

Abstract representations of events arise from mental errors in learning and memory

Christopher W. Lynn¹, Ari E. Kahn ^{2,3}, Nathaniel Nyema³ & Danielle S. Bassett ^{1,3,4,5,6,7}✉

Humans are adept at uncovering abstract associations in the world around them, yet the underlying mechanisms remain poorly understood. Intuitively, learning the higher-order structure of statistical relationships should involve complex mental processes. Here we propose an alternative perspective: that higher-order associations instead arise from natural errors in learning and memory. Using the free energy principle, which bridges information theory and Bayesian inference, we derive a maximum entropy model of people's internal representations of the transitions between stimuli. Importantly, our model (i) affords a concise analytic form, (ii) qualitatively explains the effects of transition network structure on human expectations, and (iii) quantitatively predicts human reaction times in probabilistic sequential motor tasks. Together, these results suggest that mental errors influence our abstract representations of the world in significant and predictable ways, with direct implications for the study and design of optimally learnable information sources.

¹ Department of Physics & Astronomy, College of Arts & Sciences, University of Pennsylvania, Philadelphia, PA 19104, USA. ² Department of Neuroscience, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104, USA. ³ Department of Bioengineering, School of Engineering & Applied Science, University of Pennsylvania, Philadelphia, PA 19104, USA. ⁴ Department of Electrical & Systems Engineering, School of Engineering & Applied Science, University of Pennsylvania, Philadelphia, PA 19104, USA. ⁵ Department of Neurology, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104, USA. ⁶ Department of Psychiatry, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104, USA. ⁷ Santa Fe Institute, Santa Fe, NM 87501, USA. ✉email: dsb@seas.upenn.edu

Our experience of the world is punctuated in time by discrete events, all connected by an architecture of hidden forces and causes. In order to form expectations about the future, one of the brain's primary functions is to infer the statistical structure underlying past experiences^{1–3}. In fact, even within the first year of life, infants reliably detect the frequency with which one phoneme follows another in spoken language⁴. By the time we reach adulthood, uncovering statistical relationships between items and events enables us to perform abstract reasoning⁵, identify visual patterns⁶, produce language⁷, develop social intuition⁸, and segment continuous streams of data into self-similar parcels⁹. Notably, each of these functions requires the brain to identify statistical regularities across a range of scales. It has long been known that people are sensitive to differences in individual transition probabilities such as those between words or concepts^{4,6}. In addition, mounting evidence suggests that humans can also infer abstract (or higher-order) statistical structures, including hierarchical patterns within sequences of stimuli¹⁰, temporal regularities on both global and local scales¹¹, abstract concepts within webs of semantic relationships¹², and general features of sparse data¹³.

To study this wide range of statistical structures in a unified framework, scientists have increasingly employed the language of network science¹⁴, wherein stimuli or states are conceptualized as nodes in a graph with edges or connections representing possible transitions between them. In this way, a sequence of stimuli often reflects a random walk along an underlying transition network^{15–17}, and we can begin to ask which network features give rise to variations in human learning and behavior. This perspective has been particularly useful, for example, in the study of artificial grammars¹⁸, wherein human subjects are tasked with inferring the grammar rules (i.e., the network of transitions between letters and words) underlying a fabricated language¹⁹. Complementary research in statistical learning has demonstrated that modules (i.e., communities of densely connected nodes) within transition networks are reflected in brain imaging data²⁰ and give rise to stark shifts in human reaction times²¹. Together, these efforts have culminated in a general realization that people's internal representations of a transition structure are strongly influenced by its higher-order organization^{22,23}. But how does the brain learn these abstract network features? Does the inference of higher-order relationships require sophisticated hierarchical learning algorithms? Or instead, do natural errors in cognition yield a “blurry” representation, making the coarse-grained architecture readily apparent?

To answer these questions, here we propose a single driving hypothesis: that when building models of the world, the brain is finely tuned to maximize accuracy while simultaneously minimizing computational complexity. Generally, this assumption stems from a rich history exploring the trade-off between brain function and computational cost^{24,25}, from sparse coding principles at the neuronal level²⁶ to the competition between information integration and segregation at the whole-brain level²⁷ to the notion of exploration versus exploitation²⁸ and the speed-accuracy trade-off²⁹ at the behavioral level. To formalize our hypothesis, we employ the free energy principle³⁰, which has become increasingly utilized to investigate constraints on cognitive functioning³¹ and explain how biological systems maintain efficient representations of the world around them³². Despite this thorough treatment of the accuracy-complexity trade-off in neuroscience and psychology, the prevailing intuition in statistical learning maintains that the brain is either optimized to perform Bayesian inference^{12,13}, which is inherently error free, or hierarchical learning^{10,11,16,18}, which typically entails increased rather than decreased computational complexity.

Here, we show that the competition between accuracy and computational complexity leads to a maximum entropy (or

minimum complexity) model of people's internal representations of events^{30,33}. As we decrease the complexity of our model, allowing mental errors to take effect, higher-order features of the transition network organically come into focus while the fine-scale structure fades away, thus providing a concise mechanism explaining how people infer abstract statistical relationships. To a broad audience, our model provides an accessible mapping from transition networks to human behaviors, with particular relevance for the study and design of optimally learnable transition structures—either between words in spoken and written language^{18,19,33}, notes in music³⁴, or even concepts in classroom lectures³⁵.

Results

Network effects on human expectations. In the cognitive sciences, mounting evidence suggests that human expectations depend critically on the higher-order features of transition networks^{15,16}. Here, we make this notion concrete with empirical evidence for higher-order network effects in a probabilistic sequential response task²². Specifically, we presented human subjects with sequences of stimuli on a computer screen, each stimulus depicting a row of five gray squares with one or two of the squares highlighted in red (Fig. 1a). In response to each stimulus, subjects were asked to press one or two computer keys mirroring the highlighted squares (Fig. 1b). Each of the 15 different stimuli represented a node in an underlying transition network, upon which a random walk stipulated the sequential order of stimuli (Fig. 1a). By measuring the speed with which a subject responded to each stimulus, we were able to infer their expectations about the transition structure: a fast reaction reflected a strongly anticipated transition, while a slow reaction reflected a weakly anticipated (or surprising) transition^{1,2,22,36}.

While it has long been known that humans can detect differences in transition probabilities—for instance, rare transitions lead to sharp increases in reaction times^{4,6}—more recently it has become clear that people's expectations also reflect the higher-order architecture of transition networks^{20–23,37}. To clearly study these higher-order effects without the confounding influence of variations in transition probabilities, here we only consider transition graphs with a uniform transition probability of 0.25 on each edge, thereby requiring nodes to have uniform degree $k = 4$ (Fig. 1c). Specifically, we consider two different graph topologies: a modular graph with three communities of five densely connected nodes and a lattice graph representing a 3×5 grid with periodic boundary conditions (Fig. 1d). As all transitions across both graphs have uniform probability, any systematic variations in behavior between different parts of a graph, or between the two graphs themselves, must stem from differences in the graphs' higher-order modular or lattice structures.

Regressing out the dependence of reaction times on the different button combinations (Fig. 1e), the natural quickening of reactions with time³⁸ (Fig. 1f), and the impact of stimulus recency (Methods), we identify two effects of higher-order network structure on subjects' reactions. First, in the modular graph we find that reactions corresponding to within-cluster transitions are 35 ms faster than reactions to between-cluster transitions ($p < 0.001$, F -test; Supplementary Table 1), an effect known as the cross-cluster surprisal^{22,37} (Fig. 2a). Similarly, we find that people are more likely to respond correctly for within-cluster transitions than between-cluster transitions (Supplementary Table 8). Second, across all transitions within each network, we find that reactions in the modular graph are 23 ms faster than those in the lattice graph ($p < 0.001$, F -test; Supplementary Table 2), a phenomenon that we coin the modular-lattice effect (Fig. 2b).

Thus far, we have assumed that variations in human behavior stem from people's internal expectations about the network

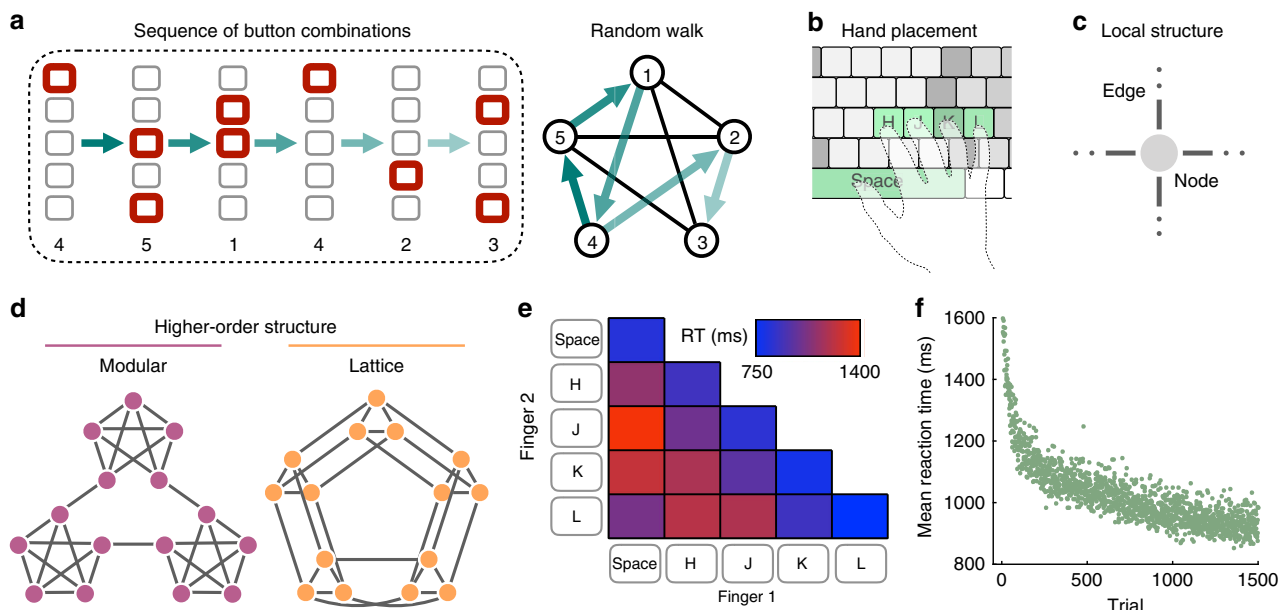


Fig. 1 Subjects respond to sequences of stimuli drawn as random walks on an underlying transition graph. **a** Example sequence of visual stimuli (left) representing a random walk on an underlying transition network (right). **b** For each stimulus, subjects are asked to respond by pressing a combination of one or two buttons on a keyboard. **c** Each of the 15 possible button combinations corresponds to a node in the transition network. We only consider networks with nodes of uniform degree $k = 4$ and edges with uniform transition probability 0.25. **d** Subjects were asked to respond to sequences of 1500 such nodes drawn from two different transition architectures: a modular graph (left) and a lattice graph (right). **e** Average reaction times for the different button combinations, where the diagonal elements represent single-button presses and the off-diagonal elements represent two-button presses. **f** Average reaction times as a function of trial number, characterized by a steep drop-off in the first 500 trials followed by a gradual decline in the remaining 1000 trials. In **e** and **f**, averages are taken over responses during random walks on the modular and lattice graphs. Source data are provided as a Source Data file.

