SIMPLICIAL HOPFIELD NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Hopfield networks are artificial neural networks which store memory patterns on the states of their neurons by choosing recurrent connection weights and update rules such that the energy landscape of the network forms attractors around the memories. How many stable, sufficiently-attracting memory patterns can we store in such a network using N neurons? The answer depends on the choice of weights and update rule. Inspired by setwise connectivity in biology, we extend Hopfield networks by adding setwise connections and embedding these connections in a simplicial complex. Simplicial complexes are higher dimensional analogues of graphs which naturally represent collections of pairwise and setwise relationships. We show that our simplicial Hopfield networks increase memory storage capacity. Surprisingly, even when connections are limited to a small random subset of equivalent size to an all-pairwise network, our networks still outperform their pairwise counterparts. Such scenarios include non-trivial simplicial topology. We also test analogous modern continuous Hopfield networks, offering a potentially promising avenue for improving the attention mechanism in Transformer models.

1 INTRODUCTION

Hopfield networks (Hopfield, 1982)[1] store memory patterns in the weights of connections between neurons. In the case of pairwise connections, these weights translate to the synaptic strength between pairs of neurons in biological neural networks. In such a Hopfield network with N neurons, there will be $\binom{N}{2}$ of these pairwise connections, forming a complete graph. Each edge is weighted by a procedure which considers P memory patterns and which, based on these patterns, seeks to minimise a defined energy function such that the network's dynamics are attracted to and ideally exactly settles in the memory pattern which is nearest to the current states of the neurons. The network therefore acts as a content addressable memory – given a partial or noise-corrupted memory, the network can update its states through recurrent dynamics to retrieve the full memory. Since its introduction, the Hopfield network has been extended and studied widely by neuroscientists (Griniasty et al., 1993; Schneidman et al., 2006; Sridhar et al., 2021; Burns et al., 2021), physicists (Amit et al., 1985; Agliari et al., 2013; Leonetti et al., 2021), and computer scientists (Widrich et al., 2020a;b; Millidge et al., 2022). Of particular interest for the machine learning community is the recent development of continuous modern Hopfield networks and their equivalence to the attention mechanism (Ramsauer et al., 2021).

An early (Amit et al., 1985; McEliece et al., 1987) and ongoing (Hillar & Tran, 2018) theme in the study of Hopfield networks has been their memory storage capacity, i.e., determining the number of memory patterns which can be reliably stored and later recalled via the dynamics. Traditional Hopfield networks with binary neuron states, in the limit of $N \to \infty$ and $P \to \infty$, maintain associative memories for up to approximately $0.14N$ patterns (Amit et al., 1985; McEliece et al., 1987), and fewer if the patterns are statistically or spatially correlated (Löwe, 1998). However, by a clever reformulation of the update rule based on the network energy, this capacity can be improved to N^{d-1}, where $d \geq 2$, and even further to $2^{N/2}$ (Krotov & Hopfield, 2016; Demircigil et al., 2017). Networks using these types of energy-based update rules are called modern Hopfield networks. Ramsauer et al. (2021) extended modern Hopfield networks to allow neurons to take continuous values – where previously they took only binary values – and this resulted in the discovery of their connection to attention. However, to the best of our knowledge, these modern Hopfield networks have

[1]Little (1974) introduced a similar Ising-based model for associative memory, prior to Hopfield (1982), however much of the research literature refers to and is more proximally inspired by Hopfield (1982).

not been extended further to include setwise connections between neurons, as has been studied and shown to improve memory capacity in traditional Hopfield networks (Peretto & Niez, 1986; Lee et al., 1986; Baldi & Venkatesh, 1987; Newman, 1988). Indeed, Krotov & Hopfield (2016), who introduced modern Hopfield networks, make a mathematical analogy between their energy-based update rule and setwise connections given their energy-based update rule can be interpreted as allowing individual pairs of pre- and post-synaptic neurons to make multiple synapses with each other – making pairwise connections mathematically as strong as equivalently-ordered setwise connections [2]. Demircigil et al. (2017) later proved this analogy to be accurate in terms of memory capacity. By adding setwise connections to modern Hopfield networks, we essentially allow all connection (pairwise and higher) to increase their strength – following the same interpretation, this can be thought of as allowing both pairwise and setwise connections between all neurons, any of which may be amplified.

Functionally, setwise connections appear in abundance in biological neural networks. What's more, these setwise interactions often modulate and interact with one another in highly complex and nonlinear fashions, adding to their potential computational expressiveness. We discuss these biological cases in Appendix A.1. There are many contemporary models in deep learning which implicitly model particular types of setwise interactions (Jayakumar et al., 2020). To explicitly model such interactions, we have multiple options. For reasons we discuss in Appendix A.2, we choose to model our setwise connections using a simplicial complex.

We therefore develop and study *Simplicial Hopfield Networks*. We weight the simplices of the simplicial complex to store memory patterns and generalise the energy functions and update rules of traditional and modern Hopfield networks. Our main contributions are:

- *We introduce extensions of various Hopfield networks with setwise connections.* In addition to generalising modern Hopfield networks to include setwise connections based on an underlying simplicial structure, we also study whether the topological features of the underlying structure influences performance.
- *We prove and discuss higher memory capacity in the general case of our networks.* For the fully-connected simplicial Hopfield network, we prove a larger memory capacity than previously shown by Newman (1988); Demircigil et al. (2017) for higher-degree Hopfield networks.
- *We empirically show improved performance under parameter constraints.* By constraining the total number of connections to that of pairwise Hopfield networks with a mixture of pairwise and setwise connections, we show our networks retain a surprising amount of improved performance over pairwise networks but without the cost of a larger number of parameters, and is robust to topological differences due to random choices of connections.

2 Simplicial Hopfield networks

2.1 Simplicial complexes

Simplicial complexes are mathematical objects which naturally represent collections of setwise relationships. Here we use the combinatorial form, called an *abstract simplicial complex*. Although, to build intuition and visualise the simplicial complex, we also refer to their geometric realisations.

Definition 2.1. Let K be a subset of $2^{[N]}$. The subset K is an abstract simplicial complex if for any $\sigma \in K$, the condition $\rho \subseteq \sigma$ gives $\rho \in K$, for any $\rho \subseteq \sigma$.

In other words, an abstract simplicial complex K is a collection of finite sets closed under taking subsets. A member of K is called a *simplex* σ. A k–dimensional simplex (or k–simplex) has cardinality $k+1$ and $k+1$ faces which are $(k-1)$–simplices (obtained by omitting one element from σ). If a simplex σ is a face of another simplex τ, we say that τ is a *coface* of σ. We denote the set of all k–simplices in K as K_k.

Geometrically, for $n = 0, 1, 2, 3$, a k–simplex is, respectively, a point, line segment, triangle, and tetrahedron. Therefore, one may visualise a simplicial complex as being constructed by gluing

[2] Work by Horn, D. & Usher, M. (1988) study almost the same system but with an slight modification to the traditional update rule, whereas Krotov & Hopfield (2016) use their modern, energy-based update rule.

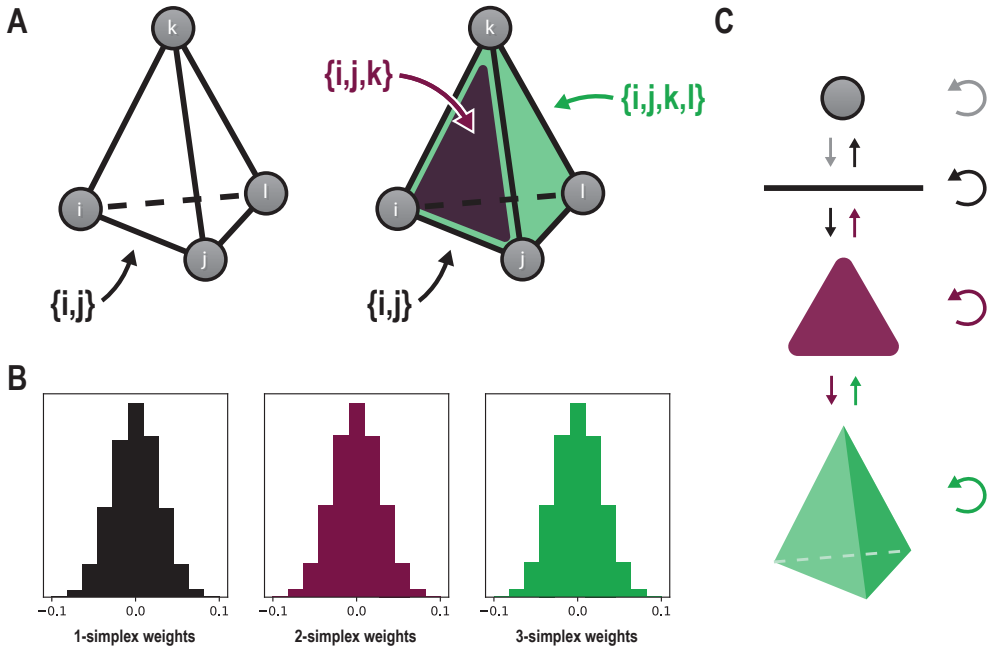

Figure 1: **A.** Comparative illustrations of connections in a pairwise Hopfield network (left) and a simplicial Hopfield network (right) with $N = 4$. In a simplicial Hopfield network, $\sigma = \{i, j\}$ is an edge (1–simplex), $\sigma = \{i, j, k\}$ is a triangle (2–simplex), $\sigma = \{i, j, k, l\}$ is a tetrahedron (3–simplex), and so on. **B.** Relative connection weight distributions of 1–, 2–, and 3–simplices in a simplicial Hopfield network. **C.** Illustration of the hierarchical relationship between simplices up to 3–simplices, with arrows indicating potential simplicial sources of weight modulation.

together simplices such that every finite set of vertices in K form the vertices of at most one higher-dimensional simplex. This structure makes it possible to associate every setwise relationship uniquely with a k–simplex identified by its elements, which in our case are neurons (see Figure 1A). Simplices in K which are contained in no higher dimensional simplices, i.e., they have no cofaces, are called the *facets* of K. The dimension of K, $\dim(K)$, is the dimension of its largest facet. A simplicial complex K wherein all facets are of the dimension k is called a pure simplicial k-complex.

2.2 Model

A network of N neurons is modelled by N spins. Let K be a simplicial complex on N vertices. In the binary neuron case, $S_j = \pm 1$. Given a set of neurons σ (which contains the neuron i and is a unique $(|\sigma| - 1)$–simplex in K), $w(\sigma)$ is the associated simplicial weight and S_σ the product of their spins. Spin configurations correspond to patterns of neural firing, with dynamics governed by a defined energy.

The traditional model is defined by the energy and weight functions

$$E = -\sum_{\sigma \in K} w(\sigma) S_\sigma \qquad\qquad w(\sigma) = \frac{1}{N} \sum_{\mu=1}^{P} \xi_\sigma^\mu, \tag{1}$$

with ξ_i^μ $(= \pm 1)$ static, independent, random variables being the P binary memory patterns stored in the simplicial weights. Similar for spins, ξ_σ^μ is the product of the random variables for the set of neurons σ for the pattern μ. Figure 1B shows examples of the resulting Gaussian distributions of weights at each simplicial dimension. We use these weights to update the state of a neuron i by

applying the traditional Hopfield update rule

$$\dot{S}_i = \Theta\left(\sum_{\sigma\in K} w(\sigma) S_{\sigma\setminus i}\right) \qquad \Theta(x) = \begin{cases} 1 & \text{if } x \geq 0 \\ -1 & \text{if } x < 0 \end{cases}. \tag{2}$$

When K is a pure simplicial 1-complex, this becomes the traditional pairwise Hopfield network (Hopfield, 1982).

