

Predictive learning rules generate a cortical-like replay of probabilistic sensory experiences

Reviewed Preprint

v2 • October 14, 2024

Revised by authors

Reviewed Preprint

v1 • August 7, 2024

Toshitake Asabuki , Tomoki Fukai 

Okinawa Institute of Science and Technology Graduate University, Onna-son, Japan • RIKEN Center for Brain Science, RIKEN ECL Research Unit, Wako, Japan • RIKEN Cluster for Pioneering Research, Wako, Japan

 https://en.wikipedia.org/wiki/Open_access
 Copyright information

Abstract

The brain is thought to construct an optimal internal model representing the probabilistic structure of the environment accurately. Evidence suggests that spontaneous brain activity gives such a model by cycling through activity patterns evoked by previous sensory experiences with the experienced probabilities. The brain's spontaneous activity emerges from internally-driven neural population dynamics. However, how cortical neural networks encode internal models into spontaneous activity is poorly understood. Recent computational and experimental studies suggest that a cortical neuron can implement complex computations, including predictive responses, through soma-dendrite interactions. Here, we show that a recurrent network of spiking neurons subject to the same predictive learning principle provides a novel mechanism to learn the spontaneous replay of probabilistic sensory experiences. In this network, the learning rules minimize probability mismatches between stimulus-evoked and internally driven activities in all excitatory and inhibitory neurons. This learning paradigm generates stimulus-specific cell assemblies that internally remember their activation probabilities using within-assembly recurrent connections. Our model contrasts previous models that encode the statistical structure of sensory experiences into Markovian transition patterns among cell assemblies. We demonstrate that the spontaneous activity of our model well replicates the behavioral biases of monkeys performing perceptual decision making. Our results suggest that interactions between intracellular processes and recurrent network dynamics are more crucial for learning cognitive behaviors than previously thought.

eLife Assessment

This **valuable** study investigates how biologically plausible learning mechanisms can support assembly formation that encodes statistics of the environment, by enabling neural sampling that is based on within-assembly connectivity strength. It **convincingly** shows that assembly formation can emerge from predictive plasticity in excitatory synapses, while two types of plasticity in inhibitory synapses are required: inhibitory homeostatic (predictive) plasticity and inhibitory competitive (anti-predictive) plasticity.

<https://doi.org/10.7554/eLife.92712.2.sa4>

Introduction

The brain is believed to construct an internal statistical model of an uncertain environment from sensory information streams for predicting the external events that are likely to occur. Evidence suggests that spontaneous brain activity learns the representation of such a model through repeated experiences of sensory events. In the cat visual cortex, spontaneously emerging activity patterns cycle through cortical states that include neural response patterns to oriented bars (Kenet et al., 2003 [↗](#)). In the ferret visual cortex, spontaneous activity gradually resembles a superposition of activity patterns evoked by natural scenes, eventually giving an optimal model of the visual experience (Berkes et al., 2011 [↗](#)). As replay activities can provide prior information for hierarchical Bayesian computation by the brain (Ernst & Banks, 2002 [↗](#); Kording & Wolpert, 2004 [↗](#); Friston, 2010 [↗](#); Fiser et al., 2010 [↗](#); Bastos et al., 2012 [↗](#); Orban et al., 2016 [↗](#); Legaspi & Toyozumi, 2019 [↗](#)), clarifying how the brain learns the spontaneous replay of optimal internal models is crucial for understanding whole-brain computing. However, the neural mechanisms underlying this modeling process are only poorly understood.

Several mechanisms of the brain's probabilistic computation have been explored (Jimenez Rezende & Gerstner, 2014 [↗](#); Li et al., 2022 [↗](#)). Models with reverberating activity are particularly interesting owing to their potential ability to generate spontaneous activity. For instance, spiking neuron networks with symmetric recurrent connections were proposed for Markov Chain Monte Carlo sampling of stochastic events (Buesing et al., 2011 [↗](#); Bill et al., 2015 [↗](#)). Spike-timing-dependent plasticity was used to organize spontaneous sequential activity patterns, providing a predictive model of sequence input (Hartmann et al., 2015 [↗](#)). However, previous models did not clarify how recurrent neural networks learn the spontaneous replay of the probabilistic structure of sensory experiences, for which these networks should learn the accurate probabilities of sensory stimuli and an appropriate excitation-inhibition balance simultaneously. Moreover, previous models assumed that each statistically salient stimulus in temporal input is already segregated and is delivered to a pre-assigned assembly of coding neurons, implying that the recurrent network, at least partly, knows the stochastic events to be modeled before learning. How the brain extracts salient events for statistical modeling has not been addressed.

Here, we present a learning principle to encode the experiences' probability structure into spontaneous network activity. To this end, we extensively use the synaptic plasticity rule proposed previously based on the hypothesis that the dendrites of a cortical neuron learn to predict its somatic responses (Urbanczik & Senn, 2014 [↗](#); Asabuki & Fukai, 2020 [↗](#)). We generalize the hypothetical predictive learning to a learning principle at the entire network level. Namely, in a recurrent network driven by external input, we ask all synapses on the dendrites of each excitatory or inhibitory neuron to learn to predict its somatic responses (although the dendrites will not be explicitly modeled). This enables the network model to simultaneously learn the events' probabilistic structure and the excitation-inhibition balance required to replay this structure. Further, our network model requires no pre-assigned cell assemblies since the model neuron can automatically segment statistically salient events in temporal input (Asabuki & Fukai, 2020 [↗](#)) - a cognitive process known as "chunking" (Fujii & Graybiel, 2003 [↗](#); Jin & Costa, 2010 [↗](#); Jin et al., 2014 [↗](#); Schapiro et al., 2013 [↗](#); Zacks et al., 2001 [↗](#)). Intriguingly, the cell assemblies generated by our model store their replay probabilities primarily in the within-assembly network structure, and intrinsic dynamical properties of membership neurons also contribute to this coding. This is in striking contrast to other network models that encode probabilities into the Markovian transition dynamics among cell assemblies (Buesing et al., 2011 [↗](#); Hartmann et al., 2015 [↗](#); Asabuki & Clopath, 2024 [↗](#)).

Our model trained on a perceptual decision-making task can replicate both unbiased and biased decision behaviors of monkeys without fine-tuning of parameters (Hanks et al., 2011 [↗](#)). In addition, in a network model consisting of distinct excitatory and inhibitory neural populations, our learning rule predicts the emergence of two types of inhibitory connections with different computational roles. We show that the emergence of the two inhibitory connection types is crucial for robust learning of an optimal internal model.

Results

Replay of probabilistic sensory experiences - A toy example

We first explain the task our model solves with a toy example. Consider a task in which the animal should decide whether a given stimulus coincides with or resembles any of two previously learned stimuli. Whether the animal learned these stimuli with a 50-50 chance or a 30-70 chance should affect the animal's anticipation of their occurrence and hence affect its decision.

It has been suggested that spontaneous activity expresses an optimal internal model of the sensory environment (Berkus et al., 2011 [↗](#)). In our toy example, the evoked activity patterns of the two stimuli should be spontaneously replayed with the same probabilities as these stimuli were experienced during learning:

$$\langle P(\text{features}|\text{input}, \text{model}) \rangle_{P(\text{input})} = P(\text{features}|\text{model}),$$

where features = {stimulus 1, stimulus 2} and the right-hand side expresses the probabilities of replayed activities. The angular brackets indicate averaging over the stimuli. According to Hebb's hypothesis, two cell assemblies should be formed to memorize the two stimuli in the toy example. Moreover, the spontaneous replay of these cell assemblies should represent the probabilities given in the right-hand side of the above equation. Below, we propose a mathematical principle of learning to achieve these requirements.

Prediction-driven synaptic plasticity for encoding an internal model

We previously proposed a learning rule for a single two-compartment neuron (Asabuki & Fukai, 2020 [↗](#)). Briefly, our previous model learns statistically salient features repeated in input sequences by minimizing the error between somatic and dendritic response probabilities without external supervision to identify the temporal locations of these features. In this study, we extend this plasticity rule to recurrent networks by asking all neurons in a network to minimize the error in response probabilities between the internally generated and stimulus-evoked activities (Fig. 1 [↗](#)). Our central interest is whether this learning principle generates spontaneous activity representing the statistical model of previous experiences.

We first introduce our learning principle using a recurrent network model (nDL model) that does not obey Dale's law for distinguishing between excitatory and inhibitory neurons (Materials and Methods). A more realistic model with distinct excitatory and inhibitory neuron pools will be shown later. The nDL model consists of Poisson spiking neurons, each receiving Poisson spike trains from all input neurons via a modifiable all-to-all afferent feedforward (FF) connection matrix **W** (Fig. 1a [↗](#)). These input neurons may be grouped into multiple input neuron groups responding to different sensory features. Due to the all-to-all connectivity, the afferent input has no specific predefined structure. Two types of all-to-all modifiable recurrent connections (REC), **M** and **G**, exist among the neurons. Matrix **M** is a mixture of excitatory and inhibitory connections, and matrix **G** represents inhibitory-only connections. Due to a minus sign for v^G , all components of **G** are positive. The firing rate of neurons are defined as a modifiable sigmoidal function of the




Figure 1.

Unsupervised prior learning in a recurrent neural network.

(a) A schematic of a network model is shown. The interconnected circles denote the model neurons, of which the activities are controlled by two types of inputs: feedforward (FF) and recurrent (REC) inputs. Colored circles indicate active neurons. Here, W denotes FF, and M and G denote REC connections. We considered two modes of activity (i.e., evoked and spontaneous activity). In the evoked mode, the membrane potential u of a network neuron was calculated as a linear combination of inputs across all different connections (v^W , v^M , and v^G). This evoked mode is considered during the learning phase, when all synapses attempt to predict the network activity, as we will explain in the main text. Once all synapses are sufficiently learned, all FF inputs are removed, and the network is driven spontaneously (spontaneous mode). Our interest lies in the statistical similarity of the network activity in these two modes. (b) The gain and threshold of output response function was controlled by a dynamic variable, h , which tracks the history of the membrane potential. (c) A schematic of the learning rule for a network neuron is shown (top). During learning, for each type of connection on a postsynaptic neuron, synaptic plasticity minimizes the error between output (gray diamond) and synaptic prediction (colored diamonds). Note that all types of synapses share the common plasticity rule, where weight updates are calculated as the multiplication of the error term and the presynaptic activities (bottom). Our hypothesis is that such plasticity rule allows a recurrent neural network to spontaneously replay the learned stochastic activity patterns without external input.

membrane potential (**Fig. 1b**), which we will explain later in detail. All types of connections, both afferent and recurrent ones, are modifiable by unsupervised learning rules derived from a common principle: on each neuron, all synapses learn to predict the neuron's response optimally (**Fig. 1c**; see Materials and Methods). In reality, all synaptic inputs may be terminated on the dendrites, although they are not modeled explicitly.

Without a teaching signal, predictive learning may suffer a trivial solution problem in which all synapses vanish, and hence all neurons become silent (Asabuki & Fukai, 2020). To avoid it, we homeostatically regulate the dynamic range of each neuron (i.e., the slope and threshold of the response function) according to the history h of its subthreshold activity (see Eqs. 10-12). When the value of h is increased, the neuron's excitability is lowered (**Fig. 1b**). The input-output curves of neurons are known to undergo homeostatic regulations through various mechanisms (Chance et al., 2002; Mitchell and Silver, 2003; Torres-Torrel et al., 2014). Though no direct experimental evidence is available for our homeostatic process via h , it mathematically avoids saturating neuronal activity.

Note that the present homeostatic regulation of intrinsic excitability differs from the homeostatic synaptic scaling mechanism. The role of homeostatic synaptic scaling in generating irregular cell-assembly activity patterns was previously studied computationally (Hiratani and Fukai, 2014; Litwin-Kumar and Doiron, 2014; Zenke et al., 2015). However, unlike the present model, the previous models did not address whether and how synaptic scaling contributes to statistical modeling by recurrent neural networks. Furthermore, unlike our model, in which neurons in the recurrent layer and input neurons are initially connected in an all-to-all manner, most previous models assumed preconfigured receptive fields for recurrent-layer neurons, implying that these models had predefined stimulus-specific cell assemblies.

Cell assembly formation for learning statistically salient stimuli

We first explain how our network segments salient stimuli and forms stimulus-specific cell assemblies via network-wide predictive learning rules. To this end, we tested a simple case in which two non-overlapping input groups are intermittently and repeatedly activated with equal probabilities. The two input patterns were separated by irregular, low-frequency, unrepeatable spike trains of all input neurons (Materials and Methods). We will consider input patterns with unequal occurrence probabilities later. After several presentations of individual input patterns, each network neuron responded selectively to one of the repeated patterns (**Figure 2a**). This result is consistent with our previous results (Asabuki & Fukai, 2020) that the plasticity of feedforward connections segments input patterns. Indeed, feedforward synapses W on each neuron were strengthened or weakened when they mediated its preferred or non-preferred stimulus, respectively (**Fig. 2b**, left; **Fig. 2c**). Inhibitory connections G grew between neurons within the same assembly but not between assemblies (**Fig. 2b**, right; **Fig. 2c**, bottom), enhancing the decorrelation of within-assembly neural activities (Asabuki & Fukai, 2020). Recurrent connections M were modified to form stimulus-specific cell assemblies, as evidenced by the self-organization of excitatory (**Fig. 2c**, top) and inhibitory (**Fig. 2c**, bottom) recurrent connections within and between cell assemblies, respectively. The inhibitory components are necessary for suppressing the simultaneous replay of different cell assemblies, as shown later.