structure. However, it is important to consider the possible confound of stimulus recency: the tendency for people to respond more quickly to stimuli that have appeared more recently^{39,40}. To ensure that the observed network effects are not simply driven by recency, we performed a separate experiment that controlled for recency in the modular graph by presenting subjects with sequences of stimuli drawn according to Hamiltonian walks, which visit each node exactly once²⁰. Within the Hamiltonian walks, we still detect a significant cross-cluster surprisal effect (Fig. 2a; Supplementary Tables 3–5). In addition, we controlled for recency in our initial random walk experiments by focusing on stimuli that previously appeared a specific number of trials in the past. Within these recency-controlled data, we find that both the cross-cluster surprisal and modular-lattice effects remain significant (Supplementary Figs. 2 and 3). Finally, for all of our analyses throughout the paper we regress out the dependence of reaction times on stimulus recency (Methods). Together, these results demonstrate that higher-order network effects on human behavior cannot be explained by recency alone.

In combination, our experimental observations indicate that people are sensitive to the higher-order architecture of transition networks. But how do people infer abstract features like community structure from sequences of stimuli? In what follows, we turn to the free energy principle to show that a possible answer lies in understanding the subtle role of mental errors.

Network effects reveal errors in graph learning. As humans observe a sequence of stimuli or events, they construct an internal representation \hat{A} of the transition structure, where \hat{A}_{ij} represents the expected probability of transitioning from node i to node j . Given a running tally n_{ij} of the number of times each transition has occurred, one might naively expect that the human brain is optimized to learn the true transition structure as accurately as

possible^{41,42}. This common hypothesis is represented by the maximum likelihood estimate⁴³, taking the simple form

$$\hat{A}_{ij}^{\text{MLE}} = \frac{n_{ij}}{\sum_k n_{ik}}. \tag{1}$$

To see that human behavior does not reflect maximum likelihood estimation, we note that Eq. (1) provides an unbiased estimate of the transition structure⁴³; that is, the estimated transition probabilities in \hat{A}^{MLE} are evenly distributed about their true value 0.25, independent of the higher-order transition structure. Thus, the fact that people’s reaction times depend systematically on abstract features of the network marks a clear deviation from maximum likelihood estimation. To understand how higher-order network structure impacts people’s internal representations, we must delve deeper into the learning process itself.

Consider a sequence of nodes (x_1, x_2) , where $x_t \in \{1, \dots, N\}$ represents the node observed at time t and N is the size of the network (here $N = 15$ for all graphs). To update the maximum likelihood estimate of the transition structure at time $t + 1$, one increments the counts n_{ij} using the following recursive rule,

$$n_{ij}(t + 1) = n_{ij}(t) + [i = x_t][j = x_{t+1}], \tag{2}$$

where the Iverson bracket $[\cdot] = 1$ if its argument is true and 0 otherwise. Importantly, we note that at each time $t + 1$, a person must recall the previous node that occurred at time t ; in other words, they must associate a cause x_t to each effect x_{t+1} that they observe. Although maximum likelihood estimation requires perfect recollection of the previous node at each step, human errors in perception and recall are inevitable^{44–46}. A more plausible scenario is that, when attempting to recall the node at time t , a person instead remembers the node at time $t - \Delta t$ with some decreasing probability $P(\Delta t)$, where $\Delta t \geq 0$. This memory distribution, in turn, generates an internal belief about which

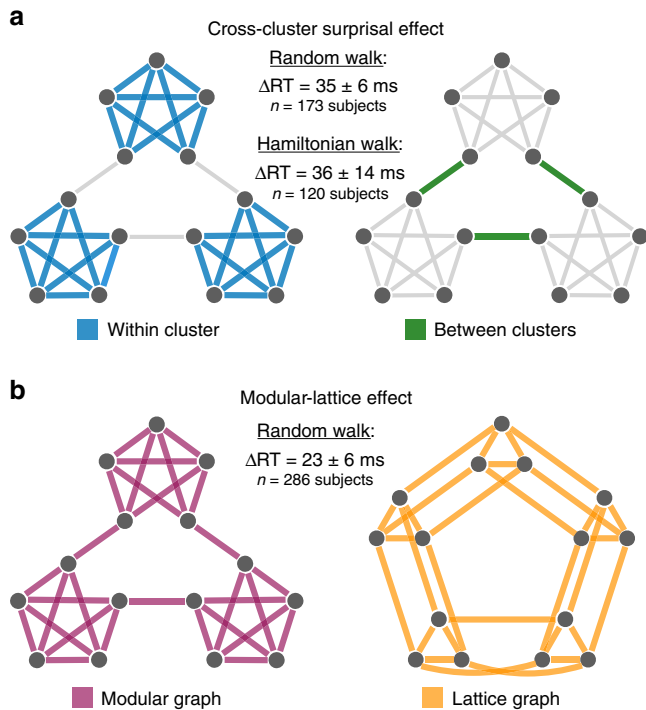


Fig. 2 The effects of higher-order network structure on human reaction times. **a** Cross-cluster surprisal effect in the modular graph, defined by an average increase in reaction times for between-cluster transitions (right) relative to within-cluster transitions (left). We detect significant differences in reaction times for random walks ($p < 0.001$, $t = 5.77$, $df = 1.61 \times 10^5$) and Hamiltonian walks ($p = 0.010$, $t = 2.59$, $df = 1.31 \times 10^4$). For the mixed effects models used to estimate these effects, see Supplementary Tables 1 and 3. **b** Modular-lattice effect, characterized by an overall increase in reaction times in the lattice graph (right) relative to the modular graph (left). We detect a significant difference in reaction times for random walks ($p < 0.001$, $t = 3.95$, $df = 3.33 \times 10^5$); see Supplementary Table 2 for the mixed effects model. Measurements were on independent subjects, statistical significance was computed using two-sided F -tests, and confidence intervals represent standard deviations. Source data are provided as a Source Data file.

node occurred at time t ,

$$B_t(i) = \sum_{\Delta t=0}^{t-1} P(\Delta t)[i = x_{t-\Delta t}]. \quad (3)$$

Updating Eq. (2) accordingly, we arrive at a learning rule that accounts for natural errors in perception and recall,

$$\tilde{n}_{ij}(t+1) = \tilde{n}_{ij}(t) + B_t(i)[j = x_{t+1}]. \quad (4)$$

Using this revised counting rule, we can begin to form more realistic predictions about people's internal estimates of the transition structure, $\hat{A}_{ij} = \tilde{n}_{ij} / \sum_k \tilde{n}_{ik}$.

We remark that $P(\Delta t)$ does not represent the forgetting of past stimuli altogether; instead, it reflects the local shuffling of stimuli in time. If one were to forget past stimuli at some fixed rate – a process that is important for some cognitive functions⁴⁷—this would merely introduce white noise into the maximum likelihood estimate \hat{A}^{MLE} (Supplementary Discussion). By contrast, we will see that, by shuffling the order of stimuli in time, people are able to gather information about the higher-order structure of the underlying transitions.

Choosing a memory distribution: the free energy principle. In order to make predictions about people's expectations, we must

choose a particular mathematical form for the memory distribution $P(\Delta t)$. To do so, we begin with a single driving hypothesis: that the brain is finely tuned to (i) minimize errors and (ii) minimize computational complexity. Formally, we define the error of a recalled stimulus to be its distance in time from the desired stimulus (i.e., Δt), such that the average error of a candidate distribution $Q(\Delta t)$ is given by $E(Q) = \sum_{\Delta t} Q(\Delta t)\Delta t$. By contrast, it might seem difficult to formalize the computational complexity associated with a distribution Q . Intuitively, we would like the complexity of Q to increase with increasing certainty. Moreover, as a first approximation we expect the complexity to be approximately additive such that the cost of storing two independent memories equals the costs of the two memories themselves. As famously shown by Shannon, these two criteria of monotonicity and additivity are sufficient to derive a quantitative definition of complexity³³—namely, the negative entropy $-S(Q) = \sum_{\Delta t} Q(\Delta t)\log Q(\Delta t)$.

Together, the total cost of a distribution Q is its free energy $F(Q) = \beta E(Q) - S(Q)$, where β is the inverse temperature parameter, which quantifies the relative value that the brain places on accuracy versus efficiency³¹. In this way, our assumption about resource constraints in the brain leads to a particular form for P : it should be the distribution that minimizes $F(Q)$, namely the Boltzmann distribution³⁰

$$P(\Delta t) = \frac{1}{Z} e^{-\beta \Delta t}, \quad (5)$$

where Z is the normalizing constant (Methods). Free energy arguments similar to the one presented here have been used increasingly to formalize constraints on cognitive functions^{31,32}, with applications from active inference⁴⁸ and Bayesian learning under uncertainty³² to human action and perception with temporal or computational limitations^{31,49,50}. Taken together, Eqs. (3)–(5) define our maximum entropy model of people's internal transition estimates \hat{A} .

To gain an intuition for the model, we consider the infinite-time limit, such that the transition estimates become independent of the particular random walk chosen for analysis. Given a transition matrix A , one can show that the asymptotic estimates in our model are equivalent to an average over walks of various lengths, $\hat{A} = \sum_{\Delta t} P(\Delta t)A^{\Delta t+1}$, which, in turn, can be fashioned into the following analytic expression,

$$\hat{A} = (1 - e^{-\beta})A(I - e^{-\beta}A)^{-1}, \quad (6)$$

where I is the identity matrix (Methods). The model contains a single free parameter β , which represents the precision of a person's mental representation. In the limit $\beta \rightarrow \infty$ (no mental errors), our model becomes equivalent to maximum likelihood estimation (Fig. 3a), and the asymptotic estimates \hat{A} converge to the true transition structure A (Fig. 3b), as expected⁵¹. Conversely, in the limit $\beta \rightarrow 0$ (overwhelming mental errors), the memory distribution $P(\Delta t)$ becomes uniform across all past nodes (Fig. 3a), and the mental representation \hat{A} loses all resemblance to the true structure A (Fig. 3b).

