

RESEARCH ARTICLE

WILEY

Computational constraints on the associative recall of spatial scenes

Kwang Il Ryom¹  | Debora Stendardi² | Elisa Ciaramelli²  | Alessandro Treves¹ 

¹SISSA - Cognitive Neuroscience, Trieste, Italy

²Dipartimento di Psicologia Renzo Canestrari, Università di Bologna, Bologna, Italy

Correspondence

Alessandro Treves, SISSA - Cognitive Neuroscience, via Bonomea 265, 34136 Trieste, Italy.
Email: ale@sisssa.it

Funding information

Ministero dell'Università e della Ricerca, Grant/Award Numbers: PRIN 20174TPEFG, "TRIPS"

Abstract

We consider a model of associative storage and retrieval of compositional memories in an extended cortical network. Our model network is comprised of Potts units, which represent patches of cortex, interacting through long-range connections. The critical assumption is that a memory, for example of a spatial view, is composed of a limited number of items, each of which has a pre-established representation: storing a new memory only involves acquiring the connections, if novel, among the participating items. The model is shown to have a much lower storage capacity than when it stores simple unitary representations. It is also shown that an input from the hippocampus facilitates associative retrieval. When it is absent, it is advantageous to cue rare rather than frequent items. The implications of these results for emerging trends in empirical research are discussed.

KEYWORDS

associative recall, computational constraints, compositional memory, Potts network, spatial scenes

1 | INTRODUCTION

In one of the first recorded attempts to dissect declarative memory, Dante Alighieri, in *canto XVII* of his *Purgatorio* (≈ 1315 , see (Alighieri, n.d.)), distinguishes three forms of increasing complexity of what he calls the *imaginative* faculty, the ability to recall facts and events not currently conveyed by sensory inputs. First, in verses 19–21, he exemplifies the ability to retrieve a single female character, from Greek mythology, from a cue comprised of a single major attribute—her metamorphosis from woman into nightingale; this can be accomplished, in principle, by a simple pattern associator (Edmund, 1998). Second, in twice as many verses, 25–30, Dante's imagination recalls an entire scene, the death of Haman from the biblical Scroll of Esther. Although concisely reported, the scene is suggested to be more vivid and rich in detail than the paired associate recall, and also *compositional* in nature: it includes elements, like the death on the cross and the Persian Shah (in this instance, Ahasueros),

which likely feature in several other unrelated scenes or events that the poet might have recalled. How can we conceive the neural mechanisms required to evoke such a compositional snapshot?

The third form of imagination, exemplified with an episode drawn from the poem by Vergil, Dante's guide, is even more complex, since it involves Princess Lavinia speaking and trying to understand what went on in her mother's mind that drove her to commit suicide. Recalling the episode therefore implies the recursive activation of theory-of-mind schemata: we need to get into Lavinia's thoughts as she attempts to comprehend the dynamical unfolding of her mother's drama by getting into her thoughts.

The ability to recall facts and events not currently conveyed by sensory inputs is integral to mind-wandering, the drifting of the mind away from current (sensory) experience toward inner contents such as memories or plans (Christoff et al., 2016; Ciaramelli & Treves, 2019). Recent research has begun to investigate the neural underpinnings of mind-wandering, and to reveal distinct patterns of

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Hippocampus* published by Wiley Periodicals LLC.

alteration of mind-wandering, following brain damage, which roughly parallel Dante's poetic intuition of three levels of complexity in the imaginative faculty. Patients with lesions in the ventro-medial prefrontal cortex (vmPFC) tend to mind-wander less than healthy and brain-damaged controls, and when they do they are more focused on the present and on the self, suggesting a deficit in activating dynamical schemata to self-project into imaginary situations different from the perceptual present, such as future events or others' perspectives (Bertossi & Ciaramelli, 2016). Hippocampal patients, on the other hand, report mind-wandering as frequently as healthy controls, but their thoughts are of a streamlined logical/semantic character, impoverished in spatial details and bereft of episodic contributions, particularly from the recent past, the last year or so of actual experiences (McCormick et al., 2018). It thus appears that vmPFC integrity is necessary for the self-initiation and unfolding of mind-wandering episodes, perhaps imbuing them also with self-related value (Rolls, 2022a), whereas hippocampus integrity is important for the composition of elements drawn from recent experience into imagined scenes that fuel mind-wandering, whether or not they closely match combinations of elements that actually occurred (Barry & Maguire, 2019; Ciaramelli et al., 2019; Ciaramelli & Treves, 2019; Hassabis et al., 2007; McCormick et al., 2018).

Why should it be so? After all, influential memory theories promote the idea that, after hippocampally driven consolidation, even episodic memories should become independent of the hippocampus (Marr, 1971; Squire et al., 2015). One such theory viewed the hippocampus as a complementary learning system, needed because the cortex, just like a back-propagation trained network, is postulated to be able to only learn slowly (McClelland et al., 1995); logically, once the cortex has taken its time, the hippocampus can be disposed of. The Multiple Trace Theory has emphasized instead the qualitative distinction between truly episodic memories that remain dependent on the hippocampus through a lifetime, and semanticized memory content that can be retrieved and utilized also without the hippocampus (Moscovitch et al., 2005). A somewhat intermediate formulation has been put forward, to try and reconcile the contradictory empirical evidence, which can be invoked in partial support of either extreme position: it holds that the hippocampus regenerates constructs that appear to be simple reactivations of the activity patterns originally encoded, but are not (Barry & Maguire, 2019). By titrating the degree of infidelity of the reactivated from the original, this proposal can satisfactorily interpolate between views that *prima facie* clash with each other.

None of the above, however, really addresses any constraints that may arise below the functional system level, that is, in the neural network mechanisms that are invoked to implement the required operations of memory storage and reactivation. An exception may be the argument that rapid neocortical learning would lead to catastrophic interference (McClelland et al., 1995), although it was later qualified that this would only happen with new content *inconsistent* with previously stored information (McClelland, 2013). Episodic memories, however, are typically neither fully consistent nor inconsistent with each other, rather, they are diverse, entailing a variably overlapping set of items.

We ask here whether there are purely computational constraints that require cooperation between the hippocampus and neocortex in the associative storage and retrieval of snapshot compositional memories, and which stem from the distinct neural network organization of the hippocampus and of vmPFC (and neocortex in general). The hippocampus has available the dentate gyrus, which can establish a new, tendentially orthogonal compressed representation for any new memory (Treves & Rolls, 1992). In the neocortex there is no dentate gyrus, but its presumably large storage capacity—particularly in humans—should allow for the associative storage of many new combinations of items, most of which are already endowed with their neuronal representations. To what extent is this the case?

