

Thresholds of descending algorithms in inference problems

Stefano Sarao Mannelli and Lenka Zdeborová

Universit Paris-Saclay, CNRS, CEA, Institut de physique thorique, 91191,
Gif-sur-Yvette, France.

E-mail: `stefano.sarao@ipht.fr`, `lenka.zdeborova@ipht.fr`

Abstract. We review recent works [1, 2, 3] on analyzing the dynamics of gradient-based algorithms in a prototypical statistical inference problem. Using methods and insights from the physics of glassy systems, these works showed how to understand quantitatively and qualitatively the performance of gradient-based algorithms. Here we review the key results and their interpretation in non-technical terms accessible to a wide audience of physicists in the context of related works.

PACS numbers: 00.00, 20.00, 42.10

Keywords: analysis of algorithms, statistical inference, spin glasses, machine learning.

1. Introduction

Machine learning has achieved astonishing success across real world problems, such as image classification, speech recognition, text processing, and physical problems, from quantum physics [4], to astrophysics [5], to high-energy physics [6]. Despite these practical successes, a large number of aspects still lacks theoretical understanding. Practitioners identified several prescriptions to construct a working machine learning applications, but it is often unclear why those recipes are effective. Consider a typical classification task, where a dataset consisting of pictures of cats and dogs is provided to the machine with the correct labels. What follows is the minimization of a cost function. Given new images of pets, the goal of the machine is to be able to correctly classify them into cats and dogs, thus successfully generalizing from what it has seen.

The optimization process itself is puzzling. In general, the cost function is high-dimensional and non-convex. Intuition would suggest that a random initialization would lead to some local spurious, non-informative, minimum with very little hope to achieve a good generalization. Instead, in practice even the use of vanilla gradient descent often leads to good generalization. Part of the computer science community analysed the problem geometrically by studying the properties of the cost function [7, 8, 9, 10, 11]. They consider generative models where a signal is observed through a noisy channel

and it is possible to tune its strength with respect to the strength of the noise, the signal to noise ratio (SNR), and change the landscape. They showed that in a variety of problems, all the minima become equally good or the spurious minima disappear as the SNR becomes sufficiently large, thus making the landscape trivial.

In this work we review the recent effort towards an understanding of the learning dynamics using the tools of disordered systems [2, 3], and we discuss the difference in performance between message passing algorithms and algorithms for sampling a high-dimensional potential [1]. The relation between the two approaches becomes apparent from the point of view of Bayesian statistics. Let \mathbf{X} be the guess on the hidden signal and \mathbf{Y} the observation, we can express plausible is to observe \mathbf{Y} given our guess, i.e. the likelihood $\mathbb{P}[\mathbf{Y}|\mathbf{X}]$. Bayes formula allows to invert the likelihood into the posterior probability $\mathbb{P}[\mathbf{X}|\mathbf{Y}] \propto \mathbb{P}[\mathbf{Y}|\mathbf{X}]\mathbb{P}[\mathbf{X}]$, that also includes prior information on the guess, such as sparsity or norm constrains. We can write an approximate expression

$$\mathbb{P}[\mathbf{X}|\mathbf{Y}] \approx_{\beta \approx 1} \frac{1}{\mathbb{P}[\mathbf{Y}]} \overbrace{\mathbb{P}[\mathbf{Y}|\mathbf{X}]^\beta}^{-cost} \mathbb{P}[\mathbf{X}] \doteq \frac{1}{\mathcal{Z}[\mathbf{Y}]} e^{-\beta \mathcal{H}(\mathbf{X}, \mathbf{Y})}. \quad (1)$$

In the last equality we identify the terms with a Gibbs distribution with inverse temperature β . Given the posterior we can estimate the signal by considering the expected value:

$$\hat{\mathbf{x}}_\beta = \mathbb{E}_{\mathbf{X}|\mathbf{Y};\beta}[\mathbf{X}]. \quad (2)$$

Observe that when the inverse temperature parameter equals 1, Eq. (1) is the posterior probability of the problem. As β tends to infinity, the cost dominates and optimizing will maximize the likelihood.

In the eyes of a statistical physicist, the expected value $\hat{\mathbf{x}}_\beta$ would rather be called \mathbf{m} as it is formally identical to the magnetization of a system under the action of the Hamiltonian. However, the exact computation a this expected value exactly is prohibitive in large dimension, in fact it is as complex as evaluating the partition function. In order to avoid such complication numerous ingenious techniques have been considered in the past to obtain an approximate estimation. Two main approaches consist of approximating the posterior, and sampling the posterior.