We then investigated whether and how spontaneous activity preserves and replays these cell assemblies in the absence of afferent input. To demonstrate this in a more complex task, we trained the network with afferent input involving five repeated patterns and then removed the input and observed post-training spontaneous network activity (**Fig. 2d**). The termination of afferent input initially lowered the activities of neurons, but their dynamic ranges gradually recovered with the excitability of the neural population (indicated by the population-averaged h value), and the network eventually started spontaneously replaying the learned cell assemblies. All plasticity rules were turned off during the recovery period (about 20 seconds from the input termination), after which the network settled in a stable spontaneous firing state (plasticity off),




Figure 2.

Formation of stimulus-selective assemblies in a recurrent network.

(a) Example dynamics of neuronal output and synaptic predictions are shown before (left) and after (right) learning. Colored bars at the top of the figures represent periods of stimulus presentations. (b) Example dynamics of feedforward connection W and inhibitory connection G are shown. W -connections onto neurons organizing to encode the same or different input patterns are shown in red and blue, respectively. Similarly, the same colors are used to represent G connections within and between assemblies. (c) Dynamics of the mean connection strengths are shown on neuron in cell assembly 1. Shaded areas represent SDs. In the schematic, triangles indicate input neurons and circles indicate network neurons. The color of each neuron indicates the stimulus preference of each neuron. (d) Example dynamics of the averaged dynamical variable $\bar{h}$ (top) and the learned network activity (bottom) are shown. The dynamical variables are averaged over the entire network. Neurons are sorted according to their preferred stimuli. During the spontaneous activity, afferent inputs to the network were removed. The inset shows the firing rate distribution of the evoked and the spontaneous activity. (e) Correlation coefficients of spontaneous activities of every pair of neurons are shown.

with firing rates lower than those of the evoked activity (inset). Then, the plasticity rules could be turned on (plasticity on) without drastically destroying the structure of spontaneous replay. Intriguingly, spontaneous neuronal activities were highly correlated within each cell assembly but were uncorrelated between different cell assemblies (**Fig. 2e**). This was because self-organized recurrent connections **M** were excitatory within each cell assembly, whereas the between-assembly recurrent connections were inhibitory, as in **Fig. 2c**.

Thus, the network model successfully segregates, remembers, and replays stimulus-evoked activity patterns in temporal input. The loss of between-assembly excitatory connections is interesting as it indicates that the present spontaneous reactivation is not due to the sequential activation of cell assemblies. This can also be seen from the relatively long intervals between consecutive cell-assembly activations: spontaneous neural activity does not propagate directly from one cell assembly to another (**Fig. 2d**). Indeed, within-assembly excitation is the major cause of spontaneous replay in this model, which we will study later in detail.

In summary, we have proposed the predictive learning rules as a novel plasticity mechanism for all types of synapses (i.e., feedforward and recurrent connections). We have shown that the plasticity rules in our model learn the segmentation of salient patterns in input sequences and form pattern-specific cell assemblies without preconfigured structures. We also showed that our model replays the learned assemblies even when external inputs were removed.

Replays of cell assemblies reflect a learned statistical model

We now turn to the central question of this study. We asked whether internally generated network dynamics through recurrent synapses (i.e., spontaneous replay of cell assemblies) can represent an optimal model of previous sensory experiences. Specifically, we examined whether the network spontaneously reactivates learned cell assemblies with relative frequencies proportional to the probabilities with which external stimuli activated these cell assemblies during learning. We addressed these questions in slightly more complex cases with increased numbers of external stimuli.

We first examined a case with five stimuli in which stimulus 1 was presented twice as often as the other four stimuli (**Fig. 3a**). Hereafter, the probability ratio refers to the relative number of times stimulus 1 is presented during learning. For instance, the case shown in **Fig. 2d** represents the probability ratio one. As in **Fig. 2d**, the network self-organized five cell assemblies to encode stimuli 1 to 5 and replayed all of them in subsequent spontaneous activity (**Fig. 3b**). We found that output neurons were activated more frequently and strongly in cell assembly 1 than in other cell assemblies. Therefore, we accessed quantitative differences in neuronal activity between different cell assemblies by varying the probability ratio.

The neuronal firing rate of cell assembly 1 relative to other cell assemblies increased approximately linearly with an increase in the probability ratio (**Fig. 3c**). Similarly, the size of cell assembly 1 relative to other cell assemblies also increased with the probability ratio (**Fig. 3d**). However, neither the relative firing rate nor the relative assembly size faithfully reflects changes in the probability ratio: scaling the probability ratio with a multiplicative factor does not scale these quantities with this factor. Therefore, we further investigated whether the assembly activity ratio, the ratio in the total firing rate of cell assembly 1 to other cell assemblies (Materials and Methods), scales faithfully with the probability ratio of cell assembly 1. This was the case: the scaling was surprisingly accurate (**Fig. 3e**).

To examine the ability of the nDL network further, we trained it with five stimuli occurring with various probabilities (**Fig. 3f** and Supplementary Fig. 1a). After learning, the spontaneous activity of the model replayed the learned cell assemblies at the desired ratios of population firing rates (**Fig. 3g** and Supplementary Fig. 1b).




Figure 3.

Priors coded in spontaneous activity.

An nDL network was trained with five probabilistic inputs. (a) Stimulus 1 appeared twice as often as the other four stimuli during learning. The example empirical probabilities of the stimuli used for learning are shown. (b) The spontaneous activity of the trained network shows distinct assembly structures. (c) The mean ratio of the population-averaged firing rate of assembly 1 to those of the other assemblies is shown for different values of the occurrence probability of stimulus 1. Vertical bars show SDs over five trials. A diagonal dashed line is a ground truth. (d) Similarly, the mean ratios of the size of assembly 1 to those of the other assemblies are shown. (e) The mean ratios of the total activities of neurons in assembly 1 to those of the other assemblies are shown. (f) Five stimuli occurring with different probabilities were used for training the nDL model. (g) The population firing rates are shown for five self-organized cell assemblies encoding the stimulus probabilities shown in (f).

We then asked whether our model would learn a prior distribution for more stimuli. To this end, we presented seven stimulus patterns to the same network with graded probabilities (Supplementary Fig. 1c). The self-organized spontaneous activity exhibited cell assemblies that well learned the graded probability distribution of these stimuli (Supplementary Fig. 1d). These results demonstrate that the trained network remembers the probabilities of repetitively experienced stimuli by the spontaneous firing rates of the encoding cell assemblies and that this dynamical coding scheme has a certain degree of scalability.

So far, we have represented external stimuli with non-overlapping subgroups of input neurons. However, in biologically realistic situations, input neuron groups may share part of their membership neurons. We tested whether the proposed model could learn the probability structure of overlapping input patterns in a case where two input neuron groups shared half of their members. The two patterns were presented with probabilities of 30% and 70%, respectively (Supplementary Fig. 2a). After sufficient learning, the network model generated two assemblies that encoded the two stimuli without sharing the coding neurons (Supplementary Fig. 2b) and replayed these assemblies with frequencies proportional to the stimulus presentation probabilities (Supplementary Fig. 2c). The results look reasonable because each neuron in the network segments one of the stimulus patterns and recurrent connections within each non-overlapping assembly can encode the probability of its replay.

Altogether, these results suggest that our model spontaneously replays learned cell assemblies with relative frequencies proportional to the probability that each cell assembly was activated during the learning phase. We have shown that the population activities of assemblies, rather than the firing rates of individual neurons, encode the occurrence probabilities of stimulus patterns.

Within-assembly recurrent connections encode probabilistic sensory experiences

To understand the mechanism underlying the statistical similarity between the evoked patterns and spontaneous activity, we then investigated whether and how biases in probabilistic sensory experiences influence the strengths of recurrent connections. To this end, we compared two cases in which two input patterns (stim 1 and stim 2) occurred with equal (50% vs. 50%) and different (30% vs. 70%) probabilities during learning (Fig. 4a). From the results shown in Fig. 3, we hypothesized that within-assembly learned connections should reflect the stimulus occurrence probabilities and hence the activation probabilities of the corresponding cell assemblies during spontaneous activity. Therefore, we calculated the total strengths of incoming recurrent synapses on each neuron within the individual cell assemblies (Fig. 4b). While the distributions of incoming synaptic strengths are similar between cell assemblies coding stimulus 1 and stimulus 2 in the 50-vs-50 case, they look different in the 30-vs-70 case (Fig. 4c).

Since incoming weights increased more significantly in the cell assembly activated by a more frequent stimulus (i.e., the assembly encoding stimulus 2 in the 30-vs-70 case), we expect that the degree of positive shifts in incoming weight distributions will reflect stimulus probabilities. To examine whether this is indeed the case, we computed the sum of total excitatory incoming weights (i.e., the sum of positive elements of M) over neurons belonging to each assembly after training. We then normalized these excitatory incoming weights over the two assemblies.

Interestingly, we found that the normalized excitatory incoming weights for the two assemblies well approximates the empirical probabilities of the two stimuli in both the 50-vs-50 and 30-vs-70 cases (Figure 4d). These analyses revealed that recurrent connections learned within assemblies encode biases in probabilistic sensory experiences. Indeed, the elimination of between-assembly excitatory connections did not significantly affect the replay probabilities, as the sampling is driven by strong within-assembly recurrent inputs after learning (Supplementary Fig. 3).




Figure 4.

Probability encoding by learned within-assembly synapses.

(a) Two input stimuli were presented in two protocols: uniform (50% vs. 50%) or biased (30% vs. 70%). (b) The total incoming synaptic strength on each neuron was calculated within each cell assembly. (c) *left*, The distributions of incoming synaptic strength are shown for the learned assemblies in the 50-vs-50 case. *right*, Same as in the left figure, but in the 30-vs-70 case. (d) *left*, The empirical probabilities of stimuli 1 and 2 and the normalized excitatory incoming weights within assemblies are compared in the 50-vs-50 case. *right*, Same as in the left figure, but in the 30-vs-70 case.

Roles of inhibitory plasticity for stabilizing cell assemblies

Experimental and computational results suggest that inhibitory synapses are more robust to spontaneous activity than excitatory synapses and are crucial for maintaining cortical circuit function (Mongillo et al., 2018 [↗](#)). To see the crucial role of the inhibitory plasticity of G for cell assembly formation, we compared the spontaneously driven activities in the learned network between two cases, plastic inhibitory connection G versus fixed G , in the 30-vs-70 case. The results show that only a single, highly active assembly self-organizes for fixed inhibitory synapses (Supplementary Fig. 4a). In contrast, such unstable dynamics do not emerge from plastic inhibitory synapses (Supplementary Fig. 4b), suggesting the crucial role of inhibitory plasticity in stabilizing spontaneous activity.

To further clarify the functional role of inhibitory plasticity in regulating spontaneous activity, we compared how the self-organized assembly structure of recurrent connections M evolves in the two simulation settings shown in Supplementary Fig. 5a. In the control model, we turned off the plasticity of G for a while after the cessation of external stimuli but again switched it on, as was previously in Fig. 2 [↗](#). The cell-assembly structure initially dissipated but eventually reached a well-defined equilibrium structure (Supplementary Fig. 5b, magenta). Consistent with this, the postsynaptic potentials mediated by connections M and G predicted the normalized firing rate of a postsynaptic excitatory neuron in the control model (Supplementary Fig. 5c). In striking contrast, the cell-assembly structure rapidly dissipated in the truncated model in which the G -plasticity was kept turned off after the cessation of external stimuli (Supplementary Fig. 5b, blue). Accordingly, the postsynaptic potentials induced by M and G , so was the normalized firing rate, evolved into trivial solutions and almost vanished in the truncated model (Supplementary Fig. 5d). Only the control model, but not the truncated model, could maintain prediction errors small and nearly constant after the termination of the stimuli (Supplementary Fig. 5e). These results indicate that maintaining the learned representations requires the continuous tuning of within-assembly inhibition.

The role of homeostatic regulation of neural activities

As indicated by the weak couplings between cell assemblies, the present mechanism of probability learning differs from the conventional sequence learning mechanisms. Consistent with this, the network trained repetitively by a fixed sequence of patterned inputs does not exhibit stereotyped sequential transitions among cell assemblies (due to the lack of strong inter-assembly excitatory connections; Supplementary Fig. 6). Indeed, the probability-encoding spontaneous activity emerges in the present model mainly from the within-assembly dynamics driven by strong within-assembly reverberating synaptic input. However, homeostatic variable h also plays a role in maintaining a stable spontaneous network activity after learning (see Fig. 2d [↗](#); activity pattern from 5 to 10 sec). This is achieved by the time evolution of h , which maintains the firing rate of each neuron in a suitable range by adjusting the threshold and gain of the somatic sigmoidal response function (Fig. 1b [↗](#)).

Therefore, we explored the role of the homeostatic variable in learning an accurate internal model of the sensory environment. In each neuron, the variable h is updated whenever the membrane potential undergoes an abrupt increase (Eq. 10 [↗](#)). Therefore, the time evolution of h monitors the activity of each neuron over the timescale of seconds, which in turn regulates the neural activity by controlling the activation function (Supplementary Fig. 7a; Eqs.8-9). When the instantaneous value of h is high, the neuron's excitability is lowered (namely, the gain and threshold of the response function are decreased or increased, respectively: see Eqs. 10 [↗](#)-12 [↗](#)). This activity regulation is crucial to avoid the trivial solution of the plasticity rules (Asabuki and Fukai, 2020 [↗](#)).




Figure 5.

Simulations of biased perception of visual motion coherence.

