



Annual Review of Biophysics

Hopfield Networks as Models of Emergent Function in Biology

Maria Yampolskaya¹ and Pankaj Mehta^{1,2,3}

¹Department of Physics, Boston University, Boston, Massachusetts, USA; email: mariay@bu.edu, pankajm@bu.edu

²Center for Regenerative Medicine of Boston University and Boston Medical Center, Boston, Massachusetts, USA

³Faculty of Computing and Data Science, and Biological Design Center, Boston University, Boston, Massachusetts, USA

Annu. Rev. Biophys. 2026. 55:323–42

The *Annual Review of Biophysics* is online at biophys.annualreviews.org

<https://doi.org/10.1146/annurev-biophys-101425-023356>

Copyright © 2026 by the author(s).
All rights reserved

Keywords

Hopfield model, modern Hopfield networks, cellular differentiation, self-assembly

Abstract

Hopfield models, originally developed to study memory retrieval in neural networks, have become versatile tools for modeling diverse biological systems in which function emerges from collective dynamics. In this review, we provide a pedagogical introduction to both classic and modern Hopfield networks from a biophysical perspective. After presenting the underlying mathematics, we build physical intuition through three complementary interpretations of Hopfield dynamics: as noise discrimination, as a geometric construction defining a natural coordinate system in pattern space, and as gradient-like descent on an energy landscape. We then survey applications of Hopfield networks in a variety of biological settings, including cellular differentiation and epigenetic memory, molecular self-assembly, and spatial neural representations.



Contents

1. INTRODUCTION	324
2. THE MATHEMATICS OF HOPFIELD MODELS	325
2.1. Overview	325
2.2. Classic Hopfield Network with Random Patterns	326
2.3. Storing Correlated Patterns with the Projection Method	328
2.4. Modern Hopfield Networks	329
3. INTERPRETATIONS OF HOPFIELD DYNAMICS	330
3.1. As Retrieving a Signal	331
3.2. As Projection onto a Subspace	332
3.3. As Traversing a Landscape	333
3.4. Other Perspectives	336
4. HOPFIELD MODELS AS A LENS FOR UNDERSTANDING BIOLOGICAL FUNCTION	337
4.1. Cell Fate Transitions and Epigenetics	338
4.2. Self-Assembly	338
4.3. Neural Representations	339
5. CONCLUSION	339

1. INTRODUCTION

From the millions of proteins in a single cell to the billions of neurons in the human brain, collective behavior is responsible for much of the complex phenomena observed in biology (4, 9, 20, 45, 48, 49). Across scales, biological systems perform specific functions as a result of many parts working together. Modeling the emergence of function from these components is part of what makes biophysics so difficult (25). One promising approach for meeting this challenge is to recast this problem in the language of high-dimensional dynamical systems. The most famous and fruitful example of this approach is the Hopfield model (22). Originally conceived as a model of memory in neural networks, the Hopfield model has now found broad applicability across biophysics, including understanding the epigenetics of cellular identity and differentiation, modeling the self-assembly of heterogeneous structures, and thinking about neural representations of spatial environments in the hippocampus.

The Hopfield model has a long pedigree, and many papers have applied analytical tools from statistical physics to calculate properties of these models, including phase diagrams and storage capacities (5–7, 14, 22, 35, 37, 41). Hopfield models have attracted renewed interest due to their close relationship with modern machine learning techniques such as transformers (for a review that highlights these developments, see 33; 52). Here, we review Hopfield networks from a biophysical perspective, focusing on Hopfield models as exemplars of complex systems exhibiting emergent function (5, 14, 24, 32, 34, 35, 52, 69). We explain the intuition behind Hopfield dynamics, why and how these models work, and how they can be used to understand disparate biophysical phenomena.

The review is organized as follows. In Section 2, we present mathematical descriptions of classic and modern Hopfield networks. In Section 3, we provide three biophysically inspired perspectives for interpreting the dynamics of Hopfield models: (a) as a system that discriminates between signal and noise, (b) as a geometric operation that projects onto a subspace, and (c) as gradient-like

dynamics on a landscape. We also include Section 3.4, which provides an overview of other, more common approaches to thinking about Hopfield models. In Section 4, we turn to applications and describe how Hopfield models have been used to study a variety of biophysical phenomena, including cell fate transitions, molecular self-assembly, and neural representations.

2. THE MATHEMATICS OF HOPFIELD MODELS

2.1. Overview

In this section, we provide a brief self-contained introduction to the mathematics of Hopfield models. We start by introducing the basic mathematical objects used to describe Hopfield models. We then specialize our discussion to three commonly used variants of the Hopfield model: the classic network introduced by Hopfield (22), the projection method (31, 47), and the exponential modern Hopfield network (14, 35, 52). Despite the many implementational differences between traditional and modern networks, they share a common underlying logic, namely (a) a dynamical update for the state of the system and (b) a set of stored patterns that act as attractors (fixed points) of the dynamics. When the dynamics are initialized sufficiently close to a stored pattern, the system flows to the corresponding fixed point, allowing the network to retrieve the stored pattern.

Hopfield models share a set of common ingredients (see **Figure 1**):

- A N -dimensional state vector $\mathbf{x}(t) = (x_1(t), \dots, x_N(t))$ whose i th component $x_i(t)$ encodes the state of the i th dynamical variable at time t . The x_i can be binary valued (e.g., $x_i = \pm 1$) or continuous valued (e.g., $x_i \in \mathbb{R}$). In the computational neuroscience and physics literature, these variables are often called neurons or spins.
- A matrix $\xi_{\mu i}$ with dimensions $P \times N$ describing the P patterns we wish to store. By construction, when the system is in a fixed point corresponding to the v th stored pattern, $x_i = \xi_{vi}$. In the spin glass literature, these stored patterns are often assumed to be random and independent. However, in the biophysical applications that we discuss below, it will be important to consider the case where stored patterns are correlated.
- An update rule or differential equation governing the dynamics, e.g., $x_i(t+1) = f_i(\mathbf{x}(t))$ for discrete time or $\frac{dx_i}{dt} = f_i(\mathbf{x}(t)) - x_i$ for continuous time (see Equation 8). In the Hopfield model, the functions $f_i(\mathbf{x}(t))$ are chosen in such a way as to ensure that all P patterns $\{\xi_v\}$ are fixed points of the dynamics when the number of patterns P is sufficiently small compared to the number of neurons N .
- A P -dimensional order parameter vector $\mathbf{m}(t) = (m_1(t), \dots, m_P(t))$ that measures the alignment of the state of the system $\mathbf{x}(t)$ at a time t with each of the P stored patterns. As discussed below, describing the systems in terms of order parameter offers a rich lens for interpreting Hopfield dynamics.
- An energy function (or, equivalently in this case, Lyapunov function) that decreases with time as a result of the update rule.

What distinguishes model variants from each other is how these basic ingredients are combined to construct the detailed network dynamics. Some models consider discrete variables where $x_i = \pm 1$, whereas others use continuous x_i . Another major difference between model variants is how the stored patterns $\xi_{\mu i}$ are constructed. The classic Hopfield model assumes that these are binary random vectors, whereas other constructions relax this assumption. The biggest difference between model variants is the choice of nonlinear function $f_i(\mathbf{x}(t))$ that defines the update rule. Here, we focus on three choices for this function corresponding to the classic Hopfield model, the projection method, and modern Hopfield networks.




