Chapter 1

Contribution to PWA90
Spin Glasses and Frustration

Scott Kirkpatrick
School of Engineering and Computer Science,
Hebrew University
91904 Jerusalem, Israel
kirk@cs.huji.ac.il

The 1975 EA paper by Sam Edwards and Phil Anderson proposed to understand a class of real materials as frustrated, glassy magnets with a novel type of weak order in time. It launched an exciting period of intense exploration of new directions for understanding random and composite materials. This work has had a lasting impact well beyond materials science, contributing powerful new methods for optimization of complex many-parameter systems as well as codes offering Shannon-bounded communications in extremely noisy channels.

1. Origins of the EA (and SK) papers.

This paper will start with a bit of personal memoir of some exciting times in condensed matter physics, the late 1960's and the 1970's, in which the paradigm for describing real materials – alloys, compounds, composites and glasses – shifted from a search for the idealized uniform "coherent potential" or "effective medium" to an effort to tease out the novel consequences of quenched disorder, such as percolation, localization, and glass transitions into amorphous but highly stable structures. These regimes, where 30+ years of conventional theory had broken down, seemed to be where Phil Anderson felt most alive, so of course he played a highly visible and central role in all of this. Other speakers cover percolation and localization; I get to talk about spin glasses. In these first sections I will stick to the period 1974 to 1978 and a bit later, in which the basic picture evolved rapidly. A final section covers the many ways in which ideas and technical methods that arose in this period of ferment are being employed today.

Alloy a few Mn atoms in a matrix of pure Cu (or almost any transition metal in a noble metal host) and you get long-ranged magnetic interactions, oscillating in sign, a Kondo effect, and a puzzling set of possibilities for the magnetic structures that might be found at higher densities. Increasing the concentration of the transition metal component to a few per cent, as crystal growers at Bell Labs, University College London, Grenoble's CRTBT and others soon did, yielded anomalous thermal and magnetic susceptibilities with cusps rather than sharp peaks, and magnetic hysteresis effects on very long time scales. For the experimental facts and debates, see Joffrin[1] for a view at the time, and Mydosh[2] for a view from a few years removed. Either Phil or Brian Coles (accounts differ on this) suggested the name "spin glass," to capture the slow hysteretic effects of these materials and others with mixtures of ferromagnetic and antiferromagnetic interactions. That name had a strong appeal.

The Edwards-Anderson (EA) paper[3] captured the main ideas which underscored that appeal. They made explicit their goal of studying the spin glass in the hopes of learning more about the transition in structural glass. First, they introduced an order parameter which measured stability as the persistence of the value of a spin, without specifying the details of the spins' arrangements:

$$q = (1/N) \sum_i << S_i >_T < S_i >_T >_{disorder}$$

Note the separation of the thermal average (the long term stability in the presence of the sources of randomness, quenched in fixed positions) from the average $< ... >_{disorder}$ over the possible positions of the magnetic atoms. The EA paper then brought together a set of technical tricks that could lead to a mean field theory for the onset of such order. The first was a replica construction in which the spin average above is performed by comparing two or more identical but non-interacting copies of the disordered material under study. Second was the trick of evaluating thermal averages of the free energy of such a quenched random system by taking the limit as $n \to 0$:

$$< < logZ >_T >_{disorder} = < lim_{n \to 0} < Z^n - 1 >_T >_{disorder} .$$

The results made it seem obvious that, at least in mean field theory, such a model could exhibit a phase transition into a new spin glass phase.

Phil later described[4] these manipulations as "a mathematical trick so hoary that its origins are lost in history." They can be found, for example, in a 1959 paper on dilute ferromagnets by Robert Brout[5] which appears to have been written while he was spending a summer at Bell Labs. (Thanks to Giorgio Parisi for the reference.)

