1
The Ising Model

Few models in theoretical physics have been studied for as long, or in as much detail, as the Ising model. It's the simplest model to display a nontrivial phase transition, and as such it plays a unique role in theoretical physics. In addition, the Ising model can be applied to a wide range of physical systems, from magnets and binary liquid mixtures, to adsorbed monolayers and superfluids, to name just a few. In this chapter, we present some of the background material that sets the stage for a detailed study of the Ising model in the chapters to come.

1.1 Magnetic Phase Transitions

The original purpose of the Ising model was to study phase transitions in magnetic systems. We'll start, then, with a brief discussion of magnetic phases and phase diagrams, and we'll introduce the Ising model itself in the next section.

A simple magnetic system, like an iron magnet for example, displays two distinct behaviors, as illustrated in Figure 1.1 (a). At high temperatures the system is *paramagnetic*. In this region, the magnetization of the iron "tracks" with the applied magnetic field, B. What this means is that if the magnetic field is positive, relative to a particular direction, the magnetization is positive as well; if the magnetic field is zero, the magnetization is zero; if the magnetic field is negative, the magnetization is negative. In fact, for small magnetic fields the magnetization is simply proportional to the field.

At low temperatures the system is *ferromagnetic*. As the magnetic field is reduced from positive to zero in this region, the magnetization decreases to a finite positive amount and remains at that value. This is a magnetization that is due to the interactions between individual electron spins, and it persists even when there is no external field to align the spins. Similarly, as the field is
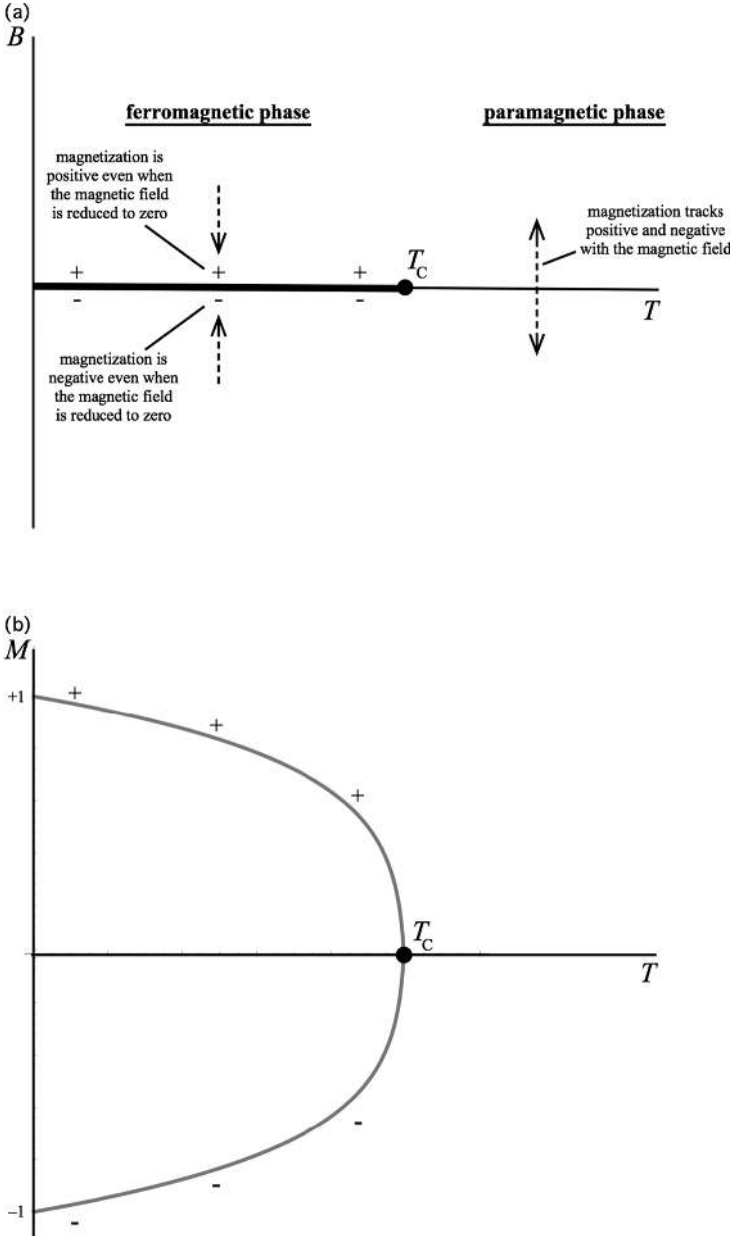
1