Figure 1

Interpretations of Hopfield models. (a) Hopfield models involve networks of neurons, and the update rule for each neuron depends on the state of every other neuron. In classic Hopfield networks, each neuron has a binary value, and neurons are connected all-to-all with pairwise interactions. (b) Hopfield models are also called associative memory models because they compare the system state, $\mathbf{x}$, to a list of memories, or stored patterns, and retrieve the memory that has the most association with the initial state. States are vectors in neuron space. In this illustration, the neurons take on binary values, though Hopfield models are flexible enough to allow continuous values as well. (c) The generalized order parameters $\{m^\mu\}$ in Equation 12 can be interpreted as coordinates defining stored pattern space. They are the result of projecting from the higher-dimensional neuron space down to the subspace spanned by stored patterns. Blue axes indicate neuron space, black axes are stored pattern space, black circles indicate stored patterns (the v th stored pattern corresponds to $m_\mu = \delta_{\mu v}$ in stored pattern space), and red circles correspond to the system state $[\mathbf{x}]$ in neuron space and $P\mathbf{x}$ (see Equation 17) for the projected component]. Panel adapted from Reference 67. (d) The dynamics of the generalized order parameters for a classic Hopfield network with 3 patterns, 200 neurons, and low temperature ($T = 0.01$; see Equation 23). The order parameters are averaged over 500 iterations of the simulation. Each simulation is started with the same initial state, which has some overlap with the 0th stored pattern, ξ_0 . As the simulation progresses, the network retrieves the 0th stored pattern, so m^0 becomes 1 (indicating perfect alignment) and $m^\mu \neq 0$ become 0 (indicating no alignment with those stored states).

2.2. Classic Hopfield Network with Random Patterns

In 1982, Hopfield (22, 23) formulated his now classic network as a model of collective computation. The network consists of N binary-valued neurons that interact with each other through a pairwise interaction matrix J_{ij} (see Figure 1). At a time t , a neuron can be firing ($x_i(t) = 1$) or silent ($x_i(t) = -1$). Hopfield dynamics can be synchronous, meaning all neurons are updated at each time step, or asynchronous, where one neuron fires at a given time step. Here we discuss the latter. At each time step, a randomly chosen neuron is updated using a dynamical update rule of the form

$$x_i(t+1) = \text{sign} \left(\sum_{j \neq i} J_{ij} x_j(t) \right). \quad 1.$$

The key insight of Hopfield was the realization that it is possible to choose the couplings J_{ij} in such a way as to ensure that a specified set of binary patterns (memories) $\{\xi_\mu\}$ are fixed points of the dynamics (see Figure 1). Hopfield originally used the case of randomly chosen $\xi_{\mu i}$, meaning each

$\xi_{\mu i}$ is independently chosen to be +1 or -1 with equal probability of being +1 or -1. Hopfield showed that choosing the couplings to be of the form

$$J_{ij} = \frac{1}{N} \sum_{\mu=1}^P \xi_{\mu i} \xi_{\mu j} \quad 2.$$

ensures that the patterns $\{\xi_{\mu}\}$ are fixed points of the dynamics as long as the number of patterns is much smaller than the number of neurons, $P \ll N$. We reproduce this argument in Section 3.1.

Hopfield also showed that these dynamics possess a Lyapunov function,

$$E = -\frac{1}{2} \sum_{ij} x_i J_{ij} x_j. \quad 3.$$

To see that E monotonically decreases under the update rule in Equation 1, consider the energy difference between two states that differ at the k th spin:

$$\begin{aligned} \Delta E &= E(x_k = +1) - E(x_k = -1) \\ &= \left(-\frac{1}{2} \sum_{i \neq k, j \neq k} x_i J_{ij} x_j - \sum_j (+1) J_{kj} x_j \right) - \left(-\frac{1}{2} \sum_{i \neq k, j \neq k} x_i J_{ij} x_j - \sum_j (-1) J_{kj} x_j \right) \\ &= -2 \sum_j J_{kj} x_j. \end{aligned}$$

If the state with $x_k = +1$ has lower energy, $\Delta E = -\sum_j J_{kj} x_j < 0$ and $\text{sign}(\sum_j J_{kj} x_j) > 0$, so the update rule returns $x_k(t+1) = +1$. If the state with $x_k = -1$ has lower energy, $\Delta E = -\sum_j J_{kj} x_j > 0$ and $\text{sign}(\sum_j J_{kj} x_j) < 0$, so the update rule returns $x_k(t+1) = -1$.

Hopfield's paper set off a flurry of activity among statistical physicists studying spin glasses (for a pedagogical and comprehensive account, see 5). One powerful concept that emerged from these perspectives was a description of these systems in terms of order parameters m_{μ} that measure the overlap between the current state of the system and stored patterns. For the classic Hopfield network, these order parameters take the form of generalized magnetizations,

$$m_{\mu} = \frac{1}{N} \sum_j \xi_{\mu j} x_j. \quad 4.$$

Both the dynamical update rule and the Lyapunov function E can be reexpressed in terms of order parameters. With Equation 2, it is easy to show that the energy function takes the simple form

$$E = -\frac{N}{2} \sum_{\mu} (m_{\mu})^2.$$

We make use of this expression for E to interpret the dynamics of the Hopfield model in terms of energy landscapes in Section 3.3.

The dynamic update rule (Equation 1) can be rewritten in terms of m_{μ} as

$$x_i(t+1) = \text{sign} \left(\sum_{\mu} \xi_{\mu i} m_{\mu}(t) \right). \quad 5.$$

This form of the dynamical update rule makes it clear that the state of the system is pushed toward stored patterns with which it has higher overlaps (i.e., patterns with large m_{μ}). The nonlinearity in the update rule—in this case, the sign function—forces the system to choose between patterns. However, the almost linear nature of this update rule often results in the presence of spurious attractors—fixed points of the dynamics that are not the stored patterns. The most common spurious attractors are odd mixtures of existing patterns of the form $\xi_{\text{mix}} = \pm \text{sign}(\pm \xi_{\mu_1} \pm \xi_{\mu_2} \pm \xi_{\mu_3})$.

Alternative learning rules with different nonlinearities can be constructed that further reduce the effect of spurious attractors by decreasing the energy of the system in those states (26, 54, 60).

Spurious attractors play an important role in the study of Hopfield networks. As the number of stored patterns P increases, the form of the dynamical update rule in Equation 5 suggests the system will be pulled in more distinct directions because the state of the system overlaps with multiple stored patterns. This frustration in the dynamics ultimately places a limit on the number of stored patterns that can be successfully recalled, giving rise to the concept of a storage capacity. In classic Hopfield models, the storage capacity $P_{\max}$ scales linearly with the number of neurons N (5, 7, 22). We provide an intuitive explanation for this result in Section 3.1.

Finally, thus far we have limited our discussion to discrete-time updates. However, it is straightforward to formulate the dynamics of Hopfield networks in continuous time. To do so, one simply approximates the discrete update rule using the derivatives as follows:

$$\frac{dx_i}{dt} \approx \frac{x_i(t + \Delta t) - x_i(t)}{\Delta t} \quad 6.$$

$$\approx \frac{1}{\tau} (x_i(t + 1) - x_i(t)) \quad 7.$$

$$= \frac{1}{\tau} \left(\text{sign} \left(\sum_{j \neq i} J_{ij} x_j(t) \right) - x_i(t) \right), \quad 8.$$

where $x_i(t + 1)$ is the discrete-time update rule, τ is a constant that sets the timescale for the dynamics, and in going from the second to the third equation we have used Equation 1.

2.3. Storing Correlated Patterns with the Projection Method

Up to this point, we have considered a classic Hopfield model with random patterns. However, when the patterns are correlated, the classic Hopfield network works poorly. The reason for this is that the system state $\mathbf{x}$ will generically have large overlaps with multiple patterns. Namely, if two patterns ξ_μ and ξ_ν are highly correlated, then Equation 4 implies that if m_μ is large, then m_ν must also be large. This interference between patterns gives rise to frustration in the dynamics because the system is pulled in multiple directions by the dynamic update rule in Equation 5. This results in increased errors in pattern retrieval and low storage capacity.

