Modulating interactions to control dynamics of neural networks

Lukas Herron
Biophysics Program and Institute for Physical Science and Technology,
University of Maryland, College Park, MD, United States
lherron@umd.edu

Pablo Sartori
Instituto Gulbenkian de Ciência,
Oeiras, Portugal
psartori@igc.gulbenkian.pt

BingKan Xue
Department of Physics, University of Florida,
Gainesville, FL, United States
b.xue@ufl.edu

Abstract

Sequential retrieval of stored patterns is a fundamental task that can be performed by neural networks. Previous models of sequential retrieval belong to a general class in which the components of the network are controlled by a slow feedback ("input modulation"). In contrast, we introduce a new class of models in which the feedback modifies the interactions among the components ("interaction modulation"). In particular, we study a model in which the symmetric interactions are modulated. We show that this model is not only capable of retrieving dynamic sequences, but it does so more robustly than a canonical model of input modulation. Our model allows retrieval of patterns with different activity levels, is robust to feedback noise, and has a large dynamic capacity. Our results suggest that interaction modulation may be a new paradigm for controlling network dynamics.

Introduction

Neural networks are capable of generating well-defined sequences of activity patterns (1; 2). These patterns can be considered metastable states of the dynamics; the ability of the network to transition from one such configuration to the next implies that the system can alter the stability of the configurations. A well-studied approach to generating sequential transitions is to modulate the activity of neurons. Such a scenario corresponds to a common approach in control theory, i.e., modulating *inputs* on a subset of variables to influence the full system (3). An alternative, less studied approach is to modify the interaction between a pair of neurons by a third neuron (4; 5). Can such interaction modulation be used to generate dynamic sequences of activity patterns in neural networks?

In this work we address this question using Hopfield networks. While the original Hopfield network was developed as a model of associative memory capable of storing and retrieving particular neural activity patterns (6), sequential transitions among patterns have long been considered (7; 8; 9; 10). These models of sequential dynamics are aimed at describing phenomena such as central pattern generation (11), counting (12) and, more recently, free association (13), and memory recall (14). Here, we model the dynamics of retrieval by introducing a small set of feedback units, which control the sequential transitions. When formulated this way, previous models of sequential dynamics are shown to fall into a class of models based on *input modulation*. We propose a new class of models that rely on *interaction modulation* to trigger autonomous transitions between configurations. Remarkably, we find that modulation of symmetric interactions allows sequential retrieval of configurations that have different activity levels, which cannot be reliably done by models that use input modulation.

Associative Memory & Hopfield Networks in 2023. NeurIPS 2023 workshop.

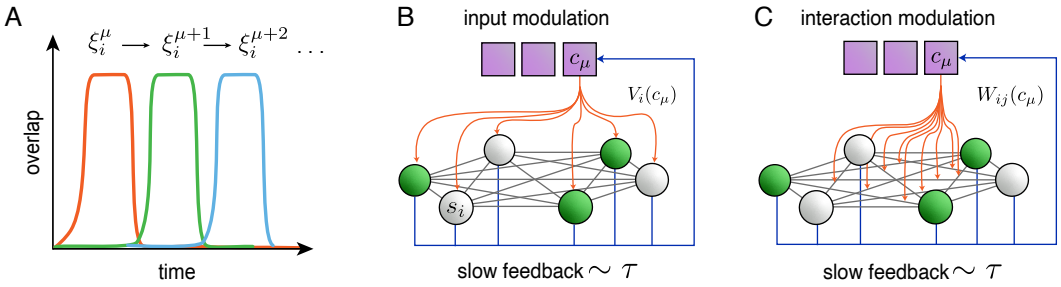

Figure 1: Sequential retrieval through input and interaction modulation. (A) The network transitions along an ordered sequence of patterns, as characterized by the overlap of the network state with each pattern over time. (B,C) A Hopfield network with units s_i (spherical nodes) is controlled by a set of feedback units c_μ (square nodes). The feedback units are updated on a slow timescale τ (blue lines), and modulate either the input V_i to the main units (B) or their interactions W_{ij} (C).

Furthermore, this model can retrieve much longer sequences than other models. Our results suggest that interaction modulation may be biologically and computationally favored over input modulation due to its robustness and large dynamic capacity.

Model

Our model is based on the classic Hopfield network (6), which can store and retrieve a given set of network configurations, or "patterns". The network is composed of N units, $\{s_i\}_{i=1}^N$, which take continuous values, $0 \leq s_i \leq 1$. The dynamics of these units follow

$$\dot{s}_i = -s_i + f(h_i) \ , \quad \text{with} \quad h_i = \sum_{j=1}^{N} W_{ij} s_j + V_i \ . \tag{1}$$

Here W_{ij} are the interactions among the units, V_i is an external input to each unit, and $f(\cdot)$ is an activation function chosen to be the Heaviside step function for simplicity. The network stores p patterns, $\{\xi^\mu\}_{\mu=1}^p$, each being a binary vector, $\xi_i^\mu \in \{0,1\}$. The average activity level of a pattern is $a_\mu = \frac{1}{N}\sum_i \xi_i^\mu$; when patterns have equal activity, as often considered, it will simply be denoted by a. The original Hopfield network corresponds to $a = \frac{1}{2}$, $W_{ij} = J_{ij}$ and $V_i = 0$, where the symmetric interactions are $J_{ij} = \frac{1}{N}\sum_{\mu=1}^p (\xi_i^\mu - a)(\xi_j^\mu - a)$.