Remarkably, for intermediate values of β , higher-order features of the transition network, such as communities of densely connected nodes, come into focus, while some of the fine-scale features, like the edges between communities, fade away (Fig. 3b). Applying Eq. (6) to the modular graph, we find that the average expected probability of within-community transitions reaches over 1.6 times the estimated probability of between-community transitions (Fig. 3c), thus explaining the cross-cluster surprisal effect^{22,37}. Furthermore, we find that the average estimated transition probabilities in the modular graph reach over 1.4 times the estimated probabilities in the lattice graph (Fig. 3d), thereby

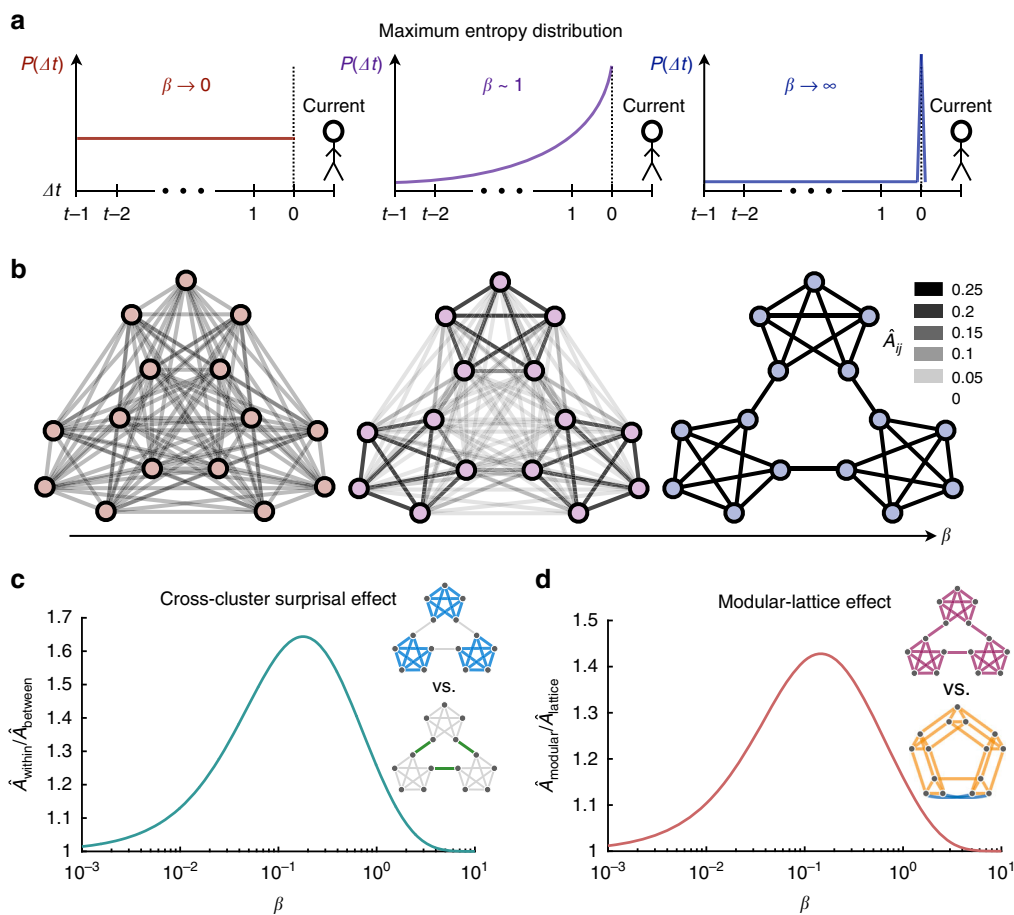


Fig. 3 A maximum entropy model of transition probability estimates in humans. **a** Illustration of the maximum entropy distribution $P(\Delta t)$ representing the probability of recalling a stimulus Δt time steps from the target stimulus (dashed line). In the limit $\beta \rightarrow 0$, the distribution becomes uniform over all past stimuli (left). In the opposite limit $\beta \rightarrow \infty$, the distribution becomes a delta function on the desired stimulus (right). For intermediate amounts of noise, the distribution drops off monotonically (center). **b** Resulting internal estimates \hat{A} of the transition structure. For $\beta \rightarrow 0$, the estimates become all-to-all, losing any resemblance to the true structure (left), while for $\beta \rightarrow \infty$, the transition estimates become exact (right). At intermediate precision, the higher-order community structure organically comes into focus (center). **c, d** Predictions of the cross-cluster surprisal effect (**c**) and the modular-lattice effect (**d**) as functions of the inverse temperature β .

predicting the modular-lattice effect. In addition to these higher-order effects, we find that the model also explains previously reported variations in human expectations at the level of individual nodes^{4,6,22} (Supplementary Fig. 1). Together, these results demonstrate that the maximum entropy model predicts the qualitative effects of network structure on human reaction times. But can we use the same ideas to quantitatively predict the behavior of particular individuals?

Predicting the behavior of individual humans. To model the behavior of individual subjects, we relate the transition estimates in Eqs. (3)–(5) to predictions about people’s reaction times. Given a sequence of nodes x_1, \dots, x_{t-1} , we note that the reaction to the next node x_t is determined by the expected probability of transitioning from x_{t-1} to x_t calculated at time $t - 1$, which we denote by $a(t) = \hat{A}_{x_{t-1},x_t}(t - 1)$. From this internal anticipation $a(t)$, the simplest possible prediction $\hat{r}(t)$ for a person’s reaction time is given by the linear relationship⁵² $\hat{r}(t) = r_0 + r_1 a(t)$, where the intercept r_0 represents a person’s reaction time with zero anticipation and the slope r_1 quantifies the strength of the relationship between a person’s reactions and their anticipation in our model⁵³.

To estimate the parameters β , r_0 , and r_1 that best describe a given individual, we minimize the root mean squared error

(RMSE) between their predicted and observed reaction times after regressing out the dependencies on button combination, trial number, and recency (Fig. 1e, f; Methods). The distributions of the estimated parameters are shown in Fig. 4a–b for random walks and in Fig. 4g–h for Hamiltonian walks. Among the 358 random walk sequences in the modular and lattice graphs (across 286 subjects; Methods), 40 were best described as performing maximum likelihood estimation ($\beta \rightarrow \infty$) and 73 seemed to lack any notion of the transition structure whatsoever ($\beta \rightarrow 0$), while among the remaining 245 sequences, the average inverse temperature was $\beta = 0.30$. Meanwhile, among the 120 subjects that responded to Hamiltonian walk sequences, 81 appeared to have a non-trivial value of β , with an average of $\beta = 0.61$. Interestingly, these estimates of β roughly correspond to the values for which our model predicts the strongest network effects (Fig. 3c, d). In the following section, we will compare these values of β , which are estimated indirectly from people’s reaction times, with direct measurements of β in an independent memory experiment.

In addition to estimating β , we also wish to determine whether our model accurately describes individual behavior. Toward this end, we first note that the average slope r_1 is large (-735 ms for random walks and -767 ms for Hamiltonian walks), suggesting that the transition estimates in our model $a(t)$ are strongly