In the modern Hopfield case, the energy function and update rule are

$$E = -\sum_{\mu=1}^{P} \sum_{\sigma\in K} F(\xi^\mu_\sigma S_\sigma) \tag{3}$$

$$\dot{S}_i = \text{sgn}\left[\sum_{\mu=1}^{P}\left(F(1\cdot\xi^\mu_i + \sum_{\sigma\in K}\xi^\mu_{\sigma\setminus i} S_{\sigma\setminus i}) - F(-1\cdot\xi^\mu_i + \sum_{\sigma\in K}\xi^\mu_{\sigma\setminus i} S_{\sigma\setminus i})\right)\right], \tag{4}$$

where the function F can be chosen to be of a polynomial $F(x) = x^n$ or exponential $F(x) = e^x$ form. When K is a pure simplicial 1-complex, this becomes the modern pairwise Hopfield network (Krotov & Hopfield, 2016).

In the continuous modern Hopfield case, spins and patterns take real values $S_j, \xi^\mu_j \in \mathbb{R}$. Patterns are arranged in a matrix $\Xi = (\xi^1, ..., \xi^P)$ and we define the *log-sum-exp function* (lse) for $T^{-1} > 0$ as

$$\text{lse}(T^{-1}, \Xi^T S, K) = \frac{1}{T^{-1}} \log\left(\sum_{\mu=1}^{P}\sum_{\sigma\in K} \exp(T^{-1}\xi^\mu_\sigma S_\sigma)\right). \tag{5}$$

The energy function is

$$E = -\text{lse}(T^{-1}, \Xi^T S, K) + \frac{1}{2}S^T S + \frac{1}{T^{-1}}\log N + \frac{1}{2}M^2. \tag{6}$$

For each simplex $\sigma \in K$, we denote the submatrix of the patterns stored on that simplex as Ξ_σ (which has dimensions $P \times \sigma$). Using the dot product to measure the similarity between the patterns and spins, the update rule is

$$\dot{S} = \text{softmax}\left(T\sum_{\sigma\in K}\left(\Xi^T_\sigma \overrightarrow{S_\sigma}\right)\right)\Xi. \tag{7}$$

In practice, however, the dot product has been found to under-perform in modern continuous Hopfield networks compared to Euclidean or Manhattan distances (Millidge et al., 2022). Transformer models in natural language tasks have also seen performance improvements by replacing the dot product with cosine similarity (Henry et al., 2020), again a measure with a more geometric flavour. However, these similarity measures generalise distances between pairs of elements rather than 'sets' of elements.

We therefore define higher-dimensional geometric similarity measures, *cumulative Euclidean distance (ced)* and *Cayley–Menger distance (cmd)*. Let d_ρ be the (Euclidean or Manhattan) distance between pattern ξ^μ_ρ and spins S_ρ for pattern μ and spins $\rho \subset \sigma$. Let K^σ_1 be the subset of K such that all elements in K^σ_1 are 1–simplex faces of σ. We define the cumulative Euclidean distance as

$$\text{ced}(\xi^\mu_\sigma, S_\sigma) = \sqrt{\sum_{\rho\in K^\sigma_1}(d_\rho)^2}. \tag{8}$$

We define $\text{cmd}(\xi^\mu_\sigma, S_\sigma)$ as the Cayley–Menger determinant of all $\rho \in K^\sigma_1$, with distances set as d_ρ.

Mixed diluted networks. A computational concern in the above models is the number of unique possible k–simplices is $\binom{N}{k+1}\binom{N}{k+1}$, e.g., with $N = 100$ there will be approximately 9.89×10^{28} possible 50–simplices, compared to just $4,950$ edges (1–simplices) found in a pairwise Hopfield network. If we allow all possible simplices for a simplicial Hopfield network with N neurons, the total number of simplices will be $\sum_{d=2}^{N}\binom{N}{d}$. Simultaneously, there is also a biological concern that high-dimensional setwise connections may be sparse and appear less frequently than low-dimensional connections.

We therefore evaluate classes of simplicial Hopfield networks which are low-dimensional, i.e., $\dim(K)$ is small, and where the total number of weighted simplices is not greater than those normally found in a pairwise Hopfield network, i.e., the number of non-zero weights is $\binom{N}{2}$. We randomly choose weights to be non-zero, with each weight of a dimension having an equal probability and according to Table 1. Such random networks have previously been studied in the traditional pairwise case as 'diluted networks' (Treves & Amit, 1988; Bovier & Gayrard, 1993a;b; Löwe & Vermet, 2011). Here we study *mixed diluted networks*, since we use a mixture of connections of different degrees. We believe we are also the first to study such networks beyond pairwise connections, as well as in modern and continuous cases.

Homology. Different collections of simplices in a simplicial complex can result in different Euler characteristics (a homotopy invariant property). Table 1 shows this from a 'functional' perspective via counting only the simplices with non-zero weights. However, even when using the same number of 1– and 2–simplices, the choices of which vertices they contain can be different according due to randomness. Therefore, the topologies of each network may vary (and so too may their subsequent dynamics and performance). One well-studied and often important topological property in the context of simplicial complexes, homology, counts the number of 'holes' in each dimension. In the 0th dimension, this is the number of connected components; in the 1st dimension, this is the number of triangles formed by edges which don't also have a 2–simplex 'filling in' the interior surface of that triangle; in the 2nd dimension, this is the number of tetrahedra formed by triangles which don't also have a 3–simplex 'filling in' the interior volume of that tetrahedron; and so on. The exact number of these holes in dimension k can be calculated by the k–th Betti number, β_k (see Appendix A.3). We therefore calculate these for our networks to study the relationship between the networks' homology and performance.

2.3 Theoretical memory capacity

Mixed networks. Much is already known about the theoretical memory capacity of various Hopfield networks, including those with explicit (Newman, 1988) or implicit (Demircigil et al., 2017) setwise connectivity. However, we wish to point out a somewhat underappreciated relationship between memory capacity and the explicit or implicit number of network connections – which, in the fully-connected network, is determined by the degree of the connections (a proof is provided in Appendix A.4).

Corollary 2.2 (Memory capacity is proportional to the number of network connections)**.** *If the connection weights in a Hopfield network are symmetric, then the order of the network's memory capacity is proportional to the number of its connections.*

What happens when there are connections between the same neurons at multiple degrees, i.e., what we call a mixed network? To the best of our knowledge, the theoretical memory capacity of such networks has not been well-studied. However, we found one classical study by Dreyfus et al. (1987) which showed, numerically, adding triplet connections to a pairwise model improved attractivity and memory capacity. Most prior formal studies have only considered connections at single higher degrees (Newman, 1988; Bengtsson, 1990). Although, higher order neural networks have historically considered such mixtures of interactions on different degrees simultaneously (Zhang, 2012), but as regular neural networks (e.g., feed-forward networks), not Hopfield networks. Higher order Boltzmann machines (HOBMs) (Sejnowski, 1986) have also been studied with mixed connections (Amari et al., 1992; Leisink & Kappen, 1999)[3]. However, HOBMs are unlike Hopfield networks in that they typically have hidden units, are trained differently, and have stochastic neural activations (despite this, there are equivalences (Leonelli et al., 2021; Marullo & Agliari, 2021; Smart & Zilman, 2021)). Modern Hopfield networks also include an implicit mixture of connections of different degrees[4] (but – and see Theorem 1 of Demircigil et al. (2017), which remains unproven – the mixture is unbalanced and not particularly natural, especially for $F(x) = x^n$ when n is odd). Therefore, we include the following result demonstrating fixed points, large basins of attraction (i.e., convergence)

[3]HOBMs also suffer the same problem as we face here, one of having many high-order parameters between the neurons to keep a track of. Possibly a factoring trick like in Memisevic & Hinton (2010) for HOBMs could helpful in simplicial Hopfield networks.

[4]Recall that $\left(\sum_i a_i\right)^b = \sum_i a_i^{b+1}$.

to those fixed points in mixed networks, and memory capacity which is linear in the number of fully-connected degrees of connections (a proof is provided in Appendix A.4).

Lemma 2.3 (Fully-connected mixed Hopfield networks)**.** *A fully-connected mixed Hopfield network based on a pure D–simplicial complex with N neurons and P patterns has, when $N \to \infty$ and P is finite: (i) fixed point attractors at the memory patterns; and (ii) dynamic convergence towards the fixed point attractors within a finite Hamming distance δ. When $P \to \infty$ with $N \to \infty$, the network has capacity to store up to $(\sum_{d=1}^{D} N^d)/(2\ ln\ N)$ memory patterns (with small retrieval errors) or $(\sum_{d=1}^{D} N^d)/(4\ ln\ N)$ (without retrieval errors).*

This naturally comports with Theorem 2 from Demircigil et al. (2017), except here we show an increased capacity in the mixed network, courtesy of Corollary 2.2.

Mixed diluted networks. As mentioned earlier, full setwise connectivity is not necessarily tractable nor realistic. Löwe & Vermet (2011) show for pairwise diluted networks constructed as Erdös-Renyi graphs (constructed by including each possible edge on the vertex set with probability p) that the memory capacity is proportional to pN. Crucial for this result is that the random graph must be asymptotically connected. This makes sense, given that if any vertex was disconnected its dynamics could never be influenced. Empirically, it does seem that a certain threshold of mean connectivity in pairwise random networks is crucial for attractor dynamics (Treves & Amit, 1988).

Remark 2.4. By a straightforward generalisation of Löwe & Vermet (2011)'s result, diluted networks constructed as pure Erdös-Renyi hypergraphs may store on the order of pN^{d-1} memory patterns, where d is the degree of the connections.

In the case of an unbounded number of allowable connections, Remark 2.4 would suggest picking as many higher-degree connections as possible when choosing between connections of lower or higher degrees in our mixed diluted networks. However, in the bounded case (our case), we are non-trivially changing the asymptotic behaviour in terms of connectivity and dynamics when we use a mixture of connection degrees. We also need to beware of asymmetries which may arise (Kanter, 1988). This makes the analysis of mixed diluted networks not particularly straightforward (also see section 4).

2.4 Numerical simulations and performance metrics

Given the large space of possible network settings, we restrict our experiments to network conditions listed in Table 1.

Table 1: List of network condition keys (top row), their number of non-negative weights for 1– and 2–simplices (second and third rows), and their 'functional' Euler characteristic (χ, bottom row). N is the number of neurons. $C = (N-1)N$. For simulation, the number of simplices at each dimension are rounded to the nearest integer.

	K1	R12(25%)	R12(50%)	R12(75%)	R2
1–simplices	$\binom{N}{2}$	$0.75\binom{N}{2}$	$0.50\binom{N}{2}$	$0.25\binom{N}{2}$	0
2–simplices	0	$0.25\binom{N}{2}$	$0.50\binom{N}{2}$	$0.75\binom{N}{2}$	$\binom{N}{2}$
χ	$N-(1/2)C$	$N-0.25C$	N	$N+0.25C$	$N+(1/2)C$

In our numerical simulations, we perform updates synchronously until E is non-decreasing or until a maximum number of steps is reached, whichever comes first. When a simulation concludes we compute the *overlap* (for binary patterns) or *mean squared error (MSE)* (for continuous patterns) of the final spin configuration with respect to all stored patterns using

$$m^\mu = \left| \frac{1}{N} \sum_{i=1}^{N} S_i \xi_i^\mu \right| \qquad \text{MSE}^\mu = \frac{1}{N} \sum_{i=1}^{N} (S_i - \xi_i^\mu)^2. \tag{9}$$

We say the network recalls (or attempts to recall) whichever pattern has the largest overlap (where $m^\mu = 1$ indicates perfect recall) or smallest MSE (where $\text{MSE}^\mu = 0$ indicates perfect recall).