2 | THE POTTS MODEL OF AN EXTENDED MEMORY NETWORK

The Potts associative network offers a tractable mathematical model (as shown from the early study by Ido Kanter (1988)) of how compositional memories may be stored over large swathes of neocortex and later retrieved through purely associative mechanisms. In contrast to the neural networks usually considered in machine learning, it is recurrent (also at large scale) rather than directional, lacking *a priori* defined input and output layers, and the connection weights are modified by a model associative rule, with no involvement of error signals. As a memory device, therefore, it operates by self-organization, without instructions or extra ingredients, as envisaged by Braitenberg (Braitenberg & Schüz, 2013). In the cortex, obviously, such self-organization, for example, what is stored where, is grounded in cortical connectivity, the overall framework for which is largely genetically determined. The Potts network does not however deal with localization issues, in order to focus solely on the interactions between patches of neocortex, mediated by axons travelling through the white matter. It subsumes the dynamics within each patch into the suitably defined Potts units. The correspondence with the underlying full-fledged cortical network has been discussed in (Naim et al., 2018).

2.1 | Potts units that model cortical patches

A Potts unit, introduced in statistical physics in 1952 (Potts, 1952), has S states pointing each along a different dimension and can be regarded for our purposes as representing a local subnetwork or cortical patch of real neurons, endowed with its set of dynamical attractors; these span different directions in activity space, and are assimilated to the states of the Potts unit. One can then define an autoassociative network of N Potts units interacting through tensor connections (Naim et al., 2018). The memories are stored in the weight matrix of the network during a learning phase, described only summarily, as in the Hopfield model (Hopfield, 1982): each memory μ is a vector or list of the states taken in the overall activity configuration to be stored by each unit i : ξ_i^μ . We label by the index k each of

the S states a Potts unit can take (assuming for simplicity S to be the same throughout the network) and we consider an extra quiescent state, $k=0$, when the unit does not participate in the activity configuration of the memory. Therefore, $k=0, \dots, S$, and each ξ_i^μ can take values in the same categorical set.

For the standard, quasi-orthogonal memories, for which it is assumed that units are assigned an activity state independently of each other and of any information already present in the network, the tensor weights read (Kanter, 1988).

$$J_{ij}^{kl} = \frac{1}{Na(1-\frac{a}{S})} \sum_{\mu=1}^p \left(\delta_{\xi_i^\mu k} - \frac{a}{S} \right) \left(\delta_{\xi_j^\mu l} - \frac{a}{S} \right) (1 - \delta_{k0})(1 - \delta_{l0}), \quad (1)$$

where i, j denote units, k, l denote states, a is the fraction of units active in each memory, and the δ 's are Kronecker symbols. This form of the tensor weight has been shown to effectively recapitulate associative, 'Hebbian' synaptic plasticity in the underlying neural network (Naim et al., 2018) and, as in Hopfield-like models, the subtraction of the mean activity per state a/S ensures a higher storage capacity (Kropff & Treves, 2005). We shall see how this *learning rule* has to be modified if we consider, in the case of compositional memories, that information about their components is already present in the network.

Locally, each patch of cortex tends to align its activity with one of the attractor states or to remain in the null state, which is expressed in the Potts model by

$$\sigma_i^k = \frac{\exp(\beta r_i^k)}{\sum_{l=1}^S \exp(\beta r_i^l) + \exp[\beta(\theta_i^0 + U)]} \quad (2)$$

and

$$\sigma_i^0 = \frac{\exp[\beta(\theta_i^0 + U)]}{\sum_{l=1}^S \exp(\beta r_i^l) + \exp[\beta(\theta_i^0 + U)]}, \quad (3)$$

where r_i^k is the input variable tending to align unit i to state k . This occurs within a time scale τ_1 , as expressed by the equations below. U and θ_i^0 are the common and unit-specific components of an effective threshold, that incorporates a variety of inhibitory effects. From Equations (2) and (3), we see that $\sum_{k=0}^S \sigma_i^k \equiv 1$, and note also that σ_i^k takes *continuous* values in the (0,1) range for each k , whereas the memories, for simplicity, are assumed discrete, implying that perfect retrieval is approached when $\sigma_i^k \simeq 1$ for $k = \xi_i^\mu$ and $\simeq 0$ otherwise. Equations (2) and (3) are SoftMax functions, and $1/\beta$ is an effective temperature parametrizing the level of noise in the cortical patch represented by individual Potts units (when $\beta \rightarrow \infty$, each unit expresses a single non-zero σ).

2.2 | Potts model dynamics

When the Potts model is studied as a model of cortical *dynamics*, for example, of latching dynamics (Russo & Treves, 2012), $U(t)$ and $\theta_i^0(t)$

are functions of time, to represent the temporal course of long-range (indirect) and local inhibition. In addition, state-specific thresholds $\theta_i^k(t)$ are used to represent firing rate adaptation effects (Ryom et al., 2021). Here, however, we are only interested in modeling associative retrieval of a content-addressable memory, that is, a simple rapid operation that succeeds or fails within the time scale τ_1 of neuronal activation. We do not consider, therefore, slow inhibition and adaptation processes, and assume for simplicity a single common component of the threshold, $U(t)$, which varies rapidly enough to regulate the overall activity in the network at a constant level,

$$\sum_{i=1}^N \sum_{k=1}^S \sigma_i^k = Na. \quad (4)$$

While this assumption is rather unrealistic, it allows assessing the associative capability of the network under ideal conditions, free of the temporal fluctuations of unregulated dynamics. The time evolution of the network is then governed by the equations

$$\tau_1 \frac{dr_i^k(t)}{dt} = h_i^k(t) - r_i^k(t). \quad (5)$$

The input that the unit i in state k receives is the sum of an *internal* component, that is, coming from the other Potts units, and a term θ_i^k , which here is used to model not adaptation (as in the latching dynamics studies (Russo & Treves, 2012)) but a sustained external signal, considered constant over the time scale of retrieval, such as that which can come from the hippocampus (we use it with a minus sign for consistency, as θ usually denotes a threshold).

$$h_i^k = \sum_{j \neq i}^N \sum_{l=1}^S J_{ij}^{kl} \sigma_j^l - \theta_i^k. \quad (6)$$

The way the patterns $\bar{\xi}^\mu$ are generated, that is, their probability distribution, has effects on the retrieval properties of the network, that is, on its ability to retrieve with good accuracy one of the stored memories, if this is partially cued. A quantitative measure of this ability of the network is the *storage capacity*, the number of memories the network model is able to store and retrieve, given values for its various parameters.

