Coherent noise enables probabilistic sequence replay in spiking neuronal networks

Younes Bouhadjar[1,2,3,*], Dirk J. Wouters[4], Markus Diesmann[1,5], and Tom Tetzlaff[1]

[1]Institute of Neuroscience and Medicine (INM-6), & Institute for Advanced Simulation (IAS-6), & JARA BRAIN Institute Structure-Function Relationships (INM-10), Jülich Research Centre, Jülich, Germany
[2]Peter Grünberg Institute (PGI-7,10), Jülich Research Centre and JARA, Jülich, Germany
[3]RWTH Aachen University, Aachen, Germany
[4]Institute of Electronic Materials (IWE 2) & JARA-FIT, RWTH Aachen University, Aachen, Germany
[5]Department of Physics, Faculty 1, & Department of Psychiatry, Psychotherapy, and Psychosomatics, Medical School, RWTH Aachen University, Aachen, Germany
[*]y.bouhadjar{at}fz-juelich.de

June 22, 2022

Abstract

Animals rely on different decision strategies when faced with ambiguous or uncertain cues. Depending on the context, decisions may be biased towards events that were most frequently experienced in the past, or be more explorative. A particular type of decision making central to cognition is sequential memory recall in response to ambiguous cues. A previously developed spiking neuronal network implementation of sequence prediction and recall learns complex, high-order sequences in an unsupervised manner by local, biologically inspired plasticity rules. In response to an ambiguous cue, the model deterministically recalls the sequence shown most frequently during training. Here, we present an extension of the model enabling a range of different decision strategies. In this model, explorative behavior is generated by supplying neurons with noise. As the model relies on population encoding, uncorrelated noise averages out, and the recall dynamics remain effectively deterministic. In the presence of locally correlated noise, the averaging effect is avoided without impairing the model performance, and without the need for large noise amplitudes. We investigate two forms of correlated noise occurring in nature: shared synaptic background inputs, and random locking of the stimulus to spatiotemporal oscillations in the network activity. Depending on the noise characteristics, the network adopts various replay strategies. This study thereby provides potential mechanisms explaining how the statistics of learned sequences affect decision making, and how decision strategies can be adjusted after learning.

Significance statement

Humans and other animals often benefit from exploring multiple alternative solutions to a given problem. Such explorative behavior is frequently attributed to noise in the neuronal dynamics. Supplying each neuron or synapse in a neuronal circuit with noise, however, does not necessarily lead to explorative dynamics. If decisions are triggered by the compound activity of ensembles of neurons or synapses, noise averages out, unless it is correlated within these ensembles. This modeling study equips a neuronal sequence-processing circuit with explorative behavior by introducing configurable coherent noise. It contributes to an understanding of the neuronal mechanisms underlying different decision strategies in the face of ambiguity, and highlights the role of coherent network activity during sequential memory recall.

Introduction

Our brains are constantly processing sequences of events, such as during listening to a song or perceiving the texture of an object. Through repeated exposure to these sequences, we effortlessly learn to predict upcoming events. In many circumstances, we have to make a decision of what elements to recall next in response to a cue. A number of previous studies have introduced spiking neuronal network implementations of sequence learning and replay (Klampfl & Maass, 2013; Klos et al., 2018; Maes et al., 2020; Cone & Shouval, 2021; Bouhadjar et al., 2021). The spiking temporal-memory (TM) model proposed in (Bouhadjar et al., 2021) provides an energy efficient sequence processing mechanism with high storage capacity by virtue of its sparse activity. It learns high-order sequences in an unsupervised, continuous manner using biological, local learning rules. After learning, the model successfully predicts upcoming sequence elements in a context dependent manner, and signals the occurrence of non-anticipated stimuli. Moreover, it can autonomously recall learned sequences in response to a cue stimulus.

In nature, cues are often incomplete or ambiguous, and it is not always clear what sequence to recall given the current context. Despite this ambiguity, we usually come to a clear decision on what sequence to recall. A key factor in decision making is reward (Cohen et al., 2007; O'Doherty et al., 2017). In this regard, the optimal decision strategy is the one that maximizes the reward, and is hence referred to as the maximization or exploitation strategy. A number of studies demonstrate that decisions are often made in an apparently suboptimal manner, such as probability matching (Vulkan, 2000; Myers, 2014). In binary choice tasks, for example, where the probability of payoff is higher for one of the two possible choices, it appears most reasonable to always decide for this high-payoff option. Instead, however, humans and other animals often decide for each of the two choices with a probability that approximately matches the payoff probability. While this behavior appears unreasonable at first glance, it may in fact be optimal when taking into account previous (pre-experiment) experiences, such as prior knowledge of changing reward contingencies. In cases where the reward probability or amplitudes change over time, a more explorative behavior is beneficial (Cohen et al., 2007; Shanks et al., 2002). Previous studies suggest that decisions are not only determined by rewards, but also by the frequency of previously experienced input patterns (Bod et al., 2003; Hansen et al., 2012). Accordingly, suboptimal decision strategies may at least partly arise as a consequence of this additional influence of occurrence frequencies.

A number of previous studies propose neuronal network models of decision making in the face of ambiguous or incomplete stimuli. The majority of these models employ some form of intrinsic stochastic dynamics or uncorrelated noise to generate explorative behavior (Buesing et al., 2011; Legenstein & Maass, 2014; Hartmann et al., 2015; Neftci et al., 2016; Jordan et al., 2019). Noise has been introduced in the form of random or non-task-related synaptic background inputs (Jordan et al., 2019), or in the form of synaptic stochasticity (Neftci et al., 2016). An alternative solution is provided in the studies of Hartmann et al. (2015) and Dold et al. (2019), where the "noise" is generated by the complex but deterministic dynamics of the functional network itself, without any additional sources of stochasticity. In most models, the noise targeting different neurons or synapses is effectively uncorrelated. Supplying each element in a neuronal circuit with uncorrelated noise, however, does not necessarily lead to exlorative dynamics: state variables arising from superpositions of many noisy, uncorrelated components become effectively deterministic as a result of noise averaging (Dold et al., 2019). The total input current of a neuron generated from superpositions of many synaptic inputs, for example, is largely unaffected by the variability in the individual synaptic responses. Similarly, in models where individual states are encoded by the activity of neuronal populations (Legenstein & Maass, 2014), the state representations become quasi deterministic if the single-neuron noise components are uncorrelated. Compensating this noise averaging effect by increasing the noise amplitude appears to be an obvious strategy, but may be hard to realize by the biological system. Alternatively, noise averaging can be avoided by employing correlated noise. As an analogy, consider a particle in a still fluid: despite the constant bombardment by surrounding molecules, a large particle will hardly undergo any Brownian motion, because the momenta of the impinging molecules largely average out. Only if the molecules move in a coherent manner, such as in a turbulent fluid, they can have a substantial influence on the particle's motion.

In biological neuronal networks, coherent noise may arise by different mechanisms: neighboring neurons typically receive inputs from partly overlapping presynaptic neuron populations. The synaptic input currents to these neurons are therefore correlated. In the literature, this type of correlation, which results from the anatomy of neurons and neuronal circuits, is referred to as shared-input correlation (Kriener et al., 2008; Tetzlaff et al., 2008). A second type of correlation in synaptic input currents arises from correlations in the presynaptic spiking activity (Renart et al., 2010; Tetzlaff et al., 2012; Helias et al., 2014). These dynamical correlations occur during stationary network states, or can be generated by different types of nonstationary activities, such as global oscillations in the population activity (Brunel & Hakim, 1999; Brunel, 2000) or traveling waves of activity propagating across the neuronal tissue (Sato et al., 2012; Takahashi et al., 2015; Roxin et al., 2005; Senk et al., 2020).

This study addresses the problem of sequential decision making in the face of ambiguity and the role of coherent noise in shaping decision strategies. We investigate how the spiking TM model of Bouhadjar et al. (2021) recalls sequences in response to ambiguous cues in the presence of coherent noise, to what extent noise averaging can be overcome by increasing the noise amplitude, and how different recall strategies can be achieved by adjusting the noise characteristics. We further explore whether shared synaptic input and random stimulus locking to spatiotemporal oscillations can serve as appropriate, natural sources of coherent noise. In Methods, we provide a detailed description of the task and the network model. For a first read, the main findings of the study can be understood without Methods.

Methods

In the following, we provide an overview of the task and the training protocol, the network model, and the task performance analysis. A detailed description of the model and parameter values is provided in Tables 1 and 2.

Learning protocol and task

During learning, the network is continuously exposed to repeated presentations of an ensemble of S sequences $s_i = \{\zeta_{i1}, \zeta_{i2}, \ldots, \zeta_{iC_i}\}$ ($C_i \in \mathbb{N}^+$, $i \in [1, \ldots, S]$) of ordered discrete items ζ_{ij}, The order of the sequence elements within a given sequence represents the temporal order of the item occurrence. Each sequence s_i is presented with a relative frequency p_i in a given set of training data, where $\sum_{i=1}^{S} p_i = 1$. After successful learning, the presentation of some sequence element leads to a context dependent prediction of the subsequent stimulus. In case the prediction is wrong the network generates a mismatch signal (see Bouhadjar et al., 2021). The network can also be configured into a replay mode where it autonomously replays learned sequences in response to a cue signal. We design the sequences such that they all start with the same two elements $\zeta_1 = \zeta_{11} = \zeta_{21} \ldots \zeta_{i1}$ and $\zeta_2 = \zeta_{12} = \zeta_{22} \ldots \zeta_{i2}$ ($i \in [1, \ldots, S]$). As a consequence, choosing the cue to be the first sequence element (ζ_1) results in an ambiguity. Here, we investigate the replay frequency of a given sequence s_i as a function of its training frequency p_i and study whether the network can choose between different replay strategies.

Network model

Network structure The network consists of a population $\mathcal{E}$ of N_E excitatory ("E") and a single inhibitory ("I") neuron. The neurons in $\mathcal{E}$ are randomly and recurrently connected, such that each neuron in $\mathcal{E}$ receives K_EE excitatory inputs from other neurons in $\mathcal{E}$. Excitatory neurons are recurrently connected to the single inhibitory neuron. The excitatory population $\mathcal{E}$ is subdivided into M non-overlapping subpopulations $\mathcal{M}_1, \ldots, \mathcal{M}_M$, each of them containing neurons with identical stimulus preference ("receptive field"). Each subpopulation $\mathcal{M}_k$ thereby represents a specific element within a sequence.

External inputs during learning The network is driven by an ensemble $\mathcal{X} = \{x_1, \ldots, x_{N_\mathsf{stim}}\}$ of M external inputs. Each of these external inputs x_k represents a specific sequence element ("A", "B", ...), and feeds all neurons in the subpopulation $\mathcal{M}_k$ with the corresponding stimulus preference. The occurrence of a specific sequence element $\zeta_{i,j}$ at time $t_{i,j}$ is modeled by a single spike $x_k(t) = \delta(t - t_{i,j})$ generated by the corresponding external source x_k. Subsequent sequence elements $\zeta_{i,j}$ and $\zeta_{i,j+1}$ within a sequence s_i are presented with an inter-stimulus interval $\Delta T = t_{i,j+1} - t_{i,j}$. Subsequent sequences s_i and s_{i+1} are separated in time by an inter-sequence time interval $\Delta T_\mathsf{seq} = t_{i+1,1} - t_{i,C_i}$.

External inputs during replay After learning the set of sequences S, we present a cue signal encoding for first sequence elements ζ_1 at time t^j for a number of trials N_t, where j denotes the trial id ($j \in [1, \ldots, N_\mathsf{t}]$). Subsequent cues are separated by an inter-trial interval $\Delta T_\mathsf{cue,j} = t^j - t^{j+1}$. In section A spiking neural network recalls sequences in response to ambiguous cues, $\Delta T_\mathsf{cue,j}$ is constant and in section Random stimulus locking to spatiotemporal oscillations as natural form of noise, $\Delta T_\mathsf{cue,j}$ is randomly and uniformly distributed between u_min ms and u_max.

During the replay, excitatory neurons are additionally driven by a background input implemented either in the form of asynchronous irregular synaptic bombardment (see A spiking neural network recalls sequences in response to ambiguous cues) or oscillatory inputs (see Random stimulus locking to spatiotemporal oscillations as natural form of noise). The first is realized using ensembles of excitatory and inhibitory spike sources $\mathcal{Q}_k$ and $\mathcal{V}_k$ ($k \in [1, \ldots, M]$), each composed of n elements. Each source is an independent realization of a Poisson point process with a rate ν. Excitatory neurons in the same subpopulation $\mathcal{M}_k$ receive K_EQ inputs with weight J_EQ from the ensemble $\mathcal{Q}_k$ and K_EV inputs with weights $-J_\mathsf{EV}$ from the ensemble $\mathcal{V}_k$. Spikes from $\mathcal{Q}_k$ and $\mathcal{V}_k$ give rise to a jump in the synaptic current of the postsynaptic cell followed by an exponential decay with a time constant τ_EQ and τ_EV, respectively. The time average input current of a neuron i is

$$\mu_i = 0 \tag{1}$$

and the variance across time

$$\sigma_i^2 = \frac{\tau_\mathsf{B}}{2} \sum_{k=0}^{2K} J_{ik}^2 \nu, \tag{2}$$

where $J = J_\mathsf{EQ} = -J_\mathsf{EV}$, $\tau_\mathsf{B} = \tau_\mathsf{EQ} = \tau_\mathsf{EV}$, and $K = K_\mathsf{EQ} = K_\mathsf{EV}$. Given that the populations of background sources are of a finite size, there is a probability that two neurons in the same subpopulation to pick a certain number of identical sources, this gives rise to the so called shared input correlation. The correlation coefficient in the input current is governed by

$$c = \frac{K}{n}. \tag{3}$$

With this relationship, we can now vary the correlation coefficient by fixing K and varying n. For the special case where c is supposed to be zero, we assume that each neuron has its own set of independent Poissonian sources. The second type of background input is implemented using an ensemble $\mathcal{G}$ of M sinusoidal current generators g_k, each with a frequency f, amplitude a, and a phase ϕ_k ($k \in [1, \ldots, M]$). Excitatory neurons in the same subpopulation M_k receive oscillatory inputs from the same source g_k.