Kanter & Sompolinsky (31) proposed a simple modification of the classic Hopfield scheme, the projection method, that allows for pattern retrieval even when patterns are correlated. The key insight of this work was to modify Equation 2, the rule for constructing neural couplings J_{ij} , to account for correlations between patterns. In the projection method, one starts by defining the correlation matrix between patterns

$$g_{\mu\nu} = \frac{1}{N} \sum_{ij} \xi_{\mu i} \xi_{\nu j}. \quad 9.$$

This matrix is equivalent to a metric tensor constructed from the basis vectors (ξ_μ, ξ_ν) of the nonorthogonal pattern subspace (see Section 3.2). Using the typical notation for a metric tensor, we use lower indices and upper indices to distinguish between covariant and contravariant vectors. This is helpful for careful tracking of coordinate-invariant quantities like scalars. In line with this

thinking, we define the inverse of the metric tensor using upper indices:

$$g^{\mu\nu} = g_{\mu\nu}^{-1}. \quad 10.$$

The dynamics are still governed by Equation 1, but now with the couplings defined according to the projection rule

$$J_{ij} = \frac{1}{N} \sum_{\mu, v=1}^P \xi_{\mu i} g^{\mu\nu} \xi_{\nu j}. \quad 11.$$

In Section 3.2 we give an intuitive explanation of the origins of this rule and why it works. The basic reasoning is that switching from uncorrelated patterns to correlated ones is equivalent to switching from an orthogonal basis to a nonorthogonal one. Using a nonorthogonal basis changes how distances are defined, which is mathematically represented with a metric tensor.

It is also helpful to define decorrelated order parameters,

$$m^\mu = \sum_v g^{\mu\nu} m_\nu, \quad 12.$$

that generalize the usual magnetizations (Equation 4). Note that we use a raised index to distinguish this decorrelated order parameter from magnetizations. When the stored patterns are orthogonal (i.e., $g_{\mu\nu} = \delta_{\mu\nu}$), the generalized order parameters are identical to the classic one ($m^\mu = m_\mu$). A straightforward calculation shows that for coupling as in Equation 11, the Lyapunov function (Equation 3) takes the simple form (31, 47)

$$\begin{aligned} E &= -\frac{N}{2} \sum_{\mu, \nu} m_\mu g^{\mu\nu} m_\nu \\ E &= -\frac{N}{2} \sum_\mu m^\mu m_\mu. \end{aligned} \quad 13.$$

Written in this form, it is clear that the energy is a dot product of the states in order parameter space, with the metric tensor correcting for the distortion of angles due to the nonorthogonality of patterns. Unless otherwise stated, we assume that patterns are correlated. For this reason, we make extensive use of both generalized order parameters and magnetizations in what follows.

2.4. Modern Hopfield Networks

The final class of models we discuss is a large class of networks that are commonly called modern Hopfield networks. Modern Hopfield networks generalize classic Hopfield models by introducing higher-order interactions between spins (18, 34, 35). Unlike classic Hopfield models where the energy function is second order in the spins (see Equation 3), in modern Hopfield networks the energy function can have terms of any order.

In this review, we restrict our discussion to exponential modern Hopfield networks where the dynamic update rule and energy function take the form (14, 37, 52)

$$\begin{aligned} x_i(t+1) &= \sum_\mu \xi_{\mu i} \sigma^\mu(\beta m^\mu(t)) \\ E(t) &= -\sum \log \left(\sum_\mu \exp(m^\mu(t)) \right) + \frac{1}{2} \sum_i x_i^2(t), \end{aligned} \quad 14.$$

Table 1 A comparison of classic and modern Hopfield networks

	Classic	Exponential modern
State variables	Typically binary ($x_i = \pm 1$)	Typically continuous ($x_i \in \mathbb{R}^N$)
Dynamics	$x_i(t+1) = \text{sign}(\sum_{\mu} \xi_{\mu i} m_{\mu}(t))$	$x_i(t+1) = \sum_{\mu} \xi_{\mu i} \sigma^{\mu}(\beta m^{\mu}(t))$
Energy function	$E = -\frac{1}{2} \sum_{ij} x_i J_{ij} x_j = -\frac{1}{2} \sum_{\mu} (m_{\mu})^2$	$E = -\sum \log(\sum_{\mu} \exp(m^{\mu})) + \frac{1}{2} \sum_i x_i^2$
Pattern storage capacity	Linear in the number of neurons ($P_{\max} \propto N$)	Exponential in the number of neurons ($P_{\max} \propto 2^N$)

where σ is the soft-max function with an inverse temperature parameter β ,

$$\sigma^{\mu}(\beta m^{\mu}) = \frac{e^{\beta m^{\mu}}}{\sum_{\nu} e^{\beta m^{\nu}}}, \quad 15.$$

and m^{μ} is as in Equation 12. A visualization of the energy function is shown in **Figure 3a**. One of the advantages of these dynamics is that unlike the classic Hopfield network, both the neurons $\mathbf{x}$ and the patterns ξ can be real valued.

For sufficiently small temperatures, it is easy to show that the stored patterns $\xi_{\mu i}$ are fixed points of the dynamics. To see this, it is helpful to rewrite the dynamic update rule entirely in terms of generalized order parameters. To do this, we multiply both sides of the dynamic update rule by $\frac{1}{N} \sum_{\nu} g^{\nu v} \xi_{\nu i}$ and sum over i to get

$$\begin{aligned} \frac{1}{N} \sum_{i,v} g^{\nu v} \xi_{\nu i} x_i(t+1) &= \frac{1}{N} \sum_{\mu,v} g^{\nu v} \sum_i \xi_{\nu i} \xi_{\mu i} \sigma^{\mu}(\beta m^{\mu}(t)) \\ \sum_v g^{\nu v} m_{\nu}(t+1) &= \sum_{\mu,v} g^{\nu v} g_{\nu \mu} \sigma^{\mu}(\beta m^{\mu}) \\ m^{\nu} &= \sigma^{\nu}(\beta m^{\nu}(t)), \end{aligned} \quad 16.$$

where in the second line we have used Equations 4 and 9, and in the third line we have used Equation 12 and the fact that $\sum_{\mu} g^{\nu v} g_{\nu \mu} = \delta_{\gamma \mu}$. Notice that in the zero temperature limit $T \rightarrow 0$, $\beta \rightarrow \infty$, any order-parameter vector $\mathbf{m} = (0, \dots, 0, 1, 0, \dots, 0)$ with $m^{\gamma} = 0$ for all $\gamma \neq \lambda$ (with $m^{\lambda} = 1$) is a fixed point of the dynamics. This is because $\sigma^{\gamma}(\beta m^{\gamma}) \approx 0$ for $\gamma \neq \lambda$ and $\sigma^{\gamma}(\beta m^{\gamma}) \approx 1$ if $\gamma = \lambda$. Just like in the classic Hopfield networks, the network retrieves the stored patterns it is closest to.

The update rule for modern Hopfield networks bears some striking similarities to the classic Hopfield construction. It involves a nonlinear function moving the system toward the state with the highest alignment m^{μ} . Now, however, the nonlinearity is a soft-max function instead of a sign function, leading to exponential storage capacity (37) (see Section 3.3 for an intuitive explanation). Additionally, the stored patterns are outside the nonlinearity in the sum, which allows for continuous-valued stored patterns. Table 1 provides a comparison of modern and classic networks.

3. INTERPRETATIONS OF HOPFIELD DYNAMICS

Section 2 introduced the mathematics of the Hopfield model. In this section, we focus on developing intuition. To do so, we present different perspectives through which one can view these systems, each of which emphasizes a different aspect of these networks.

3.1. As Retrieving a Signal

One way of interpreting Hopfield dynamics is as a system that discriminates between a target pattern—the pattern it is most aligned with and is expected to converge to—and the nontarget patterns. In other words, the system picks out the signal of the intended pattern from the noise caused by interference from other patterns. This interference is due to random overlaps between stored patterns, which become more significant when there are a large number of patterns relative to the number of neurons. For simplicity, we demonstrate this concept using the classic Hopfield model with orthogonal stored patterns, but the intuition presented here holds in more general settings.