- The idea of adapting the approximations proposed in disordered systems to computer science problems is not recent, and early works appeared in the 80s and 90s [12, 13]. Ideas from physics were transferred to problems in signal processing and optimisation, providing both theoretical understanding and practical algorithms based on Cavity Method and its variations [14, 15, 16]. Those methods have the advantage of being at the same time algorithms and analytical tools. In many problems they were proved to be asymptotically optimal [17, 18, 19], in the sense that information-theoretically they achieve the best performance in polynomial time.
- The best known algorithms that sample the posterior are Monte Carlo and the Langevin algorithm. Studies on the Langevin algorithm in disordered systems have

their root in the late 70s [20, 21, 22, 23]. Despite the dynamics was understood for some recurrent neural networks in long-time regime [24, 25], generalizing and solving the corresponding equations is very difficult even in the simplest models of statistical inference [26]. Consequently, analysis of the performance of gradient-based algorithms such as the Langevin algorithm remains an open problem. A progress on this question was recently made in a series of works [1, 2, 3] that we review here.

The paper is organized as follows: in Section 2 we introduce the model, in Section 3 we propose a comparison between sampling algorithms and approximate algorithms, in Section 4 the gradient flow algorithm is analyzed and compared with the energy landscape.

2. Spiked matrix-tensor model

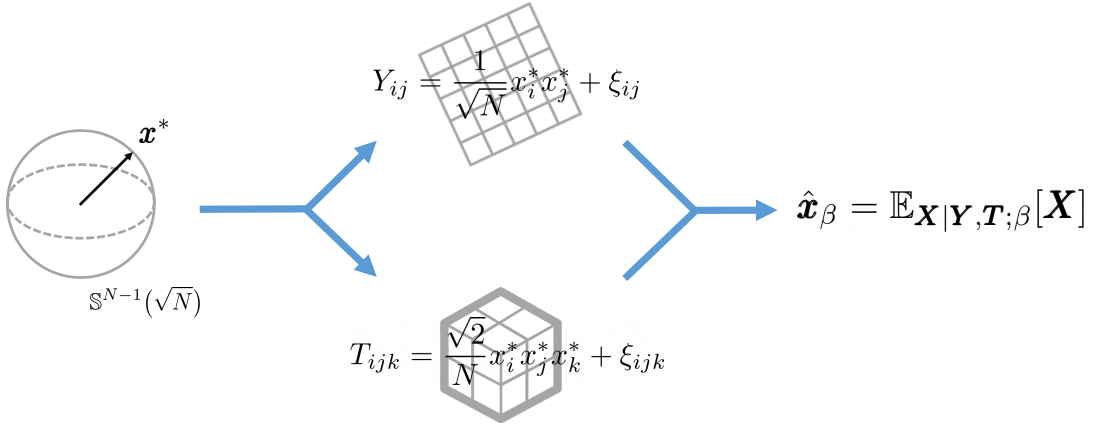


Figure 1. Cartoonish representation of the generative process. The left part of the image represents the teacher who samples the ground truth \mathbf{x}^* and uses it to generate the two observations that contain noise, $\xi_{ij} \sim \mathcal{N}(0, \Delta_2)$ and $\xi_{ijk} \sim \mathcal{N}(0, \Delta_3)$. The student, on the right part of the figure, receives the observations and constructs an estimation of the signal, $\hat{\mathbf{x}}_\beta$, by computing the expected value.

The model that we study in this report is the *spiked matrix-tensor model*, known in physics as a planted version of the spherical mixed p -spin model [27, 23, 28]. Planting is a technique introduced to study statistical inference and learning problems using the same methods as for their optimization counterparts [29]. Planting appears as an additional ferromagnetic bias towards a planted solution (or ground-truth) in the Hamiltonian. In its application to inference, planting permits to introduce a signal, the ground-truth solution, in the formulation of the problem. In the neural-network-learning language this formulation is called *teacher-student scenario*: the *teacher* knows the ground-truth and uses it to generate data, the *student* has to use the data to infer the ground-truth.

