Coexisting persistent and dynamic representations of stimulus, time and trial outcome emerge in the hippocampus-prefrontal circuitry during learning of a temporal association

Bryan C. Souza[1,†], Jan L. Klee[1], Luca Mazzucato[2,†,*] and Francesco P. Battaglia[1,†,*]

[1]Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, the Netherlands
[2]Institute of Neuroscience, University of Oregon, USA
[†]Correspondence: bryan.dacostasouza@donders.ru.nl (B.C.S.), lmazzuca@uoregon.edu (L.M.), f.battaglia@science.ru.nl (F.P.B.)
[*]These authors contributed equally to this work.

Abstract

Temporal associations between sensory stimuli separated in time rely on the interaction between the hippocampus and medial prefrontal cortex (mPFC). However, it is not known how changes in their neural activity support the emergence of temporal association learning. Here, we use simultaneous electrophysiological recordings in the hippocampal CA1 region and mPFC of mice to elucidate the neural dynamics underlying memory formation in an auditory trace conditioning task. We found that in both areas conditioned (CS+/CS-) and unconditioned stimuli (US) evoked similar temporal sequences of neural responses that progressively diverged during learning. Additionally, persistent CS representations emerged in mPFC after learning, supported by CS+ coding states whose transient reactivation reliably predicted lick onset and behavioral performance on single trials. These results show that coordination of temporal sequences in CA1 and persistent activity in mPFC may underlie temporal association learning, and that transient reactivations of engrams in mPFC predict the animal behavior.

Highlights

- Temporal representation of CS+, but not CS- and US, strengthens in CA1 after learning.
- Similarity between stimulus and reward temporal representations decrease with learning, but is recovered in error trials.
- Representation of stimulus identity is strong and stable in PFC since stimulus onset, while it only emerges in CA1 during trace period.
- Neural states defined on faster time-scales reveal the emergence of CS coding states in PFC whose onset predicts lick times and task performance.
- PFC-CA1 states are relatively less coordinated during late CS+ stimulus and trace.

Introduction

The hippocampus and the medial prefrontal cortex (mPFC) are both known to have important roles in sensory and memory representations. For instance, the hippocampus integrates multiple sensory inputs coming from the entorhinal cortex, bridging the 'what', 'where' and 'when' components of a memory (Eichenbaum (2017a)) while the mPFC is believed to create the stimulus-context associations that support decision making (Euston et al. (2012)) and to support working memory via sustained attention on relevant sensory representations (Postle (2006), Lara and Wallis (2015)). Although the interaction between those two areas are known to support different cognitive processes (Battaglia et al. (2011), Eichenbaum (2017b), Shin and Jadhav (2016), Preston and Eichenbaum (2013)), it is still unknown how the interplay between the mPFC and hippocampus can lead to the formation of internal representations linking events separated in time.

The mPFC neuronal activity can be modulated by different stimuli and reward (Starkweather et al. (2018), Otis et al. (2017)). Beyond responding to positive outcomes (Euston et al. (2012), Burton et al. (2009), Gruber et al. (2010), Pratt and Mizumori (2001)), the mPFC can also encode expectancy of both positive and negative outcomes (Baeg (2001), Gilmartin and McEchron (2005)), and reward absence (Quilodran et al. (2008)). Moreover, mPFC is implicated in acquisition of context-dependent rules (Miller and Cohen (2001), Miller (2000)) and required for the maintenance of stimulus representations during delays in working memory tasks (Goldman-Rakic (1995), Funahashi et al. (1993)), even though this representation might be stored elsewhere (Lara and Wallis (2015), Postle (2006)). The mPFC was shown to encode long memory traces of past decisions (Murakami et al. (2017)) and rewards (Bernacchia et al. (2011)) and it is required to implement trial-history biases in decision-making tasks (Murakami et al. (2017).

Similar to the mPFC, the hippocampus is also required in tasks involving delayed decision (Hattori et al. (2015), McEchron and Disterhoft (1999, 1997)). Beyond its strong link to spatial processing (Moser et al. (2015, 2008)), hippocampal activity also encodes multiple temporal features, possibly bridging task events separated in time through temporal sequences (Pastalkova et al. (2008a), MacDonald et al. (2011), Naya and Suzuki (2011),

Eichenbaum (2014)). For instance, CA1 inactivation during a trace fear condition task impaired learning by preventing the temporal binding of the CS and US stimuli in memory (Sellami et al. (2017)). During trace fear conditioning, a reorganization of CA1 ensemble activity occurred during learning of the conditioned response, linked to the emergence of sparse CS-selective responses in pyramidal cells (Ahmed et al. (2020)). However, the circuit mechanisms underlying CA1 network reorganization and the role of the CA1-mPFC coordinated interaction during learning are unknown.