2.3 | Simple memories

Simple memories in the cortex would be assigned distributed representation over several Potts units, exhibiting a higher level of (self-)organization than simple memories in the hippocampus, which are taken to be distributed over many individual neurons. Still, if across Potts units memory patterns are nearly orthogonal, that is, randomly correlated, like those assumed to be established by the dentate gyrus of the hippocampus, the Potts model equipped with Equation (1) can store and retrieve an extensive

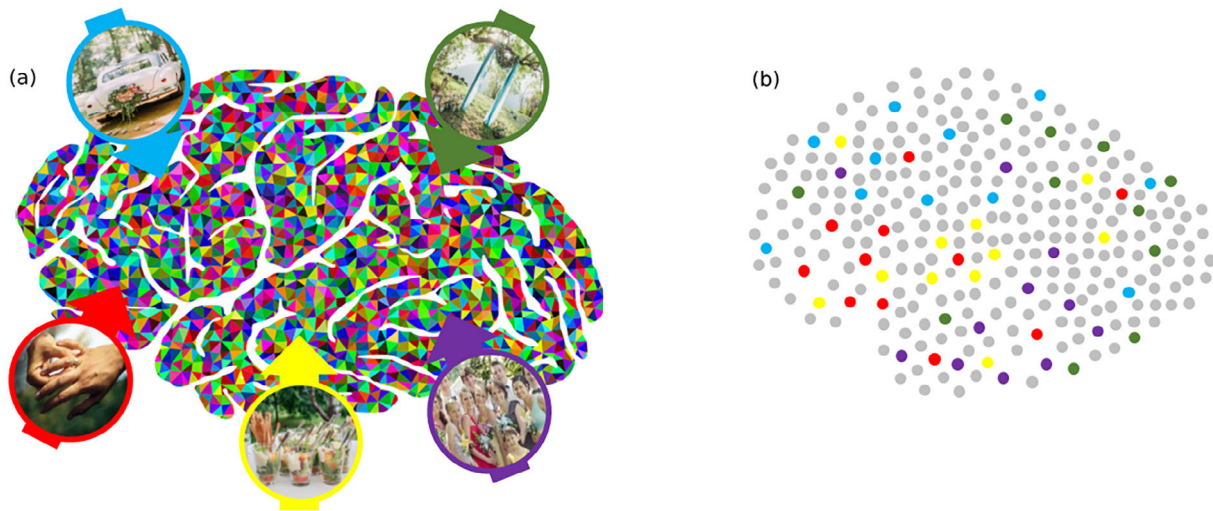


FIGURE 1 Compositional memories as represented in a Potts neural network. The episodic memory of a friends' wedding relies on a composition of items already given, at least in some instantiation, a long-term representation (a), such as the canopy, the people present, the food, the rings and the car with which they leave. In the Potts model (b), these elements are assigned sparse distributed representations over several Potts units, each loosely standing for a small patch of cortex.

number of patterns and each pattern has a large basin of attraction (Figure 2a,b).

What if memory representations have a nontrivial structure, rather than randomly correlated? In the next section, we examine the retrieval properties of Potts neural networks when memories have a semi-naturalistic internal structure, in terms of items, which are defined as percepts that are included whole in several memories. For example, a farm can be a familiar percept participating in the memory of several autobiographical events that have taken place at the farm. This implies that the representations of those memories are unlikely to be randomly correlated: they share (at least) the item farm.

2.4 | Compositional memories

We take our model compositional memories to embody a further level of organization and to include Z items, each of which has now a distributed cortical representation over several Potts units (Figure 1b). Two conceptually distinct stages of learning are, therefore, envisaged: first, the representations of the items are stored and subsequently, if a pair of items appears in a compositional memory, the tensor connections between the corresponding Potts units are strengthened. In practice, a novel item may happen to be stored only the first time it is included in an episodic memory; but here we are interested in the capacity of the model for retrieval, not in detailing the learning process. Across memories, some items may appear more frequently than others. We consider a pool of K items. Each memory can contain items with different frequency, from rare to very frequent ones.

We denote with η_i^ρ ($\rho = 1, 2, \dots, K$), the activity patterns, of sparsity of $a' = a/Z$, which represent the items. Here a is the sparsity, that is, the fraction of active units, of the memories themselves, ξ_μ^i ($\mu = 1, 2, \dots, p$), and the details of how we compose the memory

patterns from those of the items are explained in Appendix A.1. Here, we only emphasize that the network model does not discuss the *localization* of representations in distinct cortical areas, a major and relevant issue which is however beyond the scope of this study. The item activity patterns are randomly scattered across the Potts units.

The connection weights are set differently than in Equation (1), to express the notion, inherent to the compositional construct, that once an item has been encoded onto the synaptic connections, it is there and it is not stored again every time that item is present in the input,

$$J_{ij}^{kl} = (1 - \delta_{k0})(1 - \delta_{l0}) \left[\frac{1}{Na' \left(1 - \frac{a'}{S}\right)} \sum_{\rho=1}^K \left(\delta_{\eta_i^\rho k} - \frac{a'}{S} \right) \left(\delta_{\eta_j^\rho l} - \frac{a'}{S} \right) \right. \quad (7)$$

$$\left. + \frac{1}{2Na \left(1 - \frac{a}{S}\right)} \sum_{\rho=1}^K \sum_{\sigma=1}^K F_{\rho\sigma} \left(\delta_{\eta_i^\rho k} - \frac{a}{S} \right) \left(\delta_{\eta_j^\rho l} - \frac{a}{S} \right) \right],$$

where $F_{\rho\sigma}$ is 1 if a pair of items (ρ, σ) appears together in one of the p memories, and 0 otherwise. Thus, the two lines above reflect the two stages of learning envisaged. That is, while the first term of Equation (7) reflects one-shot associative learning of individual items, assumed to have occurred before, the second term likewise stores relations between items included at least once in the same compositional, episodic memory, and again the pair is stored once even if it recurs in multiple memories. Note that the prefactor with a' in the denominator makes the single-item term stronger than the pair-of-items term, as $1/a < 1/a' = Z/a$. Note also that more complex, for example, iterative and non-associative processes involved in acquiring the individual items in memory are not considered in the present model for simplicity, but they would not necessarily affect the constraints we focus on here, which are those arising from the associative storage not of items but of unique compositions of items.

2.5 | Retrieval cues versus Hippocampal input

To simply cue the network we activate a fraction f of the units active in a given memory, concentrated within some of the items of that memory, and let the network evolve without further external input. For example, for $f = 0.5$, when $Z = 5$, the cue is applied essentially to all the units active in two of the items, and to half of those active in a third (minor adjustments are due to the coincidence of some of the active units). With memories including both rare and frequent items, we consider applying a cue concentrated at either end of the frequency spectrum.

To model hippocampal inputs operating at retrieval, instead, we assume that the hippocampus has reinstated a compressed representation of the entire memory, and is able to convey a corresponding signal to all the units of the Potts network, which unlike the cue is sustained over the time course of retrieval. We express that through the state-specific thresholds, θ_i^k , by setting, for memory μ ,

$$\theta_i^k = -\gamma \delta_{\mu, k}^i \quad (8)$$

so that γ regulates the intensity of facilitation. Note that this θ_i^k is taken to be constant in time. The model thus allows contrasting two neural mechanisms for the reactivation of compositional memory: in the former, it is up to the long-range cortical connections, while in the latter, the hippocampus does it, in line with the *indexing* theory (Teyler & DiScenna, 1986), leaving to cortical connections only to retrieve the detailed content of each item.