3 NUMERICAL SIMULATIONS

3.1 BINARY MEMORY PATTERNS

After embedding random binary patterns, we tested starting the network in random initial states and recorded the final overlap of the closest pattern. Table 2 shows the final overlaps for traditional simplicial Hopfield networks ($N = 100$). Our simplicial Hopfield networks – especially K12 (75%) – significantly outperform the pairwise Hopfield networks (K1). In fact, the K12 (75%) model performs as well at $0.3N$ patterns as the the pairwise network performs on $0.05N$ patterns, a six-fold increase in the number of patterns and more than double the theoretical capacity of the pairwise network, $\sim 0.14N$ (Amit et al., 1985; McEliece et al., 1987). Surprisingly, Table 3 shows homology accounts for very little of the variance in network performance.

Table 2: Mean $\pm$ standard deviation of overlap distributions ($n = 100$) from traditional simplicial Hopfield networks with varying numbers (top row) of random binary patterns. K1 is the traditional pairwise Hopfield network. K12 (75%) significantly outperforms K1 at all tested levels (one-way t-tests $p < 10^{-11}$, $F > 50.13$). At all pattern loadings, a one-way ANOVA showed significant variance between the networks ($p < 10^{-20}$, $F > 26.35$). Box and whisker plots shown in Figure 6.

No. patterns	$0.05N$	$0.1N$	$0.15N$	$0.2N$	$0.3N$
K1	0.87 ± 0.18	0.81 ± 0.16	0.66 ± 0.10	0.65 ± 0.10	0.59 ± 0.08
R12 (25%)	0.96 ± 0.10	0.94 ± 0.14	0.82 ± 0.20	0.71 ± 0.17	0.64 ± 0.13
R12 (50%)	0.98 ± 0.10	$\mathbf{0.99 \pm 0.03}$	0.97 ± 0.10	0.91 ± 0.15	0.76 ± 0.16
R12 (75%)	$\mathbf{1 \pm 0}$	$\mathbf{0.99 \pm 0.04}$	$\mathbf{0.99 \pm 0.05}$	$\mathbf{0.98 \pm 0.08}$	$\mathbf{0.87 \pm 0.16}$
R2	$\mathbf{1 \pm 0}$	$\mathbf{0.99 \pm 0.18}$	0.94 ± 0.18	0.74 ± 0.29	0.53 ± 0.23

3.2 CONTINUOUS MEMORY PATTERNS

Energy landscape. Using Equation 3 and given a set of patterns, a simplicial complex K, and an inverse temperature T^{-1}, we may calculate the energy of network states. To inspect changes in the energy landscapes of different network conditions, we used a small network with $N = 10$ and $P = 10$ random patterns. We performed principle component analysis (PCA) to create a low dimensional projection of the patterns. Then, we generated network states evenly spaced in a 10×10 grid which spanned the projected memory patterns in the first two dimensions of PCA space. We calculated each state's energy by transforming these points from PCA space back into the N–dimensional space, across the network conditions at $T^{-1} = 1, 2, 10$ (Figure 7). At $T^{-1} = 1$, differences between the network conditions' energy landscapes are very subtle. However, at $T^{-1} = 2$ and $T^{-1} = 10$, we see a clear change: those with more 2–simplices possess more sophisticated, pattern-separating landscapes.

Recall as a function of memory loading. We tested the performance of our simplicial Hopfield networks by embedding data from the MNIST (LeCun et al., 2010), CIFAR-10 (Krizhevsky, 2009), and Tiny ImageNet (Le & Yang, 2015) datasets as memories. We followed the protocol of Millidge et al. (2022) to test recall under increasing memory load as an indication of the networks' memory capacities. To embed the memories, we normalise the pixel values between 0 and 1, and treat them as continuous-valued neurons, e.g., for MNIST we have $N = 28 \times 28 = 784$ neurons. We initialise S as one of the memory patterns corrupted by Gaussian noise with variance 0.5. After allowing the network to settle in an energy minima, we measured the performance as the fraction of correctly recalled memories (over all tested memories) of the uncorrupted patterns, where 'correct recall' was defined as a sum of the squared difference being < 50. In all tests, we used $T^{-1} = 100$.

Figure 2 compares the performance of a modern pairwise architecture – K1 – with the best performing modern simplicial architecture – K12 (75%). The performance of the K1 networks is comparable to that shown in Millidge et al. (2022). However, K12 (75%) significantly outperforms K1 across all datasets. Since the MNIST dataset is relatively simple and K1 already performs very well, the performance improvement is small, albeit still significant. However, in the CIFAR-10 and Tiny ImageNet datasets, the performance improvements are more noticeable, with most distance functions seeing improvements of $\geq 10\%$.

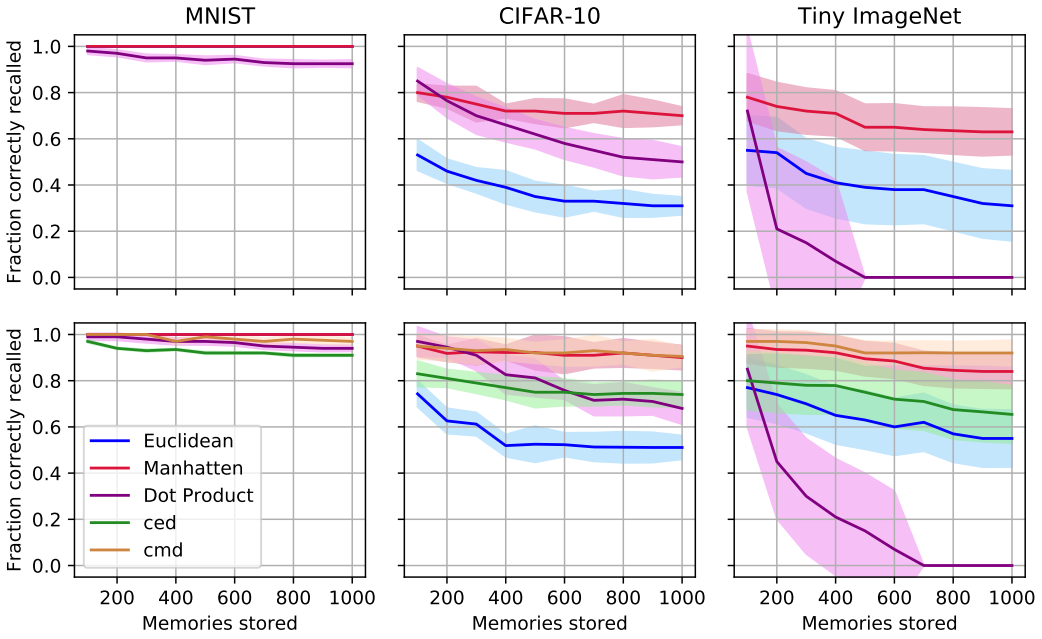

Figure 2: Recall (mean $\pm$ S.D. over 10 trials) as a function of memory loading using the MNIST, CIFAR-10, and Tiny ImageNet datasets, using different distance functions (see legend). Here we compare the performance of modern continuous pairwise networks (top row) and modern continuous simplicial networks (bottom row). The simplicial networks are R12 (75%) networks (see Table 1 for information). K12 (75%) significantly outperforms the pairwise network (K1) at all tested levels where there was not perfect recall (one-way t-tests $p < 10^{-9}$, $F > 16.01$). At all memory loadings, a one-way ANOVA showed significant variance between the networks ($p < 10^{-5}$, $F > 11.95$).

Also noticeable in the results for CIFAR-10 and Tiny ImageNet (see Figure 2) is the relatively high performance of the ced and cmd distance measures. Indeed, cmd performs as well or better than the Manhatten distance in our experiments. And both ced and cmd (along with the Euclidean and Manhatten distances) outperform the dot product in CIFAR-10 and Tiny ImageNet at high memory loadings. This further supports the intuition and results of Millidge et al. (2022), that more 'geometric' distances perform better as similarity measures in modern Hopfield networks.

4 Discussion

We have introduced a new class of Hopfield networks which generalises and extends traditional, modern, and continuous Hopfield networks. Our main finding is that mixed diluted networks can improve performance in terms of memory recall, even when there is no increase in the number of parameters. This improvement therefore comes from the topology rather than additional information in the form of parameters. We also show how distance measures of a more 'geometric flavour' can further improve performance in these networks. This simplicial framework (in diluted or undiluted forms) now opens up new avenues for researchers in neuroscience and machine learning. In neuroscience, we can now model how setwise connections, such as those provided by astrocytes, may improve memory function and may interact or form important topological structures to guide memory dynamics. In machine learning, such topological structures may now be utilised in improved attention mechanisms or Transformer models, such as in Ramsauer et al. (2021). At the intersection of these fields, we may now further study how the topology of networks in neuroscience and machine learning systems may correspond to one another and share functional characteristics, such as how the activity of 'pairwise' Transformer models have shown similarities to activities in auditory cortex (Millet et al., 2022). Could 'setwise' Transformer models correspond more closely? Or to a more diverse range of cell types? These and related questions are now open for exploration.

Convolution operations and higher-order neural networks. From the perspective of modern deep learning, considering higher order correlations between downstream inputs to a neuron is quite classical. For example, convolutional neural networks have incorporated specialised setwise operations since their inception (Fukushima, 1980; Lecun et al., 1998), and more general setwise relationships have also been introduced in higher-order neural networks (Pineda, 1987; Reid et al., 1989; Ghosh & Shun, 1992; Zhang, 2012). Although our setwise connections are not explicitly convolutional, they are in one notable sense conceptually similar: they collect information from a particular subset of neurons and only become active when those particular neurons are active in the right way. One of the main differences, however, is that – unlike typical convolution operations – we don't restrict the connection locations to some particular locations or arrangements within the input space. Our results therefore suggest that, in some cases, replacing regular feedforward connections with random convolutions may offer improved performance in some circumstances.

Improvements and extensions. Our study focusses on random choices of weighted simplices. What if we choose more carefully? Indeed, it seems quite likely biological setwise connections are not random, and are almost certainly not randomly chosen to replace random pairwise connections.

In diluted traditional Hopfield networks, and in our simplicial weights at all dimensions, there are a non-vanishing but reducing number of weights with a value of zero when the number of patterns is even as $P \to N$, but no such weights when the number of patterns is odd. Weights with a value of zero don't contribute to the network dynamics, and this is a general mathematical quirk common to all Hopfield networks. We have shown that setwise connections (even in a limited random number) can improve performance in terms of memory storage. Therefore, in cases where an even number of patterns are being stored in a Hopfield network, it may be possible to replace non-contributing weights with contributing higher-dimensional weights. In the general case, where an odd number of patterns are being stored, it might actually be advantageous to add an additional nonsense pattern in order to replace non-contributing lower-dimensional weights with contributing higher-dimensional weights.