We are interested in networks that can autonomously retrieve a sequence of patterns one after another. The system relaxes to a pattern ξ^μ, which is slowly destabilized, then relaxes to the next pattern $\xi^{\mu+1}$, and so on. Such dynamics require a separation of timescales between fast stabilization and slow destabilization (7). We therefore introduce a set of feedback units, $\{c_\mu\}_{\mu=1}^p$, which obey the dynamics

$$\dot{c}_\mu = -\frac{1}{\tau}(c_\mu - m^\mu) \ , \quad \text{with} \quad m^\mu = \frac{1}{Na(1-a)} \sum_i (\xi_i^\mu - a) s_i \ , \tag{2}$$

and a timescale $\tau \gg 1$. These feedback units c_μ will be used to destabilize the retrieved pattern by modifying either the inputs to the network, $V_i(c_\mu)$, or the interactions, $W_{ij}(c_\mu)$, as schematically depicted in Fig. 1.

We can reformulate well-studied models of sequential retrieval using our framework. For example, the Sompolinsky-Kanter (SK) model (7) is obtained by identifying $W_{ij} = J_{ij}$ and $V_i = \lambda \sum_\mu (\xi_i^{\mu+1} - a)c_\mu - \theta$. After the network retrieves a pattern ξ^ν, the feedback unit c_ν slowly activates an input that biases the system towards the subsequent pattern $\xi^{\nu+1}$. This is enough to destabilize the current pattern and induce a transition if λ is sufficiently large (Fig. 2). In this and other similar models (see Supplementary Material) the feedback units c_μ modulate the input V_i to achieve sequential retrieval.

We now consider a new class of models in which the feedback units c_μ modulate the interactions W_{ij}. While many variants are possible (see SM), we illustrate the idea using a model in which

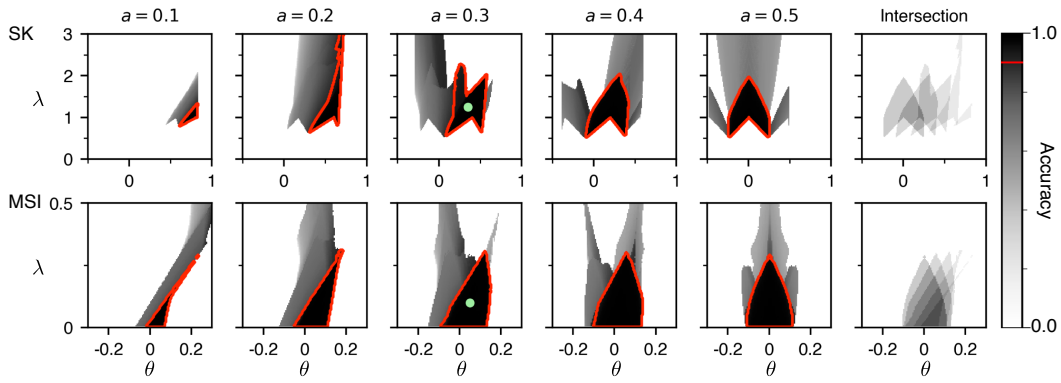

Figure 2: Parameter space of sequential retrieval. The performance of the SK and MSI models for retrieving a cyclic sequence of four orthogonal patterns is evaluated for different combinations of the bias λ and threshold θ. The parameter space is numerically swept with increments $\Delta\lambda = \Delta\theta = 0.025$. The grayscale color represents the accuracy score (see SM); the red contours represent regions of high accuracy (> 0.9). The first five columns correspond to different levels of pattern activity a, whereas the last column shows the overlay of the high-accuracy regions for these activity levels.

$W_{ij} = J_{ij}(c_\mu) + \lambda \tilde{J}_{ij}$ and $V_i = -\theta$, where

$$J_{ij}(c_\mu) = \frac{1}{N} \sum_\mu (\xi_i^{\mu+1} - a)(\xi_j^{\mu+1} - a)\, c_\mu \ . \tag{3}$$

The asymmetric interactions $\tilde{J}_{ij} = \frac{1}{N} \sum_\mu (\xi_i^{\mu+1} - a)(\xi_j^{\mu} - a)$ provide a bias towards subsequent patterns, even though the parameter λ can be small (Fig. 2). In this model, when the network has retrieved a pattern ξ^ν for a period of time, all c_μ will decay to zero except c_ν. As a consequence, all terms in J_{ij} will be turned off except for $\mu = \nu$. Therefore, only one pattern $\xi^{\nu+1}$ will be stable and subsequently retrieved by the network. We refer to this model as the "Modulation of Symmetric Interactions" (MSI) model.