I read the EA paper as soon as David Sherrington showed up at IBM Watson Labs for the summer of 1975. He had been in contact with Sam Edwards during the EA paper's gestation, and had an idea. We could keep the general approach of EA, but change the model to make the whole thing "solvable," and thus end up with a rigorous mean field limiting version of the actual problem. Mn and Fe ions in noble metals look like idealized Heisenberg spins, but Ising spins are easier to deal with. RKKY interactions are complicated, so like EA, we just replaced them with a Gaussian distribution of exchange constants, but let the range become infinite – spherical models, with N spins each interacting with the other $N - 1$

spins are solvable when ferromagnetic, right? And the classic spin glasses become ferromagnets when the transition metal concentration exceeds about 10%, so we added a term favoring ferromagnetic alignment, and ended up with the generalized Hamiltonian

$$H = \sum_{all i,j} J_{ij}(S_i S_j)/\sqrt{N} + J_0 \sum_i S_i/N$$

where the J_{ij}, from a Gaussian distribution with unit variance, are scaled as $1/(\sqrt{N})$ and J_0 by $1/N$, to keep things finite and in balance. After deriving a solution along the same lines as in EA enough times that we thought we had it right, we were left with a nonlinear equation for the free energy which could be solved at its critical points, around its ground state, and numerically in between. We made the classical assumption that evaluation of the partition function for n replicas of our model would be dominated by a single stable minimum or maximum about which a stationary phase expansion would capture the important contributions. This gave a plausible phase diagram, seen in Fig. 1, a susceptibility and specific heat with cusps, as seen in Fig. 2, and resembled much of the experimental data.[6]

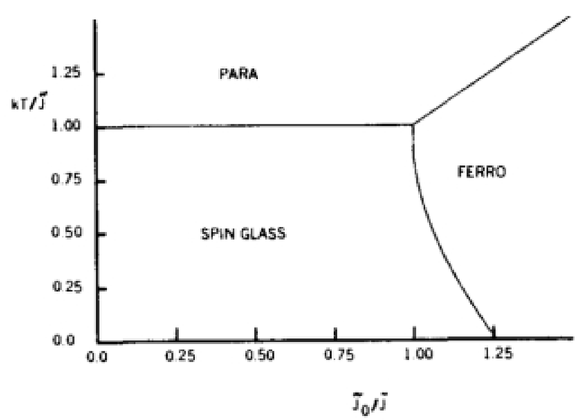

Fig. 1. Phase diagram of a "solvable" spin glass.

There was, of course, one catch. The solution gave values for the ground state energy

$$E(T = 0)/N = -1/(2\pi) \approx -0.79$$

and the entropy

$$S(T = 0)/N = -\sqrt{2/\pi} \approx -0.17.$$

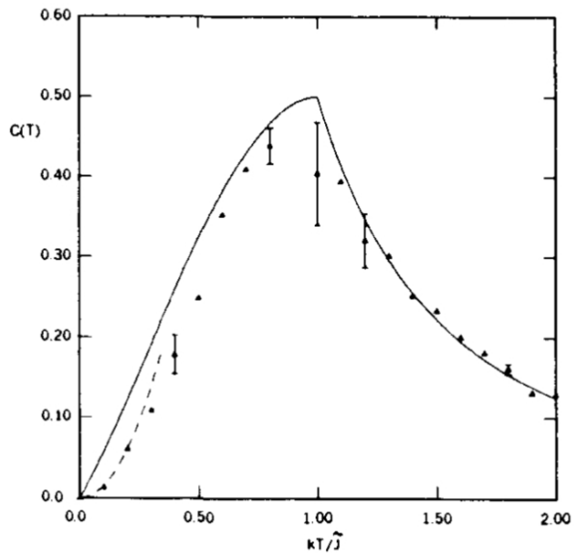

Fig. 2. Specific heat of a "solvable" spin glass (with computed values from samples with 500 spins).

