is

$$H(\sigma) = -J \sum_{\langle x,y \rangle} \delta_{\sigma_x,\sigma_y} \tag{1.7}$$

so the energy is $-J$ when σ_x and σ_y "align" and zero otherwise. The $q = 2$ case is the Ising model and $q = 1$ may be related to bond percolation on $\mathbb{Z}^d$ (via so called *Fortuin-Kasteleyn representation* leading to a *random-cluster model*).

It turns out that the above Hamiltonian can be brought to the form similar to the $O(n)$-model. Indeed, let Ω denote the set of q points uniformly distributed on the unit sphere in $\mathbb{R}^{q-1}$; we may think of these are vertices of a q-simplex (or regular q-hedron). The cases $q = 2, 3, 4$ are depicted in this figure:

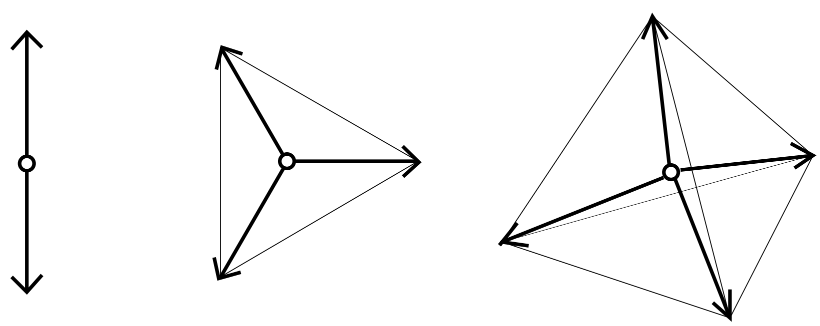

Explicitly, the elements of Ω are vectors $\hat{\mathrm{v}}_\alpha$, $\alpha = 1, \dots, q$ such that

$$\hat{\mathrm{v}}_\alpha \cdot \hat{\mathrm{v}}_\beta = \begin{cases} 1, & \text{if } \alpha = \beta, \\ -\frac{1}{q-1}, & \text{otherwise.} \end{cases} \tag{1.8}$$

(You may prove the existence of such vectors by induction on q.) It is easy to check that if S_x corresponds to σ_x and S_y to σ_y, then

$$S_x \cdot S_y = \frac{q}{q-1} \delta_{\sigma_x,\tilde{\sigma}_y} - \frac{1}{q-1} \tag{1.9}$$

and so the Potts Hamiltonian can be written as

$$H(S) = -\tilde{J} \sum_{\langle x,y \rangle} S_x \cdot S_y \tag{1.10}$$

with $\tilde{J} = J \frac{q-1}{q}$.

(4) Liquid-crystal model: There are many models that describe materials known to many of us from digital displays: liquid crystals. The distinct feature of such materials is the

presence of orientational ordering where certain grains assume distinct relative orientation despite the fact that the system as a whole is rotationally invariant. One of the simplest models describing such situations is as follows: Consider spins $S_x \in \mathbb{S}^{n-1}$ with uniform *a priori* measure. The Hamiltonian is

$$H(S) = -\tilde{J} \sum_{\langle x,y\rangle} (S_x \cdot S_y)^2 \tag{1.11}$$

The dot product implies global rotation invariance, the square takes care of the fact that reflection of any of the spins should not change the energy (i.e., only the *orientation* not the *direction* of the spin matters).

As for the Potts model, the Hamiltonian can again be brought to the form similar to the $O(n)$-model. Indeed, given a spin $S \in \mathbb{S}^{n-1}$ with Cartesian components S_α, $\alpha = 1, \dots, n$, define a $n \times n$ matrix Q by

$$Q_{\alpha\beta} = S_\alpha S_\beta - \frac{1}{n}\delta_{\alpha\beta} \tag{1.12}$$

(The subtraction of the identity is rather arbitrary at this point; the goal is to achieve zero trace and thus reduce the number of independent variables characterizing Q to $n-1$— which is exactly as many degrees of freedom as S has.) As is easy to check, if $Q \leftrightarrow S$ and $\tilde{Q} \leftrightarrow \tilde{S}$ via the above formula, then

$$\mathrm{Tr}(Q\tilde{Q}) = (S \cdot \tilde{S})^2 - \frac{1}{n}. \tag{1.13}$$

Since Q is symmetric, the trace evaluates to

$$\mathrm{Tr}(Q\tilde{Q}) = \sum_{\alpha,\beta} Q_{\alpha\beta}\tilde{Q}_{\alpha\beta} \tag{1.14}$$

which is the canonical scalar product on $n \times n$ matrices. In such language the Hamiltonian again takes the form known from the $O(n)$ model.

