


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.


In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Attractor concretion as a mechanism for the formation of context representations

Mattia Rigotti^{a,b}, Daniel Ben Dayan Rubin^{a,b}, Sara E. Morrison^a, C. Daniel Salzman^{a,c,d,e,f,g}, Stefano Fusi^{a,b,*}

^a Department of Neuroscience, Columbia University College of Physicians and Surgeons, New York, NY 10032-2695, USA

^b Institute for Neuroinformatics, University of Zürich, and ETH Zürich, 8057 Zürich, Switzerland

^c Department of Psychiatry, Columbia University Medical Center, New York, NY 10032, USA

^d W. M. Keck Center on Brain Plasticity and Cognition, Columbia University Medical Center, New York, NY 10032, USA

^e Kavli Institute for Brain Sciences, Columbia University Medical Center, New York, NY 10032, USA

^f Mahoney Center for Brain and Behavior, Columbia University Medical Center, New York, NY 10032, USA

^g New York State Psychiatric Institute, New York, NY 10032, USA

ARTICLE INFO

Article history:

Received 31 August 2009

Revised 9 January 2010

Accepted 14 January 2010

Available online 25 January 2010

Keywords:

Reinforcement learning

Computational methods

ABSTRACT

Complex tasks often require the memory of recent events, the knowledge about the context in which they occur, and the goals we intend to reach. All this information is stored in our mental states. Given a set of mental states, reinforcement learning (RL) algorithms predict the optimal policy that maximizes future reward. RL algorithms assign a value to each already-known state so that discovering the optimal policy reduces to selecting the action leading to the state with the highest value. But how does the brain create representations of these mental states in the first place? We propose a mechanism for the creation of mental states that contain information about the temporal statistics of the events in a particular context. We suggest that the mental states are represented by stable patterns of reverberating activity, which are attractors of the neural dynamics. These representations are built from neurons that are selective to specific combinations of external events (e.g. sensory stimuli) and pre-existent mental states. Consistent with this notion, we find that neurons in the amygdala and in orbitofrontal cortex (OFC) often exhibit this form of mixed selectivity. We propose that activating different mixed selectivity neurons in a fixed temporal order modifies synaptic connections so that conjunctions of events and mental states merge into a single pattern of reverberating activity. This process corresponds to the birth of a new, different mental state that encodes a different temporal context. The concretion process depends on temporal contiguity, i.e. on the probability that a combination of an event and mental states follows or precedes the events and states that define a certain context. The information contained in the context thereby allows an animal to assign unambiguously a value to the events that initially appeared in different situations with different meanings.

© 2010 Elsevier Inc. All rights reserved.

Introduction

When we execute complex tasks, we often need to store information about past events in order to decide how to react to a particular stimulus. If the task is familiar, we know what information to store and what to disregard. At the moment we make a decision about our response to a particular event, we are in a specific mental state that contains all the information that we know to be relevant to react to that event. This information is typically about our perception of the environment, our physical position, our memories, our motivation, our intentions, and all the other factors that might be relevant to reach a particular goal. In other words, every mental state is our most general disposition to behavior. In many cases the execution of a task can be considered as a series of transitions from

one mental state to the next, each triggered by the occurrence of a particular event.

In order to understand the neural mechanisms underlying the execution of complex tasks we need to answer two important questions: (1) How do we create the neural representations of mental states that contain all the relevant information to execute a task? (2) How do we learn which mental state to select in response to a particular event? Reinforcement learning (RL) algorithms (see e.g. (Sutton and Barto, 1998)) have provided an elegant theoretical framework to answer the second question. In particular, they provide prescriptions for the policy of mental state selection that maximize reward and minimize punishment. In RL algorithms, values that represent future cumulative reward are assigned to the mental states. The value increases as the agent moves closer to a pleasant outcome, like the delivery of reward. The table of values of mental states thereby determines the optimal policy. One has simply to select the action that induces a transition to the state with highest value. However, most RL algorithms presuppose that the set of mental states contain all the relevant information for performing a task, and hence

* Corresponding author. Department of Neuroscience, Columbia University College of Physicians and Surgeons, New York, NY 10032-2695, USA.

E-mail address: sf2237@columbia.edu (S. Fusi).

they do not provide an answer to the first question, on how the mental states are created.

In this paper, we make an attempt to answer this question by proposing a mechanism for the creation of mental states in context-dependent tasks, in which the optimal policy for maximizing reward is different in different contexts. In particular we will consider all the situations in which the information about temporal context can be used to create mental states that unequivocally determine the state of the environment and the actions to be executed. In other words, if the occurrence of an event in two different contexts requires different policies, we need to react to that event in two different ways that will be encoded in two different sets of mental states. We propose in our paper a mechanism that leads to the formation of different sets of mental states for different contexts.

A paradigmatic experiment

To illustrate the principles behind our proposed mechanism, we will present a neural network model that performs an extended version of an appetitive and aversive trace conditioning task used in recent neurophysiological recording experiments (Paton et al., 2006; Belova et al., 2007; Salzman et al., 2007; Belova et al., 2008; Morrison and Salzman, 2009). In those experiments, monkeys learned whether an abstract fractal image (conditioned stimulus, CS) predicted liquid reward or aversive air-puffs (unconditioned stimulus, US) after a brief time (trace) interval. Single unit recordings in the amygdala and orbitofrontal cortex (OFC) revealed the existence of cells encoding the learned value of the CSs (rewarded or punished). After a variable number of trials, the CS-reinforcement contingencies were reversed and monkeys had to learn the new contingencies. In the experiments, the CS–US associations were reversed only once. However, in principle, the two contexts defined by the sets of CS–US associations could be alternated multiple times. In this situation, it is possible that the animal at some point creates two representations corresponding to the two contexts and it can switch rapidly from the optimal policy for one context to the optimal policy for the other. This switch is qualitatively different from learning and forgetting the associations as it would not require any synaptic modification. The two independent context representations would be simultaneously encoded in the pattern of synaptic connectivity.

The model of a neural circuit performing the trace conditioning task with multiple reversals will be used to illustrate the mechanism for the formation of context representations. The single unit recordings from experiments (in which there is a single reversal), will be used to support our assumptions about the initial response properties of the model neurons.

The proposed model architecture: the Associative Network (AN) and the Context Network (CN)

In modeling data from this task, we assume that there are two interacting neural circuits. The first one is a neural circuit that we name AN (Associative Network) similar to the one proposed by (Fusi et al., 2007), that learns simple one-to-one associations between CSs and USs. The second one, the CN (Context Network) observes the activity of the AN, in particular when the AN has already learned the correct associations, and abstracts the representations of temporal contexts. These representations can then be used by the AN to predict more efficiently the value of the CSs when the context changes and the AN has to learn the new associations.

Learning and forgetting associations: the function of the AN

In every trial the AN starts from a "wait state." The CS biases a competition between the populations of neurons representing two different mental states, one that predicts the delivery of reward, and

hence has a positive value, and the other that predicts punishment and has a negative value. The delivery of reward or punishment resets the AN to the wait state. The AN encodes the CS–US associations by making CS triggered transitions to the state that represents the value of the predicted US. The CS–US associations are learned by biasing the competition between the positive and the negative state. In particular, the competition bias is learned by modifying the synaptic connections from the neurons that represent each CS and the positive and negative value coding neurons. When the associations are reversed, the learned synaptic strengths are overwritten by the new values. The AN can be in one of the three states (wait, positive, negative), and can implement only one set of CS–US mappings at a time. In the two contexts the AN implements two different "policies," as the same CS induces different transitions. Notice that the CS–US associations are learned independently for each CS. The AN does not store any information about the relations between different CS–US associations, and in particular about the fact that all associations are simultaneously modified when the context changes. This means that, for example, when the context changes, the AN cannot infer from the modification of one of the CS–US associations, the value of the other CS. This type of inference requires information about the temporal statistics of the CS–US associations which are collected by the CN. We propose that this type of inferential information is stored in the representations of the temporal contexts, which are built from the statistics of the sequence of events and mental states.

The formation of representations of temporal contexts: the main idea

In the trace conditioning task, the first context is characterized by the fact that the sequence CS A-Reward is most often followed either by itself or by CS B-Punishment. Analogously the second context is defined by the elevated probability of the transitions between CS A-Punishment and CS B-Reward. The animal observes rarely that CS A-Reward is followed by CS B-Reward or CS A-Punishment. In other words, if we look at the matrix of transitions between these sequences, we can clearly identify two clusters of sequences that are connected by significantly larger transition probabilities. These two clusters define the two relevant contexts. The idea sounds simple, but the detailed implementation turned out to be more difficult than expected because initially, the neural circuit does not know that it needs to consider the CS–US associations as the building blocks of the context representation. The neural circuit observes a series of several events and mental states and it has to abstract autonomously what is relevant for the formation of context representations.

The neural basis of the formation of context representations

In the detailed implementation, the learning process that takes place in the CN iteratively merges the neural representations of temporally contiguous events and mental states. The first compounds that are created represent short temporal sequences, or, more generally, groups of events and mental states that tend to be temporally contiguous. Compounds that often follow each other can also merge into larger compounds. In this sense the process of merging is iterative, and at every iteration the compounds grow to represent larger groups of events that are temporally contiguous to each other. In more technical terms, the temporal statistics of the compounds define a new Markov process that is a coarse grained representation of the Markov process of the previous iteration, similarly to what has been studied in (Weinan et al., 2008). The iterative process stops when all the transition probabilities between the compounds go below a certain threshold. In the specific case of the trace conditioning task, the first representations that merge are those that represent temporally contiguous events like CS A and the mental state predicting a positive value. The parameters are chosen so that the merging process stops when CS A-Reward and CS B-Punishment

belong to a single compound representing context 1, and CS B-Reward and CS A-Punishment belong to a second compound that represents context 2. The neural mechanisms that could underlie this iterative learning process are described in detail in the Methods, and the simulations are illustrated in the Results.

The initial conditions: mixed selectivity in theory and experiments

The neurons of the CN are randomly connected to the neurons of the AN and the parameters are chosen such that they respond to conjunctions of external events (CS A or B, reward, punishment) and states of the AN (neutral, positive, negative). The neurons therefore may be described as having mixed selectivity, even before the learning process of the CN starts. We observed these types of neurons both in the amygdala and in orbitofrontal cortex of monkeys performing the trace conditioning task of Paton et al. (2006). We report in this manuscript the statistics of their response properties.

From temporal sequences to context representations: attractor concretion

In the model, the temporal statistics of the mixed selectivity neurons depend on the sequence of patterns of activation of the AN. For example consider a trial in which CS A is followed by reward. The AN would start in a wait state with neutral value and CS A would steer the activity toward a positive state. The US would reset the AN activity back to the neutral state. In the CN we would observe the activation in sequence of the following populations that are selective to conjunctions of states and events: neutral-CS A, CS A-positive, positive-reward, reward-neutral. Initially, each conjunction induces a transient activation of the CN neurons. The synapses between CN neurons that are activated simultaneously are strengthened (Hebbian component of synaptic plasticity), so that the neural representations of individual conjunctions become stable self-sustaining patterns of persistent activity that are attractors of the neural dynamics (see e.g. (Hopfield, 1982; Amit, 1989)). These patterns remain active until the occurrence of the next event and the activation of a new input from the AN. A second component of the synaptic plasticity, which we call temporal sequence learning (TSL), strengthens the connections between neurons that are activated in succession, similar to what has been proposed for learning of temporal sequences (Sompolinsky and Kanter, 1986) and temporal contexts (Brunel, 1996; Griniasty et al., 1993; O'Reilly and Munakata, 2000; Rougier et al., 2005; Yakovlev et al., 1998). This component causes the merging (concretion) of attractors that are activated in a fixed temporal order, leading to the formation of the representations of temporal contexts.

