

ACCEPTED MANUSCRIPT • OPEN ACCESS

Mean-field inference methods for neural networks

To cite this article before publication: Marylou Gabri  2020 *J. Phys. A: Math. Theor.* in press <https://doi.org/10.1088/1751-8121/ab7f65>

Manuscript version: Accepted Manuscript

Accepted Manuscript is "the version of the article accepted for publication including all changes made as a result of the peer review process, and which may also include the addition to the article by IOP Publishing of a header, an article ID, a cover sheet and/or an 'Accepted Manuscript' watermark, but excluding any other editing, typesetting or other changes made by IOP Publishing and/or its licensors"

This Accepted Manuscript is   2020 The Author(s). Published by IOP Publishing Ltd.

As the Version of Record of this article is going to be / has been published on a gold open access basis under a CC BY 3.0 licence, this Accepted Manuscript is available for reuse under a CC BY 3.0 licence immediately.

Everyone is permitted to use all or part of the original content in this article, provided that they adhere to all the terms of the licence <https://creativecommons.org/licenses/by/3.0>

Although reasonable endeavours have been taken to obtain all necessary permissions from third parties to include their copyrighted content within this article, their full citation and copyright line may not be present in this Accepted Manuscript version. Before using any content from this article, please refer to the Version of Record on IOPscience once published for full citation and copyright details, as permissions may be required. All third party content is fully copyright protected and is not published on a gold open access basis under a CC BY licence, unless that is specifically stated in the figure caption in the Version of Record.

View the [article online](#) for updates and enhancements.

Mean-field inference methods for neural networks

Marylou Gabri  ^{1,2}

¹Center for Data Science, New York University

²Center for Computational Math, Flatiron Institute

Abstract

Machine learning algorithms relying on deep neural networks recently allowed a great leap forward in artificial intelligence. Despite the popularity of their applications, the efficiency of these algorithms remains largely unexplained from a theoretical point of view. The mathematical description of learning problems involves very large collections of interacting random variables, difficult to handle analytically as well as numerically. This complexity is precisely the object of study of statistical physics. Its mission, originally pointed towards natural systems, is to understand how macroscopic behaviors arise from microscopic laws. Mean-field methods are one type of approximation strategy developed in this view. We review a selection of classical mean-field methods and recent progress relevant for inference in neural networks. In particular, we remind the principles of derivations of high-temperature expansions, the replica method and message passing algorithms, highlighting their equivalences and complementarities. We also provide references for past and current directions of research on neural networks relying on mean-field methods.

Contents

1	Introduction	2
2	Machine learning with neural networks	3
2.1	Supervised learning	3
2.2	Unsupervised learning	5
3	Statistical inference and the statistical physics approach	8
3.1	Statistical inference	8
3.2	Statistical physics of disordered systems, first appearance on stage	9
4	Selected overview of mean-field treatments: free energies and algorithms	11
4.1	Naive mean-field	12
4.2	Thouless Anderson and Palmer equations	13
4.3	Belief propagation and approximate message passing	16
4.4	Replica method	24
5	Further extensions of interest for learning	27
5.1	Streaming AMP for online learning	27
5.2	Algorithms and free energies beyond i.i.d. matrices	28
5.3	Multi-value AMP	30
5.4	Model composition and multi-layer inference	31
6	Some applications	34
6.1	A brief pre-deep learning history	34
6.2	Some current directions of research	35
7	Conclusion	37
A	Index of notations and abbreviations	38

1 Introduction

B Statistical models representations	39
C Mean-field identity	41
D Georges-Yedidia expansion for generalized Boltzmann machines	42
E Vector Approximate Message Passing for the GLM	46
F Multi-value AMP derivation for the GLM	48
F.1 Approximate Message Passing	48
F.2 State Evolution	51
References	54

1 Introduction

With the continuous improvement of storage techniques, the amount of available data is currently growing exponentially. While it is not humanly feasible to treat all the data created, *machine learning*, as a class of algorithms that allows to automatically infer structure in large data sets, is one possible response. In particular, *deep learning* methods, based on neural networks, have drastically improved performances in key fields of artificial intelligence such as image processing, speech recognition or text mining. A good review of the first successes of this technology published in 2015 is [LBH15]. A few years later, the current state-of-the-art of this very active line of research is difficult to envision globally. However, the complexity of deep neural networks remains an obstacle to the understanding of their great efficiency. Made of many layers, each of which constituted of many neurons, themselves accompanied by a collection of parameters, the set of variables describing completely a typical neural network is impossible to only visualize. Instead, aggregated quantities must be considered to characterize these models and hopefully help and explain the learning process. The first open challenge is therefore to identify the relevant observables to focus on. Often enough, what seems interesting is also what is hard to calculate. In the high-dimensional regime we need to consider, exact analytical forms are unknown most of the time and numerical computations are ruled out. , ways of approximation that are simultaneously simple enough to be tractable and fine enough to retain interesting features are highly needed.

In the context where dimensionality is an issue, physicists have experimented that macroscopic behaviors are typically well described by the theoretical limit of infinitely large systems. Under this *thermodynamic* limit, the statistical physics of disordered systems offers powerful frameworks of approximation called *mean-field theories*. Interactions between physics and neural network theory already have a long history as we will discuss in Section 6.1. Yet, interconnections have been re-heightened by the recent progress in deep learning, which also brought new theoretical challenges.

Here, we wish to provide a concise methodological review of fundamental mean-field inference methods with their application to neural networks in mind. Our aim is also to provide a unified presentation of the different approximations allowing to understand how they relate and differ. Readers may also be interested in related review papers. Another methodological review is [ALG13], particularly interested in applications to neurobiology. Methods presented in the latter reference have a significant overlap with what will be covered in the following. Some elements of random matrix theory are there additionally introduced. The approximations and algorithms which will be discussed here are also largely reviewed in [ZK16]. This previous paper includes more details on spin glass theory, which originally motivated the development of the classical mean-field methods, and particularly focuses on community detection and linear estimation. Despite the significant overlap and beyond their differing motivational applications, the two previous references are also anterior to some recent exciting developments in mean-field inference covered in the present review, in particular extensions towards multi-layer networks. An older, yet very interesting, reference is the workshop proceedings [OS01], which collected both insightful introductory papers and research developments for the applications of mean-field methods in machine learning. Finally, the recent [CCC+19] covers more generally the connections between physical sciences and machine learning yet without detailing the methodologies. This review provides a very good list of references where statistical physics methods were used for learning theory, but also where machine learning helped in turn physics research.

Given the literature presented below is at the cross-roads of deep learning and disordered systems physics, we include short introductions to the fundamental concepts of both domains. These Sections 2 and 3 will help readers with one or the other background, but can be skipped by experts. In Section 4, classical mean-field inference approximations are derived on neural network examples. Section 5 covers some recent extensions of the classical methods that are of particular interest for applications to neural networks. We review in Section 6 a selection of important historical and current directions of research in neural networks leveraging mean-field methods. As a conclusion, strengths, limitations and perspectives of mean-field methods for neural networks are discussed in Section 7.

2 Machine learning with neural networks

Machine learning is traditionally divided into three classes of problems: supervised, unsupervised and reinforcement learning. For all of them, the advent of deep learning techniques, relying on deep neural networks, has brought great leaps forward in terms of performance and opened the way to new applications. Nevertheless, the utterly efficient machinery of these algorithms remains full of theoretical puzzles. This Section provides fundamental concepts in machine learning for the unfamiliar reader willing to approach the literature at the crossroads of statistical physics and deep learning. We also take this Section as an opportunity to introduce the current challenges in building a strong theoretical understanding of deep learning. A comprehensive reference is [GBC16], while [MWD⁺18] offers a broad introduction to machine learning specifically addressed to physicists.

2.1 Supervised learning

Learning under supervision Supervised learning aims at discovering systematic input to output mappings from examples. Classification is a typical supervised learning problem: for instance, from a set of pictures of cats and dogs labelled accordingly, the goal is to find a function able to predict in any new picture the species of the displayed pet.

In practice, the *training set* is a collection of P example pairs $\mathcal{D} = \{\underline{x}_{(k)}, \underline{y}_{(k)}\}_{k=1}^P$ from an input data space $\mathcal{X} \subseteq \mathbb{R}^N$ and an output data space $\mathcal{Y} \subseteq \mathbb{R}^M$. Formally, they are assumed to be i.i.d. samples from a joint distribution $p(\underline{x}, \underline{y})$. The predictor h is chosen by a training algorithm from a *hypothesis class*, a set of functions from \mathcal{X} to \mathcal{Y} , so as to minimize the error on the training set. This error is formalized as the *empirical risk*

$$\hat{\mathcal{R}}(h, \ell, \mathcal{D}) = \frac{1}{P} \sum_{k=1}^P \ell(\underline{y}_{(k)}, h(\underline{x}_{(k)})), \quad (1)$$

where the definition involves a loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ measuring differences in the output space. This learning objective nevertheless does not guarantee *generalization*, i.e. the ability of the predictor h to be accurate on inputs \underline{x} that are not in the training set. It is a surrogate for the ideal, but unavailable, *population risk*

$$\mathcal{R}(h, \ell) = \mathbb{E}_{\underline{x}, \underline{y}} [\ell(\underline{y}, h(\underline{x}))] = \int_{\mathcal{X}, \mathcal{Y}} d\underline{x} d\underline{y} p(\underline{x}, \underline{y}) \ell(\underline{y}, h(\underline{x})), \quad (2)$$

expressed as an expectation over the joint distribution $p(\underline{x}, \underline{y})$. The different choices of hypothesis classes and training algorithms yield the now crowded zoo of supervised learning algorithms.

Representation ability of deep neural networks In the context of supervised learning, deep neural networks enter the picture in the quality of a parametrized hypothesis class. Let us first quickly recall the simplest network, the *perceptron*, including only a single neuron. It is formalized as a function from \mathbb{R}^N to $\mathcal{Y} \subset \mathbb{R}$ applying an activation function f to a weighted sum of its inputs shifted by a bias $b \in \mathbb{R}$,

$$\hat{y} = h_{\underline{w}, b}(\underline{x}) = f(\underline{w}^\top \underline{x} + b) \quad (3)$$

where the weights are collected in the vector $\underline{w} \in \mathbb{R}^N$. From a practical standpoint, this very simple model can only solve the classification of linearly separable groups (see Figure 1). Yet

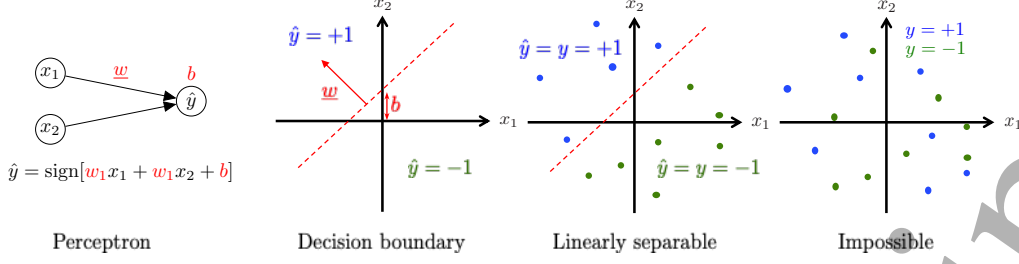


Figure 1: Let's assume we wish to classify data points $\underline{x} \in \mathbb{R}^2$ with labels $y = \pm 1$. We choose as an hypothesis class the perceptron sketched on the left with sign activation. For given weight vector \underline{w} and bias b the plane is divided by a decision boundary assigning labels. If the training data are linearly separable, then it is possible to perfectly predict all the labels with the perceptron, otherwise it is impossible.

from the point of view of learning theory, it has been the starting point of a rich statistical physics literature that will be discussed in Section 6.1.

Combining several neurons into networks defines more complex functions. The universal approximation theorem [Cyb89, Hor91] proves that the following two-layer network architecture can approximate any well-behaved function with a finite number of neurons,

$$\hat{y} = h_{\theta}(\underline{x}) = \underline{w}^{(2)\top} f(\underline{W}^{(1)} \underline{x} + \underline{b}) = \sum_{\alpha=1}^M w_{\alpha}^{(2)} f(\underline{w}_{\alpha}^{(1)\top} \underline{x} + b_{\alpha}), \quad \theta = \{\underline{w}^{(2)}, \underline{W}^{(1)}, \underline{b}\} \quad (4)$$

for f a bounded, non-constant, continuous scalar function, acting component-wise. In the language of deep learning this network has one hidden layer of M units. Input weight vectors $\underline{w}_{\alpha}^{(1)} \in \mathbb{R}^N$ are collected in a weight matrix $\underline{W}^{(1)} \in \mathbb{R}^{M \times N}$. Here, and in the following, the notation θ is used as short for the collection of adjustable parameters. The universal approximation theorem is a strong result in terms of representative power of neural networks but it is not constructive. It does not quantify the size of the network, i.e. the number M of hidden units, to approximate a given function, nor does it prescribe how to obtain the values of the parameters $\underline{w}^{(2)}$, $\underline{W}^{(1)}$ and \underline{b} for the optimal approximation. While building an approximation theory is still ongoing (see e.g. [GPEB19]). Practice, led by empirical considerations, has nevertheless demonstrated the efficiency of neural networks.

In applications, neural networks with multiple hidden layers, deep neural networks, are preferred. A generic neural network of depth L is the function

$$\hat{y} = h_{\theta}(\underline{x}) = f(\underline{W}^{(L)} f(\underline{W}^{(L-1)} \dots f(\underline{W}^{(1)} \underline{x} + \underline{b}^{(1)}) \dots + \underline{b}^{(L-1)}) + \underline{b}^{(L)}), \quad (5)$$

$$\theta = \{\underline{W}^{(l)} \in \mathbb{R}^{N_l \times N_{l-1}}, \underline{b}^{(l)} \in \mathbb{R}^{N_l}; l = 1 \dots L\}, \quad (6)$$

where $N_0 = N$ is the dimension of the input and $N_L = M$ is the dimension of the output. The architecture is fixed by specifying the number of neurons, or *width*, of the hidden layers $\{N_l\}_{l=1}^{L-1}$. The latter can be denoted $\underline{t}^{(l)} \in \mathbb{R}^{N_l}$ and follow the recursion

$$\underline{t}^{(1)} = f(\underline{W}^{(1)} \underline{x} + \underline{b}^{(1)}), \quad (7)$$

$$\underline{t}^{(l)} = f(\underline{W}^{(l)} \underline{t}^{(l-1)} + \underline{b}^{(l)}), \quad l = 2 \dots L-1, \quad (8)$$

$$\hat{y} = f(\underline{W}^{(L)} \underline{t}^{(L-1)} + \underline{b}^{(L)}). \quad (9)$$

Fixing the activation functions and the architecture of a neural network defines an hypothesis class. It is crucial that activations introduce non-linearities; the most common are the hyperbolic tangent \tanh and the rectified linear unit defined as $\text{relu}(x) = \max(0, x)$. Note that it is also possible to define stochastic neural networks by using noisy activation functions, uncommon in supervised learning applications except at training time so as to encourage generalization [PSDG14, SHK⁺14].

An originally proposed intuition for the advantage of depth is that it enables to treat the information in a hierarchical manner; either looking at different scales in different layers, or learning more and more abstract representations [BCV13]. Nevertheless, getting a clear theoretical understanding why in practice 'the deeper the better' is still an ongoing direction of research (see e.g. [Tel16, Dan17, SES19]).

Neural network training Given an architecture defining h_θ , the supervised learning objective is to minimize the empirical risk $\hat{\mathcal{R}}$ with respect to the parameters θ . This optimization problem lives in the dimension of the number of parameters which can range from tens to millions. The idea underlying the majority of training algorithms is to perform a gradient descent (GD) starting at parameters drawn randomly from an initialization distribution:

$$\theta_0 \sim p_{\theta_0}(\theta_0) \quad (10)$$

$$\theta_{t+1} \leftarrow \theta_t - \eta \nabla_{\theta} \hat{\mathcal{R}} = \theta_t - \eta \frac{1}{P} \sum_{k=1}^P \nabla_{\theta} \ell \left(y_{(k)}, h_{\theta_t}(x_{(k)}) \right). \quad (11)$$

The parameter η is the learning rate, controlling the size of the step in the direction of decreasing gradient per iteration. The computation of the gradients can be performed in time scaling linearly with depth by applying the derivative chain-rule leading to the *back-propagation* algorithm [GBC16]. A popular alternative to gradient descent is stochastic gradient descent (SGD) where the sum over the gradients for the entire training set is replaced by the sum over a small number of samples, randomly selected at each step [Rob51, Bot10].

During the training iterations, one typically monitors the *training error* (another name for the empirical risk given a training data set) and the *validation error*. The latter corresponds to the empirical risk computed on a set of points held-out from the training set, the validation set, to assess the generalization ability of the model either along the training or in order to select hyperparameters of training such as the value of the learning rate. A posteriori, the performance of the model is judged from the *generalization error*, which is evaluated on the never seen *test set*. While two different training algorithms (e.g. GD vs SGD) may achieve zero training error, they may differ in the level of generalization they typically reach.

Open questions and challenges Building on the fundamental concepts presented in the previous paragraphs, practitioners managed to bring deep learning to unanticipated performances in the automatic processing of images, speech and text (see [LBH15] for a few years old review). Still, many of the greatest successes in the field of neural network were obtained using ingenious tricks while many fundamental theoretical questions remain unresolved.

Regarding the optimization first, (S)GD training generally discovers parameters close to zero risk. Yet, gradient descent is guaranteed to converge to the neighborhood of a global minimum only for a convex function and is otherwise expected to get stuck in a local minimum. Therefore, the efficiency of gradient-based optimization is a priori a paradox given the empirical risk $\hat{\mathcal{R}}$ is non-convex in the parameters θ . Second, the generalization ability of deep neural networks trained by (S)GD is still poorly understood. The size of training data sets is limited by the cost of labelling by humans, experts or heavy computations. Thus training a deep and wide network amounts in practice to fitting a model of millions of degrees of freedom against a somehow relatively small amount of data points. Nevertheless it does not systematically lead to *overfitting*. The resulting neural networks can have surprisingly good predictions both on inputs seen during training and on new inputs [ZBH⁺17]. Results in the literature that relate the size and architecture of a network to a measure of its ability to generalize are too far from realistic settings to guide choices of practitioners. On the one hand, traditional bounds in statistics, considering worst cases, appear overly pessimistic [Vap00, BM02, SSBD14, AAKZ19]. On the other hand, historical statistical physics analyses of learning, briefly reviewed in Section 6.1, only concern simple architectures and synthetic data. This lack of theory results in potentially important waste: in terms of time lost by engineers in trial and error to optimize their solution, and in terms of electrical resources used to train and re-train possibly oversized networks while storing potentially unnecessarily large training data sets.

The success of deep learning, beyond these apparent theoretical puzzles, certainly lies in the interplay of advantageous properties of training algorithms, the neural network hypothesis class and structures in typical data (e.g. real images, conversations). Disentangling the role of the different ingredients is a very active line of research (see [VBGS17] for a review).

2.2 Unsupervised learning

Density estimation and generative modelling The goal of unsupervised learning is to directly extract structure from data. Compared to the supervised learning setting, the training data set is made of a set of example inputs $\mathcal{D} = \{x_{(k)}\}_{k=1}^P$ without corresponding

2 Machine learning with neural networks

outputs. A simple example of unsupervised learning is clustering, consisting in the discovery of unlabelled subgroups in the training data. Most unsupervised learning algorithms either implicitly or explicitly adopt a probabilistic viewpoint and implement *density estimation*. The idea is to approximate the true density $p(\underline{x})$ from which the training data was sampled by the closest (in various senses) element among a family of parametrized distributions over the input space $\{p_\theta(\cdot), \theta \in \mathbb{R}^{N_\theta}\}$. The selected p_θ is then a model of the data. If the model p_θ is easy to sample, it can be used to generate new inputs comparable to the training data points - which leads to the terminology of *generative models*. In this context, *unsupervised deep learning* exploits the representational power of deep neural networks to create sophisticated candidate p_θ .

A common formalization of the learning objective is to maximize the *likelihood*, defined as the probability of i.i.d. draws from the model p_θ to have generated the training data $\mathcal{D} = \{\underline{x}_{(k)}\}_{k=1}^P$, or equivalently its logarithm,

$$\max_{\theta} \prod_{k=1}^P p_\theta(\underline{x}_{(k)}) \iff \max_{\theta} \sum_{k=1}^P \log p_\theta(\underline{x}_{(k)}). \quad (12)$$

The second logarithmic additive formulation is generally preferred. It can be interpreted as the minimization of the Kullback-Leibler divergence between the empirical distribution $p_{\mathcal{D}}(\underline{x}) = \sum_{k=1}^P \delta(\underline{x} - \underline{x}_{(k)})/P$ and the model p_θ :

$$\min_{\theta} \text{KL}(p_{\mathcal{D}}||p_\theta) = \min_{\theta} \int d\underline{x} p_{\mathcal{D}}(\underline{x}) \log \frac{p_{\mathcal{D}}(\underline{x})}{p_\theta(\underline{x})} \iff \max_{\theta} \sum_{k=1}^P \log p_\theta(\underline{x}_{(k)}), \quad (13)$$

although considering the divergence with the discrete empirical measure is slightly abusive. The detail of the optimization algorithm here depends on the specification of p_θ . As we will see, the likelihood in itself is often intractable and learning consists in a gradient ascent on at best a lower bound, otherwise an approximation, of the likelihood.

A few years ago, an alternative strategy called adversarial training was introduced by [GPAM⁺14]. Here an additional trainable model called the discriminator, for instance parametrized by ϕ and denoted $d_\phi(\cdot)$, computes the probability for points in the input space \mathcal{X} of belonging to the training set \mathcal{D} rather than being generated by the model $p_\theta(\cdot)$. The parameters θ and ϕ are trained simultaneously such that, the generator learns to fool the discriminator and the discriminator learns not to be fooled by the generator. The optimization problem usually considered is

$$\min_{\theta} \max_{\phi} \mathbb{E}_{\mathcal{D}} [\log(d_\phi(\underline{x}))] + \mathbb{E}_{p_\theta} [\log(1 - d_\phi(\underline{x}))], \quad (14)$$

where the sum of the expected log-probabilities according to the discriminator for examples in \mathcal{D} to be drawn from \mathcal{D} and examples generated by the model not to be drawn from \mathcal{D} is maximized with respect to ϕ and minimized with respect to θ .

In the following, we present two classes of generative models based on neural networks.

Deep Generative Models A deep generative models defines a density p_θ obtained by propagating a simple distribution through a deep neural network. It can be formalized by introducing a latent variable $\underline{z} \in \mathbb{R}^N$ and a deep neural network h_θ similar to (5) of input dimension N . The generative process is then

$$\underline{z} \sim p_z(\underline{z}) \quad (15)$$

$$\underline{x} \sim p_\theta(\underline{x}|\underline{z}) = p_{\text{out}}(\underline{x}|h_\theta(\underline{z})), \quad (16)$$

where p_z is typically a factorized distribution on \mathbb{R}^N easy to sample (e.g. a standard normal distribution), and $p_{\text{out}}(\cdot|h_\theta(\underline{z}))$ is for instance a multivariate Gaussian distribution with mean and covariance that are functions of $h_\theta(\underline{z})$. The motivation to consider this class of models for joint distributions is three-fold. First the class is highly expressive. Second, it follows from the intuition that data sets leave on low dimensional manifolds, which here can be spanned by varying the latent representation \underline{z} usually much smaller than the input space dimension (for further intuition see also the reconstruction objective of the first autoencoders, see e.g. Chapter 14 of [GBC16]). Third, yet perhaps more importantly, the class can be optimized over easily using back-propagation, unlike the Restricted Boltzmann Machines presented in the next paragraph largely replaced by deep generative models. There are two main types of deep

generative models. Generative Adversarial Networks (GAN) [GPAM⁺14] trained following the adversarial objective mentioned above, and Variational AutoEncoders (VAE) [KW14, RMW14] trained to maximize a likelihood lower-bound.

Variational AutoEncoders The computation of the likelihood of one training sample $\underline{x}_{(k)}$ for a deep generative model (15)-(16) requires then the marginalization over the latent variable \underline{z} ,

$$p_{\theta}(\underline{x}) = \int d\underline{z} p_{\text{out}}(\underline{x}|\underline{h}_{\theta}(\underline{z}))p_z(\underline{z}). \quad (17)$$

This multidimensional integral cannot be performed analytically in the general case. It is also hard to evaluate numerically as it does not factorize over the dimensions of \underline{z} which are mixed by the neural network \underline{h}_{θ} . Yet a lower bound on the log-likelihood can be defined by introducing a tractable conditional distribution $q(\underline{z}|\underline{x})$ that will play the role of an approximation of the intractable *posterior* distribution $p_{\theta}(\underline{z}|\underline{x})$ implicitly defined by the model:

$$\log p_{\theta}(\underline{x}) \geq \int d\underline{z} q(\underline{z}|\underline{x}) [-\log q(\underline{z}|\underline{x}) + \log p_{\theta}(\underline{x}, \underline{z})] = \text{LB}(q, \theta, \underline{x}). \quad (18)$$

Maximum likelihood learning is then approached by the maximization of the lower bound $\text{LB}(q, \theta, \underline{x})$, which requires in practice to parametrize the tractable posterior $q = q_{\phi}$, typically with a neural network. Using the so-called re-parametrization trick [KW14, RMW14], the gradients of $\text{LB}(q_{\phi}, \theta, \underline{x})$ with respect to θ and ϕ can be approximated by a Monte Carlo, so that the likelihood lower bound can be optimized by gradient ascent.

Generative Adversarial Networks The principle of adversarial training was designed directly for a deep generative model [GPAM⁺14]. Using a deep neural network to parametrize the discriminator $d_{\phi}(\cdot)$ as well as the generator $p_{\theta}(\cdot)$, it leads to a remarkable quality of produced samples and is now one of the most studied generative model.

Restricted Boltzmann Machines Models described in the preceding paragraphs comprised only *feed forward* neural networks. In feed forward neural networks, the state or value of successive layers is determined following the recursion (7)-(9), in one pass from inputs to outputs. Boltzmann machines instead involve *undirected* neural networks which consist of stochastic neurons with symmetric interactions. The probability law associated with a neuron state is a function of neighboring neurons, themselves reciprocally function of the first neuron. Sampling a configuration therefore requires an equilibration in the place of a simple forward pass.

A Restricted Boltzmann Machine (RBM) [AHS85, Smo86] with M hidden neurons in practice defines a joint distribution over an input (or visible) layer $\underline{x} \in \{0, 1\}^N$ and a hidden layer $\underline{t} \in \{0, 1\}^M$,

$$p_{\theta}(\underline{x}, \underline{t}) = \frac{1}{\mathcal{Z}} e^{a^{\top} \underline{x} + b^{\top} \underline{t} + \underline{x}^{\top} \underline{W} \underline{t}}, \quad \theta = \{\underline{W}, a, b\}, \quad (19)$$

where \mathcal{Z} is the normalization factor, similar to the partition function of statistical physics. The parametric density model over inputs is then the marginal $p_{\theta}(\underline{x}) = \sum_{\underline{t} \in \{0, 1\}^M} p_{\theta}(\underline{x}, \underline{t})$. Although seemingly very similar to pairwise Ising models, the introduction of hidden units provides a greater representative power to RBMs as hidden units can mediate interactions between arbitrary groups of input units. Furthermore, they can be generalized to Deep Boltzmann Machines (DBMs) [SH09], where several hidden layers are stacked on top of each other.

Identically to VAEs, RBMs can represent sophisticated distributions at the cost of an intractable likelihood. Indeed the summation over 2^{M+N} terms in the partition function cannot be simplified by an analytical trick and is only realistically doable for small models. RBMs are commonly trained through a gradient ascent of the likelihood using approximated gradients. As exact Monte Carlo evaluation is a costly operation that would need to be repeated at each parameter update in the gradient ascent, several more or less sophisticated approximations are preferred: contrastive divergence (CD) [Hin02], its persistent variant (PCD) [Tie08] or even parallel tempering [DCB⁺10, CRI10].

RBMs were the first effective generative models using neural networks. They found applications in various domains including dimensionality reduction [HS06], classification [LB08], collaborative filtering [SMH07], feature learning [CNL11], and topic modeling [HS09]. Used for an unsupervised pre-training of deep neural networks layer by layer [HOT06, BLPL07], they also played a crucial role in the take-off of supervised deep learning.

3 Statistical inference and the statistical physics approach

Open questions and challenges Generative models involving neural networks such as VAE, GANs and RBMs have great expressive powers at the cost of not being amenable to exact treatment. Their training, and sometimes even their sampling requires approximations. From a practical standpoint, whether these approximations can be either made more accurate or less costly is an open direction of research. Another important related question is the evaluation of the performance of generative models [SBL⁺18]. To start with the objective function of training is very often itself intractable (e.g. the likelihood of a VAE or a RBM), and beyond this objective, the unsupervised setting does not define a priori a test task for the performance of the model. Additionally, unsupervised deep learning inherits some of the theoretical puzzles already discussed in the supervised learning section. In particular, assessing the difficulty to represent a distribution and select a sufficient minimal model and/or training data set is an ongoing effort of research.

3 Statistical inference and the statistical physics approach

To tackle the open questions and challenges surrounding neural networks mentioned in the previous Section, we need to manipulate high-dimensional probability distributions. The generic concept of statistical inference refers to the extraction of useful information from these complicated objects. Statistical physics, with its probabilistic interpretation of natural systems composed of many elementary components, is naturally interested in similar questions. We provide in this section a few concrete examples of inference questions arising in neural networks and explicit how statistical physics enters the picture. In particular, the theory of disordered systems appears here especially relevant.

3.1 Statistical inference

3.1.1 Some inference questions in neural networks for machine learning

Inference in generative models Generative models used for unsupervised learning are statistical models defining high-dimensional distributions with complex dependencies. As we have seen in Section 2.2, the most common training objective in unsupervised learning is the maximization of the log-likelihood, i.e. the log of the probability assigned by the generative model to the training set $\{\underline{x}_{(k)}\}_{k=1}^P$. Computing the probability of observing a given sample $\underline{x}_{(k)}$ is an inference question. It requires to marginalize over all the hidden representations of the problem. For instance in the RBM (19),

$$p_{\theta}(\underline{x}_{(k)}) = \frac{1}{Z} \sum_{\underline{t} \in \{0,1\}^M} e^{a^{\top} \underline{x}_{(k)} + b^{\top} \underline{t} + \underline{x}_{(k)}^{\top} W \underline{t}}. \quad (20)$$

While the numerator will be easy to evaluate, the partition function has no analytical expression and its exact evaluation requires to sum over all possible states of the network.

Learning as statistical inference: Bayesian inference and the teacher-student scenario The practical problem of training neural networks from data as introduced in Section 2 is not in general interpreted as inference. To do so, one needs to treat the learnable parameters as random variables, which is the case in Bayesian learning. For instance in supervised learning, an underlying prior distribution $p_{\theta}(\theta)$ for the weights and biases of a neural network (5)-(6) is assumed, so that Bayes rule defines a posterior distribution given the training data \mathcal{D} ,

$$p(\theta|\mathcal{D}) = \frac{p(\mathcal{D}|\theta)p_{\theta}(\theta)}{p(\mathcal{D})}. \quad (21)$$

Compared to the single output of risk minimization, we obtain an entire distribution for the learned parameters θ , which takes into account not only the training data but also some knowledge on the structure of the parameters (e.g. sparsity) through the prior. In practice, Bayesian learning and traditional empirical risk minimization may not be so different. On the one hand, the Bayesian posterior distribution is often summarized by a point estimate such as its maximum. On the other hand risk minimization is often biased towards desired properties of the weights through regularization techniques (e.g. promoting small norm) recalling the role of the Bayesian prior.

3.2 Statistical physics of disordered systems, first appearance on stage

However, from a theoretical point of view, Bayesian learning is of particular interest in the *teacher-student* scenario. The idea here is to consider a toy model of the learning problem where parameters are effectively drawn from a prior distribution. Let us use as an illustration the case of the supervised learning of the perceptron model (3). We draw a weight vector \underline{w}_0 , from a prior distribution $p_w(\cdot)$, along with a set of P inputs $\{\underline{x}_{(k)}\}_{k=1}^P$ i.i.d from a data distribution $p_x(\cdot)$. Using this *teacher* perceptron model we also draw a set of possibly noisy corresponding outputs $y_{(k)}$ from a teacher conditional probability $p(\cdot|\underline{w}_0^\top \underline{x}_{(k)})$. From the training set of the P pairs $\mathcal{D} = \{\underline{x}_{(k)}, y_{(k)}\}$, one can attempt to rediscover the teacher rule by training a *student* perceptron model. The problem can equivalently be phrased as a reconstruction inference question: can we recover the value of \underline{w}_0 from the observations in \mathcal{D} ? The Bayesian framework yields a posterior distribution of solutions

$$p(\underline{w}|\mathcal{D}) = \prod_{k=1}^P p(y_{(k)}|\underline{w}^\top \underline{x}_{(k)}) p_w(\underline{w}) / \prod_{k=1}^P p(y_{(k)}|\underline{x}_{(k)}). \quad (22)$$

Note that the terminology of teacher-student applies for a generic inference problem of reconstruction: the statistical model used to generate the data along with the realization of the unknown signal is called the *teacher*; the statistical model assumed to perform the reconstruction of the signal is called the *student*. When the two models are identical or matched, the inference is *Bayes optimal*. When the teacher model is not perfectly known, the statistical models can also be different (from slightly differing prior distributions to entirely different models), in which case they are said to be mismatched, and the reconstruction is suboptimal.

Of course in practical machine learning applications of neural networks, one has only access to an empirical distribution of the data and it is unclear whether there should exist a formal rule underlying the input-output mapping. Yet the teacher-student setting is a modelling strategy of learning which offers interesting possibilities of analysis and we shall refer to numerous works resorting to the setup in Section 6.

3.1.2 Answering inference questions

Many inference questions in the line of the ones mentioned in the previous Section have no tractable exact solution. When there exists no analytical closed-form, computations of averages and marginals require summing over configurations. Their number typically scales exponentially with the size of the system, then becoming astronomically large for high-dimensional models. Hence it is necessary to design approximate inference strategies. They may require an algorithmic implementation but must run in finite (typically polynomial) time. An important cross-fertilization between statistical physics and information sciences have taken place over the past decades around the questions of inference. Two major classes of such algorithms are Monte Carlo Markov Chains (MCMC), and mean-field methods. The former is nicely reviewed in the context of statistical physics in [Kra06]. The latter will be the focus of this short review, in the context of deep learning.

Note that representations of joint probability distributions through probabilistic graphical models and factor graphs are crucial tools to design efficient inference strategies. In Appendix B, we quickly introduce for the unfamiliar reader these two formalisms that enable to encode and exploit independencies between random variables. As examples, Figure 2 presents graphical representations of the RBM measure (19) and the posterior distribution in the Bayesian learning of the perceptron as discussed in the previous Section.