Figure 1.1 (a) A magnetic phase diagram in terms of the magnetic field, B, and temperature, T. At high temperature the system is paramagnetic, with no permanent magnetization. At low temperature the system is ferromagnetic, with up (+) and down (−) magnetization coexisting along the $B = 0$ line. The critical point at temperature T_c separates the ferromagnetic and paramagnetic regions. (b) Zero-field magnetization, M, versus temperature, T, for the system in part (a). Notice that the discontinuity in magnetization vanishes at the critical point, T_c.

reduced in magnitude from the negative side, the magnetization approaches a finite negative value. These observations are indicated by the + and – signs above and below the temperature axis in Figure 1.1 (a).

It follows that there are two coexisting phases of matter at zero magnetic field – the positive and negative magnetization phases – and that a discontinuity in the magnetization occurs as the magnetic field crosses $B = 0$. No such discontinuity is observed in the paramagnetic region of the phase diagram. The discontinuity in the magnetization decreases in magnitude with increasing temperature, and eventually vanishes at the critical temperature, T_c.

To highlight the key role played by the magnetization, consider a plot of magnetization, M, versus temperature, T, as in Figure 1.1 (b). In this plot, the applied magnetic field is zero, $B = 0$. As a result, the magnetization of the system is zero at temperatures that are higher than T_c. When the temperature is reduced to a value just below T_c, the system becomes magnetized; we refer to this as a *spontaneous magnetization*. Whether the spontaneous magnetization is positive or negative is determined randomly by the majority of spins when the temperature drops below T_c. If we reduce the temperature further, the result is an increasing magnitude of the magnetization. Eventually the magnetization saturates to its largest positive or negative value, which we normalize as $M = \pm 1$. The "jump" from the lower to the upper branch of the magnetization curve is the discontinuity that characterizes ferromagnetic behavior.

The way to think about this phase transition on a microscopic level is to picture electron spins in the iron atoms trying to align with their neighbors through exchange interactions. At high temperatures, thermal fluctuations overwhelm the magnetic interactions and the spins are oriented randomly. As the temperature is lowered, the magnetic interactions begin to win out over the thermal fluctuations, and the system spontaneously picks an orientation for the entire system. We say that the initial up–down symmetry of the system (recall that $B = 0$) has been "spontaneously broken" as a finite up or down magnetization takes over.

The magnetization is referred to as the *order parameter* for this phase transition. In general, the order parameter of a system is a quantity that is zero in the disordered phase, and finite for temperatures below T_c. In addition, the order parameter is a measure of the *amount* of order in the system; the larger the order parameter, the greater the order.

A mathematical function with a discontinuity, like the one we see in the magnetization below T_c, is referred to as a nonanalytic, or singular function – a key feature of phase transitions. It follows that the system's free energy – the function that contains all of the equilibrium information about the system – must be singular as well, as are its various derivatives. Since the magnetization is the first derivative of the free energy with respect to the magnetic field, phase

transitions like this are sometimes referred to in older literature as "first-order" phase transitions. The accepted modern terminology is to simply refer to the jump in magnetization below T_c as a "discontinuous" phase transition.

In contrast, consider the phase transition that occurs in Figure 1.1 (b) as one moves from high to low temperature along the temperature axis. In this case, the magnetization rises smoothly from zero to nonzero values as the temperature goes through T_c. This sort of phase transition, with a continuous change in the value of the order parameter, is referred to as a "continuous" phase transition.

Another example of a continuous phase transition is shown in Figure 1.2. Here we see the temperature dependence of the zero-field magnetic suscepti-bility, χ, which is defined as the derivative of the magnetization, M, with respect to the magnetic field, B:

$$\chi = \frac{\partial M}{\partial B}.$$

We see that χ is a smooth function, with no discontinuities. It is still a singular function, however; in fact, notice that it diverges to infinity at the critical temperature, T_c.

