



Robust discovery of temporal structure in multi-neuron recordings using Hopfield networks

Christopher Hillar¹ and Felix Effenberger²

¹ Redwood Center for Theoretical Neuroscience, University of California, Berkeley
Berkeley, CA 94720, USA
chillar@msri.org

² Max-Planck-Institute for Mathematics in the Sciences, Inselstr. 22
04103 Leipzig, GERMANY
felix.effenberger@mis.mpg.de

Abstract

We present here a novel method for the classical task of extracting reoccurring spatiotemporal patterns from spiking activity of large populations of neurons. In contrast to previous studies that mainly focus on synchrony detection or exactly recurring binary patterns, we perform the search in an approximate way that clusters together nearby, noisy network states in the data. Our approach is to use minimum probability flow (MPF) parameter estimation to deterministically fit very large Hopfield networks on windowed spike trains obtained from recordings of spontaneous activity of neurons in cat visual cortex. Examining the structure of the network memories over the spiking activity after training, we find that the networks robustly discover long-range temporal correlations. Specifically, the recurrent network dynamics denoise and group together windowed spike patterns, revealing underlying structure such as cell assemblies. We first demonstrate this by computing an analogy to spike triggered averages that we call memory triggered averages (MTAs). MTAs are obtained by averaging raw spike train windows that converge under the network dynamics to the same memory. The MTAs reveal prominent repeating patterns in the data that are difficult to detect with standard methods such as PCA. Additionally, when memories are collected over eight disjoint epochs in 280 seconds of windowed spiking activity from 50 neurons, their counts are nearly identical and the networks store significantly more memories than would be possible if trained on random patterns.

Keywords: neuronal population activity, parallel spike train analysis, spatiotemporal patterns, Hopfield networks, maximum entropy model, Lenz-Ising model

1 Introduction

Hopfield networks [25, 15] are a well-known model of memory and collective processing in networks of abstract McCulloch-Pitts [23] neurons, but they have rarely been used in data

science applications (although see [22]) because they have small generic pattern capacity (scaling linearly in the number of neurons) and are difficult to train, especially on noisy data. Recently, however, it has been demonstrated that super-linear storage in Hopfield [14] and Hopfield-like [11, 20, 17, 8] networks is possible for special classes of patterns and network connectivity. Moreover, advances in parameter estimation [33] allow the fitting of large models [13].

Additionally, several studies in neuroscience [30, 32, 10, 18] have shown the importance and ubiquity of the underlying discrete probabilistic model in the Hopfield network: the Lenz-Ising model of statistical physics [16]. Also, so-called “deep network” architectures, which have similar underlying maximum entropy models of data, have made a resurgence in the fields of machine learning [2] and, more recently, computer vision and image modeling [28, 19].

Here, we extend the maximum entropy approach by fitting high-dimensional Hopfield network models to electrophysiological recordings of windowed spontaneous activity in cat primary visual cortex area 18. In contrast to the majority of previously published studies [12, 27, 9, 26, 21], our method is not limited to detection of synchrony or exactly recurring patterns of neural activity, but is able to find recurring patterns in an approximate way. Before explaining our precise setup (in Section 3) and findings in detail (in Section 4), we first give theoretical background on Hopfield networks and outline the recent discoveries of efficient estimation and super-linear capacity that underlie the feasibility of our approach.

2 Background

Lenz-Ising model. Our starting point is the non-ferromagnetic *Lenz-Ising model* from statistical physics [16], more generally called a *Markov random field* in the machine learning literature, and the underlying probability distribution of a fully observable *Boltzmann machine* [1]. This discrete probability distribution has as states all length n column binary vectors, with the probability of a particular state $\mathbf{x} = (x_1, \dots, x_n) \in \{0, 1\}^n$ being

$$p_{\mathbf{x}} = \frac{1}{Z} \exp \left(\sum_{i < j} J_{ij} x_i x_j - \sum_i \theta_i x_i \right) = \frac{1}{Z} \exp(-E_{\mathbf{x}}), \quad (1)$$