3.2 Statistical physics of disordered systems, first appearance on stage

Here we re-introduce briefly fundamental concepts of statistical physics that will help to understand connections with inference and the origin of the methods presented in what follows.

The thermodynamic limit The equilibrium statistics of classical physical systems are described by the Boltzmann distribution. For a system with N degrees of freedom noted $\underline{x} \in \mathcal{X}^N$ and an energy functional $E(\underline{x})$, we have

$$p(\underline{x}) = \frac{e^{-\beta E(\underline{x})}}{\mathcal{Z}_N}, \quad \mathcal{Z}_N = \sum_{\underline{x} \in \mathcal{X}^N} e^{-\beta E(\underline{x})}, \quad \beta = 1/k_B T, \quad (23)$$

3 Statistical inference and the statistical physics approach

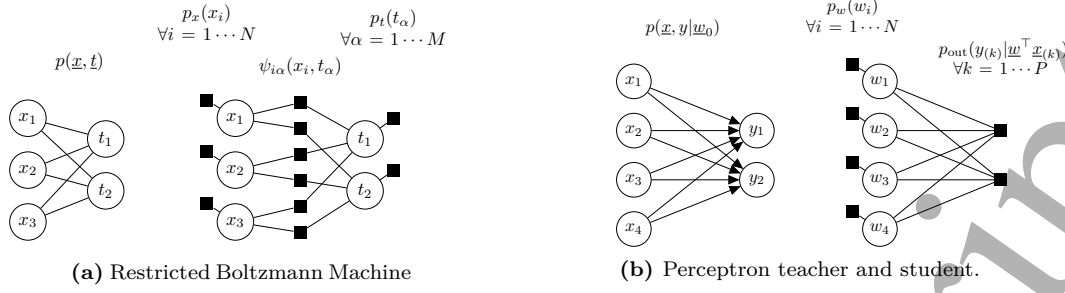


Figure 2: (a) Undirected probabilistic graphical model (left) and factor graph representation (right). (b) Left: Directed graphical model of the generative model for the training data knowing the teacher weight vector \underline{w}_0 . Right: Factor graph representation of the posterior distribution for the student $p(\underline{w} | \underline{X}, \underline{y})$, where the vector $\underline{y} \in \mathbb{R}^P$ gathers the outputs $y(k)$ and the matrix $\underline{X} \in \mathbb{R}^{N \times P}$ gathers the inputs $\underline{x}(k)$.

where we defined the partition function \mathcal{Z}_N and the inverse temperature β . To characterize the macroscopic state of the system, an important functional is the free energy

$$F_N = -\log \mathcal{Z}_N / \beta = -\frac{1}{\beta} \log \sum_{\underline{x} \in \mathcal{X}^N} e^{-\beta E(\underline{x})}. \quad (24)$$

While the number of available configurations \mathcal{X}^N grows exponentially with N , considering the *thermodynamic* limit $N \rightarrow \infty$ typically simplifies computations due to concentrations. Let $e_N = E/N$ be the energy per degree of freedom, the partition function can be re-written as a sum over the configurations of a given energy e_N

$$\mathcal{Z}_N = \sum_{e_N} e^{-N\beta f_N(e_N)}, \quad (25)$$

where we define $f_N(e_N)$ the free energy density of states of energy e_N . This rewriting implies that at large N the states of energy minimizing the free energy are exponentially more likely than any other states. Provided the following limits exist, the statistics of the system are dominated by the former states and we have the thermodynamic quantities

$$\mathcal{Z} = \lim_{N \rightarrow \infty} \mathcal{Z}_N = e^{-\beta f}, \text{ and } f = \lim_{N \rightarrow \infty} F_N / N. \quad (26)$$

The interested reader will also find a more detailed yet friendly presentation of the thermodynamic limit in Section 2.4 of [MM09].

Disordered systems Remarkably, the statistical physics framework can be applied to inhomogeneous systems with *quenched disorder*. In these systems, interactions are functions of the realization of some random variables. An iconic example is the Sherrington-Kirkpatrick (SK) model [SK75], a fully connected Ising model with random Gaussian couplings $\underline{J} = (J_{ij})$, that is where the J_{ij} are drawn independently from a Gaussian distribution. As a result, the energy functional of disordered systems is itself function of the random variables. For instance here, the energy of a spin configuration \underline{x} is then $E(\underline{x}; \underline{J}) = -\frac{1}{2} \underline{x}^\top \underline{J} \underline{x}$. In principle, system properties depend on a given realization of the disorder. In our example, the correlation between two spins $\langle x_i x_j \rangle_J$ certainly does. Yet some aggregated properties are expected to be *self-averaging* in the thermodynamic limit, meaning that they concentrate on their mean with respect to the disorder as the fluctuations are averaged out. It is the case for the free energy. As a result, here it formally verifies:

$$\lim_{N \rightarrow \infty} F_{N; \underline{J}} / N = \lim_{N \rightarrow \infty} \mathbb{E}_{\underline{J}} [F_{N; \underline{J}} / N] = f. \quad (27)$$

(see e.g. [MPV86, CCFM05] for discussions of self-averaging in spin glasses). Thus the typical behavior of complex systems is studied in the statistical physics framework by taking two important conceptual steps: averaging over the realizations of the disorder and considering the thermodynamic limit. These are starting points to design approximate inference methods. Before turning to an introduction to mean-field approximations, we stress the originality of the statistical physics approach to inference.

Statistical physics of inference problems Statistical inference questions are mapped to statistical physics systems by interpreting general joint probability distributions as Boltzmann distributions (23). Turning back to our simple examples of Section 3.1, the RBM is trivially mapped as it directly borrows its definition from statistical physics. We have

$$E(\underline{x}, \underline{t}; \underline{W}) = -\underline{a}^\top \underline{x} - \underline{b}^\top \underline{t} - \underline{x}^\top \underline{W} \underline{t}. \quad (28)$$

The inverse temperature parameter can either be considered equal to 1 or as a scaling factor of the weight matrix $\underline{W} \leftarrow \beta \underline{W}$ and bias vectors $\underline{a} \leftarrow \beta \underline{a}$ and $\underline{b} \leftarrow \beta \underline{b}$. The RBM parameters play the role of the disorder. Here the computational hardness in estimating the log-likelihood comes from the estimation of the log-partition function, which is precisely the free energy. In our second example, the estimation of the student perceptron weight vector, the posterior distribution is mapped to a Boltzmann distribution by setting

$$E(\underline{w}; \underline{y}, \underline{X}) = -\log p(\underline{y} | \underline{w}^\top \underline{X}) p_w(\underline{w}). \quad (29)$$

The disorder is here materialized by the training data. The difficulty is here to compute $p(\underline{y} | \underline{X})$ which is again the partition function in the Boltzmann distribution mapping. Relying on the thermodynamic limit, mean-field methods will provide asymptotic results. Nevertheless, experience shows that the behavior of large finite-size systems are often well explained by the infinite-size limits.

Also, the application of mean-field inference requires assumptions about the distribution of the disorder which is averaged over. Practical algorithms for arbitrary cases can be derived with ad-hoc assumptions, but studying a precise toy statistical model can also bring interesting insights. The simplest model in most cases is to consider uncorrelated disorder: in the example of the perceptron this corresponds to random input data points with arbitrary random labels. Yet, the teacher-student scenario offers many advantages with little more difficulty. It allows to create data sets with structure (the underlying teacher rule). It also allows to formalize an analysis of the difficulty of a learning problem and of the performance in the resolution. Intuitively, the definition of a ground-truth teacher rule with a fixed number of degrees of freedom sets the minimum information necessary to extract from the observations, or training data, in order to achieve perfect reconstruction. This is an *information-theoretic limit*.

Furthermore, the assumption of an underlying statistical model enables the measurement of performance of different learning algorithms over the class of corresponding problems from an average viewpoint. This is in contrast with the traditional approach of computer science in studying the difficulty of a class of problem based on the *worst* case. This conservative strategy yields strong guarantees of success, yet it may be overly pessimistic compared to the experience of practitioners. Considering a distribution over the possible problems (a.k.a different realizations of the disorder), the average performances are sometimes more informative of *typical* instances rather than worst ones. For deep learning, this approach may prove particularly interesting as the traditional bounds, based on the VC-dimension [Vap00] and Rademacher complexity [BM02, SSBD14], appear extremely loose when compared to practical examples.

Finally, we must emphasize that derivations presented here are not mathematically rigorous. They are based on ‘correct’ assumptions allowing to push further the understanding of the problems at hand, while a formal proof of the assumptions is possibly much harder to obtain.

4 Selected overview of mean-field treatments: Free energies and algorithms

Mean-field methods are a set of techniques enabling to approximate marginalized quantities of joint probability distributions by exploiting knowledge on the dependencies between random variables. They are usually said to be analytical - as opposed to numerical Monte Carlo methods. In practice they usually replace a summation exponentially large in the size of the system by an analytical formula involving a set of parameters, themselves solution of a closed set of non-linear equations. Finding the values of these parameters typically requires only a polynomial number of operations.

In this Section, we will give a selected overview of mean-field methods as they were introduced in the statistical physics and/or signal processing literature. A key take away of what follows is that closely related results can be obtained from different heuristics of derivation. We will start by deriving the simplest and historically first mean-field method. We will then

4 Selected overview of mean-field treatments: free energies and algorithms

introduce the important broad techniques that are high-temperature expansions, message-passing algorithms and the replica method. In the following Section 5 we will additionally cover the most recent extensions of mean-field methods presented in the present Section 4 that are relevant to study learning problems.

4.1 Naive mean-field

The naive mean-field method is the first and somehow simplest mean-field approximation. It was introduced by the physicists Curie [Cur95] and Weiss [Wei07] and then adopted by the different communities interested in inference [WJ08].

4.1.1 Variational derivation

The naive mean-field method consists in approximating the joint probability distribution of interest by a fully factorized distribution. Therefore, it ignores correlations between random variables. Among multiple methods of derivation, we present here the variational method: it is the best known method across fields and it readily shows that, for any joint probability distribution interpreted as a Boltzmann distribution, the rather crude naive mean-field approximation yields an upper bound on the free energy. For the purpose of demonstration we consider a Boltzmann machine without hidden units (Ising model) with variables (spins) $\underline{x} = (x_1, \dots, x_N) \in \mathcal{X} = \{0, 1\}^N$, and energy function

$$E(\underline{x}) = - \sum_{i=1}^N b_i x_i - \sum_{(ij)} W_{ij} x_i x_j = -\underline{b}^\top \underline{x} - \frac{1}{2} \underline{x}^\top \underline{W} \underline{x}, \quad \underline{b} \in \mathbb{R}^N, \quad \underline{W} \in \mathbb{R}^{N \times N}, \quad (30)$$

where the notation (ij) stands for pairs of connected spin-variables, and the weight matrix \underline{W} is symmetric. The choices for $\{0, 1\}$ rather than $\{-1, +1\}$ for the variable values, the notations \underline{W} for weights (instead of couplings), \underline{b} for biases (instead of local fields), as well as the vector notation, are leaning towards the machine learning conventions. We denote by $q_{\underline{m}}$ a fully factorized distribution on $\{0, 1\}^N$, which is a multivariate Bernoulli distribution parametrized by the mean values $\underline{m} = (m_1, \dots, m_N) \in [0, 1]^N$ of the marginals (denoted by q_{m_i}):

$$q_{\underline{m}}(\underline{x}) = \prod_{i=1}^N q_{m_i}(x_i) = \prod_{i=1}^N m_i \delta(x_i - 1) + (1 - m_i) \delta(x_i). \quad (31)$$

We look for the optimal $q_{\underline{m}}$ distribution to approximate the Boltzmann distribution $p(\underline{x}) = e^{-\beta E(\underline{x})} / \mathcal{Z}$ by minimizing the KL-divergence

$$\min_{\underline{m}} \text{KL}(q_{\underline{m}} \| p) = \min_{\underline{m}} \sum_{\underline{x} \in \mathcal{X}} q_{\underline{m}}(\underline{x}) \log \frac{q_{\underline{m}}(\underline{x})}{p(\underline{x})} \quad (32)$$

$$= \min_{\underline{m}} \sum_{\underline{x} \in \mathcal{X}} q_{\underline{m}}(\underline{x}) \log q_{\underline{m}}(\underline{x}) + \beta \sum_{\underline{x} \in \mathcal{X}} q_{\underline{m}}(\underline{x}) E(\underline{x}) + \log \mathcal{Z} \quad (33)$$

$$= \min_{\underline{m}} \beta G(q_{\underline{m}}) - \beta F \geq 0, \quad (34)$$

where the last inequality comes from the positivity of the KL-divergence. For a generic distribution q , $G(q)$ is the *Gibbs free energy* for the energy $E(\underline{x})$,

$$G(q) = \sum_{\underline{x} \in \mathcal{X}} q(\underline{x}) E(\underline{x}) + \frac{1}{\beta} \sum_{\underline{x} \in \mathcal{X}} q(\underline{x}) \log q(\underline{x}) = U(q) - H(q)/\beta \geq F, \quad (35)$$

involving the average energy $U(q)$ and the entropy $H(q)$. It is greater than the true free energy F except when $q = p$, in which case they are equal. Note that this fact also means that the Boltzmann distribution minimizes the Gibbs free energy. Restricting to factorized $q_{\underline{m}}$ distributions, we obtain the naive mean-field approximations for the mean value of the variables (or *magnetizations*) and the free energy:

$$\underline{m}^* = \arg \min_{\underline{m}} G(q_{\underline{m}}) = \langle \underline{x} \rangle_{q_{\underline{m}^*}}, \quad (36)$$

$$F_{\text{NMF}} = G(q_{\underline{m}^*}) \geq F. \quad (37)$$

4.2 Thouless Anderson and Palmer equations

The choice of a very simple family of distributions $q_{\underline{m}}$ limits the quality of the approximation but allows for tractable computations of observables, for instance the two-spins correlations $\langle x_i x_j \rangle_{q^*} = m_i^* m_j^*$ or variance of one spin $\langle x_i^2 \rangle_{q^*} - \langle x_i \rangle_{q^*}^2 = m_i^* - m_i^{*2}$.

In our example of the Boltzmann machine, it is easy to compute the Gibbs free energy for the factorized ansatz, we define functions of the magnetization vector:

$$U_{\text{NMF}}(\underline{m}) = \langle E(\underline{x}) \rangle_{q_{\underline{m}}} = -\underline{b}^\top \underline{m} - \frac{1}{2} \underline{m}^\top \underline{W} \underline{m}, \quad (38)$$

$$H_{\text{NMF}}(\underline{m}) = -\langle \log q_{\underline{m}}(\underline{x}) \rangle_{q_{\underline{m}}} = -\sum_{i=1}^N m_i \log m_i + (1 - m_i) \log(1 - m_i), \quad (39)$$

$$G_{\text{NMF}}(\underline{m}) = G(q_{\underline{m}}) = U_{\text{NMF}}(\underline{m}) - H_{\text{NMF}}(\underline{m})/\beta. \quad (40)$$

Looking for stationary points we find a closed set of non linear equations for the m_i^* ,

$$\left. \frac{\partial G_{\text{NMF}}}{\partial m_i} \right|_{\underline{m}^*} = 0 \quad \Rightarrow \quad m_i^* = \sigma(\beta b_i + \sum_{j \in \partial i} \beta W_{ij} m_j^*) \quad \forall i = 1 \cdots N, \quad (41)$$

where $\sigma(x) = (1 + e^{-x})^{-1}$. The solutions can be computed by iterating these relations from a random initialization until a fixed point is reached.

To understand the implication of the restriction to factorized distributions, it is instructive to compare this naive mean-field equation with the exact identity

$$\langle x_i \rangle_p = \langle \sigma(\beta b_i + \sum_{j \in \partial i} \beta W_{ij} x_j) \rangle_p, \quad (42)$$

derived in a few lines in Appendix C. Under the Boltzmann distribution $p(\underline{x}) = e^{-\beta E(\underline{x})}/\mathcal{Z}$, these averages are difficult to compute. The naive mean-field method is neglecting the fluctuations of the effective field felt by the variable x_i : $\sum_{j \in \partial i} W_{ij} x_j$, keeping only its mean $\sum_{j \in \partial i} W_{ij} m_j$. This incidentally justifies the name of mean-field methods.

4.1.2 When does naive mean-field hold true?

The previous derivation shows that the naive mean-field approximation allows to bound the free energy. While this bound is expected to be rough in general, the approximation is reliable when the fluctuations of the local effective fields $\sum_{j \in \partial i} W_{ij} x_j$ are small. This may happen in particular in the thermodynamic limit $N \rightarrow \infty$ in *infinite range* models, that is when weights or couplings are not only local but distributed in the entire system, or if each variable interacts directly with a non-vanishing fraction of the whole set of variables (e.g. [OS01] Section 2). The influence on one given variable of the rest of the system can then be treated as an average background. Provided the couplings are weak enough, the naive mean-field method may even become asymptotically exact. This is the case of the *Curie-Weiss* model, which is the fully connected version of the model (30) with all $W_{ij} = 1/N$ (see e.g. Section 2.5.2 of [MM09]). The sum of weakly dependent variables then concentrates on its mean by the central limit theorem. We stress that it means that for finite dimensional models (more representative of a physical system, where for instance variables are assumed to be attached to the vertices of a lattice with nearest neighbors interactions), mean-field methods are expected to be quite poor. By contrast, infinite range models (interpreted as infinite-dimensional models by physicists) are thus traditionally called *mean-field models*.

In the next Section we will recover the naive mean-field equations through a different method. The following derivation will also allow to compute corrections to the rather crude approximation we just discussed by taking into account some of the correlations it neglects.

4.2 Thouless Anderson and Palmer equations

The TAP mean-field equations [TAP77, MH76] were originally derived as an exact mean-field theory for the Sherrington-Kirkpatrick (SK) model [SK75]. The emblematic *spin glass* SK model we already mentioned corresponds to a fully connected Ising model with energy (30) and disordered couplings W_{ij} drawn independently from a Gaussian distribution with zero mean and variance W_0/N . The derivation of [TAP77] followed from arguments specific to the SK model. Later, it was shown that the same approximation could be recovered from a second order Taylor expansion at high temperature by Plefka [Ple82] and that it could be

4 Selected overview of mean-field treatments: free energies and algorithms

further corrected by the systematic computation of higher orders by Georges and Yedidia [GY99]. We will briefly present this last derivation, having again in mind the example of the generic Boltzmann machine (30).

4.2.1 Outline of the derivation

Going back to the variational formulation (34), we shall now perform a minimization in two steps. Consider first the family of distributions $q_{\underline{m}}$ enforcing $\langle \underline{x} \rangle_{q_{\underline{m}}} = \underline{m}$ for a fixed vector of magnetizations \underline{m} , but without any factorization constraint. The corresponding Gibbs free energy is

$$G(q_{\underline{m}}) = U(q_{\underline{m}}) - H(q_{\underline{m}})/\beta. \quad (43)$$

A first minimization at fixed \underline{m} over the $q_{\underline{m}}$ defines another auxiliary free energy

$$G_{\text{TAP}}(\underline{m}) = \min_{q_{\underline{m}}} G(q_{\underline{m}}). \quad (44)$$

A second minimization over \underline{m} would recover the overall unconstrained minimum of the variational problem (34) which is the exact free energy

$$F = -\log \mathcal{Z}/\beta = \min_{\underline{m}} G_{\text{TAP}}(\underline{m}). \quad (45)$$

Yet the actual value of $G_{\text{TAP}}(\underline{m})$ turns out as complicated to compute as F itself. Fortunately, $\beta G_{\text{TAP}}(\underline{m})$ can be easily approximated by a Taylor expansion around $\beta = 0$ due to interactions vanishing at high temperature, as noticed by Pfleka, Georges and Yedidia [Ple82, GY99]. After expanding, the minimization over $G_{\text{TAP}}(\underline{m})$ yields a set of self consistent equations on the magnetizations \underline{m} , called the *TAP equations*, reminiscent of the naive mean-field equations (41). Here again, the consistency equations are typically solved by iterations. Plugging the solutions \underline{m}^* back into the expanded expression yields the *TAP free energy* $F_{\text{TAP}} = G_{\text{TAP}}(\underline{m}^*)$. Note that ultimately the approximation lies in the truncation of the expansion. At first order the naive mean-field approximation is recovered. Historically, the expansion was first stopped at the second order. This choice was model dependent, it results from the fact that the mean-field theory is already exact at the second order for the SK model [MH76, TAP77, Ple82].

4.2.2 Illustration on binary Boltzmann machines and important remarks

For the Boltzmann machine (30), the TAP equations and TAP free energy (truncated at second order) are [TAP77],

$$m_i^* = \sigma \left(\beta b_i + \sum_{j \in \partial i} \beta W_{ij} m_j^* - \beta^2 W_{ij}^2 (m_j^* - \frac{1}{2})(m_i^* - m_i^{*2}) \right) \quad \forall i \quad (46)$$

$$\begin{aligned} \beta G_{\text{TAP}}(\underline{m}^*) = & -H_{\text{NMF}}(\underline{m}^*) - \beta \sum_{i=1}^N b_i m_i^* - \beta \sum_{(ij)} m_i^* W_{ij} m_j^* \\ & - \frac{\beta^2}{2} \sum_{(ij)} W_{ij}^2 (m_i^* - m_i^{*2})(m_j^* - m_j^{*2}), \end{aligned} \quad (47)$$

where the naive mean-field entropy H_{NMF} was defined in (39). For this model, albeit with $\{+1, -1\}$ variables instead of $\{0, 1\}$, several references pedagogically present the details of the derivation sketched in the previous paragraph. The interested reader should check in particular [OS01, Zam10]. We also present a more general derivation in Appendix D, see Section 4.2.3.

Onsager reaction term Compared to the naive mean-field approximation the TAP equations include a correction to the effective field called the *Onsager reaction term*. The idea is that, in the effective field at variable i , we should consider corrected magnetizations of neighboring spins $j \in \partial i$, that would correspond to the absence of variable i . This intuition echoes at two other derivations of the TAP approximation: the cavity method [MPV86] that will not be covered here and the message passing which will be discussed in the next Section.

As far as the SK model is concerned, this second order correction is enough in the thermodynamic limit as the statistics of the weights imply that higher orders will typically be

4.2 Thouless Anderson and Palmer equations

subleading. Yet in general, the correct TAP equations for a given model will depend on the statistics of interactions and there is no guarantee that there exists a finite order of truncation leading to an exact mean-field theory. In Section 5.2 we will discuss models beyond SK where a conjectured exact TAP approximation can be derived.

Single instance Although the selection of the correct TAP approximation relies on the statistics of the weights, the derivation of the expansion outlined above does not require to average over them, i.e. it does not require an average over the disorder. Consequently, the approximation method is well defined for a single instance of the random disordered model and the TAP free energy and magnetizations can be computed for a given (realization of the) set of weights $\{W_{ij}\}_{(ij)}$ as explained in the following paragraph. In other words, it means that the approximation can be used to design practical inference algorithms in finite-sized problems and not only for theoretical predictions on average over the disordered class of models. Crucially, these algorithms may provide approximations of disorder-dependent observables, such as correlations, and not only of self averaging quantities.

Finding solutions The self-consistent equations on the magnetizations (46) are usually solved by turning them into an iteration scheme and looking for fixed points. This generic recipe leaves nonetheless room for interpretation: which exact form should be iterated? How should the updates for the different equations be scheduled? Which time indexing should be used? While the following scheme may seem natural

$$m_i^{(t+1)} \leftarrow \sigma \left(\beta b_i + \sum_{j \in \partial i} \beta W_{ij} m_j^{(t)} - W_{ij}^2 \left(m_j^{(t)} - \frac{1}{2} \right) \left(m_i^{(t)} - m_i^{(t-1)} \right)^2 \right), \quad (48)$$

it typically has more convergence issues than the following alternative scheme including the time index $t-1$

$$m_i^{(t+1)} \leftarrow \sigma \left(\beta b_i + \sum_{j \in \partial i} \beta W_{ij} m_j^{(t)} - W_{ij}^2 \left(m_j^{(t)} - \frac{1}{2} \right) \left(m_i^{(t-1)} - m_i^{(t-2)} \right)^2 \right). \quad (49)$$

This issue was discussed in particular in [Kab03, Bol14]. Remarkably, this last scheme, or algorithm, is actually the one obtained by the approximate message passing derivation that will be discussed in the upcoming Section 4.3.

Solutions of the TAP equations The TAP equations can admit multiple solutions with either equal or different TAP free energy. While the true free energy F corresponds to the minimum of the Gibbs free energy, reached for the Boltzmann distribution, the TAP derivation consists in performing an effectively unconstrained minimization in two steps, but with an approximation through a Taylor expansion in between. The truncation of the expansion therefore breaks the correspondence between the discovered minimizer and the unique Boltzmann distribution, hence the possible multiplicity of solutions. For the SK model for instance, the number of solutions of the TAP equations increases rapidly as β grows [MPV86]. While the different solutions can be accessed using different initializations of the iterative scheme, it is notably hard in phases where they are numerous to find exhaustively all the TAP solutions. In theory, they should be weighted according to their free energy density and averaged to recover the thermodynamics predicted by the replica computation [DY83], another mean-field approximation discussed in Section 4.4.

4.2.3 Generalizing the Georges-Yedidia expansion

In the derivation outlined above for binary variables, $x_i = 0$ or 1 , the mean of each variable m_i was fixed. This is enough to parametrize the corresponding marginal distribution $q_{m_i}(x_i)$. Yet the expansion can actually be generalized to Potts variables (taking multiple discrete values) or even real valued variables by introducing appropriate parameters for the marginals. A general derivation fixing arbitrary real valued marginal distribution was proposed in Appendix B of [LKZ17] for the problem of low rank matrix factorization. Alternatively, another level of approximation can be introduced for real valued variables by restricting the set of marginal distributions tested to a parametrized family of distributions. By choosing a Gaussian parametrization, one recovers TAP equations equivalent to the approximate message

4 Selected overview of mean-field treatments: free energies and algorithms

passing algorithm that will be discussed in the next Section. In Appendix D, we present a derivation for real-valued Boltzmann machines with a Gaussian parametrization as proposed in [TGM⁺18].

4.3 Belief propagation and approximate message passing

Another route to rediscover the TAP equations is through the approximation of message passing algorithms. Variations of the latter were discovered multiple times in different fields. In physics they were written in a restricted version as soon as 1935 by Bethe [Bet35]. In statistics, they were developed by Pearl as methods for probabilistic inference [Pea88]. In this section we will start by introducing a case-study of interest, the Generalized Linear Model. We will then proceed by steps to outline the derivation of the Approximate Message Passing (AMP) algorithm from the Belief Propagation (BP) equations.

4.3.1 Generalized linear model

Definition We introduce the *Generalized Linear Model* (GLM) which is a fairly simple model to illustrate message passing algorithms and which is also an elementary brick for a large range of interesting inference questions on neural networks. It falls under the teacher-student set up: a student model is used to reconstruct a signal from a teacher model producing indirect observations. In the GLM, the product of an unknown signal $\underline{x}_0 \in \mathbb{R}^N$ and a known weight matrix $\underline{W} \in \mathbb{R}^{N \times M}$ is observed as \underline{y} through a noisy channel $p_{\text{out},0}$,

$$\begin{cases} \underline{W} \sim p_W(\underline{W}) \\ \underline{x}_0 \sim p_{x_0}(\underline{x}_0) = \prod_{i=1}^N p_{x_0}(x_{0,i}) \end{cases} \Rightarrow \underline{y} \sim p_{\text{out},0}(\underline{y}|\underline{W}\underline{x}_0) = \prod_{\mu=1}^M p_{\text{out},0}(y_\mu|\underline{w}_\mu^\top \underline{x}_0). \quad (50)$$

The probabilistic graphical model corresponding to this teacher is represented in Figure 3. The prior over the signal p_{x_0} is supposed to be factorized, and the channel $p_{\text{out},0}$ likewise. The inference problem is to produce an estimator $\hat{\underline{x}}$ for the unknown signal \underline{x}_0 from the observations \underline{y} . Given the prior p_x and the channel p_{out} of the student, not necessarily matching the teacher, the posterior distribution is

$$p(\underline{x}|\underline{y}, \underline{W}) = \frac{1}{\mathcal{Z}(\underline{y}, \underline{W})} \prod_{\mu=1}^M p_{\text{out}}(y_\mu | \sum_{i=1}^N W_{\mu i} x_i) \prod_{i=1}^N p_x(x_i), \quad (51)$$

$$\mathcal{Z}(\underline{y}, \underline{W}) = \int d\underline{x} p_{\text{out}}(\underline{y}|\underline{x}, \underline{W}) p_x(\underline{x}), \quad (52)$$

represented as a factor graph also in Figure 3. The difficulty of the reconstruction task of \underline{x}_0 from \underline{y} is controlled by the measurement ratio $\alpha = M/N$ and the amplitude of the noise possibly present in the channel.

Applications The generic GLM underlies a number of applications. In the context of neural networks of particular interest in this technical review, the channel p_{out} generating observations $\underline{y} \in \mathbb{R}^M$ can equivalently be seen as a stochastic activation function $g(\cdot; \underline{\epsilon})$ incorporating a noise $\underline{\epsilon} \in \mathbb{R}^M$ component-wise to the output,

$$y_\mu = g(\underline{w}_\mu^\top \underline{x}; \underline{\epsilon}_\mu). \quad (53)$$

The inference of the teacher signal in a GLM has then two possible interpretations. On the one hand, it can be interpreted as the reconstruction of the input \underline{x} of a stochastic single-layer neural network from its output \underline{y} . For example, this inference problem can arise in the maximum likelihood training of a one-layer VAE (see corresponding paragraph in Section 2.2). On the other hand, the same question can also correspond to the Bayesian learning of a single-layer neural network with a single output - the perceptron - where this time $\{\underline{W}, \underline{y}\}$ are interpreted as the collection of training input-output pairs and \underline{x}_0 plays the role of the unknown weight vector of the teacher (as cited as an example in Section 3.1.1). However, note that one of the most important applications of the GLM, Compressed Sensing (CS) [Don06], does not involve neural networks.

4.3 Belief propagation and approximate message passing

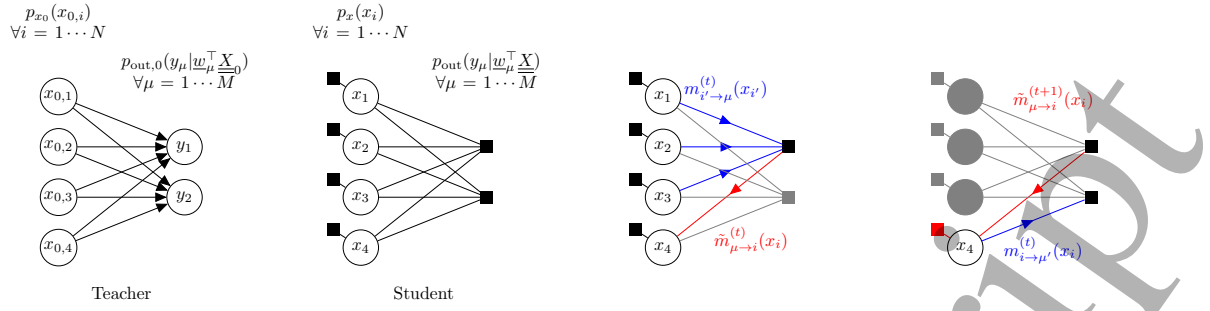


Figure 3: Graphical representations of the Generalized Linear Model. **Left:** Probabilistic graphical model of the teacher. **Middle left:** Factor graph representation of the posterior distribution on the signal \underline{x} under the student statistical model. **Middle right and right:** Belief propagation updates (67) - (68) for approximate inference.

Statistical physics treatment, random weights and scaling From the statistical physics perspective, the effective energy functional is read from the posterior (52) seen as a Boltzmann distribution with energy

$$E(\underline{x}) = -\log p_{\text{out}}(\underline{y} | \underline{x}, \underline{W}) p_x(\underline{x}) = -\sum_{\mu=1}^M \log p_{\text{out}}(y_\mu | \sum_{i=1}^N W_{\mu i} x_i) - \sum_{i=1}^N \log p_x(x_i). \quad (54)$$

The inverse temperature β has here no formal equivalent and can be thought as being equal to 1. The energy is a function of the random realizations of \underline{W} and \underline{y} , playing the role of the disorder. Furthermore, the validity of the approximation presented below require additional assumptions. Crucially, the weight matrix is assumed to have i.i.d. Gaussian entries with zero mean and variance $1/N$, much like in the SK model. The prior of the signal is chosen so as to ensure that the x_i -s (and consequently the y_μ -s) remain of order 1. Finally, the thermodynamic limit $N \rightarrow \infty$ is taken for a fixed measurement ratio $\alpha = M/N$.

4.3.2 Belief Propagation

Recall that inference in high-dimensional problems consists in marginalizations over complex joint distributions, typically in the view of computing partition functions, averages or marginal probabilities for sampling. Belief Propagation (BP) is an inference algorithm, sometimes exact and sometimes approximate as we will see, leveraging the known factorization of a distribution, which encodes the precious information of (in)dependencies between the random variables in the joint distribution. For a generic joint probability distribution p over $\underline{x} \in \mathbb{R}^N$ factorized as

$$p(\underline{x}) = \frac{1}{Z} \prod_{\mu=1}^M \psi_\mu(\underline{x}_{\partial\mu}), \quad (55)$$

ψ_μ are called potential functions taking as arguments the variables x_i -s involved in the factor μ shortened as $\underline{x}_{\partial\mu}$.

Definition of messages Let us first write the BP equations and then explain the origin of these definitions. The underlying factor graph of (55) has N nodes carrying variables x_i -s and M factors associated with the potential functions ψ_μ -s (see Appendix B for a quick reminder). BP acts on *messages* variables which are tied to the edges of the factor graph. Specifically, the sum-product version of the algorithm (as opposed to the max-sum, see e.g. [MM09]) consists in the update equations

$$\tilde{m}_{\mu \rightarrow i}^{(t)}(x_i) = \frac{1}{Z_{\mu \rightarrow i}} \int \prod_{i' \in \partial\mu \setminus i} dx_{i'} \psi_\mu(\underline{x}_{\partial\mu}) \prod_{i' \in \partial\mu \setminus i} m_{i' \rightarrow \mu}^{(t)}(x_{i'}), \quad (56)$$

$$m_{i \rightarrow \mu}^{(t+1)}(x_i) = \frac{1}{Z_{i \rightarrow \mu}} p_x(x_i) \prod_{\mu' \in \partial i \setminus \mu} \tilde{m}_{\mu' \rightarrow i}^{(t)}(x_i) \quad (57)$$

4 Selected overview of mean-field treatments: free energies and algorithms

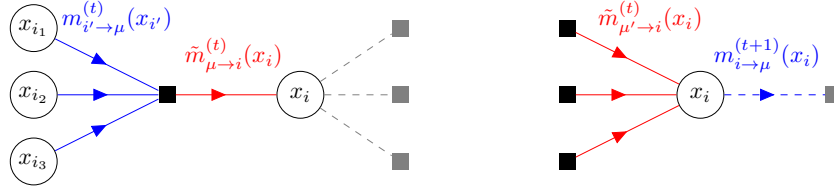


Figure 4: Representations of the neighborhood of edge i - μ in the factor graph and corresponding local BP updates. Factors are represented as squares and variable nodes as circles. **Left:** In the factor graph where all factors around x_i are removed (in gray) except for the factor μ , the marginal of x_i (in red) is updated from the messages incoming at factor μ (in blue) (56). **Right:** In the factor graph where factor μ is deleted (in gray), the marginal of x_i (in blue) is updated with the incoming messages (in red) from the rest of the factors (57).

where again the i -s index the variable nodes and the μ -s index the factor nodes. The notation $\partial\mu \setminus i$ designate the set of neighbor variables of the factor μ except the variable i (and reciprocally for $\partial i \setminus \mu$). The partition functions $\mathcal{Z}_{i \rightarrow \mu}$ and $\mathcal{Z}_{\mu \rightarrow i}$ are normalization factors ensuring that the messages can be interpreted as probabilities.