Neuron and synapse model For all types of neurons, the temporal evolution of the membrane potential is given by the leaky integrate-and-fire model (6). The total synaptic input current of excitatory neurons is composed of currents in distal dendritic branches, inhibitory currents, and currents from external sources. The inhibitory neuron receives only inputs from excitatory neurons. Individual spikes arriving at dendritic branches evoke alpha-shaped postsynaptic currents, see (8). The dendritic current includes an additional nonlinearity describing the generation of dendritic action potentials (dAPs): if the dendritic current I_{ED} exceeds a threshold θ_{dAP}, it is instantly set to the dAP plateau current I_{dAP}, and clamped to this value for a period of duration τ_{dAP}, see (12). This plateau current leads to a long lasting depolarization of the soma. Inhibitory inputs to excitatory neurons as well as excitatory inputs to inhibitory neurons trigger exponential postsynaptic currents, see (9–10). The weights J_{IE} of excitatory synapses on inhibitory neurons are chosen such that the collective firing of a subset of ρ excitatory neurons in the corresponding subpopulation causes the inhibitory neuron to fire. The weights J_{EI} of inhibitory synapses on excitatory neurons are strong such that each inhibitory spike prevents all excitatory neurons in the network that have not generated a spike yet from firing. External inputs are composed of currents resulting from the presentation of the sequence elements or currents from background inputs (see Inputs in Table 1). All synaptic time constants, delays, and weights are connection-type specific.

Plasticity Only excitatory to excitatory (EE) synapses are plastic. All other connections are static. The dynamics of the EE synaptic weights J_{ij} evolve according to a combination of an additive spike-timing-dependent plasticity (STDP) rule (Morrison et al., 2008) and a homeostatic component (Abbott & Nelson, 2000; Tetzlaff et al., 2011). During the replay mode, the plasticity is disabled and the EE weights are kept constant (see Table 1 for details about the plasticity).

Network realizations and initial conditions. For every network realization, the connectivity and the initial weights are drawn randomly and independently. All other parameters are identical for different network realizations. The initial values of all state variables are given in Table 1 and Table 2.

Simulation details The network simulations are performed in the neural simulator NEST (Gewaltig & Diesmann, 2007) under version 3.0 (Hahne et al., 2021). The differential equations and state transitions defining the excitatory neuron dynamics are expressed in the domain specific language NESTML (Plotnikov et al., 2016; Nagendra Babu et al., 2021) which generates the required C++ code for the dynamic loading into NEST. Network states are synchronously updated using exact integration of the system dynamics on a discrete-time grid with step size Δt (Rotter & Diesmann, 1999). The full source code for the implementation with a list of other software requirements is available at Zenodo: `https://doi.org/10.5281/zenodo.6378376`.

Task performance measures

Consider the set $\mathcal{S} = \{s_1, s_2, \ldots, s_S\}$ of S sequences learned by the network. Let

$$\mathcal{P} = \{\emptyset, \{s_1\}, \{s_2\}, \ldots, \{s_1, s_2\}, \{s_1, s_3\}, \ldots, \mathcal{S}\}$$

denote the power set of $\mathcal{S}$, i.e., the set of all subsets of $\mathcal{S}$, including the empty set and $\mathcal{S}$ itself. We define the relative replay frequency $f_{\mathcal{P}_k}$ of each subset $\mathcal{P}_k \in \mathcal{P}$ of sequences as the normalized number of exclusive replays of this subset $\mathcal{P}_k$, such that

$$\sum_{\mathcal{P}_k} f_{\mathcal{P}_k} = 1. \tag{4}$$

For two sequences s_1 and s_2, for example, we monitor the four different replay frequencies $f_\emptyset$ (no sequence is replayed), $f_{\{s_1\}}$ (only s_1 is replayed), $f_{\{s_2\}}$ (only s_2 is replayed), and $f_{\{s_1,s_2\}}$ (both s_1 and s_2 are replayed). In this work, we refer to $f_\emptyset$ as the "failure rate". Simultaneous replay of both sequences ($f_{\{s_1,s_2\}}$) refers to cases where the network fails at coming to a unique decision.

Summary	
Populations	excitatory neurons ($\mathcal{E}$), inhibitory neurons ($\mathcal{I}$), external spike sources ($\mathcal{X}$), background inputs in the form Poissonian sources ($\mathcal{Q}_k$ and $\mathcal{V}_k$) or sinusoidal current generators ($\mathcal{G}$). $\mathcal{E}$ composed of M disjoint subpopulations $\mathcal{M}_k$ and ($k = 1, \ldots, M$)
Connectivity	• sparse random connectivity between excitatory neurons (plastic) • local recurrent connectivity between excitatory and inhibitory neurons (static)
Neuron model	• excitatory neurons: leaky integrate-and-fire (LIF) with nonlinear input integration (dendritic action potentials) • inhibitory neurons: leaky integrate-and-fire (LIF)
Synapse model	exponential or alpha-shaped postsynaptic currents (PSCs)
Plasticity	homeostatic spike-timing dependent plasticity in excitatory-to-excitatory connections

Populations		
Name	**Elements**	**Size**
$\mathcal{E} = \cup_{i=k}^{M} \mathcal{M}_k$	excitatory (E) neurons	N_{E}
$\mathcal{I}$	inhibitory (I) neurons	N_{I}
$\mathcal{M}_k$	excitatory neurons in subpopulation k, $\mathcal{M}_k \cap \mathcal{M}_l = \emptyset$ ($\forall k \neq l \in [1, M]$)	n_{E}
$\mathcal{Q}_k$	excitatory Poisson generators, $\mathcal{Q}_k \cap \mathcal{Q}_l = \emptyset$ ($\forall k \neq l \in [1, M]$)	n
$\mathcal{V}_k$	inhibitory Poisson generators, $\mathcal{V}_k \cap \mathcal{V}_l = \emptyset$ ($\forall k \neq l \in [1, M]$)	n
$\mathcal{X} = \{x_1, \ldots, x_M\}$	external spike sources	M
$\mathcal{G} = \{g_1, \ldots, g_M\}$	sinusoidal current generators	M

Connectivity		
Source population	**Target population**	**Pattern**
$\mathcal{E}$	$\mathcal{E}$	random; fixed in-degrees $K_i = K_{\mathsf{EE}}$, delays $d_{ij} = d_{\mathsf{EE}}$, and synaptic time constants $\tau_{ij} = \tau_{\mathsf{EE}}$, plastic synaptic weights J_{ij} ($\forall i \in \mathcal{E}, \forall j \in \mathcal{E}$; "EE connections")
$\mathcal{E}$	$\mathcal{I}$	all-to-all; fixed delays $d_{ij} = d_{\mathsf{IE}}$, synaptic time constants $\tau_{ij} = \tau_{\mathsf{IE}}$, and weights $J_{ij} = J_{\mathsf{IE}}$ ($\forall i \in \mathcal{I}, \forall j \in \mathcal{E}$; "IE connections")
$\mathcal{I}$	$\mathcal{E}$	all-to-all; fixed delays $d_{ij} = d_{\mathsf{EI}}$, synaptic time constants $\tau_{ij} = \tau_{\mathsf{EI}}$, and weights $J_{ij} = J_{\mathsf{EI}}$ ($\forall i \in \mathcal{E}, \forall j \in \mathcal{I}$; "EI connections")
$\mathcal{I}$	$\mathcal{I}$	none ("II connections")
$\mathcal{Q}_k$	$\mathcal{M}_k$	random; fixed in-degrees K_i=K_{EQ}, delays $d_{ij} = d_{\mathsf{EQ}}$, synaptic time constants $\tau_{ij} = \tau_{\mathsf{EQ}}$, and weights $J_{ij} \in \{0, J_{\mathsf{EQ}}\}$ ($\forall i \in \mathcal{M}_k, j \in \mathcal{Q}_k, \forall k \in [1, M]$; "EQ connections")
$\mathcal{V}_k$	$\mathcal{M}_k$	random; fixed in-degrees K_i=K_{EV}, delays $d_{ij} = d_{\mathsf{EV}}$, synaptic time constants $\tau_{ij} = \tau_{\mathsf{EV}}$, and weights $J_{ij} \in \{0, J_{\mathsf{EV}}\}$ ($\forall i \in \mathcal{M}_k, j \in \mathcal{V}_k, \forall k \in [1, M]$; "EV connections")
$\mathcal{X}_k = x_k$	$\mathcal{M}_k$	one-to-all; fixed delays $d_{ij} = d_{\mathsf{EX}}$, synaptic time constants $\tau_{ij} = \tau_{\mathsf{EX}}$, and weights $J_{ij} = J_{\mathsf{EX}}$ ($\forall i \in \mathcal{M}_k, j \in \mathcal{X}_k, \forall k \in [1, M]$; "EX connections")
$\mathcal{G}_k = g_k$	$\mathcal{M}_k$	one-to-all; fixed synaptic weights $J_{ij} = J_{\mathsf{EG}}$ ($\forall i \in \mathcal{M}_k, j \in \mathcal{G}_k, \forall k \in [1, M]$; "EG connections")
no self-connections ("autapses"), no multiple connections ("multapses")		
all unmentioned connections $\mathcal{I} \to \mathcal{I}$, $\mathcal{V}_k \to \mathcal{V}_k$, $\mathcal{Q}_k \to \mathcal{Q}_k \ldots \mathcal{X}_k \to \mathcal{M}_l$ ($\forall k \neq l$) are absent		

Table 1: Description of the network model (continued on next page). Parameter values are given in Tab. 2.

Neuron and synapse

Neuron

Type	leaky integrate-and-fire (LIF) dynamics
Description	dynamics of membrane potential $V_i(t)$ and spiking activity $s_i(t)$ of neuron i:

- emission of the kth spike of neuron i at time t_i^k if

$$V_i(t_i^k) \geq \theta_i \tag{5}$$

with somatic spike threshold θ_i

- spike train: $s_i(t) = \sum_k \delta(t - t_i^k)$
- reset and refractoriness:

$$V_i(t) = V_\mathsf{r} \quad \forall k,\ \forall t \in \left(t_i^k,\ t_i^k + \tau_{\mathsf{ref},i}\right]$$

with refractory time $\tau_{\mathsf{ref},i}$ and reset potential V_r

- subthreshold dynamics:

$$\tau_{\mathsf{m},i}\dot{V}_i(t) = -V_i(t) + R_{\mathsf{m},i}I_i(t) \tag{6}$$

with membrane resistance $R_{\mathsf{m},i} = \dfrac{\tau_{\mathsf{m},i}}{C_{\mathsf{m},i}}$, membrane time constant $\tau_{\mathsf{m},i}$, and total synaptic input current $I_i(t)$ (see Synapse)

- excitatory neurons: $\tau_{\mathsf{m},i} = \tau_{\mathsf{m,E}}$, $C_{\mathsf{m},i} = C_\mathsf{m}$, $\theta_i = \theta_\mathsf{E}$, $\tau_{\mathsf{ref},i} = \tau_{\mathsf{ref,E}}$ $(\forall i \in \mathcal{E})$
- inhibitory neurons: $\tau_{\mathsf{m},i} = \tau_{\mathsf{m},I}$, $C_{\mathsf{m},i} = C_\mathsf{m}$, $\theta_i = \theta_\mathsf{I}$, $\tau_{\mathsf{ref},i} = \tau_{\mathsf{ref,I}}$ $(\forall i \in \mathcal{I})$

Synapse

Type	continuous, exponential, or alpha-shaped postsynaptic currents (PSCs)
Description	

- total synaptic input current

$$\begin{aligned} &\text{excitatory neurons:} \quad I_i(t) = I_{\mathsf{ED},i}(t) + I_{\mathsf{EX},i}(t) + I_{\mathsf{EI},i}(t),\ \forall i \in \mathcal{E} \\ &\text{inhibitory neurons:} \quad I_i(t) = I_{\mathsf{IE},i}(t),\ \forall i \in \mathcal{I} \end{aligned} \tag{7}$$

with dendritic, inhibitory, excitatory, and external input currents $I_{\mathsf{ED},i}(t)$, $I_{\mathsf{EI},i}(t)$, $I_{\mathsf{IE},i}(t)$, $I_{\mathsf{EX},i}(t)$ evolving according to

$$I_{\mathsf{ED},i}(t) = \sum_{j\in\mathcal{E}} (\alpha_{ij} * s_j)(t - d_{ij}) \tag{8}$$

$$\text{with } \alpha_{ij}(t) = J_{ij}\frac{e}{\tau_\mathsf{ED}} t e^{-t/\tau_\mathsf{ED}}\Theta(t) \text{ and } \Theta(t) = \begin{cases} 1 & t \geq 0 \\ 0 & \text{else} \end{cases}$$

$$\tau_\mathsf{EI}\dot{I}_{\mathsf{EI},i} = -I_{\mathsf{EI},i}(t) + \sum_{j\in\mathcal{I}} J_{ij}s_j(t - d_{ij}) \tag{9}$$

$$\tau_\mathsf{IE}\dot{I}_{\mathsf{IE},i} = -I_{\mathsf{IE},i}(t) + \sum_{j\in\mathcal{E}} J_{ij}s_j(t - d_{ij}) \tag{10}$$

$$I_{\mathsf{EX},i}(t) = I_{\mathsf{S},i}(t) + I_{\mathsf{B},i}(t) \tag{11}$$

where $I_{\mathsf{S},i}(t)$ and $I_{\mathsf{B},i}(t)$ are the stimulus and the background input, respectively (see input).

