
Efficient and optimal binary Hopfield associative memory storage using minimum probability flow

Christopher Hillar

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

Jascha Sohl-Dickstein

Redwood Center for Theoretical Neuroscience
University of California, Berkeley
Berkeley, CA 94720
jascha.sohldickstein@gmail.com

Kilian Koepsell

Redwood Center for Theoretical Neuroscience
University of California, Berkeley
Berkeley, CA 94720
kilian@berkeley.edu

Abstract

We present an algorithm to store binary memories in a Hopfield neural network using minimum probability flow, a recent technique to fit parameters in energy-based probabilistic models. In the case of memories without noise, our algorithm provably achieves optimal pattern storage (which we show is at least one pattern per neuron) and outperforms classical methods both in speed and memory recovery. Moreover, when trained on noisy or corrupted versions of a fixed set of binary patterns, our algorithm finds networks which correctly store the originals. We also demonstrate this finding visually with the unsupervised storage and clean-up of large binary fingerprint images from significantly corrupted samples.¹

Introduction. In 1982, motivated by neural modeling work of [1] and the Ising spin glass model from statistical physics [2], Hopfield introduced a method for the storage and retrieval of binary patterns in an auto-associative neural-network [3]. Even today, this model and its various extensions [4, 5] provide a plausible mechanism for memory formation in the brain. However, existing techniques for training Hopfield networks suffer either from limited pattern capacity or excessive training time, and they exhibit poor performance when trained on unlabeled, corrupted memories.

Our main theoretical contributions here are the introduction of a tractable and neurally-plausible algorithm for the optimal storage of patterns in a Hopfield network, a proof that the capacity of such a network is at least one pattern per neuron, and a novel local learning rule for training neural networks. Our approach is inspired by minimum probability flow [6], a recent technique for fitting probabilistic models that avoids computations with a partition function, the usually intractable normalization constant of a parameterized probability distribution.

We also present several experimental results. When compared with standard techniques for Hopfield pattern storage, our method is shown to be superior in efficiency and generalization. Another finding is that our algorithm can store many patterns in a Hopfield network from highly corrupted (unlabeled) samples of them. This discovery is also corroborated visually by the storage of 64×64 binary images of human fingerprints from highly corrupted versions, as in Fig. 1.

¹Supported by an NSF All-Institutes Postdoctoral Fellowship administered by the Mathematical Sciences Research Institute through its core grant DMS-0441170 (CH) and by NSF grant IIS-0917342 (KK). This work appeared at the 2012 Neural Information Processing Systems (NIPS) workshop on Discrete Optimization in Machine Learning (DISCML).




Figure 1: **Learning from corrupted samples.** We stored 80 fingerprints (64×64 binary images) in a Hopfield network with $n = 64^2 = 4096$ nodes by minimizing the MPF objective (4) over a large set of randomly generated (and unlabeled) “noisy” versions (each training pattern had a random subset of 1228 of its bits flipped; e.g., a,e). After training, all 80 fingerprints were stored as fixed-points of the network. **a.** Sample training fingerprint with 30% corruption. **b.** Fingerprint with 40% corruption. **c.** State of network after one dynamics update initialized at b. **d.** Converged network dynamics equal to original fingerprint. **e-h.** As in a-d, but for different fingerprint.

Background. A Hopfield network $\mathcal{H} = (\mathbf{J}, \theta)$ on n nodes $\{1, \dots, n\}$ consists of a symmetric weight matrix $\mathbf{J} = \mathbf{J}^\top \in \mathbb{R}^{n \times n}$ with zero diagonal and a threshold vector $\theta = (\theta_1, \dots, \theta_n)^\top \in \mathbb{R}^n$. We do not allow any row \mathbf{J}_i of the matrix \mathbf{J} to be zero. The possible states of the network are all length n binary strings $\{0, 1\}^n$, which we represent as binary column vectors $\mathbf{x} = (x_1, \dots, x_n)^\top$, each $x_i \in \{0, 1\}$ indicating the state x_i of node i . Given any state $\mathbf{x} = (x_1, \dots, x_n)^\top$, an (asynchronous) dynamical update of \mathbf{x} consists of replacing x_i in \mathbf{x} (in consecutive order starting with $i = 1$) with the value

$$x_i = H(\mathbf{J}_i \mathbf{x} - \theta_i). \quad (1)$$

Here, \mathbf{J}_i is the i th row of \mathbf{J} and H is the Heaviside function given by $H(r) = 1$ if $r > 0$ and $H(r) = 0$ if $r \leq 0$.