3 | RESULTS

3.1 | A strong constraint on compositional memories

First, for the sake of analytical clarity, we start from a simple case, in which all items appear with the same (average) frequency in the compositional memories: we vary the number of memories p , compose each by drawing from a common pool of K items, and set the other parameters at their default values, specified in Table 1, including the number of items per memory $Z = 5$. The C++ code is available in the OSF repository. Note that, when for example $p = 300$, items appear on average in 5 distinct memories each, if $K = 300$ as well, and in as many as 15 memories each, if $K = 100$. This increases the difficulty of maintaining the unique item configuration of the compositional memory, even though it is present in the full cue (Figure 2a), and once $p = 400$, compositional memories are effectively inaccessible (the overlap, which measures the correlation of the retrieved activity with the stored representation, drops to zero); whereas simple unitary

TABLE 1 Parameters of the network.

| Symbol | Meaning | Default value |
|----------|-------------------------------|---------------|
| N | Number of Potts units | 1000 |
| S | Number of states per unit | 7 |
| p | Number of memory patterns | 200 |
| a | Sparsity of patterns | 0.2 |
| Z | Number of items per memory | 5 |
| K | Number of items in total | 200 |
| B | Number of bins | 20 |
| f | Fraction of units for cuing | 0.5 |
| γ | Strength of hippocampal input | 0.1 |

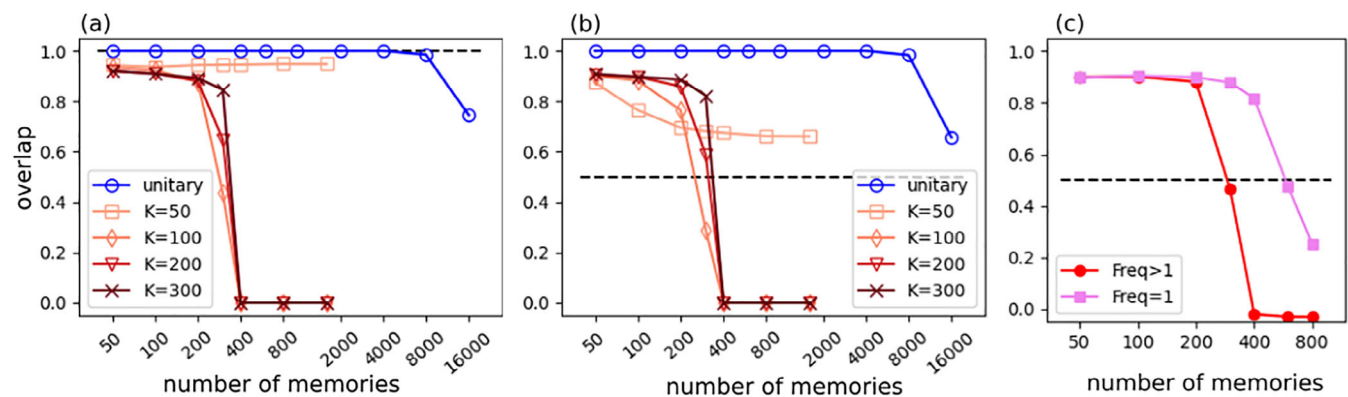


FIGURE 2 Unitary memories and compositional memories. (a) The overlap between stored patterns and retrieval states is plotted as a function of the total number p of stored memories. The network stores only one type of memory: either unitary (blue) or compositional memories with fixed frequency (red). The network starts from a perfect version of one memory pattern ($f = 1.0$) and is allowed to follow its dynamics until it reaches a retrieval state or a limiting time. From bright red to dark red, color encodes the number of items K in the pool. The blue curve is for random patterns (unitary memories). Network parameters are set at default values (see Table 1). (b) Similar to (a), but the network is partially cued by a memory pattern. The partial cue is prepared by flipping back a fraction $1 - f = 0.5$ of active units of the cued pattern into a quiescent state. (c) The same network stores two types of compositional memories: memories made of frequently used items (red) and those made of rarely used items (used once, pink). There are p (x -axis) memories in total, half in either category. The black dashed line indicates the initial value of the overlap (i.e., f).

memories (which can be conceived as comprised of non-repeated items) do not show a capacity limit, with our parameters, until $p = 16000$.

The apparent exception is, perhaps surprisingly, when the pool of items is very small, $K = 50$ —for those it appears that the network remains highly correlated with the cue, hence with the memory itself, all the way to high values of p . This is due to two effects, as clarified by Figure 2b. First, we can imagine the network as moving on a free-energy landscape (or its generalizations, the details are beyond our purposes here); for movement to be unimpeded, the landscape has to be smooth, which it is not for $K = 50$, due to the limited number of item representations dotting it. Now, the full cue does not really test the retrieval or pattern completion ability of the recurrent Potts network, but only its reluctance to drift away from the initial configuration of activity, already specified by the cue—and with a rough landscape the network is very reluctant, as it cannot effectively move. When using a partial cue, instead, for example, $f = 0.5$, the other overlap curves do not change much, but the one for $K = 50$ starts to drop already for $p > 50$. Second, if the cue maintains nevertheless activated the items it is applied to (3 out of 5, for $f = 0.5$), there is a substantial chance, if the pool is small, that also some of the remaining items will be those appearing in the memory to be retrieved. So we can consider the small K case as essentially an artifact, in any case irrelevant to human memory, which represents more than 50 items.

For larger K values on the other hand the constraint is real, and it can be understood to a first approximation by considering the individual items as robust blocs of units that can be reactivated coherently, while the challenge for the network lies in using the item pairings, stored in the connectivity, in order to retrieve the correct configuration of Z items. The challenge is tougher the more memories are stored, because more pairs of items will have been stored in the connections between the Potts units. Resorting to an argument developed many years ago for the Willshaw model (Golomb et al., 1990; Willshaw et al., 1969), we can estimate the probability that a pair of items has *not* been stored as the probability that it has not been stored in one memory, to the power p : $\text{Prob}(F_{\rho\sigma} = 0) = [1 - (Z/K)^2]^p \approx \exp(-pZ^2/K^2)$. If this probability becomes small, most item pairs will be encoded in the network, that will, therefore, find it difficult to select those in the compositional memory. This interference effect is reduced for large K , but then a complementary negative effect sets in, that the network is overloaded with items. Simulations show that the two effects complement each other and lead, irrespective of the K value, to an effective capacity much reduced with respect to that of unitary memories.

3.2 | Memories composed with frequent and with rare items

Figure 2c shows that the capacity constraints is almost as stringent also in a network that has stored compositional memories composed of frequent (hence, repeated) items, and *other* memories composed of rare items (in our model, appearing only once, hence unambiguously

individuating the episodic memory that includes them). The effective storage capacity for the latter is a bit higher, as the signal that leads from a partial cue to reactivate the complete configuration of items is clearer, but since the noise is contributed and felt by both frequent and rare memories once they share the same network, the difference is small. Note that in Figure 2c frequent items, from a pool of $K = 100$, are repeated as many times as those with $K = 200$ in Figure 2b, as they appear in half of the p memories.

We have also simulated a network storing half compositional and half unitary memories. Unitary memories can also be conceived as composed of items appearing only once; the difference with the case above is in the learning rule, which in the compositional case of Equation (7) assigns more weight to the individual items, because of the prefactor $1/d'$. Overall, however, the interference resulting from the storage of the other memories is similar, and so is the resulting storage capacity for compositional and unitary memories (not shown). Note that if the latter were alone, many more of them could be stored, but since they share the connection space with compositional memories, their effective “storage capacity” is almost the same as that of compositional memories.