in which $\mathbf{J} \in \mathbb{R}^{n \times n}$ is a real symmetric matrix with zero diagonal (the *weight matrix*), the column vector $\theta \in \mathbb{R}^n$ is a bias or *threshold* term, and $Z = \sum_{\mathbf{x}} \exp(-E_{\mathbf{x}})$ is the *partition function* (which normalizes $\mathbf{p} = (p_{\mathbf{x}})_{\mathbf{x} \in \{0,1\}^n}$ to sum to 1). Typically, the expression inside the exponential of (1) is viewed as the negative of a quadratic energy function,

$$E_{\mathbf{x}} = -\frac{1}{2} \mathbf{x}^{\top} \mathbf{J} \mathbf{x} + \theta^{\top} \mathbf{x}. \quad (2)$$

Thus, states \mathbf{x} with low energy (2) appear most often under sampling from (1). It follows from basic theory [6] that the distribution defined by (1) is the maximum entropy distribution on binary vectors given its first and second order statistics (mean and covariance).

Hopfield network. A *Hopfield network* is a recurrent network of binary nodes (representing spiking neurons) with deterministic dynamics that act to minimize an energy given by Eq. (2). Formally, the network on n nodes $\{1, \dots, n\}$ consists of a symmetric weight matrix $\mathbf{J} \in \mathbb{R}^{n \times n}$ and a threshold vector $\theta \in \mathbb{R}^n$. The possible *states* of the network are the same as those of the Lenz-Ising model.

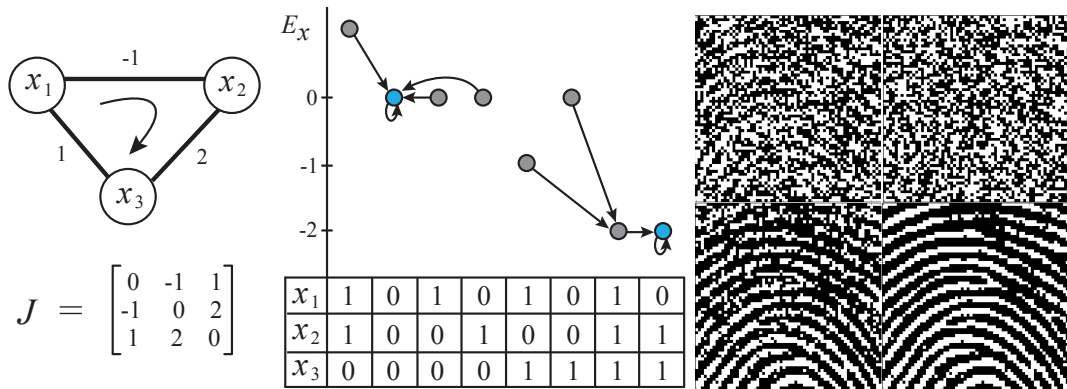


Figure 1: **Small Hopfield network / noisy learning ability.** **Left:** A 3-node Hopfield network with weight matrix \mathbf{J} and zero threshold vector θ . A state vector $\mathbf{x} = (x_1, x_2, x_3)$ has energy $E_{\mathbf{x}}$ as labeled on the y -axis of the diagram. Arrows represent one iteration of the network dynamics; i.e. x_1 , x_2 , and x_3 are updated by Eq. (3) in the order of the clockwise arrow. Resulting memories / fixed-points \mathbf{x}^* are indicated by blue circles. **Right:** Example of noisy learning in a network trained on many corrupted versions of eighty 4096-bit (64×64) binary fingerprints [13]. Clockwise from top left: One sample of 30% corrupted fingerprint shown during learning, novel 40% corrupted fingerprint shown to network after training, result of one iteration of dynamics on novel pattern, and converged memory identical to original fingerprint.

