

RESEARCH ARTICLE

Bayesian reconstruction of memories stored in neural networks from their connectivity

Sebastian Goldt^{1*}, Florent Krzakala², Lenka Zdeborová³, Nicolas Brunel^{4,5}

1 International School of Advanced Studies (SISSA), Trieste, Italy, **2** IdePHICS laboratory, Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland, **3** SPOC laboratory, Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland, **4** Department of Neurobiology, Duke University, Durham, North Carolina, United States of America, **5** Department of Physics, Duke University, Durham, North Carolina, United States of America

* sgoldt@sissa.it

OPEN ACCESS

Citation: Goldt S, Krzakala F, Zdeborová L, Brunel N (2023) Bayesian reconstruction of memories stored in neural networks from their connectivity. *PLoS Comput Biol* 19(1): e1010813. <https://doi.org/10.1371/journal.pcbi.1010813>

Editor: Lyle J. Graham, Université Paris Descartes, Centre National de la Recherche Scientifique, FRANCE

Received: May 6, 2021

Accepted: December 12, 2022

Published: January 30, 2023

Copyright: © 2023 Goldt et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data are within the manuscript.

Funding: SG and LZ acknowledge funding from the ERC under the European Union's Horizon 2020 Research and Innovation Programme Grant Agreement 714608-SMiLe. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Abstract

The advent of comprehensive synaptic wiring diagrams of large neural circuits has created the field of connectomics and given rise to a number of open research questions. One such question is whether it is possible to reconstruct the information stored in a recurrent network of neurons, given its synaptic connectivity matrix. Here, we address this question by determining when solving such an inference problem is theoretically possible in specific attractor network models and by providing a practical algorithm to do so. The algorithm builds on ideas from statistical physics to perform approximate Bayesian inference and is amenable to exact analysis. We study its performance on three different models, compare the algorithm to standard algorithms such as PCA, and explore the limitations of reconstructing stored patterns from synaptic connectivity.

Author summary

One of the central hypothesis of neuroscience is that memories are stored in synaptic connectivity. Theoretical models show how large numbers of memories can be stored in recurrent neural circuits thanks to synaptic plasticity mechanisms. Recent advances in serial block-face electron microscopy, and machine learning methods, are making it possible to fully reconstruct the synaptic connectivity of neuronal circuits of increasingly large volumes. Here, we ask the question to what extent it is possible to reconstruct memories stored in a neural circuit from the knowledge of its synaptic connectivity. We present an approximate Bayesian inference algorithm, and study its properties on specific attractor network models.

Introduction

Comprehensive synaptic wiring diagrams or “connectomes” provide a detailed map of all the neurons and their interconnections in a brain region or even an entire organism. Since the connectome of the nematode *C. elegans* was obtained using electron microscopy methods in

1986 [1], methods for data acquisition and analysis have both been scaled up and improved significantly. Today, it has become possible to provide connectomes of much more complex systems such as various *Drosophila melanogaster* circuits [2, 3], or even a large part of its brain [4, 5]; the olfactory bulb of zebrafish [6]; and various pieces of the rodent retina [7–9], hippocampus [10], and cortex [11–16]. While there still remain a number of formidable challenges on the way to the complete connectome of a mammal or even human brain [17], the data sets available today already give rise to a number of intriguing questions. At the same time, it is becoming increasingly clear that new quantitative methods must be developed to fully exploit the new troves of data that connectomics provides [18].

Here, we focus on local neural networks that store information in their synaptic connectivity. It has been hypothesised that cortical networks are optimised for this task, thanks to their extensive recurrent synaptic connectivity [19]. A popular model for these networks are attractor neural networks such as the Hopfield model [20] and various generalisations [21–24], in which memories are stored as attractor states of the dynamics. These attractor states represent learned internal representations of external stimuli that have been presented repeatedly to the network during training, inducing changes in synaptic weights of the network. One natural question to ask is then: given the knowledge of the synaptic connections between neurons in a recurrent neural network, can we reconstruct the patterns of activity that were stored in this network in the first place?

In this paper, we first give a mathematical formulation of this problem in terms of a Bayesian inference problem. The Bayesian approach has the advantage of providing a natural way to handle the uncertainty associated with estimating a large number of parameters from a large number of noisy observations, *i.e.* the reconstruction of the original stimuli from the strength of the synapses in the networks in the present case. Modelling the noise is crucial in this problem, as we cannot expect the synaptic strengths reported in connectomes to be more than rough estimates. We use tools from statistical physics to both determine when solving this inference problem is theoretically possible in a model setting, and to provide a practical algorithm to do so. We analyse the performance of the algorithm in detail on a variety of different problems, and we invite the reader to download our reference implementation of the algorithm on [GitHub](#) and to use and extend it.

The task: Reconstructing memories from network connectivity

The network model. We analyse a variant of the celebrated Hopfield model [20] for a recurrent neural network composed of N interacting neurons with state s_i , $i = 1, \dots, N$. The network is fully-connected with symmetric, bidirectional connections that have a scalar weight $J_{ij} = J_{ji} \in \mathbb{R}$. The neurons update their state at iteration $k + 1$ sequentially according to

$$s_i^{k+1} = g(a_i^k), \quad (1)$$

where $g(\cdot)$ is some non-linear activation function and $a_i^k = \sum_{j \neq i} J_{ij} s_j^k$ is the total synaptic input of the i th neuron.

The network stores P fixed patterns or memories, which are N -dimensional vectors that we collect in the matrix $\mathbf{X}^* \in \mathcal{X}^{P \times N}$. We write $\mathbf{X}_{\mu, \cdot}^*$ for the μ th pattern stored in the network, and \mathcal{X} denotes the set of values that pattern entries can take, *e.g.* $\mathcal{X} = \{\pm 1\}$ for binary patterns. Note that these patterns correspond to deviations of neuronal activity from its mean, and not neuronal activity itself, which is constrained to be non-negative. This subtraction of mean activity is expected to be performed by the plasticity rule operating at the synapse (using a ‘covariance rule’, see *e.g.* [25]). In the case of binary patterns, $X_{\mu, i} = 1$ means a neuron is active in a given pattern, while $X_{\mu, i} = -1$ means a neuron is inactive. We also assume that $X_{\mu, i}$ are i.i.d. random

variables, i.e. that stored patterns are uncorrelated. Patterns are stored in the network by choosing its weights J_{ij} such that the patterns \mathbf{X}_{μ}^* become fixed points of its dynamics.

We study this model in the thermodynamic limit $N \rightarrow \infty$, while keeping the number of patterns P of order 1. This scaling makes the resulting weight matrix of the Hopfield model a low-rank matrix. Low rank matrices have played an important role in neuroscience in recent years, in particular for the modelling of recurrent networks [26–29]. Hence, methods to estimate them from data in a principled way can help connect these theories with experimental data.

The learning rule. A classic idea for choosing the weights, or the connectivity structure of the network $\mathbf{J} = [J_{ij}] \in \mathbb{R}^{N \times N}$, is to choose the weights proportional to the empirical correlation of the patterns, $W_{ij} \sim \sum_{\mu} X_{\mu,i}^* X_{\mu,j}^*$. This prescription is also known as the Hebb rule [30] and can be written more compactly as

$$\mathbf{J} = \mathbf{W} = \frac{1}{\sqrt{N}} \mathbf{X}^* (\mathbf{X}^*)^{\top}, \tag{2}$$

where $(\mathbf{X}^*)^{\top}$ is the transpose of \mathbf{X}^* and \mathbf{W} is thus the empirical correlation matrix of the patterns, assuming that the means of the patterns are zero, as we will do throughout this work, and the choice of the $1/\sqrt{N}$ scaling is explained below. With binary neurons $s_i(t) = \pm 1$, this model corresponds to the celebrated Hopfield model [20]. In this model, the network exhibits fixed point attractors close to the stored memories, provided the number of stored patterns P is smaller than $\alpha_c N$ where $\alpha_c \sim 0.14$ [31].

The connectivity matrix Eq (2) has a number of unrealistic features that makes it inadequate for the problem we are interested in here: (i) The network is fully connected, at odds with neuronal networks in the brain; (ii) Synapses are not sign-constrained, while synapses in the brain are either excitatory (i.e. non-negative) or inhibitory (i.e. non-positive). A minimal model that satisfies both requirements is the **rectified Hopfield model**

$$J_{ij} = \Phi(W_{ij} - \tau + \zeta_{ij}) \geq 0 \tag{3}$$

where $\tau > 0$, ζ_{ij} is a noise term (see below), and we choose $\Phi(x) = \max(0, x)$. This choice ensures that weights are non-negative, effectively yielding a model of a network of excitatory neurons. This is consistent with the hypothesis that information storage occurs primarily in excitatory-to-excitatory synapses, while the job of inhibitory neurons is primarily to control the level of activity in the excitatory network. This view is consistent with a number of studies [32, 33] but has been challenged by others [34].

The noise term ζ_{ij} is taken to be a symmetric random Gaussian matrix, i.e. for $i < j$, ζ_{ij} s are i.i.d. random Gaussian variables with mean zero and standard deviation v , and $\zeta_{ji} = \zeta_{ij}$. The scalar parameter τ controls the connection probability in the network,

$$p_c = p(J_{ij} > 0) = \frac{1}{2} \operatorname{erfc}\left(\frac{\tau}{\sqrt{2}v}\right) \tag{4}$$

to leading order as $N \rightarrow \infty$. In particular, the network becomes sparse in the large τ limit. Since $P \sim O(1)$, the weights obtained from Eq (2) will have variance $1/N$, while the noise has variance v of order 1. The model we study is related to a family of connectivity matrices studied by Sompolinsky [35], and bears similarities with a model recently proposed by Mongillo *et al.* [34].

In brain networks, connectivity matrices are not symmetric. However, if we assume that the weights depend on the stimuli only via the symmetric matrix $\mathbf{X}^* (\mathbf{X}^*)^{\top}$, these asymmetries will be due to the different sources of noise in the learning process, and the fact that connectomic reconstructions will give us at best an approximation of true synaptic weights (see

[Discussion](#)). For the purpose of analysing quantities derived from the symmetric matrix \mathbf{X}^* $(\mathbf{X}^*)^\top$, as we will do here, we can hence symmetrise the matrix \mathbf{J} , or equivalently focus on the case where the noise matrix ζ_{ij} is symmetric.

A note on the $1/\sqrt{N}$ scaling. Our choice of scaling in [Eq \(2\)](#) is made with reconstructing the patterns in mind and follows from random matrix theory. Our model for the connectivity matrix is related to the spiked Wigner model [[36, 37](#)], where a random matrix \mathbf{M} is constructed as $M_{ij} = \beta u_i u_j + \zeta_{ij}$, with ζ_{ij} as above. The matrix \mathbf{M} is hence a rank-one perturbation of a random matrix with elements drawn i.i.d. from the normal distribution. The task is to reconstruct the vector \mathbf{u} , whose elements are of order 1, from the matrix \mathbf{M} . Intuitively, we want to put ourselves in the “interesting” regime, where reconstruction of the N entries of \mathbf{u} from the $\mathcal{O}(N^2)$ elements of \mathbf{M} is neither trivially easy nor impossible. A classic analysis of this model using random matrix theory [[36, 38](#)] reveals that this “interesting” regime is characterised by a phase transition in the overlap between the leading eigenvector of \mathbf{M} and the vector \mathbf{u} , which occurs precisely for a signal-to-noise ratio β that scales as $1/\sqrt{N}$. We thus choose the same scaling for our non-linear matrix model ([3](#)).

Reconstructing vs retrieving the memories. We emphasise that our focus in this paper is the *reconstruction* of stimuli from an observed connectivity matrix \mathbf{J} , which is different from the *retrieval* problem [[20, 39–41](#)], where we ask whether the patterns \mathbf{X} are stable fixed points of the dynamics of the network, [Eq \(1\)](#). We will discuss the feasibility of reconstructing patterns from the network’s dynamics at the end of the paper.