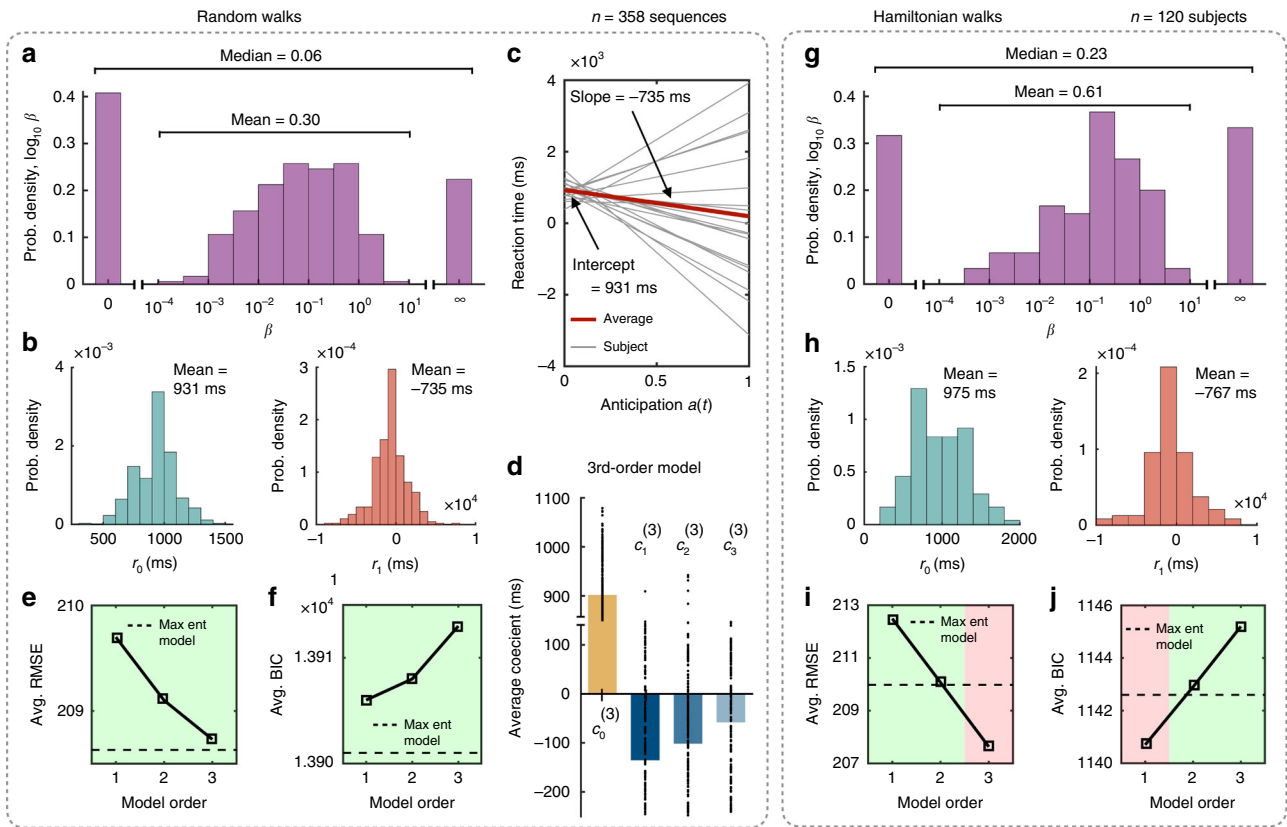


Fig. 4 Predicting reaction times for individual subjects. **a–f** Estimated parameters and accuracy analysis for our maximum entropy model across 358 random walk sequences (across 286 subjects; Methods). **a** For the inverse temperature β , 40 sequences corresponded to the limit $\beta \rightarrow \infty$, 73 corresponded to the limit $\beta \rightarrow 0$. Among the remaining 245 sequences, the average value of β was 0.30. **b** Distributions of the intercept r_0 (left) and slope r_1 (right). **c** Predicted reaction time as a function of a subject’s internal anticipation. Gray lines indicate 20 randomly selected sequences, and the red line shows the average prediction for all sequences. **d** Linear parameters for the third-order competing model; data points represent individual sequences and bars represent averages. **e, f** Comparing the performance of our maximum entropy model with the hierarchy of competing models up to third-order. Root mean squared error (RMSE; **e**) and Bayesian information criterion (BIC; **f**) of our model averaged over all sequences (dashed lines) compared to the competing models (solid lines); our model provides the best description of the data across all models considered. **g–j** Estimated parameters and accuracy analysis for our maximum entropy model across all Hamiltonian walk sequences (120 subjects). **g** For the inverse temperature β , 20 subjects were best described as performing maximum likelihood estimation ($\beta \rightarrow \infty$), 19 lacked any notion of the transition structure ($\beta \rightarrow 0$), and the remaining 81 subjects had an average value of $\beta = 0.61$. **h** Distributions of the intercept r_0 (left) and slope r_1 (right). **i** Average RMSE of our model (dashed line) compared to that of the competing models (solid line); our model maintains higher accuracy than the competing hierarchy up to the second-order model. **j** Average BIC of the maximum entropy model (dashed line) compared to that of the competing models (solid line); our model provides a better description of the data than the second- or third-order models. Source data are provided as a Source Data file.

predictive of human reaction times, and negative, confirming the intuition that increased anticipation yields decreased reaction times (Fig. 4b, h). To examine the accuracy of our model \hat{r} , we consider a hierarchy of competing models $\hat{r}^{(\ell)}$, which represent the hypothesis that humans learn explicit representations of the higher-order transition structure. In particular, we denote the ℓ th-order transition matrix by $\hat{A}_{ij}^{(\ell)} = n_{ij}^{(\ell)} / \sum_k n_{ik}^{(\ell)}$, where $n_{ij}^{(\ell)}$ counts the number of observed transitions from node i to node j in ℓ steps. The model hierarchy takes into account increasingly higher-order transitions, such that the ℓ th-order model contains perfect information about transitions up to length ℓ :

$$\begin{aligned}
 \hat{r}^{(0)}(t) &= c_0^{(0)}, \\
 \hat{r}^{(1)}(t) &= c_0^{(1)} + c_1^{(1)} a^{(1)}(t), \\
 &\vdots \\
 \hat{r}^{(\ell)}(t) &= c_0^{(\ell)} + \sum_{k=1}^{\ell} c_k^{(\ell)} a^{(k)}(t),
 \end{aligned}
 \tag{7}$$

where $a^{(k)}(t) = \hat{A}_{x_{t-1}, x_t}^{(k)}(t-1)$. Each model $\hat{r}^{(\ell)}$ contains $\ell + 1$ parameters $c_0^{(\ell)}, \dots, c_{\ell}^{(\ell)}$, where $c_k^{(\ell)}$ quantifies the predictive power of the k th-order transition structure.

Intuitively, for each model $\hat{r}^{(\ell)}$, we expect $c_1^{(\ell)}, c_2^{(\ell)}, \dots$, to be negative, reflecting a decrease in reaction times due to increased anticipation, and decreasing in magnitude, such that higher-order transitions are progressively less predictive of people’s reaction times. Indeed, considering the third-order model $\hat{r}^{(3)}$ as an example, we find that progressively higher-order transitions are less predictive of human reactions (Fig. 4d). However, even the largest coefficient ($c_1^{(3)} = -135$ ms) is much smaller than the slope in our maximum entropy model ($r_1 = -735$ ms), indicating that the representation \hat{A} is more strongly predictive of people’s reaction times than any of the explicit representations $\hat{A}^{(1)}, \hat{A}^{(2)}, \dots$. Indeed, averaging over the random walk sequences, the maximum entropy model achieves higher accuracy than the first three orders of the competing model hierarchy (Fig. 4e)—this is despite the fact that the third-order model even contains one more parameter. To account for differences in the number of

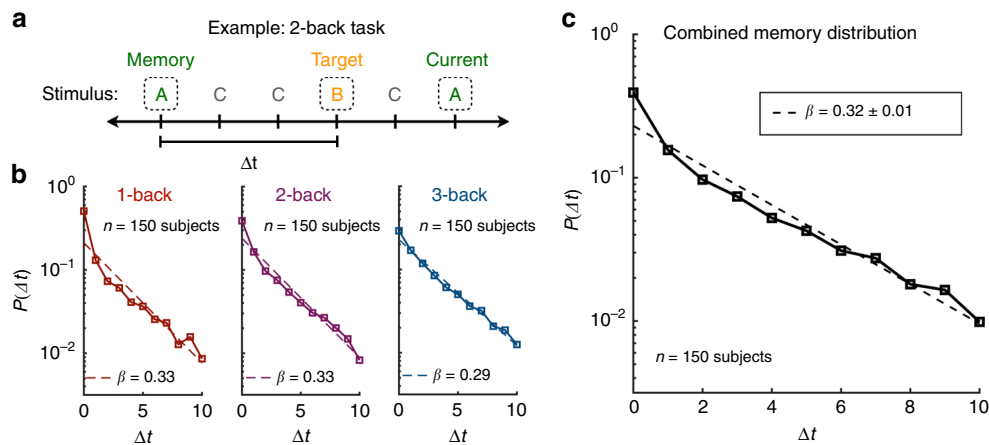


Fig. 5 Measuring the memory distribution in an n -back experiment. **a** Example of the 2-back memory task. Subjects view a sequence of stimuli (letters) and respond to each stimulus indicating whether it matches the target stimulus from two trials before. For each positive response that the current stimulus matches the target, we measure Δt by calculating the number of trials between the last instance of the current stimulus and the target. **b** Histograms of Δt (i.e., measurements of the memory distribution $P(\Delta t)$) across all subjects in the 1-, 2-, and 3-back tasks. Dashed lines indicate exponential fits to the observed distributions. The inverse temperature β is estimated for each task to be the negative slope of the exponential fit. **c** Memory distribution aggregated across the three n -back tasks. Dashed line indicates an exponential fit. We report a combined estimate of the inverse temperature $\beta = 0.32 \pm 0.01$, where the standard deviation is estimated from 1000 bootstrap samples of the combined data. Measurements were on independent subjects. Source data are provided as a Source Data file.

parameters, we additionally compare the average Bayesian information criterion (BIC) of our model with that of the competing models, finding that the maximum entropy model provides the best description of the data (Fig. 4f).

Similarly, averaging over the Hamiltonian walk sequences, the maximum entropy model provides more accurate predictions than the first two competing models (Fig. 4i) and provides a lower BIC than the second and third competing models (Fig. 4j). Notably, even in Hamiltonian walks, the maximum entropy model provides a better description of human reaction times than the second-order competing model, which has the same number of parameters. However, we remark that the first-order competing model has a lower BIC than the maximum entropy model (Fig. 4j), suggesting that humans may focus on first-order rather than higher-order statistics during Hamiltonian walks—an interesting direction for future research. On the whole, these results indicate that the free energy principle, and the resulting maximum entropy model, is consistently more effective at describing human reactions than the hypothesis that people learn explicit representations of the higher-order transition structure.

Directly probing the memory distribution. Throughout our discussion, we have argued that errors in memory shape human representations in predictable ways, a perspective that has received increasing attention in recent years^{47,54,55}. Although our framework explains specific aspects of human behavior, there exist alternative perspectives that might yield similar predictions. For example, one could imagine a Bayesian learner with a non-Markov prior that “integrates” the transition structure over time, even without sustaining errors in memory or learning. In addition, Eq. (6) resembles the successor representation in reinforcement learning^{56,57}, which assumes that, rather than shuffling the order of past stimuli, humans are instead planning their responses multiple steps in advance (Supplementary Discussion). In order to distinguish our framework from these alternatives, here we provide direct evidence for precisely the types of mental errors predicted by our model.

In the construction and testing of our model, we have developed a series of predictions concerning the shape of the

memory distribution $P(\Delta t)$, which, to recall, represents the probability of remembering the stimulus at time $t - \Delta t$ instead of the target stimulus at time t . We first assumed that $P(\Delta t)$ decreases monotonically. Second, to make quantitative predictions, we employed the free energy principle, leading to the prediction that P drops off exponentially quickly with Δt (Eq. (5)). Finally, when fitting the model to individual subjects, we estimated an average inverse temperature β between 0.30 for random walks and 0.61 for Hamiltonian walks.

To test these three predictions directly, we conducted a standard n -back memory experiment. Specifically, we presented subjects with sequences of letters on a screen, and they were asked to respond to each letter indicating whether or not it was the same as the letter that occurred n steps previously; for each subject, this process was repeated for the three conditions $n = 1, 2$, and 3. To measure the memory distribution $P(\Delta t)$, we considered all trials on which a subject responded positively that the current stimulus matched the target. For each such trial, we looked back to the last time that the subject did in fact observe the current stimulus and we recorded the distance (in trials) between this observation and the target (Fig. 5a). In this way, we were able to treat each positive response as a sample from the memory distribution $P(\Delta t)$.

The measurements of P for the 1-, 2-, and 3-back tasks are shown in Fig. 5b, and the combined measurement of P across all conditions is shown in Fig. 5c. Notably, the distributions decrease monotonically and maintain consistent exponential forms, even out to $\Delta t = 10$ trials from the target stimulus, thereby providing direct evidence for the Boltzmann distribution (Eq. (5)). Moreover, fitting an exponential curve to each distribution, we can directly estimate the inverse temperature β . Remarkably, the value $\beta = 0.32 \pm 0.1$ estimated from the combined distribution (Fig. 5c) falls within the range of values estimated from our reaction time experiments (Fig. 4a, g), nearly matching the average value $\beta = 0.30$ for random walk sequences (Fig. 4a).