For acyclic (or tree-like) factor graphs, the BP updates are guaranteed to converge to a fixed point, that is a set of time independent messages $\{m_{i \rightarrow \mu}, \tilde{m}_{\mu \rightarrow i}\}$ solution of the system of equations (56)-(57). Starting at a leaf of the tree, these messages communicate beliefs of a given node variable taking a given value based on the nodes and factors already visited along the tree. More precisely, $\tilde{m}_{\mu \rightarrow i}(x_i)$ is the marginal probability of x_i in the factor graph before visiting the factors in ∂i except for μ , and $m_{i \rightarrow \mu}(x_i)$ is equal the marginal probability of x_i in the factor graph before visiting the factor μ , see Figure 4.

Thus, at convergence of the iterations, the marginals can be computed as

$$m_i(x_i) = \frac{1}{\mathcal{Z}_i} p_x(x_i) \prod_{\mu \in \partial i} \tilde{m}_{\mu \rightarrow i}(x_i), \quad (58)$$

which can be seen as the main output of the BP algorithm. These marginals will only be exact on trees where incoming messages, computed from different part of the graph, are independent. Nonetheless, the algorithm (56)-(57), occasionally then called *loopy-BP*, can sometimes be converged on graphs with cycles and in some cases will still provide high quality approximations. For instance, graphs with no short loops are locally tree like and BP is an efficient method of approximate inference, provided correlations decay with distance (i.e. incoming messages at each node are still effectively independent). BP will also appear principled for some infinite range mean-field models previously discussed; an example of which being our case-study the GLM discussed below. While this is the only example that will be discussed here in the interest of conciseness, getting fluent in BP generally requires more than one application. The interested reader could also consult [YFW02] and [MM09] Section 14.1. for simple concrete examples.

The Bethe free energy The BP algorithm can also be recovered from a variational argument. Let's consider both the single variable marginals $m_i(x_i)$ and the marginals of the neighborhood of each factor $\tilde{m}_\mu(\underline{x}_{\partial\mu})$. On tree graphs, the joint distribution (55) can be re-expressed as

$$p(\underline{x}) = \frac{\prod_{\mu=1}^M \tilde{m}_\mu(\underline{x}_{\partial\mu})}{\prod_{i=1}^N m_i(x_i)^{n_i-1}}, \quad (59)$$

where n_i is the number of neighbor factors of the i -th variable. Abusively, we can use this form as an ansatz for loopy graph and plug it in the Gibbs free energy to derive an approximation of the free energy, similarly to the naive mean-field derivation of Section 4.1. This time the variational parameters will be the distributions m_i and \tilde{m}_μ (see e.g. [YFW02, MM09] for additional details). The corresponding functional form of the Gibbs free energy is called the Bethe free energy:

$$F_{\text{Bethe}}(m_i, \tilde{m}_\mu) = - \int d\underline{x}_{\partial\mu} \tilde{m}_\mu(\underline{x}_{\partial\mu}) \ln \psi_\mu(\underline{x}_{\partial\mu}) + (n_i - 1)H(m_i) - H(\tilde{m}_\mu), \quad (60)$$

4.3 Belief propagation and approximate message passing

where $H(q)$ is the entropy of the distribution q . Optimization of the Bethe free energy with respect to its arguments under the constraint of consistency

$$\int d\mathbf{x}_{\partial\mu \setminus i} \tilde{m}_\mu(\mathbf{x}_{\partial\mu}) = m_i(x_i) \quad (61)$$

involves Lagrange multipliers which can be shown to be related to the messages defined in (56)-(57). Eventually, one can verify that marginals defined as (58) and

$$\tilde{m}_\mu(\mathbf{x}_{\partial\mu}) = \frac{1}{\mathcal{Z}_\mu} \psi(\mathbf{x}_{\partial\mu}) \prod_{i \in \partial\mu} m_{i \rightarrow \mu}(x_i), \quad (62)$$

are stationary point of the Bethe free energy for messages that are BP solutions. In other words, the BP fixed points are consistent with the stationary point of the Bethe free energy. Using the normalizing constants of the messages, the Bethe free energy can also be re-written as

$$F_{\text{Bethe}} = - \sum_{i \in V} \log \mathcal{Z}_i - \sum_{\mu \in F} \log \mathcal{Z}_\mu + \sum_{(i\mu) \in E} \log \mathcal{Z}_{\mu i}, \quad (63)$$

with

$$\mathcal{Z}_i = \int d\mathbf{x}_i p_x(x_i) \prod_{\mu \in \partial i} \tilde{m}_{\mu \rightarrow i}(x_i), \quad (64)$$

$$\mathcal{Z}_\mu = \int \prod_{i \in \partial\mu} d\mathbf{x}_i \psi(\mathbf{x}_{\partial\mu}) \prod_{i \in \partial\mu} m_{i \rightarrow \mu}(x_i), \quad (65)$$

$$\mathcal{Z}_{\mu i} = \int d\mathbf{x}_i \tilde{m}_{\mu \rightarrow i}(x_i) m_{i \rightarrow \mu}(x_i). \quad (66)$$

As for the marginals, the Bethe free energy, will only be exact if the underlying factor graph is a tree. Otherwise it is an approximation of the free energy, that is not generally an upper bound.

Belief propagation for the GLM The writing of the BP-equations for our case-study is schematized on the right of Figure 3. There are $2 \times N \times M$ updates:

$$\tilde{m}_{\mu \rightarrow i}^{(t)}(x_i) = \frac{1}{\mathcal{Z}_{\mu \rightarrow i}} \int \prod_{i' \neq i} d\mathbf{x}_{i'} p_{\text{out}}(y_\mu | \mathbf{w}_\mu^\top \mathbf{x}) \prod_{i' \neq i} m_{i' \rightarrow \mu}^{(t)}(x_{i'}), \quad (67)$$

$$m_{i \rightarrow \mu}^{(t+1)}(x_i) = \frac{1}{\mathcal{Z}_{i \rightarrow \mu}} p_x(x_i) \prod_{\mu' \neq \mu} \tilde{m}_{\mu' \rightarrow i}^{(t)}(x_i), \quad (68)$$

for all $i - \mu$ pairs. Despite a relatively concise formulation, running BP in practice turns out intractable since for a signal \mathbf{x} taking continuous values it would entail keeping track of distributions on continuous variables. In this case, BP is approximated by the (G)AMP algorithm presented in the next section.

4.3.3 (Generalized) approximate message passing

The name of approximate message passing (AMP) was fixed by Donoho, Maleki and Montanari [DMM09] who derived the algorithm in the context of Compressed Sensing. Several works from statistical physics had nevertheless already proposed related algorithmic procedures and made connections with the TAP equations for different systems [KS98, OS01, Kab03]. The algorithm was derived systematically for any channel of the GLM by Rangan [Ran11] and became Generalized-AMP (GAMP), yet again it seems that [KU04] proposed the first generalized derivation.

The systematic procedure to write AMP for a given joint probability distribution consists in first writing BP on the factor graph, second project the messages on a parametrized family of functions to obtain the corresponding *relaxed-BP* and third close the equations on a reduced set of parameters by keeping only leading terms in the thermodynamic limit. We will quickly review and justify these steps for the GLM. Note that here a relevant algorithm for approximate inference will be derived from message passing on a fully connected graph of interactions. As

4 Selected overview of mean-field treatments: free energies and algorithms

it turns out, the high connectivity limit and the introduction of short loops does not break the assumption of independence of incoming messages in this specific case thanks to the small scale $O(1/\sqrt{N})$ and the independence of the weight entries. The statistics of the weights are here crucial.

Relaxed Belief Propagation In the thermodynamic limit $M, N \rightarrow +\infty$, one can show that the scaling $1/\sqrt{N}$ of the W_{ij} and the extensive connectivity of the underlying factor graph imply that messages are approximately Gaussian. Without giving all the details of the computation which can be cumbersome, let us try to provide some intuitions. We drop the time indices for simplicity and start with (67). Consider the intermediate reconstruction variable $z_\mu = \mathbf{w}_\mu^\top \mathbf{x} = \sum_{i' \neq i} W_{\mu i'} x_{i'} + W_{\mu i} x_i$. Under the statistics of the messages $m_{i' \rightarrow \mu}(x_{i'})$, the $x_{i'}$ are independent such that by the central limit theorem $z_\mu - W_{\mu i} x_i$ is a Gaussian random variable with respectively mean and variance

$$\omega_{\mu \rightarrow i} = \sum_{i' \neq i} W_{\mu i'} \hat{x}_{i' \rightarrow \mu}, \quad (69)$$

$$V_{\mu \rightarrow i} = \sum_{i' \neq i} W_{\mu i'}^2 C_{i' \rightarrow \mu}^x, \quad (70)$$

where we defined the mean and the variance of the messages $m_{i' \rightarrow \mu}(x_{i'})$,

$$\hat{x}_{i' \rightarrow \mu} = \int dx_{i'} x_{i'} m_{i' \rightarrow \mu}(x_{i'}), \quad (71)$$

$$C_{i' \rightarrow \mu}^x = \int dx_{i'} x_{i'}^2 m_{i' \rightarrow \mu}(x_{i'}) - \hat{x}_{i' \rightarrow \mu}^2. \quad (72)$$

Using these new definitions, (67) can be rewritten as

$$\tilde{m}_{\mu \rightarrow i}(x_i) \propto \int dz_\mu p_{\text{out}}(y_\mu | z_\mu) e^{-\frac{(z_\mu - W_{\mu i} x_i - \omega_{\mu \rightarrow i})^2}{2V_{\mu \rightarrow i}}}, \quad (73)$$

where the notation \propto omits the normalization factor for distributions. Considering that $W_{\mu i}$ is of order $1/\sqrt{N}$, the development of (73) shows that at leading order $\tilde{m}_{\mu \rightarrow i}(x_i)$ is Gaussian:

$$\tilde{m}_{\mu \rightarrow i}(x_i) \propto e^{B_{\mu \rightarrow i} x_i + \frac{1}{2} A_{\mu \rightarrow i} x_i^2} \quad (74)$$

where the details of the computations yield

$$B_{\mu \rightarrow i} = W_{\mu i} g_{\text{out}}(y_\mu, \omega_{\mu \rightarrow i}, V_{\mu \rightarrow i}) \quad (75)$$

$$A_{\mu \rightarrow i} = -W_{\mu i}^2 \partial_\omega g_{\text{out}}(y_\mu, \omega_{\mu \rightarrow i}, V_{\mu \rightarrow i}) \quad (76)$$

using the *output update functions*

$$g_{\text{out}}(y, \omega, V) = \frac{1}{\mathcal{Z}_{\text{out}}} \int dz \frac{(z - \omega)}{V} p_{\text{out}}(y | z) \mathcal{N}(z; \omega, V), \quad (77)$$

$$\partial_\omega g_{\text{out}}(y, \omega, V) = \frac{1}{\mathcal{Z}_{\text{out}}} \int dz \frac{(z - \omega)^2}{V^2} p_{\text{out}}(y | z) \mathcal{N}(z; \omega, V) - \frac{1}{V} - g_{\text{out}}(y, \omega, V)^2, \quad (78)$$

$$\mathcal{Z}_{\text{out}}(y, \omega, V) = \int dz p_{\text{out}}(y | z) \mathcal{N}(z; \omega, V). \quad (79)$$

These arguably complicated functions, again coming out of the development of (73), can be interpreted as the estimation of the mean and the variance of the gap between two different estimate of z_μ considered by the algorithm: the mean estimate $\omega_{\mu \rightarrow i}$ given incoming messages $m_{i' \rightarrow \mu}(x_{i'})$ and the same mean estimate updated to incorporate the information coming from the channel p_{out} and observation y_μ . Finally, the Gaussian parametrization (74) of $\tilde{m}_{\mu \rightarrow i}(x_i)$ serves to rewrite the other type of messages $m_{i \rightarrow \mu}(x_i)$ (68),

$$m_{i \rightarrow \mu}(x_i) \propto p_x(x_i) e^{-\frac{(\lambda_{i \rightarrow \mu} - x_i)^2}{2\sigma_{i \rightarrow \mu}}}, \quad (80)$$

4.3 Belief propagation and approximate message passing

with

$$\sigma_{i \rightarrow \mu} = \left(\sum_{\mu' \neq \mu} A_{\mu' \rightarrow i} \right)^{-1} \quad (81)$$

$$\lambda_{i \rightarrow \mu} = \sigma_{i \rightarrow \mu} \left(\sum_{\mu' \neq \mu} B_{\mu' \rightarrow i} \right). \quad (82)$$

The set of equations can finally be closed by recalling the definitions (71)-(72):

$$\hat{x}_{i \rightarrow \mu} = f_1^x(\lambda_{i \rightarrow \mu}, \sigma_{i \rightarrow \mu}) \quad (83)$$

$$C_{i \rightarrow \mu}^x = f_2^x(\lambda_{i \rightarrow \mu}, \sigma_{i \rightarrow \mu}) \quad (84)$$

with now the *input update functions*

$$\mathcal{Z}^x = \int dx p_x(x) e^{-\frac{(x-\lambda)^2}{2\sigma}}, \quad (85)$$

$$f_1^x(\lambda, \sigma) = \frac{1}{\mathcal{Z}^x} \int dx x p_x(x) e^{-\frac{(x-\lambda)^2}{2\sigma}}, \quad (86)$$

$$f_2^x(\lambda, \sigma) = \frac{1}{\mathcal{Z}^x} \int dx x^2 p_x(x) e^{-\frac{(x-\lambda)^2}{2\sigma}} - f_1^x(\lambda, \sigma)^2. \quad (87)$$

The input update functions can be interpreted as updating the estimation of the mean and variance of the signal x_i based on the information coming from the incoming messages grasped by $\lambda_{i \rightarrow \mu}$ and $\sigma_{i \rightarrow \mu}$ with the information of the prior p_x .

To sum-up, by considering the leading order terms in the thermodynamic limit, the BP equations can be self-consistently re-written as a closed set of equations over mean and variance variables (69)-(70)-(75)-(76)-(81)-(82)-(83)-(84). Eventually, r-BP can equivalently be thought of as the projection of BP onto the following parametrizations of the messages

$$\tilde{m}_{\mu \rightarrow i}^{(t)}(x_i) \propto e^{B_{\mu \rightarrow i}^{(t)} x_i + \frac{1}{2} A_{\mu \rightarrow i}^{(t)} x_i^2} \propto \int dz_\mu p_{\text{out}}(y_\mu | z_\mu) e^{-\frac{(z_\mu - W_{\mu i} x_i - \omega_{\mu \rightarrow i}^{(t)})^2}{2V_{\mu \rightarrow i}^{(t)}}}, \quad (88)$$

$$m_{i \rightarrow \mu}^{(t+1)}(x_i) \propto e^{-\frac{(\hat{x}_{i \rightarrow \mu}^{(t+1)} - x_i)^2}{2C_{i \rightarrow \mu}^{(t+1)}}} \propto p_x(x_i) e^{-\frac{(\lambda_{i \rightarrow \mu}^{(t)} - x_i)^2}{2\sigma_{i \rightarrow \mu}^{(t)}}}. \quad (89)$$

Note that, at convergence, an approximation of the marginals is recovered from the projection on the parametrization (89) of (58),

$$m_i(x_i) = \frac{1}{\mathcal{Z}_i} p_x(x_i) e^{-(\lambda_i - x_i)^2 / 2\sigma_i} = f_1^x(\lambda_i, \sigma_i), \quad (90)$$

$$\sigma_{i \rightarrow \mu} = \left(\sum_{\mu} A_{\mu \rightarrow i} \right)^{-1}, \quad (91)$$

$$\lambda_{i \rightarrow \mu} = \sigma_{i \rightarrow \mu} \left(\sum_{\mu} B_{\mu \rightarrow i} \right). \quad (92)$$

Nonetheless, r-BP is scarcely used as such as the computational cost can be readily reduced with little more approximation. Because the parameters in (88)-(89) take the form of messages on the edges of the factor graph there are still $O(M \times N)$ quantities to track to solve the self-consistency equations by iterations. Yet, in the thermodynamic limit, the messages are closely related to the marginals as the contribution of the missing message between (57) and (58) is to a certain extent negligible. Careful book keeping of the order of contributions of these small differences leads to a set of closed equations on parameters of the marginals, i.e. $O(N)$ variables, corresponding to the GAMP algorithm.

A detailed derivation and developed algorithm of r-BP for the GLM can be found for example in [ZK16] (Section 6.3.1). In Section 5.3 of the present paper, we also present the derivation in a slightly more general setting where the variables x_i and y_μ are possibly vectors instead of scalars.

4 Selected overview of mean-field treatments: free energies and algorithms

Algorithm 1 Generalized Approximate Message Passing**Input:** vector $y \in \mathbb{R}^M$ and matrix $\underline{W} \in \mathbb{R}^{M \times N}$:*Initialize:* $\hat{x}_i, C_i^x \quad \forall i$ and $g_{\text{out}\mu}, \partial_\omega g_{\text{out}\mu} \quad \forall \mu$ **repeat**1) Estimate mean $\omega_\mu^{(t)}$ and variance $V_\mu^{(t)}$ of z_μ given current $\hat{x}^{(t)}$

$$V_\mu^{(t)} = \sum_{i=1}^N W_{\mu i}^2 C_i^{x(t)} \quad (93)$$

$$\omega_\mu^{(t)} = \sum_{i=1}^N W_{\mu i} \hat{x}_i^{(t)} - \sum_{i=1}^N W_{\mu i}^2 C_i^{x(t)} g_{\text{out}\mu}^{(t-1)} \quad (94)$$

2) Estimate mean $g_{\text{out}\mu}^{(t)}$ and variance $\partial_\omega g_{\text{out}\mu}^{(t)}$ of the gap between optimal z_μ and $\omega_\mu^{(t)}$ given y_μ

$$\partial_\omega g_{\text{out}\mu}^{(t)} = \partial_\omega g_{\text{out}}(y_\mu, \omega_\mu^{(t)}, V_\mu^{(t)}) \quad (95)$$

$$g_{\text{out}\mu}^{(t)} = g_{\text{out}}(y_\mu, \omega_\mu^{(t)}, V_\mu^{(t)}) \quad (96)$$

3) Estimate mean λ and variance σ of \underline{x} given current gap between optimal \underline{z} and $\underline{\omega}$

$$\sigma_i^{(t)} = \left(- \sum_{\mu=1}^M W_{\mu i}^2 \partial_\omega g_{\text{out}\mu}^{(t)} \right)^{-1} \quad (97)$$

$$\lambda_i^{(t)} = \hat{x}_i^{(t)} + \sigma_i^{(t)} \left(\sum_{\mu=1}^M W_{\mu i} g_{\text{out}\mu}^{(t)} \right) \quad (98)$$

4) Estimate of mean $\hat{x}^{(t+1)}$ and $\underline{C}^{x(t+1)}$ variance of \underline{x} updated with information from the prior

$$C_i^{x(t+1)} = f_2^x(\lambda_i^{(t)}, \sigma_i^{(t)}) \quad (99)$$

$$\hat{x}_i^{(t+1)} = f_1^x(\lambda_i^{(t)}, \sigma_i^{(t)}) \quad (100)$$

until convergence

Generalized approximate message passing The GAMP algorithm with respect to marginal parameters, analogous to the messages parameters introduced above (summarized in (88)-(89)), is given in Algorithm 1. The origin of GAMP is again the development of the r-BP message-like equations around marginal quantities. The details of this derivation for the GLM can be found for instance in [ZK16] (Section 6.3.1). For a random initialization, the algorithm can be decomposed in 4 steps per iteration which refine the estimate of the signal \underline{x} and the intermediary variable \underline{z} by incorporating the different sources of information. Steps 2) and 4) involve the *update functions* relative to the prior and output channel defined above. Steps 1) and 3) are general for any GLM with a random Gaussian weight matrix, as they result from the consistency of the two alternative parametrizations introduced for the same messages in (88)-(89).

Relation to TAP equations Historically the main difference between the AMP algorithm and the TAP equations is that the latter was first derived for binary variables with 2-body interactions (SK model) while the former was proposed for continuous random variables with N -body interactions (Compressed Sensing). The details of the derivation (described in [ZK16] or in a more general case in Section 5.3), rely on the knowledge of the statistics of the disordered variable \underline{W} but do not require a disorder average, as in the Georges-Yedidia expansion yielding the TAP equations. By focusing on the GLM with a random Gaussian weight matrix scaling as $O(1/\sqrt{N})$ (similarly to the couplings of the SK model) we naturally obtained TAP equations at second order, with an Onsager term in the update (94) of ω_μ . Yet an advantage of the AMP derivation from BP over the high-temperature expansion is that it explicitly provides ‘correct’ time indices in the iteration scheme to solve the self consistent equations [Bol14].

4.3 Belief propagation and approximate message passing

Reconstruction with AMP AMP is therefore a practical reconstruction algorithm which can be run on a single instance (the disorder is not averaged) to estimate an unknown signal \underline{x}_0 . Note that the prior p_x and channel p_{out} used in the algorithm correspond to the student statistical model and they may be different from the true underlying teacher model that generates \underline{x}_0 and \underline{y} . In other words, the AMP algorithm may be used either in the Bayes optimal or in the mismatched setting defined in Section 3.1.1. Remarkably, it is also possible to consider a disorder average in the thermodynamic limit to study the average case computational hardness, here of the GLM inference problem, in either of these matched or mismatched configurations.

State Evolution The statistical analysis of the AMP equations for Compressed Sensing in the average case and in the thermodynamic limit $N \rightarrow \infty$ lead to another closed set of equations that was called State Evolution (SE) in [DMM09]. Such an analysis can be generalized to other problems of application of approximate message passing algorithms. The derivation of SE starts from the r-BP equations and relies on the assumption of independent incoming messages to invoke the Central Limit Theorem. It is therefore only necessary to follow the evolution of a set of means and variances parametrizing Gaussian distributions. When the different variables and factors are statistically equivalent, as it is the case of the GLM, SE reduces to a few scalar equations. The interested reader should refer to Appendix F for a detailed derivation in a more general setting.

Mismatched setting In the general mismatched setting we need to carefully differentiate the teacher and the student. We note p_{x_0} the prior used by the teacher. We also rewrite its channel $p_{\text{out},0}(y|\underline{w}^\top \underline{x})$ as the explicit function $y = g_0(\underline{w}^\top \underline{x}; \epsilon)$ assuming the noise ϵ to be distributed according to the standard normal distribution. The tracked quantities are the *overlaps*,

$$q = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \hat{x}_i^2, \quad m = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \hat{x}_i x_{0,i}, \quad q_0 = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N x_{0,i}^2 = \mathbb{E}_{p_{x_0}}[x_0^2], \quad (101)$$

along with the auxiliary V , \hat{q} , \hat{m} and $\hat{\chi}$:

$$\hat{q}^{(t)} = \int \mathcal{D}\epsilon \int d\omega dz \mathcal{N}(z, \omega; 0, \underline{\underline{Q}}^{(t)}) g_{\text{out}}(\omega, g_0(z; \epsilon), V^{(t)})^2, \quad (102)$$

$$\hat{m}^{(t)} = \int \mathcal{D}\epsilon \int d\omega dz \mathcal{N}(z, \omega; 0, \underline{\underline{Q}}^{(t)}) \partial_z g_{\text{out}}(\omega, g_0(z; \epsilon), V^{(t)}), \quad (103)$$

$$\hat{\chi}^{(t)} = - \int \mathcal{D}\epsilon \int d\omega dz \mathcal{N}(z, \omega; 0, \underline{\underline{Q}}^{(t)}) \partial_\omega g_{\text{out}}(\omega, g_0(z; \epsilon), V^{(t)}), \quad (104)$$

$$q^{(t+1)} = \int d\mathbf{x}_0 p_{x_0}(\mathbf{x}_0) \int \mathcal{D}\xi f_1^x \left((\alpha \hat{\chi}^{(t)})^{-1} \left(\sqrt{\alpha \hat{q}^{(t)}} \xi + \alpha \hat{m}^{(t)} \underline{x}_0 \right); (\alpha \hat{\chi}^{(t)})^{-1} \right)^2, \quad (105)$$

$$m^{(t+1)} = \int d\mathbf{x}_0 p_{x_0}(\mathbf{x}_0) \int \mathcal{D}\xi x_0 f_1^x \left((\alpha \hat{\chi}^{(t)})^{-1} \left(\sqrt{\alpha \hat{q}^{(t)}} \xi + \alpha \hat{m}^{(t)} \underline{x}_0 \right); (\alpha \hat{\chi}^{(t)})^{-1} \right), \quad (106)$$

$$V^{(t+1)} = \int d\mathbf{x}_0 p_{x_0}(\mathbf{x}_0) \int \mathcal{D}\xi f_2^x \left((\alpha \hat{\chi}^{(t)})^{-1} \left(\sqrt{\alpha \hat{q}^{(t)}} \xi + \alpha \hat{m}^{(t)} \underline{x}_0 \right); (\alpha \hat{\chi}^{(t)})^{-1} \right), \quad (107)$$

where we use the notation $\mathcal{N}(\cdot; \cdot, \cdot)$ for the normal distribution, $\mathcal{D}\xi$ for the standard normal measure and the covariance matrix $\underline{\underline{Q}}^{(t)}$ is given at each time step by

$$\underline{\underline{Q}}^{(t)} = \begin{bmatrix} q_0 & m^{(t)} \\ m^{(t)} & q^{(t)} \end{bmatrix}.$$

4 Selected overview of mean-field treatments: free energies and algorithms

Due to the self-averaging property, the performance of the reconstruction by the AMP algorithm on an instance of size N can be tracked along the iterations given

$$\text{MSE}(\hat{x}) = \frac{1}{N} \sum_{i=1}^N (\hat{x}_i - x_{0,i})^2 = q - 2m + q_0, \quad (108)$$

with only minor differences coming from finite-size effects. State Evolution also provides an efficient procedure to study from the theoretical perspective the AMP fixed points for a generic model, such as the GLM, as a function of some control parameters. It reports the average results for running the complete AMP algorithm on $O(N)$ variables using a few scalar equations. Furthermore, the State Evolution equations simplify further in the Bayes optimal setting.

Bayes optimal setting When the prior and channel are identical for the student and the teacher, the true unknown signal \underline{x}_0 is in some sense statistically equivalent to the estimate \hat{x} coming from the posterior. More precisely one can prove the Nishimori identities [OH91, Iba99, Nis01] (or [KKM⁺16] for a concise demonstration and discussion) implying that $q = m$, $V = q_0 - m$ and $\hat{q} = \hat{m} = \hat{\chi}$. Only two equations are then necessary to track the performance of the reconstruction:

$$\hat{q}^{(t)} = \int d\epsilon p_{\epsilon_0}(\epsilon) \int d\omega dz \mathcal{N}(z, \omega; 0, \underline{Q}^{(t)}) g_{\text{out}}(\omega, g_0(z, \epsilon), V^{(t)})^2 \quad (109)$$

$$q^{(t+1)} = \int d\mathbf{x}_0 p_{x_0}(\mathbf{x}_0) \int \mathcal{D}\xi f_1^x \left((\alpha \hat{\chi}^{(t)})^{-1} \left(\sqrt{\alpha \hat{q}^{(t)}} \xi + \alpha \hat{m}^{(t)} \underline{x}_0 \right); (\alpha \hat{\chi}^{(t)})^{-1} \right)^2. \quad (110)$$

4.4 Replica method

Another powerful technique from the statistical physics of disordered systems to examine models with infinite range interactions is the replica method. It enables an analytical computation of the quenched free energy via non-rigorous mathematical manipulations. More developed introductions to the method can be found in [MPV86, Nis01, CCFM05].

4.4.1 Steps of a replica computation

The basic idea of the replica computation is to compute the average over the disorder of $\log \mathcal{Z}$ by considering the identity $\log \mathcal{Z} = \lim_{n \rightarrow 0} (\mathcal{Z}^n - 1)/n$. First the expectation of \mathcal{Z}^n is evaluated for $n \in \mathbb{N}$, then the $n \rightarrow 0$ limit is taken by ‘analytic continuation’. Thus the method takes advantage of the fact that the average of a power of \mathcal{Z} is sometimes easier to compute than the average of a logarithm. We illustrate the key steps of the calculation for the partition function of the GLM (52).

Disorder average for the replicated system: coupling of the replicas The average of \mathcal{Z}^n for $n \in \mathbb{N}$ can be seen as the partition function of a system with $n+1$ non interacting replicas of \underline{x} indexed by $a \in \{0, \dots, n\}$, where the first replica $a = 0$ is representative of the teacher and the n other replicas are identically distributed as the student:

$$\mathbb{E}_{\underline{W}, \underline{y}, \underline{x}_0} [\mathcal{Z}^n] = \mathbb{E}_{\underline{W}} \left[\int d\underline{y} d\underline{x}_0 p_{\text{out},0}(\underline{y} | \underline{W} \underline{x}_0) p_{\underline{x}_0}(\underline{x}_0) \left(\int d\underline{x} p_{\text{out}}(\underline{y} | \underline{W} \underline{x}) p_{\underline{x}}(\underline{x}) \right)^n \right] \quad (111)$$

$$= \mathbb{E}_{\underline{W}} \left[\int d\underline{y} \prod_{a=0}^n \left(d\underline{x}_a p_{\text{out},a}(\underline{y} | \underline{W} \underline{x}_a) p_{\underline{x}_a}(\underline{x}_a) \right) \right] \quad (112)$$

$$= \mathbb{E}_{\underline{W}} \left[\int d\underline{y} \prod_{a=0}^n \left(d\underline{x}_a d\underline{z}_a \delta(\underline{z}_a - \underline{W} \underline{x}_a) p_{\text{out},a}(\underline{y} | \underline{z}_a) p_{\underline{x}_a}(\underline{x}_a) \right) \right]. \quad (113)$$

To perform the average over the disordered interactions \underline{W} we consider the statistics of $\underline{z}_a = \underline{W} \underline{x}_a$. Recall that $W_{\mu i} \sim \mathcal{N}(W_{\mu i}; 0, 1/N)$, independently for all μ and i . Consequently, the \underline{z}_a

4.4 Replica method

are jointly Gaussian in the thermodynamic limit with means and covariances

$$\mathbb{E}_{\underline{\underline{W}}}[z_{a,\mu}] = \mathbb{E}_{\underline{\underline{W}}} \left[\sum_{i=1}^N W_{\mu i} x_{a,i} \right] = 0, \quad E_{\underline{\underline{W}}}[z_{a,\mu} z_{b,\nu}] = \sum_{i=1}^N x_{a,i} x_{b,i} / N = q_{ab}. \quad (114)$$

The overlaps, that we already introduced in the SE formalism, naturally re-appear. We introduce the notation $\underline{\underline{q}}$ for the $(n+1) \times (n+1)$ overlap matrix. Integrating out the disorder $\underline{\underline{W}}$ shared by the $n+1$ replicas will therefore leave us with an effective system of now coupled replicas:

$$\begin{aligned} \mathbb{E}_{\underline{\underline{W}}, \underline{y}, \underline{x}_0} [\mathcal{Z}^n] &= \int \prod_{a,b} dN q_{ab} \int d\underline{y} \prod_{a=0}^n d\underline{z}_a p_{\text{out},a}(\underline{y} | \underline{z}_a) \\ &\exp \left(-\frac{1}{2} \sum_{\mu=1}^M \sum_{a,b} z_{a,\mu} z_{b,\mu} (\underline{\underline{q}}^{-1})_{ab} - MC(\underline{\underline{q}}, n) \right) \\ &\int \prod_{a=1}^n d\underline{x}_a p_{x_a}(\underline{x}_a) \delta(N q_{ab} - \sum_{i=1}^N x_{a,i} x_{b,i}). \end{aligned} \quad (115)$$

Change of variable for the overlaps: decoupling of the variables We consider the Fourier representation of the Dirac distribution fixing the consistency between overlaps and replicas,

$$\delta(N q_{ab} - \sum_{i=1}^N x_{a,i} x_{b,i}) = \int \frac{d\hat{q}_{ab}}{2i\pi} e^{\hat{q}_{ab}(N q_{ab} - \sum_{i=1}^N x_{a,i} x_{b,i})}, \quad (116)$$

where \hat{q}_{ab} is purely imaginary, which yields

$$\begin{aligned} \mathbb{E}_{\underline{\underline{W}}, \underline{y}, \underline{x}_0} [\mathcal{Z}^n] &= \int \prod_{a,b} dN q_{ab} \int \prod_{a,b} \frac{d\hat{q}_{ab}}{2i\pi} \exp(N \hat{q}_{ab} q_{ab}) \\ &\int d\underline{y} \prod_{a=0}^n d\underline{z}_a p_{\text{out},a}(\underline{y} | \underline{z}_a) \exp \left(-\frac{1}{2} \sum_{\mu=1}^M \sum_{a,b} z_{a,\mu} z_{b,\mu} (\underline{\underline{q}}^{-1})_{ab} - MC(\underline{\underline{q}}, n) \right) \\ &\int \prod_{a=1}^n d\underline{x}_a p_{x_a}(\underline{x}_a) \exp \left(-\hat{q}_{ab} \sum_{i=1}^N x_{a,i} x_{b,i} \right) \end{aligned} \quad (117)$$

where $C(\underline{\underline{q}}, n)$ is related to the normalization of the Gaussian distributions over the \underline{z}_a variables, and the integrals can be factorized over the i -s and μ -s. Thus we obtain

$$\mathbb{E}_{\underline{\underline{W}}, \underline{y}, \underline{x}_0} [\mathcal{Z}^n] = \int \prod_{a,b} dN q_{ab} \int \prod_{a,b} d\hat{q}_{ab} e^{N \hat{q}_{ab} q_{ab}} e^{M \log \hat{\mathcal{I}}_z(\underline{\underline{q}})} e^{N \log \hat{\mathcal{I}}_x(\underline{\underline{q}})}, \quad (118)$$

with

$$\hat{\mathcal{I}}_z(\underline{\underline{q}}) = \int d\underline{y} \prod_{a=0}^n d\underline{z}_a p_{\text{out},a}(\underline{y} | \underline{z}_a) \exp \left(-\frac{1}{2} \sum_{a,b} z_a z_b (q_{ab})^{-1} - C(\underline{\underline{q}}, n) \right), \quad (119)$$

$$\hat{\mathcal{I}}_x(\underline{\underline{q}}) = \int \prod_{a=1}^n d\underline{x}_a p_{x_a}(\underline{x}_a) \exp(-\hat{q}_{ab} x_a x_b), \quad (120)$$

where we introduce the notation $\underline{\underline{\hat{q}}}$ for the auxiliary overlap matrix with entries $(\hat{q})_{ab} = \hat{q}_{ab}$ and we omitted the factor $2i\pi$ which is eventually subleading as $N \rightarrow +\infty$. The decoupling of the x_i and the z_μ of the infinite range system yields pre-factors N and M in the exponential arguments. In the thermodynamic limit, we recall that both N and M tend to $+\infty$ while

4 Selected overview of mean-field treatments: free energies and algorithms

the ratio $\alpha = M/N$ remains fixed. Hence, the integral for the replicated average is easily computed in this limit by the saddle point method:

$$\log \mathbb{E}_{\underline{W}, \underline{y}, \underline{x}_0} [\mathcal{Z}^n] \simeq N \text{extr}_{\underline{q}, \underline{\hat{q}}} [\phi(\underline{q}, \underline{\hat{q}})], \quad \phi(\underline{q}, \underline{\hat{q}}) = \sum_{a,b} \hat{q}_{ab} q_{ab} + \alpha \hat{\mathcal{I}}_z(\underline{q}) + \hat{\mathcal{I}}_x(\underline{\hat{q}}), \quad (121)$$

where we defined the replica potential ϕ .