Solving the inference problem using statistical physics

Our aim is to reconstruct the patterns \mathbf{X}^* that were used to create the connectivity matrix \mathbf{J} of a Hopfield network using the connectivity structure ([3](#)). We will call the patterns \mathbf{X}^* the ground truth of the problem. Since the connectivity structure is stochastic, we formulate the pattern reconstruction as a probabilistic inference problem. We interpret the connectivity matrix of the network \mathbf{J} as a noisy observation of the symmetric low-rank matrix $\mathbf{W} \sim \mathbf{X}^* (\mathbf{X}^*)^\top$, which was distorted by the transformation given by the rule ([3](#)). We can characterise the conditional probability distribution of a weight J_{ij} given W_{ij} as

$$P_{\text{out}}(J_{ij} = 0 | W_{ij}) = \frac{1}{2} \operatorname{erfc}\left(\frac{W_{ij} - \tau}{\sqrt{2\nu}}\right), \tag{5a}$$

$$P_{\text{out}}(J_{ij} | W_{ij}) = \frac{\exp(-(J_{ij} - W_{ij} + \tau)^2 / 2\nu^2)}{\sqrt{2\pi\nu}} \quad \text{for } J_{ij} > 0. \tag{5b}$$

Reconstructing low-rank matrices from such noisy, distorted observations is a generic inference problem that appears in a lot of different applications, such as robust [[42](#)] and sparse [[43–45](#)] PCA, Gaussian mixture clustering [[46](#)], and community detection in dense networks [[47, 48](#)], to name but a few. Low-rank matrices have also been used extensively in neuroscience to model recurrent connectivity [[26–29](#)]. The advantage of a Bayesian approach in all these problems is that they allow for a principled and transparent integration of knowledge about the problem into the inference process, for example through the choice of prior distribution and output channel. In the following, we will assume that we know the hyper-parameters ν and τ that were used to generate the connectivity matrix \mathbf{J} . Expectation maximisation and related techniques seem natural candidates to extend our approach to cases where we would need to learn these hyper-parameters (see [Discussion](#)).

Here, we adopt a Bayesian approach [[49](#)] to the inference of the patterns given the connectivity \mathbf{J} . This means that we will consider our reconstruction of the patterns as a random

variable \mathbf{X} , whose posterior distribution $p(\mathbf{X}|\mathbf{J})$ given the connectivity matrix is given by Bayes' theorem:

$$P(\mathbf{X}|\mathbf{J}) = \frac{1}{Z(\mathbf{J})} \prod_i^N P_X(\mathbf{x}_i) \prod_{j>i}^N P_{\text{out}}(J_{ij}|W_{ij}). \quad (6)$$

where we introduced the shorthand $\mathbf{x}_i = \mathbf{X}_{:,i} \in \mathbb{R}^p$ for the i th column of \mathbf{X} . We assume that patterns are uncorrelated from each other and that we know the *a priori* distribution $P_X(\mathbf{X})$ over patterns that were stored in the network. This distribution could for example reflect the fact that we know that the memories stored in the network are binary, encoding whether a given neuron is firing or not, or that we have an idea of the probability that any given neuron is firing in a given memory.

Note that we are phrasing the problem of reconstructing the memories here as the problem of reconstructing the columns of the matrix \mathbf{X} , or equivalently, the tuning curves of each neuron. Eventually, we are interested the whole matrix \mathbf{X} , so it doesn't matter whether we reconstruct all its columns or all its rows. However, it is more convenient from an algorithmic point of view to work with the columns, which is the approach we will adopt here. The marginals of the rows of this distribution, which are N -dimensional, provide the best estimate of the patterns that can be performed [49].

Evaluating the high-dimensional integral to obtain the marginals exactly is an intractable problem. Instead, here we exploit the formal analogy between the posterior distribution (6) and certain probability distributions that arise in statistical physics to derive an algorithm called "approximate message passing" [50], which performs approximate Bayesian inference of the memories stored in the network. Furthermore, we will demonstrate that techniques from statistical physics, in particular a tool called "state evolution" (SE), can be used to analyse the behaviour and the performance of this algorithm in quite some detail. Approximate message passing algorithms can be understood as a variant of belief propagation, a general algorithm for inference in graphical models that is usually credited to Pearl [51].

Results

A Bayesian algorithm for pattern reconstruction

The inference problem considered here, where we aim to recover a symmetric low-rank matrix from noisy observations, can be solved using a class of approximate message passing (AMP) algorithms for low-rank matrix factorisation called Low-RAMP. It was derived by Lesieur et al. [52], building on previous works [44, 53, 54] that provided AMP algorithms for particular instances of low-rank matrix factorisation problems. Low-RAMP is an iterative algorithm that produces estimates for the mean $\hat{\mathbf{x}}_i$ of the marginal distribution of $p(\mathbf{x}_i)$ and their covariance matrix σ_i , where \mathbf{x}_i is in general the i th column of the low-rank matrix \mathbf{X} that we are estimating by evaluating the posterior distribution (6). In the present case, $\hat{\mathbf{x}}_i$ is the mean of the estimated 'tuning curve' of the i th neuron (see above). Using this framework, we will derive variants of the algorithm for the pattern reconstruction problem outlined in the previous section. We present the algorithm in detail in the Methods section.

We also provide a [reference implementation](#) of Low-RAMP for symmetric and bipartite matrix factorisation problems applicable to a number of different problems. It is designed to be easily extendable to other problems and also provides a number of further utility functions. All the results in this paper can be reproduced using this code.

State evolution

The AMP algorithm has the distinctive advantage over other algorithms, such as Monte Carlo methods, that its behaviour in the limit $N \rightarrow \infty$ for separable prior on the \mathbf{X}^* , random i.i.d. noise ζ_{ij} , and number of patterns $P = O(1)$, can be tracked exactly and at all times using the “state evolution” technique [50, 55]. The roots of this method go back to ideas originally introduced in physics to deal with a class of disordered systems called glasses [56, 57]. For the low-rank matrix factorisation problems we consider here, state evolution was derived and analysed in detail by Lesieur et al. [52], building on previous works that derived and analysed state evolution for other specific problems [44, 45, 53]. The last few years in particular have seen a surge of interest in using state evolution to understand the properties of approximate Bayesian inference algorithms for a number of problems [58].

Since we are adopting a probabilistic approach to estimating the patterns, we will call the reconstruction of the patterns the mean of the posterior distribution, which we denote by a hat: $\hat{\mathbf{x}}_i$. Our goal is to track the mean-squared error mse_X of the reconstruction $\hat{\mathbf{x}}_i^t$ of the true signal \mathbf{x}_i^* after t steps of the algorithm,

$$\text{mse}(t) \equiv \frac{1}{N} \sum_i^N \|\hat{\mathbf{x}}_i^t - \mathbf{x}_i^*\|_2^2, \tag{7}$$

where $\|\cdot\|_2^2$ denotes the Euclidean norm of a vector. The mse can be expressed in terms of a single matrix-valued parameter defined as

$$\mathbf{M}^t \equiv \frac{1}{N} \sum_i^N \hat{\mathbf{x}}_i^t \mathbf{x}_i^{*\top} \in \mathbb{R}^{P \times P}, \tag{8}$$

such that $\text{mse}(t) = \text{Tr}[\langle \mathbf{x}_0 \mathbf{x}_0^\top \rangle - \mathbf{M}^t]$. Here and throughout this paper we write averages with respect to the prior distribution $p_X(\mathbf{x})$ of the corresponding model as $\langle \cdot \rangle$. We write \mathbf{x}_0 with the subscript to underline that the random variable \mathbf{x}_0 is not a column of the matrix \mathbf{X} that we’re trying to evaluate; instead, it is a variable that is drawn from the prior and averaged over.

Now the goal is to find an update equation for the order parameter \mathbf{M}^t that mirrors the progress of the algorithm. This update equation is the state evolution equation [50, 55]. Remarkably, from [59] we see that the two constants defining our problem, τ and v , do not appear explicitly in the state evolution equations. Instead, the behaviour of the algorithm—and hence its performance—only depends on an effective signal-to-noise ratio (SNR) of the problem, which is a function of the threshold τ and noise variance v utilised in the connectivity structure (3). Formally, it can be expressed as the inverse of the Fisher score matrix [60] of the generative model we use to describe how the network is connected (5a) and (5b), evaluated at $W_{ij} = \mathbf{x}_i \mathbf{x}_j^\top / \sqrt{N} = 0$:

$$\frac{1}{\Delta} \equiv \mathbb{E}_{P_{\text{out}}(J|w=0)} \left(\frac{\partial \ln P_{\text{out}}(J|w)}{\partial w} \right)_{J,w=0}^2 \tag{9}$$

$$= \frac{\tau e^{-\tau^2/2v^2}}{\sqrt{2\pi}v^3} + \frac{e^{-\tau^2/v^2}}{\pi v^2 \text{erfc}(-\tau/\sqrt{2}v)} + \frac{1}{2v^2} \text{erfc}\left(\frac{\tau}{\sqrt{2}v}\right). \tag{10}$$

Here and throughout, \mathbb{E} denotes the expectation over the random variables. In fact, on the level of the algorithm, everything about the output channel (5) can be summarised in this single, scalar quantity Δ . This remarkable universality of the state evolution and hence the AMP

algorithm with respect to the output channel was first observed in [59] and dubbed “channel universality”.

State evolution provides an update equation for the order parameter M^t that mirrors the progress of the algorithm. We first define an auxiliary function

$$f(\mathbf{A}, \mathbf{b}) = \frac{1}{Z(\mathbf{A}, \mathbf{b})} \sum_{\mathbf{x} \in \mathcal{X}^P} \mathbf{x} p_{\mathcal{X}}(\mathbf{x}) \exp\left(\mathbf{b}\mathbf{x} - \frac{1}{2} \mathbf{x}^{\top} \mathbf{A} \mathbf{x}\right). \tag{11}$$

where $\mathbf{A} \in \mathbb{R}^{P \times P}$ and $\mathbf{b} \in \mathbb{R}^P$. If $\mathbf{A} = 0$ and $\mathbf{b} = 0$, this function would compute the average over the prior distribution $p_{\mathcal{X}}(\mathbf{x})$. Instead, \mathbf{b} and \mathbf{A} are estimated from the data (see the algorithm for details) so f computes an average over a distribution that contains the prior and a data-dependent part. This structure reflects the Gaussian approximation of the posterior density that we apply here, or more broadly speaking the interplay between prior information and data-dependent likelihood that is typical of Bayesian inference. Consequently, $Z(\mathbf{A}, \mathbf{b}) = \sum_{\mathbf{x} \in \mathcal{X}^P} p_{\mathcal{X}}(\mathbf{x}) \exp(\mathbf{b}\mathbf{x} - \frac{1}{2} \mathbf{x}^{\top} \mathbf{A} \mathbf{x})$ is a normalisation factor. The update equation for the order parameter M^t can be written using this auxiliary function for all the cases considered in this paper; it reads [52]

$$\mathbf{M}^{t+1} = \mathbb{E}_{\mathbf{x}_0, \mathbf{z}} \left[f\left(\frac{\mathbf{M}^t}{\Delta}, \frac{\mathbf{M}^t}{\Delta} \mathbf{x}_0 + \sqrt{\frac{\mathbf{M}^t}{\Delta}} \mathbf{z}\right) \mathbf{x}_0^{\top} \right] \tag{12}$$

where \mathbf{z} is a P -dimensional vector of Gaussian random variables with mean zero and variance 1. The average over \mathbf{x}_0 is taken with respect to the prior distribution $p_{\mathcal{X}}(\mathbf{x})$, as discussed above.

So to summarise, statistical physics gives us an algorithm to perform approximate inference of the patterns and the state evolution Eq (12) allows us to track the behaviour of the algorithm over time. We can thus analyse the performance of the algorithm in high-dimensional inference by studying the fixed points of the low-dimensional state-evolution (12). This is the key idea behind this approach, and we will now demonstrate the usefulness of this machinery by applying it to several specific cases.

Reconstructing binary patterns

As a first application of the algorithm and the analysis tools outlined so far, we consider the reconstruction of a set of binary patterns, $\mathcal{X} = \{\pm 1\}$. We will assume that both positive and negative values are equiprobable and that the components of a pattern vector are independent of each other, so the prior on a column of the matrix of stored patterns, \mathbf{x}_i , is simply

$$p_{\mathcal{X}}(\mathbf{x}_i) = \prod_j^P p_{\mathcal{X}}(X_{ij}) = \frac{1}{2^P}. \tag{13}$$

A single pattern ($P = 1$). It is instructive and helpful for the following discussions to first consider the case where $P = 1$, *i.e.* there is only a single pattern stored in the network that we are trying to recover from \mathbf{J} . The threshold function for the model then becomes $f(A, B) = \tanh(B)$, with $B \in \mathbb{R}$, and the state evolution for the now scalar parameter m^t simplifies to

$$m^{t+1} = \mathbb{E}_z \tanh\left(\frac{m^t}{\Delta} + \sqrt{\frac{m^t}{\Delta}} z\right) \tag{14}$$

where w is a scalar Gaussian random variable with zero mean and unit variance.