(a) The network model simulated perceptual decision-making of coherence in random dot motion patterns. In the network shown here, network neurons have already learned two assemblies encoding leftward or rightward movements from input neuron groups L and R. The firing rates of input neuron groups were modulated according to the coherence level Coh of random dot motion patterns (Materials and Methods). (b) The choice probabilities of monkeys (circles) and the network model (solid lines) are plotted against the motion coherence in two learning protocols with different prior probabilities. The experimental data were taken from Hanks et al. (2011). In the 50:50 protocol, moving dots in the “R” (Coh = 0.5) and “L” (Coh = -0.5) directions were presented randomly with equal probabilities, while in the 80:20 protocol, the “R” and “L” directions were trained with 80% and 20% probabilities, respectively. Shaded areas represent SDs over 20 independent simulations. The computational and experimental results show surprising coincidence without curve fitting. (c) Spontaneous and evoked activities of the trained networks are shown for the 50:50 (left) and 80:20 (right) protocols. Evoked responses were calculated for three levels of coherence: Coh = -50%, 0%, and 50%. In both protocols, the activity ratio in spontaneous activity matches the prior probability and gives the baseline for evoked responses. In the 80:20 protocol, the biased priors of “R” and “L” motion stimuli shift the activity ratio in spontaneous activity to an “R”-dominant regime.




Figure 6.

A network model with distinct excitatory and inhibitory connections.

(a) Strong excitatory connections were formed within assemblies. (b) The first type of recurrent inhibitory connections, M_{inh} , became stronger between assemblies, enhancing assembly competition. (c) The second type of inhibitory connections G were strengthened within assemblies to balance the strong excitatory inputs.

but not critical for sampling with appropriate probabilities. Actually, a model with a fixed value of h still showed spontaneous replay, although the true probability distribution was estimated less accurately (Supplementary Fig. 7b: cf. [Fig. 3f](#)).

Learning conditioned prior distributions

The predictive coding hypothesizes that top-down input from higher cortical areas provides prior knowledge about computations in lower cortical areas. This implies in the brain's hierarchical computation that the top-down input conditions the prior distributions in local cortical areas to those relevant to the given context. The proposed learning rules can account for how a conditioned input from other cortical areas conditions the prior distribution in a local cortical circuit.

The neural network consists of two mutually interacting non-overlapping subnetworks of equal sizes, where the subnetworks may represent different cortical areas (Supplementary Fig. 8a). Subnetwork A was randomly exposed to stimuli 1 and 2 (S1 and S2) with equal probabilities 1/2, whereas subnetwork B was to stimuli 3 and 4 (S3 and S4) with the conditional probabilities 1/3 and 2/3 if S1 was presented to subnetwork A and the conditional probabilities 2/3 and 1/3 if S2 was presented to subnetwork A. After learning, the network model self-organized four cell assemblies each of which responded preferentially to one of the four stimuli (Supplementary Fig. 8b). Consistent with this, the self-organized connection matrix represented strong within-assembly connections within each cell assembly and weak between-assembly connections (Supplementary Fig. 8c). Note that between-assembly connections were inhibitory between assemblies encoding mutually exclusive stimuli, i.e., S1 and S2 and S3 and S4, as they should be. Now, we turned off S3 and S4 to subnetwork B and only applied S1 or S2 to subnetwork A each at one time. Applying the same stimulus (i.e., S1 or S2) to subnetwork A activated either S3- or S4-coding cell assembly in subnetwork B in a probabilistic manner (Supplementary Fig. 8d). The cell assemblies evoked in subnetwork B by S1 or S2 to subnetwork A varied the total firing rates approximately in proportion to the conditional probabilities (e.g., $P(S3|S1) = 1/3$ vs. $P(S4|S1) = 2/3$) used during learning (Supplementary Fig. 8e). Note that S3- and S4-coding cell assemblies could become simultaneously active to represent the desired activation probabilities (e.g., a vertical arrow in Supplementary Fig. 8d). Together, these results indicate that our network can learn prior distributions conditioned by additional inputs through different pathways.

Replication of biased perceptual decision making in monkeys

Prior knowledge about the environment often biases our percept of the external world. For instance, if we know that two possible stimuli exist and that stimulus A appears more often than stimulus B, we tend to feel that a given stimulus is more likely to be stimulus A than stimulus B. Previously, a similar bias was quantitatively studied in monkeys performing a perceptual decision making task ([Hanks et al., 2011](#)). In the experiment, monkeys had to judge the direction (right or left) of the coherent motion of moving dots on a display. When both directions of coherent motion appeared randomly during learning, the monkey showed unbiased choice behaviors. However, if the frequencies of the two motion directions were different, the monkey's choice was biased toward the direction of a more frequent motion stimulus.

We constructed a network model shown in [Fig. 5a](#) to examine whether the present mechanism of spontaneous replay could account for the behavioral bias. The model comprises a recurrent network similar to that used in [Fig. 2](#) and two input neuron groups, L and R, encoding leftward or rightward coherent dot movements, respectively. We modulated the firing rates of these input neurons in proportion to the coherence of moving dots (Materials and Methods). During learning, we trained this model with external stimuli having input coherence Coh of either -0.5 or +0.5 (Materials and Methods), where all dots move leftward in the former or rightward in the latter. In so doing, we mimicked the two protocols used in the behavioral experiment of monkeys: in the




Figure 7.

The DL model of excitatory and inhibitory cell assemblies.

(a) This model consists of distinct excitatory and inhibitory neuron pools, obeying Dale's law. (b) Each inhibitory neuron projects to another neuron X through two inhibitory paths, path 1 and path 2, where the index X refers to an excitatory or an inhibitory postsynaptic neuron. Hexagons represent minimal units for prediction and learning in the neuron model and may correspond to dendrites, which were not modeled explicitly. (c) The probability ratios estimated by numerical simulations are plotted for the assembly activity ratios (purple), firing rate ratios (cyan), and assembly size ratios (green) as functions of the true probability ratio of external stimuli. Error bars indicate SEs calculated over five simulation trials with different initial states of neurons and synaptic weights in each parameter setting. (d) Inhibitory connection matrices are shown for path 1 and path 2. (e) The mean weights of self-organized synapses on excitatory and inhibitory postsynaptic neurons are shown. (f) Within-assembly and between-assembly connectivity patterns of excitatory and inhibitory neurons are shown. Colors indicate three cell assemblies self-organized. (g) The strengths of lateral inhibitions within-(W/N) and between-assemblies (B/N) are shown for paths 1 and 2. Horizontal bars show the medians and quartiles. (h) The resultant connectivity pattern suggests an effective competitive network between excitatory assemblies with self-(within-assembly) and lateral (between-assembly) inhibition.

50:50 protocol, two stimuli with $\text{Coh} = \pm 0.5$ were presented randomly with equal probabilities, while in the 80:20 protocol, stimuli with $\text{Coh} = +0.5$ and -0.5 were delivered with probabilities of 80% and 20%, respectively. In the 80:20 protocol, stimuli were highly biased toward a coherent rightward motion.

The network model could explain the biased choices of monkeys surprisingly well. In either training protocol, the recurrent network self-organized two cell assemblies responding selectively to one of the R and L input neuron groups. Then, we examined whether the responses of the self-organized network are consistent with experimental observations by stimulating it with external inputs having various degrees of input coherence. The resultant psychometric curves almost perfectly coincide with those obtained in the experiment (**Fig. 5b**). We note that the psychometric curves of the model do not significantly depend on the specific choices of parameter values as far as the network learned stable spontaneous activity. We did not perform any curve fitting to experimental data, implying that the psychometric curves are free from parameter finetuning.

Biases in the psychometric curves emerged from biased firing rates of spontaneous activity of the self-organized cell assemblies. To show this, we investigated how the activities of the two self-organized cell assemblies change before and after the onset of test stimuli in three relatively simple cases, i.e., $\text{Coh} = -0.5$, 0, and $+0.5$. **Figure 5c** shows the activity ratio AR between the R-encoding cell assembly and the entire network (Materials and Methods) in pre-stimulus spontaneous and post-stimulus evoked activity. When the network was trained in a non-biased fashion (i.e., in the 50:50 protocol), the activity ratio was close to 0.5 in spontaneous activity, implying that the two cell assemblies had similar activity levels. In contrast, when the network was trained in a biased fashion (i.e., in the 80:20 protocol), the activity ratio in spontaneous activity was close to 0.8, implying that the total spontaneous firing rate of R-encoding cell assembly was four times higher than that of L-encoding cell assembly. Our results show that the spontaneous activity generated by the proposed mechanism can account for the precise relationship between motion coherence and perceptual biases in decision making by monkeys.

Crucial roles of distinct inhibitory pathways

The model presented so far lacked biological plausibility in several key aspects. Specifically, we assumed that the recurrent connections M could change its sign through plasticity and be either excitatory or inhibitory, while the inhibitory connection G was restricted to being inhibitory only. This setting does not reflect the biological constraint that synapses maintain a consistent excitatory or inhibitory type. Furthermore, due to this unconstrained recurrent connectivity M, the original model had two types of inhibitory connections (i.e., the negative part of M and the inhibitory connection G) without providing a clear computational role for each type of inhibition.

To address these limitations and to understand the role of the two types of inhibition, we considered a novel architecture in which all recurrent connections are constrained to be either exclusively excitatory or inhibitory, maintaining their sign throughout the learning process. The refined model includes two different types of inhibitory connections (i.e., M_{inh} and G), each serving a specific computational purpose: minimizing prediction error and maintaining the excitatory-inhibitory balance. In combination with the excitatory connection M_{exc} , the M_{inh} connections are trained to minimize the prediction error between somatic and dendritic activity, as considered in the original M connection in **Figure 1**. We found that the trained M_{inh} connections introduce competition among cell assembly activities by forming strong connections between assemblies (**Fig. 6b**), allowing the network to effectively sample and replay the activities of individual assemblies. In contrast, inhibitory connections G were trained to balance network dynamics, as in the original setting. We found that the inhibitory G connections form strong intra-assembly inhibition (**Fig. 6c**), which balances the strong excitatory connections that arise within cell assemblies through plasticity (**Fig. 6a**).

In summary, the dual inhibitory mechanism allows the network to perform the reactivation of different cell assemblies while regulating their internal dynamics. The prediction-error-minimizing inhibitory connections M_{inh} facilitate selecting and activating specific assemblies through competition such that the learned probabilities are replayed. In contrast, the network-balancing inhibitory connections G prevent runaway excitation within active assemblies.

An elaborate network model with distinct excitatory and inhibitory neuron pools

The predictive learning rule performed well in training the nDL model to learn the probabilistic structure of the stimulus-evoked activity patterns. However, whether the same learning rule works in a more realistic neural network is yet to be investigated. To examine this, we constructed an elaborate network model (DL model) consisting of distinct excitatory and inhibitory neuron pools, obeying Dale's law (**Fig. 7a**). The nDL model suggested the essential roles of inhibitory plasticity in maintaining excitation-inhibition balance and generating an appropriate number of cell assemblies. To achieve these functions, inhibitory neurons in the DL model project to excitatory and other inhibitory neurons via two synaptic paths (**Fig. 7b**), motivated by the results shown in **Figure 6**. In path1, inhibitory connections alone predict the postsynaptic activity, whereas inhibitory and excitatory connections jointly predict the activity of the postsynaptic neuron in path2 (Materials and Methods). All synapses in the DL model are subject to the predictive learning rule. We trained the DL model with three input neuron groups while varying their activation probabilities. As in the nDL model, the DL model self-organized three cell assemblies activated selectively by the three input neuron groups (Supplementary Fig. 9a). Furthermore, in the absence of external stimuli, the DL model spontaneously replayed these assemblies with the assembly activity ratios in proportion to the occurrence probabilities of the corresponding stimuli during learning (**Fig. 7c**).

The two inhibitory paths divided their labors consistent with the results shown in **Figure 6**. To see this, we investigated the connectivity structures learned by these paths. In path 1, inhibitory connections were primarily found on excitatory neurons in the same assemblies (**Fig. 7d**, top). In contrast, in path 2, inhibitory connections were stronger on excitatory neurons in different assemblies than those in the same assemblies (**Fig. 7d**, bottom). On both excitatory and inhibitory neurons, the total inhibition (i.e., path 1 + path 2) was balanced with excitation (**Fig. 7e**). **Figure 7f** summarizes the connectivity structure of the DL model. Excitatory neurons in a cell assembly project to inhibitory neurons in the same assembly. Then, these inhibitory neurons project back to excitatory neurons in the same or different assemblies via paths 1 and 2. Interestingly, lateral inhibition through path 1 is more potent between excitatory neurons within each cell assembly than between different assemblies (**Fig. 7g**). In contrast, path 2 mediates equally strong within-assembly and between-assembly inhibition.

We can understand the necessity of the two inhibitory paths based on the dynamical properties of competitive neural networks. **Figure 7h** displays the effective competitive network of excitatory cell assemblies suggested by the above results. Both paths 1 and 2 contribute to within-assembly inhibition among excitatory neurons, whereas between-assembly inhibition (i.e., lateral inhibition) mainly comes from path 2. In a competitive network, the lateral inhibition to self-inhibition strength ratio determines the number of winners having non-vanishing activities: the higher the ratio is, the smaller the number of winners is (Fukai & Tanaka, 1997). Therefore, self-organizing the same number of excitatory cell assemblies as that of external stimuli requires tuning the balance between the within-assembly and between-assembly inhibitions. This tuning during learning is likely easier when the network has two independently learnable inhibitory circuits. Indeed, a network model with only one inhibitory path rarely succeeded in encoding and replaying all stimuli used in learning (Supplementary Fig. 9b, c).