3.3 | Scale-free memories

Next, we consider a more realistic case in which memories include items of different frequencies. We proceed as follows: we group items into B bins, indexed by $1, \dots, 1, \dots, B$, and each bin includes l items (Figure 3a). Then a memory is assembled by combining Z items obtained by sampling bins evenly. This results in the few items in the first bins being picked up more frequently than the many in the later bins and, as one can easily show, in an approximately *scale-free* distribution of items across memories (here, scale refers to frequency; see also the Appendix A).

Figure 3b (lightest green curve) shows that diversity in the distribution of item frequency has an adverse effect on storage capacity. A suitable comparison is between a scale-free distribution of items in $B = 20$ bins, which implies $B(B+1)/2 = 210$ items overall, the lightest green curve, and compositional memories with fixed-frequency items drawn from a pool of $K = 200$ (the red curve). The comparison indicates that the more realistic, mixed distribution of item frequencies, coexisting within the same memories, does not solve the capacity constraint imposed on compositional memories; if anything, it makes it somewhat worse.

3.4 | Hippocampal inputs

The results above indicate that memory retrieval triggered by partial cues is inherently less effective with compositional memories, in which the component items have been stored on their own, than with unitary representations. This suggests that a more effective retrieval operation could be initiated by a full cue, possibly weak but full, that is, distributed over all the component items. Such a cue could come

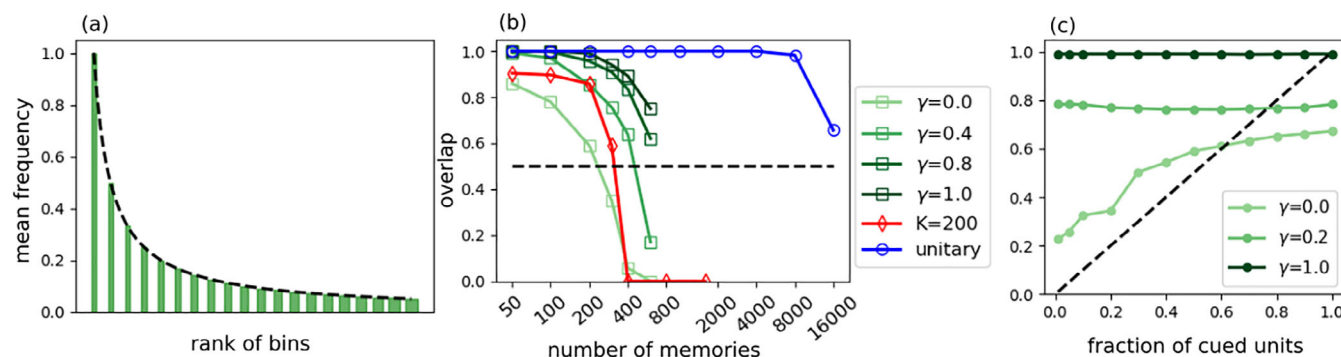


FIGURE 3 Scale-free distribution of item frequency. (a) An example of distribution of item frequency with $B = 20$ bins. Bins are arranged according to the frequency of items they include along the x-axis, with frequency indicated by bin height, while bin width alludes to the number of items per bin. (b) Retrieval for memories comprised of items following the frequency distribution given in (a). Color encodes γ values (i.e., the strength of hippocampal inputs). $f = 0.5, B = 20$. The red curve is for single frequency items, as in Figure 2b. (c) Similar to (b), the overlap is shown as a function of f for $p = 200$.

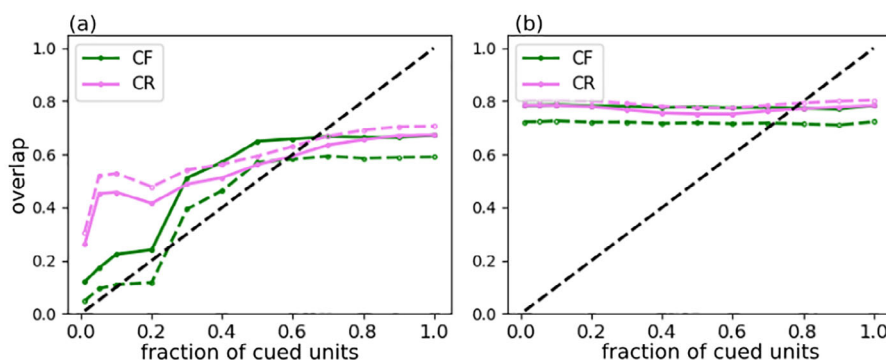


FIGURE 4 Effect of item frequency on retrieval. (a) Two methods of cueing the network with a partial version of memory patterns are compared by different colors. A fraction f of units among N_a active units for a memory pattern is chosen preferentially from frequent items (green) or from rare items (violet). $p = 200, \gamma = 0.0, B = 20$. (b) Same as (a) but for $\gamma = 0.2$, that is, with some hippocampal input. Dashed curves show results that include only a subset of memories that include very rare or very frequent items (from the last or first two bins in the distribution).

from an auxiliary compressed representation of the full memory, of the type that the hippocampus has been widely thought to store and retrieve, in turn, from partial cue.

To explore this hypothetical mechanism, we add a model hippocampal input to the compositional representation in the extended cortical network; following (Ryom et al., 2021), this is simply a sustained external contribution to the signal aligning each Potts units toward the activation state it has in the memory to be retrieved. It is parametrized by a variable γ . In formulas, using Kronecker's δ we write

$$h_i^k = \sum_{j(\neq i)=1}^N \sum_{k=1}^S J_{ij}^{kl} \delta_{\sigma_j l} + \gamma \delta_{\sigma_i^k} \quad (9)$$

Obviously, when the factor γ is large enough, successful retrieval is expected to be merely transferred from the hippocampus to the neocortex, in this model, with the latter not performing any significant role. As shown in Figure 3b, however, simulation results are complex.

On the one hand, the sustained input enhances network capacity, the more the stronger it is, but without really removing the capacity limit for compositional memories, indicated by the drop in all green curves at $p = 600$. On the other hand, also a weak hippocampal input produces a noticeable effect, when the fraction of the standard partial cue is $f = 0.5$. When $p = 200$, Figure 3c shows that even a weak sustained input, $\gamma = 0.2$, leads to retrieval to a level midway to that obtained with $\gamma = 1.0$, and as a function of f the same level is reached in the entire range $0.01 < f < 1.0$: in practice the hippocampal input requires only a minimal additional cue—and also when this is absent ($f \approx 0.0$) hippocampally triggered retrieval is effective on its own.