A *dynamics update* of state \mathbf{x} in a Hopfield network consists of replacing each x_i in \mathbf{x} with:

$$x_i = \begin{cases} 1 & \text{if } \sum_{j \neq i} J_{ij} x_j > \theta_i, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

The update given by Eq. (3) is inspired by computations exhibited in neurons [23, 29]; and a model neuron with such an update rule is called a linear threshold *McCulloch-Pitts neuron*.

A fundamental property of Hopfield networks is that asynchronous dynamics updates, Eq. (3), do not increase energy $E_{\mathbf{x}}$. Thus, after a finite number of updates, each initial state \mathbf{x} converges to a *fixed-point* \mathbf{x}^* (also called *memory*) of the dynamics. Intuitively, the dynamics is an inference technique, producing the most probable nearby memory given a noisy version.

A basic problem is to construct Hopfield networks with a given set X of binary patterns as memories (i.e. local minima of $E_{\mathbf{x}}$). Such networks are useful for memory denoising since corrupted versions of patterns in X converge through the dynamics to the originals (e.g. Fig. 1).

Minimum probability flow learning. To estimate the parameters of the Hopfield network, we use the recently developed *minimum probability flow* (MPF) technique [33, 13] for fitting parameterized distributions that avoids computation of the partition function Z .

Applied to estimating the parameters in an Lenz-Ising model / Hopfield network, Eq. (1), the minimum probability flow objective function [33, 13] is:

$$K_X(\mathbf{J}, \theta) = \sum_{\mathbf{x} \in X} \sum_{\mathbf{x}' \in \mathcal{N}(\mathbf{x})} \exp\left(\frac{E_{\mathbf{x}} - E_{\mathbf{x}'}}{2}\right). \quad (4)$$

Here, the *neighborhood* $\mathcal{N}(\mathbf{x})$ of \mathbf{x} consists of those binary vectors which are Hamming distance 1 away from \mathbf{x} (i.e. those \mathbf{x}' with exactly one bit different from \mathbf{x}).

The function in (4) is infinitely differentiable, jointly convex in the parameters, consists of only order $O(|X|n)$ terms, and can be minimized using standard gradient-based methods. Here, we use the limited-memory Broyden-Fletcher-Goldfarb-Shanno algorithm [24] to minimize (4).

Notice also that when K_X is small, the energy differences $E_{\mathbf{x}} - E_{\mathbf{x}'}$ between points \mathbf{x} in the dataset X and patterns \mathbf{x}' in single-flip neighborhoods $\mathcal{N}(\mathbf{x})$ will be negative, making \mathbf{x} a fixed-point of the Hopfield dynamics. Importantly for applications, much more is true: minimizing (4) given a storable set of patterns X will determine a Hopfield network storing those patterns as robust memories [13]. Moreover, the MPF objective function can naturally be turned into an online, neurologically plausible learning rule [13, 14].

It was found in [13] that when compared with classical techniques for Hopfield pattern storage, minimizing the MPF objective function provides superior efficiency and generalization. Another surprising discovery was that MPF learning can store patterns, unsupervised, from highly corrupted or noisy versions. This finding is illustrated visually in Fig. 1 for a “noisy fingerprints” example and in Fig. 3 for revealing patterns in a neural recording of spontaneous spiking activity from area 18 in cat primary visual cortex.

Super-linear capacity in Hopfield networks. Hopfield showed experimentally that $.15n$ binary patterns (generated uniformly at random) can be stored in an n -node network if some fixed percentage of errors in a recovered pattern were tolerated [15]. Since then, improved methods to fit Hopfield networks have been developed; e.g., [4]. Independent of the method, however, arguments of Cover [7] can be used to show that the number of generic (or “randomly generated”) patterns robustly storable in a Hopfield network with n nodes is at most $2n$, although the exact value is presently unknown (it is $< 1.7n$ from experiments in [13]). Here, “robustly stored” means that the dynamics is capable of recovering the pattern even if it is corrupted by a fixed, positive fraction of its bits.