From context representations to the creation of mental states

At the beginning of the learning process, the CN simply reflects the activity of the AN, and hence the entire AN-CN system has the same number of mental states as the sole AN (neutral, positive, negative). At the end of the learning process, the CN can represent both contexts, with one being the “active context.” Hence, the entire AN-CN system has two sets of the three AN states, one for the first context and one for the second. As the AN receives feedback from the CN, it can then easily disambiguate between the CS-US associations of the first context from those of the second. We will show that at the end of the learning process the full AN-CN system can work more efficiently than the AN alone after a context switch. Indeed, it can predict the correct value of a CS when the other is already known. We will then discuss quantitative predictions about the behavior and the neural activity that can be recorded.

Materials and methods

We first describe the details of the trace conditioning task that has been used in neurophysiological experiments and its extended version, which we used in all model simulations. We then describe

the Associative Network (AN) and the Context Network (CN). The model of the AN has already been introduced in Fusi et al. (2007). Here we summarize briefly its neural and synaptic dynamics and we show simulations of the specific case of the trace conditioning task. We then describe the novel neural and the synaptic dynamics of the CN. The description of the learning behavior and the explanation of the mechanisms are deferred to Results. Finally, we describe the details of the analysis of neurophysiological data that we use in Results to motivate the model assumptions about the initial response properties of the neurons.

The experimental protocol

The appetitive and aversive trace conditioning task described in Paton et al. (2006), Belova et al. (2007), Salzman et al. (2007), Belova et al. (2008), Morrison and Salzman (2009) uses a trace conditioning protocol (a type of Pavlovian conditioning) to train an animal to learn the relationship between abstract visual stimuli (conditioned stimuli, CS) and rewards and punishments (unconditioned stimuli, US). While a monkey centers its gaze at a fixation point, CSs are presented followed by a trace interval (a brief temporal gap), and then US delivery. In the experiments (Paton et al., 2006; Belova et al., 2007; Salzman et al., 2007; Belova et al., 2008; Morrison and Salzman, 2009), the animals demonstrated their learning of CS-US contingencies by licking a reward tube in anticipation of reward and blinking in anticipation of the delivery of the air-puff. After a variable number of trials, the CS-US associations were reversed without any notice and the animals had to learn the new contingencies.

In the version of the task that is simulated in this paper, we will consider two CSs, A and B, and two possible outcomes, reward and punishment, which are delivered after a brief time interval, as in the original task of (Paton et al., 2006). The CS-US associations are reversed multiple times, switching from context 1 in which CS A is paired with reward, and CS B with punishment, to context 2 in which CS A is paired with punishment and CS B with reward. The blocks of context 1 and context 2 trials are approximately 120 trials each, and they are alternated multiple times.

The Associative Network (AN): structure and function

The AN learns the associations between CSs and USs. This Associative Network (AN) receives feed-forward plastic inputs from the neural populations that represent external events (CSs, reward and punishment). The neurons of the AN are grouped in three different populations: two excitatory populations representing positive and negative value compete through a population of inhibitory neurons (see Fig. 1). The synaptic connections between these three populations are chosen as in Wang (2002) and Fusi et al. (2007), so that there are only three stable states: a wait state in which the AN is quiet (neutral value state), and the other two corresponding to the activation of one of the two excitatory populations (positive and negative state). The presentation of the CS generates an input that initiates and biases the competition between the positive value and the negative value coding populations of neurons. The delivery of the US brings the network back to the neutral value state (see Fig. 2 for a description of the AN dynamics). Initially, when the CS is novel and the associated US is still unknown, the CS activates an unbiased competition between positive and negative value coding populations and one of the two values is chosen randomly with equal probability. This behavior reflects the fact that the animal is already familiar with the experimental protocol and can predict that the CS will be followed in half of the cases by reward and in the other half by punishment (Fusi et al., 2007). The synaptic weights between the neurons encoding the CS and the AN neurons are modified in every trial depending on whether the prediction of the AN (positive or negative) matches the actual US that is delivered (reward or punishment) or

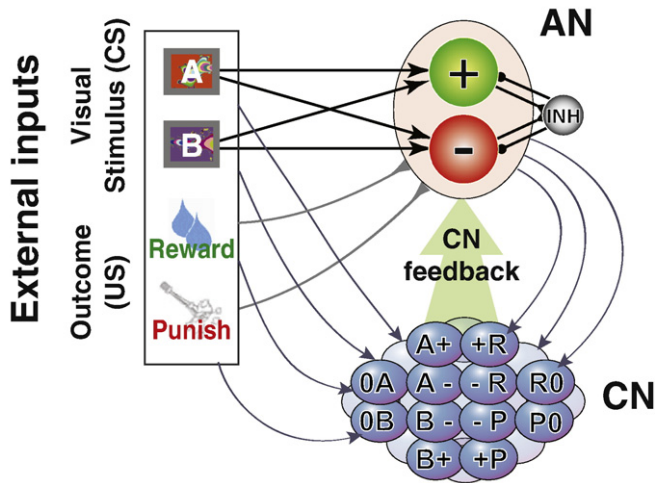


Fig. 1. The two networks of the simulated neural circuit: the Associative Network (AN, top right), and the Context Network (CN, bottom right). The AN and the CN receive inputs from the neurons encoding external events (conditioned and unconditioned stimuli). The AN network contains two populations of neurons, +, -, that encode positive and negative values respectively. These neurons are activated by external events (CSs) in anticipation of reward and punishment. The inhibitory population (INH) mediates the competition between the two populations. The connections from the CS neurons to the AN neurons are plastic and encode the associations between the CS and the predicted US. The CN neurons receive fixed random synaptic connections from both the AN and the external neurons. The neurons in the CN respond to conjunctions of external events and AN states and they are labeled accordingly. The recurrent connections within the CN are plastic and they are modified to learn context representations. After learning, the CN neurons encode the context, and they project back to the AN (described later, in Fig. 4).

not. When the prediction is correct (i.e. when the CS activates the positive state and it is followed by reward, or when the CS activates the negative states and it is followed by punishment), the synapses from the CS neurons to the correct value encoding population are strengthened (learning rate $q_+^R = 0.042$, see Fusi et al. (2007)) for the details of the synaptic dynamics), and the synapses from the CS neurons to the other value population are weakened ($q_-^R = 0.073$). These modifications reinforce the bias in the competition between positive and negative populations towards the correct prediction. When the prediction is incorrect, we assume as in Fusi et al. (2007) that all synapses from the CS neurons to the AN neurons are rapidly depressed ($q^{NR} = 0.99$). This reset forces the AN to choose randomly the value of the CS the next time it is presented. The learning rates q_s have been chosen to match the learning curves reported in Paton et al. (2006). In particular we chose them so that the simulated AN reaches 90% of the value prediction performance in 10 trials on average.

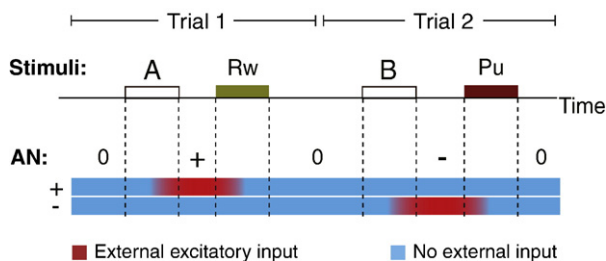


Fig. 2. Simulated activity of the AN during two trials of the trace conditioning task of Paton et al. (2006). During the first trial CS A is presented, followed by a reward. The AN network is initially in the neutral state “0” in which all populations are inactive (the activity is color coded: blue means inactive, red means active). The presentation of CS A initiates a competition between the positive coding AN population “+” and the negative coding population “-” which, in this simulation, ends with the activation of population “+”. The delivery of reward resets the AN to the “0” state. In the second trial CS B activates population “-” and punishment resets it.

The Context Network (CN): the architecture

The Context Network is made of neurons that are randomly connected to both the neurons encoding the external events and the excitatory neurons of the AN. The synapses between CN neurons are plastic and they are modified by the learning rule described below in order to create context representations. The random connections are Gauss distributed and the parameters are chosen so that a large number of the CN neurons respond to conjunctions of external events (CS A, CS B, reward, punishment, denoted by A, B, R, P, respectively) and the state of the AN (positive, negative, neutral, denoted by +, -, 0, respectively). For simplicity, we will consider in our simulations only the neurons that respond to simple conjunctions of one external event and one state. For example, some neurons would respond to CS B only when the AN switches to a negative state. We will label these neurons with “B-.” In Fig. 3 we show the simulation of a rate model neuron that behaves like a typical CN neuron with mixed selectivity. These simulations are for illustrative purposes only and to motivate our assumptions about the response properties of the CN neurons. This type of model neuron will not be used in the rest of the paper (see Materials and methods about the neural dynamics for the model neurons that will be used). The simulated neuron receives synaptic inputs with equal weights from the neurons that are activated by CS B and from the negative value coding neurons of the AN (simulated as in Fusi et al. (2007)). The firing rate is a sigmoidal function of the total synaptic input. Although the choice of equal synaptic weights might seem special, the behavior illustrated in Fig. 3 is actually the typical behavior of neurons with randomly chosen synaptic weights. The probability that a randomly connected neuron exhibits significant mixed selectivity depends on the specific model of the neuron and on the neural representations of the events, but it can be as large as 1/3 (Rigotti et al., submitted for publication).

The neurons with mixed selectivity are the building blocks of context representations and they are assumed to be present in the CN from the very beginning, before any learning process takes place. In Results, we will support this assumption with experimental data.

The neural dynamics of the CN

The activity of the AN drives the CN network, which is composed of N populations of neurons which are either active or inactive. We denote by ξ_i the activity of population i . As we consider only the randomly connected neurons that respond to simple conjunctions of external events and state of the AN, we have in total $N = 12$ different types of populations responding to the following combinations: 0A, 0B, A+, A-, B+, B-, R0, P0, +R, -R, +P, -P. Not all these combinations are necessary for the formation of context representations, but we simulate all of them for completeness, as they are all assumed to be present in the neural circuit. For simplicity we ignore neurons that respond to combinations of three or more events and AN states. We assume full connectivity within the CN. The average strength of the excitatory synaptic connections from neurons of population j to neurons of population i is denoted by J_{ij} , and it can vary from 0 to 1. Every CN population inhibits all the others through a pool of inhibitory neurons. The net effect is to include a constant negative coupling ($g_I = 0.5$) between the excitatory populations. Additionally, every population i receives a constant current θ_i and an external input h_i from the AN neurons. More quantitatively, the activity ξ_i^t of population i at time t is given by the following discrete time evolution equation:

$$\xi_i^t = \theta \left(\frac{1}{N} \sum_{j=1}^N (J_{ij} - g_I) \xi_j^{t-1} + h_i^t + \theta_i \right), \quad i = 1, \dots, N, \quad (1)$$

where θ is the Heaviside function with $\theta(x) = 1$ if $x > 0$, and $\theta(x) = 0$ otherwise (in our case $\theta_i = 0$).