Again, we see that a phase transition is characterized by singular behavior in a thermodynamic function. The susceptibility χ is the derivative of M with respect to B, and M is the derivative of the free energy with respect to B, so χ is
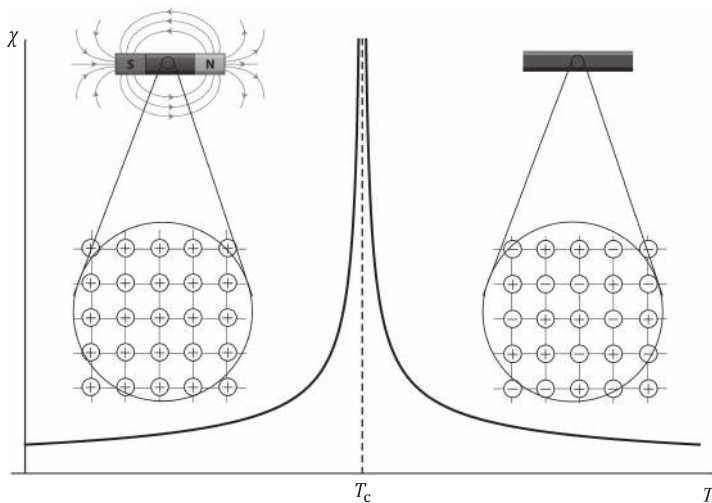

Figure 1.2 The zero-field magnetic susceptibility, $\chi = \partial M/\partial B$, as a function of temperature. The susceptibility blows up to infinity at the critical point, T_c.

the second derivative of the free energy. As a result, older literature sometimes refers to this as a "second-order" phase transition.

We shall return frequently to the topic of the free energy, which we denote by f, in our study of the Ising model. In fact, we will calculate the free energy for many different systems, and we'll see that it is the key to understanding the behavior of a system. Once we obtain the free energy, we take first derivatives to find average values, and second derivatives to find fluctuations about the average. This will be a recurring theme throughout our exploration of the Ising model.

Comparing with Other Phase Diagrams

Now that we've looked carefully at the characteristics of a magnetic phase diagram, we'd like to put it in context by comparing with the phase diagram of a familiar three-phase substance. In Figure 1.3 (a) we show a typical phase diagram for a material with solid, liquid, and gaseous phases. The bold curves denote discontinuous phase transitions, where two phases coexist along a curve in the pressure–temperature plane. As these curves are crossed, there is a discontinuity in the density of the material. The exception is at the critical
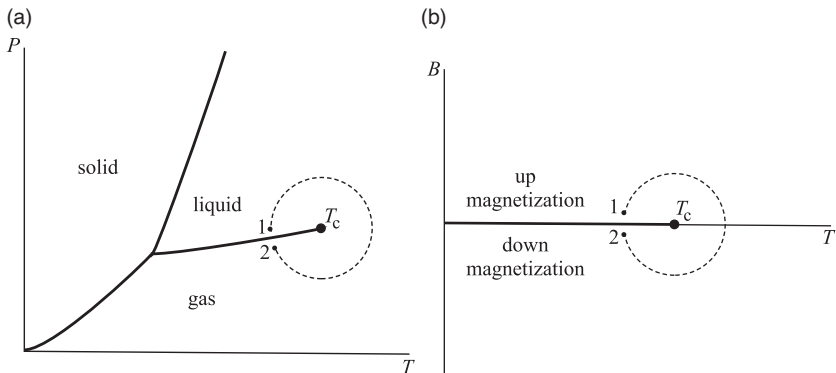

Figure 1.3 (a) Phase diagram of a simple substance with three distinct phases: solid, liquid, and gas. The bold curves denote the coexistence of the two phases on either side of the curve. The liquid–gas coexistence curve terminates at the critical point T_c, beyond which there is no distinction between liquid and gas. Points 1 and 2 can be connected without a phase transition by following a curved path, like the dashed circle, that goes around the critical point. The solid–liquid curve extends to infinity. (b) The magnetic phase diagram is similar to the liquid–gas portion of the phase diagram in part (a), with up magnetization and down magnetization playing the roles of liquid and gas, respectively. Points 1 and 2 can be connected without a phase transition.

point, T_c, where the density difference between the liquid and gas phases vanishes – similar to the vanishing magnetization difference at T_c in the magnetic phase diagram in Figure 1.3 (b).