Results

Phase space of sequential dynamics. The two models above are characterized by the same set of parameters: the bias λ and the threshold θ. This allows us to compare input and interaction modulation by examining the (λ, θ) parameter space and identifying regions where sequential retrieval is successful. To quantify the accuracy of sequential retrieval, we use a custom scoring function (see SM). In Fig. 2, shaded regions correspond to parameter combinations that produce sequential dynamics, and regions within the red contour correspond to high accuracy. Both models have compact regions of parameter space that allow sequential retrieval. However, the location and shape of these regions depend on the activity level a. The last column of Fig. 2 shows an overlay of the regions of accurate retrieval for different a values. For SK the retrieval region drifts from a positive value of θ towards 0 as a increases. Notably, for MSI the retrieval regions for different a values overlap, suggesting that MSI is able to retrieve sequences of patterns with varying activities.

Variability in pattern activity. We therefore consider patterns ξ^μ that each have a different activity level a_μ (in which case the formulae involving a are modified to have a_μ instead). As an example, we choose five patterns with a_μ equally spaced within the range $0.3 \pm 0.2r$, where the "unevenness" parameter r is varied between 0 and 1. Because the patterns have different activities, their order in the sequence can affect retrieval. Therefore, we compute the mean accuracy of retrieval over all possible permutations of the patterns (see SM). For each level of unevenness, there is a region of accurate retrieval in the parameter space (Fig. S2A). We use the area of the retrieval region (relative to that for patterns with equal activity) to measure the ability of the network to retrieve uneven patterns. As shown in Fig. 3A, the MSI model is more robust to the unevenness and ordering of the patterns.

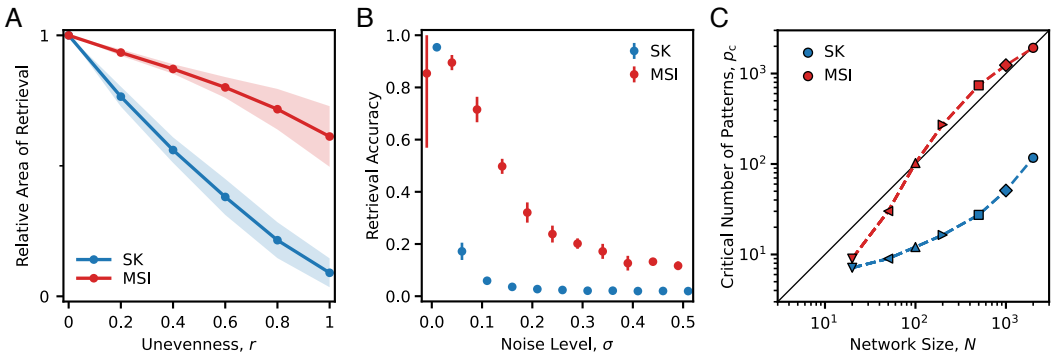

Figure 3: Dynamic capacity of sequential retrieval. (A) Relative area of the retrieval region in (λ, θ) parameter space plotted against the unevenness parameter r for the SK and MSI models. The curves show the mean of the relative area over sequence permutations (for an accuracy cutoff of 0.8; our results are not qualitatively affected by altering the accuracy cutoff, as shown in Fig. S2B), and the color shades show the standard deviations. (B) Retrieval accuracy of a cyclic sequence of $p = 20$ orthogonal patterns (with $a = 0.3$) at different noise levels σ. The mean accuracy and standard deviation are shown for 10 realizations. (C) The critical number of patterns p_c plotted against the network size N. For each network size N, the average accuracy of retrieval is estimated for different numbers of patterns ($a = 0.3$) over 100 realizations and fitted to a sigmoidal curve (Fig. S3CD); p_c is defined as where the average accuracy falls below 0.7. The parameters used to simulate each model correspond to the green dots in Fig. 2.

Noise in the feedback units. We also study the robustness of the sequential dynamics to noise in the feedback units c_μ by adding Gaussian white noise with zero mean and variance σ^2 to Eq. (2). As shown in Fig. 3B, MSI is substantially more robust than SK for the same level of noise. Such robustness is likely due to the dynamics happening across a less rugged landscape. In Hopfield networks, there are extra local minima in the slowly changing energy landscape, such as spurious states from linear combinations of patterns, which can stall the sequential dynamics. MSI suppresses all patterns except the one involved in the current transition, resulting in a smoother landscape that allows the dynamics to recover from large perturbations induced by noise.

Dynamic storage capacity. Finally, we study the storage capacity of the network in terms of the number of patterns that can be stored and sequentially retrieved by the network. We define p_c as the critical number of patterns beyond which the average accuracy drops below a threshold value. Fig. 3C shows p_c plotted against the network size N, where MSI again outperforms SK. To understand this advantage in storage capacity, we note that the retrieval of a pattern in Hopfield networks can be disrupted by crosstalk between different patterns that are not fully orthogonal to each other (15). As can be shown analytically (see SM), while both models are subject to crosstalk with a similar scaling with network size N and number of patterns p, in MSI the crosstalk is relatively suppressed. This is due to an extra factor of c_μ in the crosstalk term, which comes from interaction modulation.