- suprathreshold dynamics of dendritic currents (dAP generation):
  - emission of kth dAP of neuron i at time $t_{\mathsf{dAP},i}^k$ if $I_{\mathsf{ED},i}(t_{\mathsf{dAP},i}^k) \geq \theta_\mathsf{dAP}$
  - dAP current plateau:

$$I_{\mathsf{ED},i}(t) = I_\mathsf{dAP} \quad \forall k,\ \forall t \in \left(t_{\mathsf{dAP},i}^k, t_{\mathsf{dAP},i}^k + \tau_\mathsf{dAP}\right) \tag{12}$$

with dAP current plateau amplitude I_dAP, dAP current duration τ_dAP, and dAP activation threshold θ_dAP

  - reset: $I_{\mathsf{ED},i}(t_{\mathsf{dAP},i}^k + \tau_\mathsf{dAP}) = 0$ $(\forall k)$
  - reset and refractoriness in response to emission of lth somatic spike of neuron i at time t_i^l:

$$I_{\mathsf{ED},i}(t) = 0 \quad \forall l,\ \forall t \in \left(t_i^l,\ t_i^l + \tau_{\mathsf{ref},i}\right) \tag{13}$$

Table 1: Description of the network model (continued on next page). Parameter values are given in Tab. 2.

Plasticity	
Type	spike-timing dependent plasticity and dAP-rate homeostasis
EE synapses	• dynamics of synaptic weight $J_{ij}(t)$ (EE connections): $$J_{\text{max}}^{-1}\frac{dJ_{ij}}{dt} = \lambda_{+}\sum_{\{t_i^*\}'} x_j(t)\delta(t-[t_i^*+d_{\text{EE}}])I(t_i^*,\Delta t_{\text{min}},\Delta t_{\text{max}}) - \lambda_{-}y_iJ_{ij}\sum_{\{t_j^*\}}\delta(t-t_j^*) + \lambda_{\text{h}}\sum_{\{t_i^*\}'}\big(z^*-z_i(t)\big)\delta(t-t_i^*)I(t_i^*,\Delta t_{\text{min}},\Delta t_{\text{max}}).$$ with – list of presynaptic spike times $\{t_j^*\}$, – list of postsynaptic spike times $\{t_i^*\}' = \{t_i^*|\forall t_j^*:\ t_i^*-t_j^*+d_{\text{EE}}\geq\Delta t_{\text{min}}\}$ – indicator function $$I(t_i^*,\Delta t_{\text{min}},\Delta t_{\text{max}}) = \sum_{\{t_j^*\}} R(t_i^*-t_j^*+d_{\text{EE}}) \quad \text{with} \quad R(\tau)=\begin{cases}1 & \Delta t_{\text{min}}<\tau<\Delta t_{\text{max}}\\ 0 & \text{else,}\end{cases} \tag{14}$$ – maximum weight J_{max}, potentiation and depression rates λ_+, λ_-, homeostasis rate λ_{h}, delay d_{EE}, depression decrement y_i, minimum Δt_{min} and maximum Δt_{max} time lags between pairs of pre- and postsynaptic spikes at which synapses are potentiated, – spike trace of postsynaptic neuron i, evolving according to $$\frac{dx_j}{dt} = -\tau_+^{-1}x_j(t)+\sum_{t_j^*}\delta(t-t_j^*)$$ with presynaptic spike times t_j^* and potentiation time constant τ_+, – dAP trace of postsynaptic neuron i, evolving according to $$\frac{dz_i}{dt} = -\tau_{\text{h}}^{-1}z_i(t)+\sum_k\delta(t-t_{\text{dAP},i}^k)$$ with onset time $t_{\text{dAP},i}^k$ of the kth dAP, homeostasis time constant τ_{h}, and – target dAP activity z^*
all other synapses	non-plastic

Table 1: Description of the network model (continued on next page). Parameter values are given in Tab. 2.

Input

- prediction mode
  - stimulus
    - repetitive stimulation of the network using the same set $\mathcal{S} = \{s_1, \ldots, s_S\}$ of sequences $s_i = \{\zeta_{i,1}, \zeta_{i,2}, \ldots, \zeta_{i,C_i}\}$ of ordered discrete items $\zeta_{i,j}$ with number of sequences S and length C_i of ith sequence
    - presentation of sequence element $\zeta_{i,j}$ at time $t_{i,j}$ modeled by a single spike $x_k(t) = \delta(t - t_{i,j})$ generated by the corresponding external source x_k
    - generated current as a response to the presentation of the sequence elements:

$$\tau_{\mathsf{S}} \dot{I}_{\mathsf{S},i} = -I_{\mathsf{S},i}(t) + \sum_{j \in \mathcal{X}} J_{i,j} x_j(t - d_{ij}) \tag{15}$$

    - inter-stimulus interval $\Delta T = t_{i,j+1} - t_{i,j}$ between subsequent sequence elements $\zeta_{i,j}$ and $\zeta_{i,j+1}$ within a sequence s_i
    - inter-sequence time interval $\Delta T_{\mathsf{seq}} = t_{i+1,1} - t_{i,C_i}$ between subsequent sequences s_i and s_{i+1}
    - example sequence sets:
      - sequence set I: $\mathcal{S}$={{A,F,B,D} and {A,F,C,E}}
- replay mode
  - stimulus
    - presentation of a cue encoding for first sequence elements ζ_1 at time t^j, where j denotes the trial number ($j \in [1, \ldots, N_{\mathrm{t}}]$).
    - inter-trial interval $\Delta T_{\mathsf{cue}} = t^{j+1} - t^j$
  - background input
    - in the form of correlated Poissonian inputs

$$\tau_{\mathsf{B}} \dot{I}_{\mathsf{B},i}(t) = -I_{\mathsf{B},i}(t) + \sum_{j \in \mathcal{Q}} J_{i,j} s_j(t - d) + \sum_{j \in \mathcal{V}} J_{i,j} s_j(t - d) \tag{16}$$

      with Poissonian spike trains $s_j(t)$ of rate ν, synaptic weight $J_{i,j} \in \{0, J\}$ where $J = J_{\mathsf{EQ}} = -J_{\mathsf{EV}}$, synaptic time constant $\tau_{\mathsf{B}} = \tau_{\mathsf{EQ}} = \tau_{\mathsf{EV}}$, and delay $d = d_{\mathsf{EQ}} = d_{\mathsf{EV}}$
      - the variance of $I_{\mathsf{B},i}(t)$ across time:

$$\sigma^2 = \mathsf{Var}(I_{\mathsf{B},i}(t)) = J^2 K \nu \tau_{\mathsf{B}}, \tag{17}$$

        where $K = K_{\mathsf{EQ}} = K_{\mathsf{EV}}$ is the number of either excitatory or inhibitory Poissonian input per excitatory neuron
      - the correlation coefficient of $I_{\mathsf{B},i}(t)$ and $I_{\mathsf{B},j}(t)$ across time:

$$c = \frac{\mathrm{Cov}(I_{\mathsf{B},i}(t), I_{\mathsf{B},j})(t)}{\sigma^2} = \begin{cases} 0 & \forall i \in \mathcal{M}_k, \forall j \in \mathcal{M}_l \; (\forall k \neq l) \\ \frac{K}{n} & \forall i \in \mathcal{M}_k, \forall j \in \mathcal{M}_l \; (\forall k = l), \end{cases} \tag{18}$$

        where n is the number of either excitatory or inhibitory Poissonian sources (see Connectivity)
    - or oscillatory current

$$I_{\mathsf{B},i}(t) = J_{\mathsf{EG}} \cdot a \cdot \sin(2\pi f t + \phi_i) \tag{19}$$

      with amplitude a, frequency f, and population specific phase $\phi_i = \phi_k$ ($\forall i \in \mathcal{M}_k$)

Output

- somatic spike times $\{t_i^k | \forall i \in \mathcal{E}, k = 1, 2, \ldots\}$
- dendritic currents $I_{\mathsf{ED},i}(t)$ ($\forall i \in \mathcal{E}$)

Table 1: Description of the network model (continued on next page). Parameter values are given in Tab. 2.

Initial conditions and network realizations

- membrane potentials: $V_i(0) = V_\mathsf{r}$ ($\forall i \in \mathcal{E} \cup \mathcal{I}$)
- dendritic currents: $I_{\mathsf{ED},i}(0) = 0$ ($\forall i \in \mathcal{E}$)
- external currents: $I_{S,i}(0) = 0$ and $I_{B,i}(0) = 0$ ($\forall i \in \mathcal{E}$)
- inhibitory currents: $I_{\mathsf{EI},i}(0) = 0$ ($\forall i \in \mathcal{E}$)
- excitatory currents: $I_{\mathsf{IE},i}(0) = 0$ ($\forall i \in \mathcal{I}$)
- synaptic weights: $J_{ij}(0) \sim \mathcal{U}(J_{0,\mathsf{min}}, J_{0,\mathsf{max}})$ (uniform distribution; $\forall i, j \in \mathcal{E}$)
- spike traces: $x_i(0) = 0$ ($\forall i \in \mathcal{E}$)
- dAP traces: $z_i(0) = 0$ ($\forall i \in \mathcal{E}$)
- initial weights randomly and independently drawn for each network realization

Simulation details

- network simulations performed in NEST (Gewaltig & Diesmann, 2007) version 3.0 (Hahne et al., 2021)
- definition of excitatory neuron model using NESTML (Plotnikov et al., 2016; Nagendra Babu et al., 2021)
- synchronous update using exact integration of system dynamics on discrete-time grid with step size Δt (Rotter & Diesmann, 1999)
- source code underlying this study: `https://doi.org/10.5281/zenodo.6378376`

Table 1: Description of the network model. Parameter values are given in Tab. 2.

Model and simulation parameters

Name	Value	Description
Network		
N_E	1200	total number of excitatory neurons
N_I	1	total number of inhibitory neurons
M	8	number of excitatory subpopulations (= number of external spike sources)
n_E	$N_\mathrm{E}/M = 150$	number of excitatory neurons per subpopulation
ρ	20	(target) number of active neurons per subpopulation after learning = minimal number of coincident excitatory inputs required to trigger a spike in postsynaptic inhibitory neurons
n	$\{100, \ldots, 1000\}$	number of excitatory or inhibitory Poissonian sources
Connectivity		
K_EE	240	number of excitatory inputs per excitatory neuron (EE in-degree)
p	$K_\mathrm{EE}/N_\mathrm{E} = 0.2$	connection probability
K_EI	$N_\mathrm{I} = 1$	number of inhibitory inputs per excitatory neuron (EI in-degree)
K_IE	$N_\mathrm{E} = 1200$	number of excitatory inputs per inhibitory neuron (IE in-degree)
K_II	0	number of inhibitory inputs per inhibitory neuron (II in-degree)
K_EQ	100	number of excitatory Poissonian inputs per excitatory neuron (EQ)
K_EV	$K_\mathrm{EQ} = 100$	number of inhibitory Poissonian inputs per excitatory neuron (EV)
Excitatory neurons		
$\tau_\mathrm{m,E}$	10 ms	membrane time constant
$\tau_\mathrm{ref,E}$	20 ms	absolute refractory period
C_m	250 pF	membrane capacity
V_r	0 mV	reset potential
θ_E	20 mV (training), 7 mV (replay)	somatic spike threshold
I_dAP	200 pA	dAP current plateau amplitude
τ_dAP	60 ms	dAP duration
θ_dAP	59 pA	dAP threshold
Inhibitory neurons		
$\tau_\mathrm{m,I}$	5 ms	membrane time constant
$\tau_\mathrm{ref,I}$	2 ms	absolute refractory period
C_m	250 pF	membrane capacity
V_r	0.0 mV	reset potential
θ_I	15 mV	spike threshold

Table 2: Model and simulation parameters (continued on next page). Parameters derived from other parameters are marked in gray.

Name	Value	Description
	Synapse	
J_{IE}	$581.19\,\mathrm{pA}$ (training), $726.49\,\mathrm{pA}$ (replay)	weight of IE connections (EPSC amplitude)
J_{EI}	$-12915.49\,\mathrm{pA}$ (training), $-2260.21\,\mathrm{pA}$ (replay)	weight of EI connections (IPSC amplitude)
J_{EX}	$4112.20\,\mathrm{pA}$ (training), $1962.64\,\mathrm{pA}$ (replay)	weight of EX connections (EPSC amplitude)
J_{EQ}	$\sigma/\sqrt{K_{\mathsf{EQ}}\nu\tau_{\mathsf{EQ}}}$	weight of EQ connections (EPSC amplitude)
J_{EV}	$-J_{\mathsf{EQ}}$	weight of EV connections (EPSC amplitude)
J_{EG}	$1\,\mathrm{pA}$	weight of EG connections (EPSC amplitude)
τ_{EE}	$15\,\mathrm{ms}$	synaptic time constant of EE connections
τ_{EI}	$1\,\mathrm{ms}$	synaptic time constant of EI connections
τ_{EX}	$2\,\mathrm{ms}$	synaptic time constant of EX connection
τ_{IE}	$0.5\,\mathrm{ms}$	synaptic time constant of IE connections
τ_{EQ}	$2\,\mathrm{ms}$	synaptic time constant of EQ connections
τ_{EV}	$\tau_{\mathsf{EQ}} = 2\,\mathrm{ms}$	synaptic time constant of EV connections
d_{EE}	$2\,\mathrm{ms}$	delay of EE connections (dendritic)
d_{IE}	$0.1\,\mathrm{ms}$	delay of IE connections
d_{EI}	$0.1\,\mathrm{ms}$	delay of EI connections
d_{EX}	$0.1\,\mathrm{ms}$	delay of EX connections
d_{EQ}	$0.1\,\mathrm{ms}$	delay of EQ connections
d_{EV}	$d_{\mathsf{EQ}} = 0.1\,\mathrm{ms}$	delay of EV connections
	Plasticity	
λ_+	0.001	potentiation rate
λ_-	0.00008	depression rate
λ_{h}	0.0008	homeostasis rate
J_{min}	$0\,\mathrm{pA}$	minimum weight
J_{max}	$35\,\mathrm{pA}$	maximum weight
$J_{0,\mathsf{min}}$	$0\,\mathrm{pA}$	minimal initial weight
$J_{0,\mathsf{max}}$	$1\,\mathrm{pA}$	maximal initial weight
τ_+	$20\,\mathrm{ms}$	potentiation time constant
z^*	10.35	target dAP activity
τ_{h}	$2200\,\mathrm{ms}$	homeostasis time constant
y_i	1	depression decrement
	Input	
S	$\{2, 3\}$	number of sequences per set
C	4	number of characters per sequence
A	$\{6, 8\}$	alphabet length
ΔT	$40\,\mathrm{ms}$	inter-stimulus interval
ΔT_{seq}	$100\,\mathrm{ms}$	inter-sequence interval
ΔT_{cue}	$140\,\mathrm{ms}$ or $\sim \mathcal{U}(u_{\mathsf{min}}\, u_{\mathsf{max}})$	inter-cue interval
u_{min}	$140\,\mathrm{ms}$	minimal inter-cue interval
u_{max}	$280\,\mathrm{ms}$	maximal inter-cue interval
σ	$\{0, 15, 60\}\,\mathrm{pA}$	noise amplitude resulting from the Poissonian background inputs
c	$\frac{n}{K} = \{0, 1\}$	noise correlation
ν	$1000\,\mathrm{s}^{-1}$	rate of Poissonian background inputs
a	$\{0, 8, 16\}$	amplitude of the sinusoidal current generators
f	$\{10, 30, 70\}\,\mathrm{Hz}$	frequency of the sinusoidal current generators
	Simulation	
Δt	$0.1\,\mathrm{ms}$	time resolution
K	300	number of training episodes
N_{t}	150	number of trials

Table 2: Model and simulation parameters. Parameters derived from other parameters are marked in gray. Curly brackets depict a set of values corresponding to different experiments.

