

INTRODUCTION: WHY SPIN GLASSES?

Spin glasses are disordered magnetic materials, and it's hard to find a less promising candidate to serve as a focal point of complexity studies, much less as the object of thousands of investigations. On first inspection, they don't seem particularly exciting. Although they're a type of magnet, they're not very good at being magnetic. Metallic spin glasses are unremarkable conductors, and insulating spin glasses are fairly useless as practical insulators. So why the interest?

Well, the answer to that depends on where you're coming from. In what follows we'll explore those features of spin glasses that have attracted, in turn, condensed matter and statistical physicists, complexity scientists, and mathematicians and applied mathematicians of various sorts. In this introduction, we'll briefly touch on some of these features in order to (we hope) spark your interest. But to dig deeper and get a real sense of what's going on—that can fill a book.

Spin glass research provides mathematical tools to analyze some interesting (and hard) real-world problems.

Suppose you're given the following easily stated problem. You're shown a collection of N points on the plane, which we'll call cities. You're asked to start at one of the cities (any one will




Figure I.1. (a) An instance of a TSP problem with 19 cities. (b) One possible tour.

do), draw an unbroken line that crosses each of the other cities exactly once, and returns to the starting point. Such a line is called a tour, and an example is shown in figure I.1. All you need to do is to find the shortest possible tour.

This is an example of the Traveling Salesman Problem, or TSP for short.¹ An “instance,” or realization of the problem, is some specific placement of points on the plane (which a priori can be put anywhere). You should be able to convince yourself that the number of distinct tours when there are N cities is $(N - 1)!/2$. The factor of two in the denominator arises because a single tour can run in either direction.

Notice how quickly the number of tours increases with N : for 5 cities, there are 12 distinct tours; for 10 cities, 181,440 tours; and for 50 cities (not unusual for a sales or book tour in real life), the number of tours is approximately 3×10^{62} . The seemingly easy (i.e., lazy) way to solve this is to look at every possible tour and compute its length, a method called exhaustive search. Of course, you’re not about to do that yourself, but you have access

to a modern high-speed computer. If your computer can check out—let's be generous—a billion tours every second, it would take it 10^{46} years to come up with the answer for a 50-city tour. (For comparison, the current age of the universe is estimated at roughly 1.3×10^{10} years.) Switching to the fastest supercomputer won't help you much. Clearly, you'll need to find a much more efficient algorithm.

Does this problem seem to be of only academic interest? Perhaps it is,² but the same issues—lots of possible trial solutions to be tested and a multitude of conflicting constraints making it hard to find the best one—arise in many important real-world problems. These include airline scheduling, pattern recognition, circuit wiring, packing objects of various sizes and shapes into a physical space or (mathematically similarly) encoded messages into a communications channel, and a vast multitude of others (including problems in logic and number theory that really are mainly interesting only to academic mathematicians).

These are all examples of what are called combinatorial optimization problems, which typically, though not always, arise from a branch of mathematics called graph theory. We'll discuss these kinds of problems in chapter 6, but what should be clear for now is that they have the property that the number of possible solutions (e.g., the number of possible tours in the TSP) grows explosively as the number N of input variables (the number of cities in the TSP) increases. Finding the best solution as N gets large may or may not be possible within a reasonable time, and one often has to be satisfied with finding one of many “near-optimal,” or very good if not the best, solutions. Whichever kind of solution one seeks, it's clear that some clever programming is required. For both algorithmic and theoretical reasons, these kinds of problems have become of enormous interest to computer scientists.

What have spin glasses to do with all this? As it turns out, quite a lot. Investigations into spin glasses have turned up a

number of surprising features, one of which is that the problem of finding low-energy states of spin glasses is just another one of these kinds of problems. This led directly from studies of spin glasses to the creation of new algorithms for solving the TSP and other combinatorial optimization problems. Moreover, theoretical work trying to unravel the nature of spin glasses led to the development of analytical tools that turned out to apply nicely to these sorts of problems. So, even in the early days of spin glass research, it became clear that they could appeal to a far greater class of researchers than a narrow group of physicists and mathematicians.

Spin glasses represent a gap in our understanding of the solid state.

Why is a crystalline solid (in which constituent atoms or molecules sit in an ordered, regular array) rigid? It may be surprising to learn that it wasn't until the twentieth century that we understood the answer to this question at a deep level.

Why is window glass (which does not have crystalline structure; the atoms sit in what look to be random locations) rigid? That's an even harder question, and you may be even more surprised to learn that we still can't answer that question at a deep level.