The energy $E_{\mathbf{x}}$ of a binary pattern \mathbf{x} in a Hopfield network is defined to be

$$E_{\mathbf{x}}(\mathbf{J}, \theta) := -\frac{1}{2} \mathbf{x}^\top \mathbf{J} \mathbf{x} + \theta^\top \mathbf{x} = -\sum_{i < j} x_i x_j J_{ij} + \sum_{i=1}^n \theta_i x_i, \quad (2)$$

identical to the energy function for an Ising spin glass. In fact, the dynamics of a Hopfield network can be seen as 0-temperature Gibbs sampling of this energy function. A fundamental property of Hopfield networks is that asynchronous dynamical updates do not increase the energy (2). Thus, after a finite number of updates, each initial state \mathbf{x} converges to a fixed-point $\mathbf{x}^* = (x_1^*, \dots, x_n^*)^\top$ of the dynamics; that is, $x_i^* = H(\mathbf{J}_i \mathbf{x}^* - \theta_i)$ for each i .

Given a binary pattern \mathbf{x} , 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 with exactly one bit different from \mathbf{x}). We say that \mathbf{x} is a strict local minimum if every $\mathbf{x}' \in \mathcal{N}(\mathbf{x})$ has a strictly larger energy:

$$0 > E_{\mathbf{x}} - E_{\mathbf{x}'} = (\mathbf{J}_i \mathbf{x} - \theta_i) \delta_i, \quad (3)$$

where $\delta_i = 1 - 2x_i$ and x_i is the bit that differs between \mathbf{x} and \mathbf{x}' . It is straightforward to verify that if \mathbf{x} is a strict local minimum, then it is a fixed-point of the dynamics.

A basic problem is to construct Hopfield networks with a given set \mathcal{D} of binary patterns as fixed-points or strict local minima of the energy function (2). Such networks are useful for memory denoising and retrieval since corrupted versions of patterns in \mathcal{D} will converge through the dynamics to the originals. Traditional approaches to this problem consist of iterating over \mathcal{D} a learning rule [7] that updates a network’s weights and thresholds given a training pattern $\mathbf{x} \in \mathcal{D}$. For the purposes of

this work, we call a rule *local* when the learning updates to the three parameters J_{ij} , θ_i , and θ_j can be computed with access solely to x_i, x_j , the feedforward inputs $\mathbf{J}_i \mathbf{x}$, $\mathbf{J}_j \mathbf{x}$, and the thresholds θ_i, θ_j ; otherwise, we call the rule *nonlocal*. The locality of a rule is an important feature in a network training algorithm because of its necessity in theoretical models of computation in neuroscience.

In [3], Hopfield defined an *outer-product learning rule* (OPR) for finding such networks. OPR is a local rule since only the binary states of nodes x_i and x_j are required to update a coupling term J_{ij} during training (and only the state of x_i is required to update θ_i). Using OPR, at most $n/(4 \log n)$ patterns can be stored without errors in an n -node Hopfield network [8, 9]. In particular, the ratio of patterns storable to the number of nodes using this rule is at most $1/(4 \log n)$ memories per neuron, which approaches zero as n increases. If a small percentage of incorrect bits is tolerated, then approximately $0.15n$ patterns can be stored [3, 10].

The *perceptron learning rule* (PER) [11, 12, 7] provides an alternative method to store patterns in a Hopfield network [13]. PER is also a local rule since updating J_{ij} requires only $\mathbf{J}_i \mathbf{x}$ and $\mathbf{J}_j \mathbf{x}$ (and updating θ_i requires $\mathbf{J}_i \mathbf{x}$). Unlike OPR, it achieves optimal storage capacity, in that if it is possible for a collection of patterns \mathcal{D} to be fixed-points of a Hopfield network, then PER will converge to parameters \mathbf{J}, θ for which all of \mathcal{D} are fixed-points. However, training frequently takes many parameter update steps (see Fig. 2), and the resulting Hopfield networks do not generalize well (see Fig. 3) nor store patterns from corrupted samples (see Fig. 4).