The *spiked matrix-tensor model* was introduced in [1, 2, 3] in order to build an inference problem for which the behaviour of the gradient-based dynamics is exactly

solvable. For the sake of simplicity we will consider $p = 3$, which means that the *teacher* samples the ground truth \mathbf{x}^* and generates the data, a matrix and a order 3 tensor. The process is noisy and the data that the student receives, Y_{ij} and T_{ijk} , have an intrinsic Gaussian noise of variance Δ_2 and Δ_3 respectively. The two observations are rescaled in order to have an extensive free energy in the size of the system N . The generative process is represented in Fig. 1. Substituting the data into the posterior Eq. (1) and absorbing constant terms into the pre-factor, we obtain the Hamiltonian

$$\mathcal{H} = -\frac{1}{\Delta_2\sqrt{N}} \sum_{i<j} \xi_{ij} x_i x_j - \frac{\sqrt{2}}{\Delta_3 N} \sum_{i<j<k} \xi_{ijk} x_i x_j x_k - \frac{N}{2\Delta_2} m^2 - \frac{N}{3\Delta_3} m^3, \quad (3)$$

where $m = \frac{1}{N} \sum_i x_i x_i^*$ is the overlap with the signal. Observe that the noise terms ξ_{ij} (ξ_{ijk}) in the equation are rescaled by \sqrt{N} (N respectively) in order to have a problem that is neither impossibly hard (very high noise) nor trivially easy (very small noise). Under this choice of scaling of noise, we observe different transitions for values of Δ_2 and Δ_3 of order $O(1)$.

The spiked matrix-tensor model is a natural candidate for our analysis as it has high-dimensional non-convex energy landscape. The *algorithmic transition*, after which algorithms start to detect the signal, occurs at the same noise scaling as the *information-theoretic transition* for detection. The model is analytically tractable using different methods allowing to experiment and compare. We remark that the spiked tensor model does not have an algorithmic and information-theoretic transitions occurring in the same scaling regime of the noise, thus it is a less interesting candidate for our analysis.

3. Sampling algorithms vs approximate algorithms

We are going to consider an algorithmic version of cavity method as an example of an approximation algorithm.[14]. This algorithm was developed independently in the information theory and Bayesian inference community under the name of belief propagation [30, 31, 32]. In the case of fully connected models, belief propagation can be simplified by assuming a Gaussian structure in the beliefs, leading to the Approximate Message Passing (AMP) algorithm [33, 16]. AMP presents numerous remarkable features: it provably achieves optimal performances in many problems including the spiked matrix-tensor [17, 18, 19, 1] and its average behaviour can be analytically followed by a set of equations called *state evolution* [34]. State evolution equations allow to portrait the phase diagram of this model, see Fig. 2, it was done in [1] generalizing the results of [18] on the spiked tensor model. The phase diagram can now be used as a baseline for the behaviour of the sampling algorithms.

In order to sample from the posterior probability it is necessary to design a dynamics that has the posterior probability as its stationary measure at large times. A typical sampling algorithm with this objective is the Langevin algorithm. Given a Hamiltonian