Nonetheless, it is still possible for certain Hopfield networks to have exponentially large collections of highly structured (“non-random”) memories. In fact, theoretical and experimental evidence suggest that Hopfield networks usually have exponentially many memories (fixed-points of the dynamics). For instance, drawing weight strengths for the model randomly from a normal distribution makes n -node networks with $\approx 1.22^n$ fixed-points asymptotically [34].

Although a generic Hopfield network has exponential capacity, its basins of attraction are shallow and difficult to predetermine from network weights, leading many researchers to speculate that such *spurious minima* are to be avoided. The theorem below [14] shows however that special connectivity structures can create networks with robust memories in a large number of useful combinatorial configurations (e.g. cliques in graphs).

Theorem 1. There are n -node Hopfield networks that robustly store $\frac{2\sqrt{2n} + \frac{1}{4}}{n^{1/4}\sqrt{\pi}}$ memories.

One can also interpret the robustness aspect of Theorem 1 as saying that these networks have large “basins of attraction” around an exponential number of memories.

Given that such networks exist, it is natural to ask whether there are non-artificial signal domains which could be analyzed using Hopfield networks and MPF parameter estimation. In fact, as demonstrated in our results section, we find super-linear capacity of networks trained on windowed spontaneous multi-electrode activity in cat primary visual cortex area 18.

3 Understanding the activity of large collections of neurons using Hopfield networks

Here, we hope to shed light on the activity of spike trains in cat primary visual cortex by characterizing the windowed spontaneous activity patterns of spike trains from 50 neurons.

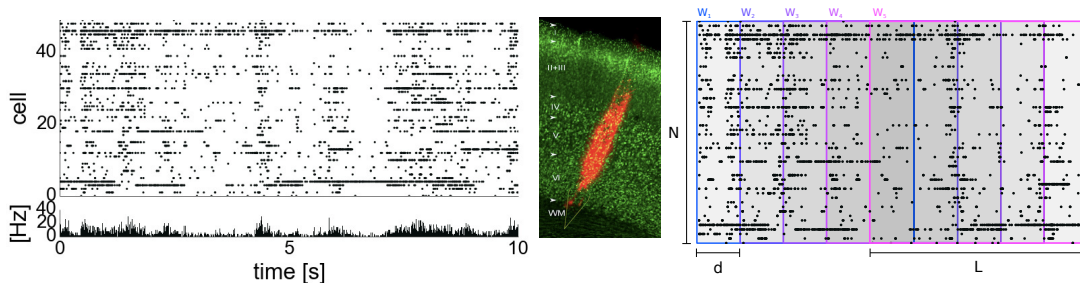


Figure 2: **Spontaneous neural activity** recorded from cat primary visual cortex. **Left:** Raster plot and PSTH showing 10 seconds of spiking activity from 50 cells (Neural data were recorded by Tim Blanche in the laboratory of Nicholas Swindale, University of British Columbia, and downloaded from the NSF-funded CRCNS Data Sharing website <http://crcns.org/>) **Middle:** Histological reconstruction of polytrode track through cortical layers (courtesy of [3]). **Right:** Windowing of binned neural activity. To generate training data for the Hopfield network, spiking data of N cells are first binned and then training vectors are extracted using a sliding window. Windows of length L are shifted by d bins (here, we take $d = 1$), resulting in training vectors of dimension $n = NL$. The above illustration shows five overlapping windows.

Such a study, however, requires the estimation of Hopfield network parameters for dauntingly high-dimensional binary patterns (the largest network fit in this work has $n = 1350$ nodes).

Instead of characterizing single time-points of neural activity as done in many previous studies, here we examine *windowed* versions of spike trains. In a first preprocessing step, the raw spiking data of N neurons is discretized by binning it into bins of 2 milliseconds. We then use a sliding window of length L to extract portions of the raw binned spiking data. Starting at the first bin of data, the window is advanced by 1 bin in each step, subsequently extracting the next L bins of raw data until all data are collected. This results in a large set of binary vectors of length $n = NL$ that we use as training data for the Hopfield networks (see Fig. 2).