In summary, we have shown the roles of distinct recurrent inhibitory connections. Using a network consisting of excitatory and inhibitory populations, we have shown that distinct inhibitory circuits are necessary to generate within- and between-assembly competition crucial to maintain the stability of learned multiple assemblies.

Discussion

Having proper generative models is crucial for accurately predicting statistical events. The brain is thought to improve the prediction accuracy of inference by learning internal generative models of the environment. These models are presumably generated through multiple mechanisms. For instance, the predictive coding hypothesizes that top-down cortical inputs provide lower sensory areas with prior information about sensory experiences (Friston, 2010 [↗](#); Bastos et al., 2012 [↗](#); Keller & Mrsic-Flogel, 2018 [↗](#)). However, experimental evidence also suggests that spontaneous activity represents an optimal model of the environment in sensory cortices. This study proposed a biologically plausible mechanism to learn such a model, or priors for experiences, with the brain's internal dynamics.

Our model adopted a single predictive learning principle for the plasticity of excitatory and inhibitory synapses to learn the replay of probabilistic experiences. On each neuron, excitatory and inhibitory synaptic weights undergo plastic changes to improve their independent predictions on the cell's firing. This was done by minimizing the mismatch between the output firing rate and the network predictions (Eqs. 13 [↗](#) and 17 [↗](#)). This simple learning rule showed excellent performance in a simplified network model and in a more realistic model obeying Dale's law. The latter model predicts a division of labor between two inhibitory paths. Intriguingly, the inhibitory path 2 of this model resembles interpyramidal inhibitory connections driven directly by nearby pyramidal cells (Ren et al., 2007 [↗](#)). In both models, inhibitory synaptic plasticity plays a crucial role in learning an accurate internal model by maintaining excitation-inhibition balance and decorrelating cell-assembly activities (Vogels et al., 2013 [↗](#); Sprekeler, 2017 [↗](#)). It should be noted that while several network models that perform error-based computations like ours exploit only inhibitory recurrent plasticity (Mikulasch et al., 2021 [↗](#); Mackwood et al., 2021 [↗](#); Hertäg and Clopath, 2022 [↗](#); Mikulasch et al., 2023 [↗](#)), our model learns to reproduce appropriate statistics by modifying both excitatory and inhibitory recurrent connections.

Various models have been proposed to account for neural mechanisms of Bayesian computation by the brain (Tully et al., 2014 [↗](#); Kappel et al., 2015 [↗](#); Hiratani & Fukai, 2018 [↗](#); Hiratani & Latham, 2020 [↗](#); Aitchison et al., 2021 [↗](#); Ma et al., 2006 [↗](#); Deneve, 2008 [↗](#); Nessler et al., 2013 [↗](#); Hiratani & Fukai, 2016 [↗](#); Huang & Rao, 2016 [↗](#); Isomura et al., 2022 [↗](#); Friston, 2010 [↗](#); Bastos et al., 2012 [↗](#); Keller & Mrsic-Flogel, 2018 [↗](#)). Typically, these models embed prior knowledge on sensory experiences into the wiring patterns of afferent (and sometimes also recurrent) synaptic inputs such that these inputs can evoke the learned activity patterns associated with the prior knowledge. The present model differs from the previous models in several aspects: i) First, the model segments repeated stimuli to be remembered in an unsupervised fashion; ii) Then it generates cell assemblies encoding the segmented stimuli; (iii) Finally, it replays these cell assemblies spontaneously with learned probabilities. Note that the same learning rules enable the network to perform all necessary computations for (i) to (iii). To our knowledge, our model is the first to perform all these steps for encoding an optimal model of the environment into spontaneous network activity.

The present mechanism of memory formation differs from the previous ones that self-organize cell assemblies through Hebbian learning rules (Vogels et al., 2011 [↗](#); Hiratani and Fukai, 2014 [↗](#); Litwin-Kumar and Doiron, 2014 [↗](#); Zenke et al., 2015 [↗](#); Triplett et al., 2018 [↗](#); Montangie 2020 [↗](#)). Firstly, these mechanisms did not aim for explicit statistical modeling of the environment. Secondly, the previous studies suggested that the orchestration of multiple plasticity rules,

including inhibitory plasticity and homeostatic synaptic scaling, enables the maintenance of cell assemblies (however, see [Manz et al., 2023](#)). For instance, in spike-timing-dependent plasticity (STDP), slight changes in the relative times of pre and postsynaptic spikes can change the polarity of synaptic modifications, implying that STDP requires a mechanism to keep synaptic weights finite ([Kempster et al., 1999](#); [Song et al., 2000](#); [Masquelier et al., 2008](#)). In contrast, our learning rule, which induces either long-term potentiation or depression according to the sign of the prediction error calculated independently within each postsynaptic neuron, does not suffer such instability.

Our model predicts a novel intracellular process that regulates the neuron's dynamic range according to the history of its subthreshold dynamics. This process plays two important roles in the statistical modeling of our model. First, it avoids the trivial solution (i.e., the zero-weight solution) of our unsupervised predictive learning by homeostatically regulating neurons' intrinsic excitability. Second, the intracellular process cooperates with reverberating synaptic inputs within each cell assembly to generate spontaneous replay activity. We have shown that the intracellular homeostasis enhances the sampling from learned distribution without relying on the recurrences among assemblies. This mechanism contrasts with the previous sampling-based models that rely on the transition dynamics between cell assemblies ([Buesing et al., 2011](#); [Bill et al., 2015](#)). How neural systems implement the proposed homeostasis is an open question.

Previous computational models have demonstrated that recurrent networks with Hebbian-like plasticity can learn appropriate Markovian statistics ([Kappel et al., 2014](#); [Asabuki and Clopath, 2024](#)). However, our model differs conceptually from these previous models. Kappel et al. showed that STDP in winner-take-all circuits can approximate online learning of hidden Markov models (HMMs). A key difference with our model is that their neural representations acquire sequences using Markovian sampling dynamics, whereas our model does not depend on such ordered sampling. Specifically, in their model, sequential sampling arises from the off-diagonal elements learned in the recurrent connections (i.e., between-assembly connections). In contrast, our network learns to generate a stochastic reactivation of cell assemblies solely by within-assembly connections. A similar argument can be made for Asabuki and Clopath paper as well. Further, while our model introduced plasticity at all types of synaptic connections, the previous model assumed plasticity only at recurrent synapses projecting onto the excitatory neurons. In addition, unlike our model, the cell assembly memberships need to be preconfigured in the previous model.

The proposed mechanism can account for the behavioral biases observed in perceptual decision making ([Hanks et al., 2011](#)). This behavioral experiment quantitatively clarified how the difference in the probability between sensory experiences during learning biases the alternative choice behavior of monkeys. In our model, two cell assemblies encoding the different stimuli are replayed at the total firing rates proportional to the corresponding occurrence probabilities. Our results suggest that the difference in spontaneous firing rates of cell assemblies is sufficient to explain the behavioral biases of monkeys. However, other mechanisms, such as biased top-down input, cannot be excluded.

What could be the advantages of coding prior distributions into spontaneous activity over other ways of probability coding? First, spontaneous replay activities in lower cortical areas may provide training data for modeling by higher cortical areas, promoting hierarchical statistical modeling in predictive coding. This is analogous to the situation where hippocampal engram cells are replayed to reinforce the activity patterns of cortical engrams for memory consolidation during sleep ([Tonegawa et al., 2018](#); [Ghandour et al., 2019](#); [Klinzing et al., 2019](#); [Takehara-Nishiuchi, 2021](#)). Memory reinforcement by activity replay has also been studied in machine intelligence ([Dayan et al., 1995](#); [Goodfellow et al., 2014](#); [Luczak et al., 2022](#)). Second, spontaneous replay of internal models may support knowledge generalization during sleep. It was recently reported that a transitive inference task requires post-learning sleep ([Kareem et al.,](#)

2021 [DOI:10.7554/eLife.92712.2](#)). In this task, mice had to infer a correct reward delivery rule in a novel behavioral situation from the outcomes of past experiences. The mice failed to generalize the learned rules if the activity of the anterior cingulate cortex was suppressed during post-learning sleep, suggesting that dynamic interactions among rule-coding cortical neurons in spontaneous activity are crucial for rule generalization. Clarifying how spontaneous brain activity generalizes the learned internal models is an intriguing open question.

Methods

The source codes used in this study are available at: https://github.com/TAsabuki/PriorNet_codes

Neural network model

Below, we first describe the model architecture and learning rule for the nDL model (i.e., single population violating the Dale's law). Details of simulation of distinct excitatory and inhibitory populations will be explained later. Unless otherwise stated, recurrent neural networks used in this study consist of $N (= 500)$ Poisson neurons, which generate spikes according to a non-stationary Poisson process with rate $\hat{\phi}(u)$, where $\hat{\phi}(\cdot)$ is a dynamics sigmoidal function, which we will explain later. The membrane potential u of neuron i at time t is given as follows:

$$u_i(t) = \sum_{k=1}^K W_{ik} x_k(t) + \sum_{k=1}^N (M_{ik} - G_{ik}) y_k(t), \quad (5)$$

where K is the number of input neurons. In some simulations, the network model had more than one input neuron group although the number of input neuron groups is not explicitly shown in Eq. 5. Three matrices $\mathbf{W} \in \mathbb{R}^{N \times K}$, $\mathbf{M} \in \mathbb{R}^{N \times N}$, and $\mathbf{G} \in \mathbb{R}^{N \times N}$ represent the weights of afferent synaptic connections, recurrent synaptic connections and inhibitory-only connections, respectively, on neurons in the recurrent network. These synaptic connections are all-to-all. In terms of the kernel function

$$\varepsilon(s) = \exp(-s/\tau) \cdot \Theta(s), \quad (6)$$

recurrent input and afferent input to neuron i are calculated as

$$x_i(t) = \sum_{t' \in t_{\text{aff}}^f} \varepsilon(t - t'), \quad (7a)$$

$$y_i(t) = \sum_{t' \in t_{\text{rec}}^f} \varepsilon(t - t'), \quad (7b)$$

where τ stands for the membrane time constant, t_{aff}^f and t_{rec}^f for the time sets of afferent and recurrent presynaptic spikes, and $\Theta(\cdot)$ for the Heaviside function.

Throughout this study, $\tau = 15$ ms.

The instantaneous firing rate $f_i(t)$ of each neuron is given as

$$f_i(t) = \hat{\phi}(u_i(t); h_i), \quad (8)$$

in terms of a dynamical sigmoidal response function $\hat{\phi}$:

$$\hat{\phi}(u_i; h_i) = \varphi_0 [1 + \exp[g\beta(h_i)(-u_i + g\theta(h_i))]]^{-1}, \quad (9)$$

with a constant value of $g = 3$ and we have dropped the explicit time dependence in our notation for the sake of simplicity. Here, the dynamical variable h is determined by the history of the membrane potential:

$$\begin{aligned} \dot{h}_i &= -\tau_h^{-1} h_i, & \text{if } h_i > u_i, \\ h_i &\leftarrow u_i, & \text{otherwise.} \end{aligned} \quad (10)$$

The maximum instantaneous firing rate ϕ_0 is 50 Hz and $\tau_h = 10$ s. Through Eq. 10, h_i tracks the maximum value of the membrane potential u_i in a time window of approximately the length τ_h in the immediate past. The value of h is utilized to regulate the gain β and threshold θ of the sigmoidal response function as follows:

$$\beta(h_i) = h_i^{-1} \beta_0, \quad (11)$$

$$\theta(h_i) = h_i \theta_0, \quad (12)$$

where the values of constant parameters are $\beta_0 = 5$, and $\theta_0 = 1$. Neuron i generates a Poisson spike train at the instantaneous firing rate of $f_i(t)$. While a small value of h leads to a steep slope of our activation function (Eq. 11), we have shown numerically that this does not lead to a problem in neural dynamics. Further, the saturation part of the sigmoidal function is crucial for stable formation of assemblies.

Learning rules

We first explain the plasticity rule for feedforward connections. Synaptic connections were modified to minimize the Kullback–Leibler divergence (KL-divergence) between two Poisson distributions associated with the neuron’s output and the feedforward activity over a sufficiently long period T :

$$\mathcal{L}_W = \int_0^T dt \sum_{i=1}^N D_{KL}[f_i(t) || \phi(v_i^W(t))], \quad (13)$$

where v_i^W is a feedforward prediction of a firing rate, defined as:

$$v_i^W = \sum_{j=1}^K W_{ij} \cdot x_j, \quad (14)$$

and f_i is the firing rate of i -th neuron. The function ϕ is a static sigmoidal function, defined as:

$$\phi(v_i) = \phi_0 [1 + \exp[g\beta_0(-v_i + g\theta_0)]]^{-1}. \quad (15)$$

The above cost function evaluates to what extent the feedforward potential predicts the activity of postsynaptic neurons (Asabuki and Fukai, 2020). We have previously shown that taking the gradient of the cost function in Eq.13 derives the online plasticity rule for the feedforward connections as:

$$\Delta W_{ij} = \eta \phi_0^{-1} \left(1 - \frac{\phi(v_i^W)}{\phi_0} \right) [f_i - \phi(v_i^W)], \quad (16)$$

where η is a learning rate and was set to $\eta = 10^{-4}$, unless otherwise specified. Here, we have dropped the explicit time dependence in our notation for the sake of simplicity.