Suppose the system is in the stored state corresponding to pattern v ($x_i = \xi_{vi}$). The correct behavior of the system is to remain in the target pattern v and not to flow to other patterns. However, if the number of stored patterns P is close to the number of neurons N , interference from nontarget patterns can cause errors in retrieval and even destabilize the fixed point corresponding to the target pattern. To see this, consider the update rule in Equation 1 with couplings given by Equation 2. When the system state is aligned with the v th pattern, $\mathbf{x}(t) = \xi_v$, the update rule takes the form

$$\begin{aligned} x_i(t+1) &= \text{sign} \left(\frac{1}{N} \sum_{\mu,j} \xi_{\mu i} \xi_{\mu j} \xi_{v j} \right) \\ &= \text{sign} \left(\frac{1}{N} \left(\sum_j \xi_{v i} \xi_{v j} \xi_{v j} + \sum_{\mu \neq v, j} \xi_{\mu i} \xi_{\mu j} \xi_{v j} \right) \right) \\ &= \text{sign} \left(\frac{1}{N} \left(N \xi_{v i} + \sum_{\mu \neq v, j} \xi_{\mu i} \xi_{\mu j} \xi_{v j} \right) \right) \\ &= \text{sign} \left(\xi_{v i} + \frac{1}{N} \sum_{\mu \neq v, j} \xi_{\mu i} \xi_{\mu j} \xi_{v j} \right). \end{aligned}$$

The first term inside the sign function is just the target pattern and can be thought of as the signal. The second term is an interference term that comes from the nontarget patterns and can be thought of as noise. Whether the dynamics remain in the correct pattern depends on the relative magnitude of these two terms. When the second term becomes nonnegligible, there will be errors in pattern storage and retrieval. In extreme cases, this noise term overwhelms the signal and the pattern is no longer stable.

Since $\xi_{\mu i}$ are drawn randomly, the interference term is a sum of the product of three random variables. Because $\xi_{\mu i} = \pm 1$, this sum is like a random walk over $(P-1)N$ random steps. The error caused by this interference term is the variance of a random walk (equivalently, it is the variance of a sum of random variables, so one can apply the central limit theorem). We can rewrite the sum in terms of the mean μ_w (which is zero for a random walk), variance σ_w , and a random normally distributed variable $z \sim N(0, 1)$:

$$\begin{aligned} \sum_{\mu \neq v, j}^{P,N} \xi_{\mu i} \xi_{\mu j} \xi_{v j} &= \mu_w + \sigma_w z \\ &= 0 + z \sqrt{(P-1)N}. \end{aligned}$$

This yields the following equation for the update rule.

$$\begin{aligned} x_i(t+1) &= \text{sign} \left(\xi_{v i} + \frac{\sqrt{(P-1)N}}{N} z \right) \\ &\approx \text{sign} \left(\xi_{v i} + \sqrt{\frac{P}{N}} z \right) \end{aligned}$$

The error in memory retrieval scales like $\sqrt{\frac{P}{N}}$. For this reason, the number of network neurons must be much larger than the number of stored patterns to keep the errors low. This expression also shows that the storage capacity scales linearly with the number of neurons, $P_{\max} \propto N$. An analogous argument can be formulated for modern Hopfield networks and one finds that the storage capacity is exponential in the number of neurons N (14, 37, 52).

3.2. As Projection onto a Subspace

Another insight into Hopfield models comes from recognizing that the transformation from the state space of the system x_i to order parameters m^μ can be viewed as change of basis from neuron space to pattern space. This way of thinking about Hopfield models is especially powerful when the number of patterns is less than the number of neurons $P \leq N$. As shown in **Figure 1b**, in this case, the patterns span a subspace of the full neuron space whose natural basis is the patterns themselves.

In particular, we can construct a projection matrix P that acts on a state $\mathbf{x}$ and projects it down to the subspace spanned by the stored patterns. From basic linear algebra, we know that this projection matrix takes the form

$$P_{ij} = \sum_{\mu, \nu} \xi_{\mu i} g^{\mu\nu} \xi_{\nu j}, \quad 17.$$

where $g^{\mu\nu}$ is the inverse of $g_{\mu\nu}$ given in Equation 9. It is easy to verify that this is indeed a projection matrix by confirming that $P^2 = P$. In terms of the projection matrix, we can decompose a vector $\mathbf{x}$ into a component that lies in the pattern subspace and a component orthogonal to this subspace,

$$\mathbf{x} = P\mathbf{x} + \mathbf{x}^\perp, \quad 18.$$

where by definition $\mathbf{x}^\perp = (I - P)\mathbf{x}$. Using Equation 12, this decomposition can be written as

$$x_i = \sum_{\mu} m^\mu \xi_{\mu i} + x_i^\perp. \quad 19.$$

This expression shows that the generalized order parameters m^μ are the natural coordinate system for describing the state of the system in terms of patterns. The projection method gets its name from the fact that the couplings J_{ij} defined in Equation 11 are exactly the projection matrix $J_{ij} = P_{ij}$. With this in mind, we see that the energy of the classic Hopfield model in Equation 13 is simply the magnitude of the orthogonal component of the current state, namely $E = \mathbf{x}^\perp \cdot \mathbf{x}^\perp$ in the case of states with constant magnitude (in other words, $\frac{1}{N} \mathbf{x} \cdot \mathbf{x} = 1$). Minimizing Hopfield energy is equivalent to minimizing the magnitude of this orthogonal component, which is the same process used in least squares linear regression.

Thus, far we have focused our discussion on classic Hopfield models. However, this perspective also sheds light on dynamics of modern Hopfield networks. By substituting Equation 16 into Equation 14, we can rewrite the dynamics of the modern Hopfield network as

$$x_i(t+1) = \sum_{\mu} \xi_{\mu i} m^\mu(t+1). \quad 20.$$

This equation emphasizes one of the key intuitions behind modern Hopfield networks: The natural space for thinking about dynamics is not neuron space, $\mathbf{x}$, but pattern space, $\mathbf{m}$. The dynamics occur in pattern space and are then converted to neuron space by the stored patterns, $\{\xi\}$. The primary constraint on the dynamical system in pattern space is that its attractors are located on the vertices of the simplex (i.e., the points $m_\mu = \delta_{\mu\nu}$ are attractors for all stored patterns ν). This is generally achieved using a nonlinear function, such as sign or soft-max, that takes m^ν as input and limits the dynamics so that $|m^\nu| \leq 1$. Section 3.3 shows how one can use these ideas to think about Hopfield models as gradient-like systems.

3.3. As Traversing a Landscape

Hopfield models are famously energy-based models: They generally have an associated Lyapunov or energy function that decreases as time passes. In other words, Hopfield dynamics descend an energy landscape and find minima corresponding to stored patterns. As previously mentioned, the energy for the classic Hopfield model with generalized order parameters is

$$E^{\text{quad}} = -\frac{1}{2N} \sum_{i,\mu,v,j} x_i \xi_{\mu i} g^{\mu v} \xi_{v j} x_j = -\frac{N}{2} \sum_{\mu} m_{\mu} m^{\mu}. \quad 21.$$

This energy has the same form as an Ising model with all-to-all pairwise interactions, a model that lives in infinite dimensions. This makes it hard to visualize this energy landscape for a generic choice of couplings J_{ij} . However, Hopfield models have a built-in dimensional reduction: the order parameters $\mathbf{m}$. In terms of m_{μ} and m^{μ} , E^{quad} is a quadratic function. As shown in **Figure 2a,b**, it is simply an inverted P -dimensional parabola centered at $\mathbf{m} = 0$. The dynamics of the Hopfield model cause the system to move down this parabola. However, the system does not descend indefinitely; it is restricted by the nonlinearity in the dynamics. The trajectories are limited to those allowed by the sign, soft-max, or other nonlinear function. One of the restrictions imposed by the nonlinearity is a bound on $\mathbf{m}$ -space, such that $|m^{\mu}| \leq 1$.

