

AN INFORMAL INTRODUCTION TO THE PARISI FORMULA

JEAN-CHRISTOPHE MOURRAT

ABSTRACT. This note is an informal presentation of spin glasses and of the Parisi formula. We also discuss some models for which the Parisi formula is not well-understood, and some partial progress that relies upon a connection with partial differential equations.

Statistical mechanics aims to model the emergent properties of systems that are made of a large number of elements. One can think of a gas made of many small particles, which gives rise to macroscopic concepts such as the density, the pressure, or the temperature of the gas. Here we will focus on a class of models of statistical mechanics called *spin glasses*. The critical feature of these models is that there is a lot of “disagreement” between the elementary units of the system, as these models aim to capture features of “complex” systems. Spin glasses have inspired many developments in a variety of topics including statistics, computer science, high-dimensional geometry, or combinatorics, and many facets of these models have been studied. Our main goal here is to discuss some of the ideas surrounding a fundamental result called the Parisi formula, which identifies the limit free energy of some of these models. The free energy is a natural quantity from a physicist’s perspective, and can be thought of as a Laplace transform of the variables of interest, so the identification of its limit yields rich insight into the behavior of the model. We will also discuss an intriguing connection between the Parisi formula and certain partial differential equations, and how this may help to address some open problems.

The outline of this note is as follows. In the first section, we introduce one of the most basic models of a spin glass called the Sherrington-Kirkpatrick model, and explain what features of this model make it indeed a spin glass. Section 2 presents the Parisi formula per se; we also discuss the crucial insight that the support of the associated Gibbs measure is approximately ultrametric. In Section 3, we discuss extensions to more general models, and showcase one example called the bipartite model that is currently less well-understood. In Section 4, we sketch how to rephrase the Parisi formula using certain partial differential equations, as well as some partial results concerning the bipartite model.

This note has a number of footnotes and of pieces of text that are within colored boxes. These are asides or more technical discussions that you can freely skip if you so wish. Reading up until the end of Section 2 should give you a good idea of what the Parisi formula is about, and you may decide to stop there. The remaining sections relate more directly to my recent research activity.

1. THE SHERRINGTON-KIRKPATRICK MODEL

To get a clearer sense of what spin glasses are, let us introduce a paradigmatic example called the Sherrington-Kirkpatrick (SK) model [72]. It may help to imagine ourselves being confronted with the following situation. Suppose that we have N individuals $\{1, \dots, N\}$ that we need to split into two groups. We can represent an assignment into two groups as a vector $\sigma \in \{\pm 1\}^N := \{-1, 1\}^N$, where σ_i indicates to which group the

individual indexed by i is assigned. For each pair (i, j) , we are given a number W_{ij} which represents how much the individual i likes the individual j ; it is positive and very large if the individual i really likes the individual j , while it is very negative if the individual i really dislikes the individual j . We want to look for an assignment $\sigma \in \{\pm 1\}^N$ that makes the following quantity as large as possible:

$$(1.1) \quad H_N(\sigma) := \frac{1}{\sqrt{N}} \sum_{i,j=1}^N W_{ij} \sigma_i \sigma_j.$$

It might have been more natural to write $\mathbf{1}_{\{\sigma_i=\sigma_j\}}$ in place of $\sigma_i \sigma_j$, but since $\mathbf{1}_{\{\sigma_i=\sigma_j\}} = \frac{1}{2}(\sigma_i \sigma_j + 1)$, this is inconsequential¹. The normalization $\frac{1}{\sqrt{N}}$ will be convenient later on.

One quickly realizes that finding the configuration $\sigma \in \{\pm 1\}^N$ that realizes the maximum of H_N is not going to be straightforward. Even when only three individuals i, j and k are involved, we may be in a situation as depicted in Figure 1.1: maybe $W_{ij} + W_{ji} > 0$ and $W_{ik} + W_{ki} > 0$, which would suggest to assign i, j and k to the same group; but if $W_{jk} + W_{kj} < 0$, then the individuals j and k would rather be in different groups, and there is no way to reconcile each of the pairwise preferences.

In other words, certain pairs will typically end up being *frustrated*. In order to find the optimal configuration, some compromises need to be made, and a close inspection of the coefficients (W_{ij}) is required. More generally, for N large, naive quick strategies² for finding a configuration σ such that $H_N(\sigma)$ is large will typically get stuck at a local maximum of the mapping H_N , and will fail to reach the true maximizer. The presence of these frustrations, and the related fact that simple methods typically do not succeed in finding the maximizing configuration, are the defining features of *glassy* systems. How this relates to the glass of our daily lives is discussed in Box 1.1. The variables $\sigma_1, \dots, \sigma_N$ are often called *spins*, because much of statistical mechanics has focused on the modeling of magnetic materials, and the model is therefore called a *spin glass*.




FIGURE 1.1. A simple situation with frustration. Here the coefficients (W_{ij}) suggest to set $\sigma_i = \sigma_j$, $\sigma_i = \sigma_k$, and $\sigma_j = -\sigma_k$, but we cannot realize these three conditions simultaneously.

Box 1.1. Why the word “glass” for these models?

In order to create glass, one starts by heating up silica (i.e. sand) and a bit of calcium and sodium carbonates (i.e. lime and soda) until they melt. One key aspect of the fabrication of glass is that the liquid mixture then needs to be cooled rapidly (“quenched” is the technical term), so as to “trap” the microscopic configuration into a disordered state inherited from the liquid phase. After the rapid quench, the material is forever trying to slowly evolve towards its energetically-preferred organized state, but encounters locally frustrated configurations that are increasingly difficult to overcome. So in this case the frustrations and slow dynamics emerge from the intricate geometry of the configuration of particles. The conflicting coefficients (W_{ij})’s in the SK model make it

¹One reason for defining the model using $\sigma_i \sigma_j$ in place of $\mathbf{1}_{\{\sigma_i=\sigma_j\}}$ has to do with other motivations coming from the modelling of magnetic alloys. From a purely mathematical perspective, it is also more pleasant to think of H_N as a polynomial function of σ , which we can think of as being defined everywhere in \mathbb{R}^N instead of just on $\{\pm 1\}^N$. Another small point is that I find it slightly more convenient not to assume that W_{ij} equals W_{ji} , but this is also a detail.

²For example, as long as there exists an index i such that changing σ_i to $-\sigma_i$ increases H_N , we look for the index that produces the greatest difference, and we iterate.

much more analytically tractable than any model that would really try to be faithful to the “geometric disorder” of a true glass, but one hopes that the two settings share certain qualitative properties.

In line with the necessary presence of these frustrations, one can show that the problem, given the coefficients (W_{ij}) , of finding a configuration $\sigma \in \{\pm 1\}^N$ that maximizes H_N , is NP-hard in general. In fact, the problem is NP-hard even if we only aim to find a configuration $\sigma \in \{\pm 1\}^N$ such that $H_N(\sigma)$ is at least a fixed positive fraction of the maximal value, no matter how small we allow the ratio to be [8]. We will however depart from this worst-case sort of analysis, by focusing instead on what a “typical” instance of the problem looks like. There are surely different possible ways to clarify what “typical” means here, but we choose to postulate that the coefficients $(W_{ij})_{i,j \leq N}$ are independent Gaussian random variables with mean zero and unit variance. The key hypothesis here is that these are independent and all have the same mean and variance; the hypothesis that they are Gaussian is a convenience but would not change the fundamental properties that will be discussed below.

To sum up, we let $(W_{ij})_{i,j \geq 1}$ be independent centered Gaussian random variables with unit variance, we define H_N according to (1.1), and we aim to study quantities such as

$$(1.2) \quad \frac{1}{N} \max_{\sigma \in \{\pm 1\}^N} H_N(\sigma),$$

in the limit of large N . More generally, we would be interested in understanding the geometry of the function H_N , for instance with a view towards finding configurations σ that essentially realize the maximum in (1.2). A particularly fruitful way to probe this is to consider a family of probability measures associated with H_N called *Gibbs measures*; see also Box 1.2. In our context, given a parameter $\beta \geq 0$, the Gibbs measure at “inverse temperature” β is the probability measure that attributes to each $\sigma \in \{\pm 1\}^N$ a probability proportional to $\exp(\beta H_N(\sigma))$.³ For each value of $\beta \geq 0$, this probability measure essentially concentrates on a level set of H_N , in the sense that H_N/N is essentially constant under the Gibbs measure; and as β is taken larger and larger, the measure becomes concentrated on near-maximizers of H_N . In order to understand this family of Gibbs measures, it is very fruitful to elucidate the behavior of the quantity

$$(1.3) \quad F_N(\beta) := \frac{1}{N} \mathbb{E} \log \sum_{\sigma \in \{\pm 1\}^N} \exp(\beta H_N(\sigma)).$$

The expectation \mathbb{E} is with respect to the randomness coming from the coefficients (W_{ij}) . The quantity $F_N(\beta)$ is usually called the *free energy* of the system⁴. Understanding this quantity or generalizations of it is very rich in insight concerning the associated Gibbs measure⁵. For mathematicians, this makes intuitive sense if we think of it as a sort of log-Laplace transform of the function H_N ; we can also see for instance that the derivative in β of F_N gives us access to the average of H_N under the Gibbs measure, averaged

³Physicists prefer to add a minus sign here, writing $\exp(-\beta H_N(\sigma))$ in place of $\exp(\beta H_N(\sigma))$, but since the laws of H_N and $-H_N$ are identical, this does not really matter and I prefer to avoid the proliferation of minus signs. The slight drawback of this convention is that the system now has a preference for *large* values of H_N , while physicists prefer to think that the energy function H_N ought to be minimized.