3.5 | Triggering retrieval from frequent or from rare items

Given the interference caused by the multiple pairings of frequent items with others, in retrieving compositional memories, one may

wonder whether the operation is more effective if triggered by the reactivation of the rarer items. This can be examined in the model simply by applying the partial cue f to the Potts units active in the representation of the rare vs. the frequent items. In Figure 4, this is done considering model scale-invariant representations produced with $B = 20$ bins, and applying the cue at either end of the frequency spectrum (solid curves); or by selecting memories with at least one item from the first two (frequent) or the last two (rare) bins, and averaging only over either restricted subset (dashed lines). Figure 4a shows that without the model hippocampal input, $\gamma = 0.0$, there is a marked effect of where the cue is applied, but only for $f \leq 0.2$, that is, effectively when a single item is cued. In that case cueing a frequent item is ineffective, while cueing a rare one is (partially) effective, although the correlation with the full memory is still far from ideal (overlap just above 0.4, or 2/5 items retrieved). With weak hippocampal input, $\gamma = 0.2$, retrieval is still incomplete, but the effect of where the partial cue was applied is virtually erased.

4 | DISCUSSION

In this study, we look at purely computational constraints for the retrieval of episodic, compositional memories, which turn out to be relatively complex to analyze, despite the artificial simplicity of the assumptions in-built in the model we have considered. To assess the range of validity of the results obtained, it is therefore useful to review the main assumptions:

- Compositional episodic memories are conceived statistically as being structured in terms of items, independently drawn from a pool of such items, with no further substructure. For example, the image of a groom can be composed with that of a thick wood as well as with that of a lawn, even though weddings more often take place on the latter. In further work we shall relax this assumption by introducing structured schemata into the model.
- Two distinct modes of content-addressing an episodic memory are envisaged. In the first, a partial cue sets in the active state the Potts units relative to a fraction f of the items composing the memory—which is intended to correspond to the initial alignment of some patches of cortex along the local attractors which represent those items, while the rest of the cortex is not aligned to anything.
- In the second mode, the hippocampus provides a sustained cue of possibly limited strength, but delivered to all relevant patches of cortex—therefore, a hippocampal index in Teyler and DiScenna's sense (Teyler & DiScenna, 1986) rather than a partial cue.
- The hippocampal representation of a compositional episodic memory, if it exists, is assumed to be unitary and not compositional, hence unrelated to the detailed semantic content of each item.
- Other simplifying assumptions are more “technical,” as they relate to the Potts neural network model.

Obviously all such assumptions are extreme, and relaxing them results in some form of interpolation. This can be regarded as a general limitation of an approach which, in the trade-off between clarity and plausibility, favors the former.

4.1 | Compositionality effectively shrinks the cortex

The model offers a number of theoretical insights. One of the main findings is that the storage capacity that had been previously calculated for unitary representations (Kanter, 1988; Kropff & Treves, 2005) is much higher and essentially irrelevant to that for compositional representations. The storage capacity for compositional representations is indeed constrained by factors that should be investigated further: the statistics of compositionality, the (long-range) connectivity, the plasticity that underlies the acquisition of compositional memories (expressed in the model by the “learning rule” adopted). The key finding, that is, the low storage capacity for compositional representations, may seem counter-intuitive: using a representation preassembled in blocks of units—the items—makes recall more difficult instead of facilitating it. The computational reason is that associative retrieval, in general, is robust to the interference of other memories if these produce uncorrelated fluctuations (i.e., the noise, in a signal-to-noise analysis) over many units or many small groups of units. If the fluctuations are coherent over large chunks of cortex, because they represent interfering items, the noise does not average out so well. It is as if compositionality nullified the key advantage of the cortex for memory—its sheer size—by obstructing the approach to the “law of large numbers,” that is, the mutual cancellation of random fluctuations, which is key to associative retrieval. The pre-assembly effectively reduces the size of the network to Z , the (average) number of items in a compositional memory; of course, only from the point of view of associative retrieval (in other respects, for example, for representational capacity, the cortex remains huge).

4.2 | Without the hippocampus, rare elements facilitate recall

Rare elements are those shared between relatively fewer memories. The effect demonstrated in Figure 4 reflects indeed the lower confusion associated with the retrieval cue coming from those items—they have established fewer strengthened connections to other items, and therefore are less likely to trigger the retrieval of multiple compatible memories. With the parameters we have adopted, the effect is not huge and limited to very partial cues (small f). Analyzing how it may scale up when cues are more detailed and the network more closely simulates a human cortex is beyond the scope of this work. We note for now that this effect, the advantage of cueing rare elements, vanishes once the hippocampus, in our model, provides a sustained full cue, even if weak, suggesting that the contribution of the

hippocampus is vital to retrieve compositional memories involving highly frequent items.

4.3 | The hippocampus helps, but by brute force

A final remark on the results is that Figure 3b indicates that the model hippocampal input does not really solve the low capacity problem. Whatever its strength γ , retrieval quality begins to decline at about the same memory load p . What happens is that in our model the hippocampus effectively takes over the retrieval task, and can send to the cortex a strong signal with its outcome, that the cortex would have been unable to get at on its own. Investigating a more significant cortical contribution, in this computational framework, probably requires a more articulated model, that we intend to analyse in future work.

5 | IMPLICATIONS FOR EMPIRICAL RESEARCH

The model can be related to a body of nascent theoretical notions and empirical data that seek to dissect the contribution of distinct brain structures to imaginative acts such as event (re)construction and mind-wandering (McCormick et al., 2018). One hypothesis is that vmPFC mediates schema-related relations among the objects in a scene, whereas the hippocampus assembles cohesive scenes (Gilboa & Marlatte, 2017; McCormick et al., 2018; Moscovitch et al., 2016). This is consistent with the evidence that constructed experience in patients with hippocampal lesions is rich in content but lacks spatial cohesiveness, whereas that of vmPFC patients also lacks (schema-based) constitutive elements (De Luca et al., 2018), and that mind-wandering is of poor episodic quality in hippocampal patients (McCormick et al., 2018) vs. severely reduced in vmPFC patients (Bertossi & Ciaramelli, 2016).

While the contribution of the hippocampus to event imagination may also be productively contrasted with that of other cortical areas, focusing on the division of labor between vmPFC and the hippocampus, a distinction that may turn out to be useful is the one analysed recently by Mullally and Maguire and involving 'Space Defining (SD)' and 'Space Ambiguous (SA) objects (Mullally & Maguire, 2011; Mullally & Maguire, 2013). Mullally and Maguire have suggested and shown empirically that SD objects promptly evoke a strong sense of a surrounding 3D space. An example SD object is a couch, which promptly evokes a sense of a surrounding 3D space compatible with a living room and not with other types of spatial layouts; SD objects define (identify) the space they fit in. A fly, an example SA object, does not. SA objects are compatible with, and shared between, many spatial layouts. Consistent with the prominent role of space processing for mental construction, SD objects are preferentially chosen as the initial building block to mentally construct a scene, and are picked last to be removed from a mental scene (Mullally & Maguire, 2013). Processing of SD and SA stimuli is associated with different activity in the parahippocampal cortex (Mullally & Maguire, 2011), the superior

temporal gyrus, and vmPFC (Monk et al., 2020), in line with the different functional properties of the two classes of items.