We can now iterate the state evolution Eq (14) with a given noise level $\Delta(\nu, \tau)$ until convergence and then compute the mse corresponding to that fixed point. The fixed point we converge to reveals information about the performance of the AMP algorithm. We plot the results on the left-hand side of Fig 1 for the two different *initialisations* of the algorithm: in blue, we plot the mse obtained by iterating SE starting with an *random initialisation*

$$m^{t=0} = 0 + \delta, \quad (15)$$

where $\delta > 0$ is a very small random number. The error obtained in this way is the one that is obtained by the AMP algorithm when initialised with a random draw from the prior distribution—in other words, a random guess for the patterns. This is confirmed by the blue crosses, which show the mean and standard deviation of the mse obtained from five independent runs of the algorithm on actual instances of the problem. The dashed orange line in Fig 1 shows the final mse obtained from an informed initialisation

$$m^{t=0} = 1 - \delta, \quad (16)$$

which would correspond to initialising the algorithm with the solution, i.e. $\hat{x}_i = x_i^*$.

In this model, we find that the AMP algorithm starting from a random guess performs just as well as the algorithm starting from the informed initialisation. This need not always be the case, and we will indeed find a different behaviour in the next sparse and skewed models we consider.

When is recovery possible?. We can see from the middle plot of Fig 1 that recovery of the memories from the connectivity J is not always possible; there exists a critical value for the effective noise Δ_c above which the mean-squared error of the solution obtained by the algorithm is the same as we would have obtained by making a random guess for the solution based on the prior distribution (13) alone, without looking at the data. We can calculate this critical noise level Δ_c using the state evolution (12). We can see from that equation that $m^t = 0$ is a trivial fixed point, in the sense that the mse corresponding to that fixed point is equal to the mse obtained by making a random guess. Expanding Eq (14) around this fixed point yields $m^{t+1} = m^t/\Delta$. There are hence two regimes for recovery, separated by a critical value

$$\Delta_c = 1 \quad (17)$$

of the effective noise (9). If $\Delta > \Delta_c$, the uniform fixed point is stable and recovery is impossible. On the other hand, for $\Delta < \Delta_c$, the uniform fixed point is unstable and hence AMP returns an estimate for the patterns that has an mse that is lower than random guessing. The phase diagram in the middle of Fig 1 delineates the easy and the impossible phase for the rectified Hopfield channel with symmetric prior (13). While there could be in principle other fixed points of the state evolution equations for other priors and channels [52], it is always one of the fixed points that is reached from either the informed or the uninformed initialisation that describes the behaviour of the algorithm.

At first sight, the impact of the additive Gaussian noise ζ_{ij} on the phase diagram in Fig 1 appears counter-intuitive. If we fix the threshold to, say, $\tau = 0.5$, reconstruction is impossible for small variances ν of ζ_{ij} . As we increase ν , i.e. as we add *more* noise to the system, recovery becomes possible. The key to understanding this behaviour is that for a single stimulus $P = 1$, a weight in the network will have one of two possible values which are symmetric around the origin, $W_{ij} = \pm 1/\sqrt{N}$. By applying the rectification, for any cut-off $\tau > W_{ij}$ the resulting weight matrix J without additive noise is trivially zero and no recovery is possible. We can only hope to detect something when an added noise ζ_{ij} pushes the value of the weight before rectification above the cut-off. Recovery then becomes possible if the added noise is large enough




Fig 1. (Left) Reconstruction and performance of the message-passing algorithm for binary patterns. We plot the mse (7) obtained by the AMP algorithm (32) as a function of the effective noise Δ (9) (blue crosses). We plot the performance of the algorithm starting from random (15) and informed (16) initialisations. Solid lines depict the prediction obtained from iterating the state evolution Eq (14). Having $\Delta/\Delta_c > 1$ corresponds to the white region in the phase diagram on the right. We also plot the mse of the reconstruction obtained by applying PCA to the weight matrix J and to the Fisher matrix S (31) (green and red, resp.) Parameters: $\tau = 0$. $N = 5000$ for AMP, $N = 20000$ for PCA. (Center) **Phase diagram for the rectified Hopfield channel with $P = 1$.** We plot whether reconstruction of the patterns better than a random guess is easy (blue) or impossible (white) using the message-passing algorithm as a function of the constant threshold τ and the variance v of the Gaussian noise appearing in the connectivity structure (3). The solid lines are the contours of the connection probability $p_C(v, \tau)$ (4). (Right) **Critical noise v^* as a function of connection probability p_C .** We plot v^* , the largest variance of the additive Gaussian noise ζ_{ij} at which reconstruction remains possible, against the probability p_C (4) that any two neurons are connected.

<https://doi.org/10.1371/journal.pcbi.1010813.g001>

that the weight without noise is larger than the cut-off $a + \zeta_{ij} > \tau$, while remaining small enough that it's significantly more likely that the noise-less weight is positive than negative. As the noise variance increases even further, its detrimental effects dominate, and recovery becomes impossible again. This mechanism is reminiscent of stochastic resonance (SR), a mechanism where a weak signal is amplified by the presence of noise. Indeed, our problem contains the three ingredients for SR (e.g. [61]): A threshold mechanism, given by the rectification in the connectivity structure: A weak signal (the stored patterns); and a noise term, ζ .

As already mentioned, when noise is too large recovery becomes impossible. We show on the right of Fig 1 the critical variance of ζ_{ij} above which reconstruction becomes impossible, v^* , as a function of the connection probability p_C , given by Eq (4). This plot can be obtained by solving, for a given value $p_C = c$, the two-dimensional system

$$\Delta = 1 \tag{18}$$

$$p_C = c \tag{19}$$

for (τ, v) . As expected, the critical variance increases with the connection probability, and it goes to zero as the connection probability goes to zero.

Comparison to principal component analysis (PCA). Principal component analysis (PCA) is another method to reconstruct the stored patterns from the network connectivity. PCA and other spectral methods have some advantages: they are non-parametric, and their implementation in the case of a single pattern is straightforward: the PCA prediction for the stored pattern is simply the leading eigenvector of J . We plot the mean-squared error (7) of this estimate with the green line on the left of Fig 1, where we see that the reconstructing error of PCA is larger than the one of AMP, especially for large values of the noise. This is also borne out by theory: the reconstruction mean-squared error of PCA can be shown to be strictly larger than the AMP estimate, since the latter is the Bayes-optimal predictor [52].

An alternative PCA algorithm can be found by linearising the AMP equations around the trivial fixed point $\hat{x} = 0$ [58, 62]. This linearisation yields an equation that can be interpreted

as PCA applied to the Fisher matrix \mathbf{S} (31) instead of \mathbf{J} . Since the Fisher matrix depends on the generative model for the data when deriving the message-passing equations, looking at its leading eigenvector offers a spectral algorithm that is more adapted to the problem at hand. Indeed, we find that its error (red line in Fig 1) is slightly lower than the error obtained from PCA on the weights directly. In either case, the performance of PCA is worse than that of AMP.

The large value of the PCA error compared to the AMP error at large noise levels in Fig 1 reveals a fundamental weakness of PCA: even at noise levels above the critical noise Δ_c , where no reconstruction is possible for any algorithm, PCA can be applied and will return a prediction—there is no concept of uncertainty in PCA. Hence the mse of PCA tends to a constant as the noise increases and the leading eigenvector of \mathbf{J} is just a random vector; for the Hopfield prior and when rescaling the eigenvectors to have the same length as draws from the Hopfield prior, this constant is 2. AMP on the other hand returns a vector full of zeros if $\Delta > \Delta_c$ (and the prior has an average of 0, as is the case for all the priors we consider). AMP thus expresses its uncertainty about the planted pattern, yielding an $mse = 1$ for inputs with $x_i = \pm 1$. The advantage of the Bayesian approach is thus that it prevents over-confident predictions in the high noise regime.

The weaker performance of PCA compared to AMP is due to the fact that spectral methods do not offer a natural way to incorporate the prior knowledge we have about the structure and distribution of stimuli into the recovery algorithm. The Bayesian framework incorporates this domain knowledge in a transparent way through the generative model of the stored patterns $p_X(\mathbf{x})$. We will see that this creates an even larger performance gap for sparse patterns and patterns with low coding level.

Many patterns ($P > 1$). For the general case of several patterns $P > 1$ with finite P , we can significantly simplify the state evolution by noticing that the matrix M^t will interpolate between a matrix full of zeros at time $t = 0$ and a suitably scaled identity matrix in the case of perfect recovery, *i.e.*

$$M^t = m^t \mathbf{I}_P / \Delta, \tag{20}$$

where \mathbf{I}_P is the identity matrix in P dimensions. In other words, for uncorrelated patterns, the different input patterns do not interact during the reconstruction, and so the off-diagonal matrix elements remain zero in the case where we only store a few patterns and the connectivity structure remains low-rank. Once we overload the model by storing many more patterns, we would have non-zero off-diagonal elements, meaning that reconstructions converge to spurious patterns, for example linear combinations of the patterns. However, in this regime the state evolution derived above also breaks down. In this case, the threshold function becomes

$$[f(A, B)]_k = \tanh\left(\frac{m^t}{\Delta} x_{0,k} + \sqrt{\frac{m^t}{\Delta}} z_k\right) \tag{21}$$

where z_k is again a standard Gaussian variable. Substituting into the state evolution gives an update equation for the parameter m^t , namely

$$m^{t+1} = \mathbb{E}_z \tanh\left(\frac{m^t}{\Delta} + \sqrt{\frac{m^t}{\Delta}} z\right) \tag{22}$$

where m^t is the overlap parameter introduced above (20). This update has the same form as the state evolution in the $P = 1$ case, Eq (14). So we find, remarkably, that recovering P distinct patterns is exactly equivalent to recovering a single pattern P times in the thermodynamic

limit where $N \rightarrow \infty$ while the number of patterns is of order $P \sim \mathcal{O}(1)$. This approximation will eventually break down in practical applications with finite network sizes, and we investigate the breaking point of this behaviour below.

Recovering many patterns with PCA poses an additional challenge. While it is easy to recover the leading rank- P subspace of the matrices \mathbf{J} or \mathbf{S} by simply computing the P leading eigenvectors, it is not clear how to recover the exact patterns from these eigenvectors, which can be any rotation of the input patterns due to the rotational symmetry of \mathbf{W} . This can be seen from the fact that the patterns \mathbf{X}^* could be multiplied by any rotational matrix \mathbf{O} with $\mathbf{O}\mathbf{O}^\top = \mathbb{I}$ without changing the resulting weight matrix \mathbf{J} , see Eq (2). The best way to recover the exact stimuli from the principal components is thus not clear *a priori* (see [63]). Other problems require combining PCA with other methods, such as k -means or gradient descent. Since we have already seen that AMP outperforms PCA on binary patterns, and we will see that this gap only increase for the other types of patterns we will study below, we do not investigate further this direction.

A first summary. Before we turn to more complicated prior distributions over the patterns, let us briefly summarise our results so far. We derived an algorithm that can reconstruct patterns from the connectivity of recurrent network whose weights are obtained from the learning rule Eq (3). Whether or not the algorithm is successful in this reconstruction depends on the noise level ν and the threshold τ . These parameters can be combined into an effective noise parameter Δ (9), which determines the performance of the message-passing algorithm. The algorithm performs well in the reconstruction task, and beats non-parametric approaches like PCA by virtue of including prior information about the distribution of the patterns in a principled way.

Sparse patterns

An interesting variation of the rectified Hopfield model is its sparse version, where only a fraction $0 \leq \rho \leq 1$ of the components x_{ij} of a pattern \mathbf{x}_i are non-zero. The prior distribution then becomes

$$p_{\mathbf{x}}(\mathbf{x}_i) = \prod_{j=1}^P p_x(x_{ij}) = \prod_{j=1}^P \left[(1 - \rho)\delta(x_{ij}) + \rho/2 \left[\delta(x_{ij} - 1) + \delta(x_{ij} + 1) \right] \right], \quad (23)$$

where $\delta(\cdot)$ is the Kronecker delta. This prior has mean $\langle x \rangle = 0_p$, where 0_p is a vector of p zeros, and covariance $\langle xx^\top \rangle = \rho \mathbf{I}_p$. We emphasize again that these patterns represent a deviation of neuronal activity from its mean—in this case ± 1 represents an increase/decrease activity, while 0 represents no change in activity in a given pattern. The state evolution will interpolate between an order parameter that is all zeros for an estimator that is drawn from the prior distribution and completely uncorrelated with the ground truth, and $M = \rho \mathbf{I}_p$ for perfect reconstruction. We delegate the mathematical details to the Methods section, and focus here on the performance of the algorithm.