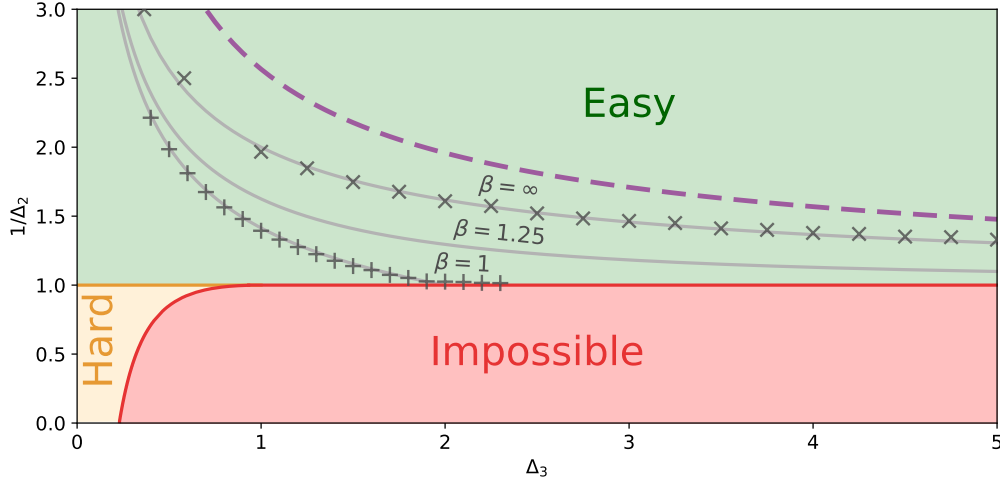


Figure 2. Phase diagram of the spiked matrix-tensor model with $p = 3$. As the variances of the noise in matrix, Δ_2 , and the noise in the tensor, Δ_3 , change different phases appear. We can distinguish the *easy* (green) phase where AMP can detect the signal, the *impossible* (red) phase where it is information-theoretically impossible to detect the signal, and the *hard* (yellow) phase where the signal can in principle be detected, but it is expected to take exponential time as it requires to jump over an energy barrier. The grey lines in the easy phase represent the algorithmic transition of the Langevin algorithm for $\beta = 1$, $\beta = 1.25$, and $\beta = \infty$. For a fixed β , the Langevin algorithm starts to detect the signal above the respective grey line. The plus marks and the cross marks are the extrapolation of the Langevin threshold from the numerical study of the dynamical equations. We can observe good agreement with the analytical prediction. The purple dashed line characterizes the trivialization transition, above that line the energy landscape does not present any spurious minima.

\mathcal{H} of a spherical system, Langevin dynamics describes the evolution of the system coupled with a thermal bath at temperature $T = 1/\beta$

$$\frac{d}{dt}x_i(t) = -\frac{\partial \mathcal{H}(\mathbf{x})}{\partial x_i}(t) - \mu(t)x_i(t) + \eta_i(t) \quad (4)$$

where $\mu(t)$ is a Lagrange multiplier that imposes the spherical constraint and $\eta_i(t)$ is the Langevin noise with $\mathbb{E}[\eta_i(t)] = 0$ and $\mathbb{E}[\eta_i(t)\eta_j(t')] = \frac{2}{\beta}\delta_{ij}\delta(t-t')$. In the late 70s, techniques [21] for the study of Langevin dynamics were adapted to disordered systems providing a set of PDEs on the evolution of few relevant observables. More recently, the results of these techniques have been proved with mathematical rigour in the mixed p -spin model [35, 36]. Those methods have been generalized to the study of planted systems [37] and applied to the present problem in [1, 3]. Two variants of the dynamical mean field theory were used to derive the corresponding equations: the dynamical cavity method was used in [1], and the generating functional formalism [3]. The equations obtained characterize the evolution of: the alignment of the system with the ground truth $m(t) = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_{\xi, \eta} \sum_i x_i(t)x_i^*$, the self-alignment at different times $C(t, t') = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_{\xi, \eta} \sum_i x_i(t)x_i(t')$, and the response to a perturbation of the

Hamiltonian at a previous time $R(t, t') = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_{\xi, \eta} \sum_i \delta x_i(t) / \delta \eta_i(t')$.

$$\begin{aligned} \frac{\partial}{\partial t} C(t, t') &= -\mu(t) C(t, t') + Q'(m(t))m(t') + \int_0^t R(t, t'') Q''(C(t, t'')) C(t', t'') dt'' \\ &+ \int_0^{t'} R(t', t'') Q'(C(t, t'')) dt'', \end{aligned} \quad (5)$$

$$\frac{\partial}{\partial t} R(t, t') = -\mu(t) R(t, t') + \int_{t'}^t R(t, t'') Q''(C(t, t'')) R(t'', t') dt'', \quad (6)$$

$$\frac{d}{dt} m(t) = -\mu(t) m(t) + Q'(m(t)) + \int_0^t R(t, t'') m(t'') Q''(C(t, t'')) dt'', \quad (7)$$

with $Q(x) = \frac{x^2}{2\Delta_2} + \frac{x^3}{3\Delta_3}$, the initial conditions $R(t, t') = 0$ for all $t < t'$, $\lim_{t' \rightarrow t^-} R(t, t') = 1 \forall t$ and $C(t, t) = 1 \forall t$ that allows to derive an additional equation for $\mu(t)$. The spiked matrix-tensor model has the nice feature of having a closed form for these equations, allowing an easier evaluation of the numerical solution by propagation from the initial conditions. In [1, 2] the limits of Langevin and gradient descent (respectively) have been evaluated numerically by extrapolation from the numerical solutions, see Fig. 2. In general the dynamical equations do not close, thus a self-consistent loop is necessary in order to evaluate a numerical solution limiting the times accessible in the numerics [38].

An alternative can be derived from the work [39] where the authors proposed an ansatz for the large time behaviour of the p -spin model, which assumes two time scales. The authors also showed that the dynamics is attracted by states - called *threshold states* - characterized by a Hessian that displays marginality, i.e. its spectrum touches the zero. In [3], these two ideas are used to derive the analytical threshold of the Langevin dynamics and gradient descent, by assuming that initially the dynamics will tend to the threshold states and at later times it will increase the alignment with the ground truth. The growth is exponential and the exponent is $\Lambda(\Delta_2, \Delta_3; \beta) = \frac{1}{\Delta_2} - \sqrt{\frac{1}{\Delta_2} + \frac{2(1-\Delta_2/\beta)}{\Delta_3}}$, the phase transition occurs when the exponent crosses the null value. Analytical and numerical results are shown in Fig. 2 giving a perfect agreement.