Discussion

We have shown that an interaction modulation model (MSI) is capable of retrieving dynamic sequences, while being robust to variation in the activity level of the patterns, noise in the feedback units, and length of the sequence. Besides these advantages, interaction modulation also exhibits potential drawbacks. A main issue is that each feedback unit has to generate $O(N^2)$ outputs to modulate all the interactions, as opposed to $O(N)$ outputs for input modulation. This requirement of large connectivity can be relaxed if we dilute the outputs of the feedback units by suppressing a random fraction of the modulated interactions (Fig. S4).

Sequence generation in neural networks has been modeled using slow feedback (7; 8; 9; 16; 13) and other approaches (17; 18; 19; 20; 21; 22). Coupling feedback to interactions is motivated by well-documented neurobiological mechanisms (23), such as heterosynpatic plasticity, by which a synapse that is not currently active can be strengthened or weakened by the firing of a third modulatory neuron

(4; 24; 25). There is also extensive literature on tripartite synapses (5; 26), which allow a synapse between two neurons to be dynamically modulated by one or more astrocytes. In line with our results, such systems have been shown to exhibit enhanced computational capabilities (27). The performance of our model may be further improved by combining interaction modulation with Dense Associative Memory models, which have recently been shown to have long sequence capacity (28; 29).

References

[1] G. V. Wallenstein, M. E. Hasselmo, and H. Eichenbaum, Trends in Neurosciences **21**, 317 (1998).

[2] H. Eichenbaum, Trends in Cognitive Sciences **17**, 81 (2013).

[3] Y.-Y. Liu and A.-L. Barabási, Reviews of Modern Physics **88** (2016).

[4] C. H. Bailey, M. Giustetto, Y.-Y. Huang, R. D. Hawkins, and E. R. Kandel, Nature Reviews Neuroscience **1**, 11 (2000).

[5] G. Perea, M. Navarrete, and A. Araque, Trends in Neurosciences **32**, 421 (2009).

[6] J. J. Hopfield, Proceedings of the National Academy of Sciences **79**, 2554 (1982).

[7] H. Sompolinsky and I. Kanter, Physical Review Letters **57**, 2861 (1986).

[8] D. Horn and M. Usher, Physical Review A **40**, 1036 (1989).

[9] S. Dehaene, J.-P. Changeux, and J.-P. Nadal, Proceedings of the National Academy of Sciences **84**, 2727 (1987).

[10] J. Buhmann and K. Schulten, Europhysics Letters **4**, 1205 (1987).

[11] D. Kleinfeld and H. Sompolinsky, Biophysical Journal **54**, 1039 (1988).

[12] D. J. Amit, Proceedings of the National Academy of Sciences **85**, 2141 (1988).

[13] E. Russo and A. Treves, Physical Review E **85**, 051920 (2012).

[14] M. Naim, M. Katkov, S. Romani, and M. Tsodyks, Physical Review Letters **124**, 018101 (2020).

[15] J. A. Hertz, Introduction to the theory of neural computation, CRC Press, 2018.

[16] L. Pantic, J. J. Torres, H. J. Kappen, and S. C. Gielen, Neural Computation **14**, 2903 (2002).

[17] S. Recanatesi, M. Katkov, and M. Tsodyks, Neural Computation **29**, 2684 (2017).

[18] P. Seliger, L. S. Tsimring, and M. I. Rabinovich, Physical Review E **67**, 4 (2003).

[19] V. S. Afraimovich, M. I. Rabinovich, and P. Varona, International Journal of Bifurcation and Chaos **14**, 1195 (2004).

[20] D. Sussillo and L. F. Abbott, Neuron **63**, 544 (2009).

[21] L. Logiaco, L. Abbott, and S. Escola, Cell Reports **35**, 109090 (2021).

[22] C. Parmelee, J. L. Alvarez, C. Curto, and K. Morrison, SIAM Journal on Applied Dynamical Systems **21**, 1597 (2022).

[23] T. F. Burns and T. Fukai, Proceedings of the 11th International Conference on Learning Representations (2023).

[24] I. R. Fiete, W. Senn, C. Z. Wang, and R. H. Hahnloser, Neuron **65**, 563 (2010).

[25] T. E. Chater and Y. Goda, Current Opinion in Neurobiology **67**, 106 (2021).

[26] R. Min, M. Santello, and T. Nevian, Frontiers in Computational Neuroscience **6** (2012).

[27] L. Kozachkov, K. V. Kastanenka, and D. Krotov, Proceedings of the National Academy of Sciences **120**, e2219150120 (2023).

[28] A. Karuvally, T. J. Sejnowski, and H. T. Siegelmann, arXiv:2212.05563 (2022).

[29] H. T. Chaudhry, J. A. Zavatone-Veth, D. Krotov, and C. Pehlevan, arXiv:2306.04532 (2023).