To see this, one can rewrite the dynamics of both classic and modern Hopfield models in terms of E^{quad} as

$$x_i^{\text{classic}}(t+1) = \text{sign} \left(-\frac{1}{N} \sum_{\mu} \xi_{\mu i} \frac{\partial E^{\text{quad}}}{\partial m_{\mu}} \right)$$

$$x_i^{\text{modern}}(t+1) = \sum_{\mu} \xi_{\mu i} \sigma^{\mu} \left(-\frac{\beta}{N} \frac{\partial E^{\text{quad}}}{\partial m_{\mu}} \right).$$

This shows that the update rule of both classic and modern networks performs a gradient-like descent on the parabolic landscape defined by E^{quad} but rectified through different nonlinear functions (a sign function for the classic model, a soft-max function for the exponential modern network). The sharper nonlinearity used in the modern Hopfield model amplifies small differences in the gradients in pattern space, allowing these networks to have a dramatically higher storage capacity than classic models. **Figure 3b,c** shows the trajectories of the exponential modern Hopfield model in $\mathbf{m}$ -space, with each point colored by the quadratic energy E^{quad} . Although the nonlinearity in dynamics changes the trajectories, in both cases the system ultimately descends the parabolic energy landscape defined by E^{quad} .

While the order parameters $\mathbf{m}$ provide a natural, lower-dimensional set of coordinates to observe the energy landscape, the large number of patterns being stored makes it difficult to visualize the full dynamics in pattern space. To get a broad view of a system with many attractors in one plot, one can use dimensional reduction methods like principal component analysis (PCA) (1, 27). **Figures 2c** and **3d-f** use PCA to visualize the energy landscape for the classic and modern models with varying numbers of stored patterns. When the Hopfield model is in the regime of stable retrieval, the quadratic energy decreases along the trajectory of the dynamics.

Up to this point, we have restricted our discussion of the classic Hopfield model to zero temperature ($T = 0$). In analogy to spin systems like the Ising model, one can introduce a nonzero temperature T that allows spins to flip stochastically via Glauber dynamics. This involves replacing the deterministic update rule in Equation 1 with a temperature-dependent stochastic rule that depends on energy differences. More precisely, at each time step, we choose a random spin i and update its state at the next time step with probabilities

$$Pr(x_i(t+1) = +1) = \frac{e^{E_{+1}/T}}{e^{E_{+1}/T} + e^{E_{-1}/T}} \quad 22.$$




Figure 2

Energy landscape of the classic Hopfield model. The energy is an inverted parabola in the space of the order parameters, $\{m_\mu\}$, as in Equation 21. The stored patterns are shown as black circles with white outlines, and the negatives of the stored patterns are shown as white circles with black outlines. (a) The energy as a function of just one order parameter, m_0 , with $m_\mu \neq 0 = 0$. The stored pattern $\xi_\mu = 0$ and its negative ($-\xi_\mu = 0$) are located at $m_0 = 1$ and $m_0 = -1$, respectively. These points act as attractor states because Hopfield dynamics descend the energy landscape but are restricted to $|m_0| \leq 1$ (this constraint is shown with a dashed line). (b) The energy landscape in two dimensions. (c, i) The phase diagram of the classic Hopfield model. The x -axis is the ratio of patterns to neurons, $\alpha = P/N$. The y -axis is the temperature of the Glauber dynamics (see Equation 23). Above the T_m curve, the model acts as a spin glass with many local minima. Below this curve, the stored patterns are metastable states. Below the T_c curve, the stored patterns are stable states and act as attractors. This is the region of stable retrieval. Subpanel adapted with permission from Reference 7; copyright 1985 American Physical Society. (ii-iv) Three principal component analysis plots of Hopfield network dynamics, $\mathbf{x}(t)$, in the three phases with varying α , at $T = 0.25$. Each plot shows many trajectories of Hopfield models, with each trajectory starting at a different random state. Each state is colored by the energy. States in the center are early in the trajectory while states toward the periphery correspond to later time points. (ii) Hopfield dynamics in the stable retrieval phase. The trajectories show an inverted parabola, as in panel b, with the stored patterns (and their negatives) at the global minima, or outer edges, of the parabola. (iii) Hopfield dynamics in the metastable retrieval phase. The trajectories still create a landscape that looks similar to panel b, but now some stored patterns are located at the inner part of the parabola rather than at the global minima. (iv) Hopfield dynamics in the spin glass phase. The energy landscape has many local minima, and stored patterns are not located at the global minima. The Hopfield model no longer retrieves the stored patterns.




Figure 3

Energy landscape of the exponential (soft-max function) Hopfield model. (a) The log-sum-exp energy of Equation 14 plotted in two dimensions as a function of order parameters $\{m^\mu\}$. Each stored pattern has its own basin of attraction, which is separated by energy barriers. The soft-max function limits the dynamics, such that $\sum_\mu m^\mu \leq 1$, so that although the stored patterns are not at the minima of the energy function, they are at the lowest energies accessible to the system. (b–f) Scatter plots of trajectories colored by the classic, quadratic energy in Equation 21 show that the soft-max Hopfield model also descends an inverted parabola, although the dynamics restrict the trajectories differently than in the classic Hopfield model. For each scatter plot, the continuous exponential Hopfield model is initialized at many random states and then evolved. In each system, the number of neurons is $N = 100$ and the temperatures are set so that the system is well within the regime of stable pattern retrieval ($\beta = 100 \log P$). (b–d) Dynamics for three stored patterns.

(b,c) Trajectories as a function of the order parameters $\{m^\mu\}$ in two and three dimensions. The order parameters provide a natural dimensionality reduction of the system, with the stored patterns forming a $P - 1$ simplex in this space. The stored patterns are the vertices of this simplex. (d–f) Principal component analysis allows visualization of the $P - 1$ simplex in two dimensions. In each case, the trajectories are restricted to the simplex while moving down the inverted parabola in Equation 21 and **Figure 2**. (d) Trajectories for $P = 3$. (e) Trajectories for $P = 10$. (f) Trajectories for $P = 50$. Even in the case of many stored patterns, the system achieves stable retrieval.

$$Pr(x_i(t+1) = -1) = \frac{e^{E_{-1}/T}}{e^{E_{+1}/T} + e^{E_{-1}/T}}, \quad 23.$$

where $E_{\pm 1} = E(x_i = \pm 1)$ is the Hopfield energy with the i th spin set to ± 1 . The addition of temperature introduces randomness. When T is small, the Hopfield model retrieves patterns in a manner similar its zero temperature behavior. When T is large, the model fails to retrieve the stored patterns and instead follows random trajectories.

The phases of the Hopfield model at varying temperatures and pattern-packing fractions $\alpha = \frac{P}{N}$ have been studied extensively (5–7, 41). The phase diagram is shown in **Figure 2**. The Hopfield model can perform accurate pattern retrieval below the T_c curve. This corresponds to

a parameter region with low temperatures and small pattern-packing fractions $\alpha \ll 1$. In this region of the phase diagram, the stored patterns are global minima of the energy landscape. This can be seen in **Figure 2c**, subpanel *ii*, which looks very similar to the quadratic zero-temperature landscape shown in **Figure 2b**.

When the number of stored patterns increases to a point between the T_c and the T_m curves, the model is in the metastable retrieval phase. Some, if not all, stored patterns are metastable states: They are local minima but no longer global minima. Even if a stored pattern is retrieved, the system may leave that state due to fluctuations. The energy landscape for the metastable phase is shown **Figure 2c**, subpanel *iii*. While the energy still resembles a quadratic potential, some stored patterns are in the inner parts of the potential, at higher energy states than other points along the trajectories. Finally, the Hopfield model's pattern retrieval ability is completely destroyed in the spin glass phase, where the number of patterns exceeds the storage capacity, the temperature is too high, or both. In this phase, stored patterns aren't minima at all. In **Figure 2c**, subpanel *iv*, the energy landscape looks less quadratic and more irregular, and the stored patterns are scattered throughout the two principal components.