The results suggest that sampling algorithms have worse algorithmic threshold than AMP. This idea was foreseen in [40], where the authors used a large deviation analysis [41] to find exponentially many atypical glassy states in the landscape. They conjectured that the presence of this atypical glassy states may block the dynamics of sampling algorithms. The same analysis was also performed in the spiked matrix-tensor model confirming their findings [1].

Another signature of the different transitions appears in the evolution of AMP and Langevin dynamics, Fig. 3. For a fixed value of Δ_2 (with $\Delta_2 < 1$), we can compare evolutions for different values of Δ_3 . As the system gets closer to the transition, the time to find the transition increases. We can thus observe that AMP maintains the same typical time to find the solution for the different values of Δ_3 , instead the typical time of the Langevin dynamics increases exponentially as Δ_3 becomes smaller. This

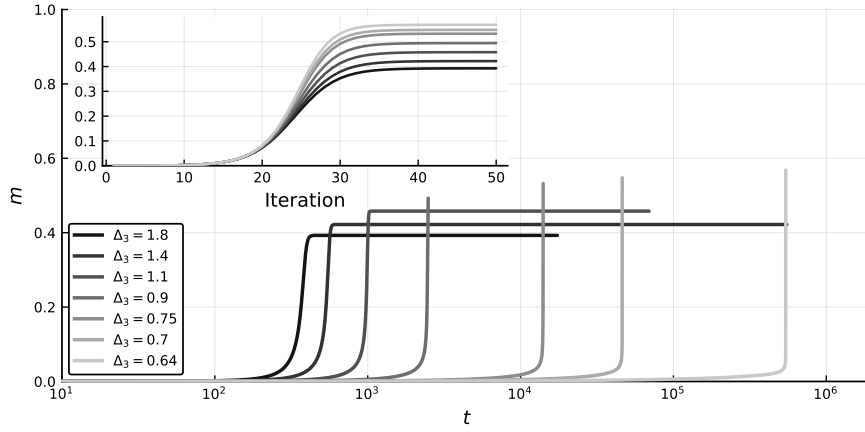


Figure 3. Comparison of the evolution of the overlap with the signal in Langevin dynamics and AMP (inset), for $\Delta_2 = 0.70$ and several values of Δ_3 .

illustrates the counter-intuitive finding that making the problem simpler by decreasing the noise in the tensor actually harms the Langevin evolution.

4. Gradient flow and geometry

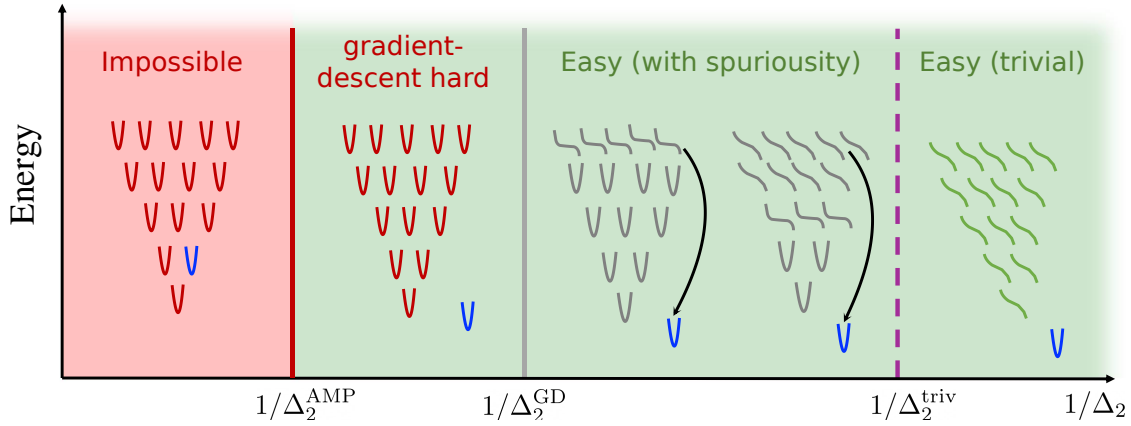


Figure 4. The cartoon represents the energy landscape for an arbitrary value of $\Delta_3 > 1$. The good minimum is drawn in blue. $1/\Delta_2$ plays the role of the SNR. Starting from low SNR, in the impossible region it is thermodynamically impossible to distinguish between good and bad minima. Increasing the SNR, the good minimum becomes energetically favored but the exponential number of the spurious minima stops the dynamics. At larger SNR the threshold minima becomes saddles pointing toward the good minimum and gradient descent starts to find the solution. Finally the SNR becomes larger than the trivialization threshold and only the good minimum survives.