Of course, at what level you're satisfied with an explanation depends on your point of view: an answer that satisfies a chemist may not satisfy a physicist (and vice versa), and mathematicians are hard to convince of anything (so they're seldom satisfied). To be fair, at some level we've understood the nature of the solid state since the nineteenth century, when modern thermodynamics and statistical mechanics were developed by Gibbs, Boltzmann, and others. The basic idea is this. Atoms and molecules at close range attract each other, but they're never isolated from the rest of the

world; consequently, the constituent particles of a system always have a random kinetic energy that we measure as temperature. At higher temperatures, entropy (roughly speaking, disorder induced by random thermal motions) wins out, and we have a liquid or gas. At lower temperatures the attractive forces win out, and the system assumes a low-energy ordered state—a crystalline solid. Liquids and crystals are two different phases of matter, and the transition from one to the other, not surprisingly, is called a phase transition.

If you've taken introductory-level physics or chemistry courses you know all this. But there are deeper issues, which enter because there are features accompanying the ordered state that aren't so easy to explain. One of these is what Philip Anderson calls "generalized rigidity" [13]: when you push on the atoms in a crystal at one end, the force propagates in a more or less uniform manner throughout the crystal so that the entire solid moves as a single entity.

This is something we all take for granted. Why is it mysterious? Well, interatomic forces are short range and typically extend only about 10^{-8} cm, whereas when you push on a solid at one end, the force you apply is transmitted in a perfectly uniform manner a billion or more times the range of the interatomic force. How does that happen? What changed from the liquid state, where the exact same forces are present? At the very least, why doesn't the solid crumple, or bend? (And for that matter, what new phenomena need to be invoked to explain crumpling and bending when they *do* happen?)

This phenomenon isn't unique to solids; the transmission of forces over long distances also occurs, for example, in liquid crystals. In fact, this property is widespread in a general sense: it occurs whenever there's a transition to an ordered state that possesses a symmetry (whose form may not always be obvious) that differs from the thermally disordered state. Without

generalized rigidity, not only would solids not be “solid” but magnets wouldn’t be magnetic, supercurrents wouldn’t flow in superconductors, and we wouldn’t be here to observe all this. All these effects have similar underlying causes, but a deep understanding based on a small set of unifying principles didn’t arise until after World War II.

At this point those of you with physics or chemistry backgrounds might feel some impatience, and protest that the answer really isn’t all that complicated. If all atoms place themselves at the same distance from their nearest neighbors—that is, form a crystal—then they’ve created a very low energy state. Deforming this state would require a large input of energy, as anyone who’s ever tried to bend, tear, or deform a solid knows. This answer is perfectly correct, and is fine as far as it goes. But it’s unsatisfying at several levels. For one thing, as we’ll see momentarily, it fails as an explanation of why glass—which is not a crystal—is rigid. Nor does it explain the sharp discontinuity in rigidity behavior at the liquid→crystal phase transition. A few thousandths of a degree above the transition, there’s no rigidity; just below, there is. Wouldn’t a gradual change in rigidity as temperature is lowered make more sense? But that’s not what happens.

So there’s much that this simple answer leaves unexplained, and many interesting phenomena that it can’t by itself predict. Additional—and deeper—principles and concepts are needed.

Generalized rigidity is one of many examples falling within the category of “emergent behavior”: when you have a system of many interacting “agents”—whether they’re physical particles exerting mutual forces, or interacting species in an ecosystem, or buyers and sellers in the stock market—new kinds of behavior arise that for the most part are not predictable or manifest at the level of the individual. In the case we just discussed, something new happens when atoms rearrange to form a crystal; the ability to transmit forces over large distances is not present in the

fundamental physical interactions—in this case, the interatomic forces—that ultimately give rise to this effect. Rigidity must somehow arise from the collective properties of *all* the particles and forces: what we call long-range order (long range, that is, on the scale of the fundamental interatomic force) and broken symmetry (more on that later).

The idea of emergence is probably the most common underlying thread in complexity science, but as this example shows, emergence is not confined to complex systems. (At least the authors have never heard of table salt referred to as a complex system.)

But back to glasses. The problem here is that as far as we know, there is no phase transition (which in physics has a very specific meaning) from liquid to glass, no obvious broken symmetry, and no obvious long-range order (though a number of speculative candidates for these last two have been proposed). A glass is a liquid that just gets more and more viscous and sluggish as it's cooled, until eventually it stops flowing on human timescales. (By the way, that old nugget about windows in thousand-year-old European cathedrals being thicker at the bottom than at the top as a result of glassy flow over a thousand-year period isn't so. If you see a window with this feature, it had some other, more prosaic cause. Flow in window glass at room temperature would take place on timescales much longer than the age of the universe. Glass really is rigid.)