Supplementary Material

Variant models of input and interaction modulation

Sequential retrieval can be achieved via feedback control that modulates either the input to the units or the interactions between the units. We summarize some of such models in Table 1. Besides SK and MSI described in the main text, there is another classic model due to Horn-Usher (HU) (8) that belongs to input modulation, two models of interaction modulation — complement of MSI (cMSI) inspired by (16) and global inhibition (GI) inspired by (17), and a new model based on modulation of asymmetric interactions (MAI) reminiscent of (9). Example dynamics for these models are shown in Fig. S1.

Model	Interaction W_{ij}	Field V_i	Feedback through c_μ
HU	$J_{ij} + \lambda\,\tilde{J}_{ij}$	$-\theta\,U_i(c_\mu)$	$U_i(c_\mu) \equiv \sum_\mu \xi_i^\mu c_\mu$
SK	J_{ij}	$\lambda\,U_i(c_\mu) - \theta$	$U_i(c_\mu) \equiv \sum_\mu (\xi_i^{\mu+1} - a) c_\mu$
MSI	$J_{ij}(c_\mu) + \lambda\,\tilde{J}_{ij}$	$-\theta$	$J_{ij}(c_\mu) \equiv \frac{1}{N}\sum_\mu (\xi_i^{\mu+1} - a)(\xi_j^{\mu+1} - a) c_\mu$
cMSI	$J_{ij}(c_\mu) + \lambda\,\tilde{J}_{ij}$	$-\theta$	$J_{ij}(c_\mu) \equiv \frac{1}{N}\sum_\mu (\xi_i^{\mu} - a)(\xi_j^{\mu} - a)(1 - c_\mu)$
GI	$J_{ij}(c_\mu) + \lambda\,\tilde{J}_{ij}$	0	$J_{ij}(c_\mu) \equiv \frac{1}{N}\sum_\mu \big((\xi_i^{\mu} - a)(\xi_j^{\mu} - a) - \theta\, f(c_\mu)\big)$
MAI	$J_{ij} + \lambda\,\tilde{J}_{ij}(c_\mu)$	$-\theta$	$\tilde{J}_{ij}(c_\mu) \equiv \frac{1}{N}\sum_\mu (\xi_i^{\mu+1} - a)(\xi_j^{\mu} - a) c_\mu$

Table 1: Models of input modulation (gray rows) and interaction modulation (white rows).

Scoring function for retrieval accuracy

To evaluate the performance of sequential retrieval over an extended period of time, we introduce a scoring scheme that first calculates instantaneous scores and then averages them over time to produce an overall accuracy for a time series. The instantaneous score function is defined for each pattern as:

$$s^\mu(\{m^\nu\}) = \frac{G(m^\mu)}{\sum_\nu G(m^\nu) + \epsilon} \tag{S1}$$

where

$$G(m^\mu) = \frac{\mathrm{expit}(m^\mu; \kappa, \rho) - \mathrm{expit}(-1; \kappa, \rho)}{\mathrm{expit}(1; \kappa, \rho) - \mathrm{expit}(-1; \kappa, \rho)} \quad \text{and} \quad \mathrm{expit}(x; \kappa, \rho) = \frac{1}{1 + \mathrm{e}^{-\kappa(x-\rho)}}. \tag{S2}$$

Our construction of $G(m^\mu)$ attenuates m^μ below some threshold ρ towards 0 and amplifies m^μ above ρ towards 1, so instances of retrieval correspond to a single high s^μ when the patterns are orthogonal. The parameters were chosen as $\kappa = 10$ and $\rho = 1 - a$ for all analyses, and ϵ was chosen to be 10^{-5} to make the instantaneous score well defined even when $m^\mu = 0$ for all μ.

Each pattern is retrieved and remains stable for a continuous interval of time, which we call an *instance* of retrieval. We identify such intervals as blocks of time when $G(m^\mu) \approx 1$, which corresponds to $m^\mu > \theta$. The score of an instance of retrieval amounts to the time average of the instantaneous scores over the interval:

$$S^\mu = \frac{1}{t_2 - t_1} \int_{t_1}^{t_2} s^\mu(\{m^\nu\})\, dt\,, \tag{S3}$$

where t_1 and t_2 are the bounds of the interval. The time series of network dynamics is typically composed of many retrieval instances, so we define the *overall accuracy* of sequential retrieval as the average score over many instances of retrieval detected within the time series. In all cases, the network is simulated with $\tau = 10$ and $\Delta t = 0.1$ for 6000 time steps. We compute the accuracy only for the latter half of each time series to avoid the transient dynamics in the beginning.

To determine an appropriate accuracy cutoff for Fig. 2, we examined the distribution of scores over parameter space (λ, θ, a) for each model, as shown in Fig. S3AB. The right-most peak (corresponding to high-accuracy retrieval) is separated from the remaining peaks by a cutoff accuracy of 0.9.