An important step and classical task in the exploratory analysis of large-scale neural datasets is the discovery of repeated activity patterns, so called *spatiotemporal patterns* [12]. Since neural spike train data is binary and typically noisy, standard methods such as principal (PCA) and independent (ICA) component analyses are not applicable, and previously published methods in this field either focus solely on synchrony detection (e.g. [27, 26, 21]) or assume a more or less noiseless scenario, seeking to classify exactly recurring STP in neuronal activity (apart from allowing some jitter, see e.g. [9]). Fitting Hopfield networks on windowed spike train data, we obtain networks that encode denoised versions of salient temporal spike patterns present in the data as memories. More specifically, the networks tend to group highly similar but distinct patterns not obvious in the raw data, removing many patterns occurring infrequently and most likely due to noise in the collection of data (see Fig. 3).

We aim to use this type of spike-pattern clustering to discover relationships between stimuli and neural recordings in this and other preparations. For instance, after learning a network and converging dynamics, one can compute a Hopfield *memory triggered average* (MTA; see Fig. 3) and relate such features of neural population activity to stimuli [35]. To compute MTAs, all spike train windows that converge to the same memory under the Hopfield network dynamics are collected and averaged.

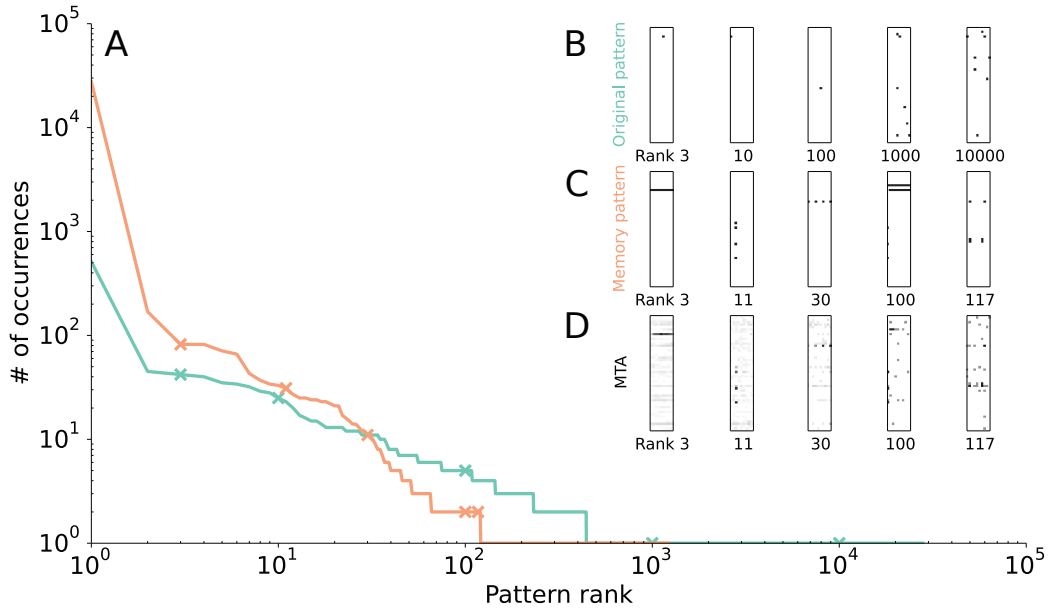


Figure 3: **Clustering multi-unit electrode recordings with Hopfield networks.** **A:** After learning an $n = 500$ -bit Hopfield network and converging dynamics (on 30,000 contiguous 2 millisecond bins of noisy spike train data recorded from 50 neurons, $L = 10$, $N = 50$; see Fig. 2), a large number of states in the data that occur infrequently are summarized by a small number of patterns in the network. **B:** Due to noise, frequently occurring patterns in the raw data do not well-capture structure. **C:** Binary memories of the Hopfield network encode denoised versions of prominent patterns present in the data. **D:** By computing *memory triggered averages* (MTAs), i.e. by averaging all raw data windows that lead to a specific fixed-point under the dynamics, we obtain a clean version of the most salient patterns in the data, e.g. revealing cell assemblies in the MTAs of rank 11 and 100.