At the point we pause to remark that all of the above Hamiltonians may be cast in the form

$$H = +\frac{1}{2}\sum_{x,y} J_{x,y}|S_x - S_y|^2 \tag{1.15}$$

This is possible because, in all cases, the norm of S_x is constant; the above formula extends the nearest-neighbor interaction to arbitrary length by introducing suitable *coupling constants* J_{xy}. The model thus obtained bears striking similarity to our last example:

(5) *Gradient free field*: Let $\Omega = \mathbb{R}$, μ_0 = Lebesgue measure and let $\mathsf{P}(x,y)$ be the transition kernel of a random walk on $\mathbb{Z}^d$. We assume that $\mathsf{P}(x,y) = \mathsf{P}(0, y-x)$. We will denote the variables by ϕ_x; the Hamiltonian is

$$H(\phi) = \frac{1}{2}\sum_{x,y} \mathsf{P}(x,y)(\phi_y - \phi_x)^2. \tag{1.16}$$

This can be rewritten as

$$H(\phi) = \big(\phi, (1-\mathsf{P})\phi\big)_{L^2(\mathbb{Z}^d)} =: \mathcal{E}_{1-\mathsf{P}}(\phi,\phi) \tag{1.17}$$

where experts on harmonic analysis of Markot chains will recognize $\mathcal{E}_{1-\mathsf{P}}(\phi,\phi)$ to be the *Dirichlet form* associated with the generator $1-\mathsf{P}$ of the above random walk. Once we introduce Gibbs measures, the joint law of the ϕ_x will be Gaussian; hence the name of the model.

Note that the only difference between (1.15) and (1.16) is that the spin variables are generally confined to a subset of a Euclidean space—which will ultimately mean their law is *not* Gaussian. One purpose of this course is to show how this similarity can be exploited to provide information on the models (1.15). The key word is *Gaussian domination* (cf the title of this course).

1.3 Gibbs formalism

To describe the statistical-mechanical properties of the above models, we resort to the formalism of Gibbs-Boltzmann distributions. First we define measure in finite volume: Given a finite set $\Lambda \subset \mathbb{Z}^d$ and a boundary condition S_{Λ^c} we define the *Gibbs measure* in Λ to be the measure on Ω^Λ given by

$$\mu_{\Lambda,\beta}^{(S_{\Lambda^c})}(\mathrm{d}S_\Lambda) = \frac{\mathrm{e}^{-\beta H_\Lambda(S)}}{Z_{\Lambda,\beta}(S_{\Lambda^c})}\prod_{x\in\Lambda}\mu_0(\mathrm{d}S_x). \tag{1.18}$$

Here $\beta \ge 0$ is the *inverse temperature*—in physics terms, $\beta = \frac{1}{\mathrm{k_B}T}$ where $\mathrm{k_B}$ is the Boltzmann constant and T is the temperature measured in Kelvins—and $Z_{\Lambda,\beta}(S_{\Lambda^c})$ is the *partition function*.

To extend this concept to infinite volume we have two options:

(1) Consider all possible weak cluster points of the family $\{\mu_{\Lambda,\beta}^{(S_{\Lambda^c})}\}$ as $\Lambda \uparrow \mathbb{Z}^d$ (with the boundary condition possibly varying with Λ).

(2) Identify a distinguishing property of Gibbs measures and use this to define infinite volume objects directly.

While approach (1) is ultimately very useful in practical problems, option (2) is more elegant at this level of generality. The requisite “distinguishing property” is as follows:

Lemma 1.1 [DLR condition] *Let $\Lambda \subset \Delta \subset \mathbb{Z}^d$ be finite sets and let $S_{\Delta^c} \in \Omega^{\Delta^c}$. Then (for $\mu_{\Delta,\beta}^{(S_{\Delta^c})}$-a.e. S_{Λ^c}),*

$$\mu_{\Delta,\beta}^{(S_{\Delta^c})}\big(\cdot\,\big|S_{\Lambda^c}\big) = \mu_{\Lambda,\beta}^{(S_{\Lambda^c})}(\cdot). \tag{1.19}$$

In simple terms, conditioning a measure in Δ on the configuration in $\Delta\setminus\Lambda$, we get the Gibbs measure in Λ with the corresponding boundary condition.