Results

A spiking neural network recalls sequences in response to ambiguous cues

In this section, we provide a brief overview of the model and the task, illustrate how the network learns overlapping sequences occurring with different frequencies during the training, and show how these occurrence frequencies are encoded in the network. We then study the network responses to ambiguous cues and the influence of the occurrence frequencies on the recall behavior in the absence or presence of noise.

Similar to (Bouhadjar et al., 2021), the model consists of a randomly and sparsely connected network of N_{E} excitatory neurons (population $\mathcal{E}$) and a single inhibitory neuron (Fig. 1A). Each excitatory neuron receives K_{EE} excitatory inputs from other randomly chosen neurons in $\mathcal{E}$. Excitatory neurons are subdivided into M subpopulations, each containing neurons with identical stimulus preference: in the absence of any additional connections, all neurons in a given subpopulation fire a spike upon presentation of a specific sequence element. The inhibitory neuron is recurrently connected to the excitatory neurons. In contrast to (Bouhadjar et al., 2021) where each excitatory subpopulation is equipped with its own inhibitory neuron, we here use a single inhibitory neuron to implement competition between the subpopulations of excitatory neurons. The network is driven by external inputs, each representing a specific sequence element ("A", "B", ...), and feeds all neurons in the subpopulation $\mathcal{M}_k$ with the corresponding stimulus preference. Neurons are modeled as point neurons with the membrane potential evolving according to the leaky integrate-and-fire dynamics. The total synaptic input current of excitatory neurons is composed of currents in distal dendritic branches, inhibitory currents, and currents from external sources. Inhibitory neurons receive only inputs from excitatory neurons. Individual spikes from other excitatory neurons arriving at dendritic branches evoke alpha-shaped postsynaptic currents. The dynamics of dendritic currents include a nonlinearity describing the generation of dendritic action potentials (dAPs). External and inhibitory inputs to excitatory neurons as well as excitatory inputs to inhibitory neurons trigger exponential postsynaptic currents. Synapses between excitatory neurons are plastic and subject to spike-timing-dependent plasticity and homeostatic control. Details on the network model are given in Methods.

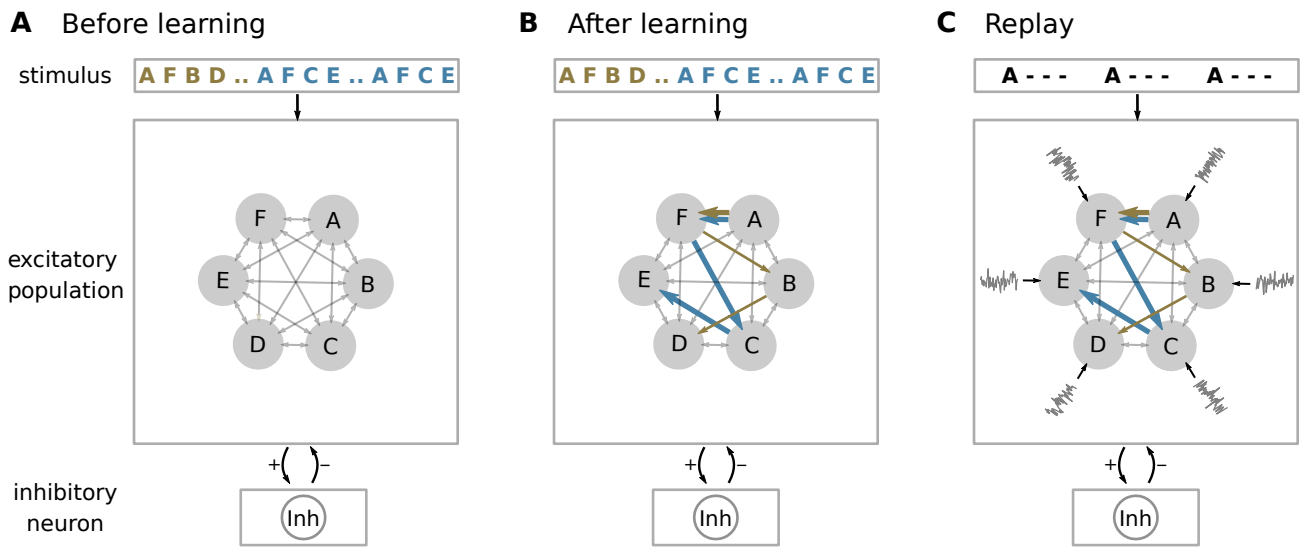

Figure 1: Network structure. A) The architecture constitutes a recurrent network of subpopulations of excitatory neurons (filled gray circles) and a single inhibitory neuron (Inh). Each excitatory subpopulation contains neurons with identical stimulus preferences. Excitatory neurons are stimulated by external sources providing sequence-element specific inputs "A","F", "B", etc. Connections between and within the excitatory subpopulations are random and sparse. The inhibitory neuron is recurrently connected to all excitatory neurons. In the depicted example, the network is repetitively presented with two sequences {A,F,C,E} (brown) and {A,F,B,D} (blue) during learning. The sequence {A,F,C,E} occurs twice as often as {A,F,B,D}. **B)** After learning, the networks forms sequence specific subnetworks (blue and brown arrows representing {A,F,B,D} and {A,F,C,E}, respectively). The connections between subpopulations representing the sequence shown more often are stronger (thick arrows). **C)** During the replay mode, the network is presented with a cue stimulus representing the first sequence element "A". In addition, the excitatory subpopulations receive input from distinct sources of background noise (gray traces).

During the learning, the network is exposed to repeated presentations of two sequences {A,F,B,D} and {A,F,C,E}, where the first sequence is shown with a relative frequency p and the second with $1-p$ (e.g., $p = 0.2$ in Fig. 2A). In the following, we refer to {A,F,B,D} as sequence 1 and to {A,F,C,E} as sequence 2. Before learning, presenting a sequence element causes all neurons in the respective subpopulation to fire. During the learning process, the repetitive sequential

presentation of sequence elements increases the strength of connections between the corresponding subpopulations to a point where the activation of a certain subpopulation by an external input generates dAPs in a specific subset of neurons in the subpopulation representing the subsequent element. The generation of the dAPs results in a long-lasting depolarization (∼50-500ms) of the soma. We refer to neurons that generate a dAP as predictive neurons. When receiving an external input, predictive neurons fire earlier as compared to non-predictive neurons. When enough predictive neurons are present within a certain subpopulation, their advanced spike initiates a fast and strong inhibitory feedback to all excitatory neurons, ultimately suppressing the firing of non predictive neurons. After learning, the model develops specific subnetworks representing the learned sequences (Fig. 1B), such that the presentation of a sequence element leads to a context dependent prediction of the subsequent element (Bouhadjar et al., 2021). As a result of Hebbian learning, the synaptic weights in the subnetwork corresponding to the most frequent sequence during learning are stronger than those for the less frequent sequence (Fig. 1B, Fig. 3A, Fig. 4A). In the prediction mode, this asymmetry in synaptic weights plays no role. For ambiguous stimuli, all potential outcomes are predicted, i.e., the network predicts both "C" and "E" simultaneously in response to stimuli "A" and "F", irrespective of the training frequencies.

The model can be configured into a replay mode, where the network autonomously replays learned sequences in response to a cue stimulus. This is achieved by changing the excitability of the neurons such that the activation of a dAP alone can cause the neurons to fire (Bouhadjar et al., 2021). In the replay mode, we present ambiguous cues and study whether the network can replay sequences following different strategies (Fig. 2B). We refer to the "maximum probability" strategy (Fig. 2B, left) as the case where the network exclusively replays the sequence with the highest occurrence frequency. "Probability matching" is referred to when the replay frequency of a sequence matches its occurrence frequency during training (Fig. 2B, middle). We call the strategy "full exploration" when all sequences are randomly replayed with the same frequency, irrespective of the occurrence frequency during training (Fig. 2B, right). In Fig. 3, we illustrate the network's decision behavior by presenting the ambiguous cue stimulus "A" three times. In the absence of noise, the network adopts the maximum probability strategy (Fig. 3B): as a result of the higher weights between the neurons representing the more frequent sequence, the dAPs get activated earlier in these neurons, which advances their somatic firing times with respect to the neurons of the less frequent sequence. This advanced response time quickly activates the inhibitory neuron, which suppresses the activity of the neurons representing the less frequent sequence.

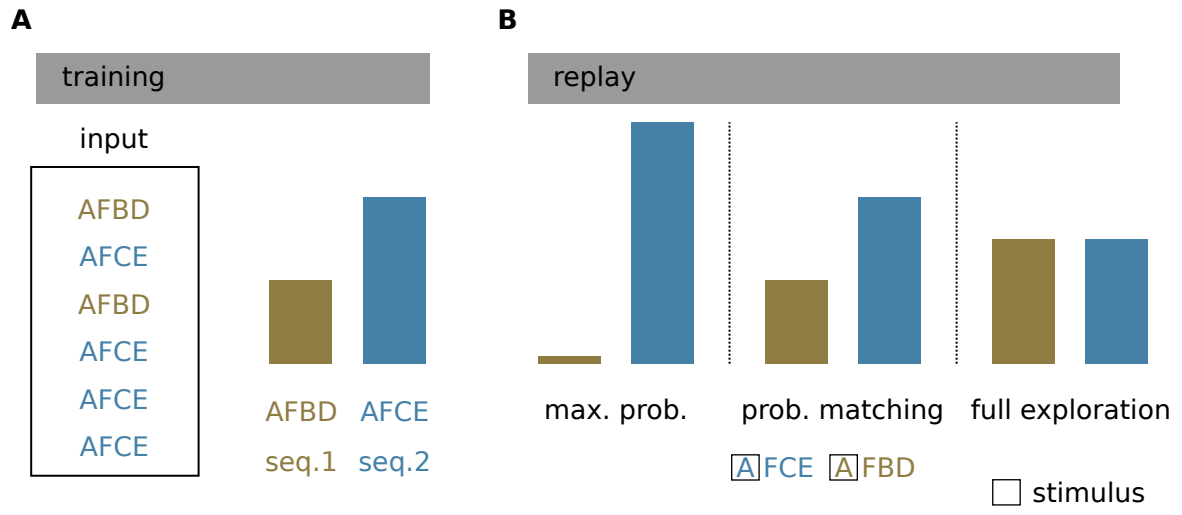

Figure 2: Task. A) During learning, the model is exposed to two (or more) competing sequences with different frequencies. Here, sequence 1 ({A,F,C,E}; blue) is shown twice as often as sequence 2 ({A,F,B,D}; brown). The respective normalized occurrence frequencies $1/3$ and $2/3$ are depicted by the histogram. **B)** During replay, the network autonomously recalls the sequences in response to an ambiguous cue signal (first sequence element; black squares) according to different strategies. Maximum probability (max-prob): Only the sequence with the highest occurrence frequency during training is replayed. Probability matching (prob. matching): The replay frequency of a sequence matches its occurrence frequency during training. Full exploration: All sequences are randomly replayed with the same frequency, irrespective of the occurrence frequency during training. Histograms represent the replay frequencies.