We believe it is natural to now study how potential online weight modulations (e.g., based on spectral theories) could offer interesting and new connections between Hopfield networks and areas such as geometric deep learning. Further, these modulations may be designed to be biologically-interpretable or be insprired by biological mechanisms. For instance, it may be possible to model associative and anti-Hebbian memory structures by strategically inserting inhibitory interactions between some simplices (Haga & Fukai, 2019; Burns et al., 2021), providing a form of higher-degree inhibitory (or disinhibitory) mechanism.

Further analytic studies. Our numerical results suggest diluted mixed networks have larger memory capacities than pairwise networks. In a fairly intuitive sense, this is not particularly surprising – we are adding degrees of freedom to the energy landscape, within which we may build more stable and nicely-behaved attractors. However, we have not yet proven this increased capacity analytically for the diluted case, only given some theoretical indications as to why this occurs and proven the undiluted case. We hypothesise it is possible to do so using generalised methods from replica-symmetric theory (Amit et al., 1985) or self-consistent signal-to-noise analysis (Shiino & Fukai, 1993), in combination with methods from structural Ramsey theory. The capacity for modern networks may be on the order of a double-exponential in the number of neurons (since, in the limit of $N \to \infty$, there is an exponential relationship in the number of multispin interactions on top of an exponential relationship in the number of intra-multispin interactions, i.e., both pair-spins and multi-spins can have higher degrees of attraction). This capacity, however, will likely scale nonlinearly with the choice of (random) dilution, e.g., there may be a steep drop in performance around a critical dilution range, likely where some important dynamical guarantees are lost due to an intolerably small number of connections of a particular order.

Even higher orders and diluted mixtures of setwise connections may also be studied. Such networks, per Lemma 2.3, will likely improve their performance as higher-degree connections are added. However, and as implied in section 2.3, the number and distribution of these connections may need to be careful chosen in highly diluted settings[5].

[5] At some level of dilution, it may also be preferable to generalise simplicial Hopfield networks as hypergraphic Hopfield networks, with or without orientation.

REFERENCES

Elena Agliari, Adriano Barra, Andrea De Antoni, and Andrea Galluzzi. Parallel retrieval of correlated patterns: From hopfield networks to boltzmann machines. *Neural Networks*, 38:52–63, 2013. ISSN 0893-6080. doi: https://doi.org/10.1016/j.neunet.2012.11.010. URL https://www.sciencedirect.com/science/article/pii/S0893608012003073.

S Amari, K Kurata, and H Nagaoka. Information geometry of boltzmann machines. *IEEE Trans Neural Netw*, 3(2):260–271, 1992.

Daniel J. Amit, Hanoch Gutfreund, and H. Sompolinsky. Storing infinite numbers of patterns in a spin-glass model of neural networks. *Phys. Rev. Lett.*, 55:1530–1533, Sep 1985. doi: 10.1103/PhysRevLett.55.1530. URL https://link.aps.org/doi/10.1103/PhysRevLett.55.1530.

Alfonso Araque, Vladimir Parpura, Rita P Sanzgiri, and Philip G Haydon. Tripartite synapses: glia, the unacknowledged partner. *Trends in neurosciences*, 22(5):208–215, 1999.

Gordon W. Arbuthnott and Jeff Wickens. Space, time and dopamine. *Trends in Neurosciences*, 30 (2):62–69, 2007. ISSN 0166-2236. doi: https://doi.org/10.1016/j.tins.2006.12.003. URL https://www.sciencedirect.com/science/article/pii/S0166223606002748.

Pierre Baldi and Santosh S. Venkatesh. Number of stable points for spin-glasses and neural networks of higher orders. *Phys. Rev. Lett.*, 58:913–916, Mar 1987. doi: 10.1103/PhysRevLett.58.913. URL https://link.aps.org/doi/10.1103/PhysRevLett.58.913.

Mats Bengtsson. *Higher Order Artificial Neural Networks*. Diane Publishing Company, 1990. ISBN 9780941375924. URL https://books.google.co.jp/books?id=FTlFeFg8HH4C.

Cristian Bodnar, Fabrizio Frasca, Nina Otter, Yuguang Wang, Pietro Liò, Guido F Montufar, and Michael Bronstein. Weisfeiler and lehman go cellular: Cw networks. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 2625–2640. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper/2021/file/157792e4abb490f99dbd738483e0d2d4-Paper.pdf.

Anton Bovier and Véronique Gayrard. Lower bounds on the memory capacity of the dilute hopfield model. In Nino Boccara, Eric Goles, Servet Martinez, and Pierre Picco (eds.), *Cellular Automata and Cooperative Systems*, pp. 55–66. Springer Netherlands, Dordrecht, 1993a. ISBN 978-94-011-1691-6. doi: 10.1007/978-94-011-1691-6_6. URL https://doi.org/10.1007/978-94-011-1691-6_6.

Anton Bovier and Véronique Gayrard. Rigorous results on the thermodynamics of the dilute hopfield model. *Journal of Statistical Physics*, 72(1):79–112, Jul 1993b. ISSN 1572-9613. doi: 10.1007/BF01048041. URL https://doi.org/10.1007/BF01048041.

Thomas F Burns, Tatsuya Haga, and Tomoki Fukai. Multiscale and extended retrieval of associative memory structures in a cortical model of local-global inhibition balance. *bioRxiv*, 2021. doi: 10.1101/2021.11.30.470555. URL https://www.biorxiv.org/content/early/2021/12/01/2021.11.30.470555.

György Buzsáki. Neural syntax: Cell assemblies, synapsembles, and readers. *Neuron*, 68(3):362–385, 2010. ISSN 0896-6273. doi: https://doi.org/10.1016/j.neuron.2010.09.023. URL https://www.sciencedirect.com/science/article/pii/S0896627310007658.

Yuhan Chen, Shengjun Wang, Claus C Hilgetag, and Changsong Zhou. Trade-off between multiple constraints enables simultaneous formation of modules and hubs in neural systems. *PLoS Comput. Biol.*, 9(3):e1002937, March 2013.

Peter H Chipman, Chi Chung Alan Fung, Alejandra Pazo Fernandez, Abhilash Sawant, Angelo Tedoldi, Atsushi Kawai, Sunita Ghimire Gautam, Mizuki Kurosawa, Manabu Abe, Kenji Sakimura, Tomoki Fukai, and Yukiko Goda. Astrocyte GluN2C NMDA receptors control basal synaptic strengths of hippocampal CA1 pyramidal neurons in the stratum radiatum. *Elife*, 10, October 2021.

Ana Covelo and Alfonso Araque. Neuronal activity determines distinct gliotransmitter release from a single astrocyte. *eLife*, 7:e32237, Jan 2018. ISSN 2050-084X. doi: 10.7554/eLife.32237. URL https://doi.org/10.7554/eLife.32237.

Mete Demircigil, Judith Heusel, Matthias Löwe, Sven Upgang, and Franck Vermet. On a model of associative memory with huge storage capacity. *Journal of Statistical Physics*, 168(2):288–299, Jul 2017. ISSN 1572-9613. doi: 10.1007/s10955-017-1806-y. URL https://doi.org/10.1007/s10955-017-1806-y.

Adi Doron, Alon Rubin, Aviya Benmelech-Chovav, Netai Benaim, Tom Carmi, Ron Refaeli, Nechama Novick, Tirzah Kreisel, Yaniv Ziv, and Inbal Goshen. Hippocampal astrocytes encode reward location. *Nature*, Aug 2022. ISSN 1476-4687. doi: 10.1038/s41586-022-05146-6. URL https://doi.org/10.1038/s41586-022-05146-6.

Gérard Dreyfus, Isabelle Guyon, Jean-Pierre Nadal, and Léon Personnaz. High order neural networks for efficient associative memory design. In D. Anderson (ed.), *Neural Information Processing Systems*, volume 0. American Institute of Physics, 1987. URL https://proceedings.neurips.cc/paper/1987/file/c7e1249ffc03eb9ded908c236bd1996d-Paper.pdf.

Stefania Ebli, Michaël Defferrard, and Gard Spreemann. Simplicial neural networks. In *NeurIPS 2020 Workshop on Topological Data Analysis and Beyond*, 2020. URL https://openreview.net/forum?id=nPCt39DVIfk.

Mária Ercsey-Ravasz, Nikola T. Markov, Camille Lamy, David C. Van Essen, Kenneth Knoblauch, Zoltán Toroczkai, and Henry Kennedy. A predictive network model of cerebral cortical connectivity based on a distance rule. *Neuron*, 80(1):184–197, 2013. ISSN 0896-6273. doi: https://doi.org/10.1016/j.neuron.2013.07.036. URL https://www.sciencedirect.com/science/article/pii/S0896627313006600.

Sarah J Etherington, Susan E Atkinson, Greg J Stuart, and Stephen R Williams. *Synaptic Integration*, chapter 1, pp. 1–15. Wiley, 2010. ISBN 9780470015902. doi: https://doi.org/10.1002/9780470015902.a0000208.pub2. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/9780470015902.a0000208.pub2.

Yifan Feng, Haoxuan You, Zizhao Zhang, Rongrong Ji, and Yue Gao. Hypergraph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 3558–3565, 2019.

Kunihiko Fukushima. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, 36(4):193–202, Apr 1980. ISSN 1432-0770. doi: 10.1007/BF00344251. URL https://doi.org/10.1007/BF00344251.

Joydeep Ghosh and Yoan Shun. Efficient higher-order neural networks for classification and function approximation. *International Journal of Neural Systems*, 03(04):323–350, 1992. doi: 10.1142/S0129065792000255. URL https://doi.org/10.1142/S0129065792000255.

Nace L. Golding, Nathan P. Staff, and Nelson Spruston. Dendritic spikes as a mechanism for cooperative long-term potentiation. *Nature*, 418(6895):326–331, Jul 2002. ISSN 1476-4687. doi: 10.1038/nature00854. URL https://doi.org/10.1038/nature00854.

I. S. Gradshteyn and I. M. Ryzhik. *Table of integrals, series, and products*. Elsevier/Academic Press, Amsterdam, seventh edition, 2007. ISBN 978-0-12-373637-6; 0-12-373637-4. Translated from the Russian, Translation edited and with a preface by Alan Jeffrey and Daniel Zwillinger, With one CD-ROM (Windows, Macintosh and UNIX).

M. Griniasty, M. V. Tsodyks, and Daniel J. Amit. Conversion of Temporal Correlations Between Stimuli to Spatial Correlations Between Attractors. *Neural Computation*, 5(1):1–17, 01 1993. ISSN 0899-7667. doi: 10.1162/neco.1993.5.1.1. URL https://doi.org/10.1162/neco.1993.5.1.1.

Zengcai V. Guo, Hidehiko K. Inagaki, Kayvon Daie, Shaul Druckmann, Charles R. Gerfen, and Karel Svoboda. Maintenance of persistent activity in a frontal thalamocortical loop. *Nature*, 545(7653):181–186, May 2017. ISSN 1476-4687. doi: 10.1038/nature22324. URL https://doi.org/10.1038/nature22324.

Tatsuya Haga and Tomoki Fukai. Extended temporal association memory by modulations of inhibitory circuits. *Phys. Rev. Lett.*, 123:078101, Aug 2019. doi: 10.1103/PhysRevLett.123.078101. URL https://link.aps.org/doi/10.1103/PhysRevLett.123.078101.

Mustafa Hajij, Kyle Istvan, and Ghada Zamzmi. Cell complex neural networks. In *NeurIPS 2020 Workshop on Topological Data Analysis and Beyond*, 2020. URL https://openreview.net/forum?id=6Tq18ySFpGU.