This leads to:

Definition 1.2 [DLR Gibbs measures] *A probability measure on $\Omega^{\mathbb{Z}^d}$ is called an* infinite volume Gibbs measure *for interaction H and inverse temperature β if for all $\Lambda \subset \mathbb{Z}^d$ and μ-a.e. S_{Λ^c},*

$$\mu\big(\cdot\,\big|S_{\Lambda^c}\big) = \mu_{\Lambda,\beta}^{(S_{\Lambda^c})}(\cdot) \tag{1.20}$$

where $\mu_{\Lambda,\beta}^{(S_{\Lambda^c})}$ is defined using the Hamiltonian H.

We will use $\mathfrak{G}_\beta$ to denote the set of all infinite volume Gibbs measures at inverse temperature β (assuming the model is clear from the context). It is clear that $\mathfrak{G}_\beta$ is convex.

Here are some straightforward, nonetheless important consequences of these definitions:

(1) As a consequence of Lemma 1.1, any weak cluster point of $(\mu_{\Lambda,\beta}^{(S_{\Lambda^c})})$ is in $\mathfrak{G}_\beta$.

(2) By the Backward Martingale Convergence, if $\Lambda_n \uparrow \mathbb{Z}^d$ and $\mu \in \mathfrak{G}_\beta$, then for μ-a.e. spin configuration S the sequence $\mu_{\Lambda_n,\beta}^{(S_{\Lambda_n^c})}$ has a weak limit.

(3) The measure μ is extremal in the simplex $\mathfrak{G}_\beta$ iff the limit of $\mu_{\Lambda_n,\beta}^{(S_{\Lambda_n^c})}$ is the same for μ-almost all spin configurations S.

Similarly direct is the proof of the following "continuity" property:

(4) Let H_n be a sequence of Hamiltonians converging—in the sup-norm on the potentials Φ_A—to Hamiltonian H, and β_n is a sequence with $\beta_n \to \beta < \infty$. Let μ_n be a sequence of corresponding Gibbs measures. Then μ_n converges to a Gibbs measure for Hamiltonian H and inverse temperature β.

Now we give a meaning to the terms often used vaguely by physicists:

Definition 1.3 [Phase coexistence] *We say that the model is at* phase coexistence *(or undergoes a* 1st-order phase transition*) whenever the parameters are such that $|\mathfrak{G}_\beta| > 1$.*

The simplest example where this happens is the Ising model. Let $\Lambda_L = \{1, \dots, L\}^d$ and consider the Ising model in Λ_L with all boundary spins set to $+1$. This is the so called *plus boundary condition*. As a consequence of stochastic domination—which we will not discuss here—$\mu_{\Lambda_L,\beta}^+ \to \mu^+$ as $L \to \infty$. Similarly, $\mu_{\Lambda_L,\beta}^- \to \mu^-$. It turns out that, in dimensions $d \geq 2$ there exists $\beta_c(d) \in (0, \infty)$ such that

$$\beta > \beta_c(d) \quad \Rightarrow \quad \mu^+ \neq \mu^- \tag{1.21}$$

while for $\beta < \beta_c(d)$, the set of all infinite volume Gibbs measures is a singleton. We will prove similar statements in all of the models introduced above.

1.4 Torus measures

In the above, we always put a boundary condition in the complement of the finite set Λ. However, it is sometimes convenient to consider other boundary conditions. One possibility is to ignore the existence of Λ^{c} complement altogether—this leads to the so called *free boundary condition*. Another possibility is to wrap Λ into a graph without boundary—typically a torus. This is the case of *periodic* or *torus boundary conditions*.