Exchange of limits: back to the quenched average The thermodynamic average of the log-partition is recovered through an a priori risky mathematical manipulation: (i) perform an analytical continuation from $n \in \mathbb{N}$ to $n \rightarrow 0$

$$\frac{1}{N} \mathbb{E}_{\underline{W}, \underline{y}, \underline{x}_0} [\log \mathcal{Z}] = \lim_{n \rightarrow 0} \frac{1}{nN} \mathbb{E}_{\underline{W}, \underline{y}, \underline{x}_0} [\mathcal{Z}^n - 1] = \lim_{n \rightarrow 0} \frac{1}{nN} \log \mathbb{E}_{\underline{W}, \underline{y}, \underline{x}_0} [\mathcal{Z}^n] \quad (122)$$

and (ii) exchange limits

$$-f = \lim_{N \rightarrow \infty} \lim_{n \rightarrow 0} \frac{1}{n} \log \mathbb{E}_{\underline{W}, \underline{y}, \underline{x}_0} [\mathcal{Z}^n] = \lim_{n \rightarrow 0} \frac{1}{n} \text{extr}_{\underline{q}, \underline{\hat{q}}} [\phi(\underline{q}, \underline{\hat{q}})]. \quad (123)$$

Despite the apparent lack of rigour in taking these last steps, the replica method has been proven to yield exact predictions in the thermodynamic limit for different problems and in particular for the GLM [RP16, BKM⁺18].

Saddle point solution: choice of a replica ansatz At this point, we are still left with the problem of computing the extrema of $\phi(\underline{q}, \underline{\hat{q}})$. To solve this optimization problem over \underline{q} and $\underline{\hat{q}}$, a natural assumption is that replicas, that are a pure artefact of the calculation, are equivalent. This is reflected in a special structure for overlap matrices between replicas that only depend on three parameters each,

$$\underline{q} = \begin{bmatrix} q_0 & m & m & m \\ m & q & q_{12} & q_{12} \\ m & q_{12} & q & q_{12} \\ m & q_{12} & q_{12} & q \end{bmatrix}, \quad \underline{\hat{q}} = \begin{bmatrix} \hat{q}_0 & \hat{m} & \hat{m} & \hat{m} \\ \hat{m} & \hat{q} & \hat{q}_{12} & \hat{q}_{12} \\ \hat{m} & \hat{q}_{12} & \hat{q} & \hat{q}_{12} \\ \hat{m} & \hat{q}_{12} & \hat{q}_{12} & \hat{q} \end{bmatrix}, \quad (124)$$

here given as an example for $n = 3$ replicas. Plugging this *replica symmetric* (RS) ansatz in the expression of $\phi(\underline{q}, \underline{\hat{q}})$, taking the limit $n \rightarrow 0$ and looking for the stationary points as a function of the parameters q, m, q_{12} and $\hat{m}, \hat{q}, \hat{q}_{12}$ recovers a set of equations equivalent to SE (7), albeit without time indices. Hence the two a priori different heuristics of BP and the replica method are remarkably consistent under the RS assumption.

Nevertheless, the replica symmetry can be spontaneously broken in the large N limit and the dominating saddle point does not necessarily correspond to the RS overlap matrix. This replica symmetry breaking (RSB) corresponds to substantial changes in the structure of the examined Boltzmann distribution. It is among the great strengths of the replica formalism to naturally capture it. Yet for inference problems falling under the teacher-student scenario, the correct ansatz is always replica symmetric in the Bayes optimal setting [Nis01, CCFM05, ZK16], and we will not investigate here this direction further. The interested reader can refer to the classical references for an introduction to replica symmetry breaking [MPV86, Nis01, CCFM05] in the context of the theory of spin-glasses.

Bayes optimal setting As in SE the equations simplify in the matched setting, where the first replica corresponding to the teacher becomes equivalent to all the others. The replica free energy of the GLM is then given as the extremum of a potential over two scalar variables:

$$-f = \text{extr}_{q, \hat{q}} \left[-\frac{1}{2} q \hat{q} + \mathcal{I}_x(\hat{q}) + \alpha \mathcal{I}_z(q_0, q) \right] \quad (125)$$

$$\mathcal{I}_x(\hat{q}) = \int \mathcal{D}\xi \, dx \, p_x(x) e^{-\frac{\hat{q}}{2} x^2 + \sqrt{\hat{q}} \xi x} \log \left(\int dx' \, p_x(x') e^{-\frac{\hat{q}}{2} x'^2 + \sqrt{\hat{q}} \xi x'} \right) \quad (126)$$

$$\mathcal{I}_z(q, q_0) = \int \mathcal{D}\xi \, dy \, dz \, p_{\text{out}}(y|z) \mathcal{N}(z; \sqrt{q} \xi, q_0 - q) \log \left(\int dz' \, p_{\text{out}}(y|z') \mathcal{N}(z'; \sqrt{q} \xi, q_0 - q) \right). \quad (127)$$

The saddle point equations corresponding to the extremization (125), fixing the values of q and \hat{q} , would again be found equivalent to the Bayes optimal SE (109) - (110). This Bayes optimal result is derived in [KMS⁺12] for the case of a linear channel and Gauss-Bernoulli prior, and can also be recovered as a special case of the low-rank matrix factorization formula (where the measurement matrix is in fact known) [KKM⁺16].

4.4.2 Assumptions and relation to other mean-field methods

A crucial point in the above derivation of the replica formula is the extensivity of the interactions of the infinite range model that allowed the factorization of the N scaling of the argument of the exponential integrand in (118). The statistics of the disorder \underline{W} and in particular the independence of all the W_{ui} was also necessary. While this is an important assumption for the technique to go through, it can be possible to relax it for some types of correlation statistics, as we will see in Section 5.2.

Note that the replica method directly enforces the disorder averaging and does not provide a prediction at the level of the single instance. Therefore it cannot be turned into a practical algorithm of reconstruction. Nonetheless, we have seen that the saddle point equations of the replica derivation, under the RS assumption, matches the SE equations derived from BP. This is sufficient to theoretically study inference questions under a teacher-student scenario in the Bayes optimal setting, and in particular predict the MSE following (108).

In the mismatched setting however, the predictions of the replica method under the RS assumption and the equivalent BP conclusions can be wrong. By introducing the symmetry breaking between replicas, the method can sometimes be corrected. It is an important endeavor of the replica formalism to grasp the importance of the overlaps and connect the form of the replica ansatz to the properties of the joint probability distribution examined. When BP fails on loopy graphs, correlations between variables are not decaying with distance, which manifests into an RSB phase. Note that there also exists message passing algorithms operating in this regime [MP01, MPZ02, MM09, SLL19, AFUZ19, AKUZ19].

5 Further extensions of interest for learning

In the previous Section we presented the classical mean-field approximations focusing on the simple and original examples of the Boltzmann machine (a.k.a. SK model) and the GLM with Gaussian i.i.d weight matrices. Along the way, we tried to emphasize how the procedures of approximation rely on structural (e.g. connectivity) and statistical properties of the model under scrutiny. In the present Section, we will see that extensions of the message passing and replica methods have now broadened the span of applicability of mean-field approximations. We focus on a selection of recent developments of particular interest to study learning problems.

5.1 Streaming AMP for online learning

In learning applications, it is sometimes advantageous for speed or generalization concerns to only treat a subset of examples at the time - making for instance the SGD algorithm the most popular training algorithm in deep learning. Sometimes also, the size of the current data sets may exceed the available memory. Methods implementing a step-by-step learning as the data arrives are referred to as *online*, *streaming* or *mini-batch* learning, as opposed to *offline* or *batch* learning.

In [MKTZ18], a mini-batch version of the AMP algorithm is proposed. It consists in a generalization of Assumed Density Filtering [OW99a, RKI16] that are fully-online, meaning that only a single example is received at once, or mini-batches have size 1. The general derivation principle is the same. On the example of the GLM, one imagines receiving at each iteration a subset of the components of \underline{y} to reconstruct \underline{x} . We denote by $\underline{y}_{(k)}$ these successive mini-batches. Bayes formula gives the posterior distribution over \underline{x} after seeing k mini-batches

$$p(\underline{x}|\underline{y}_{(k)}, \{\underline{y}_{(k-1)}, \dots, \underline{y}_{(1)}\}) = \frac{p(\underline{y}_{(k)}|\underline{x})p(\underline{x}|\{\underline{y}_{(k-1)}, \dots, \underline{y}_{(1)}\})}{\int d\underline{x} p(\underline{y}_{(k)}|\underline{x})p(\underline{x}|\{\underline{y}_{(k-1)}, \dots, \underline{y}_{(1)}\})}. \quad (128)$$

This formula suggests the iterative scheme of using as a prior on \underline{x} at iteration k the posterior obtained at iteration $k - 1$. This idea can be implemented in different approximate inference algorithms, as also noticed by [BBW⁺13] using a variational method. In the regular version

5 Further extensions of interest for learning

of AMP an effective factorized posterior is given at convergence by the input update functions (85)-(87):

$$p(\underline{x}|\underline{y}, \underline{W}) \simeq \prod_{i=1}^N \frac{1}{\mathcal{Z}_x(\lambda_i, \sigma_i)} p_x(x_i) e^{-\frac{(\lambda_i - x_i)^2}{2\sigma_i}}. \quad (129)$$

Plugging this posterior approximation in the iterative scheme yields the mini-AMP algorithm using the converged values of $\lambda_{(\ell)_i}$ and $\sigma_{(\ell)_i}$ at each anterior mini-batch $\ell < k$ to compute the prior

$$p_x^{(k)}(\underline{x}) = p(\underline{x}|\{\underline{y}_{(k-1)}, \dots, \underline{y}_{(1)}\}, \underline{W}) \simeq \prod_{i=1}^N \frac{1}{\mathcal{Z}_{x,i}} p_x(x_i) e^{-\sum_{\ell=1}^{k-1} \frac{(\lambda_{(\ell)_i} - x_i)^2}{2\sigma_{(\ell)_i}}}, \quad (130)$$

where the $\mathcal{Z}_{x,i}$ normalize each marginal factor. Compared to a naive mean-field variational approximation of the posterior, AMP takes into account more correlations and is indeed found to perform better in experiments reported by [MKTZ18]. Another advantage of the AMP based online inference is that it is amenable to theoretical analysis by a corresponding State Evolution [OW99a, RKI16, MKTZ18].

5.2 Algorithms and free energies beyond i.i.d. matrices

The derivations outlined in the previous Sections of the equivalent replica, TAP and AMP equations required the weight matrices to have Gaussian i.i.d. entries. In this case, rigorous proofs of asymptotic exactness of the mean-field solutions were found, for the SK model [Tal06] and the GLM [RP16, BKM⁺18]. Mean-field inference with different weight statistics is a priori feasible if one finds a way either to perform the corresponding disorder average in the replica computation, to evaluate the corresponding Onsager correction in the TAP equations, or to write a message passing where messages remain uncorrelated (even in the high-connectivity limit we may be interested in).

Efforts to broaden in practice the class of matrices amenable to such mean-field treatments lead to a series of works in statistical physics and signal processing with related propositions. Parisi and Potters pioneered this direction by deriving mean-field equations for orthogonal weight matrices using a high-temperature expansion [PP95]. The adaptive TAP approach proposed by Oppor and Winther [OW01a, OW01b] further allowed for inference in densely connected graphical models without prior knowledge on the weight statistics. The Onsager term of these TAP equations was evaluated using the cavity method for a given weight sample. The resulting equations were then understood to be a particular case of the Expectation Propagation (EP) [Min01] - belonging to the class of message passing algorithms for approximate inference - yet applied in densely connected models [OW05]. An associated approximation of the free energy called Expectation Consistency (EC) was additionally derived from the EP messages. Subsequently, Kabashima and collaborators [SK08, SK09, Kab08] focused on the perceptron and the GLM to propose TAP equations and a replica derivation of the free energy for the ensemble of orthogonally invariant random weight matrices. In the singular value decomposition of such weight matrices, $\underline{W} = \underline{U} \underline{S} \underline{V}^\top \in \mathbb{R}^{M \times N}$, the orthogonal basis matrices \underline{U} and \underline{V} are drawn uniformly at random from respectively $O(M)$ and $O(N)$, while the diagonal matrix of singular values \underline{S} has an arbitrary spectrum. The consistency between the EC free energy and the replica derivation for orthogonally invariant matrices was verified by [KV14] for signal recovery from linear measurements (the GLM without G). From the algorithmic perspective, Fletcher, Rangan and Schniter [RSF17, SRF16] applied the EP to the GLM to obtain the (Generalized) Vector-Approximate Message Passing (G-VAMP) algorithm. Remarkably, these authors proved that the behavior of the algorithm could be characterized in the thermodynamic limit, provided the weight matrix is drawn from the orthogonally invariant ensemble, by a set of scalar State Evolution equations similarly to the AMP algorithm. These equations are again related to the saddle point equations of the replica free energy. Concurrently, Oppor, Cakmak and Winther proposed an alternative procedure for solving TAP equations with orthogonally invariant weight matrices in Ising spin systems relying on an analysis of iterative algorithms [OCW16, QO19]. Finally, [MFC⁺19] revisits the above cited contributions and provides detailed considerations on their connections.

Below we present the aforementioned free energy as proposed by [SK08, SK09, Kab08], and the G-VAMP algorithm of [SRF16].

5.2 Algorithms and free energies beyond i.i.d. matrices

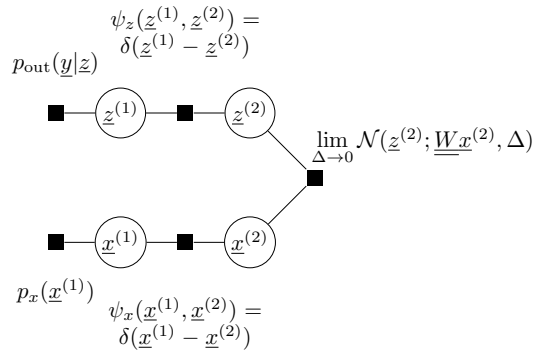


Figure 5: Factor graph representation of the GLM for the derivation of VAMP

5.2.1 Replica free energy for the GLM in the Bayes Optimal setting

Consider the ensemble of orthogonally invariant weight matrices \underline{W} with spectral density $\sum_{i=1}^N \delta(\lambda - \lambda_i)/N$ of their ‘square’ $\underline{W}\underline{W}^\top$ converging in the thermodynamic limit $N \rightarrow +\infty$ to a given density $\rho_\lambda(\lambda)$. The quenched free energy of the GLM in the Bayes optimal setting derived in [Kab08, SK09] writes

$$-f = \text{extr}_{q\hat{q}} \left[-\frac{1}{2}q\hat{q} + \mathcal{I}_x(\hat{q}) + \mathcal{J}_z(q_0, q, \alpha, \rho_\lambda) \right], \quad (131)$$

$$\mathcal{J}_z(q_0, q, \alpha, \rho_\lambda) = \text{extr}_{u\hat{u}} \left[F_{\rho_\lambda, \alpha}(q_0 - q, \hat{u}/\lambda_0) + \frac{\hat{u}q_0}{2} - \frac{\alpha\hat{u}u}{2\lambda_0} + \alpha\mathcal{I}_z(q_0\lambda_0/\alpha, u) \right], \quad (132)$$

where \mathcal{I}_x and \mathcal{I}_z were defined as (126)-(127) and the spectral density $\rho_\lambda(\lambda)$ appears via its mean $\lambda_0 = \mathbb{E}_\lambda[\lambda]$ and in the definition of

$$F_{\rho_\lambda, \alpha}(q, u) = \frac{1}{2} \text{extr}_{\Lambda_q, \Lambda_u} \left[-(\alpha - 1) \log \Lambda_u - \mathbb{E}_\lambda \log(\Lambda_u \Lambda_q + \lambda) + \Lambda_q q + \alpha \Lambda_u u \right] - \frac{1}{2}(\log q + 1) + \frac{\alpha}{2}(\log u + 1). \quad (133)$$

Gaussian random matrices are a particular case of the considered ensemble. Their singular values are characterized asymptotically by the Marcenko-Pastur distribution [MP67]. In this case, one can check that the above expression reduces to (125). More generally, note that \mathcal{J}_z generalizes \mathcal{I}_z .

5.2.2 Vector Approximate Message Passing for the GLM

The VAMP algorithm consists in writing EP [Min01] with Gaussian messages on the factor graph representation of the GLM posterior distribution given in Figure 5. The estimation of the signal \underline{x} is decomposed onto four variables, two duplicates of \underline{x} itself and two duplicates of the linear transformation $\underline{z} = \underline{W}\underline{x}$. The potential functions ψ_x and ψ_z of factors connecting copies of the same variable are Dirac distributions enforcing their equality. The factor node linking $\underline{z}^{(2)}$ and $\underline{x}^{(2)}$ is assumed Gaussian with variance going to zero. The procedure of derivation, equivalent to the projection of the BP algorithm on Gaussian messages, is recalled in Appendix B and leads to Algorithm 2. Like for AMP, the algorithm features some auxiliary variables introduced along the derivation. At convergence the duplicated \hat{x}_1, \hat{x}_2 (and \hat{z}_1, \hat{z}_2) are equal and either can be returned by the algorithm as an estimator. For readability, we omitted the time indices in the iterations that here simply follow the indicated update.

For a given instance of the GLM inference problem, i.e. a given weight matrix \underline{W} , one can always launch either the AMP algorithm or the VAMP algorithm to attempt the reconstruction. If the weight matrix has i.i.d. zero mean Gaussian entries, the two strategies are conjectured to be equivalent and GAMP can be provably convergent for certain settings [RSF14]. If the weight matrix is not Gaussian but orthogonally invariant, then only VAMP is expected to always converge. More generally, even in cases where none of these assumptions are verified, VAMP has been observed to have less convergence issues than AMP.

Like for AMP, a State Evolution can also be derived for VAMP (which was actually directly proposed for the multi-layer GLM [FRS18]). It rigorously characterizes the behavior of the

5 Further extensions of interest for learning

Algorithm 2 Vector Approximate Message Passing

Input: vector of observations $\underline{y} \in \mathbb{R}^M$ and weight matrix $\underline{W} \in \mathbb{R}^{M \times N}$;

Initialize: $\underline{A}_1^x, \underline{B}_1^x, \underline{A}_1^z, \underline{B}_1^z$

repeat

$$\hat{\underline{x}}_1 = f_1^x(B_1^x, \underline{A}_1^x), \quad \underline{C}_{\underline{1}}^x = f_2^x(B_1^x, \underline{A}_1^x) \quad (134)$$

$$\underline{A}_2^x = \underline{C}_{\underline{1}}^{x-1} - \underline{A}_1^x, \quad \underline{B}_2^x = \underline{C}_{\underline{1}}^{x-1} \hat{\underline{x}}_1 - \underline{B}_1^x \quad (135)$$

$$\hat{\underline{z}}_1 = f_1^z(B_1^z, \underline{A}_1^z), \quad \underline{C}_{\underline{1}}^z = f_2^z(B_1^z, \underline{A}_1^z) \quad (136)$$

$$\underline{A}_2^z = \underline{C}_{\underline{1}}^{z-1} - \underline{A}_1^z, \quad \underline{B}_2^z = \underline{C}_{\underline{1}}^{z-1} \hat{\underline{z}}_1 - \underline{B}_1^z \quad (137)$$

$$\hat{\underline{x}}_2 = g_1^x(B_2^x, \underline{A}_2^x, \underline{B}_2^z, \underline{A}_2^z), \quad \underline{C}_{\underline{2}}^x = g_2^x(B_2^x, \underline{A}_2^x, \underline{B}_2^z, \underline{A}_2^z) \quad (138)$$

$$\underline{A}_1^x = \underline{C}_{\underline{2}}^{x-1} - \underline{A}_2^x, \quad \underline{B}_1^x = \underline{C}_{\underline{2}}^{x-1} \hat{\underline{x}}_1 - \underline{B}_2^x \quad (139)$$

$$\hat{\underline{z}}_2 = g_1^z(B_2^z, \underline{A}_2^z, \underline{B}_2^x, \underline{A}_2^x), \quad \underline{C}_{\underline{2}}^z = g_2^z(B_2^z, \underline{A}_2^z, \underline{B}_2^x, \underline{A}_2^x) \quad (140)$$

$$\underline{A}_1^z = \underline{C}_{\underline{2}}^{z-1} - \underline{A}_2^z, \quad \underline{B}_1^z = \underline{C}_{\underline{2}}^{z-1} \hat{\underline{z}}_1 - \underline{B}_2^z \quad (141)$$

until convergence

Output: signal estimate $\hat{\underline{x}}_1 \in \mathbb{R}^N$, and estimated covariance $\underline{C}_{\underline{1}}^x \in \mathbb{R}^{N \times N}$

algorithm when \underline{W} is orthogonally invariant. One can also verify that the SE fixed points can be mapped to the solutions of the saddle point equations of the replica free energy (131) (see Section 1 of Supplementary Material of [GML⁺18]); so that the robust algorithmic procedure can advantageously be used to compute the fixed points to be plugged in the replica potential to approximate the free energy.

5.3 Multi-value AMP

A recent extension of AMP consists in treating the simultaneous reconstruction of multiple signals undergoing the same mixing step from the corresponding multiple observations. This is a situation of particular interest for learning appearing for instance in the teacher-student set-up of committee machines. The authors of [AMB⁺18] showed that the different weight vectors of these neural networks can be inferred from the knowledge of training input-output pairs introducing this extended version of AMP. Here the same matrix of training input data mixes the teacher weight vectors to produce the training output data. For a matter of consistency with the examples used in the previous sections, we here formalize the algorithm for the GLM. Nevertheless this is just a matter of rewriting of the committee algorithm of [AMB⁺18].

Concretely let's consider a GLM with P pairs of signal and observations $\{\underline{x}_{(k)}, \underline{y}_{(k)}\}_{k=1}^P$, gathered in matrices $\underline{X} \in \mathbb{R}^{N \times P}$ and $\underline{Y} \in \mathbb{R}^{M \times P}$. We are interested in the posterior distribution

$$p(\underline{X}|\underline{Y}, \underline{W}) = \frac{1}{Z(\underline{Y}, \underline{W})} \prod_{i=1}^N p(\underline{x}_i) \prod_{\mu=1}^M p_{\text{out}}(\underline{y}_\mu | \underline{w}_\mu^\top \underline{X}), \quad \underline{x}_i \in \mathbb{R}^P, \quad \underline{y}_\mu \in \mathbb{R}^P. \quad (142)$$

Compared to the traditional GLM measure (51), scalar variables are here replaced by vectors in \mathbb{R}^P . In Appendix F we present a derivation starting from BP of the corresponding AMP presented in Algorithm 3. The major difference with the scalar GLM consists in the necessity of tracking covariance matrices between the P different variables instead of simple variances.

This AMP algorithm can also be analyzed by a State Evolution. In [AMB⁺18], the teacher-student matched setting of the committee machine is examined through the replica approach and the Bayes optimal State Evolution equations are obtained as the saddle point equations of the replica free energy. In Appendix F we present the alternative derivation of the State Evolution equations from the message passing and without assuming a priori matched teacher and student, as done in [GBKZ19].

5.4 Model composition and multi-layer inference

Algorithm 3 Approximate Message Passing for multi-value GLM

Input: matrix $\underline{Y} \in \mathbb{R}^{M \times P}$ and matrix $\underline{W} \in \mathbb{R}^{M \times N}$:

Initialize: $\hat{\underline{x}}_i, \underline{C}_{\underline{i}}^x \quad \forall i$ and $\underline{g}_{\text{out}\mu}, \partial_{\omega} \underline{g}_{\text{out}\mu} \quad \forall \mu$

repeat

1) Estimate mean and variance of \underline{z}_μ given current $\hat{\underline{x}}_i$

$$\underline{V}_\mu^{(t)} = \sum_{i=1}^N \frac{W_{\mu i}^2}{N} \underline{C}_{\underline{i}}^{x(t)} \quad (143)$$

$$\underline{\omega}_\mu^{(t)} = \sum_{i=1}^N \frac{W_{\mu i}}{\sqrt{N}} \hat{\underline{x}}_i^{(t)} - \sum_{i=1}^N \frac{W_{\mu i}^2}{N} (\underline{\sigma}_i^{(t)})^{-1} \underline{C}_{\underline{i}}^{x(t)} \underline{\sigma}_i \underline{g}_{\text{out}\mu}^{(t-1)} \quad (144)$$

2) Estimate mean and variance of the gap between optimal \underline{z}_μ and $\underline{\omega}_\mu$ given \underline{y}_μ

$$\partial_{\omega} \underline{g}_{\text{out}\mu}^{(t)} = \partial_{\omega} \underline{g}_{\text{out}}(\underline{y}_\mu, \underline{\omega}_\mu^{(t)}, \underline{V}_\mu^{(t)}) \quad (145)$$

$$\underline{g}_{\text{out}\mu}^{(t)} = \underline{g}_{\text{out}}(\underline{y}_\mu, \underline{\omega}_\mu^{(t)}, \underline{V}_\mu^{(t)}) \quad (146)$$

3) Estimate mean and variance of \underline{x}_i given current optimal \underline{z}_μ

$$\underline{\sigma}_i^{(t)} = \left(- \sum_{\mu=1}^M \frac{W_{\mu i}^2}{N} \partial_{\omega} \underline{g}_{\text{out}\mu}^{(t)} \right)^{-1} \quad (147)$$

$$\underline{\lambda}_i^{(t)} = \hat{\underline{x}}_i^{(t)} + \underline{\sigma}_i^{(t)} \left(\sum_{\mu=1}^M \frac{W_{\mu i}}{\sqrt{N}} \underline{g}_{\text{out}\mu}^{(t)} \right) \quad (148)$$

4) Estimate of mean and variance of \underline{x}_i augmented of the information about the prior

$$\underline{C}_{\underline{i}}^{x(t+1)} = f_{\underline{2}}^x(\underline{\lambda}_i^{(t)}, \underline{\sigma}_i^{(t)}) \quad (149)$$

$$\hat{\underline{x}}_i^{(t+1)} = f_{\underline{1}}^x(\underline{\lambda}_i^{(t)}, \underline{\sigma}_i^{(t)}) \quad (150)$$

until convergence

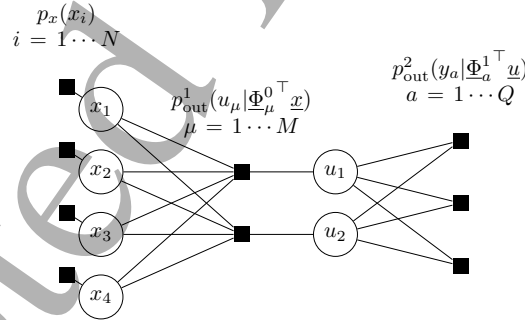


Figure 6: Factor graph representation of a generic 2-layer GLM.

5.4 Model composition and multi-layer inference

Another recent and ongoing direction of extension of mean-field methods is the combination of solutions of elementary models to tackle more sophisticated inference questions. The graphical representations of probability distributions (reintroduced briefly in Appendix B) are here of great help. In a complicated joint probability distribution, it is sometimes possible to identify well-known sub-models, such as the GLM or the RBM. Understanding how and when it is justified to plug-in different solutions is of course non-trivial and a very promising direction of research.

A particularly relevant extension in this direction is the treatment of multi-layer GLMs, or in other words multi-layer neural networks. With depth L , hidden layers noted $\underline{u}^\ell \in \mathbb{R}^{N_\ell}$,

5 Further extensions of interest for learning

and weight matrices $\underline{\Phi}^\ell \in \mathbb{R}^{N_{\ell+1} \times N_\ell}$, it formally corresponds to the statistical model

$$\underline{u}^0 = \underline{x}_0 \sim p_{x_0}(\underline{x}_0), \quad (151)$$

$$\underline{u}^\ell \sim p_{\text{out}}^\ell(\underline{u}^\ell | \underline{\Phi}^{\ell-1} \underline{u}^{\ell-1}) \quad \forall \ell = 1 \dots L-1, \quad (152)$$

$$\underline{y} \sim p_{\text{out}}^L(\underline{y} | \underline{\Phi}^{L-1} \underline{u}^{L-1}). \quad (153)$$

In [MKMZ17] a multi-layer version of AMP is derived, assuming Gaussian i.i.d weight matrices, along with a State Evolution and a replica free energy. Remarkably, the asymptotic replica prediction was mathematically proven correct in [GML⁺18]. In [FRS18], the multi-layer version of the VAMP algorithm is derived with the corresponding State Evolution for orthogonally invariant weight matrices. The matching free energies were obtained independently in [GML⁺18] by the generalization of a replica result and by [RP16] from a different argument.

In the next paragraph we sketch a derivation of the 2-layer AMP presented in Algorithm 4, it provides a good intuition of the composition ability of mean-field inference methods.

Heuristic derivation of 2-layer AMP The derivation of the multi-layer AMP follows identical steps to the derivation of the single layer presented in Section 4.3, yet for a more complicated factor graph and consequently a larger collection of messages. Without conducting the lengthy procedure, one can form an intuition for the resulting algorithm starting from the single-layer AMP. The corresponding factor graph is given on Figure 6. Compared to the single-layer case (see Figure 3), an interface with a set of $M = N_1$ hidden variables u_μ is inserted between the $N = N_0$ signals x_i and the $Q = N_2$ observations y_a . In the neighborhood of the inputs x_i the factor graph is identical to the single-layer and the input functions can be defined from a normalization partition identical to (85),

$$\mathcal{Z}^x(\lambda_i, \sigma_i) = \int dx_i p_x(x_i) e^{-\frac{(x_i - \lambda_i)^2}{2\sigma_i}}, \quad (154)$$

yielding updates (170)-(171). Similarly, the neighborhood of the observations y_a is also unchanged and the updates (160) and (161) are following from the definition of

$$\mathcal{Z}_{\text{out}}^y(\omega_a^2, V_a^2) = \int dz_a p_{\text{out}}^2(y_a | z_a) e^{-\frac{(z_a - \omega_a^2)^2}{2V_a^2}}, \quad (155)$$

identical to the single layer (79). At the interface however, the variables u_μ play the role of outputs for the first GLM and of inputs for the second GLM, which translates into a normalization partition function of mixed form

$$\mathcal{Z}_{\text{out}}^u(\omega_\mu^1, V_\mu^1, \lambda_\mu^1, \sigma_\mu^1) = \int dz_\mu \int du_\mu p_{\text{out}}^1(u_\mu | z_\mu) \times e^{-\frac{(u_\mu - \lambda_\mu^1)^2}{2\sigma_\mu^1}} e^{-\frac{(z_\mu - \omega_\mu^1)^2}{2V_\mu^1}}.$$

Updates (166) and (167) are obtained by considering that the second layer acts as an effective channel for the first layer, i.e. from the normalization interpreted as

$$\mathcal{Z}_{\text{out}}^u(\omega_\mu^1, V_\mu^1, \lambda_\mu^1, \sigma_\mu^1) = \int dz_\mu p_{\text{out}}^{\text{eff}}(z_\mu) e^{-\frac{(z_\mu - \omega_\mu^1)^2}{2V_\mu^1}}. \quad (156)$$

Finally, update equations (172) and (173) are in turn derived considering the first layer defines an effective prior for the hidden variables and rewriting the normalization as

$$\mathcal{Z}_{\text{out}}^u = \int du_\mu p_u^{\text{eff}}(u_\mu) e^{-\frac{(u_\mu - \lambda_\mu^1)^2}{2\sigma_\mu^1}}. \quad (157)$$

The rest of the algorithm updates follows as usual from the self-consistency between the different variables introduced as they correspond to different parametrization of the same marginals. The schedule of updates and the time indexing reported in Algorithm 4 results from the entire derivation starting from the BP messages. The generalization of the algorithm to an arbitrary number of layers is easily obtained repeating the heuristic arguments presented here.

