Spin Glasses, Disordered Systems, and Complexity

Daniel L. Stein*†

February 14, 2022

By the early 1970's physicists had synthesized a comprehensive mathematical and conceptual framework for understanding microscopic and bulk properties of "organized matter" (now commonly referred to as condensed matter). Condensed matter systems are assemblies of a vast number (typically of order 10^{24}-10^{25}) of interacting atoms or molecules, and include materials as diverse as crystalline solids, magnetic systems, superconductors, liquid crystals, and many others. All share the property of emergence: the collective behavior of the system as a whole cannot be readily predicted or understood by our knowledge, no matter how complete, of the individual, isolated constituents. (See the article "More is Different" elsewhere in this volume.)

This synthesis, which underlies the modern field of condensed matter physics, rests on three foundations. Two of these, quantum mechanics and thermodynamics (including statistical mechanics), are well-established subjects taught in all modern-day undergraduate physics curricula. The third — symmetry — is less well-known but has played a central role in our modern understanding of the physical universe. For our purposes, it will be sufficient to note that many of the emergent behaviors of condensed matter systems — the solidity of lead, the magnetic behavior of iron, the superconductivity of mercury — are intimately related to a change in the organizational symmetry of their microscopic constituents at a phase transition, such as that from water to ice. Each molecule in liquid water is free to travel anywhere within the liquid, and the positions of its molecules at any moment appears random; in contrast, each molecule in solid ice is confined to a very small region of space and the overall organization of molecules in ice forms a regular, repeating structure called a crystalline lattice.

Despite the (very impressive) successes of modern condensed matter physics, there remains a significant gap. The glass in your window comprises molecules that do *not* form a crystalline array: like the liquid, their positions in space appear random, but unlike the liquid, each molecule is stuck or "quenched" in a very small region of space. There is no change in symmetry of the arrangement of molecules, indeed no phase transition at all, between liquid and glass. And yet to us glass appears quite solid. Glass provides an example of "quenched disorder", so named to distinguish it from thermal disorder, such as occurs in liquids or gases where each molecule is free to fly through space. Thermal disorder, in which matter is not fully condensed, is well-understood, thanks to the development of modern thermodynamics. But quenched disorder is not[1], owing to the absence of any recognizable symmetry that can be exploited to aid understanding and analysis. This gap led Philip Anderson, one of the founders of SFI, to characterize the problem of glasses as "... the deepest and most interesting unsolved problem in solid state theory."

Glasses are the most well-known but not the only example of condensed matter systems with quenched disorder. A second category comprises a class of systems with quenched *magnetic* disorder, in which

*Department of Physics and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012 USA; NYU-ECNU Institutes of Physics and Mathematical Sciences at NYU Shanghai, 3663 Zhongshan Road North, Shanghai, 200062, China; and Santa Fe Institute, 1399 Hyde Park Rd., Santa Fe, NM USA 87501.

†Email address: daniel.stein@nyu.edu

[1] Here I refer only to the sorts of questions physicists ask. Glasses have been around for thousands of years, and are well-understood from other perspectives, such as those of chemistry, engineering, and materials science.

some of the constituent atoms (which generally live on a crystalline lattice, and are therefore confined to a small region of space) possess intrinsic magnetic moments; that is, they behave like atomic-sized bar magnets. These atomic magnets (or simply "spins" for short) interact with each other in a sufficiently complicated fashion such that at low enough temperatures the spins "freeze" in random directions. In analogy with the structural glasses described above, these are called "spin glasses", but it is important to keep in mind that in a spin glass the disorder is in the *orientations* of the atomic magnets which are fixed in space, whereas in a glass it is the *positions* of the fixed atoms that are random.

To the physicist, spin glasses possess an advantage that glasses do not. Glasses are disordered condensed matter systems which are out of thermodynamic equilibrium: that is, the properties of a piece of glass may depend not only on its present state but also on the history of its preparation. In contrast, the spin glass may be (in fact, likely is) a new phase of matter which *is* in thermodynamic equilibrium, having a sharp phase transition (analogous to the water $\rightarrow$ ice transition) from a high-temperature phase where the spins are free to (rapidly and continually) change their orientations (while the atoms themselves remain in place) to a low-temperature phase where the spins are frozen in random directions. Any solid that can become a glass can also become a crystal (which is the state of lowest energy), depending on whether the cooling from the liquid phase is fast (glass) or slow (crystal). Not so for a spin glass — as long as the atomic positions remain fixed, there is no competing, lowest-energy state where the spins either all align or else orient themselves in some ordered, repeating fashion. Hence the initial hope that the spin glass would be far more amenable to theoretical study than the ordinary structural glass. And if so, could understanding spin glasses be the key to understanding condensed matter systems with quenched disorder more generally?