So why *is* glass rigid? As of this writing, there are lots of theories and suggestions, but none that is universally accepted. The problem is that a glass is a type of disordered system—the atoms in a glass sit at random locations. But much of our current understanding of condensed matter systems—crystalline solids, ferromagnets, liquid crystals, superconductors, and so on—applies to ordered systems with well-understood symmetries that enable profound mathematical simplifications and physical insights. So it's not only our ideas on rigidity (about which we won't have much more to say) that glasses challenge.

They challenge equally our understanding of many less familiar but equally fundamental properties of the condensed state. We'll encounter many of these as we go along.

An important clarification: when we talk about disorder in glasses, we're not talking about the kind you see in liquids or gases, where at any moment atoms are also at random locations. In those higher-temperature systems, the disorder arises from thermal agitation, and the atoms (or whatever individual units constitute the system) are rapidly flitting about and changing places. That enables us to do some statistical averaging, which in turn allows us to understand the system mathematically and physically. Glasses, on the other hand, are stuck, or "quenched," in a low-temperature disordered state, and so we can't apply the same set of mathematical and physical tools that we can apply to the liquid or gaseous state. And similarly, because of the lack of any kind of obvious ordering, we can't apply the same set of tools that we utilized to understand the crystalline solid state.

Spin glasses are also systems with this sort of "quenched disorder," but here the disorder is magnetic rather than structural. We'll explain this in more detail in chapter 4, but for now it's sufficient to note that spin glasses might provide a better starting point from which to develop a theory of disordered systems than ordinary glasses. That, and the fact that there's a gaping hole in our understanding of the condensed state owing to our lack of a deep understanding of systems with quenched disorder, is the reason why spin glasses have attracted so much interest among physicists and mathematicians.

Spin glasses display features that are widespread in complex systems.

So far we've indicated why mathematicians, physicists, chemists, computer scientists, and engineers might (or should) be interested in spin glasses. But complexity studies cast a wide net,

bringing in not only workers in these fields but also biologists, economists, and other natural and social scientists of various backgrounds and interests. What about them?

This is usually the point where treatises on the subject attempt to provide a working definition of complexity. We won't attempt that here, and not only because we don't know the answer. After many years, there still is no universally accepted definition of complexity, or of how to determine whether a given system is "complex." This is not for lack of trying, and many people (including one of us) have made proposals.

But it's not our goal, and certainly not the purpose of this introduction, to concisely define complexity.³ That purpose, aside from the usual one of acquainting the reader with some basic ideas and concepts, is to convince her or him that it's worth investing some time to read the rest of the book. If you're still with us, then in your case we haven't yet obviously failed, but you may still be wondering whether all this has any relevance to your own field of interest. So we'll now take a look at some of the broader impacts and applications of spin glass theory.

We'll begin the discussion by asking, what kinds of systems are generally agreed upon (even in the absence of a definition) to be complex?

A far from exhaustive list might include a wide variety of adaptive systems or processes, systems that exhibit pattern formation, scale-free systems or networks, systems with a modular or hierarchical architecture, and systems generating or incorporating large amounts of information. Of course, some of the most interesting complex systems display several or all of these features at once.

We'll briefly discuss a few examples of each. Adaptation occurs in many contexts: biological evolution, ecological networks, the immune system, learning and cognition in biological and artificial systems, adaptive computer algorithms, economic and social

systems, and many more. Biological pattern formation occurs at the cellular level in morphogenesis, at the organismal level in zebra stripes and butterfly wings, and at the group level in schooling fish. Nonbiological patterns occur in cellular automata, or in physical systems far from equilibrium, such as the regularly spaced ripples that occur in sand dunes or in “cloud streets,” or the oscillations that occur in certain driven chemical systems. “Scale-free” systems exhibit similar-looking structure, phenomena, or behaviors on many length- or timescales, not just one or a few. The canonical example of this is the appearance of vorticity at multiple lengthscales in turbulent fluid flow, but scale-free behavior or structure in one form or another characterizes many complex systems and networks, whether physical, biological, or social.

Almost all systems regarded as complex are out of equilibrium (defined appropriately for the system in question) and maintain themselves at the boundary between rigid ordering (as in a crystal, where not much change can occur) and chaotic flow (where, so to speak, too much change is occurring, so that no coherent ordering, evolution, or adaptation can take place). This is sometimes referred to as being at the *edge of chaos* [16, 21, 22]; these systems maintain a delicate balance so that an ordered structure can be maintained while growth, evolution, and adaptation can still occur.

Many complex systems, particularly those that are the result of some kind of evolution (biological or otherwise), are hierarchically structured. A full-blown complex structure cannot spontaneously arise all at once; a modular architecture enabling a gradual increase in complexity is needed. And finally, the above discussion implies that all complex systems possess or generate a large degree of *information*, whether in the Shannon, algorithmic, or other sense.