Performance of message-passing. We first note that the critical noise level above which AMP will not be able to recover the stored pattern better than a random guess in the sparse reconstruction problem is $\Delta_c = \rho^2$. In that sense, the critical noise is a property of the AMP algorithm, and its value can again be obtained by linearising the state evolution Eq (36) around its trivial fixed point $m^t = 0$. However, it is generally believed that no other algorithm is able to recover the patterns above this noise level, which would make this threshold a property of the model rather than the algorithm; we'll come back to this point in the next paragraph. In any case, the decrease of Δ_c with ρ means that the reconstruction performance at fixed ν , τ and hence Δ decreases with pattern sparsity. However, so too does the difficulty of the problem. If




Fig 2. Reconstructing sparse, binary patterns using message passing algorithms and PCA. We plot the mse per pattern obtained by the AMP algorithm, Eq (32), as a function of the effective noise Δ (9), for random (15) and informed (16) initialisations. Lines depict the result of the state evolution, while crosses denote the performance of the AMP algorithm on an instance of the problem. While AMP performs the same starting from both initialisations for $\rho = 0.1$ and $\rho = 0.3$, there is a gap in performance for $\rho = 0.05$, which might hint at the existence of a hard phase (see main text). We also plot the mse of the reconstruction obtained by applying PCA to the weight matrix J and to the Fisher matrix S (31) (green and red, resp.) Parameters: $\tau = 0$. $N = 2000$ for AMP, $N = 20000$ for PCA.

<https://doi.org/10.1371/journal.pcbi.1010813.g002>

we normalise the noise level by the critical noise, *i.e.* if we plot the reconstruction error as a function of the Δ/Δ_c as we do in Fig 2, we see that a large fraction ρ of non-zero components X_{ij} leads to better reconstruction at small noise levels.

For $\rho = 0.1$ and $\rho = 0.3$, the performance plots in Fig 2 resemble the results obtained for the symmetric Hopfield model overall: Reconstruction is possible if the effective noise level is below the critical noise level. For these two values of the sparsity, the mse of the estimate returned by AMP (blue dashed line) matches the mse obtained from state evolution when starting from the informed initialisation (orange). However, the reconstruction errors predicted by state evolution for informed and random initial conditions *disagree* at smaller sparsity. For $\rho = 0.05$ (left), there is a range of effective noise levels Δ for which the performance of AMP does not match the performance predicted by state evolution starting from informed initial conditions. We note that this could be the signature of a so-called *hard phase*, where a better-than-chance reconstruction of the pattern is information-theoretically possible – *i.e.* there is some trace of the stored pattern in the connectivity—but AMP is not able to extract it. However, it is important to emphasise that AMP performs sub-optimally with respect to the amount of information that is in the connectivity, but not with respect to the performance of any other known algorithm. In other words, while AMP does not exploit all the information that is in the connectivity, it is broadly believed in theoretical computer science that *no algorithm* can reconstruct the patterns with non-trivial error at this level of noise in polynomial time, if this is indeed a hard phase of the reconstruction problem. PCA for example is also not able to reconstruct any trace of the pattern at this noise level (see below). We refer the interested reader to Refs. [58, 64] for recent reviews on the topic, with many other examples of problems that exhibit such a hard phase or computational-to-statistical gap, as it is also sometimes called in the literature.

Performance of PCA. Fig 2 also shows the reconstruction error of PCA applied to either the weights (green) or the Fisher matrix (red). Note that while the error of AMP is rescaled by ρ , we plot the reconstruction of PCA without rescaling (which is why we have a second y -axis in all three plots). In other words, the reconstruction error of AMP is multiplied by 20 in the

left-most plot of Fig 2, and is still below that of PCA. We rescale the AMP error in this way to ensure that the errors are comparable for different values of the sparsity, since AMP returns a sparse estimate in all cases. PCA on the other hand is agnostic about the sparsity of the patterns, and returns a dense reconstruction regardless of the value of ρ . The PCA error at high noise thus scales as $1 + \rho$. Again, we see that AMP outperforms PCA in terms of the reconstruction error, with the largest difference coming at lowest sparsity. This is to be expected, since the AMP algorithm can take information about the prior into account.

Reconstructing patterns with low coding level

As a third and final example, we consider the reconstruction of patterns with low coding level. For this, we will draw patterns from the prior distribution

$$p_X(\mathbf{x}_i) = \prod_{j=1}^P p_{x_j}(x_{ij}) = \prod_{j=1}^P [(1 - \rho)\delta(x_{ij} + \rho) + \rho\delta(x_{ij} - (1 - \rho))] \tag{24}$$

which is related to models proposed first by Tsodyks and Feigel'man [21, 65], that considered the storage of binary (0,1) patterns of activity (where 0 means inactive and 1 active), with a 'coding level' ρ (probability that a neuron is active in a given pattern). The motivation for studying the sparse $\rho \ll 1$ limit is that the activity in brain structures involved in memory is typically sparse—for instance, ρ has been estimated to be on the order of 0.01 in animals ranging from rodents to humans [66, 67]. As in the previous cases, the patterns correspond to deviations of activity from its mean, i.e. $(1 - \rho, -\rho)$ instead of $(1,0)$. This model has mean zero and a covariance matrix $\rho(1 - \rho)I_P$. We still use the channel corresponding to the connectivity structure of the rectified Hopfield model (3), so while the Fisher score matrix S stays the same, we have a new threshold function and hence a new state evolution, which we derive in *Methods*.

Here, we plot the performance of the algorithm together with the results of iterating the state evolution equation (*Methods*) for two different values of the coding level ρ in Fig 3. Note that we again rescale the reconstruction error of AMP by $(1 - \rho)\rho$ to ensure comparability of the results for different values of the coding level ρ . The error curves for PCA (green and red) in Fig 3 on the other hand are not rescaled. We note that again, AMP outperforms PCA in




Fig 3. Performance of various algorithms when reconstructing a pattern with low coding level following Tsodyks' prior. (Top) Same plot as Figs 1 and 2, for a single stored pattern drawn from Tsodyks' prior (24). For $\rho < 0.211$, the performance of the algorithm with random initialisation (15) is different from the performance with informed initialisation (16), which might be the signature of a hard phase. Parameters: $\tau = 0$, $N = 5000$ for AMP, $N = 20000$ for PCA.

<https://doi.org/10.1371/journal.pcbi.1010813.g003>

terms of the reconstruction error, which for random guessing tends to $1 + \rho(1 - \rho)$ in the case of PCA.

The hardness of recovering patterns with low coding level. In this model too, we have a uniform fixed point with $m^t = 0$. Expanding around this fixed point yields the update equation

$$m^{t+1} = \frac{(\rho - 1)^2 \rho^2}{\Delta} m^t \tag{25}$$

so we find that the critical value of the noise where the uniform or uninformative fixed point becomes unstable in this model is $\Delta < \Delta_c = (\rho - 1)^2 \rho^2$.

Recently, a closed-form sufficient criterion for the existence of a hard phase in models that have a prior with zero mean was derived in [52]; namely, a hard phase exists if the prior is “skewed” in the statistical sense, such that

$$\langle x^3 \rangle^2 > 2 \langle x^2 \rangle^3 \tag{26}$$

where the average is taken with respect to the prior distribution of this model, Eq (23). For the prior (24), this criterion predicts the existence of a first order phase transition and hence of a hard phase for $\rho > \rho_c = 1/2 - 1/\sqrt{12} \simeq 0.2113$ where we assume w.l.o.g. that $\rho < 1/2$. Note that this is a sufficient condition, and not a necessary one; in fact, for the sparse Hopfield prior, which is symmetric around zero and has $\langle x^3 \rangle = 0$, we cannot calculate the critical value of ρ using Eq (26).

Reconstructing even more patterns: How far can we go?

A natural question that arises for the algorithms we have derived is how many patterns we can reliably reconstruct. In practice, the bottleneck for reconstructing the patterns using AMP is computing the partition function of the Gaussian approximation of the posterior density of a column-vector \mathbf{x} at every step of the algorithm (see the detailed explanation of the AMP algorithm in the Methods section):

$$W(\mathbf{x}; \mathbf{A}, \mathbf{b}) = \frac{1}{Z(\mathbf{A}, \mathbf{b})} p_{\mathbf{x}}(\mathbf{x}) \exp \left(\sum_j^P b_j x_j - \frac{1}{2} \sum_{j,k} x_j A_{jk} x_k \right). \tag{27}$$

Evaluating the mean and the variance of this distribution even for the simple Hopfield prior (13) requires summing 2^P terms, so the computational cost is exponential in the number of patterns stored. We can circumvent this bottleneck by computing the function $W(\mathbf{x}; \mathbf{A}, \mathbf{b})$ using a mean-field approximation [68], which was originally proposed in [69].

The mean-field approximation. We thus approximate the posterior distribution by a factorised distribution, replacing the full covariance matrix \mathbf{A} with a vector \mathbf{a} that contains only the variance of the $j = 1, \dots, P$ elements of \mathbf{x} ,

$$\tilde{W}(\mathbf{x}; \tilde{\mathbf{a}}, \tilde{\mathbf{b}}) = \prod_j^P \frac{1}{\tilde{Z}(\tilde{a}_j, \tilde{b}_j)} p_{\mathbf{x}}(x_j) \exp \left(\tilde{b}_j x_j - \frac{1}{2} \tilde{a}_j x_j^2 \right). \tag{28}$$

where we use the tilde to denote mean-field quantities. We have implemented mean-field approximations for the models discussed thus far and we show the performance of AMP with this approximation for the three models discussed so far in Fig 4 with the state evolution prediction for the reconstruction of a single stored pattern $P = 1$ without the mean-field approximation. The picture that emerges is similar for all three models studied here: the algorithm with the mean-field approximation is able to reconstruct the stored patterns just as well as if it was looking at the reconstruction of a single pattern up to a certain noise level, beyond which




Fig 4. MSE from reconstructing with the mean-field prior approximation. We plot the final mse obtained by the AMP algorithm (no damping), using a mean-field approximation to compute the threshold function. The solid line is the mse predicted by iterating the state evolution for the scalar variable, Eq (14). We choose $N = 5000$ and set $\tau = 0$.

<https://doi.org/10.1371/journal.pcbi.1010813.g004>

performance quickly deteriorates. Intuitively, the cross-talk between the stored patterns introduces an additional source of noise for the reconstruction, which leads to failure to reconstruct the stored patterns at lower Δ than in the case $P = 1$.

Scaling of the critical number of stored patterns. Throughout this paper, we have relied on the assumption that the matrix J is approximately low-rank, in the sense that its eigenvalues can be separated into large bulk of eigenvalues, from a which only $P \sim \mathcal{O}(1)$ eigenvalues corresponding to the stored patterns detach. The mean-field approximation that we just introduced makes it now computationally feasible to run the reconstruction algorithm even with a large number of patterns. This raises an important practical question: for which number of patterns does the algorithm based on the low-rank approximation break down?