To assess the replay performance quantitatively, we present the cue stimulus "A" for N_{t} trials and examine the replay frequency of the two sequences {A,F,B,D} and {A,F,C,E} as a function of their occurrence frequencies during training. We define the sequences {A,F,B,D} or {A,F,C,E} to be successfully replayed if more than 50% of the neurons in the last subpopulations "E" or "D" have fired, respectively. In the absence of noise, the network replays only the sequence with the highest occurrence frequency p (Fig. 4E). For $p < 0.5$, the sequence {A,F,C,E} is the only

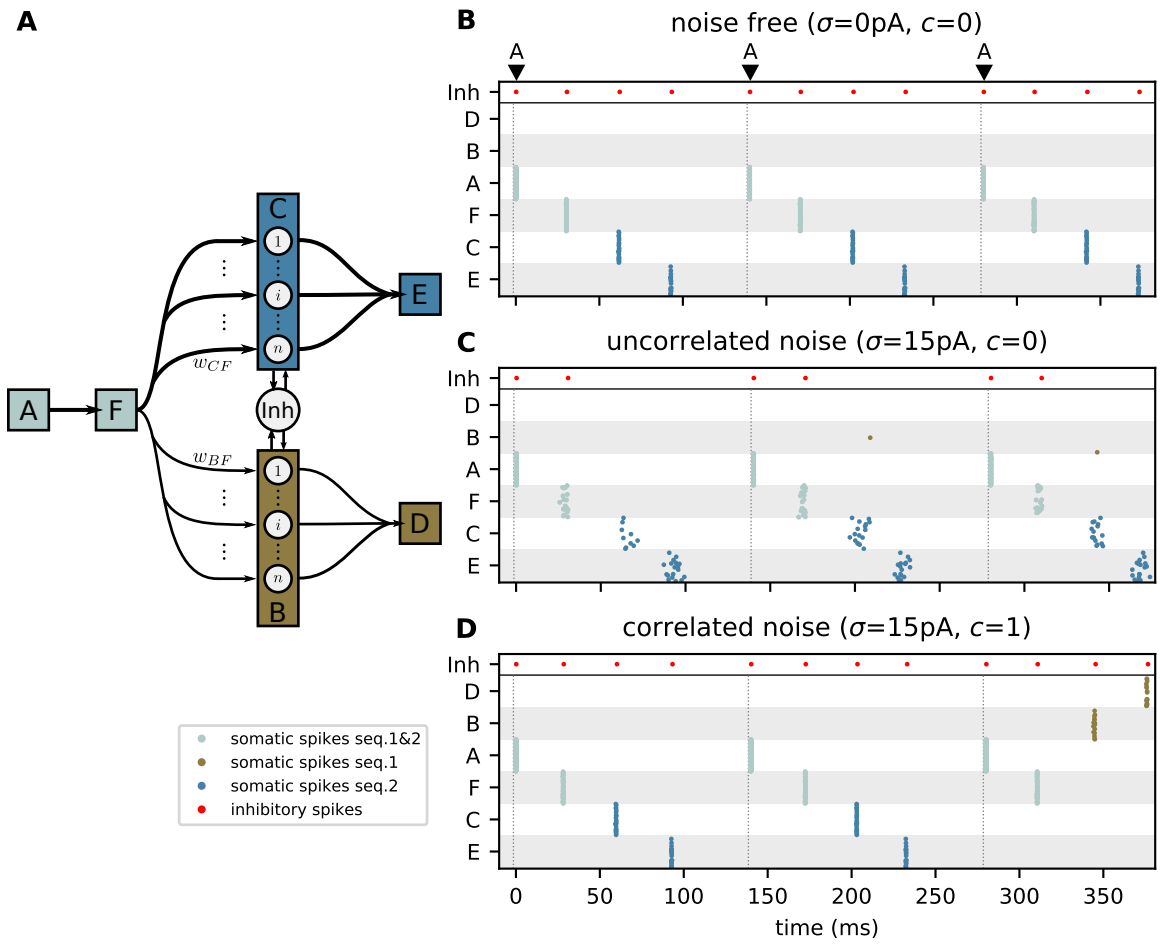

Figure 3: Correlated noise enhances exploratory behavior. A) Sketch of subpopulations of excitatory neurons (boxes) representing the elements of the two sequences {A,F,C,E} (seq. 2) and {A,F,B,D} (seq. 1). The subpopulations "C" and "B" are unfolded showing their respective neurons. The arrows depict the connections after learning the two sequences. The line thickness represents the synaptic weight. The presentation of the character "A" constitutes an ambiguous cue during replay. The inhibitory neuron (Inh) mediates competition between subpopulations (WTA). **B, C, D)** Spiking activity in the subpopulations depicted in panel A in response to three repetitions of the ambiguous cue "A" (black triangles at the top and vertical dotted lines) for three different noise configurations $\sigma = 0\,\text{pA}$, $c = 0$ (B), $\sigma = 15\,\text{pA}$, $c = 0$ (C), and $\sigma = 15\,\text{pA}$, $c = 1$ (D). Brown, blue, and silver dots mark somatic spikes of excitatory neurons corresponding to sequence 1, sequence 2, and both, respectively. Red dots mark spikes of the inhibitory neuron. Network size: number of subpopulations $M = 8$ and number of neurons per subpopulation $n_\text{E} = 150$. See Tab. 2 for remaining parameters.

one replayed in all trials, and for $p > 0.5$ the sequence {A,F,B,D} becomes the dominant one. To understand this behavior, we inspect the response latencies $t^{B/C}$ of the subpopulations "B" and "C" as a function of the occurrence frequencies of the sequences {A,F,B,D} and {A,F,C,E} (Fig. 4B). The mean of the response latency t^B or t^C is smaller for the subpopulation participating in the sequence with the higher frequency. The response latencies t^B and t^C decrease with increasing the respective occurrence frequencies. As the network is here operating in the absence of noise, the distribution of the response latencies $t^{B/C}$ across trials is narrow. Consequently, neurons representing the most frequent sequence fire earlier in all trials. This advanced response quickly activates the inhibitory neuron, which suppresses the activity of the neurons representing the less frequent sequence. For training frequencies between 0.4 and 0.6, the difference between t^B and t^C is small compared to the response latency of the WTA circuit. Hence, both sequences are occasionally replayed simultaneously (Fig. 4E).

To foster exploratory behavior, i.e., to enable occasional replay of the low-frequency sequence, we equip the excitatory neurons with noise. In this work, we investigate two different forms of noise. Here, we first consider noise provided in the form of stationary synaptic background input (see below for an alternative form of noise). Each subpopulation of excitatory neurons receives input from its private pool of independent excitatory and inhibitory Poissonian spike sources (Fig. 1C). These background inputs are parameterized by the noise amplitude σ (standard deviation of the synaptic input current arising from these background inputs) and the noise correlation c (see Fig. 1C and Methods). Only inputs to neurons of the same subpopulation are correlated by an extent parameterized by c. Neurons in different subpopulations receive uncorrelated inputs. The noise amplitude σ is chosen such the subthreshold membrane potentials of the excitatory neurons are fluctuating without eliciting additional spikes. As a consequence,

the distributions of response latencies $t^{B/C}$ across trials may be broadened and partly overlap (Fig. 4C,D). As we will show in the following, the network can adopt different replay strategies (Fig. 2B) depending on the amount of this overlap. Note that noise is injected only during replay, but not during learning. During training, the weak noise employed here hardly affects the network behavior as the external inputs (stimulus) are strong and lead to a reliable, immediate responses.

With uncorrelated noise ($c = 0$), the replay behavior remains effectively non-explorative, i.e., only the high-frequency sequence is replayed in response to the cue (Fig. 3C). This is explained by the fact that each sequence element is represented by a subset of ρ neurons. The response latency t corresponds to the population average $t = \frac{1}{\rho}\sum_{i=1}^{\rho} t_{\mathsf{s},i}$ of the response latencies $t_{\mathsf{s},i}$ (time of first spike after the cue) for each individual neuron i within this subset. The across-trial variance

$$v = \mathsf{Var}(t) = \frac{1}{\rho} v_{\mathsf{s}} + \frac{\rho - 1}{\rho} c_{\mathsf{s}} v_{\mathsf{s}} \tag{20}$$

of this population measure t is determined by the population size ρ, the population averaged spike-time variance $v_{\mathsf{s}} = \frac{1}{\rho}\sum_{i}^{\rho} \mathsf{Var}(t_{\mathsf{s},i})$, and the population averaged spike-time correlation coefficient $c_{\mathsf{s}} = \frac{1}{\rho(\rho-1)v_{\mathsf{s}}}\sum_{i}^{\rho}\sum_{j\neq i}^{\rho} \mathsf{Cov}(t_{\mathsf{s},i}, t_{\mathsf{s},j})$, with $\mathsf{Cov}(t_{\mathsf{s},i}, t_{\mathsf{s},j})$ denoting the spike-time covariance for two neurons i and j. The response-latency statistics v_{s} and c_{s} depend on the input noise statistics σ and c in a unique and monotonous manner (Goedeke & Diesmann, 2008; De la Rocha et al., 2007). In the absence of correlations ($c = c_{\mathsf{s}} = 0$), the across-trial variance v of t vanishes for large population sizes ρ. For finite population sizes, v is non-zero but small (Fig. 4C). The effect of the synaptic background noise on the variability of response latencies largely averages out. Hence, the average advance in the response of the population representing the high-frequency sequence cannot be overcome by noise; the network typically replays only the sequence with the higher occurrence frequency during training (Fig. 4F). For small differences in the occurrence frequencies (i.e., $p \approx 0.5$), the network occasionally fails to replay any sequence or replays both sequences. The mechanism underlying this behavior is explained below.

Noise averaging is efficiently avoided by introducing noise correlations. For perfectly correlated noise and, hence, perfectly synchronous spike responses ($c = c_{\mathsf{s}} = 1$), the across-trial variance v of the response latency t is identical to the across-trial variance v_{s} of the individual spike responses, i.e., $v = v_{\mathsf{s}}$, irrespective of the population size ρ; see (20). For smaller but non-zero spike correlations ($0 < c_{\mathsf{s}} < 1$), the latency variance v is reduced but doesn't vanish as ρ becomes large. Hence, in the presence of correlated noise, the across-trial response latency distributions for two competing populations have a finite width and may overlap (Fig. 4D), thereby permitting an occasional replay of the sequence observed less often during training (Fig. 4G). Replay, therefore, becomes more exploratory, such that the occurrence frequencies during training are gradually mapped to the frequencies of sequence replay. With an appropriate choice of the noise amplitude and correlation, even an almost perfect match between training and replay frequencies can be achieved (probability matching; Fig. 4G).

The results presented so far can be extended towards more than two competing sequences. As a demonstration, we train the network using three sequences {A,F,C,E}, {A,F,B,D}, and {A,F,G,H} presented with different relative frequencies. By adjusting the noise amplitude σ and correlation c, the replay frequencies can approximate the training frequencies (Fig. 5).

Noise canceling cannot be overcome by increasing noise amplitude

For subpopulations of finite size ρ, the variance v of the response latency t remains finite, and can be increased by scaling up the variance of the noise, even without correlation; see (20). However, this solution cannot be applied to network models where a decision is mediated by a fast WTA circuit. In the presence of uncorrelated noise with high amplitude, the spikes from all neurons, in all competing subpopulations, are similarly dispersed. A large dispersion in spike times prohibits a fast and reliable activation of inhibition by one of the competing subpopulations. The WTA mechanism, therefore, fails at selecting a unique sequence. Consequently, both sequences run through in most of the trials (Fig. 6A). An additional problem of the uncorrelated noise is that it impairs the propagation of the activity across the subpopulations of neurons. As our model relies on the propagation of synchronously firing neurons, the spike time dispersion resulting from the uncorrelated noise bears the risk that the spikes generated may be too dispersed to trigger dAPs in the next subpopulation (Fig. 6). As a result of these two problems, more explorative behavior cannot be achieved by increasing the amplitude of uncorrelated noise. Instead, the probability of simultaneous replay (no decision) and the failure rate increase (Fig. 6B).

Noise correlations lead to more synchronous responses, thereby reducing the overlap between the within-trial latency distributions of the two competing populations "B" and "C" (Fig. 3D). In each trial, the WTA dynamics is therefore triggered by just one of the two populations, rather than by both. Further, synchronous firing leads to a more robust activation of the subsequent subpopulation, and hence, a more robust replay. Hence, noise correlations help not only in generating more explorative behavior, but also in reducing replay failures and the chance of simultaneous activation of competing sequences (Fig. 4G).