Following experiments in the early 1970's which indicated that spin glasses exhibited some very puzzling behavior, Sam Edwards and Philip Anderson in 1975 proposed a simple model of spin glasses based on an energy function (or Hamiltonian in physics jargon) whose input is a list of the orientation of every individual spin in the entire system, and whose output is the energy of that particular configuration of spin orientations. As in real materials, in the Edwards-Anderson (or simply EA) model each spin interacts directly only with relatively nearby spins, and the interactions between pairs of spins is such as to try to force the spins to align for some pairs and anti-align for other pairs. EA further proposed that the degree of spin glass ordering can be quantified by the amount of spin freezing, i.e., by how sluggishly each spin responds to thermal agitation which can flip or alter a spin orientation. Unlike the glass, whose Hamiltonian is no different from the crystal's, in the EA model the spin glass has quenched disorder built in to its Hamiltonian from the start.

The EA model represented a major step forward, but the presence of quenched disorder in the nearest-neighbor spin interactions of the EA Hamiltonian presented a major roadblock to further theoretical progress. To get around this, one can try to construct a model which is easier to analyze. In statistical physics, such models often take the form of a "mean field theory", in which (in our case) each spin interacts equally strongly (at least in an average sense) with *every* other spin in the system. Though such an unphysical approach usually does not accurately predict behavior near a phase transition, it often provides an excellent guide to understanding the system's properties at temperatures well below the transition[2], particularly the nature of its ordering and symmetry. (For abstruse mathematical reasons, mean field theory correctly depicts the behavior at all temperatures of the system in *infinite* dimensions!)

Hence, only months after the appearance of the Edwards-Anderson paper, a mean-field version of their Hamiltonian was proposed by David Sherrington and Scott Kirkpatrick (see the corresponding article in this volume). They were able to solve their model (now known as the SK model) using a trick suggested in the Edwards-Anderson paper. The idea is to (theoretically) examine many identical copies,

[2]At least until quantum effects become important. The theory described here is purely classical, and at extremely low temperatures a classical description becomes inadequate for any system. For many laboratory spin glasses, however, quantum effects become important only at ultralow temperatures, and the classical theory described here remains valid for most temperatures below the transition.

or replicas, of the system, average over the quenched disorder in this many-replica system, and finally take the limit in which the number of replicas goes to zero (!). The reasoning behind this so-called "replica trick" is — or was initially — mathematical rather than physical; it was introduced solely as an abstract prescription for averaging over the quenched disorder in the absence of any other viable method.

Using this replica trick, along with the Edwards-Anderson conjecture on characterizing the ordering at low temperature, Sherrington and Kirkpatrick solved their model and found a sharp phase transition from a high-temperature "unfrozen"[3] phase to a low-temperature spin glass phase. The high-temperature solution looked fine, but the low-temperature solution — which describes the spin glass phase and is therefore the one of primary interest — became thermodynamically unphysical at very low temperature. As a consequence it couldn't be a correct description of the spin glass phase. Because the replica approach was new and poorly understood, it was initially supposed that the problem lay with it.

For most systems — in particular those without quenched disorder — mean field theories are relatively easy to solve once the change of symmetry in passing from a high-temperature to a low temperature phase is identified. The mean-field theory of spin glasses —that is, the SK model — presented a departure from the norm in that four years passed before the problem with SK's solution was remedied. During that period a number of attempts to solve the model were made, with two in particular having a lasting impact. The first, by David Thouless, Anderson, and Richard Palmer studied spin glass mean-field-theory using a different approach, deriving self-consistent equations that at low temperature had many solutions, potentially corresponding to many low-temperature spin glass phases with different frozen orientations of the spins. The second, by J.R.L. deAlmeida and Thouless, examined the stability of the phases found by SK and suggested that the problem could lie in the equal treatment of all replicas in SK's application of the replica trick at low temperatures. Not surprisingly, why this should be problematic and how one might treat the different replicas in an *inequivalent* fashion remained mysterious.

This was the state of affairs in 1979 when Giorgio Parisi published his solution of the SK model at low temperatures, which is the subject of this Introduction. As one can see from perusing the accompanying paper, it is highly technical mathematically, and its physical meaning and implications took another half-decade or so to be worked out, largely in a series of papers by Parisi and collaborators. One can fairly say that the mean-field-theory of spin glass took roughly a decade from the time of its introduction to be fully worked out and understood: even at the mean-field level the problem is exceedingly difficult. For this work Parisi was awarded a share of the 2021 Nobel Prize in Physics.