Despite the connection to the Ising model energy function, and the common usage of Ising spin glasses (otherwise referred to as Boltzmann machines [5]) to build probabilistic models of binary data, we are aware of no previous work on associative memories that takes advantage of a probabilistic interpretation during training. (Although probabilistic interpretations have been used for pattern recovery [14].)

Theoretical Results. We give an efficient algorithm for storing at least n binary patterns as strict local minima (and thus fixed-points) in an n -node Hopfield network, and we prove that this algorithm achieves the optimal storage capacity achievable in such a network. We also present a novel local learning rule for the training of neural networks.

Consider a collection of m binary n -bit patterns \mathcal{D} to be stored as strict local minima in a Hopfield network. Not all collections of m such patterns \mathcal{D} can so be stored; for instance, from (3) we see that no two binary patterns one bit apart can be stored simultaneously. Nevertheless, we say that the collection \mathcal{D} *can be stored as local minima* of a Hopfield network if there is some $\mathcal{H} = (\mathbf{J}, \theta)$ such that each $\mathbf{x} \in \mathcal{D}$ is a strict local minimum of the energy function $E_{\mathbf{x}}(\mathbf{J}, \theta)$ in (2).

The *minimum probability flow* (MPF) objective function [6] given the collection \mathcal{D} is

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

The function in (4) is infinitely differentiable and (generically, strictly) convex in the parameters. Notice that when $K_{\mathcal{D}}(\mathbf{J}, \theta)$ is small, the energy differences $E_{\mathbf{x}} - E_{\mathbf{x}'}$ between $\mathbf{x} \in \mathcal{D}$ and patterns \mathbf{x}' in neighborhoods $\mathcal{N}(\mathbf{x})$ will satisfy (3), making \mathbf{x} a fixed-point of the dynamics.

As the following result explains, minimizing (4) given a storable set of patterns will determine a Hopfield network storing those patterns.

Theorem 1. *If a set of binary vectors \mathcal{D} can be stored as local minima of a Hopfield network, then minimizing the convex MPF objective (4) will find such a network.*

Our next result is that at least n patterns in an n -node Hopfield network can be stored by minimizing (4) (see also [15, 16]). To make this statement mathematically precise, we introduce some notation. Let $r(m, n) < 1$ be the probability that a collection of m binary patterns chosen uniformly at random from all $\binom{2^n}{m}$ m -element subsets of $\{0, 1\}^n$ can be made local minima of a Hopfield network. The *pattern capacity* (per neuron) of the Hopfield network is defined to be the supremum of all real numbers $a > 0$ such that $\lim_{n \rightarrow \infty} r(an, n) = 1$.

Theorem 2. *The capacity of an n -node Hopfield network is at least 1 pattern per neuron.*

In other words, for any fixed $a < 1$, the fraction of all subsets of $m = an$ patterns that can be made strict local minima (and thus fixed-points) of a Hopfield network with n nodes converges to




Figure 2: **(Left)** Shows fraction of patterns made fixed-points of a Hopfield network using OPR (outer-product rule), MPF (minimum probability flow), and PER (perceptron) as a function of the number of randomly generated training patterns m . Here, $n = 64$ binary nodes and we have averaged over $t = 20$ trials. The slight difference in performance between MPF and PER is due to the extraordinary number of iterations required for PER to achieve perfect storage of patterns near the critical pattern capacity of the Hopfield network. See also Fig. 2. **(Right)** Shows time (on a log scale) to train a Hopfield network with $n = 64$ neurons to store m patterns using OPR, PER, and MPF (averaged over $t = 20$ trials).

1 as n tends to infinity. Moreover, by Theorem 1, such networks can be found by minimizing (4). Experimental evidence suggests that the capacity limit above is 1 for all $a < 1.5$, but converges to 0 for $a > 1.7$ (see Fig. 2). Although the Cover bound [17] forces $a \leq 2$, it is an open problem to determine the exact critical value of a (i.e., the exact pattern capacity of the Hopfield network).