We investigate this question in all three models numerically as follows. We fixed the noise level Δ in all experiments at 20% of the critical noise level beyond which it is information-theoretically impossible to weakly reconstruct the stored patterns. We then ran twenty reconstruction experiments with $N = 1000$ for each value of P ; the final mse for each run is shown with a dot on the left of Fig 5. For $P = 25$, all twenty runs gave an mse well below the threshold; in fact, the reconstruction error per pattern mse/P is as low as if we had stored only a single pattern in the network (blue line). As we increase the number of patterns P , the first time the algorithm is not able to recover all patterns is for $P = 29$ patterns. For $P = 36$ patterns, the algorithm did not reconstruct the patterns with an error better than chance a single time. Given the clear separation between successful runs with low error and unsuccessful runs with an error that is essentially random guessing, we set the threshold for the reconstruction mse below which we consider the algorithm successful at 20% of the trivial mse obtained by random guessing (orange line). At $P = 34$, the algorithm fails to achieve an mse below the threshold in more than 50% of the cases. We define the critical number of stored patterns P_{crit} as the largest number of patterns that can be reconstructed with an mse below the threshold in at least 50% of the runs, so $P_{crit} = 33$ in this case. On the right, we show the values of P_{crit} for all three models (binary, sparse and skewed) as a function of N . In each case, we find that $P_{crit} \sim N^\gamma$, with the exponent γ between $\gamma \approx 0.5$ for sparse patterns and $\gamma \approx 0.7$ for patterns from Tsodyk's prior. The exact values of the exponents will depend on the choices of the threshold and especially the noise level. While the separation between successful and unsuccessful runs allows for various error thresholds without changing the results, the exponents are more sensitive to the noise level, since higher noise levels induce higher fluctuations in the reconstruction errors closer to the critical noise (cf. Fig 4).




Fig 5. Scaling of the critical number of patterns that can reliably be reconstructed with mean-field message passing. (Left) For binary patterns, we plot the final reconstruction error of twenty instances of the mean-field message-passing algorithm as a function of the number of stimuli stored P . We define the critical number of patterns that can be reconstructed as the largest P at which more than half the runs yield an mse better than a threshold value (here, 20% of the trivial mse , i.e. 0.2). Parameters: $N = 1000$, $\Delta = 0.2\Delta_c$, $\tau = 0$. (Right) We plot the highest number of patterns P_{max} that could be reconstructed with an error below 20% of the trivial error in at least 50% of cases. For all three models considered, we show experimental results (crosses) and the exponent of a power law fit $P_{\text{max}} \sim N^\gamma$. In all cases, $\tau = 0$ and the noise level Δ was 20% of the critical noise level Δ_c where reconstruction becomes impossible. $\rho = 0.3$ for both the sparse Hopfield and Tsodyks's prior.

<https://doi.org/10.1371/journal.pcbi.1010813.g005>

Discussion

We have derived the conditions under which it is possible to reconstruct the patterns stored in a recurrent Hopfield-type network from knowledge of the connections alone, in a case where those weights are obtained by a learning rule which is a noisy, non-linear version of the classic Hebbian rule. We have implemented and provided practical algorithms to do so: a Bayesian approach based on message passing, and classic non-parametric approaches such as PCA on the weights and on the Fisher score matrix. The message-passing algorithm offer a principled way to reconstruct the patterns using prior information about the prior of the inputs, while the spectral methods are robust and easy-to-implement, but fail to take this extra information into account.

The performance of the algorithms can be captured by an effective noise level, which takes both the synaptic noise and the thresholding of the learning rule into account. We found that the message-passing algorithm beats the PCA reconstruction across noise levels, which is due to the principled way in which the message-passing algorithm can incorporated prior knowledge about the distribution of patterns.

While our theoretical results were obtained in the limit where the number of patterns stored is of order 1, while the number of neurons in the network tends to infinity, we have also explored a mean-field approximation that allowed us to study how the maximum number of patterns that can be reconstructed with some reliability scales with the pattern dimension. We now discuss some directions in which our work can be extended.

The case of unknown generative model

We have assumed that the hyper-parameters of the problem, such as the number of stored stimuli P or the generative model of the stimuli $p_X(x)$, are known before reconstruction. However, it is possible to extend both the message-passing algorithm and its analysis to the case of unknown hyper-parameters. Within our framework, a natural approach to learn the values of

these hyper-parameters is closely related to the well-known expectation maximisation algorithm [70]. A detailed study of expectation maximisation to obtain various hyper-parameters for message-passing algorithms was reported by Krzakala et al. [71] for the statistical estimation problem of “compressed sensing”, where one aims to reconstruct a signal from a number of measurements that is smaller than the number of unknowns. Another approach to deal with an unknown noise distribution was recently proposed by Montanari et al. [72] for the related problem of matrix denoising, where one aims to reconstruct an unknown, low-rank matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$ from observations $\mathbf{Y} = \mathbf{X} + \mathbf{W}$ when \mathbf{W} is a noise matrix with independent and identically distributed entries. Montanari et al. [72] propose an iterative approach to estimate the noise distribution from the observations \mathbf{Y} . An interesting direction for future work is to explore both strategies for the problem of reconstructing stored patterns in connectivity matrices.

From the point of view of the theoretical analysis, assuming the “wrong” rank in the algorithm in the sense that there is a mismatch between the number of patterns one wants to infer and the number of patterns stored makes the analysis much more involved. The message-passing algorithms can be modified to take this mismatch into account. While first steps towards a theoretical analysis and the development of modified message-passing have recently been made [73–75], we leave the exploration of this direction in the context of neuroscience to future work.

Finally, there could also be no (discernible) signal in the connectivity matrix, because the relative strength of the noise is too large, or because the connectivity is purely random. As we discussed above, AMP shows a clear advantage over PCA in this case: while AMP will return inputs which are close to 0, indicating that it didn’t find any structure in the data, PCA will always return the leading eigenvector, which in this case however is completely uninformative.

Reconstructing an extensive number of patterns

Another important extension of our work would be the reconstruction of an extensive number of patterns from the connectivity \mathbf{J} , *i.e.* whether the regime $P \sim \mathcal{O}(N)$ is accessible. This is essentially the problem of factorising a large matrix with extensive rank, which is also known as dictionary learning [76–78]. While we note that there have been promising signs of progress recently [79, 80], this remains a hard problem with implications that would go far beyond the application discussed here.

Reconstructing the patterns from the dynamics

It may be possible to retrieve a pattern given a distorted or noisy version of it by initialising the network with this distorted pattern and running the network dynamics (1) until convergence. Another algorithm to reconstruct the patterns is then to run the network dynamics from different initial conditions until convergence, and to take the resulting fixed points as estimates of the patterns stored in the connectivity. However, this procedure will only yield stored patterns if the initial condition is correlated with one of the patterns stored in the network. It is thus more apt to speak about “pattern completion” in this context [49]. Here, we considered the more challenging problem where such partial observations of the patterns are not available, which would be the case if the only information available is the connectivity matrix.

It is still fair to wonder how well this algorithm could do when starting from random initial conditions. In the classical Hopfield model, where the weights are given by the Hebb rule (2), any initial condition of the dynamics that has non-zero macroscopic overlap with a stored pattern is guaranteed to converge to that pattern when the number of patterns is of order one. In a finite network, one then expects random initial conditions to converge to one of the stored

patterns (depending on the random initial overlaps). While we study a learning rule that is different from Hebb's rule, we note a classic result by Sompolinsky [35] that suggests that the dynamics of recurrent networks whose connections are a nonlinear, noisy version of the Hebbian weights (2) are largely unaffected by this change.

For a single pattern, iterating the dynamics resembles the power method [81] to compute the first principal component of a symmetric matrix, and the dynamical approach to reconstruction thus closely resembles the PCA algorithm that we analysed in the case $P = 1$. For more than one pattern, there is no guarantee however that *all* the patterns will be recovered. Moreover, the performance of this algorithm is expected to be strongly model-dependent: in more realistic models such as the Tsodyks-Feigel'man model [21, 65], the quiescent state, where all neurons are at zero, is stable, and convergence to one of the stored patterns will necessitate an initial overlap that is larger than some non-zero critical value. Recovering the patterns through the dynamics is not expected to work in this case.

Similarly, this strategy is also expected to fail when reconstructing an extensive number of patterns. In this case, we need a correlation between the initial condition and each pattern which is bounded away from zero, and obtaining such a correlation for all the patterns would require exponentially many trials.

What information on synaptic weights does connectomic data provide?

EM reconstructions of neuronal circuits are of sufficiently high resolution that they can enable measurements of the volume of dendritic spines, the anatomical structures on which the vast majority of synaptic connections between pyramidal cells are formed (see e.g. [14]). The volume of dendritic spines has in turn been shown to be strongly correlated with synaptic strength, as measured by the amplitude of post-synaptic potentials triggered by pre-synaptic activity [82–84]. Noise in measurement of dendritic spines, and the lack of perfect correlation between volume and EPSP amplitudes, are two of the reasons (together with noise in the learning process) for the introduction of the noise matrix ζ_{ij} in our model.

Towards more biologically realistic models

Our focus here was to analyse the simplest model where the problem of retrieving the patterns is mathematically well-posed, and neither trivial nor impossible to achieve. We now discuss various extensions of this model that could be addressed in future work.

Asymmetry in the learned component of the connectivity matrix. We have focused here on a symmetric connectivity matrix, for multiple reasons. First, multiple *in vitro* studies in both cortex and hippocampus have shown that local networks in these structures exhibit a significant degree of symmetry, as evidenced by a much higher probability of bidirectional connections in pairs of neurons, compared to a random directed Erdos-Renyi network [85–88], with the notable exception of rodent barrel cortex where no such overrepresentation exists [89]. In addition, synaptic plasticity in area CA3 of the hippocampus has been shown to be temporally symmetric [90]. One thus expects such synaptic plasticity rules to lead to connectivity matrices with a strong degree of symmetry in this area. In cortex, plasticity rules are temporally asymmetric as a function of the timing difference between spikes of pre and post-synaptic neurons. However, plasticity depends also strongly on firing rates of pre and post-synaptic neurons, and if dependence on firing rate dominates over the dependence on spike timing as has been suggested [91], then one also expects a strongly symmetric Hebbian component in connectivity matrices. Finally, a strong degree of symmetry is also consistent with the observation of persistent activity in multiple cortical areas in rodents [92] and primates [93–96]. Consistent with this idea, Inagaki et al [92] showed using optogenetic perturbations that

the dynamics in area ALM of the mouse during a short-term memory task exhibit multiple characteristics of attractor dynamics. Of course, we do not expect connectivity matrices in all brain structures to be well captured by noisy symmetric matrix. In particular, networks storing temporal sequences [97–99] would need to contain a significant asymmetric learned component. Our methods would then need to be extended to asymmetric matrices. From the practical point of view, the application of the method proposed makes only sense when the reconstructed connectivity matrices exhibit a significant degree of symmetry.

Inhibition. Here, we have assumed inhibition is not involved in learning and simply provides a uniform inhibition to excitatory neurons, equivalent to setting a threshold for active neurons. This traditional view is consistent with the observation of high connection probabilities between a specific type of inhibitory neuron (PV positive interneurons) and pyramidal cells [100]. It is also consistent with the observation in hippocampus that inhibition is only weakly modulated by spatial location [32] (see however [101]). However, it has also been shown that synaptic connections involving inhibitory interneurons also exhibit plasticity [102], and it has been argued that plasticity of such connections could greatly expand storage capacity [34]. Our methods could be extended to the addition of inhibitory neurons, with a few caveats: estimate of the strength of such synapses might be more challenging, since synapses involving interneurons are formed directly on dendritic shafts and not on spines; Also, the connectivity matrix will then necessarily be asymmetric.

Memories stored with different strengths. In our model, as in most associative memory models, memories are stored with identical strengths. A straightforward extension of the model would be a model where memories have different embedding strengths. This class of models include ‘palimpsest’ models in which recently stored patterns gradually erase older patterns that are progressively forgotten [103–106]. Note that in this class of models, our method would be likely to infer only the most strongly embedded patterns, and memories that are on their way to becoming forgotten would not be likely to be inferred. From the point of view of message-passing algorithms, the gradual erasure would have to be modelled through the learning rule of the model, Eq (3), together with a prior over embedding strengths.

Distributions of synaptic weights. Our model leads to a truncated Gaussian distribution of non-negative synaptic weights. Distributions of experimentally recorded EPSP amplitudes [86] as well as spine volumes [107] have been fitted using log-normal distributions. Our method can easily be generalized to networks with arbitrary distributions of non-negative weights, by using a suitable non-linearity Φ in Eq (3) [34].

Binary Hebbian matrices. It has been proposed by some authors that synapses store information in a digital, not analogue fashion [35, 108, 109]. In this scenario, synapses have only a few stable states, and plasticity events correspond to transition between these states. The resulting model would then bear similarities with stochastic block models (SBMs) [110], where groups of neurons representing a particular stored memory would correspond to communities in SBMs. An important difference is that in the case of random patterns, there would be overlaps between these groups [111]. One could use similar methods as the ones proposed here to deal with such a scenario, since recovering communities in stochastic block models can be reformulated as a low-rank matrix factorisation problem, and is hence amenable to the same analysis techniques that we used here, see refs. [52, 64, 112] for examples of this rich literature.