4 Main results

Here we make a number of observations from fitting large Hopfield networks on windowed spike trains with MPF. A first interesting observation is that the number of distinct spiking patterns occurring in windows of length L in the raw data is vast and close to the maximal number possible (Fig. 4, left). This is due to the noisy nature of the data in which many patterns occur only once (see also Fig. 3) and patterns differing by a few spikes are not identified as the same.

A Hopfield network trained on these raw patterns of data forms memories of pattern clusters that have large similarity (characterized by small number of bit flips needed to obtain one from the other). Consequently, the number of converged patterns (i.e. the memories of the Hopfield network) is much lower than the number of observed patterns (see Fig. 3A). This can be explained by the network performing a denoising of the data, grouping similar patterns together and reducing the large number of patterns occurring rarely.

Although significantly smaller than the high number of distinct patterns in the raw data, the number of memories of the Hopfield network trained on these patterns is still considerably larger than the maximal number ($< 1.7NL$) of memories of a Hopfield network trained on random




Figure 4: **Number of memories in Hopfield networks trained on spontaneous cortical activity is large, and their entropies scale linearly in window size.** **Left:** Number of patterns occurring in data recorded from 280 seconds of activity from $N = 50$ cortical cells versus window length L (spike trains are produced from 2 millisecond binning). Empirical patterns are obtained by sampling windows of length L from raw spiking data chunked into eight disjoint epochs of 17,500 bins corresponding to 35 seconds of recording. Converged patterns are the memories stored in the Hopfield network that correspond to iterating recurrent dynamics on raw empirical patterns until convergence. Dashed lines show the eight disjoint epochs, with thick lines showing means. The maximal number of memories in a Hopfield network with $n = NL$ nodes when trained on random data is at most $1.7n$ and shown in orange. **Right:** Entropy of the sets of empirical and converged patterns versus window length L for data epochs of length 17,500. Dashed lines represent each epoch, thick lines show means, dots represent least-squares linear fit (r -value 0.9998, p -value $5.4486 \cdot 10^{-8}$, $\sigma = 0.0048$). The maximal possible entropy ($= \log_2 17500$) shown as a dark grey line.

(Bernoulli) data. Importantly, this demonstrates that the spiking behavior of the neurons is not well-modeled by independent Poisson processes.

Another observation (see Fig. 4, right) is that the Shannon entropy [31] calculated for each 17,500 binned time epoch (35,000 milliseconds) of raw spiking data of $N = 50$ cells is close to its maximal value even for small window sizes. The memories of the trained networks, on the other hand, have lower entropy. This is due to a denoising and grouping of similar patterns with the recurrent network dynamics. Strikingly, the entropy encoded is also nearly exactly linear in the window size L .

5 Discussion

While being valuable as an abstract functional model, Hopfield networks have long been dismissed as serious candidates for real-world applications. The main shortcomings that have prevented the practical use of Hopfield networks are threefold: First and foremost, the quantity

of generic robust memories is limited to a (small) number linear in the number n of nodes (e.g. $0.15n$ using the original Hopfield learning rule). Secondly, the usability of Hopfield networks on noisy data was limited, since the original Hopfield learning rule requires supervised learning with clean patterns, while more advanced statistical training methods require computationally intensive iterations. Thirdly, Hopfield networks are thought to be weak models of data because many datasets have higher-order correlations that might not be captured by a second-order model (the Lenz-Ising model is determined by its means and covariances). Nonetheless, as we have demonstrated here, Hopfield networks fit to noisy binary time series data obtained from spontaneous spiking activity in cortical neurons are able to extract interesting features in the data (e.g. Fig 3CD).