Effect of crosstalk on retrieval capacity

Here we provide a heuristic argument for why the MSI model is more robust than the SK model in terms of the retrieval capacity. In general, when many random patterns are stored in the network, there is crosstalk between the patterns that can interfere with the retrieval of the correct pattern. To illustrate this issue, we follow a similar analysis done for the Hopfield model and generalized for sequential retrieval (15). For simplicity, we consider the case with random unbiased patterns, i.e., with activity level $a = 0.5$, and redefine the variables as $\xi \leftarrow 2(\xi - 1)$, so that they take values between ± 1 instead of $[0, 1]$; in this unbiased case the threshold θ can be set to 0. With this convention, the symmetric interactions in the Hopfield model are simply $J_{ij} = \sum_\mu \xi_i^\mu \xi_j^\mu$. We also consider discrete dynamics, such that $s_i \leftarrow \mathrm{sgn}\left(\sum_j W_{ij} s_j + V_i\right)$, as originally studied by Hopfield (6).

Sompolinsky-Kanter (SK). Under these simplifications, the SK model can be written as

$$s_i \leftarrow \mathrm{sgn}\Big(\sum_\mu \xi_i^\mu m^\mu + \lambda \sum_\mu \xi_i^{\mu+1} c_\mu\Big) \tag{S4}$$

Suppose the network has retrieved pattern ξ^1 and remained in that state for some time, so that the feedback units are $c_1 \approx 1$ and $c_{\mu\neq 1} \approx 0$. Right before transitioning to ξ^2, the overlaps are such that $m^1 \approx 1$ and $m^{\mu\neq 1} \approx 0$. The equation above can be approximated and rearranged as

$$s_i \leftarrow \mathrm{sgn}\Big[\big(\xi_i^1 + \lambda \xi_i^2\big) + \sum_{\mu\neq 1} \big(\xi_i^\mu m^\mu + \lambda \xi_i^{\mu+1} c_\mu\big)\Big] \tag{S5}$$

The second term represents the "crosstalk" between patterns, which should vanish if the patterns are orthogonal. In the first term, $\lambda \xi_i^2$ biases the network towards the next pattern, and it has to satisfy $\lambda > 1$ to overcome the ξ_i^1 term. Then, after the network has transitioned to pattern ξ^2, we have $m^2 \approx 1$ and $m^{\mu\neq 2} \approx 0$, while the slow feedback units are still $c_1 \approx 1$ and $c_{\mu\neq 1} \approx 0$. The equation can be rearranged as

$$s_i \leftarrow \mathrm{sgn}\Big[\big(\xi_i^2 + \lambda \xi_i^2\big) + \Big(\sum_{\mu\neq 2} \xi_i^\mu m^\mu + \sum_{\mu\neq 1} \lambda \xi_i^{\mu+1} c_\mu\Big)\Big] \tag{S6}$$

The network should be stable near ξ^2 for some time as the feedback units change their values slowly. This does not give further constraints on λ as both terms $\xi_i^2 + \lambda \xi_i^2$ work to stabilize ξ^2.

Modulation of Symmetric Interactions (MSI). Similarly, the MSI model can be written as

$$s_i \leftarrow \mathrm{sgn}\Big[\sum_\mu c_\mu \xi_i^{\mu+1} m^{\mu+1} + \lambda \sum_\mu \xi_i^{\mu+1} m^\mu\Big] \tag{S7}$$

Before transitioning from ξ^1 to ξ^2, this equation can be rearranged as

$$s_i \leftarrow \mathrm{sgn}\Big[\lambda \xi_i^2 + \sum_{\mu\neq 1} \big(c_\mu \xi_i^{\mu+1} m^{\mu+1} + \lambda \xi_i^{\mu+1} m^\mu\big)\Big] \tag{S8}$$

This shows that the network is pushed towards ξ^2 even for a small λ in the absence of crosstalk. After the transition, it can be rearranged as

$$s_i \leftarrow \mathrm{sgn}\Big[(\xi_i^2 + \lambda \xi_i^3) + \Big(\sum_{\mu\neq 1} c_\mu \xi_i^{\mu+1} m^{\mu+1} + \sum_{\mu\neq 2} \lambda \xi_i^{\mu+1} m^\mu\Big)\Big] \tag{S9}$$

For a small λ, the first term would be able to stabilize the pattern ξ^2.

We expect a crosstalk term like $\sum_{\mu\neq 1} \xi_i^\mu m^\mu$ to scale as $\sim \sqrt{p/N}$. This is roughly because, when the network is in the state ξ^1, the crosstalk can be written as $\frac{1}{N} \sum_{\mu\neq 1} \sum_j \xi_i^\mu \xi_j^\mu \xi_j^1$, which is a sum over $\sim pN$ random numbers of ± 1 normalized by a factor of N, so it has zero mean and a variance of $pN/N^2 = p/N$. We also expect crosstalk terms like $\sum_{\mu\neq 1} \xi_i^{\mu+1} m^\mu$ or $\sum_{\mu\neq 1} \xi_i^{\mu+1} c_\mu$ to scale in the same way. In order for the crosstalk to overturn the leading term, it has to reach a magnitude of order 1, which implies that the critical number of patterns scales linearly as $p_c \sim N$ for both models.