Distributions of patterns. We have assumed that stored memories are random i.i.d. binary vectors. Responses of neurons to external inputs rarely follow a bimodal distributions, and can sometimes be better described by a unimodal lognormal distribution [113]. However, non-linearities associated with the synaptic plasticity rule could potentially binarize stored memories [24], which would then result in a model that is very similar to the model investigated here.

Validation

An important question is how to validate the results of such an analysis. One possibility would be simulate the dynamics of a network whose connectivity matrix is set to be the reconstructed matrix, using the inferred patterns as initial conditions, to check that the dynamics converges to fixed point attractors that are close to the inferred patterns.

Methods

A formal analogy between inference problems and statistical physics

It may come as a surprise that statistical physics can be helpful in solving and analysing inference problems like the one considered here. The connection between the two becomes more transparent if we introduce the interaction $g(\cdot) \equiv \ln P_{\text{out}}(\cdot)$ to rewrite the posterior as

$$P(\mathbf{X}|\mathbf{J}) = \frac{1}{Z(\mathbf{J})} \prod_i^N P_{\mathbf{X}}(\mathbf{x}_i) \prod_{j>i}^N \exp \left[g \left(J_{ij}, W_{ij} \right) \right]. \tag{29}$$

This distribution describes the posterior density over estimates \mathbf{X} . However, it can also be interpreted as the Gibbs or Boltzmann distribution that describes the properties of complex, disordered systems such as glasses. This analogy can be leveraged by exploiting tools from the statistical physics of disordered systems to tackle the—hard—inference problem that is inferring the patterns from the connectivity \mathbf{J} . The key ideas behind the AMP algorithm that we discuss throughout this work and AMP algorithms in general first appeared in a paper by Thouless, Anderson and Palmer [114] that dealt with physical systems described by an equilibrium distribution of the type (29). State evolution techniques were first introduced for compressed sensing problems by Donoho, Maleki and Montanari [50] based on ideas from [55], but it too is based on ideas from statistical physics often referred to as the cavity method [56]. For a much more detailed on the links between statistical physics and inference problem, see [49, 58, 115].

Approximate message passing for low-rank matrix reconstruction

We are now in a position to state the AMP algorithm for the factorisation of symmetric low-rank matrices that in this form was derived by Lesieur [52], building on the previous works deriving AMP-type algorithms for particular instances of this problem class [44, 53, 54]. We refer the interested reader to these papers for details on the derivation of this class of algorithms and their relation to belief propagation and spectral methods.

To describe the algorithm, we first define the Fisher score matrix as a transformation of the data matrix \mathbf{J} given for the inference:

$$S_{ij}(J_{ij}) \equiv \left. \frac{\partial \ln P_{\text{out}}(J_{ij}|w)}{\partial w} \right|_{w=0} \tag{30}$$

For the channel corresponding to the rectified Hopfield model (3), we find

$$S_{ij}(J_{ij}) = \begin{cases} -\frac{2e^{-\tau^2/2\nu^2}}{\sqrt{2\pi} \nu \operatorname{erfc}(-\tau/\sqrt{2\nu})}, & J_{ij} = 0 \\ \frac{J_{ij} + \tau}{\nu^2} & \text{otherwise.} \end{cases} \tag{31}$$

Low-RAMP is an iterative algorithm: at every step t , it computes a new estimate of the mean $\hat{\mathbf{x}}_i^{t+1}$ and the variance σ_i^{t+1} as

$$\hat{\mathbf{x}}_i^{t+1} = f(\mathbf{A}_i^t, \mathbf{b}_i^t) \tag{32a}$$

$$\sigma_i^{t+1} = \partial_{\mathbf{b}} f(\mathbf{A}_i^t, \mathbf{b}_i^t) \tag{32b}$$

where the threshold function was defined in Eq (11) and is repeated here for convenience:

$$f(\mathbf{A}, \mathbf{b}) = \frac{1}{Z(\mathbf{A}, \mathbf{b})} \sum_{\mathbf{x} \in \mathcal{X}^p} \mathbf{x} p_{\mathcal{X}}(\mathbf{x}) \exp\left(\mathbf{b}\mathbf{x} - \frac{1}{2} \mathbf{x}^\top \mathbf{A} \mathbf{x}\right). \tag{33}$$

There exists a set of parameters $\mathbf{A}_i \in \mathbb{R}^{p \times p}$ and $\mathbf{b}_i \in \mathbb{R}^p$ for every marginal, which in turn are updated as

$$\mathbf{b}_i^t = \frac{1}{\sqrt{N}} \sum_k^N S_{ki} \hat{\mathbf{x}}_k^t - \left(\frac{1}{N} \sum_k^N S_{ki}^2 \sigma_k^t \right) \hat{\mathbf{x}}_i^{t-1} \tag{34a}$$

$$\mathbf{A}_i^t = \frac{1}{N} \sum_k^N S_{ki}^2 \hat{\mathbf{x}}_k^t \hat{\mathbf{x}}_k^{t\top} \tag{34b}$$

To run the algorithm, we perform these steps:

1. Given the matrix \mathbf{J} , compute the Fisher score matrix \mathbf{S} using Eq (31).
2. For all $i = 1, \dots, N$, initialise the parameters \mathbf{b}_i and \mathbf{A}_i such that all entries are zero. Initialise all estimators $\hat{\mathbf{x}}_i^t$ with a random draw from the prior distribution $p_{\mathcal{X}}(\mathbf{x})$ and set $\hat{\mathbf{x}}_i^{t-1}$ to all zeros for the first step. (There is no need to initialise σ_i).
3. Compute first the update to \mathbf{A}_i^t and \mathbf{b}_i^t following Eqs (34b) and (34a), then compute the new means $\hat{\mathbf{x}}_i^{t+1}$ and their variance σ_i^{t+1} using Eqs (32a) and (32b).
4. Repeat Step 3 until the squared difference between all $\hat{\mathbf{x}}_i^t$ and $\hat{\mathbf{x}}_i^{t+1}$ is smaller than some pre-defined threshold ϵ .

We also provide an implementation of this algorithm in a [Python package](#) that was the base of all the programs written for this paper.

State evolution for reconstructing sparse patterns

We discussed in the main text that the sparse prior (23) has mean $\langle x \rangle = 0_p$ and covariance $\langle xx^\top \rangle = \rho \mathbf{I}_p$. The state evolution will interpolate between an order parameter that is all zeros at initialisation, and $M = \rho \mathbf{I}_p$ for perfect reconstruction. We are thus motivated to use the ansatz $M^t = m^t \mathbf{I}_p$. The threshold function then becomes

$$[f(\mathbf{A}, \mathbf{B})_k] = \frac{\rho e^{-A_{kk}/2} \sinh(B_k)}{1 + \rho [e^{-A_{kk}/2} \cosh(B_k) - 1]} \tag{35}$$

Substituting this form into the SE update Eq (12) yields a closed update equation for the parameter m^t ,

$$m^{t+1} = \mathbb{E}_w \frac{\rho^2 e^{-m^t/2\Delta} \sinh(m^t/\Delta + \sqrt{m^t/\Delta} z)}{1 + \rho [e^{-m^t/2\Delta} \cosh(m^t/\Delta + \sqrt{m^t/\Delta} z) - 1]}, \tag{36}$$

where z is again a scalar Gaussian random variable with zero mean and unity variance, like in Eq (14). We can recover the state evolution for the symmetric rectified Hopfield model from Eq (36), in the limit $\rho \rightarrow 1$.

State evolution for reconstructing patterns with low coding level

For Tsodyk's prior (24), we have a new threshold function (11) which reads

$$[f(A, B)]_k = 1 - \rho - \frac{1 - \rho}{1 - \rho \left(1 - e^{A_{kk}(\rho - 1/2) + B_k} \right)} \quad (37)$$

The prior distribution (24) allows us to use the same ansatz for the magnetisation that we used above to analyse sparse patterns. Setting $M^t = m^t \mathbf{I}_p$, we get the following update equation for the scalar order parameter $m^t > 0$:

$$m^{t+1} = \mathbb{E}_W \frac{(\rho - 1)^2 \rho^2 (e^{m^t/\Delta} - 1) e^{w\sqrt{m^t/\Delta}}}{\left(\rho e^{w\sqrt{m^t/\Delta}} - (\rho - 1) e^{m^t/2\Delta} \right) \left(\rho \exp\left(\frac{2w\sqrt{\Delta m^t} + m^t}{2\Delta} \right) - \rho + 1 \right)}. \quad (38)$$

Acknowledgments

SG, FK and LZ would like to thank the Department of Mathematics at Duke University for their hospitality during an extended visit.

Author Contributions

Conceptualization: Sebastian Goldt, Florent Krzakala, Lenka Zdeborová, Nicolas Brunel.

Formal analysis: Sebastian Goldt, Florent Krzakala, Lenka Zdeborová, Nicolas Brunel.

Funding acquisition: Lenka Zdeborová.

Investigation: Sebastian Goldt, Florent Krzakala, Lenka Zdeborová, Nicolas Brunel.

Methodology: Sebastian Goldt, Florent Krzakala, Lenka Zdeborová, Nicolas Brunel.

Software: Sebastian Goldt.

Writing – original draft: Sebastian Goldt, Florent Krzakala, Lenka Zdeborová, Nicolas Brunel.

Writing – review & editing: Sebastian Goldt, Florent Krzakala, Lenka Zdeborová, Nicolas Brunel.

References

1. White JG, Southgate E, Thomson JN, Brenner S. The Structure of the Nervous System of the Nematode *Caenorhabditis elegans*. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 1986; 314(1165):1–340. PMID: [22462104](https://pubmed.ncbi.nlm.nih.gov/22462104/)
2. Ohyama T, Schneider-Mizell CM, Fetter RD, Aleman JV, Franconville R, Rivera-Alba M, et al. A multi-level multimodal circuit enhances action selection in *Drosophila*. *Nature*. 2015; 520(7549):633–639. <https://doi.org/10.1038/nature14297> PMID: [25896325](https://pubmed.ncbi.nlm.nih.gov/25896325/)
3. Eichler K, Li F, Litwin-Kumar A, Park Y, Andrade I, Schneider-Mizell CM, et al. The complete connectome of a learning and memory centre in an insect brain. *Nature*. 2017; 548(7666):175–182. <https://doi.org/10.1038/nature23455> PMID: [28796202](https://pubmed.ncbi.nlm.nih.gov/28796202/)