Similarly, the recurrent connections were modified to minimize the following cost function:

$$\mathcal{L}_M = \int_0^T dt \sum_{i=1}^N D_{KL}[f_i(t) || \varphi(v_i^M(t))], \quad (17)$$

where $v_i^M = \sum_{j=1}^N M_{ij} \cdot y_j$ is a recurrent prediction. Similar to the feedforward plasticity, the gradient descent of the above cost function leads to the following plasticity rule:

$$\Delta M_{ij} = \eta \varphi_0^{-1} \left(1 - \frac{\varphi(v_i^M)}{\varphi_0} \right) [f_i - \varphi(v_i^M)]. \quad (18)$$

The derived recurrent plasticity rule suggests that the recurrent prediction learns the statistical model of the evoked activity, which in turn allows the network to replay the learned internal model.

In addition to the above plasticity rules, we defined the cost function for the inhibitory plasticity as:

$$\mathcal{L}_G = \sum_{i=1}^N [f_i(t) - \varphi(v_i^G(t))]^2, \quad (19)$$

where v_i^G is the inhibitory input onto postsynaptic neuron via inhibitory connection G :

$$v_i^G = \sum_{j=1}^N G_{ij} \cdot y_j. \quad (20)$$

Again, by taking the gradient of $\mathcal{L}_F$ with respect to G_{ij} derive the following inhibitory plasticity rule to keep the network dynamics balanced:

$$\begin{aligned} \Delta G_{ij} &\propto -\frac{\partial \mathcal{L}_G}{\partial G_{ij}} \\ &\propto [f_i - \varphi(v_i^G)] \times \frac{\partial \varphi(v_i^G)}{\partial G_{ij}} \\ &\propto \varphi(v_i^G) \left(1 - \frac{\varphi(v_i^G)}{\varphi_0} \right) \cdot [f_i - \varphi(v_i^G)] \cdot y_j. \end{aligned} \quad (21)$$

While the resultant rule is not the same as feedforward and recurrent plasticity rules, all of these rules are similar in a sense that the weight updates are proportional to the prediction error and the presynaptic activity. We therefore assumed the following rule for the inhibitory plasticity, which has the same structure with the rest of the plasticity rules that we have already explained:

$$\Delta G_{ij} = \eta \varphi_0^{-1} \left(1 - \frac{\varphi(v_i^G)}{\varphi_0} \right) [f_i - \varphi(v_i^G)]. \quad (22)$$

We have shown by numerical simulation that the rule keeps the network dynamics balanced.

Initial values of $\mathbf{W}$ and $\mathbf{M}$ are sampled from gaussian distributions with the mean 0 and variances $0.1/\sqrt{K}$ and $0.1/\sqrt{N}$, respectively. During learning, the elements of $\mathbf{W}$ and $\mathbf{M}$ can take both positive and negative values. After sufficient learning, the postsynaptic potentials v_i^W and v_i^M on neuron on neuron i converge to a common value of v_i^* . Therefore,

$\varphi(v_i^W) \approx \varphi(v_i^M) \approx \varphi(v_i^*) \approx f_i$, implying that the postsynaptic potentials of afferent and recurrent synaptic inputs to neuron i can both predict its output f_i after learning. The initial values of $\mathbf{G}$ are uniformly set to $1/\sqrt{N}$, and its elements are truncated to non-negative values during

learning. This implies that v_i^G does not become negative. After learning, $\varphi(v_i^G) \approx f_i$ is satisfied. Although some elements of $\mathbf{M}$ may give recurrent inhibitory connections, modifiable connections in $\mathbf{G}$ are necessary to encode all external inputs into specific cell assemblies.

Stimulation protocols

Feedforward input to the recurrent network consisted of K Poisson spike trains with a background firing rate of 2 Hz. The input randomly presented n non-overlapping patterns of 100 spike trains (the duration 100 ms and the mean frequency 50 Hz), one at a time, with pattern-to-pattern intervals of 100 ms. Therefore, the number of input neurons and patterns satisfy the relationship of $K = 100 \times n$. For simplicity, we simulated the constant-interval case, but using irregular intervals does not change the essential results. The value of n varies from task to task, and the values for each figure are as follows: $n = 5$ (Fig. 2c-e, Fig. 3, Fig. 6); $n = 2$ (Fig. 2a-b, Figs. 4-5); $n = 3$ (Fig. 7). The typical time length required for the convergence of learning is 1,000 s.

Measures for cell assembly activities

Here, we explain the measures used in Fig. 3. We calculated the firing rate ratio of cell assembly 1 in Fig. 3c as follows:

$$\text{Firing rate ratio} = \frac{r_i^{(1)}}{\left(\sum_{j=2}^5 \frac{1}{N_j} \sum_{i=1}^{N_j} r_i^{(j)}\right)/4}, \quad (23)$$

using the average firing rate $r_i^{(j)}$ of the i -th neuron in cell assembly j and the number N_j of neurons belonging to the cell assembly. Similarly, we defined the assembly size ratio of cell assembly 1 as

$$\text{Assembly size ratio} = \frac{N_1}{\left(\sum_{j=2}^5 N_j\right)/4}. \quad (24)$$

in Fig. 3d and assembly activity ratio of cell assembly 1 as

$$\text{Assembly activity ratio} = \frac{r_{\text{pop}}^{(1)}}{\left(\sum_{i=2}^5 r_{\text{pop}}^{(i)}\right)/4}, \quad (25)$$

in Fig. 3e. Here, $r_{\text{pop}}^{(i)}$ represents the population neural activity of cell assembly i :

$$r_{\text{pop}}^{(j)} \equiv \sum_{i=1}^{N_j} r_i^{(j)}. \quad (26)$$

Simulations of perceptual decision making

In each learning trial, we trained the network with either leftward or rightward dot movement represented by the corresponding input neurons firing at $r_{\text{max}} = 50$ Hz. In test trials, we defined input coherence as $\text{Coh} = \rho_R - 0.5$ according to Hanks et al. (2011), where ρ_R is the ratio of R input neurons to the sum of R and L input neurons in firing rate. The value of Coh ranges between -0.5 (all dots moving leftward) and +0.5 (all dots moving rightward). Then, in test trials for input coherence Coh, we generated Poisson spike trains of R and L input neurons at the rates $(\text{Coh} + 0.5)r_{\text{max}}$ and $(-\text{Coh} + 0.5)r_{\text{max}}$, respectively.

In Fig. 5c, we calculated the activity ratio (AR) as

$$\text{AR} = \frac{r_{\text{R}}^{\text{pop}}}{r_{\text{R}}^{\text{pop}} + r_{\text{L}}^{\text{pop}}}, \quad (27)$$

where r_R^{pop} and r_L^{pop} represent the average population firing rates of R-encoding and L-encoding cell assemblies, respectively. In **Fig. 5b**, we defined “choices to right” as

$$\text{Choices to right} = \text{AR} \times 100 (\%). \quad (28)$$

A network model with distinct excitatory and inhibitory synapses

Here, we explain the network model and the plasticity rules used in **Fig. 6**. The network consists of 500 neurons and the membrane potential of a neuron i at time t is given as follows:

$$u_i^{\square}(t) = \underbrace{\sum_{k=1}^K W_{ik}^{\square} x_k(t)}_{=:v_i^W} + \underbrace{\left[\sum_{n=1}^N M_{in}^{\text{exc}} y_n^{\square}(t) \right]}_{=:v_i^{M(\text{exc})}} - \underbrace{\left[\sum_{n=1}^N M_{in}^{\text{inh}} y_n^{\square}(t) \right]}_{=:v_i^{M(\text{inh})}} - \underbrace{\sum_{n=1}^N G_{in}^{\square} y_n^{\square}(t)}_{=:v_i^G}, \quad (29)$$

where $\{W_{ik}^{\square}\}$ is afferent synaptic weights, which are a mixture of excitatory and inhibitory connections as in the nDL model. The weights of recurrent excitatory synapses are $\{M_{in}^{\text{exc}}\}$. Here, we considered two types of recurrent inhibitory connections, denoted by M_{in}^{inh} and G , respectively. Here, we assumed that half of the recurrent connections were assumed to be excitatory and the remaining connections were all inhibitory, half of which were M_{in}^{inh} and the other half were G . We modified these weights according to the following equations:

$$\Delta W_{ij}^{\square} = \eta \mathcal{E}(f_i, v_i^W) x_j, \quad (30a)$$

$$\Delta M_{ij}^{\text{exc}} = \eta \mathcal{E}(f_i, v_i^{M(\text{exc})} - v_i^{M(\text{inh})}) y_j^{\square}, \quad (30b)$$

$$\Delta M_{ij}^{\text{exc}} = -\eta \mathcal{E}(f_i, v_i^{M(\text{exc})} - v_i^{M(\text{inh})}) y_j^{\square}, \quad (30c)$$

$$\Delta G_{ij}^{\square} = \eta \mathcal{E}(f_i, v_i^G) y_j^{\square}, \quad (30d)$$

where $\mathcal{E}(f_i, v_i^{\square})$ is the error term defined as:

$$\mathcal{E}(f_i, v_i^{\square}) = \varphi_0^{-1} \left(1 - \frac{\varphi(v_i^{\square})}{\varphi_0} \right) [f_i - \varphi(v_i^{\square})]. \quad (31)$$

At each time step during learning, we truncated all weights of recurrent connections to non-negative values during learning.

A network model with distinct excitatory and inhibitory neuron populations

Here, we explain the architecture of the model used in **Fig. 7**. The network consists of $N_E (= 500)$ excitatory and $N_I (= 500)$ inhibitory neurons. The membrane potential of a neuron i of a population X ($X = E$ or I) at time t is given as follows:

$$u_i^X(t) = \underbrace{\sum_{k=1}^K W_{ik}^X x_k(t)}_{=:v_i^W} + \underbrace{\left[\sum_{l=1}^{N_E} M_{il}^{\text{XE}} y_l^E(t) - \sum_{m=1}^{N_I} G_{im}^{\text{XI}(\text{path2})} y_m^I(t) \right]}_{=:v_i^{M(2)}} - \underbrace{\sum_{m=1}^{N_I} G_{im}^{\text{XI}(\text{path1})} y_m^I(t)}_{=:v_i^{M(1)}}, \quad (32)$$

where $\{W_{ik}^X\}$ is afferent synaptic weights, which are a mixture of excitatory and inhibitory connections as in the nDL model. The weights of recurrent excitatory synapses are $\{M_{il}^{XE}\}$. Here, we considered two types of recurrent inhibitory connections (i.e., path 1 and path 2), denoted by $G_{im}^{XI \text{ (path1)}}$ and $G_{im}^{XI \text{ (path2)}}$, respectively. Using the same definitions of the error term as in [Eq. 31](#), we modified these weights according to the following equations:

$$\Delta W_{ij}^X = \eta \mathcal{E}(f_i, v_i^W) x_j, \quad (33a)$$

$$\Delta M_{ij}^{XE} = \eta \mathcal{E}(f_i, v_i^{M(2)}) y_j^E, \quad (33b)$$

$$\Delta G_{ij}^{XI \text{ (path2)}} = -\eta \mathcal{E}(f_i, v_i^{M(2)}) y_j^I, \quad (33c)$$

$$\Delta G_{ij}^{XI \text{ (path1)}} = \eta \mathcal{E}(f_i, v_i^{M(1)}) y_j^I. \quad (33d)$$

To satisfy the Dale's law, we truncated all weights of recurrent connections to nonnegative values during learning.

In [Fig. 7g](#), we measured the lateral inhibition between excitatory neurons via path 1 by calculating:

$$[W_{i,j}^{LI}] = \sum_{k=1}^{N_E} G_{ik}^{EI \text{ (path1)}} M_{kj}^{IE}. \quad (34)$$

Lateral inhibition via path 2 was calculated in a similar fashion.

Simulation details

All simulations were performed in customized Python3 code written by TA with numpy 1.17.3 and scipy 0.18. Differential equations were numerically integrated using a Euler method with integration time steps of 1 ms.

Author Contributions

T.A. and T.F. conceived the study and wrote the paper.

T.A. performed the simulations and data analyses.

Competing Interest Statement

The authors declare no competing interests.

Acknowledgements

The authors express their sincere thanks to Yukiko Goda for her valuable comments on our manuscript. This work was supported by KAKENHI (nos. 18H05213 and 19H04994) to T.F.