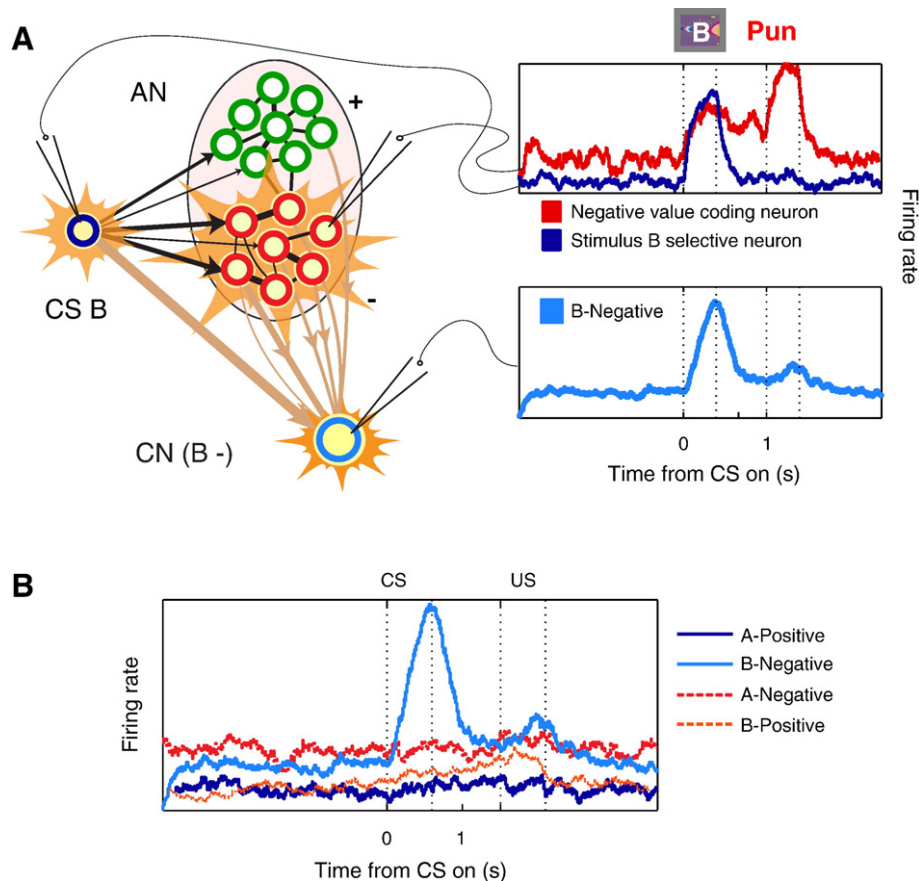


Fig. 3. Illustrative firing-rate simulations of a typical CN neuron which exhibits mixed selectivity to the conjunction of an external event (CS B) and an AN value state (negative). (A) The top panel shows the firing rate as a function of time for two simulated neurons in response to CS B. The blue trace represents the response of an external neuron encoding CS B, which is flat until the presentation of visual stimulus B. The red trace represents the response of a negative value coding neuron of the AN. CS B is already familiar and its value is correctly predicted by the AN. When CS B is shown, the negative value population is activated, and it remains active until the delivery of the US. In the bottom panel, we show the activity of a CN neuron that, by chance, is strongly connected to CS B external neurons and to the negative value coding AN neurons. The response is significantly different from spontaneous firing rate only when CS B is presented and the negative value AN state wins the competition. (B) Mixed selectivity to CS B and negative value. The cell is selective to both the value and the identity of the CS as the neuron responds only to the CS B-Negative combination and not to the other combinations (CS A-Positive, CS A-Negative, CS B-Positive).

During the learning process, the single or multiple CN populations can become attractors of the neural dynamics (stable patterns of self-sustaining activity). These patterns will be the building blocks for context representation, as illustrated in the Results. Once activated by the AN and external input, the attractors remain active indefinitely, or at least until the arrival of the next sufficiently strong input. To avoid the simultaneous activation of all CN populations, we need to guarantee that the external input can overcome the recurrent CN input and shut down previously activated stable patterns. This might be important also for weak external inputs that would normally be ignored by the CN. For this reason we complemented the described neural dynamics with a reset signal that inhibits the whole network every time an external input h^t targets at least one population which is not already activated by the recurrent input. Such a signal is important not only to reset the activity but also to learn only when the activity pattern of the CN is modified (see next section). In the time step that follows the reset signal, the activity of the network is entirely determined by the external input h .

The synaptic dynamics of the CN

The CN “observes” the activity of the neurons representing the external stimuli and the neurons of the AN. The context representations are created from the temporal statistics of the patterns of activity of the AN and the external neurons. Here we describe the equations and the details of the synaptic dynamics that lead to the formation of

the context representations, but we explain the mechanism and we show simulations only in Results.

The synapses are modified by two mechanisms: (1) a Hebbian mechanism strengthens the synapses of simultaneously active neurons, and depresses the synapses connecting an active to an inactive neuron. Analogously to the mechanism introduced in Hopfield (1982) and, more recently, in Amit and Brunel (1995), it stabilizes the CN activity that is initially imposed transiently by the external input and the AN. If the synapses between co-activated neurons become strong enough, the neurons of the activated CN population can excite each other to the point that the transient activity becomes self-sustaining (attractor of the neural dynamics). (2) the TSL (Temporal Sequence Learning) mechanism, which links together patterns of activity that are activated in sequence. This component of the synaptic dynamics is responsible for merging attractors that are often temporally contiguous. It basically strengthens the synapses between two neurons that are activated sequentially one after the other. Moreover, it depresses the synapses between active neurons and neurons that are inactivated at the next time step.

Both mechanisms are activated only when the competition between the positive and negative populations in the AN is strongly biased, indicating that the AN has already learned the associations. As the neurons have only two levels of activation, we monitor the total synaptic input to them, and we modify the synapses of the CN only when the current driving the winning population exceeds a threshold $\theta_L = 0.25$ (to be compared to the total synaptic input, i.e. the argument

of the Heaviside function in Eq. (1)). In a more realistic implementation with rate neurons, we could set a threshold for the firing rate of the AN neurons.

The Hebbian mechanism

The modifications of the synapses from population j to population i depend on the current pre and post-synaptic activity ξ_{ij}^t and on the post-synaptic recurrent synaptic input I_i^t (i.e. the input from the neurons that belong to other populations within the CN):

$$I_i^t = \frac{1}{N} \sum_{j=1}^N (J_{ij} - g_i) \xi_j^t + \theta_i. \quad (2)$$

In particular we have:

$$\Delta J_{ij}^s = \theta(\gamma^s - \xi_i^t I_i^t) \left[q_+^s (1 - J_{ij}) \xi_i^t \xi_j^t - q_-^s J_{ij} (1 - \xi_i^t) \xi_j^t \right]. \quad (3)$$

Eq. (3) describes a modified version of the perceptron learning algorithm (Rosenblatt, 1958). The synapse J_{ij} is potentiated when both the pre- and the post-synaptic neuron are simultaneously active ($\xi_i^t \xi_j^t$) and depressed when the pre-synaptic neuron is active and the post-synaptic neuron is inactive (the $(1 - \xi_i^t) \xi_j^t$ term). The two terms containing J_{ij} , and $(1 - J_{ij})$ impose a soft-bound on the synaptic weights and keep them between zero and one, as in Senn and Fusi (2005). The synapses are not updated when the post-synaptic neuron is active and the total recurrent input is sufficiently large ($\theta(\gamma^s - \xi_i^t I_i^t)$, where γ^s is a positive number which is related to the stability parameter (Gardner, 1987)). This term prevents the synapses from being updated when the recurrent input is sufficient to activate the post-synaptic neuron in the absence of the external stimulus. In other words, the synapses are not updated if the pattern of activity can already sustain itself. This term prevents the correlated parts of different attractors from dominating the neural dynamics and hence allows the network to generate attractors that are correlated (see e.g. Senn and Fusi (2005)). Moreover, if γ_s is sufficiently large, it increases the stability of the attractors (Krauth et al., 1988).

This type of learning prescription can be implemented with realistic spike-driven synaptic dynamics (Brader et al., 2007). The factors q_+^s and q_-^s are the learning rates for potentiation and depression, respectively. The parameter values are: $\gamma^s = 5 \times 10^{-4}$, $q_+^s = 7.5 \times 10^{-2}/n$ $q_-^s = 15 \times 10^{-2}/n$, where n is the number of time steps in one trial (in our simulations $n = 15$).

Temporal sequence learning (TSL)

The second learning component, temporal sequence learning (TSL), is meant to strengthen the synaptic connections between neurons that are repeatedly activated one after the other in a fixed temporal order. At every time step t we calculate how many inactive populations are activated by the new incoming external input h^t and we divide this quantity by the total number of populations N :

$$\Delta^t = \frac{1}{N} \sum_{i=1}^N \theta(h_i^t) \theta(1 - \xi_i^{t-1}). \quad (4)$$

This is a measure of the global mismatch between the pattern imposed by the external stimulation at time t and the network activity at the previous time $t - 1$. When this quantity is different from zero, it is an indication that the neural activity of the CN has been modified by the external input. In such a case a reset signal is delivered to the CN (see neural dynamics), and the synaptic connections from population

j and population i are modified according to:

$$\Delta J_{ij}^a = \Delta^t \theta(\gamma^a - \theta(h_i^t) I_i^{t-1}) \cdot \left[q_+^a (1 - J_{ij}) \xi_j^{t-1} \theta(h_i^t) - q_-^a J_{ij} \xi_j^{t-1} (1 - \theta(h_i^t)) \right], \quad (5)$$

The term in square brackets contains two terms: one potentiates the synapses and the other one depresses them. In particular, the synapses are potentiated when the post-synaptic external current h_i^t is positive and the pre-synaptic neuron was active at the previous time step (ξ_j^{t-1}). The synapses are depressed when the post-synaptic external current h_i^t is negative, and the pre-synaptic neuron was active at the previous time step. The J_{ij} dependent terms implement a soft boundary as in the case of the Hebbian term. The presence of a soft boundary is in general important to estimate probabilities (Rosenthal et al., 2001; Fusi et al., 2007; Soltani and Wang, 2006), and in our specific case to estimate the probability that a particular event is followed by another one. The parameters are: $\gamma^a = 0$, $q_+^a = 1.0$ and $q_-^a = 7.5 \times 10^{-2}$.

At every time step the synaptic weights are updated according to:

$$J_{ij} \leftarrow J_{ij} + \Delta J_{ij}^a + \Delta J_{ij}^s. \quad (6)$$

When the reset signal is delivered, first the weights are updated and then the activity is reset.