For the SK model, from Eq. (S5), the main contribution to the crosstalk scales as $\sim \lambda\sqrt{p/N}$ with $\lambda > 1$, as compared to the leading term $\sim \lambda$. Thus, the capacity is reached when $p_c/N \sim 1$. On the other hand, for the MSI model, from Eq. (S9), the first term in the crosstalk is suppressed by the extra factor c_μ, which is small for $\mu \neq 1$. The second term scales as $\lambda\sqrt{p/N}$, just like in the SK model. However, for MSI, the parameter λ is small, such as $\lambda = 0.1$ used for Fig. 3C. Therefore we expect a larger prefactor in the scaling of $p_c \sim N$, which will explain the better performance of MSI.

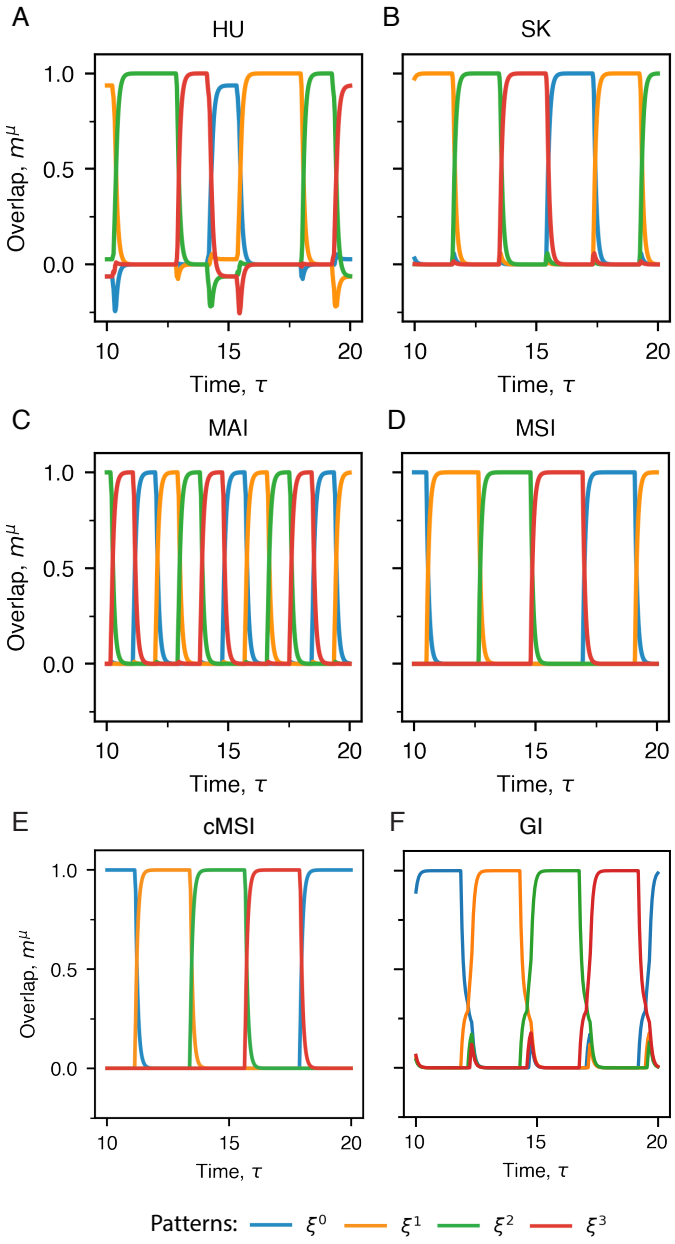

Figure S1: Examples of sequential retrieval for input and interaction modulation models. (A–B) The Horn-Usher (HU) and Sompolinsky-Kanter (SK) models that belong to the input modulation class. (C–F) Models of interaction modulation, including modulation of asymmetric interactions (MAI), modulation of symmetric interactions (MSI), the complement of MSI (cMSI), and global inhibition (GI). Each model stores a cyclic sequence of four orthogonal patterns ($p = 4$). Each color represents the overlap with a particular pattern ξ^μ, which is retrieved when the overlap m^μ approaches 1. As different overlaps sequentially increase and decrease, the patterns are retrieved one after another.