Perhaps one of the most important unifying features of complex systems, and one that doesn't get mentioned often enough, is that many of them surround us in the everyday world; sometimes they are even a part of us, such as our own brains or immune systems. Unlike many other systems investigated by scientists, they're typically not obscure or esoteric or known only to a small group of specialists or experts. And as a corollary, they are not idealized in any sense: they're real-world, messy systems, inspiring difficult questions that usually don't fit neatly into any one scientific category, such as biology, chemistry, or physics. They transcend disciplines, and consequently their understanding usually requires transdisciplinary collaborations and insights. All scientific problems are complicated, but only some are complex.

So, where do spin glasses fit into all of this? It's probably already apparent that spin glasses don't adapt in any usual sense of the word, nor do they form any obvious patterns. They don't evolve, change, or learn. Mostly, spin glasses just sit there.

In that case, how can they provide insights to those interested in any of the problems we just mentioned?

There are two broad classes of answer to this question. One class involves observed spin glass behaviors in the laboratory that are reminiscent of some of the features discussed above and that remain poorly understood theoretically. The other involves theoretical constructs that may in the end have little to do with real spin glasses in the laboratory (though they may—as we'll discuss in chapter 7, this remains a topic of controversy) but that nevertheless are very suggestive of general features of complexity. In the first case, we have experiments in search of a theory; in the second, theory in search of experiments. Both may have more relevance to complexity studies than to each other.

One of the few things that everyone agrees on is that spin glasses are systems with both quenched disorder and frustration. We've already discussed disorder, which in one form or another is common in complexity—it's hard to imagine a complex system that is perfectly regular in any simple sense. Frustration refers to the presence of numerous constraints that conflict with each other so that not all can be simultaneously satisfied. This is clearly something that should be a universal feature of complex systems, but it was in the study of spin glasses that the idea first crystallized, was put on a mathematical footing, and developed.

Disorder and frustration often go together, but they refer to distinct concepts, and neither implies the other. In many cases—even that of structural glasses—it's not so easy to derive crisp mathematical formulations of these properties, as they apply to the system at hand. In spin glasses, this can be done readily, and so they provide a well-defined mathematical laboratory for the exploration of these concepts and their possible implications and consequences.

But it wasn't just disorder and frustration that made spin glasses so useful. They also exhibited a number of other features that many complex systems display. Moreover, in spin glasses these features arise naturally and spontaneously from a minimalist starting point; they don't have to be inserted "by hand." Such features include the presence of many near-optimal solutions, which we've already seen figures prominently in combinatorial optimization problems. In the case of spin glasses, these "solutions" refer to low-lying (in energy) metastable (or possibly thermodynamic) states. They include the generation of information (in the Shannon sense) in the selection of particular outcomes when a spin glass is cooled or an external magnetic field is removed. They include a novel and exotic hierarchical ordering of states that spontaneously emerges in at least one nonphysical

(but important) model of spin glasses. And finally, spin glasses inspired the development of new mathematical techniques, which may be applicable to other kinds of complex systems, to describe and perhaps explain all this.

Consequently, starting in the early 1980s, spin glass concepts, ideas, and mathematical tools were applied to problems in neural networks, combinatorial optimization, biological evolution, protein dynamics and folding, and other topics current in biology, computer science, mathematics and applied mathematics, and the social sciences. Some applications were reasonably successful, others less so. We'll meet some of them in chapter 6.

On the experimental side, spin glasses show some very peculiar and interesting nonequilibrium behaviors. Of these, two that are potentially most relevant for other complex systems are, first, the presence of a wide range of intrinsic relaxational or equilibrational timescales, and second, the observation of memory effects: a spin glass is able to "remember" certain features of its past history in a rather remarkable way.

Many of these properties (and others that we'll encounter as we go along) are widespread throughout complex systems from many fields. What makes the spin glass so special is that these properties all seem to arise, in one way or another, from a very simple-looking energy function that can be written down in one line. It is deeply surprising that this should be so, and we have yet to understand what general complexity principles, if any, can be learned from this. But at the very least, the emergence of all these properties, and the mathematical techniques that arose to describe them, have proved useful in a wide range of studies that go far beyond the original problem of understanding an obscure class of magnetically disordered systems.

So it may well be the case that in learning something about spin glasses, you might uncover some new insights and tools to

help you better understand your own system of interest. And very possibly, even if you don't, it might still be entertaining to learn how so many new, fundamental, and useful concepts can arise from studying such an initially boring-looking and unpromising system—which is where we started the discussion.