3.4. Other Perspectives

The literature on Hopfield models is vast and this review is far from comprehensive. In this section, we present a small sampling of other ways to understand these models. Much of the literature on modern Hopfield networks is motivated by an interest in interpretable machine learning. This is because Hopfield models are a case of neural networks that can be studied by analytical methods, and they can also be reinterpreted as restricted Boltzmann machines (3, 8, 40, 56). Many tools from the statistical physics of spin glasses, such as the cavity and replica methods, have been applied to understand the behavior of Hopfield networks (2, 41). Additionally, there is a large body of work dedicated to finding the storage capacity and scaling of spurious states of Hopfield models with different interactions, interaction functions, learning rules, and so on (2, 14, 18, 35, 37, 44, 52).

Before proceeding to biophysical applications, we discuss one final useful perspective on these models inspired by the bipartite structure of restricted Boltzmann machines and the biology of neural connections (34). Instead of formulating Hopfield models in terms of all-to-all connected neurons x_i as in previous sections, this construction involves dividing the neurons into two groups: visible neurons, v_i , and hidden neurons, h_μ . Connections are allowed between hidden and visible neurons, but no connections are allowed between neurons in the same group. In this construction, correlations between visible neurons are induced entirely by hidden units.

The dynamics of these networks can be written as

$$\begin{aligned}\tau_f \frac{dv_i}{dt} &= \sum_{\mu} \xi_{\mu i} f_{\mu} - v_i + I_i \\ f_{\mu} &= \frac{\partial L_h}{\partial h_{\mu}} \\ \tau_b \frac{dh_{\mu}}{dt} &= \sum_i \xi_{\mu i} g_i - h_{\mu} \\ g_i &= \frac{\partial L_v}{\partial v_i},\end{aligned}$$

where τ_f and τ_b are constants that set the timescale for dynamics; I_i is the input current to the visible neurons (which we set to zero for the Hopfield models described above); and f_{μ} and g_i are nonlinear functions that describe the outputs of hidden neuron μ and visible neuron i , respectively,

and are defined in terms of derivatives of Lagrangian functions, denoted L_h and L_v for hidden and visible neurons, respectively. It is possible to show that these dynamics admit a Lyapunov function of the form

$$E = \left[\sum_i (v_i - I_i) g_i - L_v \right] + \left[\sum_\mu b_\mu f_\mu - L_h \right] - \sum_{\mu,i} f_\mu \xi_{\mu i} g_i.$$

Different choices of L_h and L_v give rise to different forms of Hopfield dynamics, including the classic and modern exponential models. For example, these dynamics reduce to those of the modern exponential model when the Lagrangians are chosen to be

$$L_h = \log \left(\sum_\mu e^{b_\mu} \right)$$
$$L_v = \frac{1}{2} \sum_i v_i^2,$$

with the additional assumption that the hidden neuron dynamics are fast ($\tau_h = 0$). In this limit, $\frac{db_\mu}{dt} = 0$ and the hidden neurons become equivalent to the order parameters because $b_\mu = \sum_i \xi_{\mu i} g_i$.

4. HOPFIELD MODELS AS A LENS FOR UNDERSTANDING BIOLOGICAL FUNCTION

In this section, we discuss examples of biological systems in which Hopfield networks or associative memory has been used to model complex biophysical phenomena. We focus on three applications where this approach has been especially fruitful: (a) the epigenetics of cellular identity and cell fate transitions, (b) self-assembly, and (c) neural representations in the context of spatial navigation (see Figure 4).


Neurons, state variables	Interactions between neurons	Stored patterns	Retrieval dynamics
Genes, gene expression	Gene regulation	Cell types 	Differentiation 
Molecules, molecular concentrations	Binding strength, concentrations	Structures 	Self-assembly 
Neurons, neuron firing state	Neural connections	Spatial environments 	Mental exploration/ memory retrieval 

Figure 4

Hopfield networks have been used as models of emergent function in cell fate differentiation, molecular self-assembly, and spatial memory. For each of these biological contexts, the analogies to Hopfield neurons, interactions, stored patterns, and retrieval dynamics are shown. In differentiation, stem cells specify into cell types through changes in gene expression. In self-assembly, different concentrations of promiscuous molecules lead to the assembly of different stored structures. In spatial memory retrieval, neurons known as place cells are responsible for retrieving mental maps of known environments.

4.1. Cell Fate Transitions and Epigenetics

Many animals have organs composed of highly specialized cells with distinct functions (e.g., lung secretory cells, red blood cells). These cell types arise during development via the dynamic process of differentiation. The stability and robustness of a cell's identity once it has been acquired suggest that stable cell fates correspond to dynamical attractors of the genetic and epigenetic regulatory networks governing cellular differentiation. This idea was popularized by Waddington's (61) now famous landscape metaphor, in which the developmental pathways leading from embryonic stem cells to mature types were analogized with canals or valleys in a landscape. Other works suggest that diseased states may also correspond to attractor states created by mutations or reached through a failure to transition between healthy states (38, 58, 62).

There is a growing body of work attempting to formalize the Waddington metaphor using dynamical systems theory and Hopfield models (10, 16, 17, 19, 28, 32, 42, 51, 53, 55, 57, 63). During differentiation, developmental gene regulatory networks use chemical and mechanical signals to guide the state of cells to particular cell fates. Like patterns in a Hopfield model, different cell types are stored and then retrieved by the cell (see **Figure 1**). This has inspired work on modeling cell identity using Hopfield-inspired networks (10, 32, 36, 50, 55, 68). In these models, the state of the cell, $\mathbf{x}$, encodes the global gene expression profile (or alternatively, the state of various developmental enhancers), in which x_i is the expression of the i th gene/enhancer. The stored patterns $\{\xi\}$ correspond to the gene expression profiles of stable cell fates, in which $\xi_{\mu i}$ is the expression of the i th gene in the μ th cell type. These gene expression profiles are often highly correlated, especially for cell types within the same organ, so it is important to use the generalized order parameter m^μ to describe the state of the system. In the developmental context, m^μ measures the similarity between a cell's current gene expression profile $\mathbf{x}$ and the gene expression profile characterizing cell fate μ .

The use of generalized order parameters $\{m^\mu\}$ as coordinates of cell fate space has proved extremely useful for analyzing high-dimensional gene expression data and is especially suited for making use of the vast amount of single-cell RNA-sequencing data contained in developmental atlases (21, 29, 36, 50, 67). In particular, the use of Hopfield-inspired order parameters provides a powerful and interpretable alternative to stochastic dimensional reduction techniques like t-SNE and UMAP for analyzing and visualizing gene expression data (13, 39, 59). By calculating the generalized order parameters m_μ for cells, we can assign every cell a unique coordinate in cell fate space. This coordinate system can be used for both visualization and classification of cells into cell fates. Because the stored memories are calculated directly from single-cell atlases and are fixed, order parameter-based methods also allow for robust inferences of cell identity that are consistent across samples and time (67).

Whereas the previous works focus on gene expression, work by Owen et al. (46) uses the lens of the Hopfield models to explore how the 3D organization of chromatin allows cells to implement epigenetic memory. Within the model, chromatin regions play a role analogous to neurons, with connections between regions controlled by chromatin modifications. These couplings are learned through a “mark together, park together” learning rule analogous to the “fire together, wire together” rule implied by Equation 2. Finally, the mark dynamics play a role similar to that in the dynamic update rule in Equation 1.