To further strengthen the link between mental errors and people’s internal representations, we then asked subjects to perform the original serial response task (Fig. 1), and for each subject, we estimated β using the two methods described above: (i) directly measuring β in the n -back experiment, and (ii)

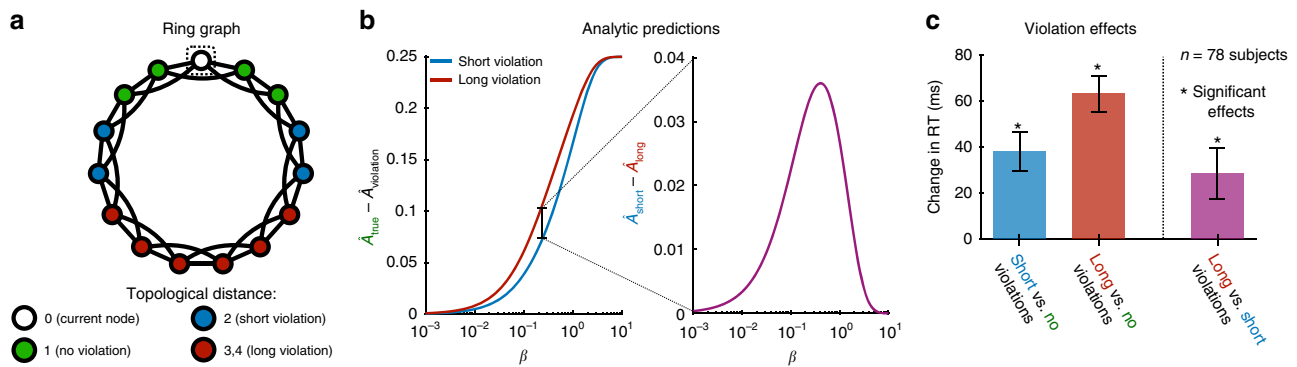


Fig. 6 Network violations yield surprise that grows with topological distance. **a** Ring graph consisting of 15 nodes, where each node is connected to its nearest neighbors and next-nearest neighbors on the ring. Starting from the boxed node, a sequence can undergo a standard transition (green), a short violation of the transition structure (blue), or a long violation (red). **b** Our model predicts that subjects' anticipations of both short (blue) and long (red) violations should be weaker than their anticipations of standard transitions (left). Furthermore, we predict that subjects' anticipations of violations should decrease with increasing topological distance (right). **c** Average effects of network violations across 78 subjects, estimated using a mixed effects model (Supplementary Tables 10 and 11), with error bars indicating one standard deviation from the mean. We find that standard transitions yield quicker reactions than both short violations ($p < 0.001$, $t = 4.50$, $df = 7.15 \times 10^4$) and long violations ($p < 0.001$, $t = 8.07$, $df = 7.15 \times 10^4$). Moreover, topologically shorter violations induce faster reactions than long violations ($p = 0.011$, $t = 2.54$, $df = 3.44 \times 10^3$), thus confirming the predictions of our model. Measurements were on independent subjects, and statistical significance was computed using two-sided F -tests. Source data are provided as a Source Data file.

indirectly estimating β in the serial response experiment. Comparing these two estimates across subjects, we find that they are significantly related with Spearman correlation $r_s = 0.28$ ($p = 0.047$, permutation test), while noting that we do not use the Pearson correlation because β is not normally distributed (Anderson–Darling test⁵⁸, $p < 0.001$ for the serial response task and $p = 0.013$ for the n -back task). Together, these results demonstrate not only the existence of the particular form of mental errors predicted by our model—down to the specific value of β —but also the relationship between these mental errors and people's internal estimates of the transition structure.

Network structure guides reactions to novel transitions. Given a model of human behavior, it is ultimately interesting to make testable predictions. Thus far, in keeping with the majority of existing research^{4,6,20–22,37}, we have focused on static transition graphs, wherein the probability A_{ij} of transitioning from state i to state j remains constant over time. However, the statistical structures governing human life are continually shifting^{59,60}, and people are often forced to respond to rare or novel transitions^{61,62}. Here we show that, when confronted with a novel transition—or a violation of the preexisting transition network—not only are people surprised, but the magnitude of their surprise depends critically on the topology of the underlying network.

We consider a ring graph where each node is connected to its nearest and next-nearest neighbors (Fig. 6a). We asked subjects to respond to sequences of 1500 nodes drawn as random walks on the ring graph, but with 50 violations randomly interspersed. These violations were divided into two categories: short violations of topological distance two and long violations of topological distances three and four (Fig. 6a). Using maximum likelihood estimation (Eq. (1)) as a guide, one would naively expect people to be equally surprised by all violations—indeed, each violation has never been seen before. In contrast, our model predicts that that surprise should depend crucially on the topological distance of a violation in the underlying graph, with topologically longer violations inducing increased surprise over short violations (Fig. 6b).

In the data, we find that all violations give rise to sharp increases in reaction times relative to standard transitions (Fig. 6c; Supplementary Table 10), indicating that people are in fact

learning the underlying transition structure. Moreover, we find that reaction times for long violations are 28 ms longer than those for short violations ($p = 0.011$, F -test; Fig. 6c; Supplementary Table 11). In addition, we confirm that the effects of network violations are not simply driven by stimulus recency (Supplementary Figs. 4 and 5). These observations suggest that people learn the topological distances between all nodes in the transition graph, not just those pairs for which a transition has already been observed^{59–62}.

Discussion

Daily life is filled with sequences of items that obey an underlying network architecture, from networks of word and note transitions in natural language and music to networks of abstract relationships in classroom lectures and literature^{5–9}. How humans infer and internally represent these complex structures are questions of fundamental interest^{10–13}. Recent experiments in statistical learning have established that human representations depend critically on the higher-order organization of probabilistic transitions, yet the underlying mechanisms remain poorly understood^{20–23}.

Here we show that network effects on human behavior can be understood as stemming from mental errors in people's estimates of the transition structure, while noting that future work should focus on disambiguating the role of recency^{39,40}. We use the free energy principle to develop a model of human expectations that explicitly accounts for the brain's natural tendency to minimize computational complexity—that is, to maximize entropy^{31,32,49}. Indeed, the brain must balance the benefits of making accurate predictions against the computational costs associated with such predictions^{24–29,50}. This competition between accuracy and efficiency induces errors in people's internal representations, which, in turn, explains with notable accuracy an array of higher-order network phenomena observed in human experiments^{20–23}. Importantly, our model admits a concise analytic form (Eq. (6)) and can be used to predict human behavior on a person-by-person basis (Fig. 4).

This work inspires directions for future research, particularly with regard to the study and design of optimally learnable network structures. Given the notion that densely connected communities help to mitigate the effects of mental errors on people's

internal representations, we anticipate that networks with high “learnability” will possess a hierarchical community structure⁶³. Interestingly, such hierarchical organization has already been observed in a diverse range of real world networks, from knowledge and language graphs⁶⁴ to social networks and the World Wide Web⁶⁵. Could it be that these networks have evolved so as to facilitate accurate representations in the minds of the humans using and observing them? Questions such as this demonstrate the importance of having simple principled models of human representations and point to the promising future of this research endeavor.

Methods

Maximum entropy model and the infinite-sequence limit. Here we provide a more thorough derivation of our maximum entropy model of human expectations, with the goal of fostering intuition. Given a matrix of erroneous transition counts \tilde{n}_{ij} , our estimate of the transition structure is given by $\hat{A}_{ij} = \tilde{n}_{ij} / \sum_k \tilde{n}_{ik}$. When observing a sequence of nodes x_1, x_2, \dots , in order to construct the counts \tilde{n}_{ij} , we assume that humans use the following recursive rule: $\tilde{n}_{ij}(t+1) = \tilde{n}_{ij}(t) + B_i(j)[x_{t+1}]$, where $B_i(i)$ denotes the belief, or perceived probability, that node i occurred at the previous time t . This belief, in turn, can be written in terms of the probability $P(\Delta t)$ of accidentally recalling the node that occurred Δt time steps from the desired node at time t : $B_i(i) = \sum_{\Delta t=0}^{t-1} P(\Delta t)[i = x_{t-\Delta t}]$.

In order to make quantitative predictions about people’s estimates of a transition structure, we must choose a mathematical form for $P(\Delta t)$. To do so, we leverage the free energy principle³¹. When building mental models, the brain is finely tuned to simultaneously minimize errors and computational complexity. The average error associated with a candidate distribution $Q(\Delta t)$ is assumed to be the average distance in time of the recalled node from the target node, denoted $E(Q) = \sum_{\Delta t} Q(\Delta t)\Delta t$. Furthermore, Shannon famously proved that the only suitable choice for the computational cost of a candidate distribution is its negative entropy³³, denoted $-S(Q) = -\sum_{\Delta t} Q(\Delta t)\log Q(\Delta t)$. Taken together, the total cost associated with a distribution $Q(\Delta t)$ is given by the free energy $F(Q) = \beta E(Q) - S(Q)$, where β , referred to as the inverse temperature, parameterizes the relative importance of minimizing errors versus computational costs. By minimizing F with respect to Q , we arrive at the Boltzmann distribution $p(\Delta t) = e^{-\beta \Delta t} / Z$, where Z is the normalizing partition function³⁰. We emphasize that this mathematical form for $P(\Delta t)$ followed directly from our free energy assumption about resource constraints in the brain.