The feedback from the Context Network to the Associative Network

The information about the current context contained in the CN activity after learning can be used by the AN to predict more efficiently the value of the stimuli. For example, when the CN-AN knows that the CS-US association has changed for CS A, it can predict that the CS-US association for CS B has changed as well. In order to do so, we need to introduce some form of feedback from the CN to the AN. In principle, CN neurons could project directly to the AN neurons, as the CN neurons contain all the information that the AN neurons need to know about the current context. However, this feedback input cannot produce the desired context-dependent bias on the AN competitive dynamics unless we introduce an intermediate population of neurons that mixes the external input and the CN activity (see Fig. 4). This is a general requirement for many systems in which there is a dependence on context (Rigotti et al., submitted for publication). Indeed, without this intermediate layer of neurons, there is no set of synaptic weights that would produce the correct prediction. The general proof is in Rigotti et al. (submitted for publication); here we give an intuitive argument for our specific case. Consider two input neurons: an external neuron that is active when CS A is presented and inactive for CS B, and a CN neuron that is active for context 1 and inactive for context 2. The AN "output" neuron encoding a negative value should be inactive for CS A+Context 1 (both input neurons active = 11), and CS B+Context 2 (00). At the same time it should be inactive for CS A+Context 2 (10), and CS B+Context 1 (01). This mapping of the input to the output is equivalent to the implementation of the logical operation "exclusive or" (XOR) and it is known that it is not possible to build a single layer network that implements it (see e.g. Minsky and Papert (1969)). The solution proposed in Rigotti et al. (submitted for publication) is to introduce an intermediate layer of randomly connected neurons. If the number of these neurons is sufficiently large, the problem equivalent to the XOR explained above can be solved. For this reason we introduced an additional population of neurons whose activity depends on the input from the CN populations and the external neurons. Analogously to the CN neurons, which are also connected by random synaptic weights to the AN, most of the neurons of the additional population respond to pairs of CN-external neurons activations. We therefore assume that the feedback population is composed of neurons responding to the

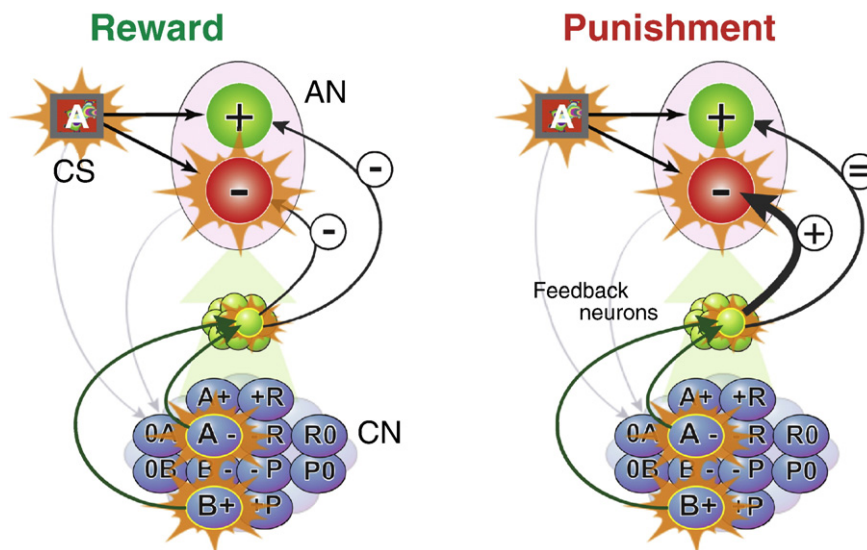


Fig. 4. The learning dynamics of the CN to AN feedback. This signal is mediated by a layer of feedback neurons selective to the CN state, encoding the context, and the external input activity. The synapses connecting the feedback neurons to the AN are modified with learning dynamics similar to the ones used for the AN synapses (see Fusi et al. (2007) and the description of the AN dynamics in Materials and methods). A circled + indicates that the synapses are potentiated. Synapses marked = remain unchanged, and \$- indicates synaptic depression. Left panel: the AN-CN system predicts a negative value for CS A, which instead is followed by Reward (context 1). The synapses connecting active feedback neurons to the AN are all depressed. Right panel: the AN-CN system correctly predicts a negative value (context 2) which is followed by punishment. The synapses from active feedback neurons to the active AN neurons are potentiated.

$2 \times N$ possible CS-CN population combinations (2 CSs multiplied by the N populations of the CN). These neurons project back to the AN neurons with plastic synapses that are modified with the same synaptic dynamics as the connections from the external neurons to the AN neurons, except that the learning rates are significantly smaller ($q_+^R = 2 \times 10^{-4}$, $q_-^R = 0$, $q_-^L = 4 \times 10^{-4}$). These feedback connections are initialized to zero, and the learning rates are chosen to be small, so that the synaptic input from the CN affects the AN dynamics only at a late learning stage, when the CN context attractors are formed and stable. The AN sees the information about the current context coming from the CN as an additional input that would operate in the same way as a constantly present contextual cue.

We stop modifying the CN synapses when the feedback input becomes too strong compared to the external input. This prevents the CN from learning from its own activity, with the danger of effects that are difficult to control. For example, if the CN learns rapidly one of the two contexts of the trace conditioning task, and it starts dominating the AN behavior, then it becomes difficult to create the representation of the second context because the AN will not have a chance to learn the CS-US associations of the second context. Indeed, it will constantly be driven by the CN, which represents and will continue to represent only the first context. In the simulations we block the CN learning dynamics when the synaptic input to the AN coming from the CN feedback is more than 2.5 times larger than the direct feedforward external input.

The analysis of recorded OFC and amygdala cells

Our assumption that the neurons of the CN are initially selective to conjunctions of external events and AN mental states (mixed selectivity) is supported by the analysis of neurophysiological recordings. We analyzed the cells recorded during the trace conditioning task with a single reversal described at the beginning of the Methods. It is reasonable to assume that this situation (i.e., the single reversal) reflects what happens in the initial or early stages of the learning process that leads to the formation of context representations. Analyses were performed on spike data from two time intervals during the trial: the CS interval (90–440 ms after image onset for monkey L; 90–390 ms after image onset for monkey R) and the trace interval (90–1500 ms after the image turned off). These time intervals were chosen because more than 90% of visual response latencies exceeded 90 ms as established by an

analysis of latencies, described previously (Paton et al., 2006; Belova et al., 2007, 2008; Morrison and Salzman, 2009).

In order to determine the degree to which neural responses are modulated by reinforcement contingencies (image value) or by the sensory characteristics of the CSs themselves, we performed a two-way ANOVA with image value and image identity as main factors. The ANOVA was performed separately on spike counts from the CS and trace intervals for each cell, as cells could encode image value at different times during the trial. If there was a significant effect of image value in either or both intervals ($p < 0.01$), the cell was classified as value-coding. We found a few cells that had opposite image value preferences in the CS and trace intervals, and these were excluded from further analysis. Neurons in OFC and the amygdala that were categorized as “non-value-coding” exhibited a variety of responses to conditioned stimuli; these included neural responses that were similar for all conditioned images, as well as responses that were strongest (or weakest) for the stimulus associated with a weak reward. In addition, a substantial proportion of OFC and amygdala neurons, both value-coding and non-value-coding, showed a significant main effect of image identity in the ANOVA, or an interaction effect of image value and image identity ($p < 0.01$).

We performed an additional analysis to determine whether in the trace conditioning task with a single reversal (presumably the situation preceding the formation of context representations) there are cells that encode the context in the same way as the CN neurons at the end of the learning process. As we will see in the simulation results, the CN neurons that represent the context after learning are selective to the context in every interval of the trial (in the presence or in the absence of external events like the CS or the US). In particular, their activity should be significantly different in the two contexts for all the individual CS-US associations. In practice, in the specific case of the trace conditioning task, there should be a threshold that separates the activity recorded in CS A-Positive and CS B-Negative trials (context 1) from the activity recorded in CS A-Negative and CS B-Positive trials (context 2). Moreover, the difference between context 1 and context 2 activities should be significant. We imposed these conditions by considering all pairs of CS-US combinations. In particular, in order to meet the criterion for a “context cell,” the average activity of the neuron in the CS A-Positive trials (μ_{A+}) had to be significantly different ($p < 0.05$, t-test) from the average activity in the CS A-Negative trials (μ_{A-}). Moreover, we

required that the differences $\mu_{A+} - \mu_{B+}$, $\mu_{B-} - \mu_{A-}$ and $\mu_{B-} - \mu_{B+}$ be significant. Additionally, we required that all four differences $\mu_{A+} - \mu_{A-}$, $\mu_{A+} - \mu_{B+}$, $\mu_{B-} - \mu_{A-}$ and $\mu_{B-} - \mu_{B+}$ have the same sign (we always subtract context 2 (reversal) epochs from context 1 (initial) epochs). A cell qualified as a context cell if it satisfied all these criteria at least in one of the two analyzed intervals (the interval during CS presentation and the trace interval).

Results

We present the results as follows: we first explain the assumptions of the model about the mixed selectivity of the neurons, and we provide experimental evidence to support them. In particular, we show how neurophysiological data recorded in the amygdala and orbitofrontal cortex of monkeys during appetitive and aversive trace conditioning supports the hypothesis that neurons have mixed selectivity to external events like the CSs and inner mental states encoding the predicted value of the stimuli. The model neurons are assumed to exhibit the same response properties before the process of creation of context representations starts. We then illustrate the proposed mechanisms underlying the formation of context representation by describing the simulations of a model neural network performing a trace conditioning task. In particular, we show how transient events can generate patterns of sustained activity that bridges the gap between two successive relevant events. We then explain the iterative process of merging of the neural representations of short temporal sequences (attractor concretion) that leads eventually to the temporal contexts. Simulations show that these representations can significantly improve the prediction of the value of a stimulus when the context changes. Finally, we use the model to make specific predictions for the patterns of neural activity that would be observed given new experimental manipulations.

Learning context representations

The initial situation: experimental evidence for neurons with mixed selectivity

In our model, we assume that the neurons of the CN have mixed selectivity to the mental states of the AN (positive, negative, neutral) and external events (CSs, USs). They should exhibit this form of selectivity to the conjunction of mental states and events even before the learning process leading to the formation of context representation starts. The assumption is based on two considerations: (1) these conjunctions contain the basic elements that characterize the contexts. For example, selectivity to CS A would not allow the network to discriminate between the two contexts, as the very same stimulus A is presented in both contexts, in which it would activate the neuron in exactly the same way. However, a neuron that responds to CS A only when the following mental state of the AN is *positive* would activate only in one context and not in the other. (2) Neurons with this form of mixed selectivity can be easily obtained with random connections, and hence without any learning procedure (see Methods and Rigotti et al. (submitted for publication)). Indeed, neurons that are connected with random synaptic weights to the neurons of the AN, which represent the mental state, and to the neurons representing the external events, are very likely to respond only to the simultaneous activation of these two populations, provided that the threshold for activating the neuron is large enough. Neural representations of patterns of activity across several randomly connected neurons are analogous to random projections and they can encode efficiently the information contained in both the external and internal inputs (e.g. the Johnson and Lindenstrauss lemma (Dasgupta and Gupta, 2002)).

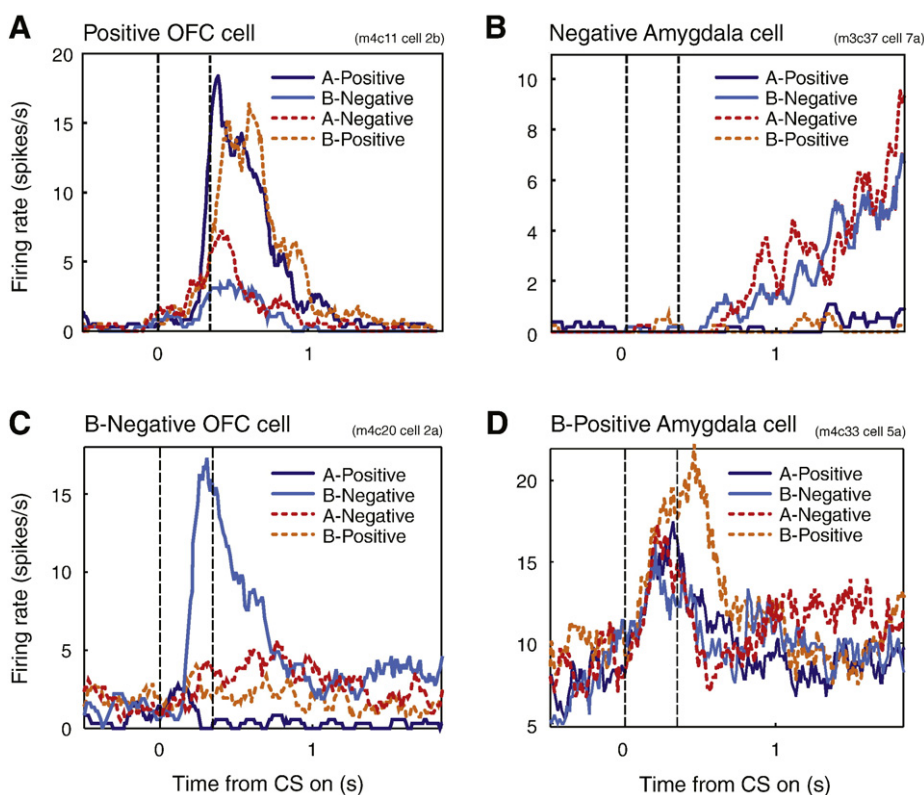


Fig. 5. Recorded activity of OFC and amygdala cells that respond as expected in AN (A,B) and CN (C, D). The activity was recorded while the monkey was performing the trace-conditioning task for the four possible CS-US pairings. The continuous traces show the activity after the monkey had learned the associations defining Context 1 (A-Positive, B-Negative). The dotted traces show the activity after learning of Context 2 (A-Negative, B-Positive). The AN cells show sustained activity during the trace interval that encodes the value of the CS. These cells have been observed both in the OFC (A) and amygdala (B). The CN cells are selective to specific combinations of CS and value, both in the OFC (C) and the amygdala (D).