We close this section by defining a new learning rule for a neural network. In words, the *minimum probability flow learning rule* (MPF) takes an input training pattern \mathbf{x} and moves the parameters (\mathbf{J}, θ) a small amount in the direction of steepest descent of the MPF objective function $K_{\mathcal{D}}(\mathbf{J}, \theta)$ with $\mathcal{D} = \{\mathbf{x}\}$. These updates for J_{ij} and θ_i take the form (where again, $\delta = \mathbf{1} - 2\mathbf{x}$):

$$\Delta J_{ij} \propto -\delta_i x_j e^{\frac{1}{2}(\mathbf{J}_i \mathbf{x} - \theta_i) \delta_i} - \delta_j x_i e^{\frac{1}{2}(\mathbf{J}_j \mathbf{x} - \theta_j) \delta_j} \quad (5)$$

$$\Delta \theta_i \propto \delta_i e^{\frac{1}{2}(\mathbf{J}_i \mathbf{x} - \theta_i) \delta_i}. \quad (6)$$

It is clear from (5) and (6) that MPF is a local learning rule.

Experimental results. We performed several experiments comparing standard techniques for fitting Hopfield networks with minimizing the MPF objective function (4). All computations were performed on standard desktop computers, and we used the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm [18] to minimize (4).

In our first experiment, we compared MPF to the two methods OPR and PER for finding 64-node Hopfield networks storing a given set of patterns \mathcal{D} . For each of 20 trials, we used the three techniques to store a randomly generated set of m binary patterns, where m ranged from 1 to 120. The results are displayed in Fig. 2 and support the conclusions of Theorem 1 and Theorem 2.

To evaluate the efficiency of our method relative to standard techniques, we compared training time of a 64-node network as in Fig. 2 with the three techniques OPR, MPF, and PER. The resulting computation times are displayed in Fig. 2 on a logarithmic scale. Notice that computation time for MPF and PER significantly increases near the capacity threshold of the Hopfield network.

For our third experiment, we compared the denoising performance of MPF and PER. For each of four values for m in a 128-node Hopfield network, we determined weights and thresholds for storing all of a set of m randomly generated binary patterns using both MPF and PER. We then flipped 0 to 64 of the bits in the stored patterns and let the dynamics (1) converge (with weights and thresholds given by MPF and PER), recording if the converged pattern was identical to the original pattern or not. Our results are shown in Fig 3, and they demonstrate the superior corrupted memory retrieval performance of MPF.




Figure 3: Shows fraction of exact pattern recovery for a perfectly trained $n = 128$ Hopfield network using rules PER (figure on the left) and MPF (figure on the right) as a function of bit corruption at start of recovery dynamics for various numbers m of patterns to store. We remark that this figure and the next do not include OPR as its performance was far worse than MPF or PER.




Figure 4: Shows fraction of patterns (shown in red for MPF and blue for PER) and fraction of bits (shown in dotted red for MPF and dotted blue for PER) recalled of trained networks (with $n = 64$ nodes each) as a function of the number of patterns m to be stored. Training patterns were presented repeatedly with 20 bit corruption (i.e., 31% of the bits flipped). (averaged over $t = 13$ trials.)

We also tested the how the efficiency of our algorithm scales with the number of nodes n . For varying n , we fit $m = n/4$ (randomly generated) patterns in n -node Hopfield networks using MPF over 50 trials and examined the training time. Our experiments show that the average training time to fit Hopfield networks with MPF is well-modeled by a polynomial $O(n^{5/2})$.

A surprising final finding in our investigation was that MPF can store patterns from highly corrupted or noisy versions on its own and without supervision. This result is explained in Fig 4. To illustrate the experiment visually, we stored $m = 80$ binary fingerprints in a 4096-node Hopfield network using a large set of training samples which were corrupted by flipping at random 30% of the original bits; see Fig. 1 for more details.