References

1. Aitchison L., Jegminat J., Menendez J.A., Pfister J.P., Pouget A., Latham P.E., et al. (2021) **Synaptic plasticity as Bayesian inference** *Nat. Neurosci* **24**:565–571
2. Asabuki T., Fukai T. (2020) **Somatodendritic consistency check for temporal feature segmentation** *Nat. Commun* **11**
3. Asabuki T., Clopath C. (2024) **Embedding stochastic dynamics of the environment in spontaneous activity by prediction-based plasticity** *eLife* **13**
4. Ashourian P., Loewenstein Y. (2011) **Bayesian inference underlies the contraction bias in delayed comparison tasks** *PLoS One* **6**
5. Baldassano C., Chen J., Zadbood A., Pillow J.W., Hasson U., Norman K.A. (2017) **Discovering event structure in continuous narrative perception and memory** *Neuron* **95**:709–721
6. Bastos A.M., Usrey W.M., Adams R.A., Mangun G.R., Fries P., Friston K.J. (2012) **Canonical microcircuits for predictive coding** *Neuron* **76**:695–711
7. Berkes P., Orban G., Lengyel M., Fiser J. (2011) **Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment** *Science* **331**:83–87
8. Bill J., Buesing L., Habenschuss S., Nessler B., Maass W., Legenstein R. (2015) **Distributed bayesian computation and self-organized learning in sheets of spiking neurons with local lateral inhibition** *PloS one* **10**
9. Buesing L., Bill J., Nessler B., Maass W. (2011) **Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons** *PLoS computational biology* **7**
10. Chance F. S., Abbott L. F., Reyes A. (2002) **D. Gain modulation from background synaptic input** *Neuron* **35**:773–782
11. Dayan P., Hinton G.E., Neal R.M., Zemel R.S. (1995) **The Helmholtz machine** *Neural Comput* **7**:889–904
12. Dehaene S., Meyniel F., Wacongne C., Wang L., Pallier C. (2015) **The neural representation of sequences: from transition probabilities to algebraic patterns and linguistic trees** *Neuron* **88**:2–19
13. Deneve S. (2008) **Bayesian spiking neurons I: inference** *Neural Comput* **20**:91–117
14. Ernst M.O., Banks M.S. (2002) **Humans integrate visual and haptic information in a statistically optimal fashion** *Nature* **415**:429–433
15. Fiser J., Berkes P., Orban G., Lengyel M. (2010) **Statistically optimal perception and learning: from behavior to neural representations** *Trends Cogn. Sci* **14**:119–130

16. Fonollosa J., Neftci E., Rabinovich M. (2015) **Learning of chunking sequences in cognition and behavior** *PLoS computational biology* **11**
17. Friston K. (2010) **The free-energy principle: a unified brain theory?** *Nat. Rev. Neurosci* **11**:127–138
18. Fritsche M., Spaak E., de Lange F.P. (2020) **A Bayesian and efficient observer model explains concurrent attractive and repulsive history biases in visual perception** *Elife* **9**
19. Fujii N., Graybiel A. M. (2003) **Representation of action sequence boundaries by macaque prefrontal cortical neurons** *Science* **301**:1246–1249
20. Fukai T., Tanaka S. (1997) **A simple neural network exhibiting selective activation of neuronal ensembles: from winner-take-all to winners-share-all** *Neural Comput* **9**:77–97
21. Ghandour K *et al.* (2019) **Orchestrated ensemble activities constitute a hippocampal memory engram** *Nat Commun* **10**
22. Goodfellow I., *et al.* (2014) **Generative Adversarial Nets** *Proceedings of the International Conference on Neural Information Processing Systems (NIPS 2014)* :2672–2680
23. Hachen I., Reinartz S., Brasselet R., Stroligo A., Diamond M.E. (2021) **Dynamics of history-dependent perceptual judgment** *Nat. Commun* **12**
24. Hanks T.D., Mazurek M.E., Kiani R., Hopp E., Shadlen M.N. (2011) **Elapsed decision time affects the weighting of prior probability in a perceptual decision task** *J. Neurosci* **31**:6339–6352
25. Hartmann C., Lazar A., Nessler B., Triesch J. (2015) **Where’s the noise? Key features of spontaneous activity and neural variability arise through learning in a deterministic network** *PLoS computational Computational bBiology* **11**
26. Hertäg L., Clopath C. (2022) **Prediction-error neurons in circuits with multiple neuron types: Formation, refinement, and functional implications** *Proceedings of the National Academy of Sciences* **119**
27. Hiratani N., Fukai T. (2016) **Hebbian wiring plasticity generates efficient network structures for robust inference with synaptic weight plasticity** *Front. Neural Circuits* **10**
28. Hiratani N., Fukai T. (2014) **Interplay between short- and long-term plasticity in cell-assembly formation** *PLoS One* **9**
29. Hiratani N., Fukai T. (2018) **Redundancy in synaptic connections enables neurons to learn optimally** *Proc. Natl. Acad. Sci. U. S. A* **115**:E6871–E6879
30. Hiratani N., Latham P.E. (2020) **Rapid Bayesian learning in the mammalian olfactory system** *Nat. Commun* **11**
31. Huang Y., Rao R.P. (2016) **Bayesian Inference and Online learning in Poisson neuronal networks** *Neural Comput* **28**:1503–1526
32. Isomura T., Shimazaki H., Friston K.J. (2022) **Canonical neural networks perform active inference** *Commun. Biol* **5**

33. Jimenez Rezende D., Gerstner W. (2014) **Stochastic variational learning in recurrent spiking networks** *Front Comput Neurosci* **8** <https://doi.org/10.3389/fncom.2014.00038>
34. Jin. X., Costa R. M. (2010) **Start/stop signals emerge in nigrostriatal circuits during sequence learning** *Nature* **466**:457–462
35. Jin. X., Tecuapetla F., Costa R. M. (2014) **Basal ganglia subcircuits distinctively encode the parsing and concatenation of action sequences** *Nat. Neurosci* **17**:423–430
36. Kappel D., Habenschuss S., Legenstein R., Maass W. (2015) **Network plasticity as Bayesian inference** *PLoS Comput. Biol* **11**
37. Abdou Kareem, Choko Kiriko, Aly Mohamed H., Okubo-Suzuki Reiko, Muramatsu Shin ichi, Inokuchi Kaoru (2021) **Inspiring cognitive inference in a cortical network during REM sleep** *bioRxiv*
38. Keller G.B., Mrosovsky T.D. (2018) **Predictive processing: a canonical cortical computation** *Neuron* **100**:424–435
39. Kempter R., Gerstner W., Van Hemmen J. L. (1999) **Hebbian learning and spiking neurons** *Physical Review E* **59**
40. Kenet T., Bibitchkov D., Tsodyks M., Grinvald A., Arieli A. (2003) **Spontaneously emerging cortical representations of visual attributes** *Nature* **425**:954–956
41. Kiebel S. J., von Kriegstein K., Daunizeau J., Friston K. J. (2009) **Recognizing sequences of sequences** *PLoS computational biology* **5**
42. Klinzing J.G., Niethard N., Born J. (2019) **Mechanisms of systems memory consolidation during sleep** *Nat. Neurosci* **22**:1598–1610
43. Kording K.P., Wolpert D.M. (2004) **Bayesian integration in sensorimotor learning** *Nature* **427**:244–247
44. Legaspi R., Toyozumi T. (2019) **A Bayesian psychophysics model of sense of agency** *Nat. Commun* **10**
45. Li A.A., Wang F., Wu S., Zhang X. (2022) **Emergence of probabilistic representation in the neural network of primary visual cortex** *iScience* **25**
46. Litwin-Kumar A., Doiron B. (2014) **Formation and maintenance of neuronal assemblies through synaptic plasticity** *Nat Commun* **5**:1–12
47. Luczak A., McNaughton B.L., Kubo Y. (2022) **Neurons learn by predicting future activity** *Nat. Mach. Intell* **4**:62–72
48. Ma W.J., Beck J.M., Latham P.E., Pouget A. (2006) **Bayesian inference with probabilistic population codes** *Nat. Neurosci* **9**:1432–1438
49. Mackwood O., Naumann L. B., Sprekeler H. (2021) **Learning excitatory-inhibitory neuronal assemblies in recurrent networks** *Elife* **10**
50. Manz P., Memmesheimer R.M. (2023) **Purely STDP-based assembly dynamics: Stability, learning, overlaps, drift and aging** *PLoS Comput Biol* **19**

51. Masquelier T., Guyonneau R., Thorpe S. J. (2008) **Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains** *PLoS one* **3**
52. Mikulasch F. A., Rudelt L., Priesemann V. (2021) **Local dendritic balance enables learning of efficient representations in networks of spiking neurons** *Proceedings of the National Academy of Sciences* **118**
53. Mikulasch F. A., Rudelt L., Wibral M., Priesemann V. (2023) **Where is the error? Hierarchical predictive coding through dendritic error computation** *Trends in Neurosciences* **46**:45–59
54. Mitchell S. J., Silver R. A. (2003) **Shunting inhibition modulates neuronal gain during synaptic excitation** *Neuron* **38**:433–445
55. Mongillo G., Rumpel S., Loewenstein Y. (2018) **Inhibitory connectivity defines the realm of excitatory plasticity** *Nat Neurosci* **21**:1463–1470
56. Montangie L., Miehl C., Gjorgjieva J. (2020) **Autonomous emergence of connectivity assemblies via spike triplet interactions** *PLoS Comput Biol* **16**
57. Nessler B., Pfeiffer M., Buesing L., Maass W. (2013) **Bayesian computation emerges in generic cortical microcircuits through spike-timing-dependent plasticity** *PLoS Comput. Biol* **9**
58. Orban G., Berkes P., Fiser J., Lengyel M. (2016) **Neural variability and sampling-based probabilistic representations in the visual cortex** *Neuron* **92**:530–543
59. Orban G., Fiser J., Aslin R.N., Lengyel M. (2008) **Bayesian learning of visual chunks by human observers** *Proc. Natl. Acad. Sci. U. S. A* **105**:2745–2750
60. Orbán G., Fiser J., Aslin R. N., Lengyel M. (2008) **Bayesian learning of visual chunks by human observers** *Proceedings of the National Academy of Sciences* **105**:2745–2750
61. Rabinovich M. I., Varona P., Tristan I., Afraimovich V. S. (2014) **Chunking dynamics: heteroclinics in mind** *Frontiers in computational neuroscience* **8**
62. Ren M., Yoshimura Y., Takada N., Horibe S., Komatsu Y. (2007) **Specialized inhibitory synaptic actions between nearby neocortical pyramidal neurons** *Science* **316**:758–761
63. Reynolds J. R., Zacks J. M., Braver T. S. (2007) **A computational model of event segmentation from perceptual prediction** *Cognitive science* **31**:613–643
64. Schapiro A. C., Rogers T. T., Cordova N. I., Turk-Browne N. B., Botvinick M. M. (2013) **Neural representations of events arise from temporal community structure** *Nat. Neurosci* **16**:486–492
65. Song S., Miller K. D., Abbott L. F. (2000) **Competitive Hebbian learning through spike-timing-dependent synaptic plasticity** *Nature neuroscience* **3**:919–926
66. Sprekeler H. (2017) **Functional consequences of inhibitory plasticity: homeostasis, the excitation-inhibition balance and beyond** *Curr. Opin. Neurobiol* **43**:198–203
67. Takehara-Nishiuchi K. (2021) **Neurobiology of systems memory consolidation** *Eur. J. Neurosci* **54**:6850–6863

68. Tonegawa S., Morrissey M.D., Kitamura T. (2018) **The role of engram cells in the systems consolidation of memory** *Nat. Rev. Neurosci* **19**:485–498
69. Torres-Torrel J., Torres B., Carrascal L. (2014) **Modulation of the input-output function by GABAA receptor-mediated currents in rat oculomotor nucleus motoneurons** *The Journal of physiology* **592**:5047–5064
70. Triplett M.A., Avitan L., Goodhill G.J. (2018) **Emergence of spontaneous assembly activity in developing neural networks without afferent input** *PLoS Comput Biol* **14**
71. Tully P.J., Hennig M.H., Lansner A. (2014) **Synaptic and nonsynaptic plasticity approximating probabilistic inference** *Front. Synaptic Neurosci* **6**
72. Urbanczik R., Senn W. (2014) **Learning by the dendritic prediction of somatic spiking** *Neuron* **81**:521–528
73. Vogels T. P., Sprekeler H., Zenke F., Clopath C., Gerstner W. (2011) **Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks** *Science* **334**:1569–1573
74. Vogels T.P., et al. (2013) **Inhibitory synaptic plasticity: spike timing-dependence and putative network function** *Front. Neural Circuits* **7**
75. Zacks J. M., et al. (2001) **Human brain activity time-locked to perceptual event boundaries** *Nat. Neurosci* **18**:449–455
76. Zenke F, Agnes EJ, Gerstner W. (2015) **Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks** *Nat Commun* **6**:1–13

Editors

Reviewing Editor

Panayiota Poirazi

FORTH Institute of Molecular Biology and Biotechnology, Heraklion, Greece

Senior Editor

Panayiota Poirazi

FORTH Institute of Molecular Biology and Biotechnology, Heraklion, Greece

Reviewer #1 (Public review):

The authors have successfully addressed most of the issues raised in the first review. Nevertheless, some of the mentioned problems require further attention, mostly regarding the formal derivation of the learning rules, as well as connections to previous research.

Regarding the derivations of learning rules: The authors have provided Goal functions for each of the plastic neural connections to give some insight into what these connections do. However, as I understand, this does not address the main concern raised in the previous review: Why do these rules lead to overall network dynamics that sample from the input distribution? Virtually all other work on neural sampling that I am aware of (e.g., from Maass Lab, Lengyel Lab, etc.) start from a single goal function for all connections that somehow quantifies the difference of network dynamics from the target distribution. In the presented

work the authors specify different goal functions for the different weights, which does not make clear how the desired network dynamics are ultimately achieved.

This becomes especially evident looking at the two different recurrent connections (M and G). M minimizes the difference between network activity f and recurrent prediction $DKL[f|\phi(My)]$, but why is this alone not enough to ensure a good sampling? G minimizes the squared error $[f-\phi(Gy)]^2$, but what does that mean? The problem is that the goal functions are self-consistent in the sense that both f and $\phi(Gy)$ depend on G, which makes an interpretation very difficult. Ultimately it's easier to interpret this by looking at the plasticity rule and see that it leads to a balance. For G the authors furthermore actually ignore the derived plasticity rule and switch to a rule similar to the one for M, meaning that the actual goal function for G is also something like $DKL[f|\phi(Gy)]$. Overall, an overarching optimization goal for the entire network is missing, which makes the interpretation very difficult. I understand that this might be very difficult to provide at this stage, but the authors should at least point out this shortcoming as an open question for the proposed framework.