It was already clear in [39] that β enters in a smooth way in the dynamical equations, thus studying the limit $\beta \rightarrow \infty$ we can derive the behaviour of gradient descent dynamics. In machine learning gradient descent and its several variations (e.g.

stochastic gradient descent) are usually used to minimize the cost function. Currently very few problems are amenable to analytical analysis of the dynamics.

In the 80s [42] and in the early 2000 [43, 44, 45] there was an effort to understand the geometrical structure of the energy landscape in disordered models. Given the number of critical points of the model, \mathcal{N}_c , the studies focused on the annealed (and quenched) complexity defined as $\log \mathbb{E}[\mathcal{N}_c]$ (and $\log \mathbb{E}[\mathcal{N}_c]$, respectively). The authors used an expression that enumerates the number of critical points, namely the *Kac-Rice formula* [46], computed using replica theory. Recently another approach for the evaluation of the Kac-Rice formula has been proposed that uses random matrix theory, giving fruitful results in the p -spin model (planted and unplanted) [47, 48, 49]. In [2] the analysis was generalized to the spiked matrix-tensor model allowing to distinguish between regions where exponentially many minima are present and regions where only the good minima appear. The line that separates them is the *trivialization transition* line. As gradient descent is run above this line, provided that the time-discretization is thin enough, we have a guarantee of finding the good minimum. Many papers [8, 9, 10, 11] have focused on this aspect showing in several problems that as the SNR is strong enough the bad minima disappear, or all the minima become equally good. In the spiked matrix-tensor model geometrical trivialization and gradient descent transition can both be pinpointed on the phase diagram. The results, Fig. 2, show that gradient descent starts to detect the signal before the trivialization transition has occurred. Although the algorithmic threshold of gradient descent can not occur after the trivialization transition, it might appear counter-intuitive to understand why the two lines do not coincide and there is a distance of order $\mathcal{O}(1)$ that separates them. In [3] the puzzle was solved, Fig. 4. The authors showed that, moving from a low SNR region where the algorithm fails, the algorithmic transition of gradient descent appears when the dominant minima (the threshold states or threshold minima) develop an instability, a Baik-BenArous-Péché instability [50], becoming saddles with a single negative direction that points toward the signal. In this region there are still exponentially many minima that do not carry information on the signal, nevertheless the dynamics is first attracted by the saddles at the threshold that shield the system from the bad minima and point in the right direction.

5. Conclusions

In this manuscript we analyze recent progresses on the understanding of the dynamics in inference problems using the tools developed in statistical physics. The attention is focused on the spiked matrix-tensor model (planted spherical mixed p -spin model) as a prototypical example of inference. The results on this model unveil unexpected behaviours of the dynamics and explain them from both a dynamical and a geometrical perspective. The techniques briefly summarized in the work can be extended to other models and some of the findings can be verified numerically. These are thrilling directions that we hope to pursue in the future.

Acknowledgments

We thank G. Biroli, C. Cammarota, F. Krzakala and P. Urbani for the collaborations that led to these results and G. Bassignana, S. Goldt and O. Scarlatella for reading the draft of the manuscript. We acknowledge funding from the ERC under the European Unions Horizon 2020 Research and Innovation Programme Grant Agreement 714608-SMiLe.