In fact, the behavior along the liquid–gas transition curve in Figure 1.3 (a) and the up-magnetization/down-magnetization line in Figure 1.3 (b) are analogous in many respects. In both cases, a discontinuity is encountered if the bold line is crossed. On the other hand, the system can be brought smoothly, and with no discontinuities at all, along the dashed paths connecting points 1 and 2. The critical points, labeled T_c, are points where a continuous phase transition is observed.

The similarity in behavior near these critical points is more than just qualitative – on close examination, the quantitative details associated with the singularities are precisely the same in both the magnetic and nonmagnetic substances. These surprisingly deep connections, which lead to the concept of *universality* in critical behavior, will be explored in detail in Chapter 5.

1.2 The Ising Model of Magnetism

Now that we're familiar with the magnetic behavior of a simple system like an iron magnet, we would like to introduce a theoretical model to describe that behavior mathematically. One approach is to produce a model as realistic as possible, including millions of iron atoms interacting with one another through overlapping wave functions, and with electrons influencing one another like tiny magnets. Simulating such a system on a computer would be difficult, but perhaps doable with powerful machines. It certainly wasn't an option in 1920 when the Ising model was introduced.

A different approach, and one that is often used in theoretical physics, is to adopt the philosophy, frequently espoused by Albert Einstein, that "Everything should be made as simple as possible, but not simpler." In terms of a theoretical model, this means we would like to have a model that is bare-bones simple, but still realistic enough to produce the desired physical behavior. Additional complications can be added later, if desired, to include other types of behavior; but it's best to study the most basic aspects of a system first. The Ising model is a good example of this approach.

Ising Variables

The Ising model is constructed with "spin" variables, s, that occupy the sites, i, of a lattice. The value of the variable at any given lattice site, s_i, is meant to

represent, in a crude way, the spin of an electron. Now, as we know, electrons are spin-1/2 particles, with magnetic moments that are either aligned or anti-aligned along a specified direction in space. In the Ising model, we represent this by saying that each spin is either up (+1) or down (–1). Thus, Ising spins are simply dimensionless numbers:

$$s_i = \pm 1.$$

More than anything else, a two-valued variable like this is what identifies a model as "Ising like." The variables can interact in various ways, or lie on the vertices of various lattices, but they always have this simple plus–minus/up–down/yes–no quality.

A microstate (or state, or configuration) of an Ising model consists of a specific assignment of +1 or –1 to each of the spins in the system. If the system is a lattice with N sites, it follows that the total number of states is 2^N. In the coming chapters, we'll carry out sums over all such states to obtain the thermodynamic functions associated with the Ising model. A specific example of spin assignments for a group of Ising spins is given in Figure 1.4, along with different ways of representing the spins pictorially.

The variables in an Ising model can be given different names, as long as they have the same basic Ising symmetry. For example, consider a binary-liquid mixture, consisting of two types of molecules, A and B. We can describe a system like this with an Ising model whose variables represent one type of molecule, s_i = A, or the other, s_i = B, as in Figure 1.5 (a). Similarly, the Ising model could represent a system of helium atoms adsorbed on a surface of graphite. In this case, each hexagonal adsorption site on the graphite surface is either occupied by a helium atom ($s_i = 1$) or empty ($s_i = 0$), as illustrated in Figure 1.5 (b).

These are just a couple examples of the many nonmagnetic systems that have an underlying Ising-like character to them, and can be studied with the Ising model. We will generally speak of the Ising model in magnetic terms, but it should be remembered that the comments can apply to other types of systems as well.