⁴even though physicists would rather divide F_N by β and add a minus sign (to match the one they would have in the exponential) before calling it that.

⁵The Gibbs measure depends itself on the sampling of the random coefficients (W_{ij}) , and we hope to describe typical properties with respect to this sorting, or average properties.

also over the coefficients (W_{ij}) . We can alternatively think of $F_N(\beta)$ as a soft version of the maximum in (1.2), so that this maximum is approximately $F_N(\beta)/\beta$ if we take β sufficiently large⁶.

Box 1.2. Gibbs measures

We may as well just think of Gibbs measures as useful mathematical tools to probe the level sets of H_N , but when H_N represents the energy function of a real physical system, one can argue from first principles that at equilibrium in an environment at inverse temperature β , the system will indeed be distributed according to this measure. In fact, this is a general phenomenon that can also show up in purely mathematical contexts. For instance, suppose that there are variables x_1, \dots, x_N taking values in a finite set $\{e_1, \dots, e_K\}$, say with $e_1 < \dots < e_K$, and let $\bar{e} \in (e_1, e_K)$. How does one of the variables, say x_1 , look like if we pick the whole vector (x_1, \dots, x_N) uniformly at random among all those that satisfy $N^{-1} \sum_{i=1}^N x_i \simeq \bar{e}$? (Here we write \simeq to allow for some wiggle room so as to make sure that there exist such vectors, for instance we can ask that the difference between the terms on the two sides of \simeq is at most ε_N , with $\varepsilon_N \rightarrow 0$ and $N\varepsilon_N \rightarrow +\infty$ as N tends to infinity.) One can show that as N tends to infinity, the probability that $x_1 = e_k$ is proportional to $\exp(-\beta e_k)$, where β is such that

$$(1.4) \quad \frac{\sum_{k=1}^K e_k \exp(-\beta e_k)}{\sum_{k=1}^K \exp(-\beta e_k)} = \bar{e}.$$

In other words, asymptotically as N tends to infinity, the law of x_1 converges weakly to the Gibbs measure at inverse temperature β , where β is defined by (1.4). One may for instance consult [29, Section 1.1] for more on this and on the relationship between the Gibbs measure, the free energy, and the entropy of the Gibbs measure.

It is not very difficult to show that the quantity in (1.2), and then also that in (1.3), stay away from zero and infinity as N tends to infinity⁷. Can we determine the limits of these quantities? Despite a somewhat lengthy introduction, I hope that you can appreciate the simplicity of the question. And yet, something I find fascinating is that the answer to this question is incredibly rich and complex.

2. THE PARISI FORMULA

An initial guess for the limit of the free energy in (1.3) was proposed by physicists in the paper that introduced the model [72], but it was already understood there that the proposed answer could not be valid for large values of β , i.e. at low temperature. In one of his most celebrated contributions, Giorgio Parisi then came up with a sophisticated non-rigorous procedure, called the replica method, that led to what is now called the *Parisi formula* [65, 66, 67, 68]. This formula is described in full in Box 2.1; here we content ourselves with the fact that it takes the form

$$(2.1) \quad \lim_{N \rightarrow \infty} F_N(\beta) = \inf_{\mu \in \text{Pr}([0,1])} \mathcal{P}(\mu),$$

for some functional \mathcal{P} , where $\text{Pr}([0,1])$ denotes the space of probability measures on $[0,1]$.

⁶The key step to justify this point is to show that this maximum is close to its expectation; see for example [29, Exercises 6.2 and 6.3].

⁷You can try! Here are some hints. For the upper bound, observe that for each $\sigma \in \{\pm 1\}^N$, the random variable $H_N(\sigma)$ is a centered Gaussian with variance N , so we can easily estimate the probability that $H_N(\sigma)$ is above CN for each σ individually. For the lower bound, one can rely on a naive greedy procedure in which, assuming we have already committed to a choice of $\sigma_1, \dots, \sigma_i$, we pick the choice of σ_{i+1} that maximizes that part of H_N that only involves $\sigma_1, \dots, \sigma_{i+1}$, and we iterate over i .

Box 2.1. The Parisi formula in full

The Parisi formula states that

$$(2.2) \quad \lim_{N \rightarrow +\infty} F_N(\beta) = \inf_{\mu \in \text{Pr}([0,1])} \left(\Phi_\mu(0,0) - \beta^2 \int_0^1 t \mu([0,t]) dt + \log(2) \right),$$

where $\Phi_\mu : [0,1] \times \mathbb{R} \rightarrow \mathbb{R}$ is the solution to the backwards parabolic equation

$$(2.3) \quad \begin{cases} -\partial_t \Phi_\mu(t,x) = \beta^2 \left(\partial_x^2 \Phi_\mu(t,x) + \mu([0,t]) (\partial_x \Phi_\mu(t,x))^2 \right) & \text{for } (t,x) \in [0,1] \times \mathbb{R}, \\ \Phi_\mu(1,x) = \log \cosh(x) & \text{for } x \in \mathbb{R}. \end{cases}$$

There are many shades of non-rigorous arguments. Some do not have all the ε 's and δ 's or discard some annoying but plausibly negligible terms, yet most mathematicians would feel rather convinced by them. Giorgio Parisi's techniques were of a different sort though, see Box 2.2. In fact, the opinions on the validity of his prediction initially varied in the physics community. Besides the rather creative nature of the arguments involved, part of the initial skepticism may have come from the fact that the Parisi formula is a minimization problem, as displayed in (2.1). For each fixed N , the free energy can be rewritten as a supremum over probability measures of an “energy” term and an “entropy” term⁸; once one gets familiar with it, this formulation feels very intuitive and appealing. One could then expect that the task of identifying the limit of the free energy boils down to understanding how to simplify this representation in the limit of large N , while preserving the main structure as a supremum. And yet the formula in (2.1) displays an infimum instead⁹.

Box 2.2. The replica method

The replica trick aims to exploit the fact that $\log x = \lim_{n \rightarrow 0} (x^n - 1)/n$ together with the calculation of what in our context would read as

$$(2.4) \quad \mathbb{E} \left[\left(\sum_{\sigma \in \{\pm 1\}^N} \exp(\beta H_N(\sigma)) \right)^n \right],$$

where n is a positive integer. The advantage of working with an integer n is that we can then expand the power and rewrite the expression in (2.4) as a sum over n copies of the variable σ of some expected value we can compute; the n variables ranging in $\{\pm 1\}^N$ are usually called “replicas” (although we will use this word in a slightly different sense further below). This replica trick was already used in the paper that introduced the model [72] and earlier for other models; a short survey of it is in [69, Appendix]. Giorgio Parisi created the art of sending n to zero in just the right way as one goes along the calculation so as to arrive at the correct answer for the limit of the free energy. To give a flavor of the manipulations involved, here is a quote from his Nobel lecture [69]. “After many trials I had an intuition: in other papers the n indices were divided into n/m groups of m elements each [...]. Everybody was assuming that m was an integer [...]. I made the bold assumption that m could be a noninteger number, more precisely a number in the interval $[0,1]$. For example, I was dividing the n replicas into $2n$ groups

⁸While we did not proceed to fully explain this in Box 2, the free energy and the entropy are in a convex duality relationship; see also [29, (1.14)] and [16, Corollaries 4.14 and 4.15] for more on this.

⁹It does not help that in the replica calculation as explained in Box 2.2, a maximization problem appears that is then flipped into a minimization problem as the positive integer n becomes smaller than 1, with not much of an explanation besides the fact that one would otherwise end up with something nonsensical. I cannot help pointing out that a rewriting of the Parisi formula that takes the form of a supremum was recently found in [54].

of 1/2 replicas each. Of course, that is crazy, but my viewpoint was I should first check if this crazy idea was leading to correct results and postpone other questions to a later stage. [...] At the end of the paper, I added the observation that one could improve the theory by dividing the n/m groups into $(n/m)/m_1$ groups of replicas, where m_1 was a new variational parameter. I was also conjecturing the correct solution was obtained when the procedure was repeated an infinite number of times. Some fancy group theory arguments also were added, arguing the permutation group of zero objects is an infinite group because it contains itself as a proper subgroup. [...] The response of the referee was remarkable. In a nutshell: *The approach does not make sense, but the numbers coming from the formulae are reasonable, so it can be published. The last observation is not worth the paper on which it is written and it should be removed.* I laughed because in the meanwhile I extended the computation to an infinite number of subdivisions". That latter calculation led to what we now call the Parisi formula.