5.4 Model composition and multi-layer inference

Algorithm 4 Generalized Approximate Message Passing for the 2-layer GLM

Input: vector $\underline{y} \in \mathbb{R}^M$ and matrices $\underline{\Phi}^0 \in \mathbb{R}^{M \times N}$, $\underline{\Phi}^1 \in \mathbb{R}^{Q \times M}$:

Initialize: $\hat{\underline{x}} \in \mathbb{R}^N$, $\underline{C}^x \in \mathbb{R}^N$, $\hat{\underline{u}} \in \mathbb{R}^M$, $\underline{C}^u \in \mathbb{R}^M$, $\underline{g}_{\text{out}}^1 \in \mathbb{R}^M$, $\partial \underline{g}_{\text{out}}^1 \in \mathbb{R}^M$, $\underline{g}_{\text{out}}^2 \in \mathbb{R}^Q$, $\partial \underline{g}_{\text{out}}^2 \in \mathbb{R}^Q$ and $t = 0$.

repeat

1) Update auxiliary variables of second layer:

$$\omega_a^{2(t)} = \sum_{\mu} \frac{\Phi_{a\mu}^1}{\sqrt{N}} \hat{u}_{\mu}^{(t)} - \sum_{\mu} \frac{(\Phi_{a\mu}^1)^2}{N} C_{\mu}^{ru(t)} g_{\text{out}a}^{2(t-1)} \quad (158)$$

$$V_a^{2(t)} = \sum_{\mu} \frac{(\Phi_{a\mu}^1)^2}{N} C_{\mu}^{ru(t)} \quad (159)$$

$$g_{\text{out}a}^{2(t)} = g_{\text{out}}^2(\underline{y}_a, \omega_a^{2(t)}, V_a^{2(t)}) \quad (160)$$

$$\partial_{\omega} g_{\text{out}a}^{2(t)} = \partial_{\omega} g_{\text{out}}^2(\underline{y}_a, \omega_a^{2(t)}, V_a^{2(t)}) \quad (161)$$

$$\lambda_{\mu}^{1(t)} = \sigma_{\mu}^1 \left(\sum_a \frac{\Phi_{a\mu}^1}{\sqrt{N}} g_{\text{out}a}^{2(t)} - \frac{(\Phi_{a\mu}^1)^2}{N} \partial_{\omega} g_{\text{out}a}^{2(t)} \hat{u}_{\mu} \right) \quad (162)$$

$$\sigma_{\mu}^{1(t)} = \left(- \sum_a \frac{(\Phi_{a\mu}^1)^2}{N} \partial_{\omega} g_{\text{out}a}^{2(t)} \right)^{-1} \quad (163)$$

1) Update auxiliary variables of first layer:

$$\omega_{\mu}^{1(t)} = \sum_i \frac{\Phi_{\mu i}^0}{\sqrt{N}} \hat{x}_i^{(t)} - \sum_i \frac{(\Phi_{\mu i}^0)^2}{N} C_{xi} g_{\text{out}\mu}^{1(t-1)} \quad (164)$$

$$V_{\mu}^{1(t)} = \sum_i \frac{(\Phi_{\mu i}^0)^2}{N} C_{xi} \quad (165)$$

$$g_{\text{out}\mu}^{1(t)} = g_{\text{out}}^1(\omega_{\mu}^{1(t)}, V_{\mu}^{1(t)}, \lambda_{\mu}^{1(t)}, \sigma_{\mu}^{1(t)}) \quad (166)$$

$$\partial_{\omega} g_{\text{out}\mu}^{1(t)} = \partial_{\omega} g_{\text{out}}^1(\omega_{\mu}^{1(t)}, V_{\mu}^{1(t)}, \lambda_{\mu}^{1(t)}, \sigma_{\mu}^{1(t)}) \quad (167)$$

$$\sigma_i^{0(t)} = \left(- \sum_{\mu} \frac{(\Phi_{\mu i}^0)^2}{N} \partial_{\omega} g_{\text{out}\mu}^{1(t)} \right)^{-1} \quad (168)$$

$$\lambda_i^0 = \sigma_i^0 \left(\sum_{\mu} \frac{\Phi_{\mu i}^0}{\sqrt{N}} g_{\text{out}\mu}^{1(t)} - \frac{(\Phi_{\mu i}^0)^2}{N} \partial_{\omega} g_{\text{out}\mu}^{1(t)} \hat{x}_i \right) \quad (169)$$

3) Update means and variances of variables of both layers, \underline{x} and \underline{u} :

$$\hat{x}_i^{(t+1)} = f_1^x(\lambda_i^{0(t)}, \sigma_i^{0(t)}) \quad (170)$$

$$C_{xi}^{x(t+1)} = f_2^x(\lambda_i^{0(t)}, \sigma_i^{0(t)}) \quad (171)$$

$$\hat{u}_{\mu}^{(t+1)} = f_1^u(\omega_{\mu}^{1(t)}, V_{\mu}^{1(t)}, \lambda_{\mu}^{1(t)}, \sigma_{\mu}^{1(t)}) \quad (172)$$

$$C_{\mu}^{ru(t+1)} = f_2^u(\omega_{\mu}^{1(t)}, V_{\mu}^{1(t)}, \lambda_{\mu}^{1(t)}, \sigma_{\mu}^{1(t)}) \quad (173)$$

$t = t + 1$

until convergence

Output: signal estimate $\hat{\underline{x}} \in \mathbb{R}^N$, and estimated variance $\underline{C}^x \in \mathbb{R}^N$

6 Some applications

6.1 A brief pre-deep learning history

The application of mean-field methods of inference to machine learning, and in particular to neural networks, already have a long history and significant contributions to their records. Here we briefly review some historical connections anterior to the deep learning revival of neural networks in the 2010s.

Statistical mechanics of learning In the 80s and 90s, a series of works pioneered the analysis of learning with neural networks through the statistical physics lense. By focusing on simple models with simple data distributions, and relying on the mean-field method of replicas, these papers managed to predict quantitatively important properties such as *capacities*: the number of training data point that could be memorized by a model, or *learning curves*: the generalization error (or population risk) of a model as a function of the size of the training set. This effort was initiated by the study of the Hopfield model [AGS85], an undirected neural network providing associative memory [Hop82]. The analysis of feed forward networks with simple architectures followed (among which [Gar87, Gar88, OH91, MZ95, OW96, MZ04], see also the reviews [SST92, WRB93, Opp95, Saa99, EV01]). The dynamics of simple learning problems was also analyzed through a mean-field framework (not covered in the previous sections) initially in the simplifying case of online learning with infinite training set [SS95a, SS95b, BS95, Saa99] but also with finite data [SB97, LM99].

Physicists, accustomed to studying natural phenomena, fruitfully brought the tradition of modelling to their investigation of learning, which translated into assumptions of random data distributions or teacher-student scenarios. Their approach was in contrast to the focus of the machine learning theorists on worst case guarantees: bounds for an hypothesis class that hold for any data distribution (e.g. Vapnik-Chervonenkis dimension and Rademacher complexity). The originality of the physicists approach, along with the heuristic character of the derivations of mean-field approximations, may nonetheless explain the minor impact of their theoretical contributions in the machine learning community at the time.

Mean-field algorithms for practitioners Along with these contributions to the statistical mechanics theory of learning, new practical training algorithms based on mean-field approximations were also proposed at the same period (see e.g. [Won95, OW96, Won97]). Yet, before the deep learning era, mean-field methods probably had a greater influence in the practice of unsupervised learning through density estimation, where we saw that approximate inference is almost always necessary. In particular the simplest method of naive mean-field, our first example in Section 4, was easily adopted and even extended by statisticians (see e.g. [WJ08] for a recent textbook and [BKM17] for a recent review). The belief propagation algorithm is another example of a well known mean-field methods by machine learners, as it was actually discovered in both communities. Yet, for both methods, early applications rarely involved neural networks and rather relied on simple probabilistic models such as mixtures of elementary distributions. They also did not take full advantage of the latest simultaneous developments in statistical physics of the mean-field theory of disordered systems.

Transferring advanced mean-field methods In this context, the inverse Ising problem has been a notable exception. The underlying question, rooted in theoretical statistical physics, is to infer the parameters of an Ising model given a set of equilibrium configurations. This is related to the unsupervised learning of the parameters of a Boltzmann machine (without hidden units) in the machine learning jargon, while it does not necessarily rely on a maximum likelihood estimation using gradients. The corresponding Boltzmann distribution, with pairwise interactions, is remarkable, not only to physicists. It is the least biased model under the assumption of fixed first and second moments in the sense that it maximizes the entropy. For this problem, physicists proposed dedicated developments of advanced mean-field methods for applications in other fields, and in particular in bio-physics (see [NZB17] for a recent review). A few works even considered the case of Boltzmann machines with hidden units, more common in the machine learning community [PA87, Gal93].

Beyond the specific case of Boltzmann machines, the language barrier between communities is undoubtedly a significant hurdle delaying the global transfer of developments in one field to the other. In machine learning, the potential of the most recent progress of mean-field approximations was advocated for in a pioneering workshop mixing communities in 1999

6.2 Some current directions of research

[OS01]. Yet the first widely-used application is possibly the Approximate Message Passing (AMP) algorithm for compressed sensing in 2009 [DMM09]. Meanwhile, in the different field of Constraint Satisfaction Problems (CSPs), there have been much tighter connections between developments in statistical physics and algorithmic solutions. The very first popular application of advanced mean-field methods outside of physics, beyond naive mean-field and belief propagation, is probably the survey propagation algorithm [MPZ02] in 2002. It borrows from the 1RSB cavity method (not treated in the present paper) to solve efficiently certain types of CSPs.

6.2 Some current directions of research

The great leap forward in the performance of machine learning with neural networks brought by deep learning algorithms, along with the multitude of theoretical and practical challenges it has opened, has re-ignited the interest of physicists for the theory of neural networks. In this Section, far from being exhaustive, we review some current directions of research leveraging mean-field approximations. Another relevant review is [CCC⁺19], which provides references both for machine learning research helped by physics methods and conversely research in physics using machine learning.

Works presented below do not necessarily implement one of the classical inference methods presented in Sections 4 and 5. In some cases, the mean-field limit corresponds to some asymptotic setting where the problem simplifies: typically some correlations weaken, fluctuations are averaged out by concentration effects and, as a result, ad-hoc methods of resolution can be designed. Thus, in the following contributions, different assumptions are considered to serve different objectives. For instance some take an infinite size limit, some assume random (instead of learned) weights or vanishing learning rates. Hence, there is no such a thing as one mean-field theory of deep neural networks. The below cited works are rather complementary pieces of solving a great puzzle.

6.2.1 Neural networks for unsupervised learning

Fundamental study of learning Given their similarity with the Ising model, Restricted Boltzmann Machines have unsurprisingly attracted a lot of interest. Studying an ensemble of RBMs with random parameters using the replica method, Toubiana and Monasson [TM17] evidenced different regimes of typical pattern of activations in the hidden units and identified control parameters as the sparsity of the weights, their strength (playing the role of an effective temperature) or the type of prior for the hidden layer. Their study contributes to the understanding of the conditions under which the RBMs can represent high-order correlations between input units, albeit without including data and learning in their model. Barra and collaborators [BGST17, BGST18], exploited the connections between the Hopfield model and RBMs to characterize RBM learning understood as an associative memory. Relying again on replica calculations, they characterize the retrieval phase of RBMs. Mézard [Mé17] also re-examined retrieval in the Hopfield model using its RBM representation and message passing, showing in particular that the addition of correlations between memorized patterns could still allow for a mean-field treatment at the price of a supplementary hidden layer in the Boltzmann Machine representation. This result remarkably draws a theoretical link between correlations in the training data and the necessity of depth in neural network models.

While the above results do not include characterization of the learning driven by data, a few others were able to discuss the dynamics of training. Huang [Hua17] studied with the replica method and TAP equations the Bayesian learning of a RBM with a single hidden unit and binary weights. Barra and collaborators [BGST17] empirically studied a teacher-student scenario of unsupervised learning by maximum likelihood on samples of an Hopfield model which they could compare to their theoretical characterization of the retrieval phase. Decelle and collaborators [DFF17, DFF18] introduced an ensemble of RBMs characterized by the spectral properties of the weight matrix and derived the typical dynamics of the corresponding order parameters during learning driven by data. Beyond RBMs, analyses of the learning in other generative models are starting to appear [WHL18].

Training algorithm based on mean-field methods Beyond bringing theoretical insights, mean-field methods are also found useful to build tractable estimators of the likelihood in generative models, which in turn serves to design novel training algorithms.

6 Some applications

For Boltzmann machines, this direction was already investigated in the 80s and 90s, [PA87, Hin89, Gal93, KR98], albeit in small models with binary units and for artificial data sets very different from modern machine learning benchmarks. More recently, a deterministic training based on naive mean-field was tested on RBMs [WH02, Tie08]. On toy deep learning data sets, the algorithm was found to perform poorly when compared to both CD and PCD, the commonly employed approximate Monte Carlo methods. However going beyond naive mean-field, considering the second order TAP expansion, allows to bridge the gap in efficiency [GTK15, TGM⁺18]. Additionally, the deterministic mean-field framework offers a tractable way of evaluating the learning success by exploiting the mean-field observables to visualize the representation learned by RBMs. Interestingly, high temperature expansions of estimators different from the maximum likelihood have also been recently proposed as efficient inference method for the inverse Ising problem [LVMC18].

By construction, variational auto-encoders (VAEs) rely on a variational approximation of the likelihood. In practice, the posterior distribution of the latent representation given an input (see Section 2.2) is typically approached by a factorized Gaussian distribution with mean and variance parametrized by neural networks. The factorization assumption relates the method to a naive mean-field approximation.

Structured Bayesian priors With the progress of unsupervised learning, the idea of using generative models as expressive priors has emerged.

For reconstruction tasks, in the event where a set of typical signals is available a priori, the latter can serve as a training set to learn a model of the underlying data distribution with a generative model. Subsequently, in the reconstruction of a new signal of the same type, the generative model can serve as a Bayesian prior. In particular, the idea to exploit RBMs in CS applications was pioneered by [DHD12] and [TDK16], who trained binary RBMs using Contrastive Divergence (to locate the support of the non-zero entries of sparse signals) and combined it with an AMP reconstruction. They demonstrated drastic improvements in the reconstruction with structured learned priors compared to the usual sparse unstructured priors. The approach, requiring to combine the AMP reconstruction for CS and the RBM TAP inference, was further generalized in [TMC⁺16, TGM⁺18] to real valued distributions. In the line of these applications, several works have also investigated using feed forward generative models for inference tasks. Using this time multi-layer VAMP inference, Rangan and co-authors [PSRF19] showed that VAEs could help for in-painting partially observed images. Note also that a different line of works, mainly considering GANs, examined the same type of applications without resorting to mean-field algorithms [BJPD17, HLV18, HV18, MV18]. Instead they performed the inference via gradient descent and back-propagation.

Another application of generative priors is to model synthetic data sets with structure. In [GML⁺18, ALM⁺19, GMKZ19], the authors designed learning problems amenable to a mean-field theoretical treatment by assuming the inputs to be drawn from a generative prior (albeit with untrained weights so far). This approach goes beyond the vanilla teacher-student scenario where input data is typically unstructured with i.i.d. components. This is a crucial direction of research as the role of structure in data appears as an important component to understand the puzzling efficiency of deep learning.

6.2.2 Neural networks for supervised learning

New results in the replica analysis of learning The classical replica analysis of learning with simple architectures, following bases set by Gardner and Derrida 30 years ago, continues to be explored. Among the most prominent results, Kabashima and collaborators [Kab08, SK08, SK09] extended the mean-field treatment of the perceptron from data matrices with i.i.d entries to random orthogonal matrices. It is a much larger class of random matrices where matrix entries can be correlated. More recently, a series of works explored in depth the specific case of the perceptron with binary weight values for classification on random inputs. Replica computations showed that the space of solutions is dominated in the thermodynamic limit by isolated solutions [HWK13, HK14], but also that subdominant dense clusters of solutions exist with good generalization properties in the teacher-student scenario case [BIL⁺15, BBC⁺16, BGK⁺18]. This observation inspired a novel training algorithm [CCS⁺17]. The simple two-layer architecture of the committee machine was also reexamined recently [AMB⁺18]. In the teacher-student scenario, a computationally hard phase of learning was evidenced by comparing a message passing algorithm (believed to be optimal) and the

replica prediction. In this work, the authors also proposed a strategy of proof of the replica prediction.

Signal propagation in depth Mean-field approximations can also help understand the role and implications of depth by characterizing signal propagation in neural networks. The following papers consider the width of each layer to go to infinity. In this limit, Sompolinsky and collaborators characterized how neural networks manage to progressively separate data manifolds fed as inputs [KS16, CLS18, CCLS19]. Another line of works focused on the initialization of neural networks (i.e. with random weights), and found an order-to-chaos transition in the signal propagation as a function of hyperparameters of training [PLR⁺16, SBG⁺17]. As a result, the authors could formulate recommendations for combinations of hyperparameters to practitioners. This type of analysis could furthermore be generalized to convolutional networks [NXL⁺19], recurrent networks [GCC⁺19] and networks with batch-normalization regularization [YPR⁺19]. The space of functions spanned by deep random networks in the infinite-size limit was also studied by [LS18, LS20], using the different but related approach of the generating functional analysis. Yet another mean-field argument, this time relying on a replica computation, allowed to compute the mutual information between layers of large non-linear deep neural networks with orthogonally invariant weight matrices [GML⁺18]. Using this method, mutual informations can be followed precisely along the learning for an appropriate teacher-student scenario. The strategy offers an experimental test bed to characterize possible links between the generalization ability of deep neural networks and information compression phases in the training (see [TZ15, SZT17, SBD⁺18]).

Dynamics of SGD learning in simple networks and generalization A number of different mean-field limits led to interesting analyses of the dynamics of gradient descent learning. In particular, the below mentioned works contribute to shed light on the generalization power of neural networks in the so-called overparametrized regimes, that is where the number of parameters exceeds largely either the number of training points or the underlying degrees of freedom of the teacher rule. In linear networks first, an exact description in the high-dimensional limit was obtained for the teacher-student setup by [AS17] using random matrix theory. The generalization was predicted to improve with the overparametrization of the student. Non-linear networks with one infinitely wide hidden layer were considered by [MMN18, RVE18, CB18b, SS18] who showed that gradient descent converges to a finite generalization error. Their results are related to others obtained in a slightly different limit of infinitely large neural networks [JGH18]. For arbitrarily deep networks, Jacot and collaborators [JGH18] showed that, in a certain setting, gradient descent was effectively performing a kernel regression with a kernel function converging to a fixed value for the entire training as the size of the layers increases. In both related limits, the absence of divergence is accounting for generalization not deteriorating despite of the explosion of the number of parameters. The relationship between the two above limits was discussed in [CB18a, MMM19, GSJW19]. Subsequent works, leveraged the formalism introduced in [JGH18]. Scaling for the generalization error as a function of network sizes were derived by [GJS⁺19]. Other authors focused on the characterization of the network output function in this limit, which takes the form of a Gaussian process [LXS⁺19]. This fact was probably first noticed by Oppor and Winther with one hidden layer [OW99b], to whom it inspired a TAP based Bayesian classification method using Gaussian processes. Finally, yet another limit was analyzed by [GAS⁺19], considering a finite number of hidden units with an infinitely wide input. Following classical works on the mean-field analysis of online learning (not covered in the previous sections [SS95a, SS95b, BS95, Saa99]), a closed set of equations can be derived and analyzed for a collection of overlaps. Note that these are the same order parameters as in replica computations. The resulting learning curves evidence the necessity of multi-layer learning to observe the improvement of generalization with overparametrization. An interplay between optimization, architecture and data sets seems necessary to explain the phenomenon.

7 Conclusion

This review aimed at presenting in a pedagogical way a selection of inference methods coming from statistical physics. In past and current lines of research that were also reviewed, these methods are sometimes turned into practical and efficient inference algorithms, or sometimes the angle stone in theoretical computations.

A Index of notations and abbreviations

What is missing There are more of these methods beyond what was covered here. In particular the cavity method [MPV86], closely related to message passing algorithms and the replica formalism, played a crucial role in the physics of spin glasses. Note also that we assumed replica symmetry, which is only guaranteed to be correct in the Bayes optimal case. References of introductions to replica symmetry breaking are [MPV86, CCFM05], and newly proposed message passing algorithms with RSB are [SLL19, AFUZ19, AKUZ19]. The methods of analysis of online learning algorithms pioneered by [SS95a, SS95b, BS95] and reviewed in [Saa99] also deserve the name of classical mean-field analysis. They are currently actively serving research efforts in deep learning theory [GAS⁺19]. Another important method is the off-equilibrium mean-field theory [CS88, CHS93, CK93], recently used for example to characterize a specific type of neural networks called graph neural networks [KTQ18] or to study properties of gradient flows [MKUZ19].

On the edge of validity We have also touched upon the limitations of the mean-field approach. To start with, the thermodynamic limit is ignoring finite-size effects. Moreover, different ways of taking the thermodynamic limit for the same problem sometimes lead to different results. Also, necessary assumptions of randomness for weights or data matrices are sometimes in clear contrast with real applications.

Thus, the temptation to apply abusively results from one field to the other can be a dangerous pitfall of the interdisciplinary approach. We could mention here the characterization of the dynamics of optimization. While physicists have extensively studied Langevin dynamics with Gaussian white noise, the continuous time limit of SGD is unfortunately not an equivalent in the general case. While some works attempt to draw insights from this analogy using strong assumptions (e.g. [CHM⁺15, JKA⁺17]), others seek precisely to understand the differences between the two dynamics in neural networks optimization (e.g. [BJSG⁺18, SSG19]). Alternatively, another good reason to consider the power of mean-field methods lies in the observation rooted in the tradition of theoretical physics that one can learn from models a priori far from the exact neural networks desired, but that retain some key properties, while being amenable to theoretical characterization. For example, [MKUZ19] studied a high-dimensional non-convex optimization problem inspired by the physics of spin glasses apparently unrelated to neural networks, but gained insights on the dynamics of gradient descent (and Langevin) that is of primal interest. Another example of this surely promising approach is [WHL18], who built and analyzed a minimal model of GANs.

Moreover, the possibility to combine well-studied simple settings to obtain a mean-field theory for more complex models, as recently demonstrated in a series of work [TDK16, TMC⁺16, TGM⁺18, MKMZ17, FRS18, GML⁺18, ALM⁺19], constitutes an exciting direction of research that should broaden considerably the limit of applications of mean-field methods.

Patching the pieces together and going further Thus the mean-field approach alone cannot to this day provide complete answers to the still numerous puzzles on the way towards a deep learning theory. Yet, considering different limits and special cases, combining solutions to approach ever more complex models, the approach should help uncover more and more corners of the big black box. Hopefully, intuition gained at the edge will help revealing the broader picture.

Acknowledgements

This paper is based on the introductory chapters of my PhD dissertation, written under the supervision of Florent Krzakala and in collaboration with Lenka Zdeborová, to whom both I am very grateful. I would like to thank also Benjamin Aubin, Cédric Gerbelot, Adrien Laversanne-Finot, and Guilhem Semerjian for their comments on the manuscript. I gratefully acknowledge the support of the ‘Chaire de recherche sur les modèles et sciences des données’ by Fondation CFM pour la Recherche-ENS and of Fondation L’Oréal For Women In Science. I also thank the Kalvi Institute for Theoretical Physics, where part of this work was written.

A Index of notations and abbreviations

$[N]$ - Set of integers from 1 to N

$\delta(\cdot)$ - Dirac distribution

$\sigma(x) = (1 + e^{-x})^{-1}$ - Sigmoid

$\text{relu}(x) = \max(0, x)$ - Rectified Linear Unit

\underline{X} - Matrix
 \underline{x} - Vector
 $\underline{I}_N \in \mathbb{R}^{N \times N}$ - Identity matrix
 $\langle \cdot \rangle$ - Average with respect to the Boltzmann distribution
 $O(N) \subset \mathbb{R}^{N \times N}$ - Orthogonal ensemble
1RSB - 1 Step Replica Symmetry Breaking
AMP - Approximate message passing
BP - Belief Propagation
cal-AMP - Calibration Approximate Message Passing
CD - Contrastive Divergence
CS - Compressed Sensing
CSP - Constrain Satisfaction Problem
DAG - Directed Acyclic Graph
DBM - Deep Boltzmann Machine
EC - Expectation Consistency
EP - Expectation Propagation
GAMP - Generalized Approximate message passing
GAN - Generative Adversarial Networks
GD - Gradient Descent
GLM - Generalized Linear Model
G-VAMP - Generalized Vector Approximate Message Passing
i.i.d. - independent identically distributed
PCD - Persistent Contrastive Divergence
r-BP - relaxed Belief Propagation
RS - Replica Symmetric
RSB - Replica Symmetry Breaking
RBM - Restricted Boltzmann Machine
SE - State Evolution
SGD - Stochastic Gradient Descent
SK - Sherrington-Kirkpatrick
TAP - Thouless Anderson Palmer
VAE - Variational Autoencoder
VAMP - Vector Approximate message passing

B Statistical models representations

Graphical representations have been developed to represent and efficiently exploit (in)dependencies between random variables encoded in joint probability distributions. They are useful tools to concisely present the model under scrutiny, as well as direct supports for some derivations of inference procedures. Let us briefly present two types of graphical representations.

Probabilistic graphical models Formally, a probabilistic graphical model is a graph $G = (V, E)$ with nodes V representing random variables and edges E representing direct interactions between random variables. In many statistical models of interest, it is not necessary to keep track of all the possible combinations of realizations of the variables as the joint probability distribution can be broken up into factors involving only subsets of the variables. The structure of the connections E reflects this factorization.

There are two main types of probabilistic graphical models: *directed graphical models* (or Bayesian networks) and *undirected graphical models* (or Markov Random Fields). They allow to represent different independencies and factorizations. In the next paragraphs we provide intuitions and remind some useful properties of graphical models, a good reference to understand all the facets of this powerful tool is [KF09].

B Statistical models representations

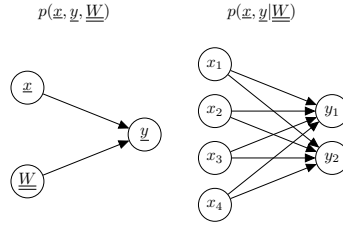


Figure 7: Left: Directed graphical model for $p(\underline{x}, y, \underline{W})$ without assumptions of factorizations for the channel and priors. Right: Directed graphical model reflecting factorization assumptions for $p(\underline{x}, y|\underline{W})$.

Undirected graphical models In undirected graphical models the direct interaction between a subset of variables $C \subset V$ is represented by undirected edges interconnecting each pair in C . This fully connected subgraph is called a *clique* and associated with a real positive *potential* function ψ_C over the variable $\underline{x}_C = \{x_i\}_{i \in C}$ carried by C . The joint distribution over all the variables carried by V , \underline{x}_V is the normalized product of all potentials

$$p(\underline{x}_V) = \frac{1}{Z} \prod_{C \in \mathcal{C}} \psi_C(\underline{x}_C). \quad (174)$$

Example (i): the Restricted Boltzmann Machine,

$$p(\underline{x}, \underline{t}) = \frac{1}{Z} e^{\underline{x}^\top \underline{W} \underline{t}} p_x(\underline{x}) p_t(\underline{t}) \quad (175)$$

with factorized p_x and p_t is handily represented using an undirected graphical model depicted in Figure 2a. The corresponding set of cliques is the set of all the pairs with one input unit (indexed by $i = 1 \dots N$) and one hidden unit (indexed by $\alpha = 1 \dots M$), joined with the set of all single units. The potential functions are immediately recovered from (175),

$$\mathcal{C} = \{\{i\}, \{\alpha\}, \{i, \alpha\}; i = 1 \dots N, \alpha = 1 \dots M\}, \quad \psi_{i\alpha}(x_i, t_\alpha) = e^{x_i W_{i\alpha} t_\alpha}, \quad (176)$$

$$p(\underline{x}, \underline{t}) = \frac{1}{Z} \prod_{\{i, \alpha\} \in \mathcal{C}} \psi_{i\alpha}(x_i, t_\alpha) \prod_{i=1}^N p_x(x_i) \prod_{\alpha=1}^M p_t(t_\alpha). \quad (177)$$

It belongs to the subclass of *pairwise* undirected graphical models for which the size of the cliques is at most two.

Undirected graphical models handily encode *conditional independencies*. Let $A, B, S \subset V$ be three disjoint subsets of nodes of G . A and B are said to be independent given S if $p(A, B|S) = p(A|S)p(B|S)$. In the graph representation it corresponds to cases where S separates A and B : there is no path between any node in A and any node in B that is not going through S .

Example (i): In the RBM, hidden units are independent given the inputs, and conversely:

$$p(\underline{t}|\underline{x}) = \prod_{\alpha=1}^M p(t_\alpha|\underline{x}), \quad p(\underline{x}|\underline{t}) = \prod_{i=1}^N p(x_i|\underline{t}). \quad (178)$$

This property is easily spotted by noticing that the graphical model (Figure 2a) is *bipartite*.

Directed graphical model A directed graphical model uses a Directed Acyclic Graph (DAG), specifying directed edges E between the random variables V . It induces an ordering of the random variables and in particular the notion of parent nodes $\pi_i \subset V$ of any given vertex $i \in V$: the set of vertices j such that $j \rightarrow i \in E$. The overall joint probability distribution factorizes as

$$p(\underline{x}) = \prod_{i \in V} p(x_i|\underline{x}_{\pi_i}). \quad (179)$$

Example (ii): The stochastic single layer feed forward network $\underline{y} = g(\underline{W}\underline{x}; \epsilon)$, where $g(\cdot; \epsilon)$ is a function applied component-wise including a stochastic noise ϵ that is equivalent to a conditional distribution $p_{\text{out}}(\underline{y}|\underline{W}\underline{x})$, and where inputs and weights are respectively drawn from distributions $p_x(\underline{x})$ and $p_W(\underline{W})$, has a joint probability distribution

$$p(\underline{y}, \underline{x}, \underline{W}) = p_{\text{out}}(\underline{y}|\underline{W}\underline{x})p_x(\underline{x})p_W(\underline{W}), \quad (180)$$

precisely following such a factorization. It can be represented with a three-node DAG as in Figure 7. Here we applied the definition at the level of vector/matrix valued random variables. By further assuming that p_{out} , p_W and p_x factorize over their components, we keep a factorization compatible with a DAG representation

$$p(\underline{y}, \underline{x}, \underline{W}) = \prod_{i=1}^N p_x(x_i) \prod_{\mu=1}^M p_{\text{out}}(y_\mu | \sum_{i=1}^N W_{\mu i} x_i) \prod_{\mu,i} p_W(W_{\mu i}). \quad (181)$$

For the purpose of reasoning it may be sometimes necessary to get to the finest level of decomposition, while sometimes the coarse grained level is sufficient.

While a statistical physicist may have never been introduced to the formal definitions of graphical models, she inevitably already has drawn a few - for instance when considering the Ising model. She also certainly found them useful to guide physical intuitions. The following second form of graphical representation is probably newer to her.

Factor graph representations Alternatively, high-dimensional joint distributions can be represented with factor graphs, that are undirected bipartite graphs $G = (V, F, E)$ with two subsets of nodes. The variable nodes V representing the random variables as in the previous section (circles in the representation) and the factor nodes F representing the interactions (squares in the representation) associated with potentials. The edge $(i\mu)$ between a variable node i and a factor node μ exists if the variable i participates in the interaction μ . We note ∂i the set of factor nodes in which variable i is involved, they are the neighbors of i in G . Equivalently we note $\partial \mu$ the neighbors of factor μ in G , they carry the arguments $\{x_i\}_{i \in \partial \mu}$, shortened as $\underline{x}_{\partial \mu}$, of the potential ψ_μ . The overall distributions is recovered in the factorized form:

$$p(\underline{x}) = \frac{1}{Z} \prod_{\mu=1}^M \psi_\mu(\underline{x}_{\partial \mu}). \quad (182)$$

Compared to an undirected graphical model, the cliques are represented by the introduction of factor nodes.

Examples: The factor-graph representation of the RBM (i) is not much more informative than the pairwise undirected graph (see Figure 2a). For the feed forward neural networks (ii) we draw the factor graph of $p(\underline{y}, \underline{x}|\underline{W})$ (see Figure 2b).

C Mean-field identity

We derive the exact identity (42) for fully connected Ising model with binary spins $\underline{x} \in \{0, 1\}^N$,

$$\begin{aligned} \langle x_i \rangle_p &= \frac{1}{Z} \sum_{\underline{x} \in \{0,1\}^N} x_i \exp \left(\beta \sum_j b_j x_j + \frac{1}{2} \sum_{jk} W_{jk} x_j x_k \right) \\ &= \frac{1}{Z} \sum_{\underline{x}_{\setminus i} \in \{0,1\}^{N-1}} \exp \left(\beta \sum_{j \neq i} b_j x_j + \frac{1}{2} \sum_{\substack{k \neq i \\ j \neq i}} W_{ij} x_i x_j \right) \sum_{x_i \in \{0,1\}} x_i e^{\beta b_i x_i + \sum_j W_{ij} x_i x_j} \end{aligned} \quad (183)$$

where $\underline{x}_{\setminus i}$ is the vector of \underline{x} without its i -th component. Yet

$$\sigma(\beta b_i + \sum_{j \in \partial i} \beta W_{ij} x_j) = \frac{\sum_{x_i \in \{0,1\}} x_i e^{\beta b_i x_i + \sum_j W_{ij} x_i x_j}}{\sum_{x_i \in \{0,1\}} e^{\beta b_i x_i + \sum_j W_{ij} x_i x_j}}, \quad (184)$$

D Georges-Yedidia expansion for generalized Boltzmann machines

so that multiplying and dividing (183) by the denominator above we obtain the identity (42) in Section 4.1

$$\langle x_i \rangle_p = \langle \sigma(\beta b_i + \sum_{j \in \partial i} \beta W_{ij} x_j) \rangle_p. \quad (185)$$

D Georges-Yedidia expansion for generalized Boltzmann machines

We here present a derivation of the Georges-Yedidia for real-valued degrees of freedom on the example of a Boltzmann machine as in [TGM⁺18]. Formally we consider $\underline{x} \in \mathbb{R}^N$ governed by the energy function and parametrized distribution

$$E(\underline{x}) = - \sum_{(ij)} W_{ij} x_i x_j - \frac{1}{\beta} \sum_{i=1}^N \log p_x(x_i; \theta_i), \quad p(\underline{x}) = \frac{1}{\mathcal{Z}} e^{\frac{\beta}{2} \underline{x}^\top \underline{W} \underline{x}} \prod_{i=1}^N p_x(x_i; \theta_i), \quad (186)$$

where $p_x(x_i; \theta_i)$ is an arbitrary prior distribution with parameter θ_i . For a Bernoulli prior with parameter $\sigma(\beta b_i)$ we recover the measure of binary Boltzmann machines. However we choose here a prior that does not depend on the temperature a priori. We now derive the expansion for this general case following the outline discussed in 4.2.1, and highlighting the differences with the binary case.

Note that inference in the generalized fully connected Boltzmann machine is somehow related to the symmetric rank-1 matrix factorization problem, which also features pairwise interactions. Similarly, inference for the bi-partite RBM maps to the asymmetric rank-1 matrix factorization. However, conversely to the Boltzmann inference, these factorizations are reconstruction problems. The mean-field techniques, derived in [LKZ16, LKZ17], allow there to compute the MMSE estimator of unknown signals from approximate marginals. Here we focus on the evaluation of the free energy.