To gain an analytic intuition for the model without referring to a particular random walk, we consider the limit of an infinitely long sequence of nodes. To begin, we consider a sequence x_1, \dots, x_T of length T . At the end of this sequence, the counting matrix takes the form

$$\begin{aligned} \tilde{n}_{ij}(T) &= \sum_{t=1}^{T-1} B_t(i)[j = x_{t+1}] \\ &= \sum_{t=1}^{T-1} \left(\sum_{\Delta t=0}^{t-1} P(\Delta t)[i = x_{t-\Delta t}] \right) [j = x_{t+1}]. \end{aligned} \tag{8}$$

Dividing both sides by T , the right-hand side becomes a time average, which by the ergodic theorem converges to an expectation over the transition structure in the limit $T \rightarrow \infty$,

$$\lim_{T \rightarrow \infty} \frac{\tilde{n}_{ij}(T)}{T} = \sum_{\Delta t=0}^{\infty} P(\Delta t) \langle [i = x_{t-\Delta t}][j = x_{t+1}] \rangle_A, \tag{9}$$

where $\langle \cdot \rangle_A$ denotes an expectation over random walks in A . We note that the expectation of an identity function is simply a probability, such that $\langle [i = x_{t-\Delta t}][j = x_{t+1}] \rangle_A = p_i(A^{\Delta t+1})_{ij}$, where p_i is the long-run probability of node i appearing in the sequence and $(A^{\Delta t+1})_{ij}$ is the probability of randomly walking from node i to node j in $\Delta t + 1$ steps. Putting these pieces together, we find that the expectation \hat{A} converges to a concise mathematical form,

$$\begin{aligned} \lim_{T \rightarrow \infty} \hat{A}_{ij}(T) &= \lim_{T \rightarrow \infty} \frac{\tilde{n}_{ij}(T)}{\sum_k \tilde{n}_{ik}(T)} \\ &= \frac{p_i \sum_{\Delta t=0}^{\infty} P(\Delta t)(A^{\Delta t+1})_{ij}}{p_i} \\ &= \sum_{\Delta t=0}^{\infty} P(\Delta t)(A^{\Delta t+1})_{ij}. \end{aligned} \tag{10}$$

Thus far, we have not appealed to our maximum entropy form for $P(\Delta t)$. It turns out that doing so allows us to write down an analytic expression for the long-time expectations \hat{A} simply in terms of the transition structure A and the inverse temperature β . Noting that $Z = \sum_{\Delta t=0}^{\infty} e^{-\beta \Delta t} = 1/(1 - e^{-\beta})$ and

$\sum_{\Delta t=0}^{\infty} (e^{-\beta} A)^{\Delta t} = (I - e^{-\beta} A)^{-1}$, we have

$$\begin{aligned} \hat{A} &= \sum_{\Delta t=0}^{\infty} P(\Delta t) A^{\Delta t+1} \\ &= \frac{1}{Z} A \sum_{\Delta t=0}^{\infty} (e^{-\beta} A)^{\Delta t} \\ &= (1 - e^{-\beta}) A (I - e^{-\beta} A)^{-1}. \end{aligned} \tag{11}$$

This simple formula for the representation \hat{A} is the basis for all of our analytic predictions (Figs. 3c, d and 6b) and is closely related to notions of communicability in complex network theory^{66,67}.

Experimental setup for serial response tasks. Subjects performed a self-paced serial reaction time task using a computer screen and keyboard. Each stimulus was presented as a horizontal row of five gray squares; all five squares were shown at all times. The squares corresponded spatially with the keys ‘Space’, ‘H’, ‘J’, ‘K’, and ‘L’, with the left square representing ‘Space’ and the right square representing ‘L’ (Fig. 1b). To indicate a target key or pair of keys for the subject to press, the corresponding squares would become outlined in red (Fig. 1a). When subjects pressed the correct key combination, the squares on the screen would immediately display the next target. If an incorrect key or pair of keys was pressed, the message ‘Error!’ was displayed on the screen below the stimuli and remained until the subject pressed the correct key(s). The order in which stimuli were presented to each subject was prescribed by either a random walk or a Hamiltonian walk on a graph of $N = 15$ nodes, and each sequence consisted of 1500 stimuli. For each subject, one of the 15 key combinations was randomly assigned to each node in the graph (Fig. 1a). Across all graphs, each node was connected to its four neighboring nodes with a uniform 0.25 transition probability. Importantly, given the uniform edge weights and homogeneous node degrees ($k = 4$), the only differences between the transition graphs lay in their higher-order structure.

In the first experiment, we presented subjects with random walk sequences drawn from two different graph topologies: a modular graph with three communities of five densely connected nodes and a lattice graph representing a 3×5 grid with periodic boundary conditions (Fig. 1c). The purpose of this experiment was to demonstrate the systematic dependencies of human reaction times on higher-order network structure, following similar results reported in recent literature^{22,37}. In particular, we demonstrate two higher-order network effects: In the cross-cluster surprisal effect, average reaction times for within-cluster transitions in the modular graph are significantly faster than reaction times for between-cluster transitions (Fig. 2a); and in the modular-lattice effect, average reaction times in the modular graph are significantly faster than reaction times in the lattice graph (Fig. 2b).

In the second experiment, we presented subjects with Hamiltonian walk sequences drawn from the modular graph. Specifically, each sequence consisted of 700 random walk trials (intended to allow each subject to learn the graph structure), followed by eight repeats of 85 random walk trials and 15 Hamiltonian walk trials (Supplementary Discussion)²⁰. Importantly, we find that the cross-cluster surprisal effect remains significant within the Hamiltonian walk trials (Fig. 2a).

In the third experiment, we considered a ring graph where each node was connected to its nearest and next-nearest neighbors in the ring (Fig. 6a). In order to study the dependence of human expectations on violations to the network structure, the first 500 trials for each subject constituted a standard random walk, allowing each subject time to develop expectations about the underlying transition structure. Across the final 1000 trials, we randomly distributed 50 network violations: 20 short violations of topological distance two and 30 long violations, 20 of topological distance three and 10 of topological distance four (Fig. 6a). As predicted by our model, we found a novel violations effect, wherein violations of longer topological distance give rise to larger increases in reaction times than short, local violations (Fig. 6b, c).

Data analysis for serial response tasks. To make inferences about subjects’ internal expectations based on their reaction times, we used more stringent filtering techniques than previous experiments when pre-processing the data²². Across all experiments, we first excluded from analysis the first 500 trials, in which subjects’ reaction times varied wildly (Fig. 1e), focusing instead on the final 1000 trials (or simply on the Hamiltonian trials in the second experiment), at which point subjects had already developed internal expectations about the transition structures. We then excluded all trials in which subjects responded incorrectly. Finally, we excluded reaction times that were implausible, either three standard deviations from a subject’s mean reaction time or below 100 ms. Furthermore, when measuring the network effects in all three experiments (Figs. 3 and 6), we also excluded reaction times over 3500 ms for implausibility. When estimating the parameters of our model and measuring model performance in the first two experiments (Fig. 4), to avoid large fluctuations in the results based on outlier reactions, we were even more stringent, excluding all reaction times over 2000 ms. Taken together, when measuring the cross-cluster surprisal and modular-lattice effects (Fig. 2), we used an average of 931 trials per subject; when estimating and evaluating our model (Fig. 4), we used an average of 911 trials per subject; and when measuring the

violation effects (Fig. 6), we used an average of 917 trials per subject. To ensure that our results are robust to particular choices in the data processing, we additionally studied all 1500 trials for each subject rather than just the final 1000, confirming that both the cross-cluster surprisal and modular-lattice effects remain significant across all trials (Supplementary Tables 6 and 7).

Measurement of network effects using mixed effects models. In order to extract the effects of higher-order network structure on subjects' reaction times, we used linear mixed effects models, which have become prominent in human research where many measurements are made for each subject^{38,68}. Put simply, mixed effects models generalize standard linear regression techniques to include both fixed effects, which are constant across subjects, and random effects, which vary between subjects. Compared with standard linear models, mixed effects models allow for differentiation between effects that are subject-specific and those that persist across an entire population. Here, all models were fit using the fitlme function in MATLAB (R2018a), and random effects were chosen as the maximal structure that (i) allowed model convergence and (ii) did not include effects whose 95% confidence intervals overlapped with zero⁶⁹. In what follows, when defining mixed effects models, we employ the standard R notation⁷⁰.

First, we considered the cross-cluster surprisal effect (Fig. 2a). As we were only interested in measuring higher-order effects of the network topology on human reaction times, it was important to regress out simple biomechanical dependencies on the target button combinations (Fig. 1d), the natural quickening of reactions with time (Fig. 1e), and the effects of recency on reaction times^{39,40}. Also, for the first experiment, since some subjects responded to both the modular and lattice graphs (Experimental Procedures), it was important to account for changes in reaction times due to which stage of the experiment a subject was in. To measure the cross-cluster surprisal effect, we fit a mixed effects model with the formula 'RT~log(Trial)*Stage+Target+Recency+Trans_Type+(1+log(Trial)*Stage+Recency+Trans_Type|ID)', where RT is the reaction time, Trial is the trial number (we found that log(Trial) was far more predictive of subjects' reaction times than the trial number itself), Stage is the stage of the experiment (either one or two), Target is the target button combination, Recency is the number of trials since the last instance of the current stimulus, Trans_Type is the type of transition (either within-cluster or between-cluster), and ID is each subject's unique ID. Fitting this mixed effects model to the random walk data in the first experiment (Supplementary Table 1), we found a 35 ms increase in reaction times ($p < 0.001$, F -test) for between-cluster transitions relative to within-cluster transitions (Fig. 2a). Similarly, fitting the same mixed effects model but without the variable Stage to the Hamiltonian walk data in the second experiment (Supplementary Table 4), we found a 36 ms increase in reaction times ($p < 0.001$, F -test) for between- versus within-cluster transitions (Fig. 2a). We note that because reaction times are not Gaussian distributed, it is fairly standard to perform a log transformation. However, for the above result as well as those that follow, we find the same qualitative effects with or without a log transformation.

Second, we studied the modular-lattice effect (Fig. 2b). To do so, we fit a mixed effects model with the formula 'RT~log(Trial)*Stage+Target+Recency+Graph+(1+log(Trial)*Stage+Recency+Graph|ID)', where Graph represents the type of transition network, either modular or lattice. Fitting this mixed effects model to the data in the first experiment (Supplementary Table 2), we found a fixed 23 ms increase in reaction times ($p < 0.001$, F -test) in the lattice graph relative to the modular graph (Fig. 2b).

Finally, we considered the effects of violations of varying topological distance in the ring lattice (Fig. 6c). We fit a mixed effects model with the formula 'RT~log(Trial)+Target+Recency+Top_Dist+(1+log(Trial)+Recency+Top_Dist|ID)', where Top_Dist represents the topological distance of a transition, either one for a standard transition, two for a short violation, or three for a long violation. Fitting the model to the data in the third experiment (Supplementary Tables 10 and 11), we found a 38 ms increase in reaction times for short violations relative to standard transitions ($p < 0.001$, F -test), a 63 ms increase in reaction times for long violations relative to standard transitions ($p < 0.001$, F -test), and a 28 ms increase in reaction times for long violations relative to short violations ($p = 0.011$, F -test). Put simply, people are more surprised by violations to the network structure that take them further from their current position in the network, suggesting that people have an implicit understanding of the topological distances between nodes in the network.