Ising Hamiltonians

The next step in constructing an Ising model is to introduce interactions involving the spin variables. These interactions produce different energies for different spin states, as described by the Hamiltonian of the system. For example, the Hamiltonian might have an energy term related to a spin aligning or antialigning with an external magnetic field. Similarly, the Hamiltonian could include a term describing the interaction of one spin with its neighbors.
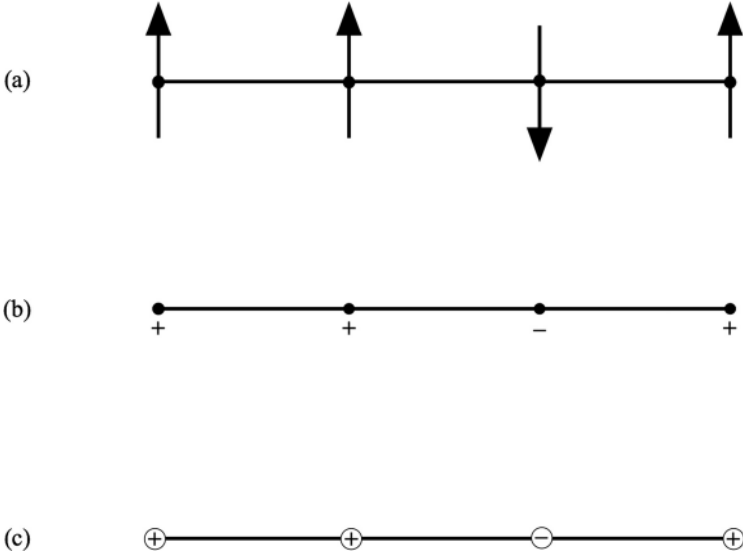
Figure 1.4 (a) Ising spins are often represented as up or down arrows indicating spins that are +1 or −1, respectively. In the state shown here, the spin values, reading from left to right, are +1, +1, −1, +1. (b) Ising spins can be represented with + or minus − signs. (c) A different way of representing the same state as in parts (a) and (b).

To begin, consider a single spin, s_1, interacting with an external magnetic field, B. The Hamiltonian H in this case can be written as follows:

$$H = -\mu_B B s_1.$$

In this expression, μ_B is the Bohr magneton, which means that $\mu_B B$ has the units of energy. The minus sign in front of μ_B ensures that a positive magnetic field, $B > 0$, favors a positive spin, $s_1 = +1$. To see why this is so, consider the energy levels of the two states of the spin:

Spin state	Energy
$s_1 = +1$	$H = -\mu_B B$
$s_1 = -1$	$H = \mu_B B$

Notice that the lowest energy for $B > 0$ corresponds to $s_1 = +1$, and so this state is favored. If $B < 0$ the lowest energy occurs when $s_1 = -1$. Thus, with this choice for the Hamiltonian, a given sign of B favors the same sign of s_1.
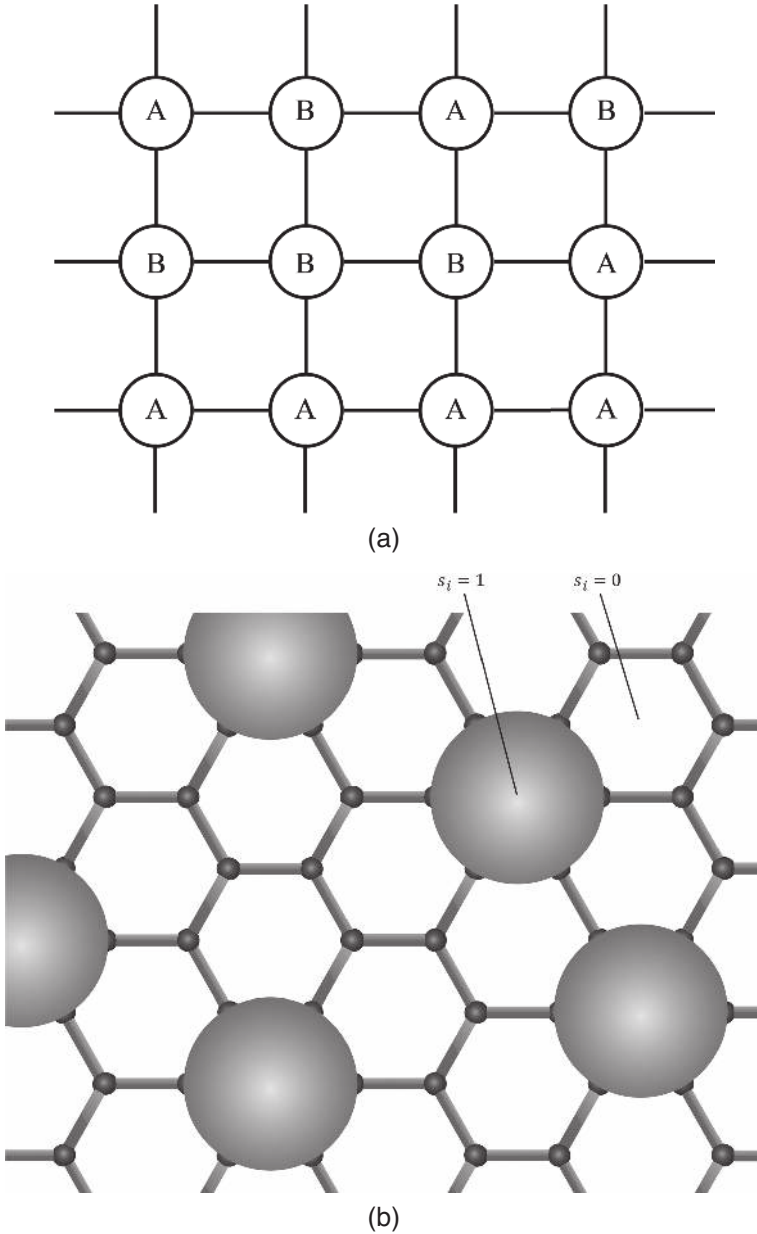
(a)

(b)

Figure 1.5 (a) A lattice model for a binary mixture, where the Ising variables take on the values $s_i = A$ or $s_i = B$. (b) The centers of the hexagons on a graphite lattice are adsorption sites that can be occupied by a helium atom, $s_i = 1$, or empty, $s_i = 0$. We can think of the s_i as Ising variables.

Next, let's consider the interaction between a spin, s_1, and a neighboring spin, s_2. If the interaction energy is J, the Hamiltonian can be written in the following form:

$$H = -Js_1s_2.$$

In this case, the minus sign ensures that $J > 0$ favors aligned spins ($s_1 = s_2$), and $J < 0$ favors antialigned spins ($s_1 = -s_2$). This can be seen in the following energy chart:

Spin state	Energy
$s_1 = s_2$	$H = -J$
$s_1 = -s_2$	$H = J$

We say that $J > 0$ is a ferromagnetic interaction, and $J < 0$ is an antiferromagnetic interaction.

Combining the two-spin interaction J with a magnetic field term B for two Ising spins yields the following Hamiltonian:

$$H = -Js_1s_2 - \mu_B B(s_1 + s_2).$$

The J term tends to align or antialign neighboring spins, and the B terms tend to align or antialign each spin with the direction of the magnetic field.

For the four-spin system shown in Figure 1.6 (a), the corresponding Hamiltonian would be

$$H = -J(s_1s_2 + s_2s_3 + s_3s_4 + s_4s_1) - \mu_B B(s_1 + s_2 + s_3 + s_4).$$

In this Hamiltonian the J terms are nearest-neighbor interactions. If we wanted to include next-nearest-neighbor interactions, with a strength J_2 for example, we would add the following terms to the Hamiltonian:

$$-J_2(s_1s_3 + s_2s_4).$$

Three- and four-spin interactions could be added as well, and we will see examples of multispin interactions in the following chapters, but in general we keep the interactions as simple as possible.

Finally, suppose we would like to apply a simple nearest-neighbor Hamiltonian with a magnetic field to the case of an infinite lattice, as in Figure 1.6 (b). In this case, the Hamiltonian is written as

$$H = -J \sum_{\langle ij \rangle} s_i s_j - \mu_B B \sum_i s_i. \tag{1.1}$$