The first was plausible (it took two years of work on computer simulations at the speeds then available to show that it is too low by about 0.03), but the second result was unacceptable for an Ising model, whose discrete excitations should be frozen out at zero temperature. So we inserted warning labels on this part of the results, wrote a very understated abstract and sent the paper off. The physics seemed right, and the evidence for a phase transition out of the high temperature paramagnetic phase into, well, something pretty interesting, was solid.

2. Cavities, TAP equations, and all that...

The two papers were well received. Experimentalists whose data looked like our curves were happy, and theorists had a fat target to shoot at. A workshop was quickly arranged to be at Aspen for summer 1976, with PWA, David Thouless and Richard Palmer at the center of things. The main objective was to find methods of solving this class of problems in mean field theory without the introduction of replicas and the very suspicious $n \to 0$ limit. Attention focused on developing "cavity methods" in which the force on the spin at a site is calculated by summing the effects of all its neighbors in the "cavity" left by removing that spin. The spin's response is then fed back into the problem for subsequent iterations. This idea for extending mean field theory goes back to Bethe, is now called "message passing" by statistical physicists, and is also known as "belief propagation" in the statistics literature.

My recollection is that in discussions over beer in Aspen the problem was solved multiple times on multiple evenings. Each evening resulted in an "Aspen preprint." In the light of the next day, each such preprint was found to have problems, calling for another day of discussion, another preprint. It was an exciting time. Phil was wont to mutter with satisfaction, as each preprint improved over its predecessor, "And gentlemen at Harvard, now abed, shall think themselves acurs'd they were not here." The results appeared in the Oct. 1976 TAP paper,[7] and gave different low temperature behavior, in rough agreement with the computer simulations on the model which were then becoming available.[8] See the dashed line for the low temperature specific heat extrapolation in Fig 2.

Shortly after the Aspen workshop ended, Thouless and de Almeida found the fatal flaw in the replica calculation. It wasn't the $n \to 0$ limit, nor the several exchanges of the order of evaluation, but simply that the n-replica calculation of the partition function was not to be calculated around a single fixed point, since that point was neither a minimum nor a maximum. Replica symmetry had broken.[9] It took about six more years and a series of attempts of increasing complexity to develop a convincing hierarchical model of replica symmetry breaking for the SK model's low temperature phase[10] with interesting consequences and twenty-plus years beyond that to obtain a rigorous justification of the main features of that proposal.

But the idea of gathering researchers and students, experimenters and theorists, interested in glasses, random magnets, percolation, localization and conduction into one community had taken hold. Phil headlined an influential European summer school that lasted 8 weeks at Les Houches in 1978, taking a break from the stress of receiving a major award over the previous winter. He covered all these subjects from the perspectives of single-particle spectra and many-particle effects in a magnificent series of lectures.[4] Because of its breadth of coverage and extreme duration, the summer school came to be known as "Camp Houches," a basic training experience for a whole generation of European (and some American) scientists still active in the study of complex and disordered systems.

3. The spin glasses of today...

At this point the further evolution of the spin glass project begins to split into several distinct routes. The high road leads upwards to the various possible forms of replica symmetry breaking and their consequences. Less rarified paths lead through problems in combinatorics that also have natural large scale limits while representing problems with real world relevance. And orthogonal directions have now used message passing techniques stimulated by the TAP equations to create powerful codes, a large subject that I won't attempt to discuss here. (But you can learn about them, described from this perspective, in the recent book by Mezard and Montanari[11]).

In a series of papers by Parisi, various collaborators along the Paris-Rome axis,