Eric Hart and Alexander C. Huk. Recurrent circuit dynamics underlie persistent activity in the macaque frontoparietal network. *eLife*, 9:e52460, May 2020. ISSN 2050-084X. doi: 10.7554/eLife.52460. URL https://doi.org/10.7554/eLife.52460. 32379044[pmid].

Alex Henry, Prudhvi Raj Dachapally, Shubham Shantaram Pawar, and Yuxuan Chen. Query-key normalization for transformers. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 4246–4253, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.379. URL https://aclanthology.org/2020.findings-emnlp.379.

Christopher J. Hillar and Ngoc M. Tran. Robust exponential memory in hopfield networks. *The Journal of Mathematical Neuroscience*, 8(1):1, Jan 2018. ISSN 2190-8567. doi: 10.1186/s13408-017-0056-2. URL https://doi.org/10.1186/s13408-017-0056-2.

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. doi: 10.1073/pnas.79.8.2554. URL https://www.pnas.org/doi/abs/10.1073/pnas.79.8.2554.

Horn, D. and Usher, M. Capacities of multiconnected memory models. *J. Phys. France*, 49(3):389–395, 1988. doi: 10.1051/jphys:01988004903038900. URL https://doi.org/10.1051/jphys:01988004903038900.

Siddhant M. Jayakumar, Wojciech M. Czarnecki, Jacob Menick, Jonathan Schwarz, Jack Rae, Simon Osindero, Yee Whye Teh, Tim Harley, and Razvan Pascanu. Multiplicative interactions and where to find them. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=rylnK6VtDH.

Marcus Kaiser and Claus C Hilgetag. Nonoptimal component placement, but short processing paths, due to long-distance projections in neural systems. *PLoS Comput. Biol.*, 2(7):e95, July 2006.

I. Kanter. Asymmetric neural networks with multispin interactions. *Phys. Rev. A*, 38:5972–5975, Dec 1988. doi: 10.1103/PhysRevA.38.5972. URL https://link.aps.org/doi/10.1103/PhysRevA.38.5972.

Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.

Dmitry Krotov and John J. Hopfield. Dense associative memory for pattern recognition. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS'16, pp. 1180–1188, Red Hook, NY, USA, 2016. Curran Associates Inc. ISBN 9781510838819.

Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.

Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.

Yann LeCun, Corinna Cortes, and Christopher Burges. MNIST handwritten digit database. http://yann.lecun.com/exdb/mnist/, 2010. URL http://yann.lecun.com/exdb/mnist/.

Y. C. Lee, Gary D. Doolen, H. H. Chen, G. Z. Sun, T. Maxwell, H. Y. Lee, and C. Lee Giles. Machine learning using higher order correlation networks. In *Proceedings of the "Evolution, Games, and Learning" Conference*, 1986.

Martijn A. R. Leisink and Hilbert J. Kappen. Learning in higher order boltzmann machines using linear response. *Neural Networks*, 13:2000, 1999.

Francesca Elisa Leonelli, Elena Agliari, Linda Albanese, and Adriano Barra. On the effective initialisation for restricted boltzmann machines via duality with hopfield model. *Neural Networks*, 143: 314–326, 2021. ISSN 0893-6080. doi: https://doi.org/10.1016/j.neunet.2021.06.017. URL https://www.sciencedirect.com/science/article/pii/S0893608021002495.

M. Leonetti, E. Hörmann, L. Leuzzi, G. Parisi, and G. Ruocco. Optical computation of a spin glass dynamics with tunable complexity. *Proceedings of the National Academy of Sciences*, 118(21): e2015207118, 2021. doi: 10.1073/pnas.2015207118. URL https://www.pnas.org/doi/abs/10.1073/pnas.2015207118.

Lek-Heng Lim. Hodge laplacians on graphs. *SIAM Review*, 62(3):685–715, 2020. doi: 10.1137/18M1223101. URL https://doi.org/10.1137/18M1223101.

W.A. Little. The existence of persistent states in the brain. *Mathematical Biosciences*, 19(1):101–120, 1974. ISSN 0025-5564. doi: https://doi.org/10.1016/0025-5564(74)90031-5. URL https://www.sciencedirect.com/science/article/pii/0025556474900315.

Matthias Löwe and Franck Vermet. The hopfield model on a sparse erdös-renyi graph. *Journal of Statistical Physics*, 143(1):205–214, Apr 2011. ISSN 1572-9613. doi: 10.1007/s10955-011-0167-1. URL https://doi.org/10.1007/s10955-011-0167-1.

Matthias Löwe. On the storage capacity of Hopfield models with correlated patterns. *The Annals of Applied Probability*, 8(4):1216 – 1250, 1998. doi: 10.1214/aoap/1028903378. URL https://doi.org/10.1214/aoap/1028903378.

Chiara Marullo and Elena Agliari. Boltzmann machines as generalized hopfield networks: A review of recent results and outlooks. *Entropy*, 23(1), 2021. ISSN 1099-4300. doi: 10.3390/e23010034. URL https://www.mdpi.com/1099-4300/23/1/34.

R. McEliece, E. Posner, E. Rodemich, and S. Venkatesh. The capacity of the hopfield associative memory. *IEEE Transactions on Information Theory*, 33(4):461–482, 1987. doi: 10.1109/TIT.1987.1057328.

Roland Memisevic and Geoffrey E. Hinton. Learning to Represent Spatial Transformations with Factored Higher-Order Boltzmann Machines. *Neural Computation*, 22(6):1473–1492, 06 2010. ISSN 0899-7667. doi: 10.1162/neco.2010.01-09-953. URL https://doi.org/10.1162/neco.2010.01-09-953.

Juliette Millet, Charlotte Caucheteux, Pierre Orhan, Yves Boubenec, Alexandre Gramfort, Ewan Dunbar, Christophe Pallier, and Jean-Remi King. Toward a realistic model of speech processing in the brain with self-supervised learning, 2022. URL https://arxiv.org/abs/2206.01685.

Beren Millidge, Tommaso Salvatori, Yuhang Song, Thomas Lukasiewicz, and Rafal Bogacz. Universal hopfield networks: A general framework for single-shot associative memory models, 2022. URL https://arxiv.org/abs/2202.04557.

Michael G Müller, Christos H Papadimitriou, Wolfgang Maass, and Robert Legenstein. A model for structured information representation in neural networks of the brain. *eNeuro*, 7(3):ENEURO.0533–19.2020, May 2020.

Charles M Newman. Memory capacity in neural network models: Rigorous lower bounds. *Neural Networks*, 1(3):223–238, 1988. ISSN 0893-6080. doi: https://doi.org/10.1016/0893-6080(88)90028-7. URL https://www.sciencedirect.com/science/article/pii/0893608088900287.

Christos H. Papadimitriou, Santosh S. Vempala, Daniel Mitropolsky, Michael Collins, and Wolfgang Maass. Brain computation by assemblies of neurons. *Proceedings of the National Academy of Sciences*, 117(25):14464–14472, 2020. doi: 10.1073/pnas.2001893117. URL https://www.pnas.org/doi/abs/10.1073/pnas.2001893117.

Gertrudis Perea, Marta Navarrete, and Alfonso Araque. Tripartite synapses: astrocytes process and control synaptic information. *Trends in neurosciences*, 32(8):421–431, 2009.

P Peretto and J J Niez. Long term memory storage capacity of multiconnected neural networks. *Biol. Cybern.*, 54(1):53–64, may 1986. ISSN 0340-1200. doi: 10.1007/BF00337115. URL https://doi.org/10.1007/BF00337115.

Giovanni Petri and Alain Barrat. Simplicial activity driven model. *Phys. Rev. Lett.*, 121:228301, Nov 2018. doi: 10.1103/PhysRevLett.121.228301. URL https://link.aps.org/doi/10.1103/PhysRevLett.121.228301.

Fernando Pineda. Generalization of back propagation to recurrent and higher order neural networks. In D. Anderson (ed.), *Neural Information Processing Systems*, volume 0. American Institute of Physics, 1987. URL https://proceedings.neurips.cc/paper/1987/file/735b90b4568125ed6c3f678819b6e058-Paper.pdf.

Hubert Ramsauer, Bernhard Schäfl, Johannes Lehner, Philipp Seidl, Michael Widrich, Lukas Gruber, Markus Holzleitner, Thomas Adler, David Kreil, Michael K Kopp, Günter Klambauer, Johannes Brandstetter, and Sepp Hochreiter. Hopfield networks is all you need. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=tL89RnzIiCd.

Max B. Reid, Lilly Spirkovska, and Ellen Ochoa. Rapid training of higher-order neural networks for invariant pattern recognition. In *In Proc. International Joint Conference on Neural Nets*, pp. 689–692, 1989.

T. Mitchell Roddenberry, Nicholas Glaze, and Santiago Segarra. Principled simplicial neural networks for trajectory prediction. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9020–9029. PMLR, 18–24 Jul 2021. URL https://proceedings.mlr.press/v139/roddenberry21a.html.

Fernando E. Rosas, Pedro A. M. Mediano, Andrea I. Luppi, Thomas F. Varley, Joseph T. Lizier, Sebastiano Stramaglia, Henrik J. Jensen, and Daniele Marinazzo. Disentangling high-order mechanisms and high-order behaviours in complex systems. *Nature Physics*, Mar 2022. ISSN 1745-2481. doi: 10.1038/s41567-022-01548-5. URL https://doi.org/10.1038/s41567-022-01548-5.

Dmitri A. Rusakov and Dimitri M. Kullmann. Extrasynaptic glutamate diffusion in the hippocampus: Ultrastructural constraints, uptake, and receptor activation. *Journal of Neuroscience*, 18(9): 3158–3170, 1998. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.18-09-03158.1998. URL https://www.jneurosci.org/content/18/9/3158.

Andrea Santoro, Federico Battiston, Giovanni Petri, and Enrico Amico. Unveiling the higher-order organization of multivariate time series, 2022. URL https://arxiv.org/abs/2203.10702.

Michael T Schaub, Austin R Benson, Paul Horn, Gabor Lippner, and Ali Jadbabaie. Random walks on simplicial complexes and the normalized hodge 1-laplacian. *SIAM Review*, 62(2):353–391, 2020.

Elad Schneidman, Michael J. Berry, Ronen Segev, and William Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012, Apr 2006. ISSN 1476-4687. doi: 10.1038/nature04701. URL https://doi.org/10.1038/nature04701.

Terrence J. Sejnowski. Higher-order boltzmann machines. *AIP Conference Proceedings*, 151(1): 398–403, 1986. doi: 10.1063/1.36246. URL https://aip.scitation.org/doi/abs/10.1063/1.36246.

Masatoshi Shiino and Tomoki Fukai. Self-consistent signal-to-noise analysis of the statistical behavior of analog neural networks and enhancement of the storage capacity. *Phys. Rev. E*, 48:867–897, Aug 1993. doi: 10.1103/PhysRevE.48.867. URL https://link.aps.org/doi/10.1103/PhysRevE.48.867.

Matthew Smart and Anton Zilman. On the mapping between hopfield networks and restricted boltzmann machines. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=RGJbergVIoO.

David I Spivak. Higher-dimensional models of networks. *arXiv preprint arXiv:0909.4314*, 2009.