The SD-SA distinction must be considered together (and not confounded) with another independent distinction, that between objects that are more or less likely to be associated with other objects or related concepts (Bar & Aminoff, 2003), and hence trigger their activation (Mullally & Maguire, 2011). Although SD objects tend to be evocative of content (associated with other objects/concepts), as in the previous example of the couch, which can easily activate, in addition to the 3D space of a living room, the image of a nearby coffee table or TV, the SD/SA and contextual richness dimensions are distinct, and dissociable from the one another both behaviorally and neurally (Mullally & Maguire, 2011). We have recently isolated, for example, SD objects high in contextual associations (e.g.: swing), SD objects low in contextual associations (e.g.: chair), SA objects high in contextual associations (e.g.: fishing rod), and SA objects low in contextual associations (e.g.: belt) to be used as cues for event construction (Stendardi et al., in preparation).

In the present model, rare items can be taken to more immediately evoke a constellation of other items, because, being rare, they have been associated strongly with a small number of other items and contexts. This is the case for SD items, especially those with low levels of contextual associations, which evoke virtually unique contexts. Cueing a rare (e.g., SD) item is likely more effective in triggering memory retrieval, as competition between memories sharing that item is less likely. Our computational findings indicate that if and only if the hippocampal input is damaged or reduced, a partial cue applied to a rare item is more effective in triggering accurate memory retrieval than one applied to a frequent item. It would be interesting to investigate, therefore, if the advantage in event construction observed for SD versus SA items is more pronounced in the case of reduced input from the hippocampus, for example testing patients with hippocampal damage or using tasks that make heavier demands on neocortical regions versus the hippocampus (e.g., priming).

As of now, it remains unclear to what extent the model captures the *spatial* nature of memories for multiple items in visual scenes (which is integral to the SD/SA distinction), especially as it does not describe earlier visual processing (Devereux et al., 2018) nor the cortical connectivity that leads to item and scene representations (Rolls, 2022b); but it is clear that it fails to consider more structured constructs, usually referred to as *schemata*. These can be elaborated in at least two different dimensions. One is to consider schemata as groups of items that often occur together as components of wider compositional scenes, irrespective of exact timing relations (Audrain & McAndrews, 2022). For example, the memory of a countryside wedding may include a makeshift religious ceremony and a festive meal, each of which includes distinct items bound together into those two schemata. A second dimension is the temporal one. If two items A and B when they co-occur do so in a fixed succession, such as the discussion of the Thesis and the friends' congratulations, proper recall would entail reactivating their representation in the same order. Ultimately, along both dimensions one moves away from the snapshot character of simple episodic memories, taking some steps toward their *semantization*.

Developing our current computational model along the first dimension involves considering some form of nested probability distributions, which opens up a very large space of possibilities, so that it is probably wise to focus on a specific set of empirical data. Along the second dimension, instead, there is a straightforward neural mechanisms that favors the ordered reactivation of the representations of two items A and B: to enhance the connections from the units active in A to those active in B, and not viceversa (Spalla et al., 2021). If a spatial relation is captured, in part, by the availability of both options, scanning $A \rightarrow B$ as well as $B \rightarrow A$, a temporal relation singles out $A \rightarrow B$. Correctly reactivating all the temporal relations in an episode that has been experienced could be challenging for the cortex, but a partial reactivation that follows several originally distinct paths, making use of self-related (Stendardi et al., 2021) and other schemata, may in fact be the substrate for the generative process envisaged by Barry and Maguire (Barry & Maguire, 2019).

Ongoing work is exploring whether assigning such temporal structure to a simple composition of items can indeed facilitate their ordered reactivation, irrespective of the model hippocampal signal. If so, the model may be useful to better understand the phenomenology presented by vmPFC patients (Ciaramelli et al., 2019) and also, perhaps, the third and most sophisticated form of the imaginative faculty, as conceived by Dante.

AUTHOR CONTRIBUTIONS

The model was developed by Kwang Il Ryom and Alessandro Treves based on ongoing empirical research by Debora Stendardi and Elisa Ciaramelli. Simulations were run by Kwang Il Ryom, and the results discussed by all coauthors, who together wrote the paper.

ACKNOWLEDGMENTS

We are grateful to Andrea Tabarroni, professor of Medieval Philosophy, who showed us the passage from Dante's Comedy on the imaginative faculty. Open Access Funding provided by Scuola Internazionale Superiore di Studi Avanzati within the CRUI-CARE Agreement.

DATA AVAILABILITY STATEMENT

The simple code that has generated the simulations that support the findings of this study are openly available in the Open Science Foundation repository at <https://osf.io/3epj/>.

ORCID

Kwang Il Ryom  <https://orcid.org/0000-0002-4580-4374>

Elisa Ciaramelli  <https://orcid.org/0000-0002-4797-8060>

Alessandro Treves  <https://orcid.org/0000-0001-7246-5673>

REFERENCES

- Alighieri, D. (n.d.). Divine Comedy. Purgatorio, canto XVII, ca. 1315. English translation by Courtney Langdon. Retrieved August 25, 2022 from <https://oll.libertyfund.org/title/langdon-the-divine-comedy-vol-2-purgatorio-english-only-trans>
- Arita, M. (2005). Scale-freeness and biological networks. *Journal of Biochemistry*, 138(1), 1–4.
- Audrain, S., & McAndrews, M. P. (2022). Schemas provide a scaffold for neocortical integration of new memories over time. *Nature Communications*, 13(1), 1–16.
- Bar, M., & Aminoff, E. (2003). Cortical analysis of visual context. *Neuron*, 38(2), 347–358.
- Barry, D. N., & Maguire, E. A. (2019). Remote memory and the hippocampus: A constructive critique. *Trends in Cognitive Sciences*, 23(2), 128–142.
- Bertossi, E., & Ciaramelli, E. (2016). Ventromedial prefrontal damage reduces mind-wandering and biases its temporal focus. *Social Cognitive and Affective Neuroscience*, 11(11), 1783–1791.
- Braitenberg, V., & Schüz, A. (2013). *Cortex: statistics and geometry of neuronal connectivity*. Springer Science & Business Media.
- Christoff, K., Irving, Z. C., Fox, K. C. R., Nathan Spreng, R., & Andrews-Hanna, J. R. (2016). Mind-wandering as spontaneous thought: A dynamic framework. *Nature Reviews Neuroscience*, 17(11), 718–731.
- Ciaramelli, E., De Luca, F., Monk, A. M., McCormick, C., & Maguire, E. A. (2019). What “wins” in vmPFC: Scenes, situations, or schema? *Neuroscience & Biobehavioral Reviews*, 100, 208–210.
- Ciaramelli, E., & Treves, A. (2019). A mind free to wander: Neural and computational constraints on spontaneous thought. *Frontiers in Psychology*, 10, 39.
- De Luca, F., McCormick, C., Mullally, S. L., Intraub, H., Maguire, E. A., & Ciaramelli, E. (2018). Boundary extension is attenuated in patients with ventromedial prefrontal cortex damage. *Cortex*, 108, 1–12.
- Devereux, B. J., Clarke, A., & Tyler, L. K. (2018). Integrated deep visual and semantic attractor neural networks predict fMRI pattern-information along the ventral object processing pathway. *Scientific Reports*, 8(1), 1–12.
- Edmund, T. (1998). Rolls and Alessandro Treves. *Neural Networks and Brain Function*.
- Gilboa, A., & Marlatte, H. (2017). Neurobiology of schemas and schema-mediated memory. *Trends in Cognitive Sciences*, 21(8), 618–631.
- Golomb, D., Rubin, N., & Sompolinsky, H. (1990). Willshaw model: Associative memory with sparse coding and low firing rates. *Physical Review A*, 41(4), 1843–1854.
- Hassabis, D., Kumaran, D., Vann, S. D., & Maguire, E. A. (2007). Patients with hippocampal amnesia cannot imagine new experiences. *Proceedings of the National Academy of Sciences*, 104(5), 1726–1731.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8), 2554–2558.
- Kanter, I. (1988). Potts-glass models of neural networks. *Physical Review A*, 37(7), 2739–2742.
- Kropff, E., & Treves, A. (2005). The storage capacity of Potts models for semantic memory retrieval. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(8), P08010.
- Lima-Mendez, G., & Van Helden, J. (2009). The powerful law of the power law and other myths in network biology. *Molecular BioSystems*, 5(12), 1482–1493.
- Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 262, 23–81.
- McClelland, J. L. (2013). Incorporating rapid neocortical learning of new schema-consistent information into complementary learning systems theory. *Journal of Experimental Psychology: General*, 142(4), 1190–1210.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3), 419.
- McCormick, C., Rosenthal, C. R., Miller, T. D., & Maguire, E. A. (2018). Mind-wandering in people with hippocampal damage. *Journal of Neuroscience*, 38(11), 2745–2754.
- Monk, A. M., Barnes, G. R., & Maguire, E. A. (2020). The effect of object type on building scene imagery—An MEG study. *Frontiers in Human Neuroscience*, 14, 592175.
- Morris Moscovitch, R., Rosenbaum, S., Gilboa, A., Addis, D. R., Westmacott, R., Grady, C., McAndrews, M. P., Levine, B., Black, S., Winocur, G., et al. (2005). Functional neuroanatomy of remote