Further work by physicists later elucidated a wealth of new information that was consistent with the formula in (2.1) [45, 46, 47]. They discovered that, as N tends to infinity (and except for some choices of (W_{ij}) of small probability), the associated Gibbs measure has a very complex and yet also very precise hierarchical structure, and that the minimizing measure in (2.1) fully describes it. One key aspect of this description is that as N becomes very large, the Gibbs measure is essentially supported on an ultrametric set; see Figure 2.1 and Box 2.3 for more on this.




FIGURE 2.1. Except for some choices of (W_{ij}) of small probability and for large N , samples from the Gibbs measure essentially behave as if they were sampled from an ultrametric space. An ultrametric space can be encoded on the leaves of a tree as in the picture above, where the blue points are at the same distance from one another; they are themselves all at the same distance from any green point; the blue, green and orange points are all at the same distance from any red point, etc.

Box 2.3. Ultrametricity of the Gibbs measure

We wish to discuss the behavior of a probability measure (the Gibbs measure) defined on $\{\pm 1\}^N$, or more generally on \mathbb{R}^N , in the limit of large N . Since the space over which this probability measure is defined changes with N , one cannot simply ask whether the sequence converges weakly for instance. There is also a lot of invariance in the model, such as by permutations of the coordinates, so we may want to discard these symmetries while keeping interesting information. In a sense, we aim to retain information about the Gibbs measure as an abstract random metric space, ignoring how it is embedded into \mathbb{R}^N and discarding small-probability parts. One very convenient way to do this is to study the convergence of the joint law of the relative distances between finitely many independent samples from the measure^a; these independent samples are called *replicas*.

Recall that, for each choice of the parameters (W_{ij}) and each $\beta \geq 0$, the Gibbs measure at inverse temperature β is the probability measure that attributes to the configuration $\sigma \in \{\pm 1\}^N$ a weight proportional to $\exp(\beta H_N(\sigma))$. It is thus a *random* probability measure, as it depends on our choice of the coefficients (W_{ij}) , and we are going to discuss properties of this probability measure that occur with high probability over the choice of these coefficients (W_{ij}) and for large N . What physicists discovered, and was later proved rigorously (up to allowing ourselves a small perturbation of the energy function that does not change the value of the limit free energy), is that the Gibbs measure is asymptotically *ultrametric*. In Euclidean space, a set S is ultrametric if it satisfies the following stronger form of the triangle inequality: for every $x, y, z \in S$, we have

$$|x - y| \leq \max(|x - z|, |y - z|).$$

This can be rephrased as the fact that whenever we draw a triangle with endpoints in S , the two largest sides of the triangle have the same length. In our context, the approximate ultrametricity of the Gibbs measure means that the following holds for every $\varepsilon > 0$ and $\beta \geq 0$ with probability tending to 1 in the W 's as N tends to infinity: if we pick three independent samples (“replicas”) σ , σ' , and σ'' from the Gibbs measure at inverse temperature β , then with Gibbs probability tending to 1 as N tends to infinity, we have

$$|\sigma - \sigma'| \leq \max(|\sigma - \sigma''|, |\sigma' - \sigma''|) + \varepsilon.$$

Any ultrametric set can be represented as the leaves of a tree, with the distance between two leaves being only allowed to depend on the depth of the most recent common ancestor of the two leaves; see Figure 2.1 (or e.g. [21, Lemma 4.2]). The Gibbs measure thus splits into a number of “pure states”^b represented by the leaves of the tree; any two configurations sampled from a pure state are at a fixed distance from one another (with high probability). These pure states organize themselves into clusters, so that the distance between two configurations sampled from two different pure states in the same cluster is essentially constant; these clusters organize themselves into super-clusters; and so on, over a potentially infinite hierarchy of clusters and super-clusters. The exact number of hierarchies depends on the specifics of the problem. There will be no such hierarchy at high temperature (low β), in some sense the tree collapses into a single point in this case. At low temperature (high β), physicists seem to expect that most “natural” problems display either one or an infinite number of levels of hierarchy, although we know how to engineer spin-glass models with an arbitrary number of them.

Beyond the fact that the asymptotic structure is ultrametric, in fact the full metric structure as well as the probability weights that the Gibbs measure attributes to each leaf are completely characterized in terms of the minimizer in (2.1). It would be too long to describe this structure precisely here (see for instance [61] or [29], including [29, Section 6.3] which explains how a 1-level hierarchy emerges from a much simpler toy model). I do want to stress though that I find this rigidity really remarkable. In particular, up to some technical caveats, we have that the hierarchical structure has depth K if and only if the support of the measure μ that minimizes the Parisi formula in (2.1) contains exactly $K + 1$ points. Moreover, if we pick two independent samples σ and σ' from the Gibbs measure, then the law of their “overlap” $\sigma \cdot \sigma' / N$ converges weakly to μ .

Whenever a non-trivial hierarchical structure appears, physicists say that there is *replica-symmetry breaking*; and if there are K levels to the hierarchy, they would say that the system has K levels of replica-symmetry breaking.

To explain the choice of language, it is useful to ask ourselves first what physicists mean when they say that a symmetry is broken. Suppose that I try to balance a pencil that I place vertically on the tip of my finger. By symmetry, the only thing that can happen is that the pencil stays perfectly vertical, right? Well, my experiments do not match this

prediction, as the pencil actually quickly falls in some direction. So one could say that the rotation symmetry is broken. Of course, if I try to be a bit diligent in my attempts at balancing the pencil, then the direction in which the pencil falls will be essentially uniformly random among all possible directions. The symmetry is therefore not really broken in a mathematical sense. The correct mathematical model of the outcome is probabilistic, but we have the impression that the symmetry is broken if we ignore this and only do one experiment.

Coming back to spin glasses and their ultrametric structure: suppose that we pick three independent samples σ , σ' and σ'' according to the Gibbs measure. Of course there is perfect permutation symmetry between these samples, so surely the distance between σ and σ' should be the same as that between σ and σ'' ? Well, this is in fact not the case (see Figure 2.1), so the permutation symmetry between the replicas is broken. The key surprising thing here is that these relative distances remain random even in the limit of large N , even though one may have thought otherwise at first, because many quantities become essentially deterministic (“self-averaging” as physicists would say) in the high-dimensional regime.

^aThis idea defines a topology on the space of (isometry classes of) random measure spaces called the Gromov-weak topology, and this topology is equivalent to that induced by the Gromov-Prokhorov metric [34].

^bThis is vaguely analogous to the decomposition of a probability measure into its ergodic components in the theory of dynamical systems.

This series of discoveries generated a lot of excitement; people had identified a toy model of a “complex” system, with a very rich and rugged energy landscape, that they could study with great precision using analytical methods. Yet, says Giorgio Parisi [69], “it was possible that the correct [solution] was different and more complex [...]. It was difficult to conclude in a definite way.” The validity of the Parisi formula then became a certainty with a series of rigorous mathematical contributions, starting with Francesco Guerra [36] who proved that $\limsup_{N \rightarrow \infty} F_N(\beta) \leq \inf_{\mu \in \text{Pr}([0,1])} \mathcal{P}(\mu)$. In a mathematical tour de force, Michel Talagrand [76] then managed to prove the converse bound and therefore give a complete justification to Parisi’s formula¹⁰. An alternative proof that covers a broader class of models was later developed by Dmitry Panchenko [60, 61]. Besides its greater generality, this alternative proof is also more conceptual, and its key step is very interesting on its own, as it consists in justifying the ultrametricity of the underlying Gibbs measure (up to a small perturbation of the energy function).

3. TOWARDS MORE GENERAL MODELS

Thanks to the insights coming from the Parisi formula and its proof, as well as further developments, we now understand many aspects of the Sherrington-Kirkpatrick model and the structure of its Gibbs measures. To give just one example, there has been much recent progress on whether one can find a polynomial-time algorithm that identifies a configuration $\sigma \in \{\pm 1\}^N$ that essentially maximizes the function H_N (see [30, 31, 32, 33, 39, 40, 50, 71, 75])¹¹.

The insights that were gained for the SK model have allowed us to make progress on a large variety of other “complex” problems that share similar features with “frustrations”, from statistics and high-dimensional geometry to computer science and combinatorics.

¹⁰This is one of the key results celebrated in the citation for Michel Talagrand’s 2024 Abel prize.

¹¹Roughly speaking, the general criterion is that this is possible if and only if the support of the minimizer $\mu \in \text{Pr}([0,1])$ in the Parisi formula (2.1) is connected for all β sufficiently large.

Examples include random constraint satisfaction problems [26, 28, 41, 43, 48, 49], the random assignment and traveling salesman problems [3, 4, 44], community detection and more general problems of large-scale statistical learning [1, 81] (see also [29, Section 4]), error correcting codes in information theory [70], and combinatorial problems such as graph coloring [23, 27, 56].