The trained Hopfield networks denoise the data, grouping similar patterns together in a way that respects the underlying statistics of the data (in contrast to simple averaging methods). They are thus able to identify prominent patterns reoccurring in the dataset, possibly corrupted by noise, and eliminate the large number of spurious patterns occurring rarely (see Fig. 3). In its memories the network encodes different structural aspects of the spiking data such as prominent temporal firing sequences and cell assemblies [5] that are very difficult if not impossible to identify in the raw data or using standard methods.

Furthermore, the number of learned memories and their entropic structure show that the spiking activity of the cells is far from random (as modeled by a Bernoulli setting), in itself we believe to be a remarkable insight. Finally, we remark that fitting such high-dimensional Hopfield networks in this study was possible on standard consumer computers because of the tractability of the minimum probability flow estimation technique applied to Hopfield networks.

6 Conclusion

We have shown that when Hopfield networks are trained on noisy binary time series obtained from recordings of spontaneous spiking cortical activity using minimum probability flow (MPF), robust temporal structure can be discovered. An important aspect of our techniques is that they are completely deterministic since the underlying parameter estimation algorithm is to minimize a convex function over the data. After training, we find that the number of memories in these networks is remarkably larger than possible with fitting random patterns. Moreover, we find a robust near-exact linear relationship (r -value = .9998) between window size and entropy. We conclude that Hopfield networks could prove a powerful tool for the efficient discovery of salient structure in many types of noisy binary data, including cortical spiking activity.

References

- [1] D.H. Ackley, G.E. Hinton, and T.J. Sejnowski. A learning algorithm for boltzmann machines. *Cognitive science*, 9(1):147–169, 1985.
- [2] Y. Bengio. Learning deep architectures for AI. *Foundations and trends in Machine Learning*, 2(1):1–127, 2009.
- [3] Timothy J Blanche, Martin A Spacek, Jamille F Hetke, and Nicholas V Swindale. Polytrodes: high-density silicon electrode arrays for large-scale multiunit recording. *Journal of Neurophysiology*, 93(5):2987–3000, 2005.
- [4] A.D. Bruce, A. Canning, B. Forrest, E. Gardner, and D.J. Wallace. Learning and memory properties in fully connected networks. In *AIP Conference Proceedings*, volume 151, page 65, 1986.
- [5] György Buzsáki. Neural syntax: cell assemblies, synapsembles, and readers. *Neuron*, 68(3):362–385, 2010.