Minimization for fixed marginals While fixing the value of the first moment is sufficient for binary variables, more than one constraint is now needed in order to minimize the Gibbs free energy at a given value of the marginals. In the same spirit of the AMP algorithm we assume a Gaussian parametrization of the marginals. We note \underline{a} the first moment of \underline{x} and \underline{c} its variance. We wish to compute the constrained minimum over the distributions q on \mathbb{R}^N

$$G(\underline{a}, \underline{c}) = \min_q [\langle E(\underline{x}) \rangle_q - H(q)/\beta \mid \langle \underline{x} \rangle_q = \underline{a}, \langle \underline{x}^2 \rangle_q = \underline{a}^2 + \underline{c}], \quad (187)$$

where the notation of squared vectors corresponds here and below to the vectors of squared entries. It is equivalent to an unconstrained problem with Lagrange multipliers $\underline{\lambda}(\underline{a}, \underline{c}, \beta)$ and $\underline{\xi}(\underline{a}, \underline{c}, \beta)$

$$G(\underline{a}, \underline{c}) = \min_q [\langle E(\underline{x}) \rangle_q - H(q)/\beta - \underline{\lambda}^\top (\langle \underline{x} \rangle_q - \underline{a})/\beta - \underline{\xi}^\top (\langle \underline{x}^2 \rangle_q - \underline{a}^2 - \underline{c})/\beta]. \quad (188)$$

The terms depending on the distribution q in the functional to minimize above can be interpreted as a Gibbs free energy for the effective energy functional

$$\tilde{E}(\underline{x}) = E(\underline{x}) - \underline{\lambda}^\top \underline{x}/\beta - \underline{\xi}^\top \underline{x}^2/\beta. \quad (189)$$

The solution of the minimization problem (188) is therefore the corresponding Boltzmann distribution

$$q_{\underline{a}, \underline{c}}(\underline{x}) = \frac{e^{-\beta \tilde{E}(\underline{x})}}{\tilde{\mathcal{Z}}} = \frac{1}{\tilde{\mathcal{Z}}} e^{-\beta E(\underline{x}) + \underline{\lambda}(\underline{a}, \underline{c}, \beta)^\top \underline{x} + \underline{\xi}(\underline{a}, \underline{c}, \beta)^\top \underline{x}^2} \quad (190)$$

and the minimum $G(\underline{a}, \underline{c})$ is

$$\begin{aligned} -\beta G(\underline{a}, \underline{c}) &= -\underline{\lambda}^\top \underline{a} - \underline{\xi}^\top (\underline{a}^2 + \underline{c}) + \log \int d\underline{x} e^{-\beta E(\underline{x}) + \underline{\lambda}^\top \underline{x} + \underline{\xi}^\top \underline{x}^2} \\ &= \log \int d\underline{x} e^{-\beta E(\underline{x}) + \underline{\lambda}^\top (\underline{x} - \underline{a}) + \underline{\xi}^\top (\underline{x}^2 - \underline{a}^2 - \underline{c})}, \end{aligned} \quad (191)$$

D.0

where the Lagrange multipliers $\lambda(\underline{a}, \underline{c}, \beta)$ and $\xi(\underline{a}, \underline{c}, \beta)$ enforcing the constraints are still implicit. Defining a functional \tilde{G} for arbitrary vectors $\tilde{\lambda} \in \mathbb{R}^N$ and $\tilde{\xi} \in \mathbb{R}^N$,

$$-\beta \tilde{G}(\underline{a}, \underline{c}, \tilde{\lambda}, \tilde{\xi}) = \log \int d\underline{x} e^{-\beta E(\underline{x}) + \tilde{\lambda}^\top (\underline{x} - \underline{a}) + \tilde{\xi}^\top (\underline{x}^2 - \underline{a}^2 - \underline{c})}, \quad (192)$$

we have

$$a_i = \langle x_i \rangle_{q_{\underline{a}, \underline{c}}} \Rightarrow -\beta \left. \frac{\partial \tilde{G}}{\partial \tilde{\lambda}_i} \right|_{\tilde{\lambda}, \tilde{\xi}} = 0, \quad -\beta \left. \frac{\partial^2 \tilde{G}}{\partial \tilde{\lambda}_i^2} \right|_{\tilde{\lambda}, \tilde{\xi}} = \langle x_i^2 \rangle_{q_{\underline{a}, \underline{c}}} - a_i^2 > 0, \quad (193)$$

$$c_i + a_i^2 = \langle x_i^2 \rangle_{q_{\underline{a}, \underline{c}}} \Rightarrow -\beta \left. \frac{\partial \tilde{G}}{\partial \tilde{\xi}_i} \right|_{\tilde{\lambda}, \tilde{\xi}} = 0, \quad -\beta \left. \frac{\partial^2 \tilde{G}}{\partial \tilde{\xi}_i^2} \right|_{\tilde{\lambda}, \tilde{\xi}} = \langle (x_i^2)^2 \rangle_{q_{\underline{a}, \underline{c}}} - (c_i + a_i^2)^2 > 0. \quad (194)$$

Hence the Lagrange multipliers are identified as minimizers of $-\beta \tilde{G}$ and

$$-\beta G(\underline{a}, \underline{c}) = -\beta \tilde{G}(\underline{a}, \underline{c}, \lambda(\underline{a}, \underline{c}, \beta), \xi(\underline{a}, \underline{c}, \beta)) = \min_{\tilde{\lambda}, \tilde{\xi}} -\beta \tilde{G}(\underline{a}, \underline{c}, \tilde{\lambda}, \tilde{\xi}). \quad (195)$$

The true free energy $F = -\log \mathcal{Z} / \beta$ would eventually be recovered by minimizing the constrained minimum $G(\underline{a}, \underline{c})$ with respect to its arguments. Nevertheless, the computation of G and \tilde{G} involves an integration over $\underline{x} \in \mathbb{R}^N$ and remains intractable. The following step of the Georges-Yedidia derivation consists in approximating these functionals by a Taylor expansion at infinite temperature where interactions are neutralized.

Expansion around $\beta = 0$ To perform the expansion we introduce the notation $A(\beta, \underline{a}, \underline{c}) = -\beta G(\underline{a}, \underline{c})$. We also define the auxiliary operator

$$U(\underline{x}; \beta) = -\frac{1}{2} \underline{x}^\top \underline{W} \underline{x} + \frac{1}{2} \langle \underline{x}^\top \underline{W} \underline{x} \rangle_{q_{\underline{a}, \underline{c}}} - \sum_{i=1}^N \frac{\partial \lambda_i}{\partial \beta} (x_i - a_i) - \sum_{i=1}^N \frac{\partial \xi_i}{\partial \beta} (x_i^2 - a_i^2 - c_i), \quad (196)$$

that allows to write concisely for any observable O the derivative of its average with respect to β ,

$$\frac{\partial \langle O(\underline{x}; \beta) \rangle_{q_{\underline{a}, \underline{c}}}}{\partial \beta} = \left\langle \frac{\partial O(\underline{x}; \beta)}{\partial \beta} \right\rangle_{q_{\underline{a}, \underline{c}}} - \langle U(\underline{x}; \beta) O(\underline{x}; \beta) \rangle_{q_{\underline{a}, \underline{c}}}. \quad (197)$$

To compute the derivatives of λ and ξ with respect to β we note that

$$\frac{\partial A}{\partial a_i} = -\beta \frac{\partial \tilde{G}}{\partial a_i} = -\lambda_i(\beta, \underline{a}, \underline{c}) - 2a_i \xi_i(\beta, \underline{a}, \underline{c}), \quad (198)$$

$$\frac{\partial A}{\partial c_i} = -\beta \frac{\partial \tilde{G}}{\partial c_i} = -\xi_i(\beta, \underline{a}, \underline{c}), \quad (199)$$

where we used that $\partial \tilde{G} / \partial \tilde{\lambda}_i = 0$ and $\partial \tilde{G} / \partial \tilde{\xi}_i = 0$ when evaluated for $\lambda(\underline{a}, \underline{c}, \beta)$ and $\xi(\underline{a}, \underline{c}, \beta)$. Consequently,

$$\frac{\partial \xi_i}{\partial \beta} = -\frac{\partial}{\partial c_i} \frac{\partial A}{\partial \beta}, \quad \frac{\partial \lambda_i}{\partial \beta} = -\frac{\partial}{\partial a_i} \frac{\partial A}{\partial \beta} + 2a_i \frac{\partial \xi_i}{\partial \beta}. \quad (200)$$

We can now proceed to compute the first terms of the expansion that will be performed for the functional A .

Zeroth order Substituting $\beta = 0$ in the definition of A we have

$$A(0, \underline{a}, \underline{c}) = -\lambda(0, \underline{a}, \underline{c})^\top \underline{a} - \xi(0, \underline{a}, \underline{c})^\top (\underline{a}^2 + \underline{c}) + \log \tilde{\mathcal{Z}}^0(\lambda(0, \underline{a}, \underline{c}), \xi(0, \underline{a}, \underline{c})), \quad (201)$$

with

$$\tilde{\mathcal{Z}}^0(\lambda(0, \underline{a}, \underline{c}), \xi(0, \underline{a}, \underline{c})) = \int d\underline{x} e^{\lambda(0, \underline{a}, \underline{c})^\top \underline{x} + \xi(0, \underline{a}, \underline{c})^\top \underline{x}^2} \prod_{i=1}^N p_x(x_i; \theta_i) \quad (202)$$

$$= \prod_{i=1}^N \int dx_i e^{\lambda_i(0, \underline{a}, \underline{c}) x_i + \xi_i(0, \underline{a}, \underline{c}) x_i^2} p_x(x_i; \theta_i). \quad (203)$$

D Georges-Yedidia expansion for generalized Boltzmann machines

At infinite temperature the interaction terms of the energy do not contribute so that the integral in \tilde{Z}^0 factorizes and can be evaluated numerically in the event that it does not have a closed-form.

First order We compute the derivative of A with respect to β . We use again that $\underline{\lambda}(\underline{a}, \underline{c}, \beta)$ and $\underline{\xi}(\underline{a}, \underline{c}, \beta)$ are stationary points of \tilde{G} to write

$$\frac{\partial A}{\partial \beta} = -\beta \frac{\partial \tilde{G}}{\partial \beta} = \frac{\partial}{\partial \beta} \left[\log \int d\underline{x} e^{-\beta E(\underline{x}) + \underline{\lambda}(\underline{a}, \underline{c}, \beta)^\top (\underline{x} - \underline{a}) + \underline{\xi}(\underline{a}, \underline{c}, \beta)^\top (\underline{x}^2 - \underline{a}^2 - \underline{c})} \right] \quad (204)$$

$$= \left\langle \frac{\partial}{\partial \beta} (-\beta E(\underline{x})) + \frac{\partial \underline{\lambda}}{\partial \beta}^\top (\underline{x} - \underline{a}) + \frac{\partial \underline{\xi}}{\partial \beta}^\top (\underline{x}^2 - \underline{a}^2 - \underline{c}) \right\rangle_{q_{\underline{a}, \underline{c}}} \quad (205)$$

$$= \frac{1}{2} \langle \underline{x}^\top \underline{W} \underline{x} \rangle_{q_{\underline{a}, \underline{c}}}. \quad (206)$$

At infinite temperature the average over the product of variables becomes a product of averages so that we have

$$\left. \frac{\partial A}{\partial \beta} \right|_{\beta=0} = \frac{1}{2} \underline{a}^\top \underline{W} \underline{a} = \sum_{(ij)} W_{ij} a_i a_j. \quad (207)$$

Second order Using the first order derivative of A we can compute the derivatives of the Lagrange parameters (200) and the auxillary operator at infinite temperature,

$$\left. \frac{\partial \xi_i}{\partial \beta} \right|_{\beta=0} = 0, \quad \left. \frac{\partial \lambda_i}{\partial \beta} \right|_{\beta=0} = - \sum_{j \in \partial i} W_{ij} a_j, \quad U(\underline{x}; 0) = - \sum_{(ij)} W_{ij} (x_i - a_i)(x_j - a_j).$$

The second order derivative is then easily computed at infinite temperature

$$\left. \frac{\partial^2 A}{\partial \beta^2} \right|_{\beta=0} = \frac{1}{2} \frac{\partial}{\partial \beta} \left(\langle \underline{x}^\top \underline{W} \underline{x} \rangle_{q_{\underline{a}, \underline{c}}} \right) \Big|_{\beta=0} = - \frac{1}{2} \langle U(\underline{x}; 0) (\underline{x}^\top \underline{W} \underline{x}) \rangle_{q_{\underline{a}, \underline{c}}}^{\beta=0} \quad (208)$$

$$= \sum_{(ij)} W_{ij}^2 \langle (x_i - a_i) x_i (x_j - a_j) \rangle_{q_{\underline{a}, \underline{c}}}^{\beta=0} = \sum_{(ij)} W_{ij}^2 c_i c_j. \quad (209)$$

TAP free energy for the generalized Boltzmann machine Stopping at the second order of the systematic expansion, and gathering the different terms derived above we have

$$\begin{aligned} -\beta G(\underline{a}, \underline{c}) &= -\underline{\lambda}(0, \underline{a}, \underline{c})^\top \underline{a} - \underline{\xi}(0, \underline{a}, \underline{c})^\top (\underline{a}^2 + \underline{c}) + \log \tilde{Z}^0(\underline{\lambda}(0, \underline{a}, \underline{c}), \underline{\xi}(0, \underline{a}, \underline{c})) \\ &\quad + \beta \sum_{(ij)} W_{ij} a_i a_j + \frac{\beta^2}{2} \sum_{(ij)} W_{ij}^2 c_i c_j, \end{aligned} \quad (210)$$

where the values of the parameters $\underline{\lambda}(0, \underline{a}, \underline{c})$ and $\underline{\xi}(0, \underline{a}, \underline{c})$ are implicitly defined through the stationary conditions (193)-(194). The TAP approximation of the free energy also requires to consider the stationary points of the expanded expression as a function of \underline{a} and \underline{c} .

This second condition yields the relations

$$-2\xi_i(0, \underline{a}, \underline{c}) = -\beta^2 \sum_{j \in \partial i} W_{ij}^2 c_j = A_i \quad (211)$$

$$\lambda_i(0, \underline{a}, \underline{c}) = A_i a_i + \beta \sum_{j \in \partial i} W_{ij} a_j = B_i \quad (212)$$

where we define new variables A_i and B_i . While the extremization with respect to the Lagrange multipliers gives

$$a_i = \frac{1}{\tilde{Z}_i^x} \int d x_i x_i p_x(x_i; \theta_i) e^{-\frac{A_i}{2} x_i^2 + B_i x_i} = f_1^x(B_i, A_i; \theta_i), \quad (213)$$

$$c_i = \frac{1}{\tilde{Z}_i^x} \int d x_i x_i^2 p_x(x_i; \theta_i) e^{-\frac{A_i}{2} x_i^2 + B_i x_i} - a_i^2 = f_2^x(B_i, A_i; \theta_i), \quad (214)$$

E.0

where we introduce update functions f_1^x and f_2^x with respect to the partition function

$$\mathcal{Z}_i^x(B_i, A_i; \theta_i) = \int dx_i p_x(x_i; \theta_i) e^{-\frac{A_i}{2} x_i^2 + B_i x_i}. \quad (215)$$

Finally we can rewrite the TAP free energy as

$$\begin{aligned} -\beta G(\underline{a}, \underline{c}) = & -\underline{B}^\top \underline{a} + \underline{A}^\top (\underline{a}^2 + \underline{c})/2 + \sum_{i=1}^N \log \tilde{\mathcal{Z}}_i^x(B_i, A_i; \theta_i) \\ & + \beta \sum_{(ij)} W_{ij} a_i a_j + \frac{\beta^2}{2} \sum_{(ij)} W_{ij}^2 c_i c_j, \end{aligned} \quad (216)$$

with the values of the parameters set by the self-consistency conditions (211), (212), (213) and (214), which are the TAP equations of the generalized Boltzmann machine at second order. Note that the naive mean-field equations are recovered by ignoring the second order terms in β^2 .

Relation to message passing The TAP equations obtained above must correspond to the fixed points of the Approximate Message Passing (AMP) following the derivation from Belief Propagation (BP) that is presented in Section 4.3.3. In the Appendix B of [TGM⁺18] the relaxed-BP equations are derived for the generalized Boltzmann machine:

$$B_{i \rightarrow j}^{(t)} = \sum_{k \in \partial i \setminus j} \beta W_{ik} a_{k \rightarrow i}^{(t)}, \quad A_{i \rightarrow j}^{(t)} = - \sum_{k \in \partial i \setminus j} \beta^2 W_{ik}^2 c_{k \rightarrow i}^{(t)}, \quad (217)$$

$$a_{i \rightarrow j}^{(t)} = f_1^x(B_{i \rightarrow j}^{(t-1)}, A_{i \rightarrow j}^{(t-1)}; \theta_i), \quad c_{i \rightarrow j}^{(t)} = f_2^x(B_{i \rightarrow j}^{(t-1)}, A_{i \rightarrow j}^{(t-1)}; \theta_i), \quad (218)$$

To recover the TAP equations for them we define

$$B_i^{(t)} = \sum_{k \in \partial i} \beta W_{ik} a_{k \rightarrow i}^{(t)}, \quad A_i^{(t)} = - \sum_{k \in \partial i} \beta^2 W_{ik}^2 c_{k \rightarrow i}^{(t)}, \quad (219)$$

$$a_i^{(t)} = f_1^x(B_i^{(t-1)}, A_i^{(t-1)}; \theta_i), \quad c_i^{(t)} = f_2^x(B_i^{(t-1)}, A_i^{(t-1)}; \theta_i). \quad (220)$$

As $B_i^{(t)} = B_{i \rightarrow j}^{(t)} + \beta W_{ij} a_{j \rightarrow i}^{(t)}$ and $A_i^{(t)} = A_{i \rightarrow j}^{(t)} - \beta^2 W_{ij}^2 c_{j \rightarrow i}^{(t)}$ we have by developing f_2^x that $c_i^{(t)} = c_{i \rightarrow j}^{(t)} + O(\beta)$ so that

$$A_i^{(t)} = -\beta^2 \sum_{j \in \partial i} W_{ij}^2 c_j^{(t)} + o(\beta^2). \quad (221)$$

By developing f_1^x we also have

$$a_k^{(t)} = f_1^x(B_{k \rightarrow j}^{(t-1)} + \beta W_{kj} a_{j \rightarrow k}^{(t-1)}, A_{k \rightarrow j}^{(t-1)} - \beta^2 W_{kj}^2 c_{j \rightarrow k}^{(t-1)}; \theta_k) \quad (222)$$

$$= a_{k \rightarrow j}^{(t)} + \frac{\partial f_1^x}{\partial B_k} \beta W_{kj} a_{j \rightarrow k}^{(t-1)} + O(\beta^2), \quad (223)$$

with $\frac{\partial f_1^x}{\partial B_k}(B_k^{(t-1)}, A_k^{(t-1)}; \theta_k) = c_k^{(t)}$. Finally, by replacing in the definition of B_i the messages we obtain

$$B_i^{(t)} = \sum_{k \in \partial i} \beta W_{ik} a_{k \rightarrow i}^{(t)} = \sum_{k \in \partial i} \beta W_{ik} a_k^{(t)} - \beta W_{ki} c_k^{(t)} a_{i \rightarrow k}^{(t-1)}. \quad (224)$$

As $a_{i \rightarrow k}^{(t-1)} = a_i^{(t-1)} + O(\beta)$ and using the definition of $A_i^{(t)}$, we finally recover

$$B_i^{(t)} = \sum_{k \in \partial i} \beta W_{ik} a_k^{(t)} + A_i^{(t)} a_i^{(t-1)}. \quad (225)$$

Hence we indeed recover the TAP equations as the AMP fixed points in (220), (221) and (225). Beyond the possibility to cross-check our results, the message passing derivation also specifies a scheme of updates to solve the self-consistency equations obtained by the Georges-Yedidia expansion. In the applications we consider below we should resort to this time indexing with good convergence properties [Bol14].

E Vector Approximate Message Passing for the GLM

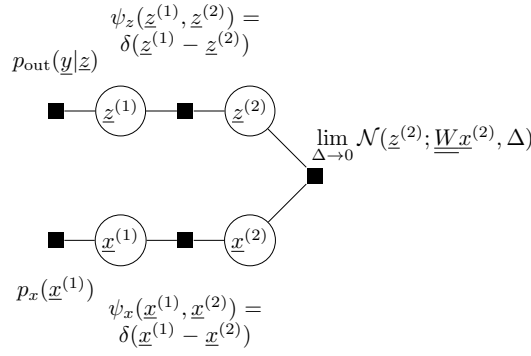


Figure 8: Factor graph representation of the GLM for the derivation of VAMP (reproduction of Figure 5 to help following the derivation described here in the Appendix).

Solutions of the TAP equations As already discussed in Section 4.2, the TAP equations do not necessarily admit a single solution. In practice, different fixed points are reached when considering different initializations of the iteration of the self-consistent equations.

E Vector Approximate Message Passing for the GLM

We recall here a possible derivation of G-VAMP discussed in Section 4 (Algorithm 2). We consider a projection of the BP equations for the factor graph Figure 8.

Gaussian assumptions We start by parametrizing marginals as well as messages coming out of the Dirac factors. For $a = 1, 2$:

$$m_{x,(a)}(\underline{x}^{(a)}) = \mathcal{N}(\underline{x}^{(a)}, \hat{\underline{x}}^{(a)}, \underline{C}^{x(a)}), \quad m_{z,(a)}(\underline{z}^{(a)}) = \mathcal{N}(\underline{z}^{(a)}, \hat{\underline{z}}^{(a)}, \underline{C}^{z(a)}), \quad (226)$$

and

$$\tilde{m}_{\psi_x \rightarrow \underline{x}^{(a)}}(\underline{x}^{(a)}) \propto e^{-\frac{1}{2} \underline{x}^{(a)\top} \underline{A}^{(a)} \underline{x}^{(a)} + \underline{B}_x^{(a)\top} \underline{x}^{(a)}}, \quad (227)$$

$$\tilde{m}_{\psi_z \rightarrow \underline{z}^{(a)}}(\underline{z}^{(a)}) \propto e^{-\frac{1}{2} \underline{z}^{(a)\top} \underline{A}^{(a)} \underline{z}^{(a)} + \underline{B}_z^{(a)\top} \underline{z}^{(a)}}. \quad (228)$$

Self consistency of the parametrizations at Dirac factor nodes Around ψ_x the message passing equations are simply

$$\tilde{m}_{\psi_x \rightarrow \underline{x}^{(2)}}(\underline{x}^{(2)}) = m_{\underline{x}^{(1)} \rightarrow \psi_x}(\underline{x}^{(2)}), \quad \tilde{m}_{\psi_x \rightarrow \underline{x}^{(1)}}(\underline{x}^{(1)}) = m_{\underline{x}^{(2)} \rightarrow \psi_x}(\underline{x}^{(1)}) \quad (229)$$

and similarly around ψ_z . Moreover, considering that messages are marginals to which the contribution of the opposite message is retrieved we have

$$m_{\underline{x}^{(1)} \rightarrow \psi_x}(\underline{x}^{(1)}) \propto m_{x,(1)}(\underline{x}^{(1)}) / \tilde{m}_{\psi_x \rightarrow \underline{x}^{(1)}}(\underline{x}^{(1)}), \quad (230)$$

$$m_{\underline{x}^{(2)} \rightarrow \psi_x}(\underline{x}^{(2)}) \propto m_{x,(2)}(\underline{x}^{(2)}) / \tilde{m}_{\psi_x \rightarrow \underline{x}^{(2)}}(\underline{x}^{(2)}). \quad (231)$$

Combining this observation along with (229) leads to updates (139) and (135). The same reasoning can be followed for the messages around ψ_z leading to updates (141) and (137).

Input and output update functions The update functions of means and variances of the marginals are deduced from the parametrized message passing. For the variable $\underline{x}^{(1)}$

taking into account the prior p_x , the updates are very similar to GAMP input functions:

$$\hat{\underline{x}}^{(1)} \propto \int d\underline{x}^{(1)} \underline{x}^{(1)} p_x(\underline{x}^{(1)}) \tilde{m}_{\psi_x \rightarrow \underline{x}^{(1)}}(\underline{x}^{(1)}) \quad (232)$$

$$= \frac{1}{\mathcal{Z}_x^{(1)}} \int d\underline{x}^{(1)} \underline{x}^{(1)} p_x(\underline{x}^{(1)}) e^{-\frac{1}{2} \underline{x}^{(1)\top} \underline{A}_x^{(1)} \underline{x}^{(1)} + \underline{B}_x^{(1)\top} \underline{x}^{(1)}} = f_1^x(\underline{B}_x^{(1)}, \underline{A}_x^{(1)}), \quad (233)$$

$$\underline{\underline{C}}_x^{(1)} = \frac{1}{\mathcal{Z}_x^{(1)}} \int d\underline{x}^{(1)} \underline{x}^{(1)} \underline{x}^{(1)\top} p_x(\underline{x}^{(1)}) e^{-\frac{1}{2} \underline{x}^{(1)\top} \underline{A}_x^{(1)} \underline{x}^{(1)} + \underline{B}_x^{(1)\top} \underline{x}^{(1)}} \quad (234)$$

$$= f_1^x(\underline{B}_x^{(1)}, \underline{A}_x^{(1)}) f_1^x(\underline{B}_x^{(1)}, \underline{A}_x^{(1)})^\top \quad (235)$$

where $\mathcal{Z}_x^{(1)}$ is as usual the partition ensuring the normalization.

Similarly for the variable $\underline{z}^{(1)}$, the update functions are very similar to the GAMP output functions including the information coming from the observations:

$$\hat{\underline{z}}^{(1)} \propto \int d\underline{z}^{(1)} p_{\text{out}}(\underline{y}|\underline{z}^{(1)}) \tilde{m}_{\psi_z \rightarrow \underline{z}^{(1)}}(\underline{z}^{(1)}) \quad (236)$$

$$= \frac{1}{\mathcal{Z}_z^{(1)}} \int d\underline{z}^{(1)} p_{\text{out}}(\underline{y}|\underline{z}^{(1)}) e^{-\frac{1}{2} \underline{z}^{(1)\top} \underline{A}_z^{(1)} \underline{z}^{(1)} + \underline{B}_z^{(1)\top} \underline{z}^{(1)}} = f_1^z(\underline{B}_z^{(1)}, \underline{A}_z^{(1)}), \quad (237)$$

$$\underline{\underline{C}}_z^{(1)} = \frac{1}{\mathcal{Z}_z^{(1)}} \int d\underline{z}^{(1)} \underline{z}^{(1)} \underline{z}^{(1)\top} p_{\text{out}}(\underline{y}|\underline{z}^{(1)}) e^{-\frac{1}{2} \underline{z}^{(1)\top} \underline{A}_z^{(1)} \underline{z}^{(1)} + \underline{B}_z^{(1)\top} \underline{z}^{(1)}} \quad (238)$$

$$= f_1^z(\underline{B}_z^{(1)}, \underline{A}_z^{(1)}) f_1^z(\underline{B}_z^{(1)}, \underline{A}_z^{(1)})^\top \quad (239)$$

Linear transformation For the middle factor node we consider the vector variable concatenating $\underline{x} = [\underline{x}^{(2)} \underline{z}^{(2)}] \in \mathbb{R}^{N+M}$. The computation of the corresponding marginal with the message passing then yields

$$m_{\bar{x}}(\underline{\bar{x}}) \propto \lim_{\Delta \rightarrow 0} \mathcal{N}(\underline{z}^{(2)}; \underline{W}\underline{x}^{(2)}, \underline{\Delta}\underline{I}_{\underline{M}}) e^{-\frac{1}{2} \underline{x}^{(2)\top} \underline{A}_x^{(2)} \underline{x}^{(2)} + \underline{B}_x^{(2)\top} \underline{x}^{(2)}} e^{-\frac{1}{2} \underline{z}^{(2)\top} \underline{A}_z^{(2)} \underline{z}^{(2)} + \underline{B}_z^{(2)\top} \underline{z}^{(2)}}. \quad (240)$$

The means of $\underline{x}^{(2)}$ and $\underline{z}^{(2)}$ are then updated through

$$\hat{\underline{x}}^{(2)}, \hat{\underline{z}}^{(2)} = \arg \min_{\underline{x}, \underline{z}} \left[\|\underline{W}\underline{x} - \underline{z}\|^2 / \Delta + \underline{x}^\top \underline{A}_x^{(2)} \underline{x} - 2\underline{B}_x^{(2)\top} \underline{x} + \underline{z}^\top \underline{A}_z^{(2)} \underline{z} - 2\underline{B}_z^{(2)\top} \underline{z} \right], \quad (241)$$

at $\Delta \rightarrow 0$. At this point it is advantageous in terms of speed to consider the singular value decomposition $\underline{W} = \underline{U}\underline{S}\underline{V}^\top$ and to simplify the form of the variance matrices by taking them proportional to the identity, i.e. $\underline{A}_z^{(2)} = A_z^{(2)} \underline{I}_M$ etc. Under this assumption the solution of the minimization problem is

$$\hat{\underline{x}}^{(2)} = g_1^x(\underline{B}_x^{(2)}, A_x^{(2)}, \underline{B}_z^{(2)}, A_z^{(2)}) = \underline{V}\underline{D} \left(A_z^{(2)-2} \underline{S}\underline{U}^\top \underline{B}_z^{(2)} + A_x^{(2)-2} \underline{V}^\top \underline{B}_x^{(2)} \right), \quad (242)$$

$$\hat{\underline{z}}^{(2)} = g_1^z(\underline{B}_x^{(2)}, A_x^{(2)}, \underline{B}_z^{(2)}, A_z^{(2)}) = \underline{W}g_1^x(\underline{B}_x^{(2)}, A_x^{(2)}, \underline{B}_z^{(2)}, A_z^{(2)}), \quad (243)$$

with \underline{D} a diagonal matrix with entries $D_{ii} = (A_z^{(2)-1} S_{ii}^2 + A_x^{(2)-1})^{-1}$. The scalar variances are then updated using the traces of the Jacobians with respect to the $\underline{B}^{(2)}$ -s

$$\underline{\underline{C}}_x^{(2)} = \frac{A_x^{(2)}}{N} \text{tr} \left(\partial g_2^x / \partial \underline{B}_x^{(2)} \right) \underline{I}_N = \frac{1}{N} \sum_{i=1}^N (A_z^{(2)-1} S_{ii}^2 + A_x^{(2)-1})^{-1} \underline{I}_N \quad (244)$$

$$= g_2^x(\underline{B}_x^{(2)}, A_x^{(2)}, \underline{B}_z^{(2)}, A_z^{(2)}) \quad (245)$$

$$\underline{\underline{C}}_z^{(2)} = \frac{A_z^{(2)}}{M} \text{tr} \left(\partial g_2^z / \partial \underline{B}_z^{(2)} \right) \underline{I}_M = \frac{1}{M} \sum_{i=1}^N S_{ii} (A_z^{(2)-1} S_{ii}^2 + A_x^{(2)-1})^{-1} \underline{I}_M \quad (246)$$

$$= g_2^z(\underline{B}_x^{(2)}, A_x^{(2)}, \underline{B}_z^{(2)}, A_z^{(2)}). \quad (247)$$

F Multi-value AMP derivation for the GLM

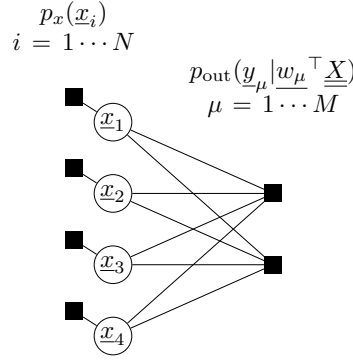


Figure 9: Factor graph of the Generalized Linear Model (GLM) on vector variables corresponding to the joint distribution (248).

F Multi-value AMP derivation for the GLM

We here present the derivation of the multi-value AMP and its SE motivated in Section 5.3, focusing on the multi-value GLM. These derivations also appear in [GBKZ19].

F.1 Approximate Message Passing

The systematic procedure to write AMP for a given joint probability distribution consists in first writing BP on the factor graph, second project the messages on a parametrized family of functions to obtain the corresponding relaxed-BP and third close the equations on a reduced set of parameters by keeping only leading terms in the thermodynamic limit.

For the generic multi-value GLM the posterior measure we are interested in is

$$p(\underline{X}|\underline{Y}, \underline{W}) = \frac{1}{\mathcal{Z}(\underline{Y}, \underline{W})} \prod_{i=1}^N p(\underline{x}_i) \prod_{\mu=1}^M p_{\text{out}}(\underline{y}_\mu | \underline{w}_\mu^\top \underline{X} / \sqrt{N}), \quad \underline{x}_i \in \mathbb{R}^P, \quad \underline{y}_\mu \in \mathbb{R}^P. \quad (248)$$

where the known entries of matrix \underline{W} are drawn i.i.d from a standard normal distribution (the scaling in $1/\sqrt{N}$ is here made explicit). The corresponding factor graph is given on Figure 9. We are considering the simultaneous reconstruction of P signals $\underline{x}_{0,(k)} \in \mathbb{R}^N$ and therefore write the message passing on the variables $\underline{x}_i \in \mathbb{R}^P$. The major difference with the scalar version ($P=1$) of AMP is that we will consider covariance matrices between variables coming from the P observations instead of scalar variances.

Belief propagation (BP) We start with BP on the factor graph of Figure 9. For all pairs of index $i - \mu$, we define the update equations of messages function

$$\tilde{m}_{\mu \rightarrow i}^{(t)}(\underline{x}_i) = \frac{1}{\mathcal{Z}_{\mu \rightarrow i}} \int \prod_{i' \neq i} d\underline{x}_{i'} p_{\text{out}} \left(\underline{y}_\mu | \sum_j \frac{W_{\mu j}}{\sqrt{N}} \underline{x}_j \right) \prod_{i' \neq i} m_{i' \rightarrow \mu}^{(t)}(\underline{x}_{i'}) \quad (249)$$

$$m_{i \rightarrow \mu}^{(t+1)}(\underline{x}_i) = \frac{1}{\mathcal{Z}_{i \rightarrow \mu}} p_x(\underline{x}_i) \prod_{\mu' \neq \mu} \tilde{m}_{\mu' \rightarrow i}^{(t)}(\underline{x}_i), \quad (250)$$

where $\mathcal{Z}_{\mu \rightarrow i}$ and $\mathcal{Z}_{i \rightarrow \mu}$ are normalization function that allow to interpret messages as probability distributions. To improve readability, we drop the time indices in the following derivation, and only specify them in the final algorithm.