Neurons with the assumed mixed selectivity (in the CN) and with the expected response properties for the AN have been observed in various areas of the brain. For example, the CSs are assumed to evoke value-dependent sustained activity in the AN. We observed neurons with these response properties both in the OFC and in the amygdala while the animal was performing the trace conditioning experiment (see Figs. 5A, B). The activity can be sustained throughout the trace interval (see the cells of Paton et al. (2006)), for a limited time (Fig. 5A), or it can ramp up in anticipation of reinforcement (Fig. 5B).

The CN neurons are assumed to respond to conjunctions of events and the internal states of the AN. In the trace conditioning experiment, we expect to observe in the first context neural representations of the mixtures, such as CS A-Positive or CS B-Negative, whereas in the second context, the patterns represent CS A-Negative or CS B-Positive. We have often observed neural responses reflecting this type of mixed selectivity in the amygdala and OFC. We recorded and analyzed 216 cells in OFC and 222 from the amygdala (recorded from two monkeys). We used a two-way ANOVA to determine whether image value, image identity or an interaction between image value and identity accounted for a significant portion of the variance in neural firing. For this analysis, we excluded the first five trials of the experiment of each type of trial, as well as the first five trials of each trial type after reversal. We did this to exclude trials in which neurons were changing their firing rate during learning about reinforcement contingencies. Of particular interest to our proposed model, a substantial number of neurons in both the amygdala and OFC had a significant effect of the interaction between image identity and value (66/216 OFC neurons, and 87/222 amygdala neurons, $p < 0.01$, 2-way ANOVA). Neurons with a significant interaction term indicate that responses to images are modulated in an unequal manner by reinforcement contingencies, which is the precise type of response profile postulated by the model. Two examples of these types of neurons are depicted in Figs. 5C, D. Notice that in each case, neurons represent the conjunction between a particular image and a particular reinforcement. In these two cases, image identity and image value do not have a significant effect on responses, but many of the cells with significant interactive effects also show significant effects of a main factor in the ANOVA.

The first learning phase: from transient conjunctions of events and mental states to attractors

We assumed that the CN initially receives a strong excitatory input from the AN and the neurons representing the external events only when particular events (the CSs A and B, the USs Reward and Punishment) are preceded or followed by specific states of the Associative Network, AN (neutral, positive, negative). For example, consider a trial in which CS A is associated with reward (first trial in Fig. 6A). We assume that the AN has already learned the correct association, and the presentation of CS A induces a transition from a state with neutral value (0) to a state with a positive value (+) (see Fig. 6B). The neurons encoding a positive value have sustained activity until reward is delivered and the activity of the AN is shut down. The CN neurons observe the following sequence of AN states and events: Neutral-A, A-Positive, Positive-Reward, Reward-Neutral. The corresponding populations of neurons are transiently activated in the same order (0A, A+, +R, R0, see Fig. 6C). Analogously, in a trial in which CS B is associated with punishment, we have Neutral-B, B-Negative, Negative-Punishment, Punishment-Neutral (see Fig. 6, second trial). For simplicity, we assumed in our simulations that each conjunction like Punishment-Neutral activates a single population of the CN.

The Hebbian component of learning strengthens the synaptic connections between the CN neurons that are repeatedly co-activated. Moreover, it depresses the synapses from active to inactive neurons (see Methods for the detailed equations). As a consequence, all patterns of activity of the CN that are activated a sufficient number of

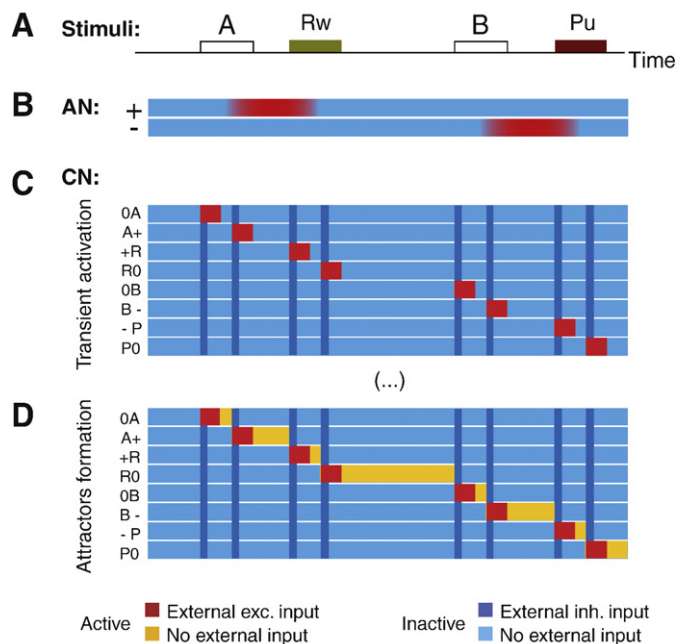


Fig. 6. First learning phase: from transient events to attractors. (A) The scheme of two consecutive trials. In the first trial the presentation of CS A is followed, after a delay, by the delivery of reward. In the second one, CS B is followed by punishment. (B) Color coded activity of the AN (red = active, blue = inactive) as a function of time in response to the events depicted in panel A. The simulation starts in an inactive state with neutral value (0). The presentation of CS A induces a transition to a state in which the neurons encoding positive value (+) have self-sustained activity. The activity is shut down by the delivery of reward. Analogously for the CS B-Punishment case. (C) Color coded activity of the CN populations as a function of time (red = active in the presence of external input, blue = inactive because of the strong inhibitory input generated by a reset signal). Each row represents the activity of one population that is labeled according to its selectivity (e.g. 0A is a neuron that responds only when the AN is the neutral state and CS A is presented). The external events together with the activation of positive and negative states of the AN activate the populations of the CN (red bars). Every time a different population is activated a reset signal is delivered (blue stripe). (D) First CN attractors: the synapses within each repeatedly activated population are strengthened to the point that the activity self-sustains also after the event terminates (yellow bars).

times become attractors of the neural dynamics. At the end of the first phase, we have the situation illustrated in Fig. 6D. Each conjunction of events and AN states activates a population of the CN, and the activity remains elevated even after the event terminates. However, every time the AN activates a population of CN neurons that was inactive, or deactivates a population of CN neurons that was active, we assumed that a reset signal is delivered and the previous pattern of reverberating activity is shut down by a strong inhibitory input (blue stripes in the figure).

The second learning phase: concretion of temporally contiguous attractors

Now that the representations of the conjunction of events and AN states are attractors of the neural dynamics, the time gap between one event and the next one is bridged by the self-sustained patterns of reverberating activity. Two successive conjunctions of events belonging to the same or different trials become temporally contiguous. This enables the temporal sequence learning (TSL) mechanism to modify synaptic weights and to link two successive patterns of activity, so that the process of attractor concretion can start. The TSL mechanism operates only when the pattern of activity of the CN changes because it is modified by the AN input, and it strengthens the synapses between an active pre-synaptic neuron and a neuron that is activated at the successive time step. Moreover it depresses the synapses between active neurons and neurons that are inactivated at the successive time step. If the synapses between two populations, say *a*

and *b*, are sufficiently potentiated, then the activation of *a* causes also the activation of *b*, leading to the merging of the two attractors (attractor concretion).

The process of formation of the context representations requires a few iterations, and the typical phases that we observe in the simulations are illustrated in Fig. 7. The iterative process generates representations of progressively longer temporal sequences. To illustrate this process, consider the same two trials considered in Fig. 6. The CN dynamics are now simulated at different stages of the learning process (Figs. 6B–E). A scheme representing the temporal statistics of the activation of the CN populations that are relevant for the concretion process is shown in the right column. Although this scheme does not allow us to make quantitative predictions about the detailed neural dynamics, it is useful to describe the dynamics of the concretion process, and in particular to understand how the temporal statistics of the events and mental states are related to the probability that two attractors merge into a single representation. Each arrow links two CN populations, and its thickness represents the propensity of these populations to merge. The propensity depends on both the parameters of the learning dynamics and the temporal statistics of the activation of the two populations. In particular, the arrow connecting two generic populations, say *a* to *b*, is proportional to the probability that *a* is followed by *b*, multiplied by the number of times that *a* is

activated within a certain time interval, which depends on the parameters of the learning dynamics. This is motivated by the fact that the synapses between *a* and *b* are potentiated by the TSL mechanism every time that *a* is activated, and then it is followed by *b*. They are depressed when *a* is followed by a population different from *b*. The stronger a synapse becomes, the higher the probability that *a* activates *b*, and hence that the two populations merge into a single representation. The propensity to concretion depends also on other details of the TSL mechanism (see the Methods for the description of the full dynamics), but also and more strongly on the effects of the Hebbian mechanism. In particular, the Hebbian component of synaptic dynamics strengthens the synapses within population *a* every time *a* is activated, and it depresses the synapses from *a* to all the other populations, including *b*. This effect of stabilization of *a* increases with the time that the CN spends in a state in which *a* is active. This means that the strength of the connections between *a* and *b*, and hence the propensity to concretion, should decrease with the time that *a* is active. This is valid only when *a* and *b* are not already co-activated, because in such a case the Hebbian term actually strengthens their connections. For these reasons, the propensity is inversely proportional to the fraction of time that the CN spends in *a* when *b* is inactive, multiplied by the ratio between the Hebbian learning rate and the TSL learning rate. Summarizing, the propensity