Consider the torus $\mathbb{T}_L$ which we define as $(\mathbb{Z}/L\mathbb{Z})^d$ endowed with the corresponding (periodicized) nearest-neighbor relation. For nearest-neighbor interactions, the corresponding Hamiltonian is defined easily, but some care is needed for interactions that can be of arbitrary range. If $S \in \Omega^{\mathbb{T}_L}$ we define the *torus Hamiltonian* $H_L(S)$ by

$$H_L(S) = H_{\Lambda_L}(\text{periodic extension of } S \text{ to } \mathbb{Z}^d) \tag{1.22}$$

where we recall $\Lambda_L = \{1, \dots, L\}^d$. For $H(S) = \frac{1}{2}\sum_{x,y} J_{x,y} S_x \cdot S_y$ we thus get

$$H_L(S) = \frac{1}{2}\sum_{x,y} J^{(L)}_{x,y}\, S_x \cdot S_y \tag{1.23}$$

where $J^{(L)}_{x,y}$ are the periodicized coupling constants

$$J^{(L)}_{x,y} = \sum_{z\in\mathbb{Z}^d} J_{x,y+Lz}. \tag{1.24}$$

The Gibbs measure on $\Omega^{\mathbb{T}_L}$ is then defined accordingly:

$$\mu_{L,\beta}(\mathrm{d}S) = \frac{\mathrm{e}^{-\beta H_L(S)}}{Z_{L,\beta}} \prod_{x\in\mathbb{T}_L} \mu_0(\mathrm{d}S_x) \tag{1.25}$$

where $Z_{L,\beta}$ is the torus partition function. The following holds:

Lemma 1.4 *Every (weak) cluster point of $(\mu_{L,\beta})_{L\ge 1}$ lies in $\mathfrak{G}_\beta$.*

There *is* something to prove here because, due to (1.24), the interaction depends on L.

1.5 Some thermodynamics

For historical, and also practical reasons, many accounts of statistical mechanics start with the notion of free energy. We will need this notion only tangentially—it suffices to think of the free energy as a cumulant generating function—in the proofs of phase coexistence. The relevant statement is as follows:

Theorem 1.5 *For $x \in \mathbb{Z}^d$ let τ_x be the shift-by-x which is defined by $(\tau_x S)_y = S_{y-x}$. Let $g\colon \Omega^{\mathbb{Z}^d} \to \mathbb{R}$ be a bounded, local function and let $\mu_{L,\beta}$ be the torus measures. Then:*

(1) The limit

$$f(h) = \lim_{L\to\infty} \frac{1}{L^d} \log E_{\mu_{L,\beta}}\Big\{\exp\Big(h \sum_{x\in\mathbb{T}_L} g\circ\tau_x\Big)\Big\} \tag{1.26}$$

exists for all $h \in \mathbb{R}$ and is convex in h.

(2) If $\mu \in \mathfrak{G}_\beta$ is translation invariant, then

$$\frac{\partial f}{\partial h^-}\Big|_{h=0} \le E_\mu(g) \le \frac{\partial f}{\partial h^+}\Big|_{h=0}. \tag{1.27}$$

(3) There exist translation-invariant, ergodic measures $\mu^\pm \in \mathfrak{G}_\beta$ such that

$$E_{\mu^\pm}(g) = \frac{\partial f}{\partial h^\pm}\Big|_{h=0}. \tag{1.28}$$

Proof of (1). The existence of the limit follows by standard subbadditivity arguments. In fact, for compact state-spaces and bounded interactions, the measure $\mu_{L,\beta}$ could be replaced by any sequence of Gibbs measures in Λ_L with (even variable) boundary conditions. The convexity of f follows by Hölder inequality. □

Proof of (2). Let $\mu \in \mathfrak{G}_\beta$ be translation invariant and abbreviate

$$Z_L(h) = E_\mu\Big\{\exp\Big(h \sum_{x\in\Lambda_L} g\circ\tau_x\Big)\Big\} \tag{1.29}$$

Since $\log Z_L$ is convex in h, we have for any $h > 0$ that

$$\begin{aligned}\log Z_L(h) - \log Z_L(0) &\ge h\frac{\partial}{\partial h}\log Z_L(h)\Big|_{h=0}\\ &= hE_\mu\Big(\sum_{x\in\Lambda_L} g\circ\tau_x\Big) = h|\Lambda_L|E_\mu(g).\end{aligned} \tag{1.30}$$