Vivek H. Sridhar, Liang Li, Dan Gorbonos, Máté Nagy, Bianca R. Schell, Timothy Sorochkin, Nir S. Gov, and Iain D. Couzin. The geometry of decision-making in individuals and collectives. *Proceedings of the National Academy of Sciences*, 118(50):e2102157118, 2021. doi: 10.1073/pnas.2102157118. URL https://www.pnas.org/doi/abs/10.1073/pnas.2102157118.

A Treves and D J Amit. Metastable states in asymmetrically diluted hopfield networks. *Journal of Physics A: Mathematical and General*, 21(14):3155–3169, jul 1988. doi: 10.1088/0305-4470/21/14/016. URL https://doi.org/10.1088/0305-4470/21/14/016.

Michael Widrich, Bernhard Schäfl, Milena Pavlović, Hubert Ramsauer, Lukas Gruber, Markus Holzleitner, Johannes Brandstetter, Geir Kjetil Sandve, Victor Greiff, Sepp Hochreiter, and Günter Klambauer. Modern hopfield networks and attention for immune repertoire classification. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 18832–18845. Curran Associates, Inc., 2020a. URL https://proceedings.neurips.cc/paper/2020/file/da4902cb0bc38210839714ebdcf0efc3-Paper.pdf.

Michael Widrich, Bernhard Schäfl, Milena Pavlović, Hubert Ramsauer, Lukas Gruber, Markus Holzleitner, Johannes Brandstetter, Geir Kjetil Sandve, Victor Greiff, Sepp Hochreiter, and Günter Klambauer. Modern hopfield networks and attention for immune repertoire classification. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 18832–18845. Curran Associates, Inc., 2020b. URL https://proceedings.neurips.cc/paper/2020/file/da4902cb0bc38210839714ebdcf0efc3-Paper.pdf.

Chenxin Xu, Maosen Li, Zhenyang Ni, Ya Zhang, and Siheng Chen. Groupnet: Multiscale hypergraph neural networks for trajectory prediction with relational reasoning, 2022. URL https://arxiv.org/abs/2204.08770.

Ming Zhang. *Artificial higher order neural networks for modeling and simulation*. IGI Global, 2012. ISBN 9781466621756. doi: 10.4018/978-1-4666-2175-6.

A APPENDIX

A.1 SETWISE CONNECTIONS AND MODULATIONS ARE BOUNTIFUL IN BIOLOGY

Setwise connections are not limited to the case, as one might expect, of multiple synaptic contacts between pairs or sets of cells. Other examples (some of which are illustrated in Figure 3) include the wide spatial dispersion of certain neurotransmitters (Rusakov & Kullmann, 1998; Arbuthnott & Wickens, 2007), persistently-connected neuronal assembly structures or 'synapsembles' (Buzsáki, 2010; Papadimitriou et al., 2020), distributed persistent activity during activities such as motor planning and working memory (Guo et al., 2017; Hart & Huk, 2020), neuroglial modulations of neurotransmitter release probabilities across multiple neurons or synapses (Covelo & Araque, 2018; Chipman et al., 2021), 'tripartite' astrocyte–neuron synapses (Araque et al., 1999; Perea et al., 2009), astrocytic coding Doron et al. (2022), and during synaptic integration at the level of individual neurons (Golding et al., 2002; Etherington et al., 2010). Modulations of such connections are illustrated in Figure 4.

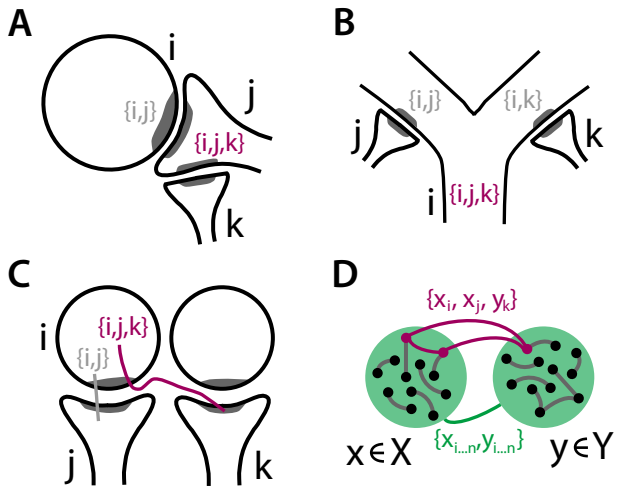

Figure 3: **A.** Multi-synaptic bouton. **B.** Dendritic integration. **C.** Extra-synaptic neurotransmitter diffusion. **D.** Functional connections between neural assemblies.

A.2 OPTIONS FOR MODELLING SETWISE CONNECTIVITY IN NEURAL NETWORKS AND WHY WE CHOOSE SIMPLICIAL COMPLEXES

In geometric and topological deep learning, recent advances have been realised by utilising higher dimensional analogues of graphs such as simplicial complexes (Ebli et al., 2020; Roddenberry et al., 2021), cell complexes (Hajij et al., 2020; Bodnar et al., 2021), and hypergraphs (Feng et al., 2019; Xu et al., 2022). Unlike graphs, these structures can naturally represent higher-degree, setwise relationships. However, not all structures are appropriate for all systems (Spivak, 2009; Rosas et al., 2022).

Why do we choose to model our setwise weights as simplicial complexes and not, say, general cell complexes or hypergraphs? There are three main reasons:

1. *Simplicial complexes allow all possible setwise relations to exist.* Simplices, by construction, may span any number of vertices. This means any possible combination of neurons may share a common, setwise weight. This is also possible in hypergraphs, but not possible in all complexes, e.g., cube complexes may not include triangles. The 'natural' complex for modelling all possible setwise relationships is therefore a simplicial one. Since we also wish for our setwise weights to be symmetrical, i.e., have the same value when updating the spin with respect to each constituent neuron, it is unnecessary to include any more than one unique object per setwise relationship, which also makes the choice of a simplex suitable

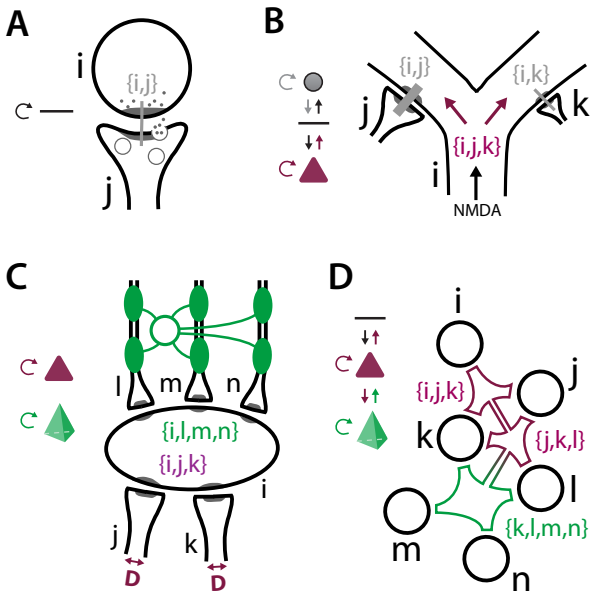

Figure 4: **A.** Neurotransmitter depletion at a synapse. **B.** NMDA-spikes and unequal synaptic strengths in dendritic integration. **C.** Transmission speed plasticity using myelination (top) and axon diameter (bottom) to affect 'temporal'/functional setwise influence in a post-synaptic cell. **D.** Astrocytic messaging, between both neurons and astrocytes.

(since there can only be one unique simplex for a given set of vertices – which is also the case for edges in undirected hypergraphs, but we find these are less suitable).

2. *The sub-edge problem of hypergraphs makes them less suitable.* Hypergraphs are graphs where edges are sets containing any number of unique vertices from the vertex set. In a sense, these are a more general structure than simplicial complexes and lack downward closure, i.e., if the edge $\{1, 2, 3\}$ exists, edges such as $\{2, 3\}$ or $\{1\}$ do not necessarily exist, whereas if $\{1, 2, 3\}$ was a simplex in a simplicial complex, simplices $\{2, 3\}$ and $\{1\}$ exist. However, hypergraphs do not have well-defined 'sub-edges' (described as the 'sub-edge problem' in Remark 3.5 of (Spivak, 2009)). This has the consequence of defining interactions between 'levels' of hyperedges (setwise relationships) in hypergraphs slightly awkward. In contrast, simplicial complexes have a well-defined hierarchy of setwise relationships, partly due to the downward closure condition.

3. *Downward closure of setwise connections is biologically plausible.* Another benefit of downward closure in simplicial complexes is the fact that it currently seems to be better supported from the perspective of biological plausibility (also see Appendix A.1). For example, although it can happen that a setwise connection (anatomical or functional) between neurons could exist without any underlying pairwise connections, the typical machinery used to create such setwise connections is sufficiently local to assume that, because of the functional local modularity of connections in the brain (Kaiser & Hilgetag, 2006; Chen et al., 2013; Müller et al., 2020; Ercsey-Ravasz et al., 2013), there is a high probability of these neurons having a pairwise connection simply due to proximity.

As an additional practical benefit, simplicial complexes are currently more well-studied than structures such as hypergraphs (at least in some areas, e.g., spectral theories or (co)homology, which are of natural interest here), meaning that we can also take advantage of the relative maturity of the field in those areas – admittedly, we use very few advanced methods or properties in an essential way in this study, although we hope to do so in future studies having now introduced an initial interpretation of simplicial Hopfield networks and begun exploring some of their potential benefits. However, it will also be interesting to see what differences can be found between hypergraphic and simplicial

Hopfield networks, and perhaps which provides a closer approximation to biology or which shows improved performance on certain tasks.

Among other possibilities, the weight of a simplex $w(\sigma)$ could be modulated by the local energy of its spins S_σ, its coface's spins, or of those of simplices in the same dimension as σ which are 'neighbouring' (in the Hodge Laplacian sense (Lim, 2020; Schaub et al., 2020)). These interactions could take many different mathematical forms (Petri & Barrat, 2018; Ebli et al., 2020; Roddenberry et al., 2021; Rosas et al., 2022; Santoro et al., 2022). Neurobiologically, these interactions could represent neural–glia interactions, glia–glia interactions, nonlinear dendritic integration (especially dendritic spikes and shunting), neurotransmitter–neuromodulator interactions, or hierarchical assembly operations and dynamics, to name a few (see Appendix A.1 for illustrations and further information).

The downward closure of simplicial complexes could be seen as a disadvantage. For example, when including simplices of high dimension, we are also forced to limit our choices of simplices if we wish to maintain the simplicial structure. Again, whenever a n–simplex exists in K, all its faces must also exist, e.g., if a triangle exists, so must its surrounding edges and their surrounding vertices. If any constituent simplex is missing, the structure of the simplicial complex is broken and in our case would become an undirected, weighted hypergraph where there are additional interactions between the hyperedges. Instead, in our simulations, we prefer to interpret 'missing simplices' as merely functionally 'missing-in-action' by setting their weights to zero if we do not wish to include them in the model. This has the consequence of having no mathematical effect on our update rules while retaining the benefits of a simplicial structure, which is convenient for defining potential weight modulations.

A.3 Simplicial homology

Simplicial homology allows us to precisely count the number of 'holes' in each dimension of a simplicial complex by calculating the k–dimensional Betti number, β_k. A related topological property which is particularly useful when studying low-dimensional objects (e.g., classification of surfaces) is the Euler characteristic, which for a simplicial complex can be calculated by $\chi(K) = \sum_{i=0}^{\infty}(-1)^k|K_k|$, i.e., it is an alternating sum which 'balances' out the number of holes in odd and even dimensions. It is related to the Betti numbers insofar as the Euler characteristic is also given by $\chi(K) = \sum_{i=0}^{\infty}(-1)^k\beta_k$. Note, however, that $|K_k| \neq \beta_k$. As such, although the Euler characteristic can be used for comparing the topologies of two simplicial complexes, it is not as informative (with respect to holes) as homology (although the latter is more costly to compute, as we will now see).