- [6] Thomas Cover and Joy. Thomas. *Elements of information theory*. Wiley-Interscience, Hoboken, NJ, second edition, 2006.
- [7] T.M. Cover. Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *Electronic Computers, IEEE Transactions on*, (3):326–334, 1965.
- [8] C. Curto, V. Itskov, K. Morrison, Z. Roth, and J.L. Walker. Combinatorial neural codes from a mathematical coding theory perspective. *Neural computation*, 25(7):1891–1925, 2013.
- [9] K. S. Gansel and W. Singer. Detecting multineuronal temporal patterns in parallel spike trains. *Frontiers in neuroinformatics*, 6(May):18, 2012.
- [10] E. Granot-Atedgi, G. Tkačik, R. Segev, and E. Schneidman. Stimulus-dependent maximum entropy models of neural population codes. *PLoS computational biology*, 9(3):e1002922, 2013.
- [11] V. Gripon and C. Berrou. Sparse neural networks with large learning diversity. *Neural Networks, IEEE Transactions on*, 22(7):1087–1096, 2011.
- [12] S. Grün and S. Rotter. *Analysis of parallel spike trains*. Springer, 2010.
- [13] C. Hillar, J. Sohl-Dickstein, and K. Koepsell. Efficient and optimal Little-Hopfield auto-associative memory storage using minimum probability flow. In *4th Neural Information Processing Systems (NIPS) workshop on Discrete Optimization in Machine Learning (DISCML)*, 2012.
- [14] C. Hillar and N. M. Tran. Robust exponential memory in Hopfield networks. *ArXiv e-prints: nlin.AO 1411.4625*, 2014.
- [15] J.J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, 1982.
- [16] E. Ising. Beitrag zur Theorie des Ferromagnetismus. *Zeitschrift für Physik*, 31:253–258, 1925.
- [17] Amin Karbasi, Amir Hesam Salavati, and Amin Shokrollahi. Iterative learning and denoising in convolutional neural associative memories. In *Proceedings of The 30th International Conference on Machine Learning*, pages 445–453, 2013.
- [18] Urs Köster, Jascha Sohl-Dickstein, Charles M Gray, and Bruno A Olshausen. Modeling higher-order correlations within cortical microcolumns. *PLoS computational biology*, 10(7):e1003684, 2014.
- [19] A. Krizhevsky, I. Sutskever, and G. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems 25*, pages 1106–1114, 2012.
- [20] K.R. Kumar, A.H. Salavati, and A. Shokrollahi. Exponential pattern retrieval capacity with non-binary associative memory. In *Information Theory Workshop (ITW), 2011 IEEE*, pages 80–84. IEEE, 2011.
- [21] V. Lopes-dos Santos, S. Ribeiro, and A. B. L. Tort. Detecting cell assemblies in large neuronal populations. *Journal of neuroscience methods*, 220(2):149–66, 2013.
- [22] N. Martine and T. Jean-Bernard. Neural approach for TV image compression using a Hopfield type network. In D.S. Touretzky, editor, *Advances in Neural Information Processing Systems 1*, pages 264–271. 1989.
- [23] W.S. McCulloch and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *Bulletin of mathematical biology*, 5(4):115–133, 1943.
- [24] J. Nocedal. Updating quasi-newton matrices with limited storage. *Mathematics of computation*, 35(151):773–782, 1980.
- [25] LA Pastur and AL Figotin. Exactly soluble model of a spin glass. *Sov. J. Low Temp. Phys.*, 3:378–383, 1977.
- [26] D. Picado-Muñoz, C. Borgelt, D. Berger, G. Gerstein, and S. Grün. Finding neural assemblies with frequent item set mining. *Frontiers in neuroinformatics*, 7(May):9, 2013.
- [27] G. Pipa, D. W. Wheeler, W. Singer, and D. Nikolić. NeuroXidence: reliable and efficient analysis of an excess or deficiency of joint-spike events. *Journal of Computational Neuroscience*, 25(1):64–88, 2008.
- [28] M.A. Ranzato, A. Krizhevsky, and G.E. Hinton. Factored 3-way restricted boltzmann machines

- for modeling natural images. In *International Conference on Artificial Intelligence and Statistics*, pages 621–628, 2010.
- [29] Frank Rosenblatt. *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Spartan Books, Washington, D.C., 1961.
- [30] E. Schneidman, M.J. Berry, R. Segev, and W. Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–12, 2006.
- [31] C. Shannon. A mathematical theory of communication. *Bell Sys. Tech. J.*, 27:379–423, 623–656, 1948.
- [32] J. Shlens, G.D. Field, J.L. Gauthier, M.I. Grivich, D. Petrusca, A. Sher, A.M. Litke, and E.J. Chichilnisky. The structure of multi-neuron firing patterns in primate retina. *J. Neurosci.*, 26(32):8254–66, 2006.
- [33] J. Sohl-Dickstein, P.B. Battaglino, and M.R. DeWeese. New method for parameter estimation in probabilistic models: minimum probability flow. *Physical review letters*, 107(22):220601, 2011.
- [34] F. Tanaka and S.F. Edwards. Analytic theory of the ground state properties of a spin glass. I. ising spin glass. *Journal of Physics F: Metal Physics*, 10:2769, 1980.
- [35] G. Tkacik, O. Marre, D. Amodei, E. Schneidman, W. Bialek, and M.J. Berry II. Searching for collective behavior in a large network of sensory neurons. *PLoS computational biology*, 10(1):e1003408, 2014.