References

- [1] Stefano Sarao Mannelli, Giulio Biroli, Chiara Cammarota, Florent Krzakala, Pierfrancesco Urbani, and Lenka Zdeborová. Marvels and pitfalls of the Langevin algorithm in noisy high-dimensional inference. *arXiv preprint arXiv:1812.09066*, 2018.
- [2] Stefano Sarao Mannelli, Florent Krzakala, Pierfrancesco Urbani, and Lenka Zdeborová. Passed & spurious: Descent algorithms and local minima in spiked matrix-tensor models. In *International Conference on Machine Learning*, pages 4333–4342, 2019.
- [3] Stefano Sarao Mannelli, Giulio Biroli, Chiara Cammarota, Florent Krzakala, and Lenka Zdeborová. Who is afraid of big bad minima? analysis of gradient-flow in spiked matrix-tensor models. In *Advances in Neural Information Processing Systems*, pages 8676–8686, 2019.
- [4] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, 2017.
- [5] Nicholas M Ball and Robert J Brunner. Data mining and machine learning in astronomy. *International Journal of Modern Physics D*, 19(07):1049–1106, 2010.
- [6] Alexander Radovic, Mike Krz Williams, David Rousseau, Michael Kagan, Daniele Bonacorsi, Alexander Himmel, Adam Aurisano, Kazuhiro Terao, and Taritree Wongjirad. Machine learning at the energy and intensity frontiers of particle physics. *Nature*, 560(7716):41, 2018.
- [7] Afonso S Bandeira, Nicolas Boumal, and Vladislav Voroninski. On the low-rank approach for semidefinite programs arising in synchronization and community detection. In *Conference on learning theory*, pages 361–382, 2016.
- [8] Kenji Kawaguchi. Deep learning without poor local minima. In *Advances in Neural Information Processing Systems*, pages 586–594, 2016.
- [9] Rong Ge, Jason D Lee, and Tengyu Ma. Matrix completion has no spurious local minimum. In *Advances in Neural Information Processing Systems*, pages 2973–2981, 2016.
- [10] Rong Ge, Chi Jin, and Yi Zheng. No spurious local minima in nonconvex low rank problems: A unified geometric analysis. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1233–1242, 2017.
- [11] Simon S Du, Jason D Lee, Yuandong Tian, Aarti Singh, and Barnabas Poczos. Gradient descent learns one-hidden-layer CNN: Dont be afraid of spurious local minima. In *International Conference on Machine Learning*, pages 1338–1347, 2018.
- [12] Marc Mézard, Giorgio Parisi, and Miguel-Angel Virasoro. *Spin glass theory and beyond*. World Scientific Publishing, 1987.
- [13] Andreas Engel and Christian Van den Broeck. *Statistical mechanics of learning*. Cambridge University Press, 2001.
- [14] Marc Mézard and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009.
- [15] Lenka Zdeborová and Florent Krzakala. Statistical physics of inference: thresholds and algorithms. *Advances in Physics*, 65(5):453–552, 2016.
- [16] Thibault Lesieur, Florent Krzakala, and Lenka Zdeborová. Constrained low-rank matrix estimation: Phase transitions, approximate message passing and applications. *Journal of Statistical Mechanics: Theory and Experiment*, 2017(7):073403, 2017.