Regarding the relation to previous work the authors have provided a lot more detailed discussion, which very much clears up the contributions and novel ideas in their work. Still, there are some claims that are not consistent with the literature. Especially, in lines 767 ff. the authors state that Kappel et al "assumed plasticity only at recurrent synapses projecting onto the excitatory neurons. In addition, unlike our model, the cell assembly memberships need to be preconfigured in the [...] model." This is not correct, as Kappel et al learn both the feed-forward and recurrent connections, hence the main difference is that in Kappel et al sampling is sequential and not random. This is why I mentioned this work in the first review, as it speaks against the authors claims of novelty (719 ff.), which should be adjusted accordingly.

<https://doi.org/10.7554/eLife.92712.2.sa3>

Reviewer #2 (Public review):

Summary:

The paper reconsiders the formation of Hebbian-type assemblies, with their spontaneous reactivation representing the statistics of the sensory inputs, in the light of predictive synaptic plasticity. It convincingly shows that not all plasticity rules can be predictive in the narrow sense. While plasticity for the excitatory synapses (the forward projecting and recurrent ones) are predictive, two types of plasticity in the recurrent inhibition is required: a homeostatic and competitive one.

Details:

Besides the excitatory forward and recurrent connections that are learned based on predictive synaptic plasticity, two types of inhibitory plasticity are considered. A first type of inhibition is homeostatic and roughly balances excitation within the cell assemblies. Plasticity in this type 1 inhibition is also predictive, analogous to the plasticity of the excitatory synapses. However, plasticity in type 2 inhibition is competitive and has a switched sign. Both types of inhibitory plasticity, the predictive (homeostatic) and the anti-predictive (competitive) one, work together with the predictive excitatory plasticity to form cell assemblies representing sensory stimuli. Only if the two types of homeostatic and competitive inhibitory plasticity are present, will the spontaneous replay of the assemblies reflect the statistics of the stimulus presentation.

Critical review:

The simulations include Dale's law, making them more biologically realistic. The paper emphasizes predictive plasticity and introduces type 1 inhibitory plasticity that, by construction, tries to fully explain away the excitatory input. In the absence of external inputs, however, due to the symmetry between the excitatory and inhibitory-type-1 plasticity rules, excitation and inhibition tend to fully cancel each other. Multiple options may solve the dilemma:

(1) As other predictive dendritic plasticity models assume, the presynaptic source for recurrent inhibition is typically less informative than the presynaptic source of excitation, so that inhibition is not able to fully explain away excitation.

(2) Beside the inhibitory predictive plasticity that mirrors the analogous excitatory predictive plasticity, and additional competitive plasticity can be introduced.

The paper chooses solution (2) and suggests an additional inhibitory recurrent pathway that is not predictive, but instead anti-predictive with a reversed sign. The combination of the two types of inhibitory plasticities lead to a stable formation of cell assemblies. The stable target activity of the plasticity rules in a memory recall is not anymore 0, as it would be with only type-1-inhibitory plasticity.

Instead, the target activity of plasticity is now enhanced within a winning assembly, and also positive but reduced in the losing assemblies.

<https://doi.org/10.7554/eLife.92712.2.sa2>

Reviewer #3 (Public review):

Summary:

The work shows how learned assembly structure and its influence on replay during spontaneous activity can reflect the statistics of stimulus input. In particular, stimuli that are more frequent during training elicit stronger wiring and more frequent activation during replay. Past works (Litwin-Kumar and Doiron, 2014; Zenke et al., 2015) have not addressed this specific question, as classic homeostatic mechanisms forced activity to be similar across all assemblies. Here, the authors use a dynamic gain and threshold mechanism to circumnavigate this issue and link this mechanism to a cellular monitoring of membrane potential history.

Strengths:

(1) This is an interesting advance, and the authors link this to experimental work in sensory learning in environments with non-uniform stimulus probabilities.

(2) The authors consider their mechanism in a variety of models of increasing complexity (simple stimuli, complex stimuli; ignoring Dale's law, incorporating Dale's law).

(3) Links a cellular mechanism of internal gain control (their variable h) to assembly formation and the non-uniformity of spontaneous replay activity. Offers a promise of relating cellular and synaptic plasticity mechanisms under a common goal of assembly formation.

Weaknesses:

(1) However, while the manuscript does show that assembly wiring does follow stimulus likelihood, it is not clear how the assembly specific statistics of h reflect these likelihoods. I find this to be a key issue.

(2) The authors model does take advantage of the sigmoidal transfer function, and after learning an assembly is either fully active or near fully silent (Fig. 2a). This somewhat artificial saturation may be the reason that classic homeostasis is not required, since runaway activity is not as damaging to network activity.

(3) Classic mechanisms of homeostatic regulation (synaptic scaling, inhibitory plasticity) try to ensure that firing rates match a target rate (on average). If the target rate is the same for all neurons then having elevated firing rates for one assembly compared to others during spontaneous activity would be difficult. If these homeostatic mechanisms were incorporated, how would they permit the elevated firing rates for assemblies that represent more likely stimuli?

<https://doi.org/10.7554/eLife.92712.2.sa1>

Author response:

The following is the authors' response to the previous reviews.

Public Reviews:

Reviewer #1 (Public Review):

In their manuscript, the authors propose a learning scheme to enable spiking neurons to learn the appearance probability of inputs to the network. To this end, the neurons rely on error-based plasticity rules for feedforward and recurrent connections. The authors show that this enables the networks to spontaneously sample assembly activations according to the occurrence probability of the input patterns they respond to. They also show that the learning scheme could explain biases in decision-making, as observed in monkey experiments. While the task of neural sampling has been solved before in other models, the novelty here is the proposal that the main drivers of sampling are within-assembly connections, and not between-assembly (Markov chains) connections as in previous models. This could provide a new understanding of how spontaneous activity in the cortex is shaped by synaptic plasticity.

The manuscript is well written and the results are presented in a clear and understandable way. The main results are convincing, concerning the spontaneous firing rate dependence of assemblies on input probability, as well as the replication of biases in the decision-making experiment. Nevertheless, the manuscript and model leave open several important questions. The main problem is the unclarity, both in theory and intuitively, of how the sampling exactly works. This also makes it difficult to assess the claims of novelty the authors make, as it is not clear how their work relates to previous models of neural sampling.

We agree with the reviewer that our previous manuscript was not clear regarding the mechanism of the model. We have performed additional simulations and included a derivation of the learning rule to address this, which we explain below.

Regarding the unclarity of the sampling mechanism, the authors state that within-assembly excitatory connections are responsible for activating the neurons according to stimulus probability. However, the intuition for this process is not made clear anywhere in the manuscript. How do the recurrent connections lead to the observed effect of sampling? How exactly do assemblies form from feedforward plasticity? This intuitive unclarity is accompanied by a lack of formal justification for the plasticity rules. The authors refer to a previous publication from the same lab, but it is

difficult to connect these previous results and derivations to the current manuscript. The manuscript should include a clear derivation of the learning rules, as well as an (ideally formal) intuition of how this leads to the sampling dynamics in the simulation.

We have included a derivation of our plasticity rules in lines 871-919 in the revised manuscript. Consistent with our claim that predictive plasticity updates the feedforward and the recurrent synapses to predict output firing rates, we have shown that the corresponding cost function measures the discrepancy among the recurrent prediction, feedforward prediction, and the output firing rate. The resultant feedforward plasticity is the same with our previous rule (Asabuki and Fukai, 2020), which segments the salient patterns embedded in the input sequence. The recurrent plasticity rule suggests that the recurrent prediction learns the statistical model of the evoked activity, enabling the network to replay the learned internal model.

Similarly, for the inhibitory plasticity, we defined a cost function that evaluates the difference between the firing rate and inhibitory potential within each neuron. This rule is crucial for maintaining balanced network dynamics. See our response below for more details on the role of inhibitory plasticity.

Some of the model details should furthermore be cleared up. First, recurrent connections transmit signals instantaneously, which is implausible. Is this required, would the network dynamics change significantly if, e.g., excitation arrives slightly delayed? Second, why is the homeostasis on h required for replay? The authors show that without it the probabilities of sampling are not matched, but it is not clear why, nor how homeostasis prevents this. Third, G and M have the same plasticity rule except for G being confined to positive values, but there is no formal justification given for this quite unusual rule. The authors should clearly justify (ideally formally) the introduction of these inhibitory weights G , which is also where the manuscript deviates from their previous 2020 work. My feeling is that inhibitory weights have to be constrained in the current model because they have a different goal (decorrelation, not prediction) and thus should operate with a completely different plasticity mechanism. The current manuscript doesn't address this, as there is no overall formal justification for the learning algorithm.

First, while the reviewer's suggestion to test with delayed excitation is intriguing and crucial for a more biologically detailed spiking neuron model, we have chosen to maintain the current model configuration. Our use of Poisson spiking neurons, which generate spikes based on instantaneous firing rates, does not heavily depend on precise spike timing information. Therefore, to preserve the simplicity of our results, we kept the model unchanged.

Second, we agree that our previous claim regarding the importance of the memory trace h for sampling may have been confusing. As shown in Supplementary Figure 7b in the revised manuscript, when we eliminated the dynamics of the memory trace, sampling performance did indeed decrease. However, we also observed that the assembly activity ratio continued to show a linear relationship with stimulus probabilities. Based on these findings, we have revised our claim in the manuscript to clarify that the memory trace is primarily critical for firing rate homeostasis, rather than directly influencing sampling within the learned network. We have explained this in ll. 446-448 in the revised manuscript.

Third, we explored a new architecture where all recurrent connections are either exclusively excitatory or inhibitory, keeping their sign throughout the learning process. This change addresses the reviewer's concern about our initial assumption that only the inhibitory connection G was constrained to non-negative values. We found that inhibition plays a crucial role in decorrelation and prediction, helping activate specific assemblies through

competition while preventing runaway excitation within active assemblies. We have explained this in ll.560-593 in the revised manuscript.

Finally, the authors should make the relation to previous models of sampling and error-based plasticity more clear. Since there is no formal derivation of the sampling dynamics, it is difficult to assess how they differ exactly from previous (Markov-based) approaches, which should be made more precise. Especially, it would be important to have concrete (ideally experimentally testable) predictions on how these two ideas differ. As a side note, especially in the introduction (line 90), this unclarity about the sampling made it difficult to understand the contrast to Markovian transition models.

As the reviewer pointed out, previous computational models have demonstrated that recurrent networks with Hebbian-like plasticity can learn appropriate Markovian statistics (Kappel et al., 2014; Asabuki and Clopath, 2024). However, our model differs conceptually from these previous models. While Kappel et al. showed that STDP in winner-take-all circuits can approximate online learning of hidden Markov models (HMMs), a key difference with our model is that their neural representations acquire sequences using Markovian sampling dynamics, whereas our model does not depend on such ordered sampling. Specifically, in their model, sequential sampling arises from learned structures in the off-diagonal elements of the recurrent connections (i.e., between-assembly connections). In contrast, our network learns to stochastically generate recurrent cell assemblies by relying solely on within-assembly connections. A similar argument can be made for Asabuki and Clopath paper as well. Further, while our model introduced plasticity rule for all types of connections, Asabuki and Clopath paper introduced plasticity for recurrent synapses projecting on the excitatory neurons only and the cell assembly memberships were preconfigured unlike our model. We have added additional clarifying sentences in ll. 757-772 of the revised manuscript to elaborate on this point.

There are also several related models that have not been mentioned and should be discussed. In 663 ff. the authors discuss the contributions of their model which they claim are novel, but in Kappel et al (STDP Installs in Winner-Take-All Circuits an Online Approximation to Hidden Markov Model Learning) similar elements seem to exist as well, and the difference should be clarified. There is also a range of other models with lateral inhibition that make use of error-based plasticity (most recently reviewed in Mikulasch et al, Where is the error? Hierarchical predictive coding through dendritic error computation), and it should be discussed how the proposed model differs from these.

We have clarified the difference from previously proposed recurrent network model to perform Markovian sampling. Please see our reply above.

We have also included additional sentence in ll. 704-709 in the revised manuscript to discuss how our model differs from similar predictive learning models: “It should be noted that while several network models that perform errorbased computations like ours exploit only inhibitory recurrent plasticity (Mikulasch et al., 2021; Mackwood et al., 2021; Hertäg and Clopath., 2022; Mikulasch et al., 2023), our model learns the structured spontaneous activity to reproduce the evoked statistics by modifying both excitatory and inhibitory recurrent connections.”

Reviewer #2 (Public Review):

Summary:

The paper considers a recurrent network with neurons driven by external input. During the external stimulation predictive synaptic plasticity adapts the forward and recurrent weights. It is shown that after the presentation of constant stimuli, the network

spontaneously samples the states imposed by these stimuli. The probability of sampling stimulus $x^{(i)}$ is proportional to the relative frequency of presenting stimulus $x^{(i)}$ among all stimuli $i=1, \dots, 5$.

Methods:

Neuronal dynamics:

For the main simulation (Figure 3), the network had 500 neurons, and 5 nonoverlapping stimuli with each activating 100 different neurons where presented. The voltage u of the neurons is driven by the forward weights W via input rates x , the inhibitory recurrent weights G , are restricted to have non-negative weights (Dale's law), and the other recurrent weights M had no sign-restrictions. Neurons were spiking with an instantaneous Poisson firing rate, and each spike-triggered an exponentially decaying postsynaptic voltage deflection. Neglecting time constants of the postsynaptic responses, the expected postsynaptic voltage reads (in vectorial form) as

$$u = Wx + (M - G)f \text{ (Eq. 5)}$$

where $f = \phi(u)$ represents the instantaneous Poisson rate, and ϕ a sigmoidal nonlinearity. The rate f is only an approximation (symbolized by $\approx$) of $\phi(u)$ since an additional regularization variable h enters (taken up in Point 4 below). The initialisation of W and M is Gaussian with mean 0 and variance $1/\sqrt{N}$, N the number of neurons in the network. The initial entries of G are all set to $1/\sqrt{N}$.