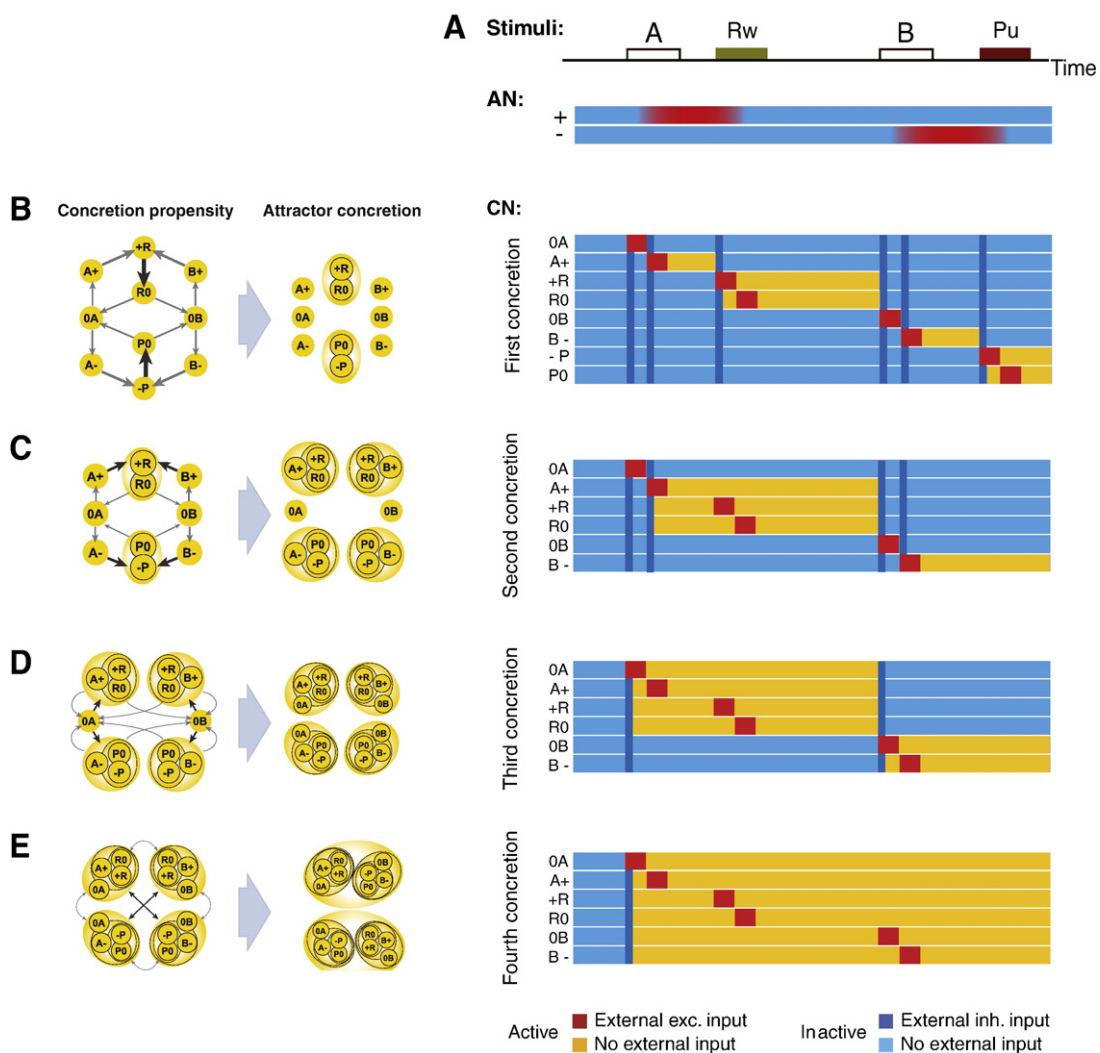


Fig. 7. Second learning phase: attractor concretion. (A) Scheme of two trials and color coded activity of the AN as a function of time as in Fig. 6. (B, C, D, E) From left to right: Scheme of propensities to concretion, scheme of attractors following concretion, Color coded activity of the CN populations as a function of time as in Fig. 6 following the concretion. B, C, D, E describe different iterations of the concretion process (see the text for a detailed description).

of a to merge with b is high when a is activated repeatedly and it is often followed by b . However it is reduced if the CN spends a large fraction of time in a , or if the CN is often driven to states other than b .

The largest propensities drive the first concretions. For example Fig. 7B shows that +R and R0 are the first population to merge into a single attractor. Indeed +R is consistently followed by R0 in both contexts. Notice that also A+ is always followed by +R, but the propensity to merge is smaller because A+ is activated on average half of the times that +R is activated. The result of this first concretion is illustrated in the simulations of Fig. 7B. The activation of +R now turns on also the R0 population, and from now on, the two populations will always co-activate since they are part of a new compound.

The next iteration is again driven by the concretion propensities; however, now there are new attractor states in the CN, and all the propensities must be recalculated. The new scheme of propensities is shown in Fig. 7C. The next concretions are again predicted correctly by the propensities. The same process is iterated in Fig. 7D and in E, where we finally obtain the representations of the two contexts. Notice that at every iteration the width of the arrows progressively decreases. This is due to the fact that the CN spends more and more time in the new attractors and hence the propensity to concretion with other attractors decreases because of the Hebbian component of learning. At some point the process of concretion stops because the propensity is too small to induce a concretion. We chose the learning rates to stop the process as soon as we have the representations of the two contexts (but see also Discussion for a different choice of parameters). The full simulations of the learning process are shown in Fig. 8.

Predicting context-dependent values: the expected behavior

After the learning process described in the previous section, the CN contains a representation of the current context. When the CS-US associations are reversed, the first time a CS is presented, the value is predicted incorrectly by the AN. This resets the synapses from the neurons representing the external events to the AN, and the transitions to the positive and negative states become random with equal probability (Fusi et al., 2007). At the same time a “surprise” signal is generated in the CN and the attractor representing the context is reset. If the AN selects by chance the wrong value state, then it keeps selecting the state randomly. As soon as the AN selects the correct value, then it also activates the correct context in the CN, and the AN-CN system starts predicting the correct values for all CSs. Although it is not possible for our neural circuit to switch with 100% probability from one context to the other in one trial, it is still possible to harness the information contained in the context representation of

the CN in order to improve the prediction of the US. Indeed, as soon as the AN guesses the correct value for one CS, say CS A, the CN also selects the correct context and then it is possible to predict the US that follows CS B with certainty. Summarizing, as soon as the context changes, the AN predicts the wrong value. The surprise signal resets the CN context, and then the AN-CN system selects randomly one of the two possible contexts until it guesses the correct one. This strategy is less efficient than switching to the alternative context as soon as one knows that the context has changed, but it is still more efficient than learning independently the associations, as it allows the AN-CN system to predict correctly the value of all the CSs once it knows the new value of one CS.

This mechanism is implemented in our neural circuit by the feedback from the CN to the AN, as described in Materials and methods. The CN and the external neurons project to a population of randomly connected neurons which represent mixtures of the CN context representations and the external events. These neurons contain the information about the current context and the occurring event. The synapses between these neurons and the AN are plastic, and they are modified in the same way as the synapses from external neurons to the AN, which encode simple Pavlovian associations. As soon as the context is correctly determined by the CN, it is simple to predict the values of both CSs, as the AN sees the CN as an additional input that represents explicitly the current context. Indeed, Fig. 9 shows the percentage of correct predictions of the value of one CS when the neural circuit has already guessed correctly the value of the other CS. In other words, we quantify the ability of the neural circuit to use the context information to infer the value of one CS once the value of the other CS is known. In the absence of context information, this percentage is at a significantly lower level, which depends on the specific sequence of events (left). In the presence of the CN feedback, this percentage is close to 100% (right). This behavior is in principle observable in an experiment, and the plot of Fig. 9 provides us with a behavioral criterion for establishing the existence of context representations.

Experimental predictions about the response properties of recorded cells

As the process of concretion takes place, the neural representations evoked by events like the presentations of the CSs or the delivery of the USs become progressively more similar in the CN. For example, the activation of CS B and Negative (B-) should eventually evoke the activation of the attractor representation of context 1. Neurons that initially were activated by CS B and Negative only are predicted to be activated also by CS A and Positive. In particular, as the process of attractor concretion starts from the events that follow each other with the highest probability, the first compounds that form are likely to be [+R, R0], [-P, P0], followed by [0A, A+], [0B, B-], for context 1, and [0B, B+], [0A, B-] for context 2 (see Fig. 10).

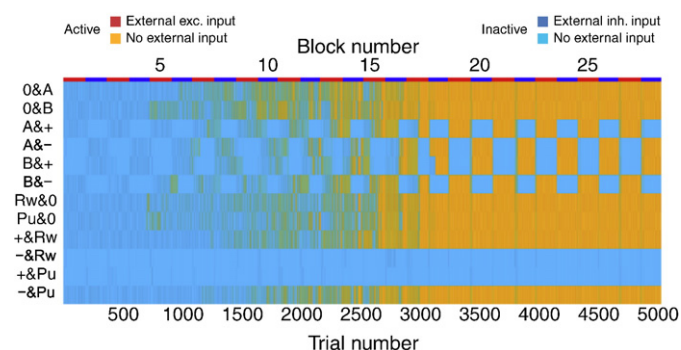


Fig. 8. Full learning simulation. Color coded activity of the CN populations as a function of time as in Fig. 6. Red and blue bars above the plot indicate context 1 and 2, respectively. Temporally contiguous attractors merge into single representations of short temporal sequences (attractor concretion). Eventually, the context representations emerge, and they are demonstrated by the coactivation of the attractors representing all conjunctions of events and AN states in each context (e.g. 0A, A+, +R, R0, 0B, B-, -P, P0 for context 1).

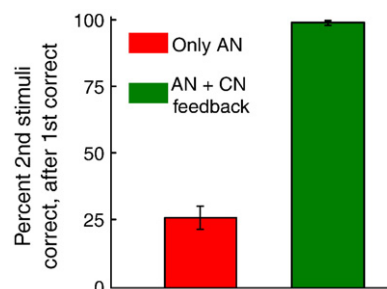


Fig. 9. Harnessing the feedback from the CN to the AN: percent of correct predictions of the value of one CS when the new value of the other CS is already known. The performance is estimated immediately after a context switch. In the absence of the context information provided by the CN, the performance is significantly worse (left) than in the presence of CN feedback, when the performance is close to 100%.

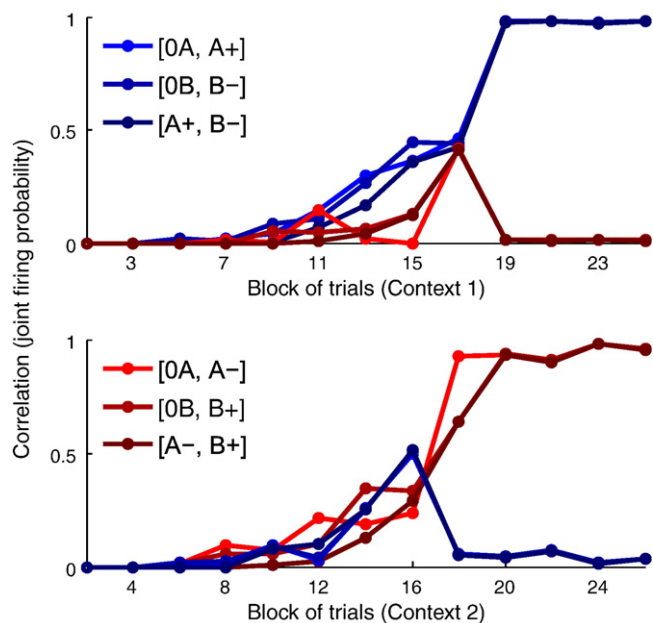


Fig. 10. Predictions on the correlations between neurons that respond to conjunctions of events. The probability that two populations of neurons of CN are co-activated is computed by running the simulation several times and it is plotted as a function of the number of blocks of context 1 (top) and context 2 (bottom) trials. Different colors denote different co-activated populations. Initially the probability is zero, as we assumed that in the CN there are only populations that respond to simple conjunctions of events and AN states. As learning progresses, the probability of co-activation of populations that represent events of the context increases. For example, in context 1, neurons that initially respond only to CS A and Positive (A+), after 15 reversals (block 16), respond also to CS B and Negative (B-). The corresponding compound is denoted by [A+, B-].