4. Xu CS, Januszewski M, Lu Z, Takemura Sy, Hayworth KJ, Huang G, et al. A Connectome of the Adult *Drosophila* Central Brain. *bioRxiv*. 2020;. <https://doi.org/10.7554/eLife.57443> PMID: 32880371
5. Scheffer LK, Xu CS, Januszewski M, Lu Z, Takemura Sy, Hayworth KJ, et al. A Connectome and Analysis of the Adult *Drosophila* Central Brain. *bioRxiv*. 2020;. <https://doi.org/10.7554/eLife.57443> PMID: 32880371
6. Wanner AA, Friedrich RW. Whitening of odor representations by the wiring diagram of the olfactory bulb. *Nat Neurosci*. 2020; 23(3):433–442. <https://doi.org/10.1038/s41593-019-0576-z> PMID: 31959937
7. Briggman KL, Helmstaedter M, Denk W. Wiring specificity in the direction-selectivity circuit of the retina. *Nature*. 2011; 471(7337):183–188. <https://doi.org/10.1038/nature09818> PMID: 21390125
8. Helmstaedter M, Briggman KL, Turaga SC, Jain V, Seung HS, Denk W. Connectomic reconstruction of the inner plexiform layer in the mouse retina. *Nature*. 2013; 500(7461):168–174. <https://doi.org/10.1038/nature12346> PMID: 23925239
9. Kim JS, Greene MJ, Zlateski A, Lee K, Richardson M, Turaga SC, et al. Space–time wiring specificity supports direction selectivity in the retina. *Nature*. 2014; 509(7500):331–336. <https://doi.org/10.1038/nature13240> PMID: 24805243
10. Mishchenko Y, Hu T, Spacek J, Mendenhall J, Harris KM, Chklovskii DB. Ultrastructural Analysis of Hippocampal Neuropil from the Connectomics Perspective. *Neuron*. 2010; 67(6):1009–1020. <https://doi.org/10.1016/j.neuron.2010.08.014> PMID: 20869597
11. Kasthuri N, Hayworth KJ, Berger DR, Schalek RL, Conchello JA, Knowles-Barley S, et al. Saturated reconstruction of a volume of neocortex. *Cell*. 2015; 162(3):648–661. <https://doi.org/10.1016/j.cell.2015.06.054> PMID: 26232230
12. Lee WCA, Bonin V, Reed M, Graham BJ, Hood G, Glattfelder K, et al. Anatomy and function of an excitatory network in the visual cortex. *Nature*. 2016; 532(7599):370–374. <https://doi.org/10.1038/nature17192> PMID: 27018655
13. Motta A, Berning M, Boergens KM, Staffler B, Beining M, Loomba S, et al. Dense connectomic reconstruction in layer 4 of the somatosensory cortex. *Science*. 2019; p. eaay3134. <https://doi.org/10.1126/science.aay3134> PMID: 31649140
14. Dorkenwald S, Turner NL, Macrina T, Lee K, Lu R, Wu J, et al. Binary and analog variation of synapses between cortical pyramidal neurons. *bioRxiv*. 2019;.
15. Macrina T, Lee K, Lu R, Turner NL, Wu J, Popovych S, et al. Petascale neural circuit reconstruction: automated methods. *bioRxiv*. 2021.
16. Consortium M, Alexander Bae J, Baptiste M, Bodor AL, Brittain D, Buchanan J, et al. Functional connectomics spanning multiple areas of mouse visual cortex. *bioRxiv*. 2021.
17. Motta A, Schurr M, Staffler B, Helmstaedter M. Big data in nanoscale connectomics, and the greed for training labels. *Current Opinion in Neurobiology*. 2019; 55:180–187. <https://doi.org/10.1016/j.conb.2019.03.012> PMID: 31055238
18. Litwin-Kumar A, Turaga SC. Constraining computational models using electron microscopy wiring diagrams. *Current Opinion in Neurobiology*. 2019; 58:94–100. <https://doi.org/10.1016/j.conb.2019.07.007> PMID: 31470252
19. Brunel N. Is cortical connectivity optimized for storing information? *Nature Neuroscience*. 2016; 19(5):749–755. PMID: 27065365
20. Hopfield JJ. Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci U S A*. 1982; 79(8):2554–2558. <https://doi.org/10.1073/pnas.79.8.2554> PMID: 6953413
21. Tsodyks MV, Feigel'man MV. The Enhanced Storage Capacity in Neural Networks with Low Activity Level. *Europhys Lett*. 1988; 6(2):101–105. <https://doi.org/10.1209/0295-5075/6/2/002>
22. Amit DJ. The Hebbian paradigm reintegrated: Local reverberations as internal representations. *Behavioral and Brain Sciences*. 1995; 18(4):617–626. <https://doi.org/10.1017/S0140525X00040164>
23. Amit DJ, Brunel N. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral cortex (New York, NY)*. 1991; 7(3):237–252. PMID: 9143444
24. Pereira U, Brunel N. Attractor dynamics in networks with learning rules inferred from in vivo data. *Neuron*. 2018; 99(1):227–238. <https://doi.org/10.1016/j.neuron.2018.05.038> PMID: 29909997
25. Dayan P, Abbott L. *Theoretical neuroscience*. MIT press, Cambridge; 2001.
26. Mastrogiuseppe F, Ostojic S. Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. *Neuron*. 2018; 99(3):609–623. <https://doi.org/10.1016/j.neuron.2018.07.003> PMID: 30057201

27. Schuessler F, Mastrogiuseppe F, Dubreuil A, Ostojic S, Barak O. The interplay between randomness and structure during learning in RNNs. *Advances in neural information processing systems*. 2020; 33:13352–13362.
28. Sussillo D, Abbott LF. Generating coherent patterns of activity from chaotic neural networks. *Neuron*. 2009; 63(4):544–557. <https://doi.org/10.1016/j.neuron.2009.07.018> PMID: 19709635
29. Logiaco L, Abbott L, Escola S. Thalamic control of cortical dynamics in a model of flexible motor sequencing. *Cell reports*. 2021; 35(9):109090. <https://doi.org/10.1016/j.celrep.2021.109090> PMID: 34077721
30. Hebb DO. *The organization of behavior: A neuropsychological approach*. John Wiley & Sons; 1949.
31. Amit DJ, Gutfreund H, Sompolinsky H. Statistical mechanics of neural networks near saturation. *Annals of physics*. 1987; 173(1):30–67. [https://doi.org/10.1016/0003-4916\(87\)90092-3](https://doi.org/10.1016/0003-4916(87)90092-3)
32. Grienberger C, Milstein AD, Bittner KC, Romani S, Magee JC. Inhibitory suppression of heterogeneously tuned excitation enhances spatial coding in CA1 place cells. *Nat Neurosci*. 2017; 20:417–426. <https://doi.org/10.1038/nn.4486> PMID: 28114296
33. Lim S, McKee JL, Woloszyn L, Amit Y, Freedman DJ, Sheinberg DL, et al. Inferring learning rules from distributions of firing rates in cortical neurons. *Nature neuroscience*. 2015; 18(12):1804. <https://doi.org/10.1038/nn.4158> PMID: 26523643
34. Mongillo G, Rumpel S, Loewenstein Y. Inhibitory connectivity defines the realm of excitatory plasticity. *Nature neuroscience*. 2018; 21(10):1463–1470. <https://doi.org/10.1038/s41593-018-0226-x> PMID: 30224809
35. Sompolinsky H. Neural networks with nonlinear synapses and a static noise. *Physical Review A*. 1986; 34(3):2571. <https://doi.org/10.1103/PhysRevA.34.2571>
36. Féral D, Pécché S. The largest eigenvalue of rank one deformation of large Wigner matrices. *Communications in mathematical physics*. 2007; 272(1):185–228. <https://doi.org/10.1007/s00220-007-0209-3>
37. Potters M, Bouchaud JP. *A First Course in Random Matrix Theory: For Physicists, Engineers and Data Scientists*. Cambridge University Press; 2020.
38. Benaych-Georges F, Nadakuditi RR. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*. 2011; 227(1):494–521. <https://doi.org/10.1016/j.aim.2011.02.007>
39. Amit DJ, Gutfreund H, Sompolinsky H. Storing infinite numbers of patterns in a spin-glass model of neural networks. *Phys Rev Lett*. 1985; 55(14):1530–1533. <https://doi.org/10.1103/PhysRevLett.55.1530> PMID: 10031847
40. Amit DJ, Gutfreund H, Sompolinsky H. Spin-glass models of neural networks. *Phys Rev A*. 1985; 32(2):1007–1018. <https://doi.org/10.1103/PhysRevA.32.1007>
41. Nishimori H. *Statistical Physics of Spin Glasses and Information Processing*. 1st ed. Oxford: Oxford University Press; 2001.
42. Wright J, Ganesh A, Rao S, Peng Y, Ma Y. Robust Principal Component Analysis: Exact Recovery of Corrupted Low-Rank Matrices via Convex Optimization. In: Bengio Y, Schuurmans D, Lafferty JD, Williams CKI, Culotta A, editors. *Advances in Neural Information Processing Systems 22*. Curran Associates, Inc.; 2009. p. 2080–2088.
43. Zou H, Hastie T, Tibshirani R. Sparse Principal Component Analysis. *Journal of Computational and Graphical Statistics*. 2006; 15(2):265–286. <https://doi.org/10.1198/106186006X113430>
44. Deshpande Y, Montanari A. Information-theoretically optimal sparse PCA. In: 2014 IEEE Int. Symp. Inf. Theory. IEEE; 2014. p. 2197–2201. <https://doi.org/10.1109/ISIT.2014.6875223>
45. Lesieur T, Krzakala F, Zdeborová L. Phase transitions in sparse PCA. In: *IEEE International Symposium on Information Theory—Proceedings*; 2015. p. 1635–1639.
46. Matsushita R, Tanaka T. Low-rank matrix reconstruction and clustering via approximate message passing. In: *Advances in Neural Information Processing Systems*; 2013. p. 917–925.
47. Fortunato S. Community detection in graphs. *Physics Reports*. 2010; 486(3-5):75–174. <https://doi.org/10.1016/j.physrep.2009.11.002>
48. Decelle A, Krzakala F, Moore C, Zdeborová L. Inference and Phase Transitions in the Detection of Modules in Sparse Networks. *Physical Review Letters*. 2011; 107(6):065701. <https://doi.org/10.1103/PhysRevLett.107.065701> PMID: 21902340
49. MacKay DJ. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press; 2003.
50. Donoho DL, Maleki A, Montanari A. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*. 2009; 106(45):18914–18919. <https://doi.org/10.1073/pnas.0909892106> PMID: 19858495

51. Pearl J. Fusion, propagation, and structuring in belief networks. *Artificial Intelligence*. 1986; 29(3):241–288. [https://doi.org/10.1016/0004-3702\(86\)90072-X](https://doi.org/10.1016/0004-3702(86)90072-X)
52. Lesieur T, Krzakala F, Zdeborová L. Constrained low-rank matrix estimation: phase transitions, approximate message passing and applications. *J Stat Mech Theory Exp*. 2017; 2017(7):073403. <https://doi.org/10.1088/1742-5468/aa7284>
53. Fletcher AK, Rangan S. Iterative Reconstruction of Rank-One Matrices in Noise. *Inf Inference A J IMA*. 2012;.
54. Matsushita R, Tanaka T. Low-rank matrix reconstruction and clustering via approximate message passing. In: Burges CJ, Bottou L, Welling M, Ghahramani Z, Weinberger KQ, editors. *Advances in Neural Information Processing Systems*. vol. 26. Curran Associates, Inc.; 2013. p. 1–8. Available from: <https://proceedings.neurips.cc/paper/2013/file/5b69b9cb83065d403869739ae7f0995e-Paper.pdf>.
55. Bolthausen E. An Iterative Construction of Solutions of the TAP Equations for the Sherrington–Kirkpatrick Model. *Commun Math Phys*. 2014; 325(1):333–366. <https://doi.org/10.1007/s00220-013-1862-3>
56. Mézard M, Parisi G, Virasoro MA. SK Model: The Replica Solution without Replicas. *Europhys Lett*. 1986; 1(2):77–82. <https://doi.org/10.1209/0295-5075/1/2/006>
57. Mézard M, Parisi G, Virasoro MA. *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*. vol. 9. World Scientific Publishing Company; 1987.
58. Zdeborová L, Krzakala F. Statistical physics of inference: thresholds and algorithms. *Adv Phys*. 2016; 65(5):453–552. <https://doi.org/10.1080/00018732.2016.1211393>
59. Lesieur T, Krzakala F, Zdeborová L. MMSE of probabilistic low-rank matrix estimation: Universality with respect to the output channel. In: 2015 53rd Annu. Allert. Conf. Commun. Control. Comput. IEEE; 2015. p. 680–687. <https://doi.org/10.1109/ALLERTON.2015.7447070>
60. Cover TM, Thomas JA. *Elements of Information Theory*. John Wiley & Sons; 2006.
61. Gammaitoni L, Hänggi P, Jung P, Marchesoni F. Stochastic resonance. *Reviews of modern physics*. 1998; 70(1):223. <https://doi.org/10.1103/RevModPhys.70.223>
62. Krzakala F, Moore C, Mossel E, Neeman J, Sly A, Zdeborová L, et al. Spectral redemption in clustering sparse networks. *Proceedings of the National Academy of Sciences*. 2013; 110(52):20935–20940. <https://doi.org/10.1073/pnas.1312486110> PMID: 24277835
63. Blumenfeld B. An algorithm for the analysis of temporally structured multidimensional measurements. *Frontiers in Computational Neuroscience*. 2010; 3:28. <https://doi.org/10.3389/neuro.10.028.2009> PMID: 20179787
64. Bandeira AS, Perry A, Wein AS. Notes on computational-to-statistical gaps: predictions using statistical physics. *arXiv preprint arXiv:180311132*. 2018;.
65. Tsodyks MV. Associative Memory in Asymmetric Diluted Network with Low Level of Activity. *Europhys Lett*. 1988; 7(3):203–208. <https://doi.org/10.1209/0295-5075/7/3/003>
66. Lee JS, Briguglio JJ, Cohen JD, Romani S, Lee AK. The Statistical Structure of the Hippocampal Code for Space as a Function of Time, Context, and Value. *Cell*. 2020; 183:620–635. <https://doi.org/10.1016/j.cell.2020.09.024> PMID: 33035454
67. Waydo S, Kraskov A, Quiñero R, Fried I, Koch C. Sparse representation in the human medial temporal lobe. *J Neurosci*. 2006; 26:10232–10234. <https://doi.org/10.1523/JNEUROSCI.2101-06.2006> PMID: 17021178
68. Barber D. *Bayesian Reasoning and Machine Learning*. Cambridge University Press; 2012.
69. Manoel A, Krzakala F, Tramel EW, Zdeborova L. Streaming Bayesian inference: Theoretical limits and mini-batch approximate message-passing. In: 2017 55th Annu. Allert. Conf. Commun. Control. Comput. IEEE; 2017. p. 1048–1055. <https://doi.org/10.1109/ALLERTON.2017.8262853>
70. Dempster AP, Laird NM, Rubin DB. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*. 1977; 39(1):1–22.
71. Krzakala F, Mézard M, Sausset F, Sun Y, Zdeborová L. Probabilistic reconstruction in compressed sensing: algorithms, phase diagrams, and threshold achieving matrices. *Journal of Statistical Mechanics: Theory and Experiment*. 2012; 2012(08):P08009. <https://doi.org/10.1088/1742-5468/2012/08/P08009>
72. Montanari A, Ruan F, Yan J. Adapting to unknown noise distribution in matrix denoising. *arXiv preprint arXiv:181002954*. 2018;.
73. Antenucci F, Franz S, Urbani P, Zdeborová L. Glassy nature of the hard phase in inference problems. *Physical Review X*. 2019; 9(1):011020. <https://doi.org/10.1103/PhysRevX.9.011020>