In this presentation, I will single out a seemingly modest generalization of the SK model in which the variables are of two different types. We can visualize this by thinking of the variables as being organized into two layers as depicted in Figure 3.1, with each layer being made of variables of a single type; we now postulate that there are direct interactions between variables of different types only. I find this model interesting because it is closely related to several classical models of artificial neural networks. One is the Hopfield model, which is a model of memory storage and retrieval with a long history [5, 6, 7, 9, 38, 42, 78]. Another one is the so-called restricted Boltzmann machine, which is an artificial neural-network architecture that was popular until around ten years ago for learning data distributions and then generating new samples [37, 73]; in this case our model is only capturing the initial stage of the network before any learning occurs, but physicists have already managed to gain very interesting insight from there [13, 79, 80]¹².




FIGURE 3.1. Elementary units are organized into two layers and only interact across layers.

To formalize the model precisely, we can represent the variables as a pair $\sigma = (\sigma_1, \sigma_2) = (\sigma_{1,1}, \dots, \sigma_{1,N}, \sigma_{2,1}, \dots, \sigma_{2,N}) \in \{\pm 1\}^N \times \{\pm 1\}^N$, and we set

$$(3.1) \quad H_N^{\text{bip}}(\sigma) := \frac{1}{\sqrt{N}} \sum_{i,j=1}^N W_{i,j} \sigma_{1,i} \sigma_{2,j}.$$

For simplicity we have imposed σ_1 and σ_2 to be vectors of the same length N , but this is not fundamental; the important point is to make sure that the respective sizes of the two layers remain proportional to one another as we consider larger and larger systems. We have also retained the idea that each of the variables takes values in $\{\pm 1\}$, but this can also be relaxed¹³. We will call this model the *bipartite model*.

This model may seem barely distinguishable from the SK model at first sight. Yet, to this day, we do not know what the limit of the free energy is in this case; in fact, even the fact that the free energy converges as N tends to infinity is not known. To be clear, the free energy here is

$$\frac{1}{N} \mathbb{E} \log \sum_{\sigma \in \{\pm 1\}^N \times \{\pm 1\}^N} \exp(\beta H_N^{\text{bip}}(\sigma)).$$

The asymptotic behavior of the maximum

$$\frac{1}{N} \max_{\sigma \in \{\pm 1\}^N \times \{\pm 1\}^N} H_N^{\text{bip}}(\sigma)$$

is also poorly understood. The problem here goes beyond that of fixing some technical part in the proof of the Parisi formula. Indeed, one may imagine several possible ways

¹²In the time since this paragraph was written, the Nobel committee awarded the 2024 Physics prize to John Hopfield and Geoffrey Hinton, citing in particular their works on these models.

¹³By sticking with the simplest possible version of the model, we have artificially introduced a symmetry between the two layers, but we want to insist on devising analysis techniques that do not rely on this symmetry; a simple way to break the symmetry while staying within our context and notation is to add a term $h \sum_{i=1}^N \sigma_{1,i}$ into the definition of $H_N^{\text{bip}}(\sigma)$, for some parameter $h \neq 0$.

for the Parisi formula to generalize to this bipartite model, however it can be shown that none of those candidates for the limit are valid [51, Section 6].

In order to clarify what distinguishes the bipartite model from the SK model at the technical level, it is best to change a bit our viewpoint on the definition of these random fields H_N and H_N^{bip} . Instead of writing them down explicitly as in (1.1) and (3.1), an equivalent way to define them is to specify that they are centered Gaussian fields, and to display their covariance. For the SK model, we have for every $\sigma, \tau \in \{\pm 1\}^N$ that

$$(3.2) \quad \mathbb{E}[H_N(\sigma)H_N(\tau)] = N \left(\frac{\sigma \cdot \tau}{N} \right)^2,$$

where the dot in $\sigma \cdot \tau$ denotes the usual scalar product in \mathbb{R}^N . So instead of writing the formula in (1.1), we could also have said “let $(H_N(\sigma))_{\sigma \in \{\pm 1\}^N}$ be the centered Gaussian vector whose covariance is given by (3.2)”. The one advantage of the explicit formula, besides its possibly more intuitive appeal, is that it makes it transparent that such a random vector exists. More generally, one could consider centered Gaussian fields $(H_N(\sigma))_{\sigma \in \mathbb{R}^N}$ such that, for some smooth function $\xi : \mathbb{R} \rightarrow \mathbb{R}$, we have for every $\sigma, \tau \in \mathbb{R}^N$ that

$$(3.3) \quad \mathbb{E}[H_N(\sigma)H_N(\tau)] = N \xi \left(\frac{\sigma \cdot \tau}{N} \right);$$

the SK model corresponds to the case when $\xi(r) = r^2$. For the choice of $\xi(r) = r^3$, we can find a Gaussian field that satisfies (3.3) by setting

$$H_N(\sigma) := \frac{1}{N} \sum_{i,j,k=1}^N J_{i,j,k} \sigma_i \sigma_j \sigma_k,$$

where $(J_{i,j,k})$ are independent centered Gaussians with unit variance. One can realize $\xi(r) = r^p$ for any positive integer p by proceeding similarly. Multiplying a random field by a factor of λ transforms the function ξ into $\lambda^2 \xi$. And by adding independent versions of fields with possibly different functions ξ , we can also create a new field whose covariance is the sum of those functions ξ . So we see that any function ξ that can be written in the form

$$(3.4) \quad \xi(r) = \sum_{p=0}^{+\infty} a_p r^p,$$

with $a_p \geq 0$ going to zero sufficiently rapidly as p tends to infinity, will be a valid covariance function. One can show that these are all the admissible functions (see [53, Proposition 6.6] for a statement that also covers the case of models with multiple types of spins).

For the bipartite model, we have instead that, for every $\sigma, \tau \in \{\pm 1\}^N \times \{\pm 1\}^N$,

$$(3.5) \quad \mathbb{E}[H_N^{\text{bip}}(\sigma)H_N^{\text{bip}}(\tau)] = N \left(\frac{\sigma_1 \cdot \tau_1}{N} \right) \left(\frac{\sigma_2 \cdot \tau_2}{N} \right).$$

The key technical difference between the SK and the bipartite models is that here the relevant function that shows up on the right side of (3.5) is the mapping $(x, y) \mapsto xy$, which is *not convex*. To be precise, for models with only one type of spins, i.e. of the form in (3.3), what is crucial is that the function ξ is convex over \mathbb{R}_+ ; as one can see from (3.4), this is in fact always true! This convexity property can however break down as soon as we consider models with two or more types of spins. In general, we may consider models with a fixed number D of types of spins, say $\sigma = (\sigma_1, \dots, \sigma_D) \in (\mathbb{R}^N)^D$, with a covariance

such that, for every $\sigma, \tau \in (\mathbb{R}^N)^D$,

$$\mathbb{E}[H_N(\sigma)H_N(\tau)] = N\xi\left(\left(\frac{\sigma_d \cdot \sigma_{d'}}{N}\right)_{1 \leq d, d' \leq D}\right),$$

where ξ is some (admissible) function from $\mathbb{R}^{D \times D}$ to \mathbb{R} . Those models for which we can write down and rigorously prove a Parisi formula for the limit free energy are exactly those for which the function ξ is convex over the space of symmetric positive semidefinite matrices (see [21, Theorem 1], which crucially builds upon [10, 60, 62, 63, 64]).

4. A CONNECTION WITH PARTIAL DIFFERENTIAL EQUATIONS

So what can we do for the bipartite model? As discussed in the previous section, in this case it is not even clear what one is supposed to show, as there is no clear way to extend the Parisi formula into some reasonable candidate for the limit free energy. I will explain one possible route which my collaborators and I have been exploring. The idea is in two steps. First, we enrich the free energy, adding some (hopefully not too complicated) terms to the energy function, so that we end up with a free energy that depends on additional parameters besides the inverse temperature. Next, one hopes to find a partial differential equation that this free energy solves approximately, and to characterize the limit free energy as the unique solution to this equation. This approach works well for simpler models such as the Curie-Weiss model and its generalizations, or for some problems of statistical inference such as community detection, see [29, Chapters 1 to 4] for a detailed presentation. Perhaps good pedagogic practice would require that I discuss these models first. Instead I will try to directly address the more complicated case of spin glasses, but some aspects of the discussion will then be rather sketchy.

The fact that the limit free energy of a model of statistical mechanics may solve a partial differential equation is an old observation going back at least to [19, 57] (see also [14] for a recent survey on related topics). In the context of spin glasses, these connections were first explored in [2, 11, 12, 35] under simplifying assumptions. The fact that the Parisi formula can be recast as the value of the solution to some partial differential equation is from [22, 52, 55] (see also [29]).

In order to keep the notation simple, we will first discuss the approach in the case of the SK model. Let us start by defining a new free energy with an additional parameter, and then discuss motivations for this choice. For every $t \geq 0$ and $h \geq 0$, we set

$$(4.1) \quad F_N(t, h) := -\frac{1}{N} \mathbb{E} \log \sum_{\sigma \in \{\pm 1\}^N} \exp(\sqrt{2t}H_N(\sigma) - Nt + \sqrt{2h}z \cdot \sigma - Nh),$$

where $z = (z_1, \dots, z_N)$ is a vector of independent centered Gaussians with unit variance, independent of H_N , and we recall that the function H_N for the SK model is defined in (1.1).