Estimating parameters and making quantitative predictions. Given an observed sequence of nodes x_1, \dots, x_{t-1} , and given an inverse temperature β , our model predicts the anticipation, or expectation, of the subsequent node x_t to be $a(t) = \hat{A}_{x_{t-1}, x_t}(t-1)$. In order to quantitatively describe the reactions of an individual subject, we must relate the expectations $a(t)$ to predictions about a person's reaction times $\hat{r}(t)$ and then calculate the model parameters that best fit the reactions of an individual subject. The simplest possible prediction is given by the linear relation $\hat{r}(t) = r_0 + r_1 a(t)$, where the intercept r_0 represents a person's reaction time with zero anticipation and the slope r_1 quantifies the strength with which a person's reaction times depend on their internal expectations.

In total, our predictions $\hat{r}(t)$ contain three parameters (β , r_0 , and r_1), which must be estimated from the reaction time data for each subject. Before estimating these parameters, however, we first regress out the dependencies of each subject's reaction times on the button combinations, trial number, and recency using a

mixed effects model of the form 'RT~log(Trial)*Stage+Target+Recency+(1+log(Trial)*Stage+Recency|ID)', where all variables were defined in the previous section. Then, to estimate the model parameters that best describe an individual's reactions, we minimize the RMS prediction error with respect to each subject's observed reaction times, $RMSE = \sqrt{\frac{1}{T} \sum (r(t) - \hat{r}(t))^2}$, where T is the number of trials. We note that, given a choice for the inverse temperature β , the linear parameters r_0 and r_1 can be calculated analytically using standard linear regression techniques. Thus, the problem of estimating the model parameters can be restated as a one-dimensional minimization problem; that is, minimizing RMSE with respect to the inverse temperature β . To find the global minimum, we began by calculating RMSE along 100 logarithmically spaced values for β between 10^{-4} and 10. Then, starting at the minimum value of this search, we performed gradient descent until the gradient fell below an absolute value of 10^{-6} . For a derivation of the gradient of the RMSE with respect to the inverse temperature β , we point the reader to the Supplementary Discussion. Finally, in addition to the gradient descent procedure described above, for each subject we also manually checked the RMSE associated with the two limits $\beta \rightarrow 0$ and $\beta \rightarrow \infty$. The resulting model parameters are shown in Fig. 4a, b for random walk sequences and Fig. 4g, h for Hamiltonian walk sequences.

Experimental setup for n -back memory task. Subjects performed a series of n -back memory tasks using a computer screen and keyboard. Each subject observed a random sequence of the letters 'B', 'D', 'G', 'T', and 'V', wherein each letter was randomly displayed in either upper or lower case. The subjects responded on each trial using the keyboard to indicate whether or not the current letter was the same as the letter that occurred n trials previously. For each subject, this task was repeated for the conditions $n = 1, 2$, and 3, and each condition consisted of a sequence of 100 letters. The three conditions were presented in a random order to each subject. After the n -back task, each subject then performed a serial response task (equivalent to the first experiment described above) consisting of 1500 random walk trials drawn from the modular graph.

Data analysis for n -back memory task. In order to estimate the inverse temperature β for each subject from their n -back data, we directly measured their memory distribution $P(\Delta t)$. As described in the main text, we treated each positive response indicating that the current stimulus matched the target stimulus as a sample of $P(\Delta t)$ by measuring the distance in trials Δt between the last instance of the current stimulus and the target (Fig. 5a). For each subject, we combined all such samples across the three conditions $n = 1, 2$, and 3 to arrive at a histogram for Δt . In order to generate robust estimates for the inverse temperature β , we generated 1000 bootstrap samples of the Δt histogram for each subject. For each sample, we calculated a linear fit to the distribution $P(\Delta t)$ on log-linear axes within the domain $0 \leq \Delta t \leq 4$ (note that we could not carry the fit out to $\Delta t = 10$ because the data is much sparser for individual subjects). To ensure that the logarithm of $P(\Delta t)$ was well defined for each sample—that is, to ensure that $P(\Delta t) > 0$ for all Δt —we added one count to each value of Δt . We then estimated the inverse temperature β for each sample by calculating the negative slope of the linear fit of $\log P(\Delta t)$ versus Δt . To arrive at an average estimate of β for each subject, we averaged β across the 1000 bootstrap samples. Finally, we compared these estimates of β from the n -back experiment with estimates of β from subjects' reaction times in the subsequent serial response task, as described above. We found that these two independent estimates of people's inverse temperatures are significantly correlated (excluding subjects for which $\beta = 0$ or $\beta \rightarrow \infty$), with a Spearman coefficient $r_s = 0.28$ ($p = 0.047$, permutation test). We note that we do not use the Pearson correlation coefficient because the estimates for β are not normally distributed for either the reaction time task ($p < 0.001$) nor the n -back task ($p = 0.013$) according to the Anderson-Darling test⁵⁸. This non-normality can be clearly seen in the distributions of β in Fig. 4a, g.

Experimental procedures. All participants provided informed consent in writing and experimental methods were approved by the Institutional Review Board of the University of Pennsylvania. In total, we recruited 634 unique participants to complete our studies on Amazon's Mechanical Turk. For the first serial response experiment, 101 participants only responded to sequences drawn from the modular graph, 113 participants only responded to sequences drawn from the lattice graph, and 72 participants responded to sequences drawn from both the modular and lattice graphs in back-to-back (counter-balanced) sessions for a total of 173 exposures to the modular graph and 185 exposures to the lattice graph. For the second experiment, we recruited 120 subjects to respond to random walk sequences with Hamiltonian walks interspersed, as described in the Supplementary Discussion. For the third experiment, we recruited 78 participants to respond to sequences drawn from the ring graph with violations randomly interspersed. For the n -back experiment, 150 subjects performed the n -back task and, of those, 88 completed the subsequent serial response task. Worker IDs were used to exclude duplicate participants between experiments, and all participants were financially remunerated for their time. In the first experiment, subjects were paid up to \$11 for up to an estimated 60 min: \$3 per network for up to two networks, \$2 per network for correctly responding on at least 90% of the trials, and \$1 for completing the entire task. In the second and third experiments, subjects were paid up to \$7.50 for

an estimated 30 min: \$5.50 for completing the experiment and \$2 for correctly responding on at least 90% of the trials. In the n -back experiment, subjects were paid up to \$8.50 for an estimated 45 min: \$7 for completing the entire experiment and \$1.50 for correctly responding on at least 90% of the serial response trials.

At the beginning of each experiment, subjects were provided with the following instructions: “In a few minutes, you will see five squares shown on the screen, which will light up as the experiment progresses. These squares correspond with keys on your keyboard, and your job is to watch the squares and press the corresponding key when that square lights up”. For the 72 subjects that responded to both the modular and lattice graphs in the first experiment, an additional piece of information was also provided: “This part will take around 30 minutes, followed by a similar task which will take another 30 minutes”. Before each experiment began, subjects were given a short quiz to verify that they had read and understood the instructions. If any questions were answered incorrectly, subjects were shown the instructions again and asked to repeat the quiz until they answered all questions correctly. Next, all subjects were shown a 10-trial segment that did not count towards their performance; this segment also displayed text on the screen explicitly telling the subject which keys to press on their keyboard. Subjects then began their 1500-trial experiment. For the subjects that responded to both the modular and lattice graphs, a brief reminder was presented before the second graph, but no new instructions were given. After completing each experiment, subjects were presented with performance information and their bonus earned, as well as the option to provide feedback.

Reporting summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

Source data for Fig. 1 are provided in Supplementary Data File 1. Source data for Fig. 2, Supplementary Figs. 2 and 3, and Supplementary Tables 1–9 are provided in Supplementary Data File 2. Source data for Fig. 5 are provided in Supplementary Data File 3. Source data for Fig. 6, Supplementary Figs. 4 and 5, and Supplementary Tables 10 and 11 are provided in Supplementary Data File 4. Source data from Supplementary Fig. 1 are provided in Supplementary Data File 5.

Code availability

The code that supports the findings of this study is available from the corresponding author upon reasonable request.

Received: 9 October 2019; Accepted: 13 February 2020;