74. Antenucci F, Krzakala F, Urbani P, Zdeborová L. Approximate survey propagation for statistical inference. *Journal of Statistical Mechanics: Theory and Experiment*. 2019; 2019(2):023401. <https://doi.org/10.1088/1742-5468/aafa7d>
75. Lucibello C, Saglietti L, Lu Y. Generalized Approximate Survey Propagation for High-Dimensional Estimation. In: Chaudhuri K, Salakhutdinov R, editors. *Proceedings of the 36th International Conference on Machine Learning*. vol. 97 of *Proceedings of Machine Learning Research*. PMLR; 2019. p. 4173–4182. Available from: <https://proceedings.mlr.press/v97/lucibello19a.html>.
76. Mallat SG, Zhang Z. Matching pursuits with time-frequency dictionaries. *IEEE Transactions on signal processing*. 1993; 41(12):3397–3415. <https://doi.org/10.1109/78.258082>
77. Elad M, Aharon M. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image processing*. 2006; 15(12):3736–3745. <https://doi.org/10.1109/TIP.2006.881969> PMID: 17153947
78. Mairal J, Bach F, Ponce J, Sapiro G. Online learning for matrix factorization and sparse coding. *Journal of Machine Learning Research*. 2010; 11(1).
79. Barbier J, Macris N. Statistical limits of dictionary learning: random matrix theory and the spectral replica method. *arXiv preprint arXiv:210906610*. 2021;.
80. Maillard A, Krzakala F, Mézard M, Zdeborová L. Perturbative construction of mean-field equations in extensive-rank matrix factorization and denoising. *arXiv preprint arXiv:211008775*. 2021;.
81. Horn RA, Johnson CR. *Matrix analysis*. Cambridge University Press; 2013.
82. Matsuzaki M, Ellis-Davies GC, Nemoto T, Miyashita Y, Iino M, Kasai H. Dendritic spine geometry is critical for AMPA receptor expression in hippocampal CA1 pyramidal neurons. *Nature neuroscience*. 2001; 4(11):1086–1092. <https://doi.org/10.1038/nn736> PMID: 11687814
83. Noguchi J, Nagaoka A, Watanabe S, Ellis-Davies GC, Kitamura K, Kano M, et al. In vivo two-photon uncaging of glutamate revealing the structure–function relationships of dendritic spines in the neocortex of adult mice. *The Journal of physiology*. 2011; 589(10):2447–2457. <https://doi.org/10.1113/jphysiol.2011.207100> PMID: 21486811
84. Holler S, Köstinger G, Martin KA, Schuhknecht GF, Stratford KJ. Structure and function of a neocortical synapse. *Nature*. 2021; 591(7848):111–116. <https://doi.org/10.1038/s41586-020-03134-2> PMID: 33442056
85. Markram H, Lubke J, Frotscher M, Roth A, Sakmann B. Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex. *J Physiol (London)*. 1997; 500:409–440. <https://doi.org/10.1113/jphysiol.1997.sp022031> PMID: 9147328
86. Song S, Sjostrom PJ, Reigl M, Nelson S, Chklovskii DB. Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol*. 2005; 3:e68. <https://doi.org/10.1371/journal.pbio.0030068> PMID: 15737062
87. Wang Y, Markram H, Goodman PH, Berger TK, Ma J, Goldman-Rakic PS. Heterogeneity in the pyramidal network of the medial prefrontal cortex. *Nat Neurosci*. 2006; 9:534–542. <https://doi.org/10.1038/nn1670> PMID: 16547512
88. Guzman SJ, Schlogl A, Frotscher M, Jonas P. Synaptic mechanisms of pattern completion in the hippocampal CA3 network. *Science*. 2016; 353:1117–1123. <https://doi.org/10.1126/science.aaf1836> PMID: 27609885
89. Lefort S, Tomm C, Floyd Sarria JC, Petersen CC. The excitatory neuronal network of the C2 barrel column in mouse primary somatosensory cortex. *Neuron*. 2009; 61:301–316. <https://doi.org/10.1016/j.neuron.2008.12.020> PMID: 19186171
90. Mishra RK, Kim S, Guzman SJ, Jonas P. Symmetric spike timing-dependent plasticity at CA3-CA3 synapses optimizes storage and recall in autoassociative networks. *Nat Commun*. 2016; 7:11552. <https://doi.org/10.1038/ncomms11552> PMID: 27174042
91. Graupner M, Wallisch P, Ostojic S. Natural Firing Patterns Imply Low Sensitivity of Synaptic Plasticity to Spike Timing Compared with Firing Rate. *J Neurosci*. 2016; 36:11238–11258. <https://doi.org/10.1523/JNEUROSCI.0104-16.2016> PMID: 27807166
92. Inagaki HK, Fontolan L, Romani S, Svoboda K. Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature*. 2019; 566:212–217. <https://doi.org/10.1038/s41586-019-0919-7> PMID: 30728503
93. Fuster JM, Alexander G. Neuron activity related to short-term memory. *Science*. 1971; 173:652–654. <https://doi.org/10.1126/science.173.3997.652> PMID: 4998337
94. Miyashita Y, Chang HS. Neuronal correlate of pictorial short-term memory in the primate temporal cortex. *Nature*. 1988; 331:68–70. <https://doi.org/10.1038/331068a0> PMID: 3340148

95. Funahashi S, Bruce CJ, Goldman-Rakic PS. Mnemonic coding of visual space in the monkey's dorso-lateral prefrontal cortex. *J Neurophysiol.* 1989; 61:331–349. <https://doi.org/10.1152/jn.1989.61.2.331> PMID: 2918358
96. Romo R, Brody CD, Hernández A, Lemus L. Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature.* 1999; 399:470–474. <https://doi.org/10.1038/20939> PMID: 10365959
97. Kleinfeld D, Sompolinsky H. Associative neural network model for the generation of temporal patterns. Theory and application to central pattern generators. *Biophys J.* 1988; 54:1039–1051. [https://doi.org/10.1016/S0006-3495\(88\)83041-8](https://doi.org/10.1016/S0006-3495(88)83041-8) PMID: 3233265
98. Fiete IR, Senn W, Wang CZ, Hahnloser RH. Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity. *Neuron.* 2010; 65:563–576. <https://doi.org/10.1016/j.neuron.2010.02.003> PMID: 20188660
99. Gillett M, Pereira U, Brunel N. Characteristics of sequential activity in networks with temporally asymmetric Hebbian learning. *Proc Natl Acad Sci U S A.* 2020; 117:29948–29958. <https://doi.org/10.1073/pnas.1918674117> PMID: 33177232
100. Fino E, Yuste R. Dense inhibitory connectivity in neocortex. *Neuron.* 2011; 69:1188–1203. <https://doi.org/10.1016/j.neuron.2011.02.025> PMID: 21435562
101. Rolotti SV, Ahmed MS, Szoboszlai M, Geiller T, Negrean A, Blockus H, et al. Local feedback inhibition tightly controls rapid formation of hippocampal place fields. *Neuron.* 2022; 110:783–794. <https://doi.org/10.1016/j.neuron.2021.12.003> PMID: 34990571
102. Froemke RC. Plasticity of cortical excitatory-inhibitory balance. *Annu Rev Neurosci.* 2015; p. 195–219. <https://doi.org/10.1146/annurev-neuro-071714-034002> PMID: 25897875
103. Mézard M, Nadal JP, Toulouse G. Solvable models of working memories. *J Physique.* 1986; 47:1457–. <https://doi.org/10.1051/jphys:019860047090145700>
104. Parisi G. A memory which forgets. *J Phys A: Math Gen.* 1986; 19:L617. <https://doi.org/10.1088/0305-4470/19/10/011>
105. Amit DJ, Fusi S. Dynamic learning in neural networks with material synapses. *Neural Computation.* 1994; 6:957–982. <https://doi.org/10.1162/neco.1994.6.5.957>
106. Lahiri S, Ganguli S. A memory frontier for complex synapses. In: Burges CJC, Bottou L, Welling M, Ghahramani Z, Weinberger KQ, editors. *Advances in Neural Information Processing Systems.* vol. 26. Curran Associates, Inc.; 2013. p. 1034–1042.
107. Loewenstein Y, Kuras A, Rumpel S. Multiplicative dynamics underlie the emergence of the log-normal distribution of spine sizes in the neocortex in vivo. *J Neurosci.* 2011; 31:9481–9488. <https://doi.org/10.1523/JNEUROSCI.6130-10.2011> PMID: 21715613
108. Petersen CC, Malenka RC, Nicoll RA, Hopfield JJ. All-or-none potentiation at CA3-CA1 synapses. *Proc Natl Acad Sci USA.* 1998; 95:4732–4737. <https://doi.org/10.1073/pnas.95.8.4732> PMID: 9539807
109. O'Connor DH, Wittenberg GM, Wang SSH. Graded bidirectional synaptic plasticity is composed of switch-like unitary events. *Proc Natl Acad Sci U S A.* 2005; 102:9679–9684. <https://doi.org/10.1073/pnas.0502332102> PMID: 15983385
110. Holland PW, Laskey KB, Leinhardt S. Stochastic blockmodels: First steps. *Social Networks.* 1983; 5(2):109–137. [https://doi.org/10.1016/0378-8733\(83\)90021-7](https://doi.org/10.1016/0378-8733(83)90021-7)
111. Airoldi EM, Blei D, Fienberg S, Xing E. Mixed Membership Stochastic Blockmodels. In: Koller D, Schuurmans D, Bengio Y, Bottou L, editors. *Advances in Neural Information Processing Systems.* vol. 21. Curran Associates, Inc.; 2008. p. 1–8. Available from: <https://proceedings.neurips.cc/paper/2008/file/8613985ec49eb8f757ae6439e879bb2a-Paper.pdf>.
112. Decelle A, Krzakala F, Moore C, Zdeborová L. Inference and Phase Transitions in the Detection of Modules in Sparse Networks. *Physical Review Letters.* 2011; 107(6):065701. <https://doi.org/10.1103/PhysRevLett.107.065701> PMID: 21902340
113. Lim S, McKee JL, Woloszyn L, Amit Y, Freedman DJ, Sheinberg DL, et al. Inferring learning rules from distributions of firing rates in cortical neurons. *Nat Neurosci.* 2015; 18:1804–1810. <https://doi.org/10.1038/nn.4158> PMID: 26523643
114. Thouless DJ, Anderson PW, Palmer RG. Solution of 'Solvable model of a spin glass'. *Philosophical Magazine.* 1977; 35(3):593–601. <https://doi.org/10.1080/14786437708235992>
115. Mézard M, Montanari A. *Information, Physics and Computation.* Cambridge University Press; 2009.