Several comments need to be made to explain some of the choices that have been made in this new definition of F_N . First, we replaced β with $\sqrt{2t}$. The ultimate reason for this is that the subsequent formulas will look nicer in this way, but one way to sense that this may be so is to recall that $H_N(\sigma)$ is a Gaussian random variable, and so the scaling $\sqrt{2t}H_N(\sigma)$ is such that the variance of the Gaussian scales linearly, as with Brownian motion; it is as if we were continuously adding new independent copies of $H_N(\sigma)$ homogeneously in time. The same goes for $\sqrt{2h}$ in front of $z \cdot \sigma$. The factor $-Nt$ is also ultimately here just to make the subsequent formulas a bit nicer; it is half the variance of the Gaussian random variable $\sqrt{2t}H_N(\sigma)$. People who are familiar with stochastic calculus will recognize a stochastic exponential here: in particular, we have

that $\mathbb{E}[\exp(\sqrt{2t}H_N(\sigma) - Nt)] = 1$. So if we were to exchange the expectation and the logarithm in (4.1), we would end up with summands that are all equal to 1. By Jensen's inequality, the quantity with expectation and logarithm interchanged is smaller than the original one, and we set things up so as to monitor the defect in this Jensen's inequality¹⁴.

From the formula (4.1) we have so far mostly discussed simple changes of variables, or tried to give some excuse for the presence of the silly-looking terms Nt or Nh , but the important part is to get a feeling as to why we are adding this term $\sqrt{2h}z \cdot \sigma$. That we are looking for a partial differential equation for F_N means that we hope to be able to compensate small variations of t with small variations in h . So we aim to find some term that looks like $H_N(\sigma)$ in some sense, while being also simpler to analyse. Perhaps one way to think that the term $\sqrt{2h}z \cdot \sigma$ is not an unreasonable choice is to write

$$H_N(\sigma) = \frac{1}{\sqrt{N}} \sum_{i=1}^N \left(\sum_{j=1}^N J_{ij} \sigma_j \right) \sigma_i,$$

and to venture the guess that maybe the random variables $(\sum_{j=1}^N J_{ij} \sigma_j)$ could be substituted with equivalent independent Gaussians, because, well, at least for each fixed σ they are Gaussian after all. A more detailed explanation is beyond the scope of this note; for me the best heuristic is that discussed in [29, Exercise 6.5 and solution].

Let us be pragmatic here and just calculate the derivatives of F_N to see if something interesting happens. In order to express these derivatives nicely, we introduce some notation for the Gibbs measure. For any function f , we write

$$(4.2) \quad \langle f(\sigma) \rangle := \frac{\sum_{\sigma \in \{\pm 1\}^N} f(\sigma) \exp(H_N(t, h, \sigma))}{\sum_{\sigma \in \{\pm 1\}^N} \exp(H_N(t, h, \sigma))},$$

where we have set $H_N(t, h, \sigma) := \sqrt{2t}H_N(\sigma) - Nt + \sqrt{2h}z \cdot \sigma - Nh$. In the notation on the left side of (4.2), the bracket $\langle \cdot \rangle$ thus stands for the expectation with respect to the Gibbs measure, and we think of σ as a random variable that is sampled accordingly. We also write σ' to denote an independent copy of σ under the Gibbs measure, that is, we write

$$\langle f(\sigma, \sigma') \rangle := \frac{\sum_{\sigma, \sigma' \in \{\pm 1\}^N} f(\sigma, \sigma') \exp(H_N(t, h, \sigma) + H_N(t, h, \sigma'))}{\sum_{\sigma, \sigma' \in \{\pm 1\}^N} \exp(H_N(t, h, \sigma) + H_N(t, h, \sigma'))}.$$

A calculation¹⁵ gives us that

$$(4.3) \quad \partial_t F_N(t, h) = \mathbb{E} \left\langle \left(\frac{\sigma \cdot \sigma'}{N} \right)^2 \right\rangle \quad \text{and} \quad \partial_h F_N(t, h) = \mathbb{E} \left\langle \frac{\sigma \cdot \sigma'}{N} \right\rangle.$$

¹⁴In general, we also do not restrict ourselves to models defined on $\{\pm 1\}^N$. For a model whose covariance is given by (3.3), we thus write

$$F_N(t, h) := -\frac{1}{N} \mathbb{E} \log \int \exp \left(\sqrt{2t}H_N(\sigma) - Nt\xi(|\sigma|^2/N) + \sqrt{2h}z \cdot \sigma - h|\sigma|^2 \right) dP_N(\sigma),$$

where $P_N = P_1^{\otimes N}$ is the N -fold tensor product of a probability measure P_1 on \mathbb{R} with compact support. We recover the SK model by choosing $\xi(r) = r^2$ and $P_1 = (\delta_1 + \delta_{-1})/2$, up to the addition of a trivial factor of $\log 2$. Notice that with this definition, Jensen's inequality tells us that $F_N \geq 0$, thanks to the extra minus sign that we added there.

¹⁵You can try! The only thing to have in mind is that if G is a centered Gaussian of unit variance and f is a nice enough function (say C^1 with reasonable growth of f and f' at infinity), then

$$\mathbb{E}[Gf(G)] = \mathbb{E}[f'(G)].$$

To see this, you just need to write the expectation explicitly as an integral against the Gaussian density, and integrate by parts.

The dependence on t and h in the right-hand sides of the identities above is hidden in the definition of the Gibbs average $\langle \cdot \rangle$. For more general models as in (3.3), we would find the exact same expression for $\partial_h F_N$ as in (4.3) (for the corresponding definition of the Gibbs measure), while for the derivative in t , we would find that

$$(4.4) \quad \partial_t F_N(t, h) = \mathbb{E} \left\langle \xi \left(\frac{\sigma \cdot \sigma'}{N} \right) \right\rangle.$$

Continuing with the SK model for now, we thus obtain that

$$(4.5) \quad \partial_t F_N - (\partial_h F_N)^2 = \mathbb{E} \left\langle \left(\frac{\sigma \cdot \sigma'}{N} \right)^2 \right\rangle - \left(\mathbb{E} \left\langle \frac{\sigma \cdot \sigma'}{N} \right\rangle \right)^2.$$

The right-hand side of this identity is the variance of the random variable $\sigma \cdot \sigma' / N$ under $\mathbb{E} \langle \cdot \rangle$. Now, if you have read Section 2 carefully, you know that this variance will not be small in general. Yet this would have been a reasonable first guess, since $\sigma \cdot \sigma'$ is a sum of a large number of terms, so we could at first anticipate that this will have small fluctuations — and in fact, this intuition turns out to be valid as long as t is small. Just to see what happens (or by deciding that we restrict to small t), let us get along a little with this hypothesis that the variance of $\sigma \cdot \sigma' / N$ tends to zero as N tends to infinity. We are then led to the expectation that F_N may converge to a limit function f that solves the equation

$$(4.6) \quad \partial_t f - (\partial_h f)^2 = 0.$$

Moreover, as we set $t = 0$ in the definition of F_N in (4.1), we find that we can easily calculate the result by writing $z \cdot \sigma = \sum_{i=1}^N z_i \sigma_i$, writing the exponential of the sum as a sum of exponentials, and realizing that we can then factorize the summation into a product of N simple sums over $\{\pm 1\}$. In short, we see that

$$(4.7) \quad F_N(0, h) = F_1(0, h).$$

So if we believe that the random variable $\sigma \cdot \sigma' / N$ has vanishingly small fluctuations in the limit of large N , then we are tempted towards the conjecture that F_N converges to the function f that solves (4.6) with initial condition $f(0, \cdot) = F_1(0, h)$. In fact, this guess for the limit free energy that we thus obtain is exactly the one that was proposed (using other arguments) in the paper that introduced the SK model [72]. For the model with the covariance as in (3.3), we would instead have guessed the partial differential equation

$$(4.8) \quad \partial_t f - \xi(\partial_h f) = 0.$$

Partial differential equations of this sort are called Hamilton-Jacobi equations.

The hypothesis that the random variable $\sigma \cdot \sigma' / N$ has vanishingly small fluctuations in the limit of large N is just of the sort that works very well for simpler models, but as discussed earlier, here it is only valid at high temperature, that is for small values of t .

Even though things do not really work out, the idea we have explored here seems promising, so we are going to build upon it. The point is that we have not fully succeeded in “closing the equation” with the introduction of this parameter h ; but maybe if we were introducing more parameters, then we would be able to do better.

The full construction of the free energy that incorporates these additional parameters would take too much space to explain here (one can consult [29, Section 6.4]), but the basic idea is that, since we anticipate that the true system may have a complicated ultrametric structure as discussed in the previous section, we want to replace our naive “external field” $\sqrt{2h}z$, which we were “dotting” against σ , with a more refined object that

has such an ultrametric structure embedded into it. This ultrametric structure is encoded using a non-decreasing function $q : [0, 1] \rightarrow [0, \infty)$.