Predictive synaptic plasticity:

The 3 types of synapses were each adapted so that they individually predict the postsynaptic firing rate f , in matrix form

$$\Delta W \approx (f - \phi(Wx))x^T$$

$$\Delta M \approx (f - \phi(Mf))f^T$$

$$\Delta G \approx (f - \phi(Mf))f^T \text{ but confined to non-negative values of } G \text{ (Dale's law).}$$

The T tells us to take the transpose, and the $\approx$ again refers to the fact that the ϕ entering in the learning rule is not exactly the ϕ determining the rate, only up to the regularization (see Point 4).

Main formal result:

As the authors explain, the forward weight W and the unconstrained weight M develop such that, in expectations,

$$f \approx \phi(Wx) \approx \phi(Mf) \approx \phi(Gf),$$

consistent with the above plasticity rules. Some elements of M remain negative. In this final state, the network displays the behaviour as explained in the summary.

Major issues:

Point 1: Conceptual inconsistency

The main results seem to arise from unilaterally applying Dale's law only to the inhibitory recurrent synapses G , but not to the excitatory recurrent synapses M .

In fact, if the same non-negativity restriction were also imposed on M (as it is on G), then their learning rules would become identical, likely leading to $M=G$. But in this case, the

network becomes purely forward, $u = Wx$, and no spontaneous recall would arise. Of course, this should be checked in simulations.

Because Dale's law was only applied to G, however, M and G cannot become equal, and the remaining differences seem to cause the effect.

Predictive learning rules are certainly powerful, and it is reasonable to consider the same type of error-correcting predictive learning rule, for instance for different dendritic branches that both should predict the somatic activity. Or one may postulate the same type of error-correcting predictive plasticity for inhibitory and excitatory synapses, but then the presynaptic neurons should not be identical, as it is assumed here. Both these types of error-correcting and error-forming learning rules for same-branches and inhibitory/excitatory inputs have been considered already (but with inhibitory input being itself restricted to local input, for instance).

The model presented above lacked biological plausibility in several key aspects. Specifically, we assumed that the recurrent connection M could change sign through plasticity and be either excitatory or inhibitory, while the inhibitory connection G was restricted to being inhibitory only. This initial setting does not reflect the biological constraint that synapses typically maintain a consistent excitatory or inhibitory type. Furthermore, due to this unconstrained recurrent connectivity M, the original model had two types of inhibitory connections (i.e., the negative part of M and the inhibitory connection G) without providing a clear computational role for each type of inhibition.

To address these limitations and to understand the role of the two types of inhibition, we explored a new architecture where all recurrent connections are either exclusively excitatory or inhibitory, keeping their sign throughout the learning process. This change addresses the reviewer's concern about our initial assumption that only the inhibitory connection G was constrained to non-negative values. We found that inhibition plays a crucial role in prediction and decorrelation, helping activate specific assemblies through competition while preventing runaway excitation within active assemblies. We have explained this in ll. 561593 in the revised manuscript.

Point 2: Main result as an artefact of an inconsistently applied Dale's law?

The main result shows that the probability of a spontaneous recall for the 5 nonoverlapping stimuli is proportional to the relative time the stimulus was presented. This is roughly explained as follows: each stimulus pushes the activity from 0 up towards $f = \phi(Wx)$ by the learning rule (roughly). Because the mean weights W are initialized to 0, a stimulus that is presented longer will have more time to push W up so that positive firing rates are reached (assuming x is non-negative). The recurrent weights M learn to reproduce these firing rates too, while the plasticity in G tries to prevent that (by its negative sign, but with the restriction to non-negative values). Stimuli that are presented more often, on average, will have more time to reach the positive target and hence will form a stronger and wider attractor. In spontaneous recall, the size of the attractor reflects the time of the stimulus presentation. This mechanism so far is fine, but the only problem is that it is based on restricting G, but not M, to non-negative values.

As mentioned above, we have included an additional simulation where all weights are non-negative. We have demonstrated the new results in Figure 6 before presenting the two-population model in the revised manuscript (Figure 7), so that readers can follow the importance of two pathways of inhibitory connections.

Point 3: Comparison of rates between stimulation and recall.

The firing rates with external stimulations will be considerably larger than during replay (unless the rates are saturated).

This is a prediction that should be tested in simulations. In fact, since the voltage roughly reads as $u = Wx + (M - G)f$, and the learning rules are such that eventually $M \approx G$, the recurrences roughly cancel and the voltage is mainly driven by the external input x . In the state of spontaneous activity without external drive, one has $u = (M - G)f$, and this should generate considerably smaller instantaneous rates $f \approx \phi(u)$ than in the case of the feedforward drive (unless f is in both cases at the upper or lower ceiling of ϕ). This is a prediction that can also be tested.

Because the figures mostly show activity ratios or normalized activities, it was not possible for me to check this hypothesis with the current figures. So please show non-normalized activities for comparing stimulation and recall for the same patterns.

We agree with the reviewer that the activity levels of spontaneous and induced activity should be compared. We have shown the distributions of activity level of these activities in our new Figure 2d. As expected, we found that the evoked activity showed stronger activity compared to the spontaneous activity.

Point 4: Unclear definition of the variable h .

The formal definition of $h = h_i$ is given by (suppressing here the neuron index i and the h -index of τ)

$\tau \frac{dh}{dt} = -h$ if $h > u$, (Eq. 10) $h = u$ otherwise.

But if it is only Equation 10 (nothing else is said), h will always become equal to u , or will vanish, i.e. either $h = u$ or $h = 0$ after some initial transient. In fact, as soon as $h > u$, h is decaying to 0 according to the first line. If $u > 0$, then it stops at $u = h$ according to the second line. No reason to change $h = u$ further. If $u \leq 0$ while $h > u$, then h is converging to 0 according to the first line and will stay there. I guess the authors had issues with the recurrent spiking simulations and tried to fix this with some regularization. However as presented, it does not become clear how their regulation works.

We apologize for the reviewer that our definition of h was unclear. As the reviewer pointed out, since the memory trace is always positive and larger than (or equal to) the membrane potential, it is possible that the membrane potential becomes always negative and the memory trace reach to 0 constantly. However, since the network is always balanced between excitatory and inhibitory inputs, and it does not happen that the membrane potential always diverges negatively. In fact, we trained without any manipulations other than the memory trace described in the manuscript, and the network was able to learn the assembly structure stably.

BTW: In Eq. 11 the authors set the gain β to $\beta = \beta_0/h$ which could become infinite and, putatively more problematic, negative, depending on the value of h . Maybe some remark would convince a reader that no issues emerge from this.

We have mentioned in ll. 864-866 in the revised manuscript that no issues emerge from the slope parameter.

Added from discussions with the editor and the other reviewers:

Thanks for alerting me to this Supplementary Figure 8. Yes, it looks like the authors did apply there Dale's law for both the excitatory and inhibitory synapses. Yet, they also

introduced two types of inhibitory pathways converging both to the excitatory and inhibitory neurons. For me, this is a confirmation that applying Dale's law to both excitatory and inhibitory synapses, with identical learning rules as explained in the main part of the paper, does not work.

Adding such two pathways is a strong change from the original model as introduced before, and based on which all the Figures in the main text are based. Supplementary Figure 8 should come with an analysis of why a single inhibitory pathway does not work. I guess I gave the reason in my Points 1-3. Some form of symmetry breaking between the recurrent excitation and recurrent inhibition is required so that, eventually, the recurrent excitatory connection will dominate.

Making the inhibitory plasticity less expressive by applying Dale's law to only those inhibitory synapses seems to be the answer chosen in the Figures of the main text (but then the criticism of unilaterally applying Dale's law).

Applying Dale's law to both types of synapses, but dividing the labor of inhibition into two strictly separate and asymmetric pathways, and hence asymmetric development of excitatory and inhibitory weights, seems to be another option. However, introducing such two separate inhibitory pathways, just to rescue the fact that Dale's law is applied to both types of synapses, is a bold assumption. Is there some biological evidence of such two pathways in the inhibitory, but not the excitatory connections? And what is the computational reasoning to have such a separation, apart from some form of symmetry breaking between excitation and inhibition? I guess, simpler solutions could be found, for instance by breaking the symmetry between the plasticity rules for the excitatory and inhibitory neurons. All these questions, in my view, need to be addressed to give some insights into why the simulations do work.

The reviewer's intuition is correct. To effectively learn cell assembly structures and replay their activities, our model indeed requires two types of inhibitory connections. Please refer to our response above for further details.

Overall, Supplementary Figure 8 seems to me too important to be deferred to the Supplement. The reasoning behind the two inhibitory pathways should appear more prominently in the main text. Without this, important questions remain. For instance, when thinking in a rate-based framework, the two inhibitory pathways twice try to explain the somatic firing rate away. Doesn't this lead to a too strong inhibition? Can some steady state with a positive firing rate caused by the recurrence, in the absence of an external drive, be proven? The argument must include the separation into Path 1 and Path 2. So far, this reasoning has not been entered.

In fact, it might be that, in a spiking implementation, some sparse spikes will survive. I wonder whether at least some of these spikes survive because of the other rescuing construction with the dynamic variable h (Equation 10, which is not transparent, and that is not taken up in the reasoning either, see my Point 4)

Perhaps it is helpful for the authors to add this text in the reply to them.

We have moved the former Supplemental Figure 8 to the main Figure 7. Please see our response above about the role of dual inhibitory connection types.

Reviewer #3 (Public Review):

Summary:

The work shows how learned assembly structure and its influence on replay during spontaneous activity can reflect the statistics of stimulus input. In particular, stimuli that are more frequent during training elicit stronger wiring and more frequent activation during replay. Past works (Litwin-Kumar and Doiron, 2014; Zenke et al., 2015) have not addressed this specific question, as classic homeostatic mechanisms forced activity to be similar across all assemblies. Here, the authors use a dynamic gain and threshold mechanism to circumnavigate this issue and link this mechanism to cellular monitoring of membrane potential history.

Strengths:

(1) This is an interesting advance, and the authors link this to experimental work in sensory learning in environments with non-uniform stimulus probabilities.

(2) The authors consider their mechanism in a variety of models of increasing complexity (simple stimuli, complex stimuli; ignoring Dale's law, incorporating Dale's law).

(3) Links a cellular mechanism of internal gain control (their variable h) to assembly formation and the non-uniformity of spontaneous replay activity. Offers a promise of relating cellular and synaptic plasticity mechanisms under a common goal of assembly formation.

Weaknesses:

(1) However, while the manuscript does show that assembly wiring does follow stimulus likelihood, it is not clear how the assembly-specific statistics of h reflect these likelihoods. I find this to be a key issue.

We agree that our previous claim regarding the importance of the memory trace h for sampling may have been confusing. As shown in Supplementary Figure 7b, when we eliminated the dynamics of the memory trace, sampling performance did indeed decrease. However, we also observed that the assembly activity ratio continued to show a linear relationship with stimulus probabilities. Based on these findings, we revised our claim in the manuscript to clarify that the memory trace is primarily critical for learning to avoid trivial solutions, rather than directly influencing sampling within the learned network. We have explained this in ll. 446-448 in the revised manuscript.

(2) The authors' model does take advantage of the sigmoidal transfer function, and after learning an assembly is either fully active or nearly fully silent (Figure 2a). This somewhat artificial saturation may be the reason that classic homeostasis is not required since runaway activity is not as damaging to network activity.

The reviewer's intuition is correct. The saturating nonlinearity is important for the network to form stable assembly structures. We have added an additional sentence in ll. 866-868 to mention this.

(3) Classic mechanisms of homeostatic regulation (synaptic scaling, inhibitory plasticity) try to ensure that firing rates match a target rate (on average). If the target rate is the same for all neurons then having elevated firing rates for one assembly compared to others during spontaneous activity would be difficult. If these homeostatic mechanisms were incorporated, how would they permit the elevated firing rates for assemblies that represent more likely stimuli?

LIF neurons) may solve this problem by utilizing spike-timing statistics.

Recommendations for the authors:

Reviewer #1 (Recommendations For The Authors):

Minor issues:

Figure 1: It would be helpful to display the equation for output rate here as well.

We have included the equation in the revised Figure 1a.

Figure 3c: Typo "indivisual neurons".

We have modified the typo. We thank the reviewer for their careful review.

Line 325: Do you mean Figure 3f,g?

We repeated the task with different numbers of stimuli in Supplementary Figure 1c,d.

Line 398: Winner-take-all can be misunderstood, as it typically stands for competition in inference, not in learning.

We have rephrased it as “unstable dynamics” in l. 400

Line 429: Are intra-assembly and within-assembly the same? If so please use consistent terminology.

We have made the terminology consistent.

Line 792 ff.: Please mention that (t) was left away.

We have included a sentence to mention it in ll. 847-848 in the revised manuscript.

Line 817: Should u_i be v_i ?

We have modified the term.

Methods: What is the value of τ_h ?

We have used $\tau_h = 10$ s, which is mentioned in l. 853

<https://doi.org/10.7554/eLife.92712.2.sa0>