4.2. Self-Assembly

A growing body of work has begun to reveal a deep and surprising connection between the physics of self-assembly and the dynamics of associative memory in neural networks (11, 12, 15, 43, 69). At the heart of this connection lies a simple but powerful idea: Just as a Hopfield network retrieves

a stored memory from a partial cue, a multicomponent molecular system can self-organize into a target structure in response to incomplete or noisy chemical inputs. Though these works focus largely on simpler in vitro biophysical and soft-matter systems, they suggest that the uncanny ability of biological systems to reuse the same components to form different structures may also be governed by Hopfield-like design principles.

In these systems, the roles of neurons and couplings are played by physical components and their interactions. The state of the system at any time can be described by the concentration of molecular components, for example, different species of proteins, DNA strands, chemical building blocks, or concentration of a liquid component. Couplings between neurons are replaced by pairwise interaction strengths between components. A major new source of complication in these systems is nonspecific binding between components. This nonspecific binding causes interference between target structures analogous to interference between patterns in the classic Hopfield model, ultimately limiting the capacity of these networks to form disparate structures.

4.3. Neural Representations

Finally, we briefly note that there has been a renewed interest in Hopfield models in the context of neural representations of spatial structures in the hippocampus (for a review of these developments, see 64; 65, 66). In this context, the stored patterns correspond to internal maps of different locations that have been learned by the individual. Retrieval of a pattern corresponds to mentally exploring an internal map when presented with initial conditions similar to a known environment. While a detailed account of this literature is beyond the scope of this review, there are several interesting developments in the debates on cognitive maps that are likely to be of interest to the broader biophysical community.

These works suggest that Hopfield models can generalize beyond memorizing patterns in order to learn latent representations, which correspond to a broader manifold of states than patterns alone. This raises the question of how Hopfield-like models generalize to new memories that were not explicitly a priori specified (64, 66). This topic has also been explored from a statistical physics perspective in References 30 and 44. In many biophysical applications, we would like the system to function for a whole class of patterns, structures, memories, or cell types that are drawn from a probabilistic ensemble with shared characteristics. How can we generalize Hopfield models to account for this by learning latent representations in addition to memorizing patterns?

Another development concerns the ability to extend Hopfield models to include additional sensory inputs and control signals that shape stored memories (64). In the context of spatial neural representations, this often involves debates about if and how input from the cortex binds together memories stored by the hippocampus. This raises the more general question of how one can modify Hopfield models to account for landscapes where memories are dynamic, with attractors becoming stable and unstable depending on external cues. A similar question has been explored by Yampolskaya et al. (68) in the context of cell fate identity where developmental signals play a role analogous to cortex signals. Both areas represent interesting directions for further investigation.

5. CONCLUSION

Hopfield models offer a striking example of how simple dynamical rules in high-dimensional systems can give rise to robust, emergent function in biology. These models provide a powerful lens through which to understand phenomena ranging from cell fate decisions to molecular self-assembly to neural representation. In this review, we have presented a self-contained and pedagogical introduction to this rich body of work, with an emphasis on building intuition by exploring multiple perspectives through which to view these models. While the mathematical

details may vary across systems, this conceptual lens reveals common biophysical principles that underlie emergent function in diverse biological contexts. More broadly, Hopfield models serve as a compelling example of how a well-posed mathematical model can connect the properties of microscopic interactions to macroscopic behaviors. As we continue to search for organizing principles that govern collective dynamics in biology, the Hopfield model remains an essential conceptual tool for understanding emergent function in biophysical systems.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

The authors would like to thank our collaborators, members of the Mehta and Kondev groups, the Kotton lab, and Laertis Ikonomou for many useful discussions that influence this review. This work was funded by a Chan Zuckerberg Initiative (CZI) investigator grant in theory and biology to P.M. and a National Institutes of Health/National Institute of General Medical Sciences grant 1R35GM119461 to P.M.

LITERATURE CITED

1. Abdi H, Williams LJ. 2010. Principal component analysis. *WIREs Comput. Stat.* 2(4):433–59
2. Agliari E, Albanese L, Alemanno F, Alessandrelli A, Barra A, et al. 2023. Dense Hebbian neural networks: a replica symmetric picture of supervised learning. *Phys. A Stat. Mech. Appl.* 626:129076
3. Agliari E, De Marzo G. 2020. Tolerance versus synaptic noise in dense associative memories. *Eur. Phys. J. Plus* 135:883
4. Alberts B. 1998. The cell as a collection of protein machines: preparing the next generation of molecular biologists. *Cell* 92(3):291–94
5. Amit DJ, Amit DJ. 1989. *Modeling Brain Function: The World of Attractor Neural Networks*. Cambridge University Press
6. Amit DJ, Gutfreund H, Sompolinsky H. 1985. Spin-glass models of neural networks. *Phys. Rev. A* 32(2):1007–18
7. Amit DJ, Gutfreund H, Sompolinsky H. 1985. Storing infinite numbers of patterns in a spin-glass model of neural networks. *Phys. Rev. Lett.* 55(14):1530
8. Barra A, Beccaria M, Fachechi A. 2018. A new mechanical approach to handle generalized Hopfield neural networks. *Neural Netw.* 106:205–22
9. Bialek W. 2012. *Biophysics: Searching for Principles*. Princeton University Press
10. Boukacem NE, Leary A, Thériault R, Gottlieb F, Mani M, François P. 2024. Waddington landscape for prototype learning in generalized Hopfield networks. *Phys. Rev. Res.* 6(3):033098
11. Braz Teixeira R, Carugno G, Neri I, Sartori P. 2024. Liquid Hopfield model: retrieval and localization in multicomponent liquid mixtures. *PNAS* 121(48):e2320504121
12. Chalk C, Buse S, Shrinivas K, Murugan A, Winfree E. 2024. Learning and inference in a lattice model of multicomponent condensates. Paper presented at the 30th International Conference on DNA Computing and Molecular Programming (DNA 30), Sep. 16–20
13. Chari T, Pachter L. 2023. The specious art of single-cell genomics. *PLOS Comput. Biol.* 19(8):e1011288
14. Demircigil M, Heusel J, Löwe M, Upgang S, Vermet F. 2017. On a model of associative memory with huge storage capacity. *J. Stat. Phys.* 168:288–99
15. Evans CG, O'Brien J, Winfree E, Murugan A. 2024. Pattern recognition in the nucleation kinetics of non-equilibrium self-assembly. *Nature* 625(7995):500–7
16. Fard AT, Srihari S, Mar JC, Ragan MA. 2016. Not just a colourful metaphor: modelling the landscape of cellular development using Hopfield networks. *NPJ Syst. Biol. Appl.* 2(1):16001