Relaxed BP (r-BP) The second step of the derivation is to develop messages keeping only terms up to order $O(1/N)$ as we take the thermodynamic limit $N \rightarrow +\infty$ (at fixed $\alpha = M/N$). At this order, we will find that it is consistent to consider the messages to be

F.1 Approximate Message Passing

approximately Gaussian, i.e. characterized by their means and co-variances. Thus we define

$$\hat{\underline{x}}_{i \rightarrow \mu} = \int d\underline{x} \underline{x} m_{i \rightarrow \mu}(\underline{x}) \quad (251)$$

$$\underline{\underline{C}}_{i \rightarrow \mu}^x = \int d\underline{x} \underline{x} \underline{x}^T m_{i \rightarrow \mu}(\underline{x}) - \hat{\underline{x}}_{i \rightarrow \mu} \hat{\underline{x}}_{i \rightarrow \mu}^T \quad (252)$$

and

$$\underline{\omega}_{\mu \rightarrow i} = \sum_{i' \neq i} \frac{W_{\mu i'}}{\sqrt{N}} \hat{\underline{x}}_{i' \rightarrow \mu} \quad (253)$$

$$\underline{V}_{\mu \rightarrow i} = \sum_{i' \neq i} \frac{W_{\mu i'}^2}{N} \underline{\underline{C}}_{i' \rightarrow \mu}^x, \quad (254)$$

where $\underline{\omega}_{\mu \rightarrow i}$ and $\underline{V}_{\mu \rightarrow i}$ are related to the intermediate variable $\underline{z}_\mu = \underline{w}_\mu^T \underline{X}$.

Expansion of $\tilde{m}_{\mu \rightarrow i}$ - We defined the Fourier transform \hat{p}_{out} of $p_{\text{out}}(\underline{y}_\mu | \underline{z}_\mu)$ with respect to its argument $\underline{z}_\mu = \underline{w}_\mu^T \underline{X}$,

$$\hat{p}_{\text{out}}(\underline{y}_\mu | \underline{\xi}_\mu) = \int d\underline{z}_\mu \hat{p}_{\text{out}}(\underline{y}_\mu | \underline{z}_\mu) e^{-i \underline{\xi}_\mu^T \underline{z}_\mu}. \quad (255)$$

Using reciprocally the Fourier representation of $p_{\text{out}}(\underline{y}_\mu | \underline{z}_\mu)$,

$$p_{\text{out}}(\underline{y}_\mu | \underline{z}_\mu) = \frac{1}{(2\pi)^M} \int d\underline{\xi}_\mu \hat{p}_{\text{out}}(\underline{y}_\mu | \underline{\xi}_\mu) e^{i \underline{\xi}_\mu^T \underline{z}_\mu}, \quad (256)$$

we decouple the integrals over the different $\underline{x}_{i'}$ in (249),

$$\tilde{m}_{\mu \rightarrow i}(\underline{x}_i) \propto \int d\underline{\xi}_\mu \hat{p}_{\text{out}}(\underline{y}_\mu | \underline{\xi}_\mu) e^{i \frac{W_{\mu i}}{\sqrt{N}} \underline{\xi}_\mu^T \underline{x}_i} \prod_{i' \neq i} \int d\underline{x}_{i'} m_{i' \rightarrow \mu}(\underline{x}_{i'}) e^{i \frac{W_{\mu i'}}{\sqrt{N}} \underline{x}_i \underline{\xi}_\mu^T \underline{x}_{i'}} \quad (257)$$

$$\propto \int d\underline{\xi}_\mu \hat{p}_{\text{out}}(\underline{y}_\mu | \underline{\xi}_\mu) e^{i \underline{\xi}_\mu^T \left(\frac{W_{\mu i}}{\sqrt{N}} \underline{x}_i + \underline{\omega}_{\mu \rightarrow i} \right) - \frac{1}{2} \underline{\xi}_\mu^T \underline{V}_{\mu \rightarrow i}^{-1} \underline{\xi}_\mu} \quad (258)$$

where developing the exponentials of the product in (249) allows to express the integrals over the $\underline{x}_{i'}$ as a function of the definitions (253)-(254), before re-exponentiating to obtain the final result (258). Now reversing the Fourier transform and performing the integral over $\underline{\xi}$ we can further rewrite

$$\tilde{m}_{\mu \rightarrow i}(\underline{x}_i) \propto \int d\underline{z}_\mu p_{\text{out}}(\underline{y}_\mu | \underline{z}_\mu) e^{-\frac{1}{2} \left(\underline{z}_\mu - \frac{W_{\mu i}}{\sqrt{N}} \underline{x}_i - \underline{\omega}_{\mu \rightarrow i} \right)^T \underline{V}_{\mu \rightarrow i}^{-1} \left(\underline{z}_\mu - \frac{W_{\mu i}}{\sqrt{N}} \underline{x}_i - \underline{\omega}_{\mu \rightarrow i} \right)} \quad (259)$$

$$\propto \int d\underline{z}_\mu \mathbb{P}_{\text{out}}(\underline{z}_\mu; \underline{\omega}_{\mu \rightarrow i}, \underline{V}_{\mu \rightarrow i}) e^{\left(\underline{z}_\mu - \underline{\omega}_{\mu \rightarrow i} \right)^T \underline{V}_{\mu \rightarrow i}^{-1} \frac{W_{\mu i}}{\sqrt{N}} \underline{x}_i - \frac{W_{\mu i}^2}{2N} \underline{x}_i^T \underline{V}_{\mu \rightarrow i}^{-1} \underline{x}_i}, \quad (260)$$

where we are led to introduce the *output update functions*,

$$\mathbb{P}_{\text{out}}(\underline{z}_\mu; \underline{\omega}_{\mu \rightarrow i}, \underline{V}_{\mu \rightarrow i}) = p_{\text{out}}(\underline{y}_\mu | \underline{z}_\mu) \mathcal{N}(\underline{z}_\mu; \underline{\omega}_{\mu \rightarrow i}, \underline{V}_{\mu \rightarrow i}), \quad (261)$$

$$\underline{\mathcal{Z}}_{\text{out}}(\underline{y}_\mu, \underline{\omega}_{\mu \rightarrow i}, \underline{V}_{\mu \rightarrow i}) = \int d\underline{z}_\mu p_{\text{out}}(\underline{y}_\mu | \underline{z}_\mu) \mathcal{N}(\underline{z}_\mu; \underline{\omega}_{\mu \rightarrow i}, \underline{V}_{\mu \rightarrow i}), \quad (262)$$

$$\underline{g}_{\text{out}}(\underline{y}_\mu, \underline{\omega}_{\mu \rightarrow i}, \underline{V}_{\mu \rightarrow i}) = \frac{1}{\underline{\mathcal{Z}}_{\text{out}}} \frac{\partial \underline{\mathcal{Z}}_{\text{out}}}{\partial \underline{\omega}} \quad \text{and} \quad \partial_{\underline{\omega}} \underline{g}_{\text{out}} = \frac{\partial \underline{g}_{\text{out}}}{\partial \underline{\omega}}, \quad (263)$$

where $\mathcal{N}(\underline{z}; \underline{\omega}, \underline{V})$ is the multivariate Gaussian distribution of mean $\underline{\omega}$ and covariance \underline{V} . Further expanding the exponential in (260) up to order $O(1/N)$ leads to the Gaussian parametrization

$$\tilde{m}_{\mu \rightarrow i}(\underline{x}_i) \propto 1 + \frac{W_{\mu i}}{\sqrt{N}} \underline{g}_{\text{out}} \underline{x}_i + \frac{W_{\mu i}^2}{2N} \underline{x}_i^T (\underline{g}_{\text{out}} \underline{g}_{\text{out}}^T + \partial_{\underline{\omega}} \underline{g}_{\text{out}}) \underline{x}_i \quad (264)$$

$$\propto e^{\underline{B}_{\mu \rightarrow i}^T \underline{x}_i - \frac{1}{2} \underline{x}_i^T \underline{A}_{\mu \rightarrow i} \underline{x}_i}, \quad (265)$$

F Multi-value AMP derivation for the GLM

with

$$\underline{B}_{\mu \rightarrow i} = \frac{W_{\mu i}}{\sqrt{N}} \underline{g}_{\text{out}}(\underline{y}_{\mu}, \omega_{\mu \rightarrow i}, \underline{V}_{\mu \rightarrow i}) \quad (266)$$

$$\underline{A}_{\mu \rightarrow i} = -\frac{W_{\mu i}^2}{N} \partial_{\omega} \underline{g}_{\text{out}}(\underline{y}_{\mu}, \omega_{\mu \rightarrow i}, \underline{V}_{\mu \rightarrow i}). \quad (267)$$

Consistency with $m_{i \rightarrow \mu}$ - Inserting the Gaussian approximation of $\tilde{m}_{\mu \rightarrow i}$ in the definition of $m_{i \rightarrow \mu}$, we get the parametrization

$$m_{i \rightarrow \mu}(\underline{x}_i) \propto p_x(\underline{x}_i) \prod_{\mu' \neq \mu} e^{\underline{B}_{\mu' \rightarrow i}^T \underline{x}_i - \frac{1}{2} \underline{x}_i^T \underline{A}_{\mu' \rightarrow i} \underline{x}_i} \propto p_x(\underline{x}_i) e^{-\frac{1}{2} (\underline{x}_i - \underline{\lambda}_{i \rightarrow \mu})^T \underline{\sigma}_{i \rightarrow \mu}^{-1} (\underline{x}_i - \underline{\lambda}_{i \rightarrow \mu})} \quad (268)$$

with

$$\underline{\lambda}_{i \rightarrow \mu} = \underline{\sigma}_{i \rightarrow \mu} \left(\sum_{\mu' \neq \mu} \underline{B}_{\mu' \rightarrow i} \right) \quad (269)$$

$$\underline{\sigma}_{i \rightarrow \mu} = \left(\sum_{\mu' \neq \mu} \underline{A}_{\mu' \rightarrow i} \right)^{-1}. \quad (270)$$

Closing the equations - Ensuring the consistency with the definitions (251)-(252) of mean and covariance of $m_{i \rightarrow \mu}$ we finally close our set of equations by defining the *input update functions*,

$$\underline{Z}^x = \int d\underline{x} p_x(\underline{x}) e^{-\frac{1}{2} (\underline{x} - \underline{\lambda})^T \underline{\sigma}^{-1} (\underline{x} - \underline{\lambda})} \quad (271)$$

$$\underline{f}_1^x(\underline{\lambda}, \underline{\sigma}) = \frac{1}{\underline{Z}^x} \int d\underline{x} \underline{x} p_x(\underline{x}) e^{-\frac{1}{2} (\underline{x} - \underline{\lambda})^T \underline{\sigma}^{-1} (\underline{x} - \underline{\lambda})} \quad (272)$$

$$\underline{f}_2^x(\underline{\lambda}, \underline{\sigma}) = \frac{1}{\underline{Z}^x} \int d\underline{x} \underline{x} \underline{x}^T p_x(\underline{x}) e^{-\frac{1}{2} (\underline{x} - \underline{\lambda})^T \underline{\sigma}^{-1} (\underline{x} - \underline{\lambda})} - \underline{f}_1^x(\underline{\lambda}, \underline{\sigma}) \underline{f}_1^x(\underline{\lambda}, \underline{\sigma})^T, \quad (273)$$

so that

$$\underline{\hat{x}}_{i \rightarrow \mu} = \underline{f}_1^x(\underline{\lambda}_{i \rightarrow \mu}, \underline{\sigma}_{i \rightarrow \mu}) \quad (274)$$

$$\underline{C}_{i \rightarrow \mu}^x = \underline{f}_2^x(\underline{\lambda}_{i \rightarrow \mu}, \underline{\sigma}_{i \rightarrow \mu}). \quad (275)$$

The closed set of equations (253), (254), (266) (267), (269), (270), (274) and (275), with restored time indices, defines the r-BP algorithm. At convergence of the iterations, we obtain the approximated marginals

$$m_i(\underline{x}_i) = \frac{1}{\underline{Z}_i} p_x(\underline{x}_i) e^{-\frac{1}{2} (\underline{x}_i - \underline{\lambda}_i)^T \underline{\sigma}_i^{-1} (\underline{x}_i - \underline{\lambda}_i)} \quad (276)$$

with

$$\underline{\lambda}_i = \underline{\sigma}_i \left(\sum_{\mu=1}^M \underline{B}_{\mu \rightarrow i} \right) \quad (277)$$

$$\underline{\sigma}_i = \left(\sum_{\mu} \underline{A}_{\mu \rightarrow i} \right)^{-1}. \quad (278)$$

As usual, while BP requires to follow iterations over $M \times N$ message distributions over vectors in \mathbb{R}^P , r-BP only requires to track $O(M \times N \times P)$ variables, which is a great simplification. Nonetheless, r-BP can be further reduced to the more practical GAMP algorithm, given the scaling of the weights in $O(1/\sqrt{N})$.

Approximate message passing We define parameters $\underline{\omega}_\mu$, \underline{V}_μ and $\underline{\hat{x}}_i$, $\underline{C}_{\mu i}^x$, likewise $\underline{\lambda}_i$ and $\underline{\sigma}_i$ defined above and consider their relations to the original $\underline{\lambda}_{i \rightarrow \mu}$, $\underline{\sigma}_{i \rightarrow \mu}$, $\underline{\omega}_{\mu \rightarrow i}$, $\underline{V}_{\mu \rightarrow i}$, $\underline{\hat{x}}_{i \rightarrow \mu}$ and $\underline{C}_{i \rightarrow \mu}^x$. As a result we obtain the vectorized AMP for the GLM presented in Algorithm 3. Note that, similarly to GAMP, relaxing the Gaussian assumption on the weight matrix entries to any distribution with finite second moment yields the same algorithm using the Central Limit Theorem.

F.2 State Evolution

We consider the limit $N \rightarrow +\infty$ at fixed $\alpha = M/N$ and a quenched average over the disorder (here the realizations of \underline{X}_0 , \underline{s}_0 , \underline{Y} and \underline{W}), to derive a State Evolution analysis of the previously derived AMP. To this end, our starting point will be the r-BP equations.

F.2.1 State Evolution derivation in mismatched prior and channel setting

Definition of the overlaps The important quantities to follow the dynamic of iterations and fixed points of AMP are the overlaps. Here, they are the $P \times P$ matrices

$$\underline{q} = \frac{1}{N} \sum_{i=1}^N \underline{\hat{x}}_i \underline{\hat{x}}_i^T, \quad \underline{m} = \frac{1}{N} \sum_{i=1}^N \underline{\hat{x}}_i \underline{x}_{0,i}^T, \quad \underline{q}_0 = \frac{1}{N} \sum_{i=1}^N \underline{x}_{0,i} \underline{x}_{0,i}^T. \quad (279)$$

Output parameters Under independent statistics of the entries of \underline{W} and under the assumption of independent incoming messages, the variable $\underline{\omega}_{\mu \rightarrow i}$ defined in (253) is a sum of independent variables and follows a Gaussian distribution by the Central Limit Theorem. Its first and second moments are

$$\mathbb{E}_{\underline{W}} [\underline{\omega}_{\mu \rightarrow i}] = \frac{1}{\sqrt{N}} \sum_{i' \neq i} \mathbb{E}_{\underline{W}} [W_{\mu i'}] \underline{\hat{x}}_{i' \rightarrow \mu} = 0, \quad (280)$$

$$\mathbb{E}_{\underline{W}} [\underline{\omega}_{\mu \rightarrow i} \underline{\omega}_{\mu \rightarrow i}^T] = \frac{1}{N} \sum_{i' \neq i} \sum_{i'' \neq i} \mathbb{E}_{\underline{W}} [W_{\mu i'} W_{\mu i''}] \underline{\hat{x}}_{i' \rightarrow \mu} \underline{\hat{x}}_{i'' \rightarrow \mu}^T \quad (281)$$

$$\begin{aligned} &= \frac{1}{N} \sum_{i' \neq i} \mathbb{E}_{\underline{W}} [W_{\mu i'}^2] \underline{\hat{x}}_{i' \rightarrow \mu} \underline{\hat{x}}_{i' \rightarrow \mu}^T = \frac{1}{N} \sum_{i'=1}^N \underline{\hat{x}}_{i' \rightarrow \mu} \underline{\hat{x}}_{i' \rightarrow \mu}^T + O(1/N) \\ &= \frac{1}{N} \sum_{i=1}^N \underline{\hat{x}}_i \underline{\hat{x}}_i^T - \partial_{\underline{\lambda}} f_{\underline{\lambda}}^x \underline{\sigma}_{\underline{\lambda}} B_{\mu \rightarrow i} \underline{\hat{x}}_i^T - \left(\partial_{\underline{\lambda}} f_{\underline{\lambda}}^x \underline{\sigma}_{\underline{\lambda}} B_{\mu \rightarrow i} \underline{\hat{x}}_i^T \right)^T + O(1/N) \\ &= \frac{1}{N} \sum_{i=1}^N \underline{\hat{x}}_i \underline{\hat{x}}_i^T + O(1/\sqrt{N}) \end{aligned} \quad (282)$$

where we used the facts that the $W_{\mu i}$ -s are independent with zero mean, and that $B_{\mu \rightarrow i}$, defined in (266), is of order $O(1/\sqrt{N})$. Similarly, the variable $\underline{z}_{\mu \rightarrow i} = \sum_{i' \neq i} \frac{W_{\mu i'}}{\sqrt{N}} \underline{x}_{0,i'}$ is Gaussian with first and second moments

$$\mathbb{E}_{\underline{W}} [\underline{z}_{\mu \rightarrow i}] = \frac{1}{\sqrt{N}} \sum_{i' \neq i} \mathbb{E}_{\underline{W}} [W_{\mu i'}] \underline{x}_{0,i'} = 0, \quad (283)$$

$$\mathbb{E}_{\underline{W}} [\underline{z}_{\mu \rightarrow i} \underline{z}_{\mu \rightarrow i}^T] = \frac{1}{N} \sum_{i'=1}^N \underline{x}_{0,i'} \underline{x}_{0,i'}^T + O(1/\sqrt{N}). \quad (284)$$

Furthermore, their covariance is

$$\mathbb{E}_{\underline{W}} [\underline{z}_{\mu \rightarrow i} \underline{\omega}_{\mu \rightarrow i}^T] = \frac{1}{N} \sum_{i' \neq i} \mathbb{E}_{\underline{W}} [W_{\mu i'}^2] \underline{x}_{0,i'} \underline{\hat{x}}_{i' \rightarrow \mu}^T = \frac{1}{N} \sum_{i'=1}^N \underline{x}_{0,i'} \underline{\hat{x}}_{i' \rightarrow \mu}^T + O(1/N) \quad (285)$$

$$= \frac{1}{N} \sum_{i=1}^N \underline{x}_{0,i} \underline{\hat{x}}_i^T - \underline{x}_{0,i} \partial_{\underline{\lambda}} f_{\underline{\lambda}}^x \underline{\sigma}_{\underline{\lambda}} B_{\mu \rightarrow i}^T + O(1/N) \quad (286)$$

$$= \frac{1}{N} \sum_{i=1}^N \underline{x}_{0,i} \underline{\hat{x}}_i^T + O(1/\sqrt{N}). \quad (287)$$

F Multi-value AMP derivation for the GLM

Hence we find that for all μ -s and all i -s, $\omega_{\mu \rightarrow i}$ and $z_{\mu \rightarrow i}$ are approximately jointly Gaussian in the thermodynamic limit following a unique distribution $\mathcal{N}\left(z_{\mu \rightarrow i}, \omega_{\mu \rightarrow i}; \underline{0}, \underline{Q}\right)$ with the block covariance matrix

$$\underline{Q} = \begin{bmatrix} \underline{q}_0 & \underline{m} \\ \underline{m}^\top & \underline{q} \end{bmatrix}. \quad (288)$$

For the variance message $\underline{V}_{\mu \rightarrow i}$, defined in (254), we have

$$\mathbb{E}_{\underline{W}} \left[\underline{V}_{\mu \rightarrow i} \right] = \sum_{i' \neq i} \mathbb{E}_{\underline{W}} \left[\frac{W_{\mu i'}}{N} \right]^2 \underline{C}_{i' \rightarrow \mu}^x = \sum_{i'=1}^N \frac{1}{N} \underline{C}_{i' \rightarrow \mu}^x + O(1/N) \quad (289)$$

$$= \sum_{i'=1}^N \frac{1}{N} \underline{C}_{i'}^x + O(1/\sqrt{N}), \quad (290)$$

where using the developments of $\underline{\lambda}_{i \rightarrow \mu}$ and $\underline{\sigma}_{i \rightarrow \mu}$ (269)-(270), along with the scaling of $\underline{B}_{\mu \rightarrow i}$ in $O(1/\sqrt{N})$ we replaced

$$\underline{C}_{i \rightarrow \mu}^x = f_{\underline{2}}^x(\underline{\lambda}_{i \rightarrow \mu}, \underline{\sigma}_{i \rightarrow \mu}) = f_{\underline{2}}^x(\underline{\lambda}_i, \underline{\sigma}_i) - \partial_{\underline{\lambda}_{\underline{2}}^x} f_{\underline{2}}^x \underline{\sigma}_i \underline{B}_{\mu \rightarrow i}^T = f_{\underline{2}}^x(\underline{\lambda}_i, \underline{\sigma}_i) + O(1/\sqrt{N}). \quad (291)$$

Futhermore, we can check that

$$\lim_{N \rightarrow +\infty} \mathbb{E}_{\underline{W}} \left[\underline{V}_{\mu \rightarrow i}^2 - \mathbb{E}_{\underline{W}} \left[\underline{V}_{\mu \rightarrow i} \right]^2 \right] = 0, \quad (292)$$

meaning that all $\underline{V}_{\mu \rightarrow i}$ concentrate on their identical mean in the thermodynamic limit, which we note

$$\underline{V} = \sum_{i=1}^N \frac{1}{N} \underline{C}_{\underline{i}}^x. \quad (293)$$

Input parameters Here we use the re-parametrization trick to express \underline{y}_μ as a function $g_0(\cdot)$ taking a noise $\epsilon_\mu \sim p_\epsilon(\epsilon_\mu)$ as inputs: $\underline{y}_\mu = g_0(\underline{w}_\mu^\top \underline{X}_{\underline{0}}, \epsilon_\mu)$. Following (267)-(266) and (276),

$$\underline{\sigma}_i^{-1} \underline{\lambda}_i = \sum_{\mu=1}^M \frac{W_{\mu i}}{\sqrt{N}} g_{\text{out}} \left(\underline{y}_\mu, \omega_{\mu \rightarrow i}, \underline{V}_{\mu \rightarrow i} \right) \quad (294)$$

$$= \sum_{\mu=1}^M \frac{W_{\mu i}}{\sqrt{N}} g_{\text{out}} \left(g_0 \left(\sum_{i' \neq i} \frac{W_{\mu i'}}{\sqrt{N}} x_{0, i'} + \frac{W_{\mu i}}{\sqrt{N}} x_{0, i}, \epsilon_\mu \right), \omega_{\mu \rightarrow i}, \underline{V}_{\mu \rightarrow i} \right) \quad (295)$$

$$= \sum_{\mu=1}^M \frac{W_{\mu i}}{\sqrt{N}} g_{\text{out}} \left(g_0 \left(\sum_{i' \neq i} \frac{W_{\mu i'}}{\sqrt{N}} x_{0, i'}, \epsilon_\mu \right), \omega_{\mu \rightarrow i}, \underline{V}_{\mu \rightarrow i} \right) + \sum_{\mu=1}^M \frac{W_{\mu i}^2}{N} \partial_z g_{\text{out}} \left(g_0 \left(z_{\mu \rightarrow i}, \epsilon_\mu \right), \omega_{\mu \rightarrow i}, \underline{V}_{\mu \rightarrow i} \right) x_{0, i}. \quad (296)$$

The first term is again a sum of independent random variables, given the $W_{\mu i}$ are i.i.d. with zero mean, of which the messages of type $\mu \rightarrow i$ are assumed independent. The second term has non-zero mean and can be shown to concentrate. Finally recalling that all $\underline{V}_{\mu \rightarrow i}$ also concentrate on \underline{V} we obtain the distribution

$$\underline{\sigma}_i^{-1} \underline{\lambda}_i \sim \mathcal{N} \left(\underline{\sigma}_i^{-1} \underline{\lambda}_i; \alpha \hat{m} x_{0, i}, \sqrt{\alpha \hat{q}} I_P \right) \quad (297)$$

with

$$\hat{q} = \int d\epsilon p_\epsilon(\epsilon) ds_0 p_{s_0}(s_0) \int d\omega d\underline{z} \mathcal{N}(\underline{z}, \omega; \underline{0}, \underline{Q}) \underline{g}_{\text{out}}(g_0(\underline{z}, \epsilon), \omega, \underline{V}) \times \quad (298)$$

$$\underline{g}_{\text{out}}(g_0(\underline{z}, \epsilon), \omega, \underline{V})^T, \quad (299)$$

$$\hat{m} = \int d\epsilon p_\epsilon(\epsilon) ds_0 p_{s_0}(s_0) \int d\omega d\underline{z} \mathcal{N}(\underline{z}, \omega; \underline{0}, \underline{Q}) \partial_{\underline{z}} \underline{g}_{\text{out}}(g_0(\underline{z}, \epsilon), \omega, \underline{V}). \quad (299)$$

For the inverse variance $\underline{\sigma}_i^{-1}$ one can check again that it concentrates on its mean

$$\underline{\sigma}_i^{-1} = \sum_{\mu=1}^M \frac{W_{\mu i}^2}{N} \partial_{\omega} \underline{g}_{\text{out}}(\underline{y}_\mu, \omega_{\mu \rightarrow i}, \underline{V}_{\mu \rightarrow i}) \simeq \alpha \hat{\chi}, \quad (300)$$

$$\hat{\chi} = - \int d\epsilon p_\epsilon(\epsilon) ds_0 p_{s_0}(s_0) \int d\omega d\underline{z} \mathcal{N}(\underline{z}, \omega; \underline{0}, \underline{Q}) \partial_{\omega} \underline{g}_{\text{out}}(g_0(\underline{z}, \epsilon), \omega, \underline{V}). \quad (301)$$

Closing the equations These statistics of the input parameters must ensure that consistently

$$\underline{V} = \frac{1}{N} \sum_{i=1}^N \underline{C}_{\underline{=i}}^x = \mathbb{E}_{\lambda, \underline{\sigma}} \left[\underline{f}_{\underline{=2}}^x(\lambda, \underline{\sigma}) \right], \quad (302)$$

$$\underline{q} = \frac{1}{N} \sum_{i=1}^N \hat{\underline{x}}_i \hat{\underline{x}}_i^\top = \mathbb{E}_{\lambda, \underline{\sigma}} \left[\underline{f}_{\underline{=1}}^x(\lambda, \underline{\sigma}) \underline{f}_{\underline{=1}}^x(\lambda, \underline{\sigma})^\top \right], \quad (303)$$

$$\underline{m} = \frac{1}{N} \sum_{i=1}^N \hat{\underline{x}}_i \underline{x}_{0,i}^\top = \mathbb{E}_{\lambda, \underline{\sigma}} \left[\underline{f}_{\underline{=1}}^x(\lambda, \underline{\sigma}) \underline{x}_{0,i}^\top \right], \quad (304)$$

which gives upon expressing the computation of the expectations

$$\underline{V} = \int d\underline{x}_0 p_{x_0}(\underline{x}_0) \int \mathcal{D}\xi \underline{f}_{\underline{=2}}^x \left((\alpha \hat{\chi})^{-1} \left(\sqrt{\alpha \hat{q}} \xi + \alpha \hat{m} \underline{x}_0 \right); (\alpha \hat{\chi})^{-1} \right), \quad (305)$$

$$\underline{m} = \int d\underline{x}_0 p_{x_0}(\underline{x}_0) \int \mathcal{D}\xi \underline{f}_{\underline{=1}}^x \left((\alpha \hat{\chi})^{-1} \left(\sqrt{\alpha \hat{q}} \xi + \alpha \hat{m} \underline{x}_0 \right); (\alpha \hat{\chi})^{-1} \right) \underline{x}_0^\top, \quad (306)$$

$$\underline{q} = \int d\underline{x}_0 p_{x_0}(\underline{x}_0) \int \mathcal{D}\xi \underline{f}_{\underline{=1}}^x \left((\alpha \hat{\chi})^{-1} \left(\sqrt{\alpha \hat{q}} \xi + \alpha \hat{m} \underline{x}_0 \right); (\alpha \hat{\chi})^{-1} \right) \times \quad (307)$$

$$\underline{f}_{\underline{=1}}^x \left((\alpha \hat{\chi})^{-1} \left(\sqrt{\alpha \hat{q}} \xi + \alpha \hat{m} \underline{x}_0 \right); (\alpha \hat{\chi})^{-1} \right)^\top.$$

The State Evolution analysis of the GLM on the vector variables finally consists in iterating alternatively the equations (298), (299), (301), and the equations (305), (306) (307) until convergence.

Performance analysis The mean squared error (MSE) on the reconstruction of \underline{X} by the AMP algorithm is then predicted by

$$\text{MSE}(\underline{X}) = q - 2m + q_0, \quad (308)$$

where the scalar values used here correspond to the (unique) value of the diagonal elements of the corresponding overlap matrices. This MSE can be computed throughout the iterations of State Evolution. Remarkably, the State Evolution MSEs follow precisely the MSE of the cal-AMP predictors along the iterations of the algorithm provided the procedures are initialized consistently. A random initialization of $\hat{\underline{x}}_i$ in cal-AMP corresponds to an initialization of zero overlap $m = 0$, $\nu = 0$, with variance of the priors $q = q_0$ in the State Evolution.

F References

F.2.2 Bayes optimal State Evolution

The SE equations can be greatly simplified in the Bayes optimal setting where the statistical model used by the student (priors p_x and p_s , and channel p_{out}) is known to match the teacher. In this case, the true unknown signal \underline{x}_0 is in some sense statistically equivalent to the estimate $\hat{\underline{x}}$ coming from the posterior. More precisely one can prove the Nishimori identities [OH91, Iba99, Nis01] (or [KKM⁺16] for a concise demonstration and discussion) implying that

$$\underline{q} = \underline{m}, \quad \underline{V} = \underline{q} - \underline{m}, \quad \underline{\hat{q}} = \underline{\hat{m}} = \underline{\hat{x}} \quad \text{and} \quad r = \nu. \quad (309)$$

As a result the State Evolution reduces to a set of two equations

$$\underline{m} = \int d\underline{x}_0 p_{x_0}(\underline{x}_0) \int \mathcal{D}\underline{\xi} f_{-1}^x \left((\alpha \underline{\hat{m}})^{-1} \left(\sqrt{\alpha \underline{\hat{m}}} \underline{\xi} + \alpha \underline{\hat{m}} \underline{x}_0 \right); (\alpha \underline{\hat{m}})^{-1} \right) \underline{x}_0^\top \quad (310)$$

$$\underline{\hat{m}} = \int d\underline{\epsilon} p_\epsilon(\underline{\epsilon}) d s_0 p_{s_0}(s_0) \int d\underline{\omega} d\underline{z} \mathcal{N}(\underline{z}, \underline{\omega}; \underline{0}, \underline{Q}) g_{\text{out}} \left(g_0(\underline{z}, \underline{\epsilon}), \underline{\omega}, \underline{q} - \underline{m} \right) \times \quad (311)$$

$$g_{\text{out}} \left(g_0(\underline{z}, \underline{\epsilon}), \underline{\omega}, \underline{q} - \underline{m} \right),$$

with the block covariance matrix

$$\underline{Q} = \begin{bmatrix} \underline{q} & \underline{m} \\ \underline{m}^\top & \underline{m} \end{bmatrix}. \quad (312)$$

References

- [AAKZ19] Alia Abbata, Benjamin Aubin, Florent Krzakala, and Lenka Zdeborová. Rademacher complexity and spin glasses: A link between the replica and statistical theories of learning. *arXiv preprint*, 1912.02729, 2019.
- [AFUZ19] Fabrizio Antenucci, Silvio Franz, Pierfrancesco Urbani, and Lenka Zdeborová. Glassy Nature of the Hard Phase in Inference Problems. *Physical Review X*, 9(1):11020, 2019.
- [AGS85] Daniel J. Amit, Hanoach Gutfreund, and Haim Sompolinsky. Storing infinite numbers of patterns in a spin-glass model of neural networks. *Physical Review Letters*, 55(14):1530–1533, 1985.
- [AHS85] David H Ackley, Geoffrey E Hinton, and J Sejnowski. A Learning Algorithm for Boltzmann Machine. *Cognitive Science*, 9:147–169, 1985.
- [AKUZ19] Fabrizio Antenucci, Florent Krzakala, Pierfrancesco Urbani, and Lenka Zdeborová. Approximate survey propagation for statistical inference. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(2):023401, feb 2019.
- [ALG13] Madhu Advani, Subhaneil Lahiri, and Surya Ganguli. Statistical mechanics of complex neural systems and high dimensional data. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(03):P03014, mar 2013.
- [ALM⁺19] Benjamin Aubin, Bruno Loureiro, Antoine Maillard, Florent Krzakala, and Lenka Zdeborová. The spiked matrix model with generative priors. *arXiv preprint*, 1905.12385, 2019.
- [AMB⁺18] Benjamin Aubin, Antoine Maillard, Jean Barbier, Florent Krzakala, Nicolas Macris, and Lenka Zdeborová. The committee machine: Computational to statistical gaps in learning a two-layers neural network. In *Neural Information Processing Systems 2018*, number NeurIPS, pages 1–44, 2018.
- [AS17] Madhu S. Advani and Andrew M. Saxe. High-dimensional dynamics of generalization error in neural networks. *arXiv preprint*, 1710.03667:1–32, 2017.
- [BBC⁺16] Carlo Baldassi, Christian Borgs, Jennifer T Chayes, Alessandro Ingrosso, Carlo Lucibello, Luca Saglietti, and Riccardo Zecchina. Unreasonable effectiveness of learning neural networks: From accessible states and robust ensembles to basic algorithmic schemes. *Proceedings of the National Academy of Sciences of the United States of America*, 113(48):E7655–E7662, 2016.