As just noted, it took roughly half a decade to decode Parisi's solution from mathematics to physics. As originally conceived, the replica trick involved the (theoretical) creation of n independent, identical replicas of a system and then taking $n \to 0$. This procedure implicitly assumes that there is nothing to distinguish one replica from another: all are equivalent. More specifically, suppose there are n replicas. Then one can construct an $n \times n$ matrix $Q_{\alpha\beta}$ whose off-diagonal elements contain information on the similarity between replica α and replica β. Because these replicas (in those early days) were considered nothing more than a mathematical artifice, there was no reason to expect the relation between any pair of replicas to be any different from the relation between any other pair. Mathematically this meant that all off-diagonal elements of $Q_{\alpha\beta}$ should be equal; the invariance of the elements of $Q_{\alpha\beta}$ with respect to permutation of its indices is now referred to as "replica symmetry". Parisi's (mathematical) solution was to "break" the symmetry among the replicas in a hierarchical fashion: in the first step $Q_{\alpha\beta}$ is broken into submatrices, or "blocks", such that within each block, the off-diagonal matrix elements in the diagonal subblocks have one value while those in the off-diagonal subblocks have a different value, and in each subsequent step each subblock is further broken down in a similar fashion. This led to the characterization of Parisi's solution as *replica symmetry breaking*, or RSB for short.

[3]Referring to the spin orientations, not to the atomic positions themselves, which remain fixed in space.

This mathematical procedure was an inspired guess with (at the time) no physical justification. For the first time, however, a solution was found that behaved sensibly at all temperatures.

But what does this solution tell us about the nature of ordering in spin glasses? Physically, it means the following: first, rather than there being a unique[4] spin glass phase, there are an *infinite* number (for a system comprising an infinite number of spins) of possible low-temperature spin glass phases accessible to a single system, each corresponding to a different configuration of frozen spin orientations. Instead of trying to characterize each of these, one studies the *overlaps* between randomly chosen *pairs* of configurations. The overlap between two configurations is their degree of similarity: roughly speaking, for a system in which the spins can point in only two directions (say north and south), it is the number of spins whose time-averaged orientation points in the same direction in both replicas minus the number pointing in the opposite direction. Most surprisingly, the organization of all of these different possible spin glass configurations is highly structured: they are arranged in a hierarchical fashion, the overlap or similarity between configurations organized much the same as kinship relations in a family tree.

This is all rather astounding, for any number of reasons. Most significantly, the mean-field spin glass exhibits a wholly unanticipated and exotic kind of ordering — or symmetry breaking, as physicists call it, emphasizing the change in symmetry in going from the simple, unordered high-temperature phase to the low-temperature one. Especially surprising is the intricate organization in the relations among the individual spin glass states in the low-temperature phase: a rich and complex structure arises from a simple-looking, featureless Hamiltonian with built-in quenched disorder. (Indeed, part of the Nobel Committee's scientific press release on the occasion of the awarding of the 2021 Physics Nobel Prize celebrated the arising of "order from disorder" in Parisi's solution of the SK model.)

Once the SK model had been solved, the natural question became, is replica symmetry breaking a universal feature of the physics of quenched disorder? Despite decades of intensive research by numerous workers the scientific community has not yet arrived at a consensus, even for spin glasses themselves. Recall that mean-field theory is an accurate description of a system in infinite dimensions, though it usually provides a good guide to the structure of the low-temperature phase in finite dimensions as well. But systems with quenched disorder present any number of thorny mathematical issues that are not well understood, and there is no guarantee that they will behave as "nicely" as non-disordered systems. There are several competing pictures describing the possible thermodynamic structure of real three-dimensional spin glasses in the laboratory, and while each have their supporters, none have received universal acceptance. The problem remains open at present.

What about ordinary glasses? Here too there has been progress over the past few decades, but just as in spin glasses, no universal agreement on the correct picture. It should be noted, though, that one prominent picture of glasses (called the random first-order transition theory) has attracted considerable attention and incorporates many of the ideas and techniques of replica symmetry breaking.

But problems involving quenched randomness are not limited to physical systems. Starting in the early 1980's and continuing to the present, replica symmetry breaking and, more generally, spin glass physics as a whole, found an enormous array of applications in a wide variety of fields outside of physics, in particular computer science and biology. These applications include new algorithms, advances in combinatorial optimization problems (e.g. airline scheduling, encoding information in communication channels, and many others), neural networks, machine learning, protein structure and dynamics, and much more. (See the papers by Hopfield and Kauffman and Levin in this volume.) This spin glass "cornucopia", as Anderson once called it, was noted by the 2021 Nobel Committee as a central contribution of Parisi's work, and justifies characterizing the 2021 Physics Nobel Prize as one that recognizes, possibly for the first time, the importance and centrality of complex systems research in the modern scientific era.

[4]Technical note for physicists: up to a global symmetry transformation.