***k*-chains and boundaries.** The group of k-chains is a free Abelian group with the basis of K_k,

$$C_k = C_k(K) := \mathbb{Z}K_k := \left\{ \sum_{\sigma \in K_k} \alpha_\sigma \sigma \mid \alpha_\sigma \in \mathbb{Z} \right\}.$$

The boundary (difference between the 'end points') of a face σ in dimension k is

$$\partial_k(\sigma) := \sum_{j \in \sigma} \mathrm{sign}(j, \sigma)(\sigma \backslash j).$$

where $\mathrm{sign}(j, \sigma) = (-1)^{i-1}$, where j is the i-th element of σ (ordered) and $\sigma \backslash j := \sigma \backslash \{j\}$.

Example 1. Consider the simplicial complex

$$K = \{\{1,2,3\}, \{1,2\}, \{1,3\}, \{2,3\}, \{1\}, \{2\}, \{3\}, \emptyset\}. \tag{10}$$

For $k = 2$, we have $K_2 = \{\{1,2,3\}\}$ and $C_2 = \{\{1 \cdot \{1,2,3\}\}\}$. Let us calculate the boundary of $\sigma = \{1,2,3\}$. First, we calculate the respective sign functions:

$$\begin{aligned}
\mathrm{sign}(1, \sigma) &= (-1)^{i-1} = (-1)^{1-1} = (-1)^0 = 1 \\
\mathrm{sign}(2, \sigma) &= (-1)^{i-1} = (-1)^{2-1} = (-1)^1 = -1 \\
\mathrm{sign}(3, \sigma) &= (-1)^{i-1} = (-1)^{3-1} = (-1)^2 = 1.
\end{aligned}$$

Using these values, we may calculate the boundary by

$$\begin{aligned}\partial_2(\sigma) &= \operatorname{sign}(1,\sigma)(\sigma\backslash 1) && + \operatorname{sign}(2,\sigma)(\sigma\backslash 2) && + \operatorname{sign}(3,\sigma)(\sigma\backslash 3)\\ &= (1)(\sigma\backslash 1) && + (-1)(\sigma\backslash 2) && + (1)(\sigma\backslash 3)\\ &= (1)(\{2,3\}) && + (-1)(\{1,3\}) && + (1)(\{1,2\})\\ &= \{2,3\} && - \{1,3\} && + \{1,2\}.\end{aligned}$$

The boundary of $\{1,2,3\}$ is $\{2,3\} - \{1,3\} + \{1,2\}$. Notice, this is a cycle.

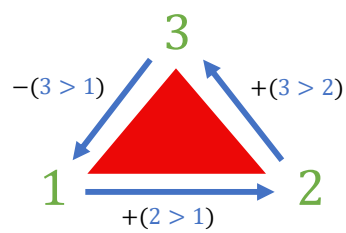

Figure 5: Geometric realisation of the simplicial complex in equation 10 and the boundary of $\{1,2,3\}$.

Chain complex. The k-th boundary mapping, ∂_k, is the map $C_k(K) \to C_{k-1}(K)$. If $k > m-1$ or $k < -1$, then $C_k(K) := 0$ and $\partial_k := 0$, where m is the number of 0–simplices. Based on this, the chain complex of K is

$$0 \to C_{n-1}(K) \xrightarrow{\partial_{n-1}} \dots \xrightarrow{\partial_2} C_1(K) \xrightarrow{\partial_1} C_0(K) \xrightarrow{\partial_0} C_{-1}(K) \to 0.$$

We define $\partial^2 := \partial \circ \partial = 0$. For example, $\partial_{i-1} \circ \partial_i = 0$. This has the consequence of making the boundary of, say, a solid tetrahedron (3–simplex) a set of oriented triangles (2–simplices) with a 'net flow' of 0 (which calculus students may find similar to Stokes' curl theorem).

Example 2. Consider the following simplicial complex and its chain complex:

$$K = \{\{1,2\},\{1\},\{2\},\{3\},\{4\},\emptyset\},$$

$$0 \to C_1(K) \xrightarrow{\partial_1} C_0(K) \xrightarrow{\partial_0} C_{-1}(K) = 0.$$

The boundary map ∂_1 is $\{1,2\} \mapsto \{2\} - \{1\}$ and all faces in K_0 are mapped to the empty set by ∂_0.

Homology. We can now see that our k-cycles are $Z_k = \ker \partial_k$ (k-chains α where $\partial_k(\alpha) = 0$). Whereas, the k-boundaries are $B_k = \operatorname{im} \partial_{k+1}$ (k-chains in the image of ∂_{k+1}). Notice, $B_k \subset Z_k$.

The (reduced) k-homology of K is the Abelian group

$$\widetilde{H}_k(K) := Z_k / B_k,$$

and we define $\widetilde{H_{m-1}}(K) := \ker \partial_{m-1}$ for $k > m-1$ and $\widetilde{H}_k(K) := 0$ for $k < 0$.

The k-th Betti number (number of topological holes) is

$$\begin{aligned}\beta_k &= \dim\left(\widetilde{H}_k\right)\\ &= \dim(Z_k) - \dim(B_k)\\ &= \operatorname{nullity}(\partial_k) - \operatorname{rank}(\partial_{k+1}).\end{aligned}$$

Note, $\operatorname{nullity} \partial_0 = \dim(C_0)$.

Example 3. Consider the simplicial complex

$$K = \{\{1,2\},\{1,3\},\{2,3\},\{1\},\{2\},\{3\},\emptyset\},$$

which is the same as that depicted in Figure 5 but without the filled-in triangle. The chain complex is

$$0 \to \mathbb{Z}^3 \xrightarrow[\partial_1]{\begin{array}{c|ccc} & \{1,2\} & \{1,3\} & \{2,3\} \\ \hline \{1\} & -1 & -1 & 0 \\ \{2\} & 1 & 0 & -1 \\ \{3\} & 0 & 1 & 1 \end{array}} \mathbb{Z}^3 \xrightarrow[\partial_0]{\text{(0 map)}} 0.$$

We may compute its Betti numbers by

$$\begin{aligned} \beta_0 &= \text{nullity}\,(\partial_0) - \text{rank}\,(\partial_{0+1}) \\ &= 3 - 2 = 1 \\ \beta_1 &= \text{nullity}\,(\partial_1) - \text{rank}\,(\partial_{1+1}) \\ &= 1 - 0 = 1, \end{aligned}$$

and, by definition, $\beta_{>1} = 0$.

A.4 PROOFS

The following are our proofs for statements in the main text.

Corollary A.1 (Memory capacity is proportional to the number of network connections)**.** *If the connection weights in a Hopfield network are symmetric, then the order of the network's memory capacity is proportional to the number of its connections.*

Proof. Let d be the degree of connections in a Hopfield network with N neurons. The explicit or implicit number of connections in such a network is N^d. By a simple counting argument, the number of repeated connections between any set of d neurons is interpreted as $d!$. By Newman (1988); Demircigil et al. (2017), the order of such a network's memory capacity is N^{d-1}. So the following relationship holds:

$$\frac{N^{d-1}}{N^d/d!} \cdot \frac{N}{d!} = \frac{d!}{N} \cdot \frac{N}{d!} = 1.$$

□

Lemma A.2 (Fully-connected mixed Hopfield networks)**.** *A fully-connected mixed Hopfield network based on a pure D–simplicial complex with N neurons and P patterns has, when $N \to \infty$ and P is finite: fixed point attractors at the memory patterns and dynamic convergence towards the fixed point attractors within a finite Hamming distance δ. When $P \to \infty$ with $N \to \infty$, the network has capacity to store up to $(\sum_{d=1}^{D} N^d)/(2\ ln\ N)$ memory patterns (with small retrieval errors) or $(\sum_{d=1}^{D} N^d)/(4\ ln\ N)$ (without retrieval errors).*

Proof. Let K be a pure D–simplicial complex. Let S_i be the spin of neuron i, and let the spin correspond to a stored pattern, $S_i = \xi_i^1$. To be in a fixed point, the local field h_i applied to i must satisfy the inequality $S_i h_i > 0$, meaning the local field being applied to the neuron must be of the same sign as the present spin.

In the case of the mixed network, based on the simplicial Hopfield equations 1–2, the local field is

$$\begin{aligned} h_i\{S\} &= \frac{1}{|K_1|} \sum_{\sigma \in K_1}^{|K_1|} \sum_{\mu=1}^{P} \xi_i^\mu \xi_{\sigma \setminus i}^\mu S_{\sigma \setminus i} + \frac{1}{|K_2|} \sum_{\sigma \in K_2}^{|K_2|} \sum_{\mu=1}^{P} \xi_i^\mu \xi_{\sigma \setminus i}^\mu S_{\sigma \setminus i} + \dots \\ &= \sum_{d=1}^{D} \left(\frac{1}{|K_d|} \sum_{\sigma \in K_d}^{|K_d|} \sum_{\mu=1}^{P} \xi_i^\mu \xi_{\sigma \setminus i}^\mu S_{\sigma \setminus i} \right). \end{aligned} \tag{11}$$

Because K is a pure D–simplicial complex, each dimension K_d will have $\binom{N}{d}$ elements. Using Stirling's Approximation for the binomial coefficient, $\binom{N}{d} \approx N^d/d!$ (for $N \gg d$, which holds if $\dim(K)$ is small, which we argue it should normally be for both computational and biological reasons), we can simplify equation 11 slightly

$$h_i\{S\} = \sum_{d=1}^{D} \left(\frac{d!}{N^d} \sum_{\sigma \in K_d} \sum_{\mu=1}^{P} \xi_i^\mu \xi_{\sigma \setminus i}^\mu S_{\sigma \setminus i} \right). \tag{12}$$

To analyse the stability of a pattern, we set $S_i = \xi_i^1$, where the choice of 1 is arbitrary (since the weights are symmetric). Substituting ξ^1 for S and equation 12 for h, the inequality we must satisfy for $i = 1$ becomes

$$\xi_1^1 h_1 = \sum_{d=1}^{D} \left(\frac{d! \prod_m^d (N-m)}{N^d} + \frac{d!}{N^d} \sum_{\sigma \in K_d} \sum_{\mu=2}^{P} \xi_1^1 \xi_1^\mu \xi_{\sigma \setminus i}^\mu \xi_{\sigma \setminus i}^1 \right) > 0. \tag{13}$$

Notice in equation 13 we decomposed the summation over patterns into *signal* terms (for the pattern we are analysing) and *noise* terms (for the contribution of all other patterns). In the limit of $N \to \infty$, the signal terms are fixed numbers (of order 1) and, by the Central Limit Theorem, since the noise terms are sums of random numbers (essentially, random walks), they will have means of 0 and standard deviations of

$$\frac{d!}{N^d} \sqrt{(P-1) \prod_m^d (N-m)}, \tag{14}$$

which we can approximate as $\sqrt{P/N^d}$. We can see that as d increases, the noise terms reduce in variance. However, if P remains fixed and N is sufficiently large, the noise terms become negligible compared to the signal terms. This therefore guarantees that every pattern will be a *fixed point*.