So this construction leads us to a new definition of the free energy F_N , which we now think of as a function of t and q in place of t and h . We may call this extended function $F_N(t, q)$ the “enriched” free energy. The version in (1.1) corresponds to the choice of the constant function $q = h$ in this new definition; in particular, we are computing the free energy of the “vanilla” model when we compute $F_N(\beta^2/2, 0)$, up to remembering the extra minus sign and the term Nt in (4.1).

Let us denote by \mathcal{Q} the space of non-decreasing functions $q : [0, 1] \rightarrow [0, \infty)$. The notion of differentiation for functions $g : \mathcal{Q} \rightarrow \mathbb{R}$ that we use is as follows. We say that g is differentiable at $q \in \mathcal{Q}$, and in this case denote by $\partial_q g(q, \cdot) \in L^2([0, 1]; \mathbb{R})$ the derivative at q , if for every $q' \in \mathcal{Q}$, we have

$$g((1 - \varepsilon)q + \varepsilon q') = g(q) + \varepsilon \int_0^1 (q' - q)(u) \partial_q g(q, u) du + o(\varepsilon) \quad (\varepsilon \rightarrow 0).$$

Equipped with this, one can find a nice expression for the derivative of F_N with respect to q , which resembles what we have found in the second part of (4.3) but also involves some other variables that play a role in the ultrametric structure of the new external field. One is then led to form the quantity

$$\partial_t F_N(t, q) - \int_0^1 (\partial_q F_N(t, q, u))^2 du = \text{small?}$$

What replaces the “small” right-hand side above is a conditional variance of $\sigma \cdot \sigma' / N$, in particular it is indeed smaller than the variance of $\sigma \cdot \sigma' / N$, so we are indeed making progress. For the model with the covariance as in (3.3), we are led to hope that

$$\partial_t F_N(t, q) - \int_0^1 \xi(\partial_q F_N(t, q, u)) du = \text{small?}$$

While we did not explain the construction of this new external field with an ultrametric structure, it turns out that it is still “factorizable” in the sense that the nice property that we had in (4.7) is still valid in this more general setting, that is, we have for every $q \in \mathcal{Q}$ that

$$(4.9) \quad F_N(0, q) = F_1(0, q).$$

It turns out that our more sophisticated hopes are now correct. For notational convenience, we define $\psi(q) := F_1(0, q)$.

Theorem 4.1 (The Parisi formula as a Hamilton-Jacobi equation [22, 52, 55]). *The enriched free energy $F_N : \mathbb{R}_+ \times \mathcal{Q} \rightarrow \mathbb{R}$ for the SK model converges pointwise to the function $f : \mathbb{R}_+ \times \mathcal{Q} \rightarrow \mathbb{R}$ that solves*

$$(4.10) \quad \begin{cases} \partial_t f - \int_0^1 (\partial_q f)^2 = 0 & \text{on } \mathbb{R}_+ \times \mathcal{Q}, \\ f(0, \cdot) = \psi & \text{on } \mathcal{Q}. \end{cases}$$

More generally, if $F_N : \mathbb{R}_+ \times \mathcal{Q} \rightarrow \mathbb{R}$ now stands for the enriched free energy associated with a model whose covariance is given by (3.3), then F_N converges to the function f that solves

$$(4.11) \quad \begin{cases} \partial_t f - \int_0^1 \xi(\partial_q f) = 0 & \text{on } \mathbb{R}_+ \times \mathcal{Q}, \\ f(0, \cdot) = \psi & \text{on } \mathcal{Q}. \end{cases}$$

Part of the task of making sense of this theorem is that one needs to find a good notion of solution for the Hamilton-Jacobi equations in (4.10) and (4.11). This is based on the

notion of viscosity solutions (see [29, Chapter 3] for an introduction that is tailored to our context).

How does this formulation relate to the Parisi formula? The key point is that when the non-linearity (the square function in (4.10)) is convex, the solution admits a variational representation known as the Hopf-Lax formula. The solution to (4.10) can be written as

$$(4.12) \quad f(t, q) = \sup_{q' \in \mathcal{Q}} \left(\psi(q + q') - \frac{1}{4t} \int_0^1 (q')^2 \right).$$

For the solution to the more general equation in (4.11), we have

$$(4.13) \quad f(t, q) = \sup_{q' \in \mathcal{Q}} \left(\psi(q + q') - t \int_0^1 \xi^* \left(\frac{q'}{t} \right) \right),$$

where ξ^* is the convex dual of ξ defined, for every $s \in \mathbb{R}$, by

$$\xi^*(s) := \sup_{r \geq 0} (rs - \xi(r)).$$

One can then bridge this representation back to the Parisi formula by setting $q = 0$ in (4.12), making some change of variables, and doing some explicit calculations involving the function ψ .¹⁶

All this is nice, but we already knew the answer in these cases. The really interesting things come about when we consider models with different types of spins such as the bipartite model. Recall that in this case the energy function H_N^{bip} is defined in (3.1). Let us first see how a simpler guess analogous to (4.6) or (4.8) would look like then. The important point is that since there are now two types of spins, we would like to make sure that we have one additional variable (one additional “external field”) to act upon each of these two types. So we set, for every $t \geq 0$, $h = (h_1, h_2) \in \mathbb{R}_+^2$, and $\sigma = (\sigma_1, \sigma_2) \in \{\pm 1\}^N \times \{\pm 1\}^N$,

$$H_N(t, h, \sigma) := \sqrt{2t} H_N^{\text{bip}}(\sigma) - Nt + \sqrt{2h_1} z_1 \cdot \sigma_1 - Nh_1 + \sqrt{2h_2} z_2 \cdot \sigma_2 - Nh_2,$$

(we drop the superscript ^{bip} from now on for ease of notation), and then for every $t \geq 0$ and $h = (h_1, h_2) \in \mathbb{R}_+^2$, we set the new free energy to be

$$F_N(t, h) := -\frac{1}{N} \mathbb{E} \log \sum_{\sigma \in \{\pm 1\}^N \times \{\pm 1\}^N} \exp(H_N(t, h, \sigma)).$$

The definition of the Gibbs average $\langle \cdot \rangle$ is hopefully clear, it is essentially as in the formula in (4.2), except that now the variable h is in \mathbb{R}_+^2 and the summation variable σ ranges in $\{\pm 1\}^N \times \{\pm 1\}^N$. In place of (4.3) or (4.4), we obtain that

$$\partial_t F_N(t, h) = \mathbb{E} \left\langle \left(\frac{\sigma_1 \cdot \sigma'_1}{N} \right) \left(\frac{\sigma_2 \cdot \sigma'_2}{N} \right) \right\rangle,$$

with still

$$\partial_{h_1} F_N(t, h) = \mathbb{E} \left\langle \frac{\sigma_1 \cdot \sigma'_1}{N} \right\rangle \quad \text{and} \quad \partial_{h_2} F_N(t, h) = \mathbb{E} \left\langle \frac{\sigma_2 \cdot \sigma'_2}{N} \right\rangle.$$

So now if we expect that the random variables $\sigma_1 \cdot \sigma'_1 / N$ and $\sigma_2 \cdot \sigma'_2 / N$ do not fluctuate much in the limit of large N , we are led to the belief that F_N converges to f solving

$$\partial_t f - \partial_{h_1} f \partial_{h_2} f = 0.$$

The initial condition $f(0, \cdot)$ is easy to compute as it satisfies (4.7) again.

¹⁶Our formulas here are with a supremum, unlike in (2.1), but this is only for the trivial reason that we have added a minus sign in the definition in (4.1).

The problem is that this guess is too naive again of course, so instead we need to replace our pair of simple variables $h = (h_1, h_2) \in \mathbb{R}_+^2$ by a pair of functions $q = (q_1, q_2) \in \mathcal{Q}^2$ that encodes the ultrametric structure of the external fields that act upon σ_1 and σ_2 respectively. This defines for us a more sophisticated free energy F_N , now defined on $\mathbb{R}_+ \times \mathcal{Q}^2$ in place of $\mathbb{R}_+ \times \mathbb{R}_+^2$. Again we have that (4.9) holds (now with q ranging in \mathcal{Q}^2); let us set $\psi := F_1(0, \cdot)$ for convenience. By going through similar calculations as for the SK model, one thus ends up expecting that:

Conjecture 4.2. *Maybe the enriched free energy $F_N : \mathbb{R}_+ \times \mathcal{Q}^2 \rightarrow \mathbb{R}$ for the bipartite model converges to the function $f : \mathbb{R}_+ \times \mathcal{Q}^2 \rightarrow \mathbb{R}$ that solves*

$$(4.14) \quad \begin{cases} \partial_t f - \int_0^1 \partial_{q_1} f \partial_{q_2} f = 0 & \text{on } \mathbb{R}_+ \times \mathcal{Q}^2, \\ f(0, \cdot) = \psi & \text{on } \mathcal{Q}^2 \quad ? \end{cases}$$

In this case, since the non-linearity in the equation (4.14) (i.e. the mapping $(x, y) \mapsto xy$) is neither convex nor concave, there is no standard way to represent the solution to (4.14) variationally. So there is no clear way to rewrite this candidate for the limit free energy into a form that would resemble the Parisi formula for the SK model.