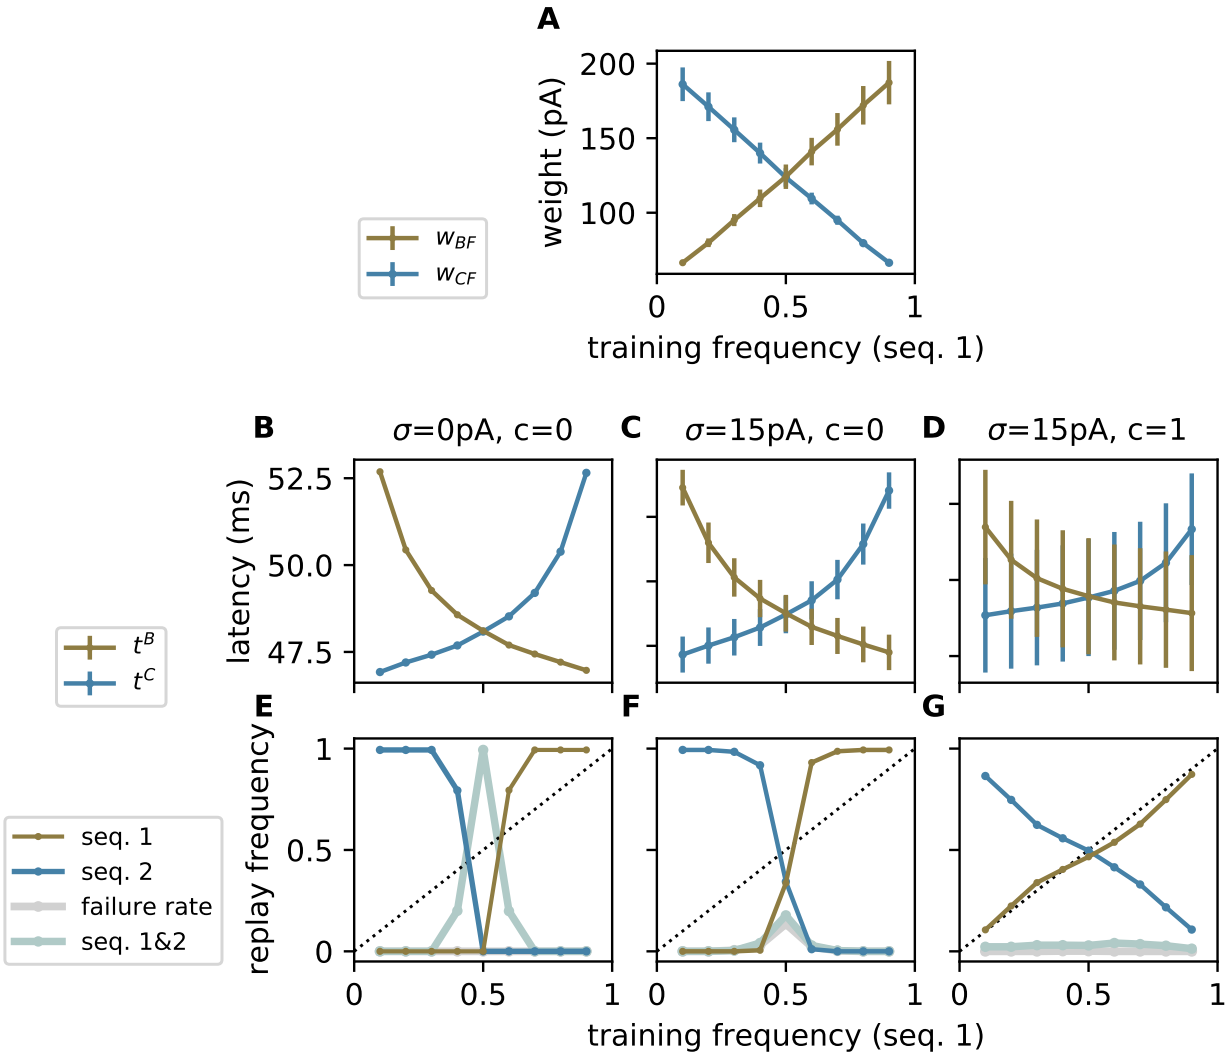

Figure 4: Uncorrelated noise averages out in population based encoding. Dependence of **A)** the compound weights (PSC amplitudes) w_{BF} (brown) and w_{CF} (blue; see Fig. 3A), **B–D)** the population averaged response latencies t^B and t^C (subpopulation averaged time of first spike after the cue) for subpopulations "B" (brown) and "C" (blue), and **E–G)** the relative replay frequencies of sequences 1 (brown) and 2 (blue) along with the failure rate (gray) and the joint probability of replaying both sequences (silver) on the occurrence frequency of sequence 1 during training. Panels B–G depict results for three different noise parameterization $\sigma = 0\,\mathrm{pA}$, $c = 0$ (B,E), $\sigma = 15\,\mathrm{pA}$, $c = 0$ (C,F), and $\sigma = 15\,\mathrm{pA}$, $c = 1$ (D,G). In panel A, circles and error bars depict the mean and the standard deviation across different network realizations. In panels B–D, circles and error bars represent the mean and the standard deviation across $N_{\mathrm{t}} = 150$ trials (cue repetitions), averaged across 5 different network realizations. In panels E–G, circles represent the mean across $N_{\mathrm{t}} = 150$ trials, averaged across 5 different network realizations. See Tab. 2 for remaining parameters.

Noise amplitude and level of correlation control replay strategy

We know from psychophysics experiments that humans and animals are able to flexibly change the strategy by which they select probabilistic outcomes (Shanks et al., 2002; Cohen et al., 2007). Our model is also able to express these different strategies. For small noise amplitudes and irrespective of the correlation level, the model replays deterministically the frequent sequence (max-prob, see Fig. 7A). With increasing the noise amplitude and sufficiently large correlation, the replay frequency approximately matches the training frequency (probability matching, Fig. 7B). Increasing the noise further leads to an even more explorative replay, where the replay frequencies become less dependent on the training frequencies (Fig. 7C). Full exploration requires a large amount of noise to overwrite the asymmetry in the synaptic weights. In the presence of high-amplitude noise, neurons may become strongly hyperpolarized during stimulus arrival, and hence, fail to respond. Further, strong noise induced depolarizations may trigger spurious bursts of synchronous spikes.

In our work, it is also possible to change the replay strategy by changing the correlation level (see Fig. S1). If the noise amplitude is well adjusted, a low correlation level makes the dynamics non-explorative. With increasing the correlation level dynamics become more explorative. This suggests, therefore, that the replay strategy can be controlled by changing the noise amplitudes or the correlation level. Adjusting noise amplitudes can be achieved in

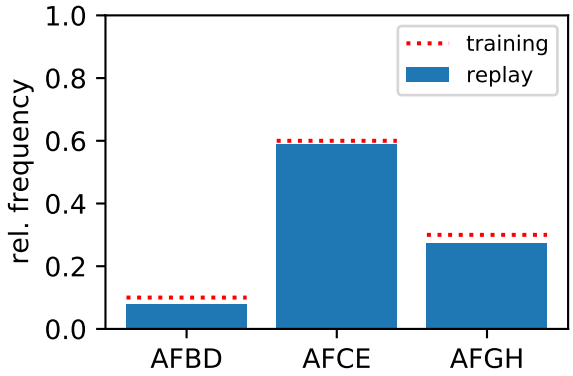

Figure 5: Multiple competing sequences are learned and replayed according to their occurrence frequencies (probability matching). During learning, three partly overlapping sequences {A,F,C,E}, {A,F,B,D}, and {A,F,G,H} are repetitively presented with relative frequencies 0.2, 0.5, 0.3, respectively (dotted red lines). After learning, the network autonomously replays the learned sequences in response to the ambiguous cue "A" with frequencies depicted by the blue bars. Noise parameters: $\sigma = 15\,\text{pA}$, $c = 1$. See Tab. 2 for remaining parameters.

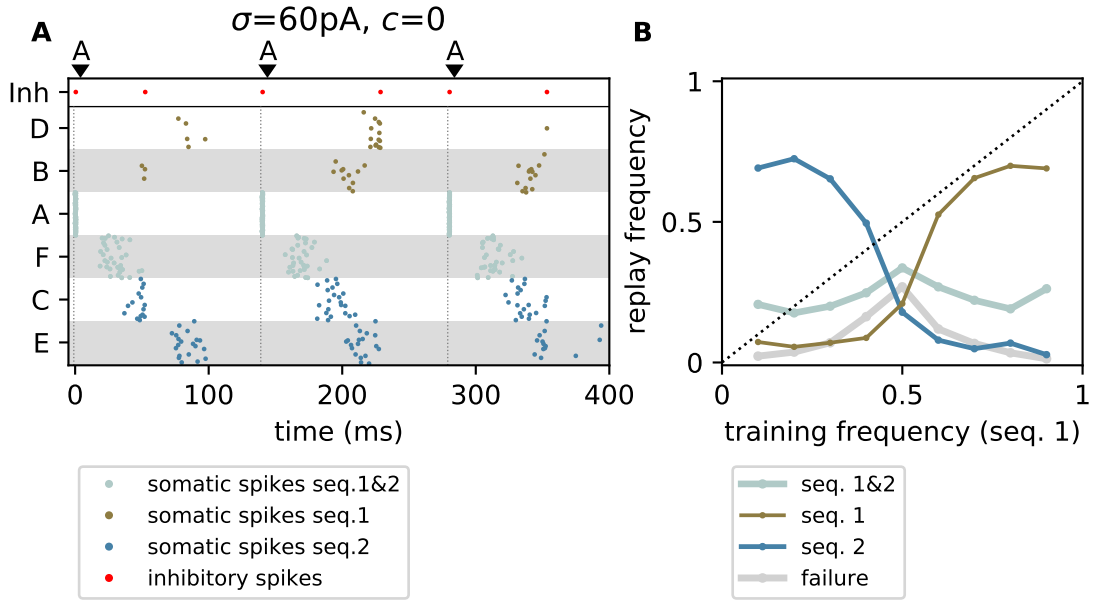

Figure 6: Winner-take-all mechanism fails when increasing the amplitude of the uncorrelated noise. A) Brown, blue, and silver dots mark somatic spikes of excitatory neurons belonging to sequence {A,F,B,D} (seq. 1), sequence {A,F,C,E} (seq. 2), or both, respectively. Red dots mark spikes from inhibitory neurons. Each trial is initiated by stimulating the first element in the sequence ("A", see dark arrows and vertical dashed lines). During training, the sequences 1 and 2 are shown with relative frequencies 0.2 and 0.8, respectively. **B)** The relative replay frequency of sequence 1 (brown) and sequence 2 (blue) along with the failure rate (gray) and the joint probability of replaying both sequences (silver) are plotted as a function of the relative training frequency of sequence 1. Circles represent the mean across $N_\text{t} = 150$ trials averaged across 5 network realizations. Parameters: $\sigma = 60\,\text{pA}$ and $c = 0$. See Tab. 2 for the remaining parameters.

biology by changing the rate of the background input or by changing the effective weights, which can be modulated by neuromodulation (Atherton et al., 2015) or attention (Baluch & Itti, 2011). So far, we studied correlations induced by shared presynaptic inputs. To achieve explorative dynamics, shared input correlations need to be high ($c \sim 1$). However, shared input correlations resulting from the cortical anatomy are rather low (Abeles, 1991; Braitenberg & Schüz, 1998; Shadlen & Newsome, 1998; Song et al., 2005). In the next section, we therefore propose an alternative form of noise, where correlations are generated by the network dynamics.

Random stimulus locking to spatiotemporal oscillations as natural form of noise

So far, we have used shared synaptic background input in the form of stationary Poissonian spikes as a source of correlated noise. Locally coherent noise may however also be provided in the form of randomness in stimulus timing with respect to spatiotemporal background oscillations. In the presence of traveling cortical waves, for example, nearby neurons in a given subpopulation share the same oscillation phase, whereas distant neurons belonging to

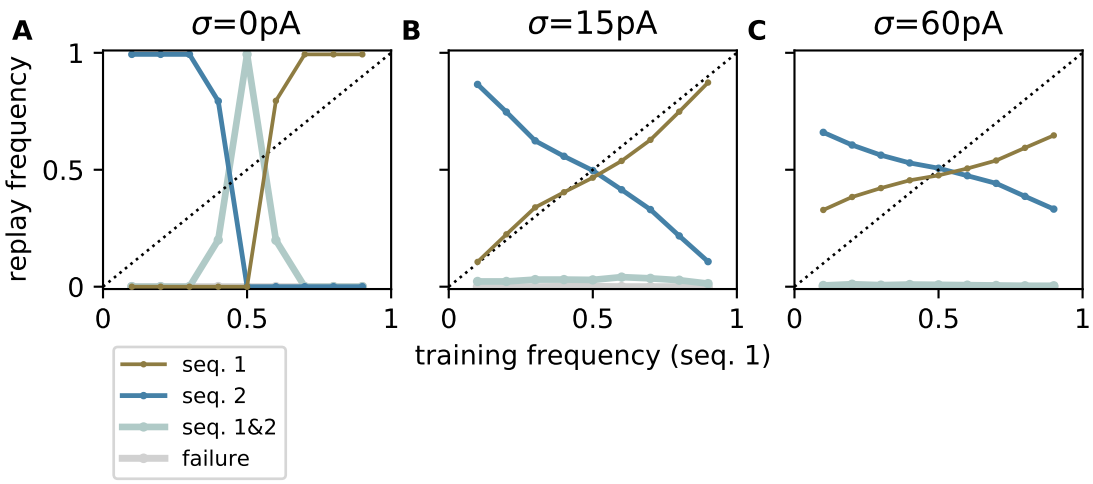

Figure 7: Different replay strategies achieved by increasing the noise amplitude. The relative replay frequency of sequence 1 (brown) and sequence 2 (blue) along with the failure rate (gray) and the joint probability of replaying both sequences (silver) are plotted as a function of the relative training frequency of sequence 1 for different noise amplitudes $\sigma = 0\,\text{pA}$ (**A**), $\sigma = 15\,\text{pA}$ (**B**), and $\sigma = 60\,\text{pA}$ (**C**) with correlation coefficient $c = 1$. Circles represent the mean across $N_\text{t} = 150$ trials, averaged across 5 different network realizations. See Tab. 2 for remaining parameters.

different subpopulations experience different phases (Fig. 8). At the time of stimulus arrival, the neurons in the up phase are more excitable and tend to fire earlier than neurons in a down phase. Spatiotemporal oscillations in cortical activity are ubiquitous and occur in different forms (Nauhaus et al., 2009; Muller & Destexhe, 2012; Sato et al., 2012). In nature, external stimuli are usually not consistently locked to such oscillations. It is therefore reasonable to assume that the stimulus onset times are random with respect to the oscillation phase.

Here, we exploit this form of randomness to increase the trial-to-trial variability during replay. To investigate its effect on the replay performance, we first train the network in the absence of any background input using two sequences {A,F,C,E} and {A,F,B,D} with relative training frequencies p and $1 - p$, respectively. During replay, we inject an oscillating background current with amplitude a and frequency f into all excitatory neurons (see Methods). Neurons within a given subpopulation share the same oscillation phase. Phases for different subpopulations are randomly drawn from a uniform distribution between 0 and 2π. The replay performance of the network is assessed by monitoring the network responses to repetitive presentations of an external cue "A" with random, uniformly distributed inter-cue intervals $\Delta T_\text{cue} \sim \mathcal{U}(u_\text{min}\, u_\text{max})$. We define sequences {A,F,B,D} or {A,F,C,E} to be successfully replayed if more than 50% of the neurons in the last subpopulations "E" or "D" have fired, respectively. The analysis is repeated for a range of training frequencies p, oscillation amplitudes a, and frequencies f.

Depending on the choice of the oscillation amplitude a and frequency f, the network replicates different replay strategies (Fig. 9). For low-amplitude oscillations, the model replays only the sequence with the higher training frequency (max-prob). With increasing oscillation amplitude, the model becomes more explorative and occasionally replays the less frequent sequence. By adjusting the oscillation amplitude, the replay frequency can be closely matched to the training frequency. In line with experimental recording (Buzsáki, 2006; Buzsáki & Draguhn, 2004), this behavior is observed for a range of physiological oscillation frequencies: alpha ($10\,\text{Hz}$), beta ($30\,\text{Hz}$), gamma ($70\,\text{Hz}$). Due to the low-pass filtering of the neuronal membranes and synapses, higher oscillation frequencies have a smaller effect. Consequently, increasing the oscillation frequency leads to a more reliable replay of the frequent sequence. For slow oscillations with long periods that are large compared to the average inter-cue interval, the network responses in subsequent trials are more correlated. For sufficiently many trials, however, the network can still explore different solutions.