- episodic, semantic and spatial memory: a unified account based on multiple trace theory. *Journal of Anatomy*, 207(1), 35–66.
- Moscovitch, M., Cabeza, R., Winocur, G., & Nadel, L. (2016). Episodic memory and beyond: the hippocampus and neocortex in transformation. *Annual Review of Psychology*, 67, 105–134.
- Mullally, S. L., & Maguire, E. A. (2011). A new role for the parahippocampal cortex in representing space. *Journal of Neuroscience*, 31(20), 7441–7449.
- Mullally, S. L., & Maguire, E. A. (2013). Exploring the role of space-defining objects in constructing and maintaining imagined scenes. *Brain and Cognition*, 82(1), 100–107.
- Naim, M., Boboeva, V., Kang, C. J., & Treves, A. (2018). Reducing a cortical network to a potts model yields storage capacity estimates. *Journal of Statistical Mechanics: Theory and Experiment*, 2018(4), 043304.
- Potts, R. B. (1952). Some generalized order-disorder transformations. In *Mathematical proceedings of the cambridge philosophical society* (Vol. 48, pp. 106–109). Cambridge University Press.
- Rolls, E. T. (2022a). The hippocampus, ventromedial prefrontal cortex, and episodic and semantic memory. *Progress in Neurobiology*, 102334.
- Rolls, E. T. (2022b). Hippocampal spatial view cells for memory and navigation, and their underlying connectivity in humans. *Hippocampus*. <https://doi.org/10.1002/hipo.23467>
- Ruderman, D., & Bialek, W. (1993). Statistics of natural images: Scaling in the woods. *Advances in Neural Information Processing Systems*, 6, 551–558.
- Russo, E., & Treves, A. (2012). Cortical free-association dynamics: Distinct phases of a latching network. *Physical Review E*, 85(5), 51920.
- Ryom, K. I., Boboeva, V., Soldatkina, O., & Treves, A. (2021). Latching dynamics as a basis for short-term recall. *PLoS Computational Biology*, 17(9), e1008809.
- Spalla, D., Cornacchia, I. M., & Treves, A. (2021). Continuous attractors for dynamic memories. *Elife*, 10, e69499.
- Squire, L. R., Genzel, L., Wixted, J. T., & Morris, R. G. (2015). Memory consolidation. *Cold Spring Harbor Perspectives in Biology*, 7(8), a021766.
- Stendardi, D., Biscotto, F., Bertossi, E., & Ciaramelli, E. (2021). Present and future self in memory: the role of vmPFC in the self-reference effect. *Social Cognitive and Affective Neuroscience*, 16(12), 1205–1213.
- Teyler, T. J., & DiScenna, P. (1986). The hippocampal memory indexing theory. *Behavioral Neuroscience*, 100(2), 147.
- Treves, A., & Rolls, E. T. (1992). Computational constraints suggest the need for two distinct input systems to the hippocampal ca3 network. *Hippocampus*, 2(2), 189–199.
- Willshaw, D. J., Peter Buneman, O., & Longuet-Higgins, H. C. (1969). Non-holographic associative memory. *Nature*, 222(5197), 960–962.

How to cite this article: Ryom, K. I., Stendardi, D., Ciaramelli, E., & Treves, A. (2023). Computational constraints on the associative recall of spatial scenes. *Hippocampus*, 1–11. <https://doi.org/10.1002/hipo.23511>

APPENDIX A

A.1 | Making memory representations

We construct representations of compositional memories in two steps. In the first step, we assign Z items to each memory. This is done either by sampling items evenly, so that on average they all occur with the same frequency, or unevenly, as described in the text, for example, with the quasi-scale-free procedure discussed below, and represented in Figure 3a. In the second step, we write a representation of each memory by merging representations of its Z items. The only issue in doing so is that there are some units that are shared by more than one item. This would lead to representations with sparsity (fraction of active units) less than a . In order to constrain all memories to have the same sparsity a , we compute the “fields” h_i^k of all units, by assuming that the Z items of a particular memory are activated. Then we select the Na units (and their states) which receive the largest field, to define them as the representation of this particular memory.

A.2 | Scale-free item frequency

Scale-free distributions have been invoked as a simple description of many natural phenomena, and there is considerable controversy as to the ideas that have been put forward (Arita, 2005; Lima-Mendez & Van Helden, 2009). There has been also considerable work on the scale-invariant distribution of objects of different sizes in natural scenes, which is closer to being relevant for the compositionality of memory for scenes (Ruderman & Bialek, 1993). Here our intent is merely practical, however: to generate a simple distribution of frequencies, which does not involve an extra arbitrary parameter. The distribution described in the text is approximately scale free, because no such parameter is introduced explicitly, although implicitly the number B of bins sets the upper and lower ends of the frequency range with which items are assigned to memories: from about pZ/B to pZ/B^2 times. Within this range, each “frequency scale” is approximately represented evenly.

A.3 | Default parameters

When not varied systematically, parameters of the Potts model are set as in Table 1. The C++ simulation code is available in the Open Science Foundation repository, <https://osf.io/z3epj/>.