As was said already, we do not know if Conjecture 4.2 is valid or not. One thing we know is that any subsequential limit of F_N must satisfy the equation in (4.14) “almost everywhere” in $\mathbb{R}_+ \times \mathcal{Q}^2$, for a suitable interpretation of “almost everywhere” [21]. This is not sufficient to fully characterize the solution to (4.14) though; even in finite dimensions, there may be many functions that solve a Hamilton-Jacobi equation almost everywhere. Another thing that we know is that the solution to (4.14) is a lower bound [22, 51], that is, for every $t \geq 0$ and $q \in \mathcal{Q}^2$, we have

$$\liminf_{N \rightarrow \infty} F_N(t, q) \geq f(t, q),$$

where f denotes the unique viscosity solution to (4.14). The result that I find most interesting, and which essentially matches what the physicists say using other language, is one that I will only describe informally. In order to do so, let us first discuss a classical method to solve (4.14) for short times called the method of characteristics (see [29, Section 3.5] for more precision). It turns out that if one does some formal calculations assuming that the solution to (4.14) is twice differentiable, one can guess a simple formula for what the value of the solution ought to be along each of a spanning family of lines called the characteristics. In our case the characteristic that starts at q is the straight line

$$t \mapsto q - t \nabla \xi(\partial_q \psi(q)),$$

where we denoted by $\nabla \xi : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+^2$ the function mapping (x, y) to (y, x) (in line with the idea that the relevant “function ξ ” for the bipartite model is $(x, y) \mapsto xy$), and $\partial_q \psi = (\partial_{q_1} \psi, \partial_{q_2} \psi)$. This guess provided by characteristics is actually really valid for short times, in the sense that it really gives us the value of the viscosity solution there. The problem is that for large times, these lines may start to intersect each other, and they do not agree on what they predict (and the viscosity solution may then take a value that is not even among those multiple options). One of the main results of [21] is that if one assumes that the enriched free energy $F_N : \mathbb{R}_+ \times \mathcal{Q}^2 \rightarrow \mathbb{R}$ of the bipartite model converges to some limit function f , then there is always *some* characteristic line which prescribes the correct value for f . The only remaining source of ambiguity is that we are not able to say which of the characteristic lines is the correct one, in case multiple lines arrive at a given point. See also Figure 4.1 for an illustration.




FIGURE 4.1. Each characteristic line offers us a “prediction” for what the limit free energy is. For small t , each point (t, q) (e.g. the orange point) is reached by exactly one characteristic line, so we know the value of the free energy then. For large t , it could happen that multiple characteristics reach a point, as happens for the black point. We know that the limit free energy must be as prescribed by one of those characteristics, but we do not know which one.

5. CLOSING THOUGHTS

Spin glasses are models of statistical mechanics that are relatively simple to define. Yet, one aspect I find striking about them is that they display a very rich and interesting mathematical structure. Despite their apparent simplicity, they may give us some insight into the behavior of various “complex” systems across disciplines.

I hope that this informal overview of the Parisi formula and of some related open questions has piqued your interest. I want to stress though that this is only one facet of the study of spin glasses, and there surely is something for everyone there! Books on spin glasses include [15, 17, 18, 20, 24, 25, 29, 47, 58, 59, 61, 74, 77, 78]. The book [29] puts particular emphasis on the PDE point of view discussed here in Section 4 (and contains an overview of the research literature). For a multi-disciplinary perspective and applications of spin glasses to other areas, one can consult the recent book [20] on “spin glass theory and far beyond”.

REFERENCES

- [1] Emmanuel Abbe. Community detection and stochastic block models: recent developments. *J. Mach. Learn. Res.*, 18(1):6446–6531, 2017.
- [2] Elena Agliari, Adriano Barra, Raffaella Burioni, and Aldo Di Biasio. Notes on the p-spin glass studied via Hamilton-Jacobi and smooth-cavity techniques. *J. Math. Phys.*, 53(6):063304, 29, 2012.
- [3] David Aldous. Asymptotics in the random assignment problem. *Probab. Theory Related Fields*, 93(4):507–534, 1992.
- [4] David J. Aldous. The $\zeta(2)$ limit in the random assignment problem. *Random Structures Algorithms*, 18(4):381–418, 2001.
- [5] Shun’ichi Amari. Learning patterns and pattern sequences by self-organizing nets of threshold elements. *IEEE Transactions on computers*, 100(11):1197–1206, 1972.
- [6] Daniel J. Amit, Hanoach Gutfreund, and Haim Sompolinsky. Spin-glass models of neural networks. *Phys. Rev. A*, 32(2):1007, 1985.
- [7] Daniel J. Amit, Hanoach Gutfreund, and Haim Sompolinsky. Statistical mechanics of neural networks near saturation. *Annals of Physics*, 173(1):30–67, 1987.

- [8] Sanjeev Arora, Eli Berger, Hazan Elad, Guy Kindler, and Muli Safra. On non-approximability for quadratic programs. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 206–215. IEEE, 2005.
- [9] Adriano Barra, Alberto Bernacchia, Enrica Santucci, and Pierluigi Contucci. On the equivalence of Hopfield networks and Boltzmann machines. *Neural Networks*, 34:1–9, 2012.
- [10] Adriano Barra, Pierluigi Contucci, Emanuele Mingione, and Daniele Tantari. Multi-species mean field spin glasses. Rigorous results. *Ann. Henri Poincaré*, 16(3):691–708, 2015.
- [11] Adriano Barra, Gino Del Ferraro, and Daniele Tantari. Mean field spin glasses treated with PDE techniques. *Eur. Phys. J. B*, 86(7):Art. 332, 10, 2013.
- [12] Adriano Barra, Aldo Di Biasio, and Francesco Guerra. Replica symmetry breaking in mean-field spin glasses through the Hamilton-Jacobi technique. *J. Stat. Mech. Theory Exp.*, (9):P09006, 22, 2010.
- [13] Adriano Barra, Giuseppe Genovese, Peter Sollich, and Daniele Tantari. Phase diagram of restricted Boltzmann machines and generalized Hopfield networks with arbitrary priors. *Phys. Rev. E*, 97(2):022310, 2018.
- [14] Roland Bauerschmidt, Thierry Bodineau, and Benoit Dagallier. Stochastic dynamics and the Polchinski equation: an introduction. *Preprint, arXiv:2307.07619*, 2023.
- [15] Erwin Bolthausen and Anton Bovier, editors. *Spin glasses*, volume 1900 of *Lecture Notes in Mathematics*. Springer, Berlin, 2007.
- [16] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities*. Oxford University Press, Oxford, 2013.
- [17] Anton Bovier. *Statistical mechanics of disordered systems: a mathematical perspective*, volume 18 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 2006.
- [18] Anton Bovier and Pierre Picco, editors. *Mathematical aspects of spin glasses and neural networks*, volume 41 of *Progress in Probability*. Birkhäuser Boston, Inc., Boston, MA, 1998.
- [19] J. G. Brankov and V. A. Zagrebnov. On the description of the phase transition in the Husimi-Temperley model. *J. Phys. A*, 16(10):2217–2224, 1983.
- [20] Patrick Charbonneau, Enzo Marinari, Marc Mézard, Giorgio Parisi, Federico Ricci-Tersenghi, Gabriele Sicuro, and Francesco Zamponi. *Spin glass theory and far beyond*. World Scientific, 2023.
- [21] Hong-Bin Chen and Jean-Christophe Mourrat. On the free energy of vector spin glasses with non-convex interactions. *Probab. Math. Phys.*, to appear.
- [22] Hong-Bin Chen and Jiaming Xia. Hamilton-Jacobi equations from mean-field spin glasses. *arXiv preprint arXiv:2201.12732*, 2022.
- [23] Amin Coja-Oghlan. Phase transitions in discrete structures. In *European Congress of Mathematics*, pages 599–618. Eur. Math. Soc., Zürich, 2018.
- [24] Pierluigi Contucci and Cristian Giardinà. *Perspectives on spin glasses*. Cambridge University Press, Cambridge, 2013.
- [25] Cirano De Dominicis and Irene Giardinà. *Random fields and spin glasses: a field theory approach*. Cambridge University Press, New York, 2006.
- [26] Jian Ding, Allan Sly, and Nike Sun. Proof of the satisfiability conjecture for large k . In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 59–68, 2015.
- [27] Jian Ding, Allan Sly, and Nike Sun. Maximum independent sets on random regular graphs. *Acta Math.*, 217(2):263–340, 2016.
- [28] Jian Ding, Allan Sly, and Nike Sun. Satisfiability threshold for random regular NAE-SAT. *Comm. Math. Phys.*, 341(2):435–489, 2016.
- [29] Tomas Dominguez and Jean-Christophe Mourrat. *Statistical mechanics of mean-field disordered systems: a Hamilton-Jacobi approach*. Zurich Lectures in Advanced Mathematics. European Mathematical Society, Zürich, 2024.
- [30] Ahmed El Alaoui, Andrea Montanari, and Mark Sellke. Optimization of mean-field spin glasses. *Ann. Probab.*, 49(6):2922–2960, 2021.
- [31] David Gamarnik. The overlap gap property: a topological barrier to optimizing over random structures. *Proc. Natl. Acad. Sci. USA*, 118(41):e2108492118, 2021.
- [32] David Gamarnik and Aukosh Jagannath. The overlap gap property and approximate message passing algorithms for p -spin models. *Ann. Probab.*, 49(1):180–205, 2021.
- [33] David Gamarnik, Christopher Moore, and Lenka Zdeborová. Disordered systems insights on computational hardness. *J. Stat. Mech. Theory Exp.*, (11):Paper No. 114015, 41, 2022.
- [34] Andreas Green, Peter Pfaffelhuber, and Anita Winter. Convergence in distribution of random metric measure spaces (Λ -coalescent measure trees). *Probab. Theory Related Fields*, 145(1-2):285–322, 2009.