- [BBW⁺13] Tamara Broderick, Nicholas Boyd, Andre Wibisono, Ashia C. Wilson, and Michael I. Jordan. Streaming Variational Bayes. *Neural Information Processing Systems 2013*, pages 1–9, 2013.
- [BCV13] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation Learning: A Review and New Perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, aug 2013.
- [Bet35] Hans Bethe. Statistical Theory of Superlattices. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 150(871):552, 1935.
- [BGK⁺18] Carlo Baldassi, Federica Gerace, Hilbert J. Kappen, Carlo Lucibello, Luca Saglietti, Enzo Tartaglione, and Riccardo Zecchina. Role of Synaptic Stochasticity in Training Low-Precision Neural Networks. *Physical Review Letters*, 120(26):268103, 2018.
- [BGST17] Adriano Barra, Giuseppe Genovese, Peter Sollich, and Daniele Tantari. Phase transitions in restricted Boltzmann machines with generic priors. *Physical Review E*, 96(4):1–5, 2017.
- [BGST18] Adriano Barra, Giuseppe Genovese, Peter Sollich, and Daniele Tantari. Phase diagram of restricted Boltzmann machines and generalized Hopfield networks with arbitrary priors. *Physical Review E*, 97(2), 2018.
- [BIL⁺15] Carlo Baldassi, Alessandro Ingrosso, Carlo Lucibello, Luca Saglietti, and Riccardo Zecchina. Subdominant Dense Clusters Allow for Simple Learning and High Computational Performance in Neural Networks with Discrete Synapses. *Physical Review Letters*, 115(12):1–5, 2015.
- [BJPD17] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G. Dimakis. Compressed Sensing using Generative Models. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 537–546, 2017.
- [BJSG⁺18] Marco Baity-Jesi, Levent Sagun, Mario Geiger, Stefano Spigler, G. Ben Arous, Chiara Cammarota, Yann LeCun, Matthieu Wyart, and Giulio Biroli. Comparing Dynamics: Deep Neural Networks versus Glassy Systems. *Proceedings of the 35th International Conference on Machine Learning*, PMLR(80):314–323, 2018.
- [BKM17] David M. Blei, Alp Kucukelbir, and Jon D. McAuliffe. Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- [BKM⁺18] Jean Barbier, Florent Krzakala, Nicolas Macris, Léo Miolane, and Lenka Zdeborová. Phase Transitions, Optimal Errors and Optimality of Message-Passing in Generalized Linear Models. *Proceedings of the 31st Conference On Learning Theory*, PMLR 75:728–731, aug 2018.
- [BLPL07] Yoshua Bengio, Pascal Lamblin, Dan Popovici, and Hugo Larochelle. Greedy Layer-Wise Training of Deep Networks. *Advances in neural information processing systems*, 19(1):153, 2007.
- [BM02] Peter L. Bartlett and Shahar Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- [Bol14] Erwin Bolthausen. An iterative construction of solutions of the tap equations for the Sherrington-Kirkpatrick model. *Communications in Mathematical Physics*, 325(1):333–366, 2014.
- [Bot10] Léon Bottou. Large-Scale Machine Learning with Stochastic Gradient Descent. In *Proceedings of COMPSTAT’2010*, pages 177–186. Physica-Verlag HD, Heidelberg, 2010.
- [BS95] Michael Biehl and H. Schwarze. Learning by on-line gradient descent. *J. Phys. A. Math. Gen.*, 28(3):643–656, 1995.
- [CB18a] Lénaïc Chizat and Francis Bach. A Note on Lazy Training in Supervised Differentiable Programming. *Arxiv preprint arXiv11021808*, 812.07956, 2018.
- [CB18b] Lénaïc Chizat and Francis Bach. On the global convergence of gradient descent for overparameterized models using optimal transport. In *Advances in Neural Information Processing Systems 31*, pages 3040–3050, 2018.

F References

- [CCC⁺19] Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, Laurent Daudet, Maria Schuld, Naftali Tishby, Leslie Vogt-Maranto, and Lenka Zdeborová. Machine learning and the physical sciences. *Reviews of Modern Physics*, 91(4), 2019.
- [CCFM05] Tommaso Castellani, Andrea Cavagna, Dipartimento Fisica, and Piazzale Aldo Moro. Spin-Glass Theory for Pedestrians. *Journal of Statistical Mechanics: Theory and Experiment*, (5):215–266, 2005.
- [CCLS19] Uri Cohen, SueYeon Chung, Daniel D Lee, and Haim Sompolsky. Separability and Geometry of Object Manifolds in Deep Neural Networks. *bioRxiv*, 10.1101/64, 2019.
- [CCS⁺17] Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann Lecun, Carlo Baldassi, Christian Borgs, Jennifer Chayes, Levent Sagun, and Riccardo Zecchina. Entropy SGD: Biasing gradient descent into wide valleys. *International Conference on Learning Representations (ICLR)*, pages 1–19, 2017.
- [CHM⁺15] Anna Choromanska, Mikael Henaff, Michael Mathieu, Gérard Ben Arous, and Yann LeCun. The Loss Surfaces of Multilayer Networks. In *Artificial Intelligence and Statistics*, volume 38 PMLR, pages 192–204., 2015.
- [CHS93] A. Crisanti, H. Horner, and H. J. Sommers. The spherical spin interaction spin-glass model. *Zeitschrift für Physik B Condensed Matter*, 92(2):257–271, jun 1993.
- [CK93] L. F. Cugliandolo and J. Kurchan. Analytical solution of the off-equilibrium dynamics of a long-range spin-glass model. *Physical Review Letters*, 71(1):173–176, jul 1993.
- [CLS18] Sueyeon Chung, Daniel D. Lee, and Haim Sompolsky. Classification and Geometry of General Perceptual Manifolds. *Physical Review X*, 8(3):31003, 2018.
- [CNL11] Adam Coates, Andrew Y Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *International Conference on Artificial Intelligence and Statistics*, pages 215–223, 2011.
- [ÇO19] Burak Çakmak and Manfred Oppert. Memory-free dynamics for the Thouless-Anderson-Palmer equations of Ising models with arbitrary rotation-invariant ensembles of random coupling matrices. *Physical Review E*, 99(6), 2019.
- [CRI10] KyungHyun Cho, Tapani Raiko, and Alexander Ilin. Parallel tempering is efficient for learning restricted Boltzmann machines. *Proceedings of the International Joint Conference on Neural Networks*, 2010.
- [CS88] A. Crisanti and Haim Sompolsky. Dynamics of spin systems with randomly asymmetric bonds: Ising spins and Glauber dynamics. *Physical Review A*, 37(12):4865–4874, jun 1988.
- [Cur95] Pierre Curie. Lois expérimentales du magnétisme. Propriétés magnétiques des corps à diverses températures. *Ann Chim Phys.*, 5:289–405, 1895.
- [Cyb89] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems*, 2(4):303–314, dec 1989.
- [Dan17] Amit Daniely. Depth Separation for Neural Networks, 2017.
- [DCB⁺10] Guillaume Desjardins, Aaron Courville, Yoshua Bengio, Pascal Vincent, and Olivier Delalleau. Parallel Tempering for Training of Restricted Boltzmann Machines. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 9, pages 145–152, 2010.
- [DFF17] Aurélien Decelle, Giancarlo Fissore, and Cyril Furtlehner. Spectral dynamics of learning in restricted Boltzmann machines. *EPL (Europhysics Letters)*, 119(6):60001, sep 2017.
- [DFF18] A. Decelle, G. Fissore, and C. Furtlehner. Thermodynamics of Restricted Boltzmann Machines and Related Learning Dynamics. *Journal of Statistical Physics*, 172(6):1576–1608, 2018.
- [DHD12] A. Dremeau, Cédric Herzet, and Laurent Daudet. Boltzmann Machine and Mean-Field Approximation for Structured Sparse Decompositions. *IEEE Transactions on Signal Processing*, 60(7):3425–3438, jul 2012.

- [DMM09] 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.
- [Don06] D.L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, apr 2006.
- [DY83] C De Dominicis and A P Young. Weighted averages and order parameters for the infinite range Ising spin glass. *Journal of Physics A: Mathematical and General*, 16(9):2063–2075, 1983.
- [EV01] A Engel and C Van den Broeck. *Statistical Mechanics of Learning*. Cambridge University Press, Cambridge, 2001.
- [FRS18] Alyson K. Fletcher, Sundeep Rangan, and Philip Schniter. Inference in Deep Networks in High Dimensions. *2018 IEEE International Symposium on Information Theory (ISIT)*, 1(8):1884–1888, jun 2018.
- [Gal93] Conrad C Galland. The limitations of deterministic Boltzmann machine learning. *Network*, 4:355–379, 1993.
- [Gar87] Elisabeth Gardner. Maximum Storage Capacity in Neural Networks. *Europhysics Letters (EPL)*, 4(4):481–485, aug 1987.
- [Gar88] Elisabeth Gardner. The space of interactions in neural network models. *Journal of Physics A: General Physics*, 21(1):257–270, 1988.
- [GAS⁺19] Sebastian Goldt, Madhu S. Advani, Andrew M. Saxe, Florent Krzakala, and Lenka Zdeborová. Dynamics of stochastic gradient descent for two-layer neural networks in the teacher-student setup. In *Neural Information Processing Systems*, 2019.
- [GBC16] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- [GBKZ19] Marylou Gabrié, Jean Barbier, Florent Krzakala, and Lenka Zdeborová. Blind calibration for compressed sensing: State evolution and an online algorithm. *arXiv preprint*, 1910.00285, 2019.
- [GCC⁺19] Dar Gilboa, Bo Chang, Minmin Chen, Greg Yang, Samuel S. Schoenholz, Ed H. Chi, and Jeffrey Pennington. Dynamical Isometry and a Mean Field Theory of LSTMs and GRUs. *arXiv preprint*, 1901.08987, 2019.
- [GJS⁺19] Mario Geiger, Arthur Jacot, Stefano Spigler, Franck Gabriel, Levent Sagun, Stéphane D’Ascoli, Giulio Biroli, Clément Hongler, and Matthieu Wyart. Scaling description of generalization with number of parameters in deep learning. *arXiv preprint*, 1901.01608:1–22, 2019.
- [GMKZ19] Sebastian Goldt, Marc Mézard, Florent Krzakala, and Lenka Zdeborová. Modelling the influence of data structure on learning in neural networks. *arXiv preprint*, 1909.11500, 2019.
- [GML⁺18] Marylou Gabrié, Andre Manoel, Clément Luneau, Jean Barbier, Nicolas Macris, Florent Krzakala, and Lenka Zdeborová. Entropy and mutual information in models of deep neural networks. In S Bengio, H Wallach, H Larochelle, K Grauman, N Cesa-Bianchi, and R Garnett, editors, *Advances in Neural Information Processing Systems 31*, number Nips, pages 1826—1836. Curran Associates, Inc., mar 2018.
- [GPAM⁺14] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks. *Neural Information Processing Systems*, pages 1–9, 2014.
- [GPEB19] Philipp Grohs, Dmytro Perekrestenko, Dennis Elbrächter, and Helmut Bölcskei. Deep Neural Network Approximation Theory. *arXiv preprint*, 1901.02220, 2019.
- [GSJW19] Mario Geiger, Stefano Spigler, Arthur Jacot, and Matthieu Wyart. Disentangling feature and lazy learning in deep neural networks: an empirical study. *arXiv preprint*, 1906.08034, 2019.
- [GTK15] Marylou Gabrié, Eric W. Tramel, and Florent Krzakala. Training Restricted Boltzmann Machines via the Thouless-Anderson-Palmer Free Energy. *Advances in Neural Information Processing Systems 28*, pages 640—648, jun 2015.

F References

- [GY99] A Georges and J S Yedidia. How to expand around mean-field theory using high-temperature expansions. *Journal of Physics A: Mathematical and General*, 24(9):2173–2192, 1999.
- [Hin89] Geoffrey E. Hinton. Deterministic Boltzmann learning performs steepest descent in weight-space. *Neural computation*, 1(1):143–150, 1989.
- [Hin02] Geoffrey E. Hinton. Training products of experts by minimizing Contrastive divergence. *Neural computation*, 14:1771–1800, 2002.
- [HK14] Haiping Huang and Yoshiyuki Kabashima. Origin of the computational hardness for learning with binary synapses. *Physical Review E*, 15:052813, 2014.
- [HLV18] Paul Hand, Oscar Leong, and Vladislav Voroninski. Phase Retrieval Under a Generative Prior. In *Neural Information Processing Systems 2018*, number NeurIPS, 2018.
- [Hop82] J J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America*, 79(8):2554–2558, 1982.
- [Hor91] Kurt Hornik. Approximation Capabilities of Multilayer Neural Network. *Neural Networks*, 4(1989):251–257, 1991.
- [HOT06] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. A fast learning algorithm for deep belief nets. *Neural Computation*, 18:1527–1554, 2006.
- [HS06] Geoffrey E. Hinton and Ruslan R Salakhutdinov. Reducing the Dimensionality of Data with Neural Networks. *Science*, 313(5786):504–507, jul 2006.
- [HS09] Geoffrey E Hinton and Ruslan R Salakhutdinov. Replicated softmax: an undirected topic model. In *Advances in neural information processing systems*, pages 1607–1614, 2009.
- [Hua17] Haiping Huang. Statistical mechanics of unsupervised feature learning in a restricted Boltzmann machine with binary synapses. *Journal of Statistical Mechanics: Theory and Experiment*, 2017(5), 2017.
- [HV18] Paul Hand and Vladislav Voroninski. Global Guarantees for Enforcing Deep Generative Priors by Empirical Risk. In *Proceedings of the 31st Conference On Learning Theory*, volume 75, pages 970–978, 2018.
- [HWK13] Haiping Huang, K. Y. Michael Wong, and Yoshiyuki Kabashima. Entropy landscape of solutions in the binary perceptron problem. *Journal of Physics A: Mathematical and Theoretical*, 46(37), 2013.
- [Iba99] Yukito Iba. The Nishimori line and Bayesian statistics. *Journal of Physics A: Mathematical and General*, 32(21):3875–3888, may 1999.
- [JGH18] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural Tangent Kernel: Convergence and Generalization in Neural Networks. In *Advances in neural information processing systems*, number 5, 2018.
- [JKA⁺17] Stanisław Jastrzębski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. Three Factors Influencing Minima in SGD. *arXiv preprint*, 1711.04623:1–21, 2017.
- [Kab03] Yoshiyuki Kabashima. A CDMA multiuser detection algorithm on the basis of belief propagation. *Journal of Physics A: Mathematical and General*, 36(43):11111–11121, oct 2003.
- [Kab08] Yoshiyuki Kabashima. Inference from correlated patterns: A unified theory for perceptron learning and linear vector channels. *Journal of Physics: Conference Series*, 95(1):012001, jan 2008.
- [KF09] Daphne Koller and Nir Friedman. *Probabilistic Graphical Models Principles and Techniques*. MIT Press, 2009.
- [KKM⁺16] Yoshiyuki Kabashima, Florent Krzakala, Marc Mézard, Ayaka Sakata, and Lenka Zdeborová. Phase Transitions and Sample Complexity in Bayes-Optimal Matrix Factorization. *IEEE Transactions on Information Theory*, 62(7):4228–4265, 2016.

F.2 References

- [KMS⁺12] Florent Krzakala, Marc Mézard, Francois Sausset, Yifan Sun, and Lenka Zdeborová. Probabilistic reconstruction in compressed sensing: algorithms, phase diagrams, and threshold achieving matrices. *J. Stat. Mech*, 2012.
- [KR98] Hilbert J Kappen and Francisco De Borja Rodríguez. Boltzmann Machine Learning Using Mean Field Theory and Linear Response Correction. *Advances in Neural Information Processing Systems 10*, pages 280–286, 1998.
- [Kra06] Werner Krauth. Statistical mechanics: algorithms and computations. *OUP Oxford*, 13, 2006.
- [KS98] Yoshiyuki Kabashima and David Saad. Belief propagation vs. TAP for decoding corrupted messages. *Europhysics Letters*, 44(5):668–674, 1998.
- [KS16] Jonathan Kadmon and Haim Sompolinsky. Optimal architectures in a solvable model of deep networks. In *Advances in Neural Information Processing Systems 29*, number Nips, 2016.
- [KTO18] Tatsuro Kawamoto, Masashi Tsubaki, and Tomoyuki Obuchi. Mean-field theory of graph neural networks in graph partitioning. In *Advances in Neural Information Processing Systems*, 2018.
- [KU04] Yoshiyuki Kabashima and Shinsuke Uda. A BP-based algorithm for performing Bayesian inference in large perceptron-type networks, 2004.
- [KV14] Yoshiyuki Kabashima and Mikko Vehkaperä. Signal recovery using expectation consistent approximation for linear observations. *IEEE International Symposium on Information Theory - Proceedings*, pages 226–230, 2014.
- [KW14] Diederik P Kingma and Max Welling. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations (ICLR)*, number Ml, pages 1–14, dec 2014.
- [LB08] Hugo Larochelle and Yoshua Bengio. Classification using discriminative restricted Boltzmann machines. In *Proceedings of the 25th international conference on Machine learning*, pages 536–543. ACM, 2008.
- [LBH15] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [LKZ16] Thibault Lesieur, Florent Krzakala, and Lenka Zdeborová. MMSE of probabilistic low-rank matrix estimation: Universality with respect to the output channel. *2015 53rd Annual Allerton Conference on Communication, Control, and Computing, Allerton 2015*, pages 680–687, 2016.
- [LKZ17] 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):1–64, 2017.
- [LM99] S. Li and K. Y. Michael Wong. Statistical dynamics of batch learning. *Advances in Neural Information Processing Systems*, pages 286–292, 1999.
- [LS18] Bo Li and David Saad. Exploring the Function Space of Deep-Learning Machines. *Physical Review Letters*, 120(24), 2018.
- [LS20] Bo Li and David Saad. Large deviation analysis of function sensitivity in random deep neural networks. *Journal of Physics A: Mathematical and Theoretical*, 53(10):104002, mar 2020.
- [LVMC18] Andrey Y. Lokhov, Marc Vuffray, Sidhant Misra, and Michael Chertkov. Optimal structure and parameter learning of Ising models. *Science Advances*, 4(3):1–21, 2018.
- [LXS⁺19] Jaehoon Lee, Lechao Xiao, Samuel S. Schoenholz, Yasaman Bahri, Jascha Sohl-Dickstein, Jeffrey Pennington, and M L Feb. Wide Neural Networks of Any Depth Evolve as Linear Models Under Gradient Descent. *arXiv preprint*, 1902.06720, 2019.
- [Méz17] Marc Mézard. Mean-field message-passing equations in the Hopfield model and its generalizations. *Physical Review E*, 95(2):1–15, 2017.
- [MFC⁺19] Antoine Maillard, Laura Foini, Alejandro Lage Castellanos, Florent Krzakala, Marc Mézard, and Lenka Zdeborová. High-temperature Expansions and Message Passing Algorithms. *arXiv preprint*, (1906.08479), 2019.

F References

- [MH76] T. Morita and T. Horiguchi. Exactly solvable model of a quantum spin glass. *Solid State Communications*, 19(2):833–835, 1976.
- [Min01] Thomas P Minka. A family of algorithms for approximate Bayesian inference. *PhD Thesis*, 2001.
- [MKMZ17] Andre Manoel, Florent Krzakala, Marc Mézard, and Lenka Zdeborová. Multi-layer generalized linear estimation. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pages 2098–2102. IEEE, jun 2017.
- [MKTZ18] Andre Manoel, Florent Krzakala, Eric W. Tramel, and Lenka Zdeborová. Streaming Bayesian inference: Theoretical limits and mini-batch approximate message-passing. *55th Annual Allerton Conference on Communication, Control, and Computing, Allerton 2017*, 2018-Janua(i):1048–1055, 2018.
- [MKUZ19] Stefano Sarao Mannelli, Florent Krzakala, Pierfrancesco Urbani, and Lenka Zdeborová. Passed & Spurious: analysing descent algorithms and local minima in spiked matrix-tensor model. In *Proceedings of the 36th International Conference on Machine Learning*, pages PMLR 97:4333–4342, 2019.
- [MM09] Marc Mézard and Andrea Montanari. *Information, Physics, and Computation*. Oxford University Press, 2009.
- [MMM19] Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. *arXiv preprint*, 1902.06015, 2019.
- [MMN18] Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 15(33):E7665–E7671, 2018.
- [MP67] V A Marčenko and L A Pastur. DISTRIBUTION OF EIGENVALUES FOR SOME SETS OF RANDOM MATRICES. *Mathematics of the USSR-Sbornik*, 1(4):457–483, apr 1967.
- [MP01] Marc Mézard and Giorgio Parisi. The Bethe lattice spin glass revisited. *The European Physical Journal B*, 20(2):217–233, mar 2001.
- [MPV86] Marc Mézard, Giorgio Parisi, and Miel Virasoro. *Spin Glass Theory and Beyond*, volume 9 of *World Scientific Lecture Notes in Physics*. WORLD SCIENTIFIC, nov 1986.
- [MPZ02] Marc Mézard, Giorgio Parisi, and Riccardo Zecchina. Analytic and Algorithmic Solution of Random Satisfiability Problems. *Science*, 297(5582):812–815, aug 2002.
- [MV18] Dustin G. Mixon and Soledad Villar. SUNLayer: Stable denoising with generative networks. *arXiv preprint*, 1803.09319, 2018.
- [MWD⁺18] Pankaj Mehta, Ching-hao Wang, Alexandre G. R. Day, Clint Richardson, Charles K. Fisher, Unlearn Ai, San Francisco, David J. Schwab, Fifth Ave, New York, Marin Bukov, Ching-hao Wang, Alexandre G. R. Day, Clint Richardson, Charles K. Fisher, and David J. Schwab. A high-bias, low-variance introduction to Machine Learning for physicists. *arXiv preprint*, 1803.08823, 2018.
- [MZ95] Rémi Monasson and Riccardo Zecchina. Weight Space Structure and Internal Representations: a Direct Approach to Learning and Generalization in Multilayer Neural Networks. *Physical Review Letters*, 75(12):2432–2435, 1995.
- [MZ04] Rémi Monasson and Riccardo Zecchina. Learning and Generalization Theories of Large Committee-Machines. *Modern Physics Letters B*, 09(30):1887–1897, 2004.
- [Nis01] Hidetoshi Nishimori. *Statistical Physics of Spin Glasses and Information Processing: An Introduction*. Clarendon Press, 2001.
- [NXL⁺19] Roman Novak, Lechao Xiao, Jaehoon Lee, Yasaman Bahri, Greg Yang, Jiri Hron, Daniel A. Abolafia, Jeffrey Pennington, and Jascha Sohl-Dickstein. Bayesian Deep Convolutional Networks with Many Channels are Gaussian Processes. In *International Conference on Learning Representations*, 2019.
- [NZB17] H. Chau Nguyen, Riccardo Zecchina, and Johannes Berg. Inverse statistical problems: from the inverse Ising problem to data science. *Advances in Physics*, 66(3):197–261, 2017.

- [OÇW16] Manfred Oppor, Burak Çakmak, and Ole Winther. A theory of solving TAP equations for Ising models with general invariant random matrices. *Journal of Physics A: Mathematical and Theoretical*, 49(11):114002, mar 2016.
- [OH91] Manfred Oppor and David Haussler. Calculation of the learning curve of Bayes optimal classification algorithm for learning a perceptron with noise. In *COLT '91 Proceedings of the fourth annual workshop on Computational learning theory*, pages 75–87, 1991.
- [Opp95] Manfred Oppor. Statistical Mechanics of Learning : Generalization. *The Handbook of Brain Theory and Neural Networks*, page 20, 1995.
- [OS01] Manfred Oppor and David Saad. *Advanced mean field methods: Theory and practice*. MIT press, 2001.
- [OW96] Manfred Oppor and Ole Winther. Mean field approach to bayes learning in feed-forward neural networks. *Physical Review Letters*, 76(11):1964–1967, 1996.
- [OW99a] Manfred Oppor and Ole Winther. A Bayesian approach to on-line learning. In *On-line learning in neural networks*, pages 363–378. Cambridge University Press, 1999.
- [OW99b] Manfred Oppor and Ole Winther. Mean field methods for classification with Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 2–8, 1999.
- [OW01a] Manfred Oppor and Ole Winther. Adaptive and self-averaging Thouless-Anderson-Palmer mean-field theory for probabilistic modeling. *Physical Review E*, 64(5):056131, oct 2001.
- [OW01b] Manfred Oppor and Ole Winther. Tractable approximations for probabilistic models: The adaptive Thouless-Anderson-Palmer mean field approach. *Physical Review Letters*, 86(17):3695–3699, 2001.
- [OW05] Manfred Oppor and Ole Winther. Expectation consistent free energies for approximate inference. *Advances in Neural Information Processing Systems*, 17:1001–1008, 2005.
- [PA87] Carsten Peterson and James R. Anderson. A mean field theory learning algorithm for neural networks. *Complex systems*, 1:995–1019, 1987.
- [Pea88] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems*. Elsevier, 1988.
- [Ple82] T Plefka. Convergence condition of the TAP equation for the infinite-ranged Ising spin glass model. *Journal of Physics A: Mathematical and General*, 15(6):1971–1978, 1982.
- [PLR⁺16] Ben Poole, Subhaneil Lahiri, Maithra Raghu, Jascha Sohl-dickstein, and Surya Ganguli. Exponential expressivity in deep neural networks through transient chaos. *Advances in Neural Information Processing Systems*, (Nips):1–9, jun 2016.
- [PP95] Giorgio Parisi and Marc Potters. Mean-field equations for spin models with orthogonal interaction matrices. *Journal of Physics A: Mathematical and General*, 28(18):5267–5285, sep 1995.
- [PSDG14] Ben Poole, Jascha Sohl-Dickstein, and Surya Ganguli. Analyzing noise in autoencoders and deep networks. *arXiv preprint*, 1406.1831, 2014.
- [PSRF19] Parthe Pandit, Mojtaba Sahraee, Sundeep Rangan, and Alyson K. Fletcher. Asymptotics of MAP Inference in Deep Networks. *arXiv preprint*, 1903.01293, 2019.
- [Ran11] Sundeep Rangan. Generalized Approximate Message Passing for Estimation with Random Linear Mixing. In *2011 IEEE International Symposium on Information Theory Proceedings*, pages 2168–2172. IEEE, jul 2011.
- [RKI16] Paulo V. Rossi, Yoshiyuki Kabashima, and Jun Ichi Inoue. Bayesian online compressed sensing. *Physical Review E*, 94(2):022137, aug 2016.
- [RMW14] D J Rezende, S Mohamed, and D Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *Proceedings of The 31st International Conference in Machine Learning, Beijing China*, 32:1278–, 2014.
- [Rob51] S Robbins, H Monro. A stochastic approximation method. *Statistics*, pages 102–109, 1951.

F References

- [RP16] Galen Reeves and Henry D Pfister. The replica-symmetric prediction for compressed sensing with Gaussian matrices is exact. In *2016 IEEE International Symposium on Information Theory (ISIT)*, pages 665–669. IEEE, jul 2016.
- [RSF14] Sundeep Rangan, Philip Schniter, and Alyson Fletcher. On the convergence of approximate message passing with arbitrary matrices. *IEEE International Symposium on Information Theory - Proceedings*, pages 236–240, 2014.
- [RSF17] Sundeep Rangan, Philip Schniter, and Alyson K. Fletcher. Vector approximate message passing. In *2017 IEEE International Symposium on Information Theory (ISIT)*, volume 1, pages 1588–1592. IEEE, jun 2017.
- [RVE18] Grant M. Rotskoff and Eric Vanden-Eijnden. Parameters as interacting particles: long time convergence and asymptotic error scaling of neural networks. In *Advances in neural information processing systems 31*, pages 7146–7155, 2018.
- [Saa99] David Saad. *On-Line Learning in Neural Networks*. Cambridge University Press, jan 1999.
- [SB97] P. Sollich and D. Barber. On-line learning from finite training sets: An Analytical Case Study. *Europhysics Letters (EPL)*, 38(6):477–482, may 1997.
- [SBD⁺18] Andrew M. Saxe, Yamini Bansal, Joel Dapello, Madhu S. Advani, Artemy Kolchinsky, Brendan D Tracey, and David D Cox. On The Information Bottleneck Theory of Deep Learning. *International Conference on Learning Representations 2018*, pages 1–27, 2018.
- [SBG⁺17] Samuel S. Schoenholz, Google Brain, Justin Gilmer, Google Brain, Surya Ganguli, Jascha Sohl-Dickstein, and Google Brain. Deep Information Propagation. In *International Conference on Learning Representations 2017*, pages 1–18, 2017.
- [SBL⁺18] Mehdi S. M. Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, Sylvain Gelly, Olivier Bachem, and Mario Lucic. Assessing Generative Models via Precision and Recall. *Neural Information Processing Systems 2018, (NeurIPS)*:1–10, 2018.
- [SES19] Itay Safran, Ronen Eldan, and Ohad Shamir. Depth Separations in Neural Networks: What is Actually Being Separated? In *32nd Conference on Learning Theory*, volume 99, pages 1–3. PMLR, 2019.
- [SH09] Ruslan Salakhutdinov and Geoffrey Hinton. Deep Boltzmann Machines. *Artificial Intelligence and Statistics*, 5(2):448–455, 2009.
- [SHK⁺14] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014.
- [SK75] David Sherrington and Scott Kirkpatrick. Solvable Model of a Spin-Glass. *Physical Review Letters*, 35(26):1792–1796, dec 1975.
- [SK08] Takashi Shinzato and Yoshiyuki Kabashima. Perceptron capacity revisited: classification ability for correlated patterns. *Journal of Physics A: Mathematical and Theoretical*, 41(32):324013, aug 2008.
- [SK09] Takashi Shinzato and Yoshiyuki Kabashima. Learning from correlated patterns by simple perceptrons. *Journal of Physics A: Mathematical and Theoretical*, 42(1):015005, jan 2009.
- [SLL19] Luca Saglietti, Yue M. Lu, and Carlo Lucibello. Generalized Approximate Survey Propagation for High-Dimensional Estimation. In *Proceedings of the 36th International Conference on Machine Learning*, pages 4173–4182, 2019.
- [SMH07] Ruslan Salakhutdinov, Andriy Mnih, and Geoffrey Hinton. Restricted Boltzmann machines for collaborative filtering. In *Proceedings of the 24th international conference on Machine learning*, pages 791–798. ACM, 2007.
- [Smo86] Paul Smolensky. Information Processing in Dynamical Systems: Foundations of Harmony Theory. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume I: Foundations*, chapter 6. MIT Press, Cambridge, Massachusetts, 186.
- [SRF16] Philip Schniter, Sundeep Rangan, and Alyson K. Fletcher. Vector approximate message passing for the generalized linear model. In *2016 50th Asilomar Conference on Signals, Systems and Computers*, pages 1525–1529. IEEE, nov 2016.

- [SS95a] David Saad and Sara A. Solla. Exact solution for on-line learning in multilayer neural networks. *Physical Review Letters*, 74(21):4337–4340, 1995.
- [SS95b] David Saad and Sara A. Solla. On-line learning in soft committee machines. *Physical Review E*, 52(4):4225–4243, 1995.
- [SS18] Justin Sirignano and Konstantinos Spiliopoulos. Mean Field Analysis of Neural Networks. *arXiv preprint*, 1805.01053, 2018.
- [SSBD14] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, Cambridge, 2014.
- [SSG19] Umut Simsekli, Levent Sagun, and Mert Gurbuzbalaban. A Tail-Index Analysis of Stochastic Gradient Noise in Deep Neural Networks. 2019.
- [SST92] H. S. Seung, Haim Sompolinsky, and Naftali Tishby. Statistical mechanics of learning from examples. *Physical Review A*, 45(8):6056–6091, 1992.
- [SZT17] Ravid Schwartz-Ziv and Naftali Tishby. Opening the Black Box of Deep Neural Networks via Information. *arXiv preprint*, 1703.00810, 2017.
- [Tal06] Michel Talagrand. The Parisi formula. *Annals of Mathematics*, 163(1):221–263, jan 2006.
- [TAP77] D. J. Thouless, P. W. Anderson, and R. G. Palmer. Solution of ‘Solvable model of a spin glass’. *Philosophical Magazine*, 35(3):593–601, 1977.
- [TDK16] Eric W. Tramel, Angélique Drémeau, and Florent Krzakala. Approximate message passing with restricted Boltzmann machine priors. *Journal of Statistical Mechanics: Theory and Experiment*, 2016(7):073401, jul 2016.
- [Tel16] Matus Telgarsky. Benefits of depth in neural networks. In *29th Annual Conference on Learning Theory*, volume 49, pages 1517–1539, 2016.
- [TGM⁺18] Eric W. Tramel, Marylou Gabrié, Andre Manoel, Francesco Caltagirone, and Florent Krzakala. Deterministic and Generalized Framework for Unsupervised Learning with Restricted Boltzmann Machines. *Physical Review X*, 8(4):041006, oct 2018.
- [Tie08] Tijmen Tieleman. Training restricted Boltzmann machines using approximations to the likelihood gradient. *ICML; Vol. 307*, page 7, 2008.
- [TM17] Jérôme Tubiana and Rémi Monasson. Emergence of Compositional Representations in Restricted Boltzmann Machines. *Physical Review Letters*, 118(13), 2017.
- [TMC⁺16] Eric W. Tramel, Andre Manoel, Francesco Caltagirone, Marylou Gabrié, and Florent Krzakala. Inferring sparsity: Compressed sensing using generalized restricted Boltzmann machines. In *2016 IEEE Information Theory Workshop (ITW)*, pages 265–269. IEEE, sep 2016.
- [TZ15] Naftali Tishby and Noga Zaslavsky. Deep Learning and the Information Bottleneck Principle. *2015 IEEE Information Theory Workshop, ITW 2015*, 2015.
- [Vap00] Vladimir Vapnik. *The Nature of Statistical Learning Theory*. Springer New York, New York, NY, 2000.
- [VBGS17] Rene Vidal, Joan Bruna, Raja Giryes, and Stefano Soatto. Mathematics of Deep Learning. *arXiv preprint*, 1712.04741, dec 2017.
- [Wei07] Pierre Weiss. L’hypothèse du champ moléculaire et la propriété ferromagnétique. *Journal de Physique Théorique et Appliquée*, 6(1):661–690, 1907.
- [WH02] Max Welling and Ge Hinton. A new learning algorithm for mean field Boltzmann machines. *Artificial Neural Networks—ICANN 2002*, pages 351–357, 2002.
- [WHL18] Chuang Wang, Hong Hu, Yue M. Lu, and John A Paulson. A Solvable High-Dimensional Model of GAN. *arXiv preprint*, 1805.08349, 2018.
- [WJ08] Martin J. Wainwright and Michael I. Jordan. Graphical Models, Exponential Families, and Variational Inference. *Foundations and Trends® in Machine Learning*, 1:1–305, 2008.
- [Won95] K. Y. Michael Wong. Microscopic equations and stability conditions in optimal neural networks. *Europhysics Letters (EPL)*, 30(4):245–250, 1995.

F References

- [Won97] K. Y. Michael Wong. Microscopic equations in rough energy landscape for neural networks. *Advances in Neural Information Processing Systems*, pages 302–308, 1997.
- [WRB93] Timothy L.H. Watkin, Albrecht Rau, and Michael Biehl. The statistical mechanics of learning a rule. *Reviews of Modern Physics*, 65(2):499–556, 1993.
- [YFW02] Jonathan S Yedidia, William T Freeman, and Yair Weiss. Understanding Belief Propagation and its Generalizations. *Intelligence*, 8:236–239, 2002.
- [YPR⁺19] Greg Yang, Jeffrey Pennington, Vinay Rao, Jascha Sohl-Dickstein, and Samuel S. Schoenholz. A Mean Field Theory of Batch Normalization. *arXiv preprint*, 1902.08129, 2019.
- [Zam10] Francesco Zamponi. Mean field theory of spin glasses. *arXiv preprint*, 1008.4844, 2010.
- [ZBH⁺17] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding Deep Learning Requires ReThinking Generalization. *International Conference on Learning Representations 2017*, pages 1–15, 2017.
- [ZK16] Lenka Zdeborová and Florent Krzakala. Statistical physics of inference: Thresholds and algorithms. *Advances in Physics*, 65(5):453–552, sep 2016.