- [17] Léo Miolane. Fundamental limits of low-rank matrix estimation: the non-symmetric case. *arXiv preprint arXiv:1702.00473*, 2017.
- [18] Thibault Lesieur, Léo Miolane, Marc Lelarge, Florent Krzakala, and Lenka Zdeborová. Statistical and computational phase transitions in spiked tensor estimation. In *Information Theory (ISIT), 2017 IEEE International Symposium on*, pages 511–515. IEEE, 2017.
- [19] Jean Barbier, Florent Krzakala, Nicolas Macris, Léo Miolane, and Lenka Zdeborová. Optimal errors and phase transitions in high-dimensional generalized linear models. *Proceedings of the National Academy of Sciences*, 116(12):5451–5460, 2019.
- [20] Paul Cecil Martin, ED Siggia, and HA Rose. Statistical dynamics of classical systems. *Physical Review A*, 8(1):423, 1973.
- [21] C De Dominicis. Dynamics as a substitute for replicas in systems with quenched random impurities. *Physical Review B*, 18(9):4913, 1978.
- [22] Theodore R Kirkpatrick and Devarajan Thirumalai. Dynamics of the structural glass transition and the p-spin interaction spin-glass model. *Physical review letters*, 58(20):2091, 1987.
- [23] A Crisanti, H Horner, and H-J Sommers. The spherical p -spin interaction spin-glass model. *Zeitschrift für Physik B Condensed Matter*, 92(2):257–271, 1993.
- [24] Haim Sompolinsky, Andrea Crisanti, and Hans-Jurgen Sommers. Chaos in random neural networks. *Physical review letters*, 61(3):259, 1988.
- [25] ACC Coolen. Statistical mechanics of recurrent neural networks II. Dynamics. *arXiv preprint cond-mat/0006011*, 2000.
- [26] Elisabeth Agoritsas, Giulio Biroli, Pierfrancesco Urbani, and Francesco Zamponi. Out-of-equilibrium dynamical mean-field equations for the perceptron model. *Journal of Physics A: Mathematical and Theoretical*, 51(8):085002, 2018.
- [27] David J Gross and Marc Mézard. The simplest spin glass. *Nuclear Physics B*, 240(4):431–452, 1984.
- [28] Andrea Crisanti and Luca Leuzzi. Spherical 2+ p spin-glass model: An exactly solvable model for glass to spin-glass transition. *Physical review letters*, 93(21):217203, 2004.
- [29] Florent Krzakala and Lenka Zdeborová. Hiding quiet solutions in random constraint satisfaction problems. *Physical review letters*, 102(23):238701, 2009.
- [30] Robert Gallager. Low-density parity-check codes. *IRE Transactions on information theory*, 8(1):21–28, 1962.
- [31] Judea Pearl. *Reverend Bayes on inference engines: A distributed hierarchical approach*. Cognitive Systems Laboratory, School of Engineering and Applied Science, University of California, Los Angeles, 1982., 1982.
- [32] Judea Pearl. Fusion, propagation, and structuring in belief networks. *Artificial intelligence*, 29(3):241–288, 1986.
- [33] David L Donoho, Arian Maleki, and Andrea Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, Nov 2009.
- [34] Adel Javanmard and Andrea Montanari. State evolution for general approximate message passing algorithms, with applications to spatial coupling. *Information and Inference: A Journal of the IMA*, 2(2):115–144, 2013.
- [35] Gérard Ben Arous, Amir Dembo, and Alice Guionnet. Cugliandolo-Kurchan equations for dynamics of spin-glasses. *Probability theory and related fields*, 136(4):619–660, 2006.
- [36] Amir Dembo and Eliran Subag. Dynamics for spherical spin glasses: disorder dependent initial conditions. *arXiv preprint arXiv:1908.01126*, 2019.
- [37] Chiara Cammarota and Giulio Biroli. Aging and relaxation near random pinning glass transitions. *EPL (Europhysics Letters)*, 98(1):16011, 2012.
- [38] Felix Roy, Giulio Biroli, Guy Bunin, and Chiara Cammarota. Numerical implementation of dynamical mean field theory for disordered systems: application to the Lotka-Volterra model of ecosystems. *Journal of Physics A: Mathematical and Theoretical*, 2019.

- [39] Leticia F Cugliandolo and Jorge Kurchan. Analytical solution of the off-equilibrium dynamics of a long-range spin-glass model. *Physical Review Letters*, 71(1):173, 1993.
- [40] Fabrizio Antenucci, Silvio Franz, Pierfrancesco Urbani, and Lenka Zdeborová. Glassy nature of the hard phase in inference problems. *Physical Review X*, 9(1):011020, 2019.
- [41] Rémi Monasson. Structural glass transition and the entropy of the metastable states. *Physical review letters*, 75(15):2847, 1995.
- [42] AJ Bray and MA Moore. Metastable states, internal field distributions and magnetic excitations in spin glasses. *Journal of Physics C: Solid State Physics*, 14(19):2629, 1981.
- [43] Andrea Crisanti, Luca Leuzzi, Giorgio Parisi, and Tommaso Rizzo. Complexity in the sherrington-kirkpatrick model in the annealed approximation. *Physical Review B*, 68(17):174401, 2003.
- [44] Andrea Cavagna, Irene Giardinà, Giorgio Parisi, and Marc Mézard. On the formal equivalence of the tap and thermodynamic methods in the sk model. *Journal of Physics A: Mathematical and General*, 36(5):1175, 2003.
- [45] A Crisanti, L Leuzzi, G Parisi, and T Rizzo. Quenched computation of the dependence of complexity on the free energy in the sherrington-kirkpatrick model. *Physical Review B*, 70(6):064423, 2004.
- [46] Robert J Adler and Jonathan E Taylor. *Random fields and geometry*. Springer Science & Business Media, 2009.
- [47] Antonio Auffinger, Gérard Ben Arous, and Jiří Černý. Random matrices and complexity of spin glasses. *Communications on Pure and Applied Mathematics*, 66(2):165–201, 2013.
- [48] Gérard Ben Arous, Song Mei, Andrea Montanari, and Mihai Nica. The landscape of the spiked tensor model. *arXiv preprint arXiv:1711.05424*, 2017.
- [49] Valentina Ros, Gérard Ben Arous, Giulio Biroli, and Chiara Cammarota. Complex energy landscapes in spiked-tensor and simple glassy models: Ruggedness, arrangements of local minima, and phase transitions. *Physical Review X*, 9(1):011003, 2019.
- [50] Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability*, 33(5):1643–1697, 2005.