This prediction can be tested experimentally by single unit recordings. We predict that the fraction of recorded cells that are activated by events that belong to the same context should increase monotonically with the total number of trials that the animal experiences. Notice that in the simulations of Fig. 10 we assumed that there are initially no neurons that already respond to conjunctions of more than two events. In the brain there might be neurons that from the very beginning respond already to those conjunctions of events that represent the contexts in the CN. The probability of finding those neurons is predicted to be significantly smaller than the probability of finding neurons that respond to simple conjunctions of events, like those that we simulated, but in general we cannot exclude that they are already present before the learning process starts. Consistent with our predictions, in the trace conditioning experiment we observed 11/216 cells in OFC and 15/222 cells in the amygdala that are significantly selective to the conjunctions of events that represent the context (see [Materials and methods](#) for the details of the analysis). Our analysis shows that these cells may only be significantly selective by chance, or they may simply be a very small fraction of the recorded cells. In both cases, our assumption that the majority of cells responds to simple conjunctions of events is correct.

Discussion

We proposed attractor concretion as a possible mechanism for creating representations of contexts. The formation of these representations leads to the generation of new mental states that were not present in the initial AN-CN system. Indeed, the CN, at the end of the learning process, operates as an additional input to the AN that contains information about the active context. When the AN-CN system is considered, the number of mental states eventually doubles, as for each AN state there are two states of the CN. In practice, two new sets of

mental states are generated as soon as the CN starts encoding the two contexts. In the intermediate stages of the iterative process that leads to the formation of the new states, the CN neurons are activated by short sequences, which generate a number of mental states that is larger than the final one. However, the temporal statistics of these sequences are not correlated with any useful information about the contexts. In particular, they do not allow the AN to disambiguate between the CS-US associations of the two different contexts. Indeed in the first three stages (in Figs. 7B, C, D) there is no evidence of the existence of temporal contexts, both in the transition probabilities and in the propensities. It is only at the end of the concretion iterative process that the CN can detect a clustered structure in the transition probabilities (i.e., large probabilities to make a transition within the context, and small ones to switch to a different one). We believe that the kind of merging behavior we observed in the case we analyzed is very general and it applies to a wide variety of tasks in which the information about temporal contexts is important to improve the performance.

In our manuscript we analyzed a particularly difficult case, because the CSs and the USs appear in the two contexts in a perfectly symmetric way, and the contexts are solely defined by the temporal statistics of the events. Our model would certainly work and generate quantitative predictions in simpler cases, in which, for example, a particular CS-US identifies unequivocally one context, or in which an explicit contextual cue appears in all or in some trials to signal which context is active. The concretion rules illustrated in the schemes of Fig. 7 allow us to make predictions about the typical behavior of our model network, although it is always important to run a full simulation to analyze the behavior of the concretion process. For example, the representation of a contextual cue that appears in every trial would have the highest propensity to merge with CS-US associations, and it would operate as a kernel around which the context representations are built. The details of this learning process obviously depend on the specific task protocol, on the representations of the CSs and the USs and on the parameters of the model, but the model can certainly generate quantitative predictions in a wide variety of experiments.

A hierarchy of context representations

The final context representations are determined not only by the temporal statistics of the events, but also by the parameters of the neuronal and synaptic dynamics. In a system with multiple neural circuits, each characterized by different dynamical parameters, we expect the creation of a hierarchy of contexts that correspond to different processes of attractor concretion. Some neural circuits can represent only single events, as the TSL component is not strong enough; some others might represent general contexts that correspond to conjunctions of many events, stimuli and internal conditions. For example, there could be cases with no merging, in which the network simply represents the individual events A-Reward, B-Punishment, A-Punishment and B-Reward. For the parameters used in our simulations, the patterns of activity corresponding to A-Reward and B-Punishment merge into the representation of context 1, whereas A-Punishment and B-Reward merge into context 2. For other sets of parameters, there could be a unique, large compound comprising all possible events of the context corresponding to the general task. Such a compound would link together A-Reward, B-Punishment, A-Punishment and B-Reward. All these neural circuits with different neural and synaptic parameters are likely to be simultaneously present in the brain, either in one particular area, or spread across different areas. They would provide the brain with a hierarchy of contexts at many different levels that all together determine a general mental state. A population of cells reading out all these context representations could easily encode the value of the current state, and such a value would in general depend on which context pattern is activated in every level of the hierarchy.

The use of a hierarchy of contexts generated by the heterogeneity of the network is a new possibility that will be investigated systematically in future theoretical and experimental studies. In particular, it should be possible to detect to which hierarchical level a recorded neuron belongs. This could be done by manipulating the transition probabilities between events. For example, we considered an experiment in which there are only two CSs and the probability to make a transition from a trial with one CS to a trial with a different CS within the same context is simply $1/2$. In the case of $2p$ CSs, this probability would be $1/2p$. The probability to switch to a different context should be accordingly reduced when all the CSs are presented at least once in each context ($\ll 1/2p$). Hence the clustered structure of transition probabilities that defines the temporal context would still be detectable for any $p > 2$, but the transition probabilities would all be rescaled down when p increases. The parameters of the neural circuits that determine the propensity to concretion and the maximal number of CSs are always related. In particular, neural circuits with a small propensity to concretion would generate context representations only if p is small enough. $p = 2$ maximizes the probability that we observe what we described in the experimental prediction section. However, we also have the additional prediction that for any neural circuit that shows concretion for $p = 2$, there is always a p that is large enough so that no concretion should be observed. Of course the entire brain would still be able to create context representations because there would be a different neural circuit with the proper parameters to generate the context representations in a different situation.

Previous experimental and theoretical works on temporal context representation

Neural signals that could provide a representation of temporal context have been observed in several experiments. For example, Miyashita investigated the representations of sequences of visual stimuli (Miyashita and Chang, 1988). In these experiments, a monkey was trained to perform a Delayed Match to Sample task in which the sample stimuli were presented in a fixed temporal order. Single neuron activity recorded in inferotemporal cortex revealed that cells activated by one particular stimulus were likely to be activated also by the neighboring stimuli in the temporal sequence. In other words, the spatial patterns of neural activity across multiple cells, induced by temporally contiguous stimuli, were highly correlated, reflecting the order of presentation of the sensory events. This work inspired several theoretical works on the neural mechanisms underlying context representation in the brain. For example, Griniasty et al (1993) interpreted this data as being an expression of the recurrent dynamics of cortical circuits. In the model the authors proposed, these circuits initially produce different patterns of stable, reverberating activity in response to individual sensory stimuli. If the stimuli are repeatedly presented in a fixed temporal order, the synaptic connections are modified so that a sensory stimulus activates a pattern of activity that is correlated not only to its representation when presented individually, but also to the representations of the stimuli that surround it in the temporal sequence. This pattern of reverberating activity may be a neural representation of the context in which the sensory stimulus appears. A more detailed model of the learning process that is responsible for tuning the synaptic weights has been proposed by (Brunel, 1996) and some of the predictions have been verified in experiments (Yakovlev et al., 1998).

In all these models, each stimulus evokes a different representation and the context is encoded in the correlations between the representations. Highly correlated neural patterns correspond to stimuli that belong to the same context. In our model, attractor concretion leads to new inseparable “entities” that represent contexts. Different events, like visual stimuli, activate the very same pattern of activity if they belong to the same context. This is different from the representations in which the contexts are encoded in the correlations

because, in that case, each event still activates a characteristic pattern of activity that is unique, though similar to the patterns elicited by the events of the same context. Our approach is similar to what has been proposed in (Brunel, 1996) for pair associates. One of the advantages of creating new entities that represent context is compositionality. Context representations can easily merge to generate new compounds even if they are highly structured and complex. In other words, when a new entity is created, there is a significant reduction of the dimensionality of the state space. Indeed, if two populations are always coactivated, they behave as a single one, reducing the effective number of independent populations, and hence the dimensions of the state space. This reduction greatly simplifies the process of representation of complex contexts. This is not true in the case in which each individual event contains the information about the correlation with all the other population of neurons (i.e. with all the other dimensions). One of the disadvantages of creating new entities is that the information about individual events is lost, unless multiple systems are considered as discussed in the previous section of Discussion.

Alternative approaches to the creation of context representations

As we briefly discussed in the Introduction, one of the limitations of Reinforcement Learning (RL) algorithms is that a representation of mental states is assumed, yet the algorithms do not provide for how these representations are created. This is a major limitation when the environment is only partially observable, as in the case of limited sensory data, or when agents have limited computational resources to process all the details of the sensory stimuli. In all of these cases, we often might be induced to think that we are in the same situation when we actually are not and we would need to select different actions. For example, when we drive in a forest, we often arrive at two similar crossings, which can lead to confusion if we do not take into account other information, such as where we have been recently. Thus, in many circumstances it is possible to decide our actions if we remember some of our previous experiences. For example, one crossing might be preceded by a pump station and the other might be preceded by a level crossing. In this situation, we need to create two distinct mental states, each corresponding to a different temporal context.

We proposed a simple mechanism and a biologically plausible neural network model that autonomously generates context representations. Alternative and complementary approaches have been proposed to solve analogous problems in different fields. For example, Hidden Markov Models (HMM) have been widely used to predict complex temporal series where the next event might depend on a long sequence of previous observations (see e.g. (Rabiner, 1989)). In these models and in some of their extensions to decision processes (e.g. POMDP, Partially Observable Markov Decision Processes (Kaelbling et al., 1998)), the number and the meaning of the states of the agent are not known a priori. The algorithms usually start from a large number of hidden states that are randomly linked to the observed states of the environment. The states acquire a meaning by iterating a recursive algorithm that estimates both the probabilities that a hidden state is related to an observed state, and that it is followed by another state (see e.g. Viterbi or Baum-Welch algorithms). Although very efficient in many applications like speech recognition, these algorithms suffer from many limitations. They require knowing the number of states a priori. If there are not enough hidden states, then it is possible to add more, but estimating all the transition probabilities between hidden states becomes rapidly an intractable problem as the number of states increases. Moreover, all probabilities should be reevaluated every time a new state is introduced, which makes the system rather inflexible. Finally, the convergence of the recursive algorithms to the global optimal solution is not guaranteed and the final scheme of hidden states and transition probabilities depends strongly on the initial condition. Some of these limitations might be

related to the fact that the hidden states are initially chosen in a completely random fashion, and they are not constructed on the basis of the temporal statistics of the events. Although we still do not have a general theory and a convergence theorem as in the case of HMMs, we believe that our approach does not suffer from these limitations, and it is closer to the mechanisms that the brain might be using to create context representations.

Where are the cells of the AN and the CN in the brain?

As shown in Fig. 5, neurons with mixed selectivity to the relevant conjunction of events for defining the two contexts have been observed both in prefrontal cortex (PFC), in particular in OFC, and in the amygdala. As these two areas are strongly interacting (Salzman and Fusi, *in press*), it is likely that the neural circuits of the AN and the CN are distributed across these brain regions, and probably also other areas (e.g. in other subareas of PFC, as well as in the hippocampus and related structures).

When temporal contiguity is broken by intervening distractors

In many realistic situations, there are contexts in which the temporal contiguity between relevant events is broken by distractors, e.g. by the presentation of a random visual stimulus. In our model, these distractors would disrupt the process of formation of context representations. There are at least two possible solutions to this problem. The first one has been proposed in (O'Reilly and Frank, 2006), and it is based on a gating system that can learn to ignore the irrelevant events. In these models the irrelevant events are simply “gated,” and hence not represented in the AN and in the CN. The second possibility is based on short term synaptic mechanisms that could preserve the memory of relevant events even when distractors are presented (see Mongillo et al. (2008) for a possible mechanism based on short term synaptic facilitation). For example, the TSL mechanism could be implemented by tagging the synapses at time t , and then modifying them in the next time interval τ , creating links between the event occurring at t and all events occurring the next time interval τ . Such a mechanism would suffer from the introduction of the inherent time constant τ of synaptic tagging, whereas our mechanism can work and generalize with almost any timing between successive events. Notice that synaptic tagging could in principle allow us to eliminate the first phase of learning, in which the attractor representations of individual events are created to bridge the temporal gap between events that are separated in time.