- [35] Francesco Guerra. Sum rules for the free energy in the mean field spin glass model. *Fields Institute Communications*, 30(11), 2001.
- [36] Francesco Guerra. Broken replica symmetry bounds in the mean field spin glass model. *Comm. Math. Phys.*, 233(1):1–12, 2003.
- [37] Geoffrey E. Hinton. A practical guide to training restricted Boltzmann machines. In *Neural networks: Tricks of the trade*, pages 599–619. Springer, 2012.
- [38] John J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA*, 79(8):2554–2558, 1982.
- [39] Brice Huang and Mark Sellke. Tight Lipschitz hardness for optimizing mean field spin glasses. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science—FOCS 2022*, pages 312–322. IEEE Computer Soc., Los Alamitos, CA, 2022.
- [40] David Jekel, Juspreet Singh Sandhu, and Jonathan Shi. Potential Hessian ascent: The Sherrington-Kirkpatrick model. *Preprint, arXiv:2408.02360*, 2024.
- [41] Florent Krzakala, Andrea Montanari, Federico Ricci-Tersenghi, Guilhem Semerjian, and Lenka Zdeborová. Gibbs states and the set of solutions of random constraint satisfaction problems. *Proc. Natl. Acad. Sci. USA*, 104(25):10318–10323, 2007.
- [42] William A Little. The existence of persistent states in the brain. *Mathematical biosciences*, 19(1-2):101–120, 1974.
- [43] Marc Mézard and Andrea Montanari. *Information, physics, and computation*. Oxford Graduate Texts. Oxford University Press, Oxford, 2009.
- [44] Marc Mézard and Giorgio Parisi. A replica analysis of the travelling salesman problem. *Journal de physique*, 47(8):1285–1296, 1986.
- [45] Marc Mézard, Giorgio Parisi, Nicolas Sourlas, Gérard Toulouse, and Miguel Virasoro. Replica symmetry breaking and the nature of the spin glass phase. *Journal de Physique*, 45(5):843–854, 1984.
- [46] Marc Mézard, Giorgio Parisi, Nicolas Sourlas, Gérard Toulouse, and Miguel Virasoro. Nature of the spin-glass phase. *Phys. Rev. Lett.*, 52(13):1156, 1984.
- [47] Marc Mézard, Giorgio Parisi, and Miguel Virasoro. *Spin glass theory and beyond*, volume 9 of *World Scientific Lecture Notes in Physics*. World Scientific Publishing Co., Inc., Teaneck, NJ, 1987.
- [48] Marc Mézard, Giorgio Parisi, and Riccardo Zecchina. Analytic and algorithmic solution of random satisfiability problems. *Science*, 297(5582):812–815, 2002.
- [49] Rémi Monasson, Riccardo Zecchina, Scott Kirkpatrick, Bart Selman, and Lidror Troyansky. Determining computational complexity from characteristic ‘phase transitions’. *Nature*, 400(6740):133–137, 1999.
- [50] Andrea Montanari. Optimization of the Sherrington-Kirkpatrick Hamiltonian. *SIAM Journal on Computing*, (0):FOCS19–1, 2021.
- [51] Jean-Christophe Mourrat. Nonconvex interactions in mean-field spin glasses. *Probab. Math. Phys.*, 2(2):281–339, 2021.
- [52] Jean-Christophe Mourrat. The Parisi formula is a Hamilton-Jacobi equation in Wasserstein space. *Canad. J. Math.*, 74(3):607–629, 2022.
- [53] Jean-Christophe Mourrat. Free energy upper bound for mean-field vector spin glasses. *Ann. Inst. Henri Poincaré Probab. Stat.*, 59(3):1143–1182, 2023.
- [54] Jean-Christophe Mourrat. Un-inverting the Parisi formula. *Ann. Inst. H. Poincaré Probab. Stat.*, to appear.
- [55] Jean-Christophe Mourrat and Dmitry Panchenko. Extending the Parisi formula along a Hamilton-Jacobi equation. *Electron. J. Probab.*, 25:Paper No. 23, 17, 2020.
- [56] Roberto Mulet, Andrea Pagnani, Martin Weigt, and Riccardo Zecchina. Coloring random graphs. *Phys. Rev. Lett.*, 89(26):268701, 2002.
- [57] Charles Newman. Percolation theory: A selective survey of rigorous results. In *Advances in multiphase flow and related problems*. SIAM, 1986.
- [58] Hidetoshi Nishimori. *Statistical physics of spin glasses and information processing*, volume 111 of *International Series of Monographs on Physics*. Oxford University Press, New York, 2001.
- [59] Manfred Opper and David Saad, editors. *Advanced mean field methods*. Neural Information Processing Series. MIT Press, Cambridge, MA, 2001.
- [60] Dmitry Panchenko. The Parisi ultrametricity conjecture. *Ann. of Math. (2)*, 177(1):383–393, 2013.
- [61] Dmitry Panchenko. *The Sherrington-Kirkpatrick model*. Springer Monographs in Mathematics. Springer, New York, 2013.
- [62] Dmitry Panchenko. The free energy in a multi-species Sherrington-Kirkpatrick model. *Ann. Probab.*, 43(6):3494–3513, 2015.
- [63] Dmitry Panchenko. Free energy in the Potts spin glass. *Ann. Probab.*, 46(2):829–864, 2018.

- [64] Dmitry Panchenko. Free energy in the mixed p -spin models with vector spins. *Ann. Probab.*, 46(2):865–896, 2018.
- [65] Giorgio Parisi. Infinite number of order parameters for spin-glasses. *Phys. Rev. Lett.*, 43(23):1754, 1979.
- [66] Giorgio Parisi. The order parameter for spin glasses: a function on the interval 0-1. *J. Phys. A*, 13(3):1101, 1980.
- [67] Giorgio Parisi. A sequence of approximated solutions to the S-K model for spin glasses. *J. Phys. A*, 13(4):L115–L121, 1980.
- [68] Giorgio Parisi. Order parameter for spin-glasses. *Phys. Rev. Lett.*, 50(24):1946, 1983.
- [69] Giorgio Parisi. Nobel lecture: Multiple equilibria. *Reviews of Modern Physics*, 95(3):030501, 2023.
- [70] Tom Richardson and Ruediger Urbanke. *Modern coding theory*. Cambridge university press, 2008.
- [71] Mark Sellke. Optimizing mean field spin glasses with external field. *Electron. J. Probab.*, 29:Paper No. 4, 47, 2024.
- [72] David Sherrington and Scott Kirkpatrick. Solvable model of a spin-glass. *Phys. Rev. Lett.*, 35(26):1792, 1975.
- [73] Paul Smolensky. Information processing in dynamical systems: foundations of harmony theory. In *Parallel distributed processing: explorations in the microstructure of cognition*, volume 1, pages 194–281. MIT Press, 1986.
- [74] Daniel L. Stein and Charles M. Newman. *Spin glasses and complexity*. Primers in Complex Systems. Princeton University Press, Princeton, NJ, 2013.
- [75] Eliran Subag. Following the ground states of full-RSB spherical spin glasses. *Comm. Pure Appl. Math.*, 74(5):1021–1044, 2021.
- [76] Michel Talagrand. The Parisi formula. *Ann. of Math. (2)*, 163(1):221–263, 2006.
- [77] Michel Talagrand. *Mean field models for spin glasses. Volume I*, volume 54 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer-Verlag, Berlin, 2011.
- [78] Michel Talagrand. *Mean field models for spin glasses. Volume II*, volume 55 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer, Heidelberg, 2011.
- [79] Jérôme Tubiana. *Restricted Boltzmann machines: from compositional representations to protein sequence analysis*. PhD thesis, Université Paris Sciences et Lettres, 2018.
- [80] Jérôme Tubiana and Rémi Monasson. Emergence of compositional representations in restricted boltzmann machines. *Phys. Rev. Lett.*, 118(13):138301, 2017.
- [81] Lenka Zdeborová and Florent Krzakala. Statistical physics of inference: Thresholds and algorithms. *Advances in Physics*, 65(5):453–552, 2016.

(Jean-Christophe Mourrat) DEPARTMENT OF MATHEMATICS, ENS LYON AND CNRS, LYON, FRANCE