17. Ferrell JE. 2012. Bistability, bifurcations, and Waddington's epigenetic landscape. *Curr. Biol.* 22(11):R458–66
18. Gardner E. 1987. Multiconnected neural network models. *J. Phys. A Math. Gen.* 20(11):3453–64
19. Guo J, Zheng J. 2017. HopLand: single-cell pseudotime recovery using continuous Hopfield network-based modeling of Waddington's epigenetic landscape. *Bioinformatics* 33(14):i102–9
20. Haddow A. 1945. *What Is Life? The Physical Aspect of the Living Cell*. By Erwin Schrödinger. Cambridge: University Press. 1944, viii+ 91 pages. *Cancer Res.* 5(11):670–72
21. Herriges MJ, Yampolskaya M, Thapa BR, Lindstrom-Vautrin J, Wang F, et al. 2023. Durable alveolar engraftment of PSC-derived lung epithelial cells into immunocompetent mice. *Cell Stem Cell* 30(9):1217–34
22. Hopfield JJ. 1982. Neural networks and physical systems with emergent collective computational abilities. *PNAS* 79(8):2554–58
23. Hopfield JJ. 1984. Neurons with graded response have collective computational properties like those of two-state neurons. *PNAS* 81(10):3088–92
24. Hopfield JJ. 2010. Neurodynamics of mental exploration. *PNAS* 107(4):1648–53
25. Hopfield JJ. 2014. Whatever happened to solid state physics? *Annu. Rev. Condens. Matter Phys.* 5:1–13
26. Hopfield JJ, Feinstein DI, Palmer RG. 1983. 'Unlearning' has a stabilizing effect in collective memories. *Nature* 304(5922):158–59
27. Hotelling H. 1933. Analysis of a complex of statistical variables into principal components. *J. Educ. Psychol.* 24(6):417–41
28. Huang S. 2012. The molecular and mathematical basis of Waddington's epigenetic landscape: a framework for post-Darwinian biology? *Bioessays* 34(2):149–57
29. Ikonomidou L, Herriges MJ, Lewandowski SL, Marsland R III, Villacorta-Martin C, et al. 2020. The in vivo genetic program of murine primordial lung epithelial progenitors. *Nat. Commun.* 11(1):635
30. Kalaj S, Lauditi C, Perugini G, Lucibello C, Malatesta EM, Negri M. 2025. Random features Hopfield networks generalize retrieval to previously unseen examples. *Phys. A Stat. Mech. Appl.* 678:130946
31. Kanter I, Sompolinsky H. 1987. Associative recall of memory without errors. *Phys. Rev. A* 35(1):380–92
32. Karin O. 2024. EnhancerNet: a predictive model of cell identity dynamics through enhancer selection. *Development* 151(19):dev202997
33. Krotov D. 2023. A new frontier for Hopfield networks. *Nat. Rev. Phys.* 5(7):366–67
34. Krotov D, Hopfield J. 2021. Large associative memory problem in neurobiology and machine learning. Preprint, arXiv:2008.06996v3 [q-bio.NC]
35. Krotov D, Hopfield JJ. 2016. Dense associative memory for pattern recognition. *Adv. Neural Inform. Proc. Syst.* 29:1172–80
36. Lang AH, Li H, Collins JJ, Mehta P. 2014. Epigenetic landscapes explain partially reprogrammed cells and identify key reprogramming genes. *PLOS Comput. Biol.* 10(8):e1003734
37. Lucibello C, Mézard M. 2024. The exponential capacity of dense associative memories. *Phys. Rev. Lett.* 132:077301
38. Maetschke SR, Ragan MA. 2014. Characterizing cancer subtypes as attractors of Hopfield networks. *Bioinformatics* 30(9):1273–79
39. McInnes L, Healy J, Melville J. 2020. UMAP: uniform manifold approximation and projection for dimension reduction. Preprint, arXiv:1802.03426v3 [stat.ML]
40. Mehta P, Bukov M, Wang CH, Day AG, Richardson C, et al. 2019. A high-bias, low-variance introduction to machine learning for physicists. *Phys. Rep.* 810:1–124
41. Mézard M. 2017. Mean-field message-passing equations in the Hopfield model and its generalizations. *Phys. Rev. E* 95(2):022117
42. Moris N, Pina C, Arias AM. 2016. Transition states and cell fate decisions in epigenetic landscapes. *Nat. Rev. Genet.* 17(11):693–703
43. Murugan A, Zeravcic Z, Brenner MP, Leibler S. 2015. Multifarious assembly mixtures: systems allowing retrieval of diverse stored structures. *PNAS* 112(1):54–59
44. Negri M, Lauditi C, Perugini G, Lucibello C, Malatesta E. 2023. Storage and learning phase transitions in the random-features Hopfield model. *Phys. Rev. Lett.* 131(25):257301



45. Nelson PC, Nelson P. 2004. *Biological Physics*. WH Freeman
46. Owen JA, Osmanović D, Mirny L. 2023. Design principles of 3D epigenetic memory systems. *Science* 382(6672):eadg3053
47. Personnaz L, Guyon I, Dreyfus G. 1985. Information storage and retrieval in spin-glass like neural networks. *J. Phys. Lett.* 46(8):359–65
48. Phillips R, Kondev J, Theriot J, Garcia HG. 2012. *Physical Biology of the Cell*. Garland Science
49. Phillips R, Quake SR. 2006. The biological frontier of physics. *Phys. Today* 59(5):38–43
50. Pusuluri ST, Lang AH, Mehta P, Castillo HE. 2017. Cellular reprogramming dynamics follow a simple 1D reaction coordinate. *Phys. Biol.* 15(1):016001
51. Raju A, Siggia ED. 2023. A geometrical perspective on development. *Dev. Growth Differ.* 65(5):245–54
52. Ramsauer H, Schäfl B, Lehner J, Seidl P, Widrich M, et al. 2021. Hopfield networks is all you need. Preprint, arXiv:2008.02217v3 [cs.NE]
53. Rand DA, Raju A, Sáez M, Corson F, Siggia ED. 2021. Geometry of gene regulatory dynamics. *PNAS* 118(38):e2109729118
54. Serricchio L, Bocchi D, Chilin C, Marino R, Negri M, et al. 2025. Daydreaming Hopfield networks and their surprising effectiveness on correlated data. *Neural Netw.* 186:107216
55. Smart M, Zilman A. 2023. Emergent properties of collective gene-expression patterns in multicellular systems. *Cell Rep. Phys. Sci.* 4(2):101247
56. Smolensky P. 1986. Information processing in dynamical systems: foundations of harmony theory. In *Parallel Distributed Processing*, Vol. 1: *Explorations in the Microstructure of Cognition: Foundations*, ed. DE Rumelhart, JL McClelland, PDP Research Group. MIT Press
57. Szedlak A, Sims S, Smith N, Paternostro G, Piermarocchi C. 2017. Cell cycle time series gene expression data encoded as cyclic attractors in Hopfield systems. *PLOS Comput. Biol.* 13(11):e1005849
58. Taherian Fard A, Ragan MA. 2017. Modeling the attractor landscape of disease progression: a network-based approach. *Front. Genet.* 8:48
59. Van der Maaten L, Hinton G. 2008. Visualizing data using t-SNE. *J. Mach. Learn. Res.* 9(11):2579–605
60. Van Hemmen J, Ioffe L, Kühn R, Vaas M. 1990. Increasing the efficiency of a neural network through unlearning. *Phys. A Stat. Mech. Appl.* 163(1):386–92
61. Waddington CH. 2014. *The Strategy of the Genes*. Routledge
62. Wang F, Ting C, Riemondy KA, Douglas M, Foster K, et al. 2023. Regulation of epithelial transitional states in murine and human pulmonary fibrosis. *J. Clin. Investig.* 133(22):e165612
63. Wang J, Zhang K, Xu L, Wang E. 2011. Quantifying the Waddington landscape and biological paths for development and differentiation. *PNAS* 108(20):8257–62
64. Whittington JC, McCaffary D, Bakermans JJ, Behrens TE. 2022. How to build a cognitive map. *Nat. Neurosci.* 25(10):1257–72
65. Whittington JC, Muller TH, Mark S, Chen G, Barry C, et al. 2020. The Tolman-Eichenbaum machine: unifying space and relational memory through generalization in the hippocampal formation. *Cell* 183(5):1249–63
66. Whittington JC, Warren J, Behrens TE. 2022. Relating transformers to models and neural representations of the hippocampal formation. Preprint, arXiv:2112.04035v2 [cs.NE]
67. Yampolskaya M, Herriges MJ, Ikonou L, Kotton DN, Mehta P. 2023. scTOP: physics-inspired order parameters for cellular identification and visualization. *Development* 150(21):dev201873
68. Yampolskaya M, Ikonou L, Mehta P. 2025. Finding signatures of low-dimensional geometric landscapes in high-dimensional cell fate transitions. Preprint, bioRxiv. <https://www.biorxiv.org/content/10.1101/2025.05.27.656406v2.full>
69. Zhong W, Schwab DJ, Murugan A. 2017. Associative pattern recognition through macro-molecular self-assembly. *J. Stat. Phys.* 167:806–26