Fig. 3. Resting at Camp Houches, 1978

and others, increasingly elaborate schemes for breaking replica symmetry and evaluating $<< \log Z >>$ were explored, starting with one-step symmetry breaking in which the replicas were divided into m groups, with one order parameter found between different groups and another within a group. This was not good enough for the SK model, although it seems to work for certain problems in combinatorics and constraint satisfaction. The ultimate form of replica order parameter for the SK model is a continuous function $q(x)$ on the line $0 \leq x \leq 1$, representing a continuous distribution of quasistable states of low energy. A surprising consequence emerged from this full RSB approach. Correlations between these "pure" states are ultrametric. If we calculate the pairwise overlaps between three states, A, B and C, e.g.:

$$Q^{A,B} = (1/N) \sum_i S_i^A S_i^B$$

then at least two of the overlaps, $Q^{A,B}$, $Q^{A,C}$, and $Q^{B,C}$, are equal. This is most simply understood as suggesting a clustered, tree-like structure in which sets of "pure" states are derived from one another. Since in a glass or a spin glass, we generally can observe only one state at a time, the experimental consequences of ultrametricity are unclear. But in large scale problems of searching or optimizing networked objects, multiple solutions can be obtained, so the predictions of clustering or ultrametricity prove important for establishing heuristics or characterizing the space of possible solutions and algorithm performance in that space.

The cavity approach, and the directed message passing suggested by the TAP equations, has had perhaps its widest application in the optimization and searching of complex systems, such as information networks. One simple problem has proven to be replica symmetric. Consider the simplest possible metric-free travelling salesman problem (TSP). Let the upper triangle of the NxN matrix of distances $d_{i,j}$ between each pair of N points be filled by random numbers from some distribution. The uniform distribution on the interval $[0, 1]$ proves most convenient. We symmetrize all distances so that $d_{i,j} = d_{j,i}$ and then ask for the shortest continuous tour of links connecting all sites without visiting any site more than once. A greedy solution, at each step choosing the link to the closest remaining unvisited site, leads to a tour length that diverges as $\log N$. Both a replica-symmetric Ansatz[10] and a cavity approach,[12] however, predict a constant limiting tour length as $N \to \infty$. The cavity solution is more easily evaluated, and predicts that the optimal tour will have length 2.0415... in the large N limit, in agreement with numerical experiments.

The environment in which message-passing methods have most obviously burgeoned is search on the Internet. Its scale today is set by the earth's population, a few billion. In the near future, an Internet filled by information from the world's sensors may increase in scale by many more decades. Finding the answers most relevant to a particular query in this vast space requires pre-conditioning each search with information passed selectively along the same network that is being searched. Today's search engines rely upon several families of methods for accumulating these hints, whether by identifying "hubs" (sources of many useful links) and "authorities" (generators of links most likely to be relevant to the specific query), or by ranking pages by the degree to which other pages link to them. For a good overall review of these approaches see Easley and Kleinberg's recent book.[13]

Searches for groups of similar objects or for communities with similar interests within this large space are important in building systems that make reasonable reccomendations or that establish trust. They may need to operate on scales of order N (the world's or a nation's population), $\sqrt{N}$ (a few hundred thousands) or even smaller. A classic problem to use in idealizing such a search is the problem of finding the largest clique (completely connected subgraph) in a simply defined random graph, such as the Erdos-Rényi $G(N, p)$ with N nodes and a fraction, p, of its bonds present. The realization that for, say, $G(N, 1/2)$ as $N \to \infty$, the largest clique that will be found has exactly $2 \log_2 N$ nodes was an early example of the recognition that sharp phase transition-like behavior occurs in large mathematical objects as well as in materials.[14] Finding even one such clique (there will be many) is extremely hard, since a greedy search will run out of nodes to search among once the clique found is of order $\log_2 N$, and no stronger search methods are known. So in this problem, an SK-like model has been identified and extensively studied. One simply "plants" a clique of size $\sqrt{N}$ in the graph and asks for an algorithm that can find it, preferably in time proportional to the number of links in the graph. At the moment, the best algorithm for this search uses a very simple form of message-passing. Deshpande and Montanari[15] use the power method, raising the graph's adjacency matrix to a small power, to find its largest eigenvalue and the associated eigenvector. The large components of that eigenvector are the elements in the planted clique. I am currently seeing if this trick of using messages to amplify the weak signal that a clique gives off can be pushed down to the $\log_2 N$ scale where naturaly occurring cliques are found, by combining the power method trick with a steady reduction in the scale of focus.