Dividing by $|\Lambda_L|$, passing to $L \to \infty$ and applying independence of f on the boundary condition, we get

$$f(h) - f(0) \ge hE_\mu(g). \tag{1.31}$$

Divide by h and let $h \downarrow 0$ to get one half of (1.27). The other half is proved analogously. □

Proof of (3). A variant of proof of (2) shows that if μ_h is a translation-invariant Gibbs measure for the Hamiltonian modified by adding the term $-(h/\beta)\sum_x g\circ\tau_x$, then

$$\frac{\partial f}{\partial h^-} \le E_{\mu_h}(g) \le \frac{\partial f}{\partial h^+}. \tag{1.32}$$

In particular, if $h > 0$ we have

$$E_{\mu_h}(g) \ge \frac{\partial f}{\partial h^-} \underset{h>0}{\ge} \frac{\partial f}{\partial h^+}\Big|_{h=0} \tag{1.33}$$

by the monotonicity of derivatives of convex functions. Taking $h \downarrow 0$ and extracting a weak limit from μ_h, we get a Gibbs measure $\mu^+ \in \mathfrak{G}_\beta$ such that

$$E_{\mu^+}(g) \ge \frac{\partial f}{\partial h^+}\Big|_{h=0}. \tag{1.34}$$

(The expectations converge because g is a local—and thus continuous, in the product topology—function.) Applying (2) we verify (1.28) for μ^+.

The measure μ^+ is translation invariant and so it remains to show that μ^+ can actually be chosen ergodic. To that end let us first prove that

$$\frac{1}{|\Lambda_L|}\sum_{x\in\Lambda_L} g\circ\tau_x \underset{L\to\infty}{\longrightarrow} E_{\mu^+}(g), \quad \text{in } \mu^+\text{-probability} \tag{1.35}$$

The random variables on the left are bounded by the norm of g and have expectation $E_{\mu^+}(g)$ so it suffices to prove that the limsup is no larger than the expectation. However, if that weren't the case, we would have

$$\mu^+\Bigl(\sum_{x\in\Lambda_L} g\circ\tau_x > \bigl(E_{\mu^+}(g)+\epsilon\bigr)|\Lambda_L|\Bigr) > \epsilon \tag{1.36}$$

for some $\epsilon > 0$ and some sequence of L's. But then for all $h > 0$,

$$E_{\mu^+}\Bigl\{\exp\Bigl(h\sum_{x\in\Lambda_L} g\circ\tau_x\Bigr)\Bigr\} \ge \epsilon\mathrm{e}^{|\Lambda_L|h[E_{\mu^+}(g)+\epsilon]}. \tag{1.37}$$

This implies

$$f(h) \ge h\bigl(E_{\mu^+}(g)+\epsilon\bigr) \tag{1.38}$$

which cannot hold for all $h > 0$ should the right-derivative of f at $h = 0$ be equal $E_{\mu^+}(g)$. Hence (1.35) holds.

By the Pointwise Ergodic Theorem, the limit in (1.35) occurs μ^+-almost surely. This implies that the same must be true for any measure in the decomposition of μ^+ into ergodic components. By classic theorems from Gibbs-measure theory, every measure in this decomposition is also in $\mathfrak{G}_\beta$ and so we can choose μ^+ ergodic. $\square$

The above theorem is very useful for the proofs of phase coexistence. Indeed, one can often prove some estimates that via (1.27) imply that f is not differentiable at $h = 0$. Then one applies (1.28) to infer the existence of two distinct, ergodic Gibbs measures saturating the bounds in (1.27). Examples of the approach will be discussed momentarily.

1.6 Literature remarks

This chapter contains only the absolute minimum we need for understanding the rest of the course. For a comprehensive treatment of Gibbs-measure theory, we refer to books by Israel, Simon and Georgii. The acronym DLR stands for Dobrushin and Lanford-Ruelle who first introduced the idea of conditional definition of infinite volume Gibbs measures. The $O(n)$ model goes back to Heisenberg, the Ising model was introduced by Lenz and given to Ising as a thesis problem. An excellent reference for liquid crystals is the classic monograph by de Gennes. The tetrahedral representation of the Potts model can be found in Wu's review article on the Potts model; the matrix representation of the liquid-crystal model goes back to (at least) Zagrebnov.