Discussion. We have presented a novel technique for the storage of patterns in a Hopfield associative memory. The first step of the method is to fit an Ising model using minimum probability flow learning to a discrete distribution supported equally on a set of binary target patterns. Next, we use the learned Ising model parameters to define a Hopfield network. We show that when the set of target patterns is storable, these steps result in a Hopfield network that stores all of the patterns as fixed-points. We have also demonstrated that the resulting (convex) algorithm outperforms current techniques for training Hopfield networks. We have shown improved recovery of memories from noisy patterns and improved training speed as compared to training by PER. We have demonstrated optimal storage capacity in the noiseless case, outperforming OPR. We have also demonstrated the unsupervised storage of memories from heavily corrupted training data. Furthermore, the learn-

ing rule that results from our method is local; that is, updating the weights between two units requires only their states and feedforward input. As MPF allows the fitting of large Hopfield networks quickly, new investigations into the structure of Hopfield networks are possible [19]. It is our hope that the robustness and speed of this learning technique will enable practical use of Hopfield associative memories in computational neuroscience, computer science, and scientific modeling.

References

- [1] W. Little, "The existence of persistent states in the brain," *Mathematical Biosciences*, vol. 19, no. 1, pp. 101–120, 1974.
- [2] E. Ising, "Beitrag zur Theorie des Ferromagnetismus," *Zeitschrift fur Physik*, vol. 31, pp. 253–258, Feb. 1925.
- [3] J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 79, no. 8, p. 2554, 1982.
- [4] M. Cohen and S. Grossberg, "Absolute stability of global pattern formation and parallel memory storage by competitive neural networks." *IEEE Transactions on Systems, Man, & Cybernetics*, 1983.
- [5] G. Hinton and T. Sejnowski, "Learning and relearning in boltzmann machines," *Parallel distributed processing: Explorations in the microstructure of cognition*, vol. 1, pp. 282–317, 1986.
- [6] J. Sohl-Dickstein, P. B. Battaglino, and M. R. DeWeese, "New method for parameter estimation in probabilistic models: minimum probability flow," *Physical review letters*, vol. 107, no. 22, p. 220601, 2011.
- [7] J. Hertz, A. Krogh, and R. Palmer, *Introduction to the theory of neural computation*. Westview press, 1991, vol. 1.
- [8] G. Weisbuch and F. Fogelman-Soulié, "Scaling laws for the attractors of Hopfield networks," *Journal de Physique Lettres*, vol. 46, no. 14, pp. 623–630, 1985.
- [9] R. McEliece, E. Posner, E. Rodemich, and S. Venkatesh, "The capacity of the Hopfield associative memory," *Information Theory, IEEE Transactions on*, vol. 33, no. 4, pp. 461–482, 1987.
- [10] D. Amit, H. Gutfreund, and H. Sompolinsky, "Statistical mechanics of neural networks near saturation," *Annals of Physics*, vol. 173, no. 1, pp. 30–67, 1987.
- [11] F. Rosenblatt, "The perceptron: a perceiving and recognizing automation (projet para), cornell aeronautical laboratory report," 1957.
- [12] M. Minsky and S. Papert, *Perceptrons*. MIT press, 1988.
- [13] M. Jinwen, "The asymmetric hopfield model for associative memory," in *Neural Networks, 1993. IJCNN'93-Nagoya. Proceedings of 1993 International Joint Conference on*, vol. 3. IEEE, 1993, pp. 2611–2614.
- [14] F. Sommer and P. Dayan, "Bayesian retrieval in associative memories with storage errors," *Neural Networks, IEEE Transactions on*, vol. 9, no. 4, pp. 705–713, 1998.
- [15] L. Personnaz, I. Guyon, and G. Dreyfus, "Collective computational properties of neural networks: New learning mechanisms," *Physical Review A*, vol. 34, no. 5, p. 4217, 1986.
- [16] I. Kanter and H. Sompolinsky, "Associative recall of memory without errors," *Physical Review A*, vol. 35, no. 1, p. 380, 1987.
- [17] T. Cover, "Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition," *Electronic Computers, IEEE Transactions on*, no. 3, pp. 326–334, 1965.
- [18] J. Nocedal, "Updating quasi-newton matrices with limited storage," *Mathematics of computation*, vol. 35, no. 151, pp. 773–782, 1980.
- [19] C. Hillar and N. Tran, "Robust exponential memory in Hopfield networks," *ArXiv e-prints*.