Finally, an important indirect consequence of the spin glass project (to me, at least) is the realization of simulated annealing as an archetype of a wide range of physically-motivated methods of searching for optimal solutions to complex non-convex problems with many parameters. While spending a year or so searching among the SK model's low energy configurations to determine ground state energies and spin relaxation dynamics, it became perfectly natural for me to let samples evolve at finite temperatures using the Metropolis procedure for the dynamics. Heating the sample up and slowly cooling it was a natural way to improve on ground state energy estimates and to obtain good statistics. In subsequent years, it was equally natural for a small group of us at IBM interested in designing good automatic tools for placing and routing computer components during the design process to view our problems as medium sized spin glasses. There was frustration due to the conflict between the need to optimize performance by putting circuits as close to one another as possible, and the constraints of minimizing heat generation and making wiring possible, both of which tended to force circuits apart. When we sent off a paper summarizing half a dozen such examples and stating this approach as generally as possible, we encountered a certain amount of resistance to our unconventional subject and intuitive approach. Phil, however, fully appreciated the ideas, liked the paper, and was instrumental in getting it accepted in Science.[16]

Acknowledgments

The spin glass community is very much alive today, well into its third generation. While preparing this article I had the pleasure of taking part in 2014's "Spin Glasses: an old tool for new problems," at the Institut d'Études Scientifique de Cargèse, organized by Lenka Zdeborova, Federico Ricci-Tersenghi, Florent Krzakala, and Giorgio Parisi. Work reported there ran the gamut from studies of glasses (jamming and dynamics in hard sphere models) to information theory in engineering, computation and biology to rigorous proofs, some recently obtained, of replica solutions to the originally-stated problems.

4. Bibliography

References

1. J. Joffrin, Disordered Systems – Experimental Viewpoint, in eds. R. Balian, R. Maynard and G. Toulouse, *Les Houches 1978: Ill-Condensed Matter* (North-Holland, Amsterdam, 1978) 63-157.
2. J. A. Mydosh, *Spin Glasses: An Experimental Introduction* (Taylor and Francis, London, 1993).
3. S. F. Edwards and P. W. Anderson, *J. Phys. F5*(1975)965.

4. P. W. Anderson, Lectures on Amorphous Systems, in eds. R. Balian, R. Maynard and G. Toulouse, *Les Houches 1978: Ill-Condensed Matter* (North-Holland, Amsterdam, 1978) 159-261.
5. R. Brout, *Phys. Rev 115* (1959) 824-835.
6. D. Sherrington and S. Kirkpatrick, *Phys. Rev. Letts. 35* (1975) 1792.
7. D. J. Thouless, P. W. Anderson, and R. J. Palmer, *Philos. Mag. 35* (1976) 593-601.
8. S. Kirkpatrick and D. Sherrington, *Phys. Rev. 17* (1977) 4384-4403 .
9. J. R. L. de Almeida and D. J. Thouless, *J. Phys. A 11* (1978) 983-990.
10. M. Mezard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond,* (World Scientific, Singapore, 1987).
11. M. Mezard and A. Montanari, *Information, Physics, and Computation,* (Oxford University Press, Oxford, UK, 2009).
12. W. Krauth and M. Mezard, *Europhys. Letts. 8* (1989) 213-218.
13. D. Easley and J. Kleinberg, *Networks, Crowds and Markets: Reasoning about a Highly Connected World,* (Cambridge University Press, Cambridge, UK 2010).
14. D. W. Matula, *Proc. Chapel Hill Conference on Combinatorial Math. and its Applications, Univ. North Carolina* (1970) 356-369.
15. Y. Deshpande and A. Montanari, ArXiv:1304.7047 (2013).
16. S. Kirkpatrick, C. D. Gelatt, Jr. and M. P. Vecchi, *Science 220* (1983) 671-680