Published online: 08 May 2020

References

- Hyman, R. Stimulus information as a determinant of reaction time. *J. Exp. Psychol.* **45**, 188 (1953).
- Sternberg, S. Memory-scanning: mental processes revealed by reaction-time experiments. *Am. Sci.* **57**, 421–457 (1969).
- Johnson-Laird, P. N. Mental models in cognitive science. *Cogn. Sci.* **4**, 71–115 (1980).
- Saffran, J. R., Aslin, R. N. & Newport, E. L. Statistical learning by 8-month-old infants. *Science* **274**, 1926–1928 (1996).
- Bousfield, W. A. The occurrence of clustering in the recall of randomly arranged associates. *J. Gen. Psychol.* **49**, 229–240 (1953).
- Fiser, J. & Aslin, R. N. Statistical learning of higher-order temporal structure from visual shape sequences. *J. Exp. Psychol.* **28**, 458 (2002).
- Friederici, A. D. Neurophysiological markers of early language acquisition: from syllables to sentences. *Trends Cogn. Sci.* **9**, 481–488 (2005).
- Tompson, S. H., Kahn, A. E., Falk, E. B., Vettel, J. M. & Bassett, D. S. Individual differences in learning social and non-social network structures. *J. Exp. Psychol. Learn. Mem. Cogn.* **45**, 253–271 (2019).
- Reynolds, J. R., Zacks, J. M. & Braver, T. S. A computational model of event segmentation from perceptual prediction. *Cogn. Sci.* **31**, 613–643 (2007).
- Meyniel, F. & Dehaene, S. Brain networks for confidence weighting and hierarchical inference during probabilistic learning. *Proc. Natl. Acad. Sci. USA* **114**, E3859–E3868 (2017).
- Dehaene, S., Meyniel, F., Wacongne, C., Wang, L. & Pallier, C. The neural representation of sequences: From transition probabilities to algebraic patterns and linguistic trees. *Neuron* **88**, 2–19 (2015).
- Piantadosi, S. T., Tenenbaum, J. B. & Goodman, N. D. Bootstrapping in a language of thought: A formal model of numerical concept learning. *Cognition* **123**, 199–217 (2012).
- Tenenbaum, J. B., Griffiths, T. L. & Kemp, C. Theory-based Bayesian models of inductive learning and reasoning. *Trends Cogn. Sci.* **10**, 309–318 (2006).
- Newman, M. E. The structure and function of complex networks. *SIAM Rev.* **45**, 167–256 (2003).
- Gómez, R. L. Variability and detection of invariant structure. *Psychol. Sci.* **13**, 431–436 (2002).
- Newport, E. L. & Aslin, R. N. Learning at a distance I. Statistical learning of non-adjacent dependencies. *Cogn. Psychol.* **48**, 127–162 (2004).
- Garvert, M. M., Dolan, R. J. & Behrens, T. E. A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *Elife* **6**, e17086 (2017).
- Cleeremans, A. & McClelland, J. L. Learning the structure of event sequences. *J. Exp. Psychol. Gen.* **120**, 235–253 (1991).
- Gomez, R. L. & Gerken, L. Artificial grammar learning by 1-year-olds leads to specific and abstract knowledge. *Cognition* **70**, 109–135 (1999).
- Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B. & Botvinick, M. M. Neural representations of events arise from temporal community structure. *Nat. Neurosci.* **16**, 486–492 (2013).
- Karuz, E. A., Thompson-Schill, S. L. & Bassett, D. S. Local patterns to global architectures: Influences of network topology on human learning. *Trends Cogn. Sci.* **20**, 629–640 (2016).
- Kahn, A. E., Karuz, E. A., Vettel, J. M. & Bassett, D. S. Network constraints on learnability of probabilistic motor sequences. *Nat. Hum. Behav.* **2**, 936–947 (2018).
- Karuz, E. A., Kahn, A. E. & Bassett, D. S. Human sensitivity to community structure is robust to topological variation. *Complexity* **2019**, 1–8 (2019).
- Tversky, A. & Kahneman, D. Judgment under uncertainty: Heuristics and biases. *Science* **185**, 1124–1131 (1974).
- de Camp Wilson, T. & Nisbett, R. E. The accuracy of verbal reports about the effects of stimuli on evaluations and behavior. *Soc. Psychol.* **41**, 118–131 (1978).
- Vinje, W. E. & Gallant, J. L. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* **287**, 1273–1276 (2000).
- Tononi, G., Sporns, O. & Edelman, G. M. A measure for brain complexity: Relating functional segregation and integration in the nervous system. *Proc. Natl. Acad. Sci. USA* **91**, 5033–5037 (1994).
- Cohen, J. D., McClure, S. M. & Angela, J. Y. Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **362**, 933–942 (2007).
- Wickelgren, W. A. Speed-accuracy tradeoff and information processing dynamics. *Acta Psychol.* **41**, 67–85 (1977).
- Jaynes, E. T. Information theory and statistical mechanics. *Phys. Rev.* **106**, 620 (1957).
- Ortega, P. A. & Braun, D. A. Thermodynamics as a theory of decision-making with information-processing costs. *Proc. R. Soc. A* **469**, 20120683 (2013).
- Friston, K., Kilner, J. & Harrison, L. A free energy principle for the brain. *J. Physiol. Paris* **100**, 70–87 (2006).
- Shannon, C. E. A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379–423 (1948).
- Brown, G. J. & Cooke, M. Computational auditory scene analysis. *Comput. Speech Lang.* **8**, 297–336 (1994).
- Lake, B. M., Salakhutdinov, R. & Tenenbaum, J. B. Human-level concept learning through probabilistic program induction. *Science* **350**, 1332–1338 (2015).
- McCarthy, G. & Donchin, E. A metric for thought: a comparison of p300 latency and reaction time. *Science* **211**, 77–80 (1981).
- Karuz, E. A., Kahn, A. E., Thompson-Schill, S. L. & Bassett, D. S. Process reveals structure: How a network is traversed mediates expectations about its architecture. *Sci. Rep.* **7**, 12733 (2017).
- Baayen, R. H., Davidson, D. J. & Bates, D. M. Mixed-effects modeling with crossed random effects for subjects and items. *J. Mem. Lang.* **59**, 390–412 (2008).
- Murdock, B. B. Jr The serial position effect of free recall. *J. Exp. Psychol.* **64**, 482–488 (1962).
- Baddeley, A. D. & Hitch, G. The recency effect: Implicit learning with explicit retrieval? *Mem. Cogn.* **21**, 146–155 (1993).
- Stachenfeld, K. L., Botvinick, M. M. & Gershman, S. J. The hippocampus as a predictive map. *Nat. Neurosci.* **20**, 1643–1653 (2017).
- Momennejad, I. et al. The successor representation in human reinforcement learning. *Nat. Hum. Behav.* **1**, 680–692 (2017).
- Boas, M. L. *Mathematical Methods in the Physical Sciences* (Wiley, 2006).
- Gregory, R. L. Perceptions as hypotheses. *Philos. Trans. R. Soc. Lond. B* **290**, 181–197 (1980).
- Howard, M. W. & Kahana, M. J. A distributed representation of temporal context. *J. Math. Psychol.* **46**, 269–299 (2002).
- Howard, M. W. & Kahana, M. J. Contextual variability and serial position effects in free recall. *J. Exp. Psychol. Learn. Mem. Cogn.* **25**, 923 (1999).
- Richards, B. A. & Frankland, P. W. The persistence and transience of memory. *Neuron* **94**, 1071–1084 (2017).
- Friston, K., Samothrakis, S. & Montague, R. Active inference and agency: Optimal control without cost functions. *Biol. Cybern.* **106**, 523–541 (2012).

49. Ortega, P. A. & Stocker, A. A. In *Advances in Neural Information Processing Systems* 100–108 (Curran Association, Red Hook, NY, 2016).
50. Gershman, S. J., Horvitz, E. J. & Tenenbaum, J. B. Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science* **349**, 273–278 (2015).
51. Grimmett, G. & Stirzaker, D. *Probability and Random Processes* (Oxford University Press, 2001).
52. Neter, J., Kutner, M. H., Nachtsheim, C. J. & Wasserman, W. *Applied Linear Statistical Models* Vol. 4 (Irwin, Chicago, 1996).
53. Seber, G. A. & Lee, A. J. *Linear Regression Analysis*, Vol. 329 (John Wiley & Sons, 2012).
54. Collins, A. G. & Frank, M. J. How much of reinforcement learning is working memory, not reinforcement learning? A behavioral, computational, and neurogenetic analysis. *Eur. J. Neurosci.* **35**, 1024–1035 (2012).
55. Collins, A. G. & Frank, M. J. Within- and across-trial dynamics of human EEG reveal cooperative interplay between reinforcement learning and working memory. *Proc. Natl Acad. Sci. USA* **115**, 2502–2507 (2018).
56. Dayan, P. Improving generalization for temporal difference learning: The successor representation. *Neural Comput.* **5**, 613–624 (1993).
57. Gershman, S. J., Moore, C. D., Todd, M. T., Norman, K. A. & Sederberg, P. B. The successor representation and temporal context. *Neural Comput.* **24**, 1553–1568 (2012).
58. Stephens, M. A. EDF statistics for goodness of fit and some comparisons. *J. Am. Stat. Assoc.* **69**, 730–737 (1974).
59. Whitehead, S. D. & Lin, L.-J. Reinforcement learning of non-Markov decision processes. *Artif. Intell.* **73**, 271–306 (1995).
60. Wang, X. & McCallum, A. Topics over time: A non-markov continuous-time model of topical trends. In *SIGKDD* 424–433 (ACM, 2006).
61. Wolfe, J. M., Horowitz, T. S. & Kenner, N. M. Cognitive psychology: Rare items often missed in visual searches. *Nature* **435**, 439 (2005).
62. Triá, F., Loreto, V., Servedio, V. D. P. & Strogatz, S. H. The dynamics of correlated novelties. *Sci. Rep.* **4**, 5890 (2014).
63. Arenas, A., Diaz-Guilera, A. & Pérez-Vicente, C. J. Synchronization reveals topological scales in complex networks. *Phys. Rev. Lett.* **96**, 114102 (2006).
64. Guimera, R., Danon, L., Diaz-Guilera, A., Giralt, F. & Arenas, A. Self-similar community structure in a network of human interactions. *Phys. Rev. E* **68**, 065103 (2003).
65. Ravasz, E. & Barabási, A.-L. Hierarchical organization in complex networks. *Phys. Rev. E* **67**, 026112 (2003).
66. Estrada, E. & Hatano, N. Communicability in complex networks. *Phys. Rev. E* **77**, 036111 (2008).
67. Estrada, E., Hatano, N. & Benzi, M. The physics of communicability in complex networks. *Phys. Rep.* **514**, 89–119 (2012).
68. Schall, R. Estimation in generalized linear models with random effects. *Biometrika* **78**, 719–727 (1991).
69. Hox, J. J., Moerbeek, M. & van de Schoot, R. *Multilevel Analysis: Techniques and Applications* (Routledge, 2017).
70. Bates, D. et al. Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* **67**, 1–48 (2015).

Acknowledgements

We thank Pedro Ortega for discussions. D.S.B., C.W.L., and A.E.K. acknowledge support from the John D. and Catherine T. MacArthur Foundation, the Alfred P. Sloan

Foundation, the ISI Foundation, the Paul Allen Foundation, the Army Research Laboratory (W911NF-10-2-0022), the Army Research Office (Bassett-W911NF-14-1-0679, Grafton-W911NF-16-1-0474, DCIST- W911NF-17-2-0181), the Office of Naval Research, the National Institute of Mental Health (2-R01-DC-009209-11, R01-MH112847, R01-MH107235, R21-MH-106799), the National Institute of Child Health and Human Development (1R01HD086888-01), National Institute of Neurological Disorders and Stroke (R01 NS099348), and the National Science Foundation (BCS-1441502, BCS-1430087, NSF PHY-1554488 and BCS-1631550). The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies.

Author contributions

C.W.L., A.E.K., and D.S.B. conceived the project. C.W.L. designed the model and performed the analysis. C.W.L., A.E.K., and D.S.B. planned the experiments and discussed the results. C.W.L., A.E.K., and N.N. performed the experiments. C.W.L. wrote the manuscript and Supplementary Information. A.E.K., N.N., and D.S.B. edited the manuscript and Supplementary Information.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41467-020-15146-7>.

Correspondence and requests for materials should be addressed to D.S.B.

Peer review information *Nature Communications* thanks the anonymous reviewer(s) for their contribution to the peer review of this work.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020