Discussion

This work proposes a spiking neural network model able to recall sequences in response to ambiguous cues following different strategies. In this model, explorative recall strategies are achieved by providing the network with external coherent noise. We explore two forms of coherent noise implemented either in the form of shared synaptic input or a random stimulus locking to global spatiotemporal oscillations in the neuronal activity. The model can switch between different replay strategies by adjusting the noise characteristics such as the noise or oscillation amplitude, as well as the noise correlation or oscillation frequency.

The sequence processing model proposed here relies on a form of population encoding. In the absence of correlations, noise injected to single neurons therefore largely averages out and leads to a quasi-deterministic and

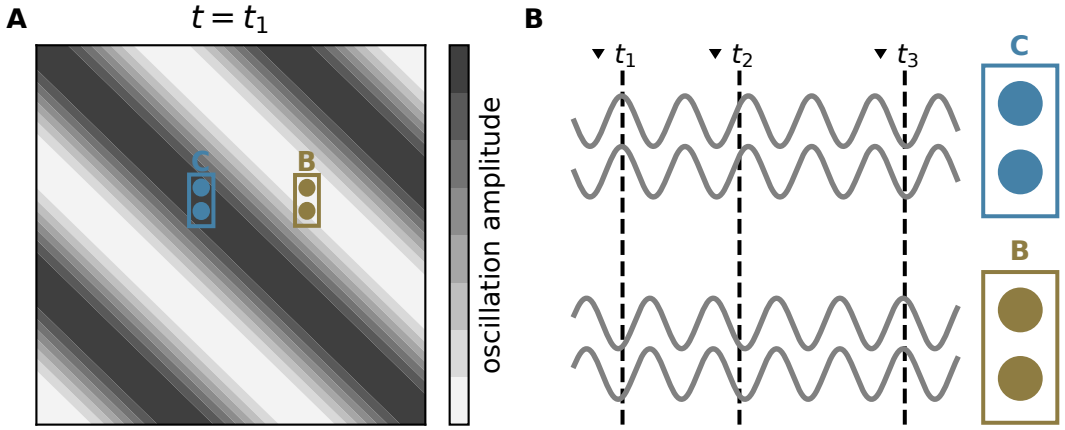

Figure 8: Random locking of stimulus to global oscillations as a form of noise. A) Snapshot of a wave of activity traveling across a cortical region at time t_1 of the 1st stimulus onset. Grayscale depicts wave amplitudes in different regions. Brown and blue rectangles mark populations of neurons with stimulus preferences "B" and "C", respectively. **B)** Background inputs to neurons in populations "B" and "C" at different times. Background inputs to each population "B" and "C" at different times. Background inputs to neurons within each population are in phase due to their spatial proximity. Background inputs to different populations are phase shifted. Arrows on the top depict stimulus onset times. The times $t_1, t_2, \ldots$ indicate input arrival to populations "B" and "C" (dashed vertical lines are random not locked to the background activity).

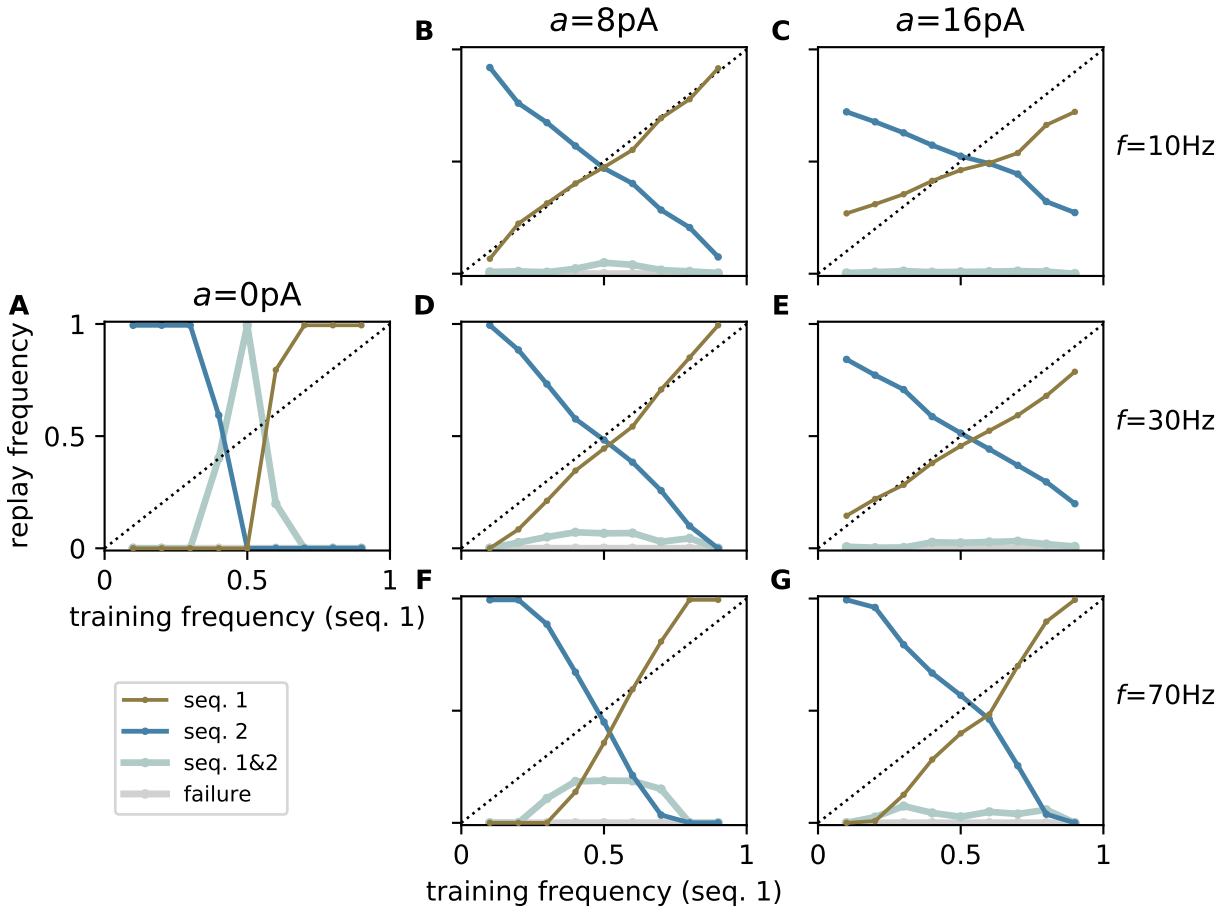

Figure 9: Background oscillations as a mean for changing the replay strategies. The relative replay frequency of sequence 1 (brown) and sequence 2 (blue) along with the failure rate (gray) and the joint probability of replaying both sequences (silver) are plotted as a function of the relative training frequency of sequence 1 for different amplitudes $a \in \{0, 8, 16\}$ and frequencies of the background oscillations: $f = 10\,\text{Hz}$ (**B,C**), $f = 30\,\text{Hz}$ (**A,D,E**), and $f = 70\,\text{Hz}$ (**F,G**). Circles represent the mean across $N_\text{t} = 150$ trials, averaged across 5 network realizations. See Tab. 2 for remaining parameters.

non-exploratory behavior. Locally correlated noise, in contrast, permits an explorative recall behavior where the sequence frequency during learning can be gradually mapped to the recall frequency. Furthermore, noise correlations foster synchronization between neurons within subpopulations, and thereby lead to a more robust context-specific activation of sequences during recall. The problem of noise averaging and the proposed solution are not unique to the model presented here, but are generic for all systems where relevant state variables arise from superpositions of many noisy, uncorrelated components. Fluctuations in the total input current of a single neuron resulting from superpositions of thousands of synaptic inputs, for example, can be efficiently controlled by the level of correlation in the presynaptic activity (Salinas & Sejnowski, 2001). Similarly, explorative behavior in other models of population based probabilistic computing (e.g. Legenstein & Maass, 2014) can be enhanced by equipping neurons within each population with correlated noise.

Correlation in neuronal firing can originate from both anatomical constraints or network dynamics (Tetzlaff et al., 2012; Helias et al., 2014). In this study, we investigate both types. The first type of noise is implemented in the form of irregular synaptic background input (Faisal et al., 2008; Fellous et al., 2004; Destexhe et al., 2001; Holt et al., 1996), where the correlation between neurons of the same subpopulation is resulting from shared presynaptic sources (Stroeve & Gielen, 2001; Kriener et al., 2008). From an anatomical perspective, this is reasonable as neighboring neurons indeed receive a considerable amount of inputs from identical presynaptic neurons. However, we show that the level of shared-input correlation required for an effective avoidance of noise averaging and maintenance of near synchronous activity is rather high, which contradicts anatomical studies reporting small connection probabilities in local cortical circuits, and hence, low levels of shared input correlation (Abeles, 1991; Braitenberg & Schüz, 1998; Song et al., 2005; Shadlen & Newsome, 1998). We therefore propose a second, biologically more plausible type of coherent noise resulting from a random stimulus locking to an intrinsic spatiotemporal coherent activity pattern on a large spatial scale, such as waves of cortical activity. Coherent spatiotemporal activity patterns in the cortex are observed in many different forms and under various conditions, including different sleep states, but also in awake behaving animals (Buzsáki, 2006; Buzsáki & Draguhn, 2004; Sato et al., 2012; Denker et al., 2018). In such states, neighboring neurons receive coherent input with identical phase, whereas distant neurons are exposed to different phases. Natural external stimuli such as cue signals are usually not systematically locked to this type of intrinsic activity (unless stimuli are presented in closed-loop experiments). It is therefore reasonable to assume that the stimulus onset is random with respect to the internal state. In this study, we employ activity waves as a specific form of coherent spatiotemporal activity, and show that explorative behavior is generated for a range of plausible oscillation frequencies. We propose that a similar behavior can be achieved for other non-oscillatory forms of coherent activity, such as transient propagating wave fronts or bumps (Ermentrout, 1998; Coombes, 2005; Muller et al., 2018), as well as by other factors modulating the excitability of neighboring neurons in a coherent manner, such as transient neuromodulatory signals.

By changing the noise characteristics (such as the amplitude or frequency of the background activity, or the level of correlation), the model proposed in this study can replay competing sequences according to different strategies. For low levels of noise, the network systematically replays the sequence that occurred most often during learning (max-prob). For higher noise levels, it can match the replay frequency to the occurrence frequency during training (probability matching), or become even more explorative. This offers a potential mechanistic explanation of how animals can adjust their decision strategy based on environmental conditions (Cohen et al., 2007). In the living brain, the noise properties could be controlled by neuromodulatory signals or by inputs from other brain areas (e.g., during attention; Cohen & Maunsell, 2009). Our and many other studies predict that, in cases where the decision strategy is shifted towards exploration, more energy needs to be provided for noise generation. In line with this prediction, Daw et al. (2006) show that explorative behavior is accompanied by an increase in the BOLD signal amplitude in cortical areas associated with decision making.

A number of previous studies suggest that synaptic stochasticity, i.e., the variability in postsynaptic responses including synaptic failure (Branco & Staras, 2009), may constitute an efficient source of noise for probabilistic computations in neuronal circuits (Maass, 2014; Neftci et al., 2016). The total input to a neuron resulting from large ensembles of synapses, however, is likely to be subject to noise averaging, unless the variability of synaptic responses is correlated across synapses. To date, it remains unclear how such correlations could potentially arise. Localized neuromodulatory signals or shared presynaptic spike histories may play a role in this.

Overall, our work ties together concepts from sequence processing and decision making in the face of ambiguity. It demonstrates that locally coherent noise is a potential mechanism underlying exploratory behavior, and shows that a random stimulus locking to macroscopic coherent activity patterns, e.g., traveling waves, can constitute such a form of noise.

Acknowledgments

The authors thank Abigail Morrison and Alexander René for valuable discussions on the project. This project was funded by the Helmholtz Association Initiative and Networking Fund (project number SO-092, Advanced Computing Architectures), and the European Union's Horizon 2020 Framework Programme for Research and Innovation under

the Specific Grant Agreement No. 785907 (Human Brain Project SGA2) and No. 945539 (Human Brain Project SGA3). Open access publication was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation; grant 491111487).

Author contributions

All authors conceived and designed the work. YB performed the simulations, and analyzed and visualized the data. YB wrote the first draft of the manuscript. All authors reviewed the manuscript and approved it for publication. YB was supervised by TT and DJW.

References

Abbott, L. F., & Nelson, S. B. (2000). Synaptic plasticity: taming the beast. *Nat. Neurosci. 3*, 1178–1183.

Abeles, M. (1991). *Corticonics: Neural Circuits of the Cerebral Cortex* (1st ed.). Cambridge: Cambridge University Press.

Atherton, L. A., Dupret, D., & Mellor, J. R. (2015). Memory trace replay: the shaping of memory consolidation by neuromodulation. *Trends Neurosci. 38*(9), 560–570.

Baluch, F., & Itti, L. (2011). Mechanisms of top-down attention. *Trends Neurosci. 34*(4), 210–224.

Bod, R., Hay, J., & Jannedy, S. (2003). *Probabilistic linguistics*. Mit Press.

Bouhadjar, Y., Wouters, D. J., Diesmann, M., & Tetzlaff, T. (2021). Sequence learning, prediction, and replay in networks of spiking neurons. *ArXiv*.

Braitenberg, V., & Schüz, A. (1998). *Cortex: Statistics and Geometry of Neuronal Connectivity* (2nd ed.). Berlin: Springer-Verlag.

Branco, T., & Staras, K. (2009). The probability of neurotransmitter release: variability and feedback control at single synapses. *Nat. Rev. Neurosci. 10*(5), 373–383.

Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J. Comput. Neurosci. 8*(3), 183–208.

Brunel, N., & Hakim, V. (1999). Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Comput. 11*(7), 1621–1671.

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 Comput. Biol. 7*, e1002211.

Buzsáki, G. (2006). *Rhythms of the Brain*. Oxford University Press.

Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science 304*, 1926–1929.

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should i stay or should i go? how the human brain manages the trade-off between exploitation and exploration. *Philos. Trans. R. Soc. B 362*(1481), 933–942.

Cohen, M. R., & Maunsell, J. H. R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci. 12*, 1594–1600.

Cone, I., & Shouval, H. Z. (2021). Learning precise spatiotemporal sequences via biophysically realistic learning rules in a modular, spiking network. *eLife 10*, e63751.

Coombes, S. (2005). Waves, bumps, and patterns in neural field theories. *Biol. Cybern. 93*, 91–108.

Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature 441*(7095), 876–879.

De la Rocha, J., Doiron, B., Shea-Brown, E., Kresimir, J., & Reyes, A. (2007). Correlation between neural spike trains increases with firing rate. *Nature 448*(16), 802–807.

Denker, M., Zehl, L., Kilavik, B. E., Diesmann, M., Brochier, T., Riehle, A., & Grün, S. (2018). LFP beta amplitude is linked to mesoscopic spatio-temporal phase patterns. *Sci. Rep. 8*(1), 1–21.

Destexhe, A., Rudolph, M., Fellous, J.-M., & Sejnowski, T. (2001). Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. *Neuroscience 107*(1), 13–24.

Dold, D., Bytschok, I., Kungl, A. F., Baumbach, A., Breitwieser, O., Senn, W., Schemmel, J., Meier, K., & Petrovici, M. A. (2019). Stochasticity from function—why the bayesian brain may need no noise. *Neural Netw. 119*, 200–213.

Ermentrout, B. (1998). Neural networks as spatio-temporal pattern-forming systems. *Rep. Prog. Phys. 61*(4), 353–430.

Faisal, A. A., Selen, L. P., & Wolpert, D. M. (2008). Noise in the nervous system. *Nat. Rev. Neurosci. 9*(4), 292–303.

Fellous, J., Tiesinga, P., Thomas, P., & Sejnowski, T. (2004). Discovering spike patterns in neuronal responses. *J. Neurosci. 12*(24), 2989–3001.

Gewaltig, M.-O., & Diesmann, M. (2007). NEST (NEural Simulation Tool). *Scholarpedia J. 2*(4), 1430.

Goedeke, S., & Diesmann, M. (2008). The mechanism of synchronization in feed-forward neuronal networks. *New. J. Phys. 10*, 015007.

Hahne, J., Diaz, S., Patronis, A., Schenck, W., Peyser, A., Graber, S., Spreizer, S., Vennemo, S. B., Ippen, T., Mørk, H., Jordan, J., Senk, J., Konradi, S., Weidel, P., Fardet, T., Dahmen, D., Terhorst, D., Stapmanns, J., Trensch, G., van Meegen, A., Pronold, J., Eppler, J. M., Linssen, C., Morrison, A., Sinha, A., Mitchell, J., Kunkel, S., Deepu, R., Hagen, E., Vierjahn, T., Kamiji, N. L., de Schepper, R., Machado, P., Albers, J., Klijn, W., Myczko, A., Mayner, W., Nagendra Babu, P., Jiang, H., Billaudelle, S., Vogler, B. S., Miotto, G., Kusch, L., Antonietti, A., Morales-Gregorio, A., Dolderer, J., Bouhadjar, Y., & Plesser, H. E. (2021). Nest 3.0.

Hansen, K. A., Hillenbrand, S. F., & Ungerleider, L. G. (2012). Effects of prior knowledge on decisions made under perceptual vs. categorical uncertainty. *Front. Neurosci. 6*.

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 Comput. Biol. 11*(12), e1004640.

Helias, M., Tetzlaff, T., & Diesmann, M. (2014). The correlation structure of local cortical networks intrinsically results from recurrent dynamics. *PLOS Comput. Biol. 10*(1), e1003428.

Holt, G. R., Softky, W. R., Koch, C., & Douglas, R. J. (1996). Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. *J. Neurophysiol. 75*(5), 1806–1814.

Jordan, J., Petrovici, M. A., Breitwieser, O., Schemmel, J., Meier, K., Diesmann, M., & Tetzlaff, T. (2019). Deterministic networks for probabilistic computing. *Sci. Rep. 9*, 18303.

Klampfl, S., & Maass, W. (2013). Emergence of dynamic memory traces in cortical microcircuit models through stdp. *J. Neurosci. 33*(28), 11515–11529.

Klos, C., Miner, D., & Triesch, J. (2018). Bridging structure and function: A model of sequence learning and prediction in primary visual cortex. *PLOS Comput. Biol. 14*(6), e1006187.

Kriener, B., Tetzlaff, T., Aertsen, A., Diesmann, M., & Rotter, S. (2008). Correlations and population dynamics in cortical networks. *Neural Comput. 20*, 2185–2226.

Legenstein, R., & Maass, W. (2014). Ensembles of spiking neurons with noise support optimal probabilistic inference in a dynamically changing environment. *PLOS Comput. Biol. 10*(10), e1003859.

Maass, W. (2014). Noise as a resource for computation and learning in networks of spiking neurons. *Proc. IEEE 102*(5), 860–880.

Maes, A., Barahona, M., & Clopath, C. (2020). Learning spatiotemporal signals using a recurrent spiking network that discretizes time. *PLOS Comput. Biol. 16*(1), e1007606.

Morrison, A., Diesmann, M., & Gerstner, W. (2008). Phenomenological models of synaptic plasticity based on spike-timing. *Biol. Cybern. 98*(6), 459–478.

Muller, L., Chavane, F., Reynolds, J., & Sejnowski, T. J. (2018). Cortical travelling waves: mechanisms and computational principles. *Nat. Rev. Neurosci. 19*(5), 255–268.

Muller, L., & Destexhe, A. (2012). Propagating waves in thalamus, cortex and the thalamocortical system: Experiments and models. *J. Physiol. 106*(5-6), 222–238.

Myers, J. L. (2014). Probability learning and sequence learning. *Handbook of Learning and Cognitive Processes, ed. WK Estes*, 171–205.

Nagendra Babu, P., Linssen, C., Eppler, J. M., Schulte to Brinke, T., Ziaeemehr, A., Fardet, T., Bouhadjar, Y., Duarte, R., Rumpe, B., & Morrison, A. (2021). Nestml 4.0.

Nauhaus, I., Busse, L., Carandini, M., & Ringach, D. L. (2009). Stimulus contrast modulates functional connectivity in visual cortex. *Nat. Neurosci. 12*, 70–76.

Neftci, E. O., Pedroni, B. U., Joshi, S., Al-Shedivat, M., & Cauwenberghs, G. (2016). Stochastic synapses enable efficient brain-inspired learning machines. *Front. Neurosci. 10*, 241.

O'Doherty, J. P., Cockburn, J., & Pauli, W. M. (2017). Learning, reward, and decision making. *Annu. Rev. Psychol. 68*(1), 73–100.

Plotnikov, D., Blundell, I., Ippen, T., Eppler, J. M., Rumpe, B., & Morrison, A. (2016). NESTML: a modeling language for spiking neurons. In A. Oberweis & R. Reussner (Eds.), *Modellierung 2016*, Volume P-254 of *Lecture Notes in Informatics (LNI)*, pp. 93–108. Modellierung 2016, Karlsruhe (Germany), 17 Mar 2016 - 19 Mar 2016: Gesellschaft für Informatik e.V. (GI).

Renart, A., De La Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., & Harris, K. D. (2010). The asynchronous state in cortical circuits. *Science 327*, 587–590.

Rotter, S., & Diesmann, M. (1999). Exact digital simulation of time-invariant linear systems with applications to neuronal modeling. *Biol. Cybern. 81*(5-6), 381–402.

Roxin, A., Brunel, N., & Hansel, D. (2005). The role of delays in shaping spatio-temporal dynamics of neuronal activity in large networks. *Phys. Rev. Lett. 94*(23), 238103.

Salinas, E., & Sejnowski, T. J. (2001). Correlated neuronal activity and the flow of neural information. *Nat. Rev. Neurosci. 2*(8), 539–550.

Sato, T. K., Nauhaus, I., & Carandini, M. (2012). Traveling waves in visual cortex. *Neuron 75*(2), 218–229.

Senk, J., Korvasová, K., Schuecker, J., Hagen, E., Tetzlaff, T., Diesmann, M., & Helias, M. (2020). Conditions for wave trains in spiking neural networks. *Phys. Rev. Res. 2*(2).

Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *J. Neurosci. 18*(10), 3870–3896.

Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *J. Behav. Decis. Mak. 15*(3), 233–250.

Song, S., Sjöström, P., Reigl, M., Nelson, S., & Chklovskii, D. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLOS Biol. 3*(3), e68.

Stroeve, S., & Gielen, S. (2001). Correlation between uncoupled conductance-based integrate-and-fire neurons due to common and synchronous presynaptic firing. *Neural Comput. 13*(9), 2005–2029.

Takahashi, K., Kim, S., Coleman, T. P., Brown, K. A., Suminski, A. J., Best, M. D., & Hatsopoulos, N. G. (2015). Large-scale spatiotemporal spike patterning consistent with wave propagation in motor cortex. *Nat. Commun. 6*(7169).

Tetzlaff, C., Kolodziejski, C., Timme, M., & Wörgötter, F. (2011). Synaptic scaling in combination with many generic plasticity mechanisms stabilizes circuit connectivity. *Front. Comput. Neurosci. 5*, 47.

Tetzlaff, T., Helias, M., Einevoll, G. T., & Diesmann, M. (2012). Decorrelation of neural-network activity by inhibitory feedback. *PLOS Comput. Biol. 8*(8), e1002596.

Tetzlaff, T., Rotter, S., Stark, E., Abeles, M., Aertsen, A., & Diesmann, M. (2008). Dependence of neuronal correlations on filter characteristics and marginal spike-train statistics. *Neural Comput. 20*(9), 2133–2184.

Vulkan, N. (2000). An economist's perspective on probability matching. *J. Econ. Surv. 14*(1), 101–118.

Supporting information

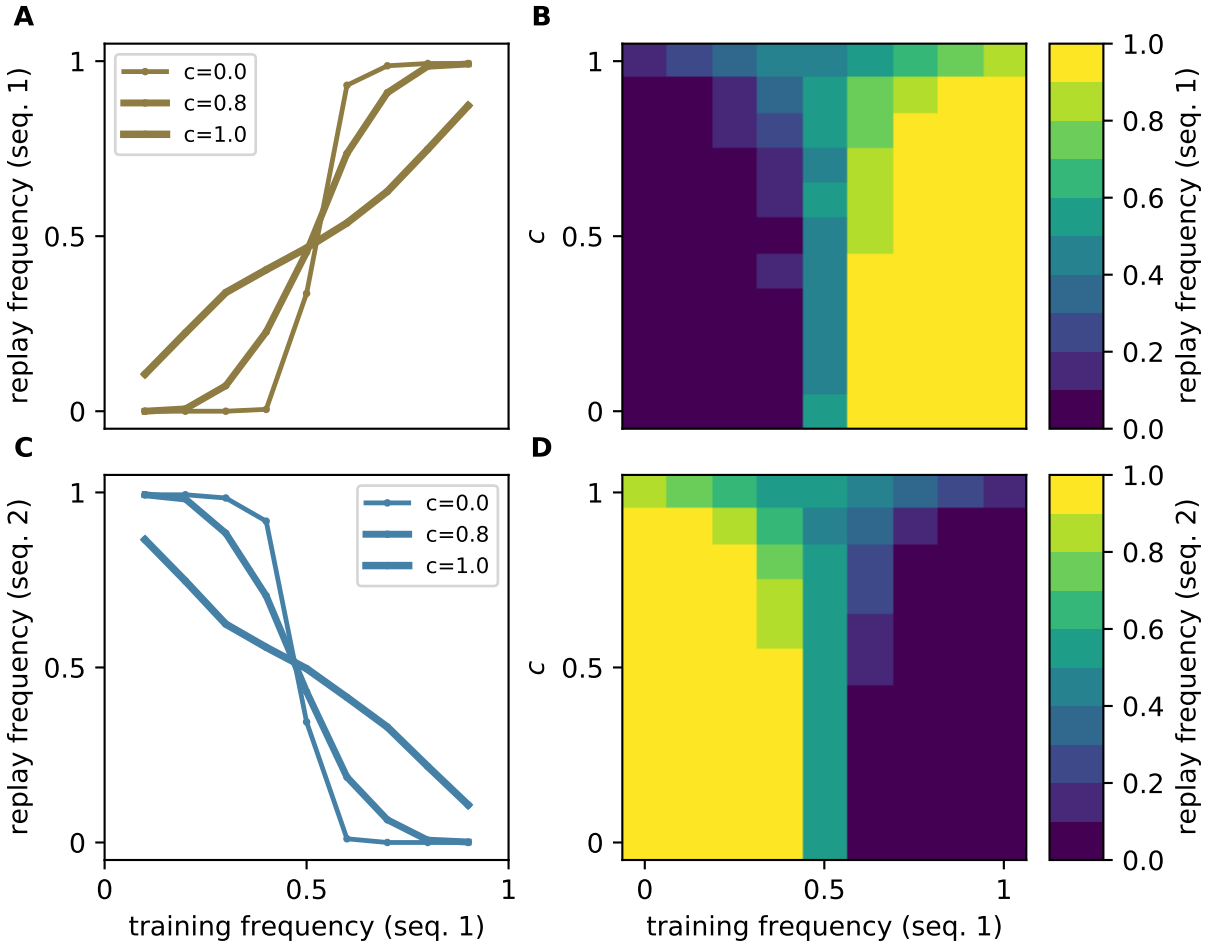

Figure S1: Adjusting level of correlation permits different replay strategies. The relative replay frequencies of sequence 1 (brown, **A**) and 2 (blue, **C**) as a function of the occurrence frequency of sequence 1 during training plotted for different correlation levels (different markers: $c = 0$, $c = 0.8$, and $c = 1$). Contour plot showing the dependence of the relative replay frequencies of sequences 1 (**B**) and 2 (**D**) on the training frequency of sequence 1 and different correlation levels. Noise amplitude $\sigma = 15\,\text{pA}$. The replay frequencies are computed as the mean across $N_\text{t} = 150$ trials, averaged across 5 different network realizations. See Tab. 2 for remaining parameters.