More general mental states and operant tasks

We focused on a trace conditioning task to illustrate the proposed mechanism for the formation of context representations. However, the same mechanism can be applied to other experiments and to operant tasks. For example in (Asaad et al., 1998; Pasupathy and Miller, 2005) the monkeys are trained to associate saccadic movements to visual responses. The AN model has already been used to reproduce quantitatively the observed behavior of the monkeys when they learn and forget visuo-motor associations (Fusi et al., 2007). The positive and negative states are equivalent to the decisions of the monkey about the direction of the saccade (left or right). In one context, stimulus A is associated with left and B with right, and in the second context the associations are reversed. Although in the experiments the monkeys never learned to switch from one context to another immediately, it is possible that in other conditions (see the discussion of Fusi et al. (2007)) they would be able to create the representations of the contexts. The model and the predictions would be the same as for the trace conditioning task. Notice that the two contexts correspond to two simple rules that could be expressed as “whenever A is associated with left, B is associated with right” and

“whenever A is associated with right, B is associated with left.” In this case the representation of the temporal context would be equivalent to the representation of a rule. In recent years, investigators accumulated evidence that the activity of prefrontal neurons can encode abstract rules (Genovesio et al., 2005; Mansouri et al., 2007, 2006; Wallis et al., 2001). These rules allow the animal to respond to the same sensory stimulus in different ways depending on the strategy or on a sequence of sensory cues preceding the stimulus. Hence, they are all analogous to the temporal contexts that we studied here. In one of the cited experiments, Mansouri et al. (2006) observed sustained activity in the inter-trial intervals that encodes the rule in effect when the monkey was performing a simplified version of the Wisconsin Card Sorting Test. This rule was determined by the temporal context of monkey choices, and the rule-selective inter-trial activity therefore corresponds to an active representation of context (O'Reilly and Munakata, 2000; Loh and Deco, 2005; Deco and Rolls, 2005; Rigotti et al., 2008; Rougier et al., 2005). In all the models and the experiments that we described, a large proportion of neurons exhibit mixed-selectivity. Interestingly, it has been shown (Dayan, 2007) that mixed selectivity neurons implemented with multilinear functions can actually play an important role in neural systems that implement both habits and rules during the process of learning of complex cognitive tasks. Multilinearity implements conditional maps between the sensory input, the working memory state, and an output representing the motor response.

Temporal contiguity can also be important in the creation of invariant representations. Some investigators propose that invariant representations of objects can be generated by linking temporally contiguous views (Rolls and Milward, 2000; Miyashita and Chang, 1988; Li and DiCarlo, 2008). This is yet another example of an abstraction process that relies on temporal contiguity to create the internal representations. As shown in our manuscript, temporal contiguity can also be an important aspect of the statistics of the environment in all the cases in which behavior depends on information about temporal context. We believe that it is particularly important to study the neural mechanisms that allow the animal to encode complex patterns of temporal contiguity, as these processes might underlie the neural basis of cognitive functions that range from the creation of invariant representations to rule abstraction.

Acknowledgments

We would like to thank X-J. Wang and S. Ostojic for many discussions, ideas and suggestions, and for reading the manuscript. We are grateful to E. Curti who was involved in the early states of this work. This work was supported by the SNF grant PP00A-106556, by DARPA SyNAPSE, by the NIMH grant 2R01MH58754, by the Gatsby and the Swartz Foundations. CDS was supported by the James S. McDonnell and Gatsby Foundations, NIDA (R01 DA020656), NIMH (R01 MH082017), and NEI (R24 EY015634). SEM received support from a National Science Foundation graduate fellowship and from an individual NIMH National Research Service Award (F31 MH081620).

References

- Amit, D.J., 1989. *Modeling Brain Function*. Cambridge Univ. Press.
- Amit, D., Brunel, N., 1995. Learning internal representations in an attractor neural network with analogue neurons. *Network: Comput. Neural Syst.* 6 (3), 359–388.
- Asaad, W.F., Rainer, G., Miller, E.K., 1998. Neural activity in the primate prefrontal cortex during associative learning. *Neuron* 21 (6), 1399–1407.
- Belova, M.A., Paton, J.J., Morrison, S.E., Salzman, C.D., 2007. Expectation modulates neural responses to pleasant and aversive stimuli in primate amygdala. *Neuron* 55 (6), 970–984.
- Belova, M.A., Paton, J.J., Salzman, C.D., 2008. Moment-to-moment tracking of state value in the amygdala. *J. Neurosci.* 28 (40), 10023–10030.
- Brader, J.M., Senn, W., Fusi, S., 2007. Learning real-world stimuli in a neural network with spike-driven synaptic dynamics. *Neural Comput.* 19 (11), 2881–2912.
- Brunel, N., 1996. Hebbian learning of context in recurrent neural networks. *Neural Comput.* 8 (8), 1677–1710.

- Dasgupta, S., Gupta, A., 2002. An elementary proof of the Johnson–Lindenstrauss lemma. *Random Struct. Algorithms* 22, 60–65.
- Dayan, P., 2007. Bilinearity, rules, and prefrontal cortex. *Front. Comput. Neurosci.* 1, 1.
- Deco, G., Rolls, E.T., 2005. Synaptic and spiking dynamics underlying reward reversal in the orbitofrontal cortex. *Cereb. Cortex* 15 (1), 15–30.
- Weinan, E., Li, T., Vanden-Eijnden, E., 2008. Optimal partition and effective dynamics of complex networks. *Proc. Natl. Acad. Sci. U. S. A.* 105, 7907–7912.
- Fusi, S., Asaad, W.F., Miller, E.K., Wang, X.J., 2007. A neural circuit model of flexible sensorimotor mapping: learning and forgetting on multiple timescales. *Neuron* 54 (2), 319–333.
- Gardner, E., 1987. Maximum storage capacity in neural networks. *Europhys. Lett.* 4, 481–485.
- Genovesio, A., Brasted, P.J., Mitz, A.R., Wise, S.P., 2005. Prefrontal cortex activity related to abstract response strategies. *Neuron* 47 (2), 307–320.
- Griniasty, M., Tsodyks, M.V., Amit, D.J., 1993. Conversion of temporal correlations between stimuli to spatial correlations between attractors. *Neural Comput.* 5, 1–17.
- Hopfield, J.J., 1982. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U. S. A.* 79 (8), 2554–2558.
- Kaelbling, L., Littman, M., Cassandra, A., 1998. Planning and acting in partially observable stochastic domains. *Artif. Intell. J.* 101, 99–134.
- Krauth, W., Nadal, J.P., Mézard, M., 1988. The roles of stability and symmetry in the dynamics of neural networks. *J. Phys. A: Math. Gen.* 21, 2995–3011.
- Li, N., DiCarlo, J.J., 2008. Unsupervised natural experience rapidly alters invariant object representation in visual cortex. *Science* 321 (5895), 1502–1507.
- Loh, M., Deco, G., 2005. Cognitive flexibility and decision-making in a model of conditional visuomotor associations. *Eur. J. Neurosci.* 22, 2927–2936.
- Mansouri, F.A., Matsumoto, K., Tanaka, K., 2006. Prefrontal cell activities related to monkeys' success and failure in adapting to rule changes in a Wisconsin card sorting test analog. *J. Neurosci.* 26 (10), 2745–2756.
- Mansouri, F.A., Buckley, M.J., Tanaka, K., 2007. Mnemonic function of the dorsolateral prefrontal cortex in conflict-induced behavioral adjustment. *Science* 318 (5852), 987–990.
- Minsky, M.L., Papert, S., 1969. *Perceptrons*.
- Miyashita, Y., Chang, H.S., 1988. Neuronal correlate of pictorial short-term memory in the primate temporal cortex. *Nature* 331 (6151), 68–70.
- Mongillo, G., Barak, O., Tsodyks, M., 2008. Synaptic theory of working memory. *Science* 319 (5869), 1543–1546.
- Morrison, S., Salzman, C.D., 2009. The convergence of information about rewarding and aversive stimuli in single neurons. *J. Neurosci.* 29 (37), 11471–11483.
- O'Reilly, R., Munakata, Y., 2000. *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. The MIT Press.
- O'Reilly, R.C., Frank, M.J., 2006. Making working memory work: a computational model of learning in the prefrontal cortex and basal ganglia. *Neural Comput.* 18 (2), 283–328.
- Pasupathy, A., Miller, E.K., 2005. Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature* 433 (7028), 873–876.
- Paton, J.J., Belova, M.A., Morrison, S.E., Salzman, C.D., 2006. The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature* 439 (7078), 865–870.
- Rabiner, L.R., 1989. A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE* 77 (2), 257–286.
- Rigotti, M., Ben Dayan Rubin, D., Wang, X.J., Fusi, S., 2008. Mixed neuronal selectivity is important in recurrent neural networks implementing context dependent tasks. *Society for Neuroscience Annual Meeting (Washington, DC, Society for Neuroscience)*, p. 929.3/TT11.
- Rigotti, M., Ben Dayan Rubin, D., Wang, X.J., and Fusi, S., (submitted for publication). Internal representation of task rules by recurrent dynamics: the importance of the diversity of neural responses. Submitted to *Nature Neuroscience*.
- Rolls, E.T., Milward, T., 2000. A model of invariant object recognition in the visual system: learning rules, activation functions, lateral inhibition, and information-based performance measures. *Neural Comput.* 12 (11), 2547–2572.
- Rosenblatt, F., 1958. The perceptron: A probabilistic model for information storage and organization in the brain, volume 65.
- Rosenthal, O., Fusi, S., Hochstein, S., 2001. Forming classes by stimulus frequency: behavior and theory. *Proc. Natl. Acad. Sci. U. S. A.* 98 (7), 4265–4270.
- Rougier, N.P., Noelle, D.C., Braver, T.S., Cohen, J.D., O'Reilly, R.C., 2005. Prefrontal cortex and flexible cognitive control: rules without symbols. *Proc. Natl. Acad. Sci. U. S. A.* 102 (20), 7338–7343.
- Salzman, C.D. and Fusi, S., (in press). Emotion, cognition and mental state representation in amygdala and prefrontal cortex. *Annual Review of Neuroscience*.
- Salzman, C.D., Paton, J.J., Belova, M.A., Morrison, S.E., 2007. Flexible neural representations of value in the primate brain. *Ann. N. Y. Acad. Sci.* 1121, 336–354.
- Senn, W., Fusi, S., 2005. Learning only when necessary: better memories of correlated patterns in networks with bounded synapses. *Neural Comput.* 17 (10), 2106–2138.
- Soltani, A., Wang, X.J., 2006. A biophysically based neural model of matching law behavior: melioration by stochastic synapses. *J. Neurosci.* 26 (14), 3731–3744.
- Sompolinsky, H., Kanter, I., 1986. Temporal association in asymmetric neural networks. *Phys. Rev. Lett.* 57 (22), 2861–2864.
- Sutton, R., Barto, A., 1998. *Introduction to reinforcement learning*. MIT Press Cambridge, MA, USA.
- Wallis, J.D., Anderson, K.C., Miller, E.K., 2001. Single neurons in prefrontal cortex encode abstract rules. *Nature* 411 (6840), 953–956.
- Wang, X.J., 2002. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* 36 (5), 955–968.
- Yakovlev, V., Fusi, S., Berman, E., Zohary, E., 1998. Inter-trial neuronal activity in inferior temporal cortex: a putative vehicle to generate long-term visual associations. *Nat. Neurosci.* 1 (4), 310–317.