Furthermore, these fixed points will remain highly stable against random noise. Suppose we randomly flip a finite number δ of spins away from a pattern $\xi = S$ (a fixed point). The signal terms' strengths are reduced by 2δ but still of order 1, whereas the noise terms remain of order $N^{-1/2}$. Therefore, states within Hamming distance δ away from ξ will *converge* to the fixed point.

Now let $P \to \infty$ with N. The total variance of the noise terms is

$$v = \sum_{d=1}^{D} \sqrt{P/N^d}. \tag{15}$$

The probability of stability of neuron i in equation 13 is the probability that the noise terms are larger than -1; at ≤ -1 they will overcome the signal terms. This probability is

$$\Pr(\xi_1^1 h_1 > 0) = \frac{1}{\sqrt{2\pi v^2}} \int_{-1}^{\infty} dx \exp\left(-\frac{x^2}{2v^2}\right) = \frac{1}{2}\left[1 + \text{erf}\left(\sqrt{\frac{1}{2v^2}}\right)\right], \tag{16}$$

where

$$\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x dt e^{-t^2}. \tag{17}$$

For small v (which this is, especially as d increases and relative to the signal), the value of the error function in equation 16 will be large and can therefore be approximated as (Gradshteyn & Ryzhik, 2007)

$$\text{erf}(x) \approx 1 - \frac{1}{\sqrt{\pi} x} e^{-x^2}. \tag{18}$$

We can now approximate equation 16 as

$$\Pr(\xi_1^1 h_1 > 0) \approx 1 - \sqrt{\frac{z}{2\pi}} \exp\left(-\frac{1}{2z}\right), \tag{19}$$

where

$$z = \sum_{d=1}^{D} P/N^d. \tag{20}$$

We can now calculate the probability of a stable pattern, i.e., that the inequality $\xi_i^1 h_i > 0$ is satisfied for all i, with

$$\begin{aligned} \Pr(\text{stable pattern}) &\approx \left[1 - \sqrt{\frac{z}{2\pi}} \exp\left(-\frac{1}{2z}\right)\right]^N \\ &\approx 1 - N\sqrt{\frac{z}{2\pi}} \exp\left(-\frac{1}{2z}\right). \end{aligned} \tag{21}$$

Since $N \to \infty$, equation 21 will be close to 1 as the second term will be negligible. This will always be true if

$$z = \frac{1}{2\ln N}. \tag{22}$$

Therefore, since we have N neurons with $\sum_{d=1}^{D} N^d$ connections between them, the maximum number of patterns we may store (while accepting small errors) is

$$p_c = \frac{\sum_{d=1}^{D} N^d}{2 \ln N}. \tag{23}$$

Or, if we cannot accept errors,

$$p_c = \frac{\sum_{d=1}^{D} N^d}{4 \ln N}. \tag{24}$$

□

An additional informal perspective to consider regarding memory capacity is to notice that if P scales with N, the ratio between the signal and noise terms will be constant. And, since $P \propto N^d$, the theoretical memory capacity scales polynomially with N and linearly with D, and theoretical capacity is approximately $\sum_{d=1}^{D} c_d \cdot N^d$, where c_d is a constant which depends on d.

A.5 Supplementary figures and tables

Table 3: Pearson correlation coefficients (r) between overlap and β_1 for mixed diluted networks from Table 2. Bolded values are significant at $\alpha = 0.05$ (without multiple comparisons adjustment). With multiple comparisons adjustment, there are no significant correlations. Given their construction, all networks have $\beta_0 = 1$ (although for completeness we also verified this numerically), and although there is a small chance of 2–dimensional holes in some networks, we found that $\beta_{\geq 2} = 0$ for all simulated networks.

No. patterns	$0.05N$	$0.1N$	$0.15N$	$0.2N$	$0.3N$
R12 (50%)	0.02	0.09	**0.21**	-0.15	-0.01
R12 (75%)	N/A	-0.05	-0.1	0.01	-0.08

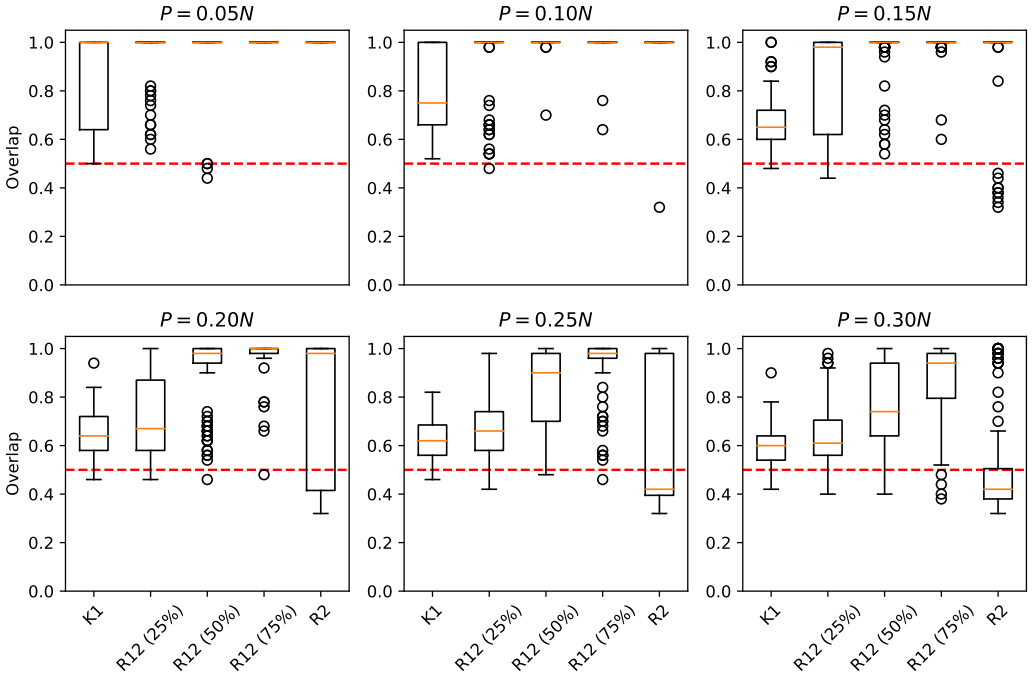

Figure 6: Box and whisker plots of final overlap distributions from traditional simplicial Hopfield networks with varying numbers of embedded patterns. Orange lines indicate the median. The red dashed line indicates an overlap of 0.5, chance.

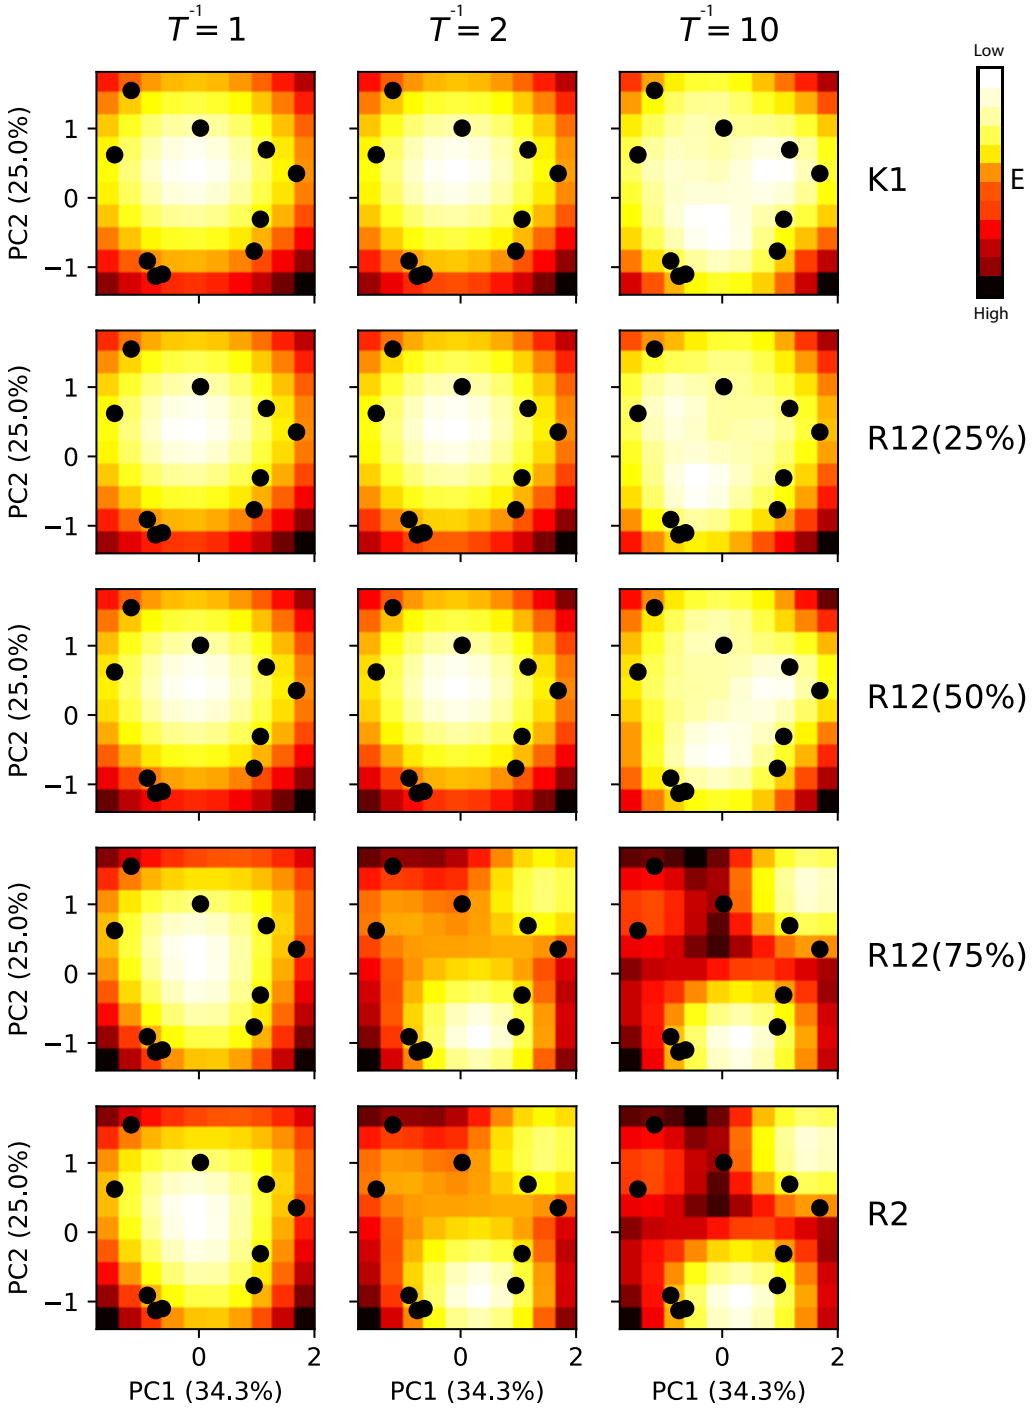

Figure 7: Relative energies of continuous modern networks in a 10×10 grid plane of the first two dimensions of PCA space, as computed using the memory patterns. Each network has $N = 10$ and the same $P = 10$ memory patterns are embedded. Network conditions vary by row and inverse temperatures vary by column. Black dots are projections of the 10 embedded patterns in the PCA space. The combined explained variance of the first two principle components is 59.3%.