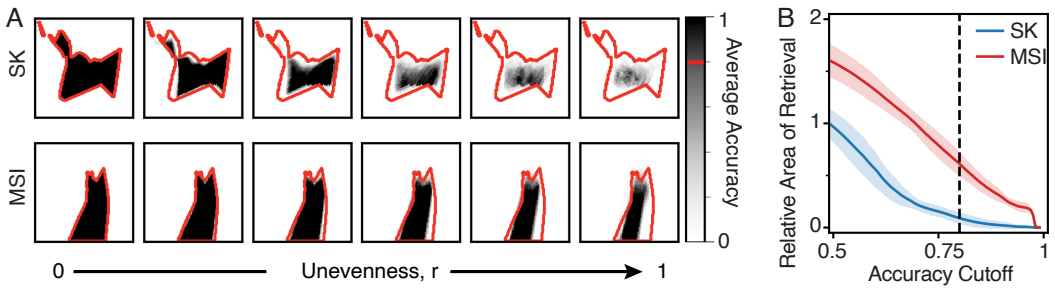

Figure S2: Sequential retrieval of patterns with different activities and ordering. (A) Phase diagrams for each model at different unevenness r of activity levels. The parameter ranges are the same as for Fig. 2 in the main text. The retrieval accuracy at each point (θ, λ) is first binarized according to a cutoff of 0.8, and then averaged over all permutations. This average accuracy is colored using a grayscale, and the red contours represent the shaded area for $r = 0$. The ratio of the gray area to the red contoured region is the relative area plotted against unevenness r in Fig. 3A. (B) The relative area of retrieval for unevenness $r = 1$ plotted against different choices of accuracy cutoff. The dashed line corresponds to a cutoff of 0.8 used in panel (A).

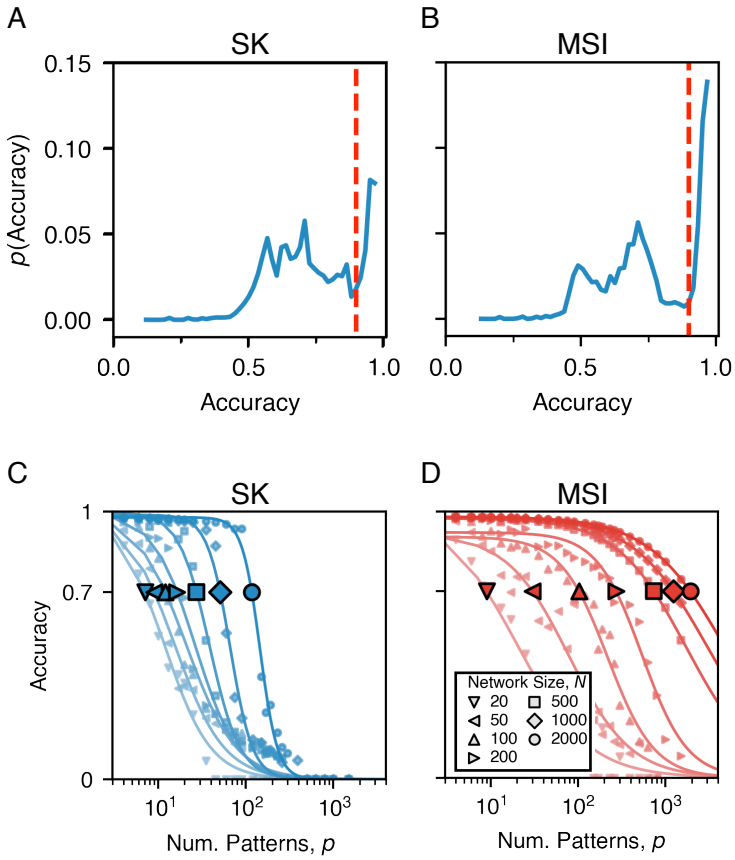

Figure S3: (A–B) The distributions of accuracies over the parameter space in Fig. 2 marginalized over a (scores of zero are omitted for clarity). The rightmost peak, corresponding to accurate retrieval, is separated by an accuracy cutoff of 0.9. (C–D) The average accuracy of retrieval plotted against the number of patterns p. The accuracy is calculated for a sequence of p randomly generated patterns with $a = 0.3$, averaged over 100 realizations. A sigmoidal curve is fitted for each N, and the critical number of patterns p_c is defined as the value of p where the average accuracy falls below 0.7, indicated by the marker. The parameters used to simulate each model are the same as for Fig. S1.

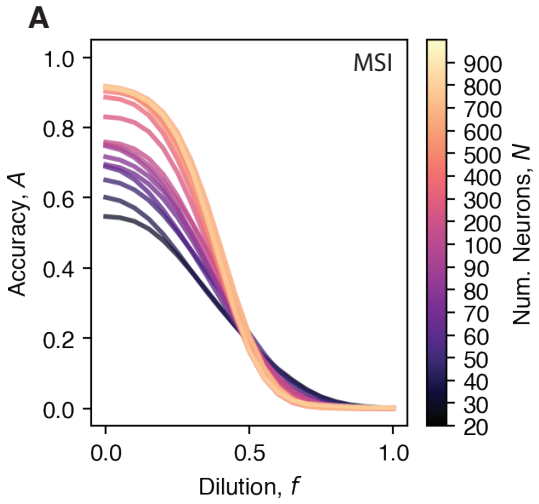

Figure S4: Retrieval accuracy as a function of the dilution fraction of modulated interactions. (A) Different curves represent different network sizes; the number of patterns is $p = \alpha N$ where $\alpha = 0.2$. Each curve is averaged over 100 trials and smoothed with a Gaussian kernel of width 0.1. The simulation parameters are $(\lambda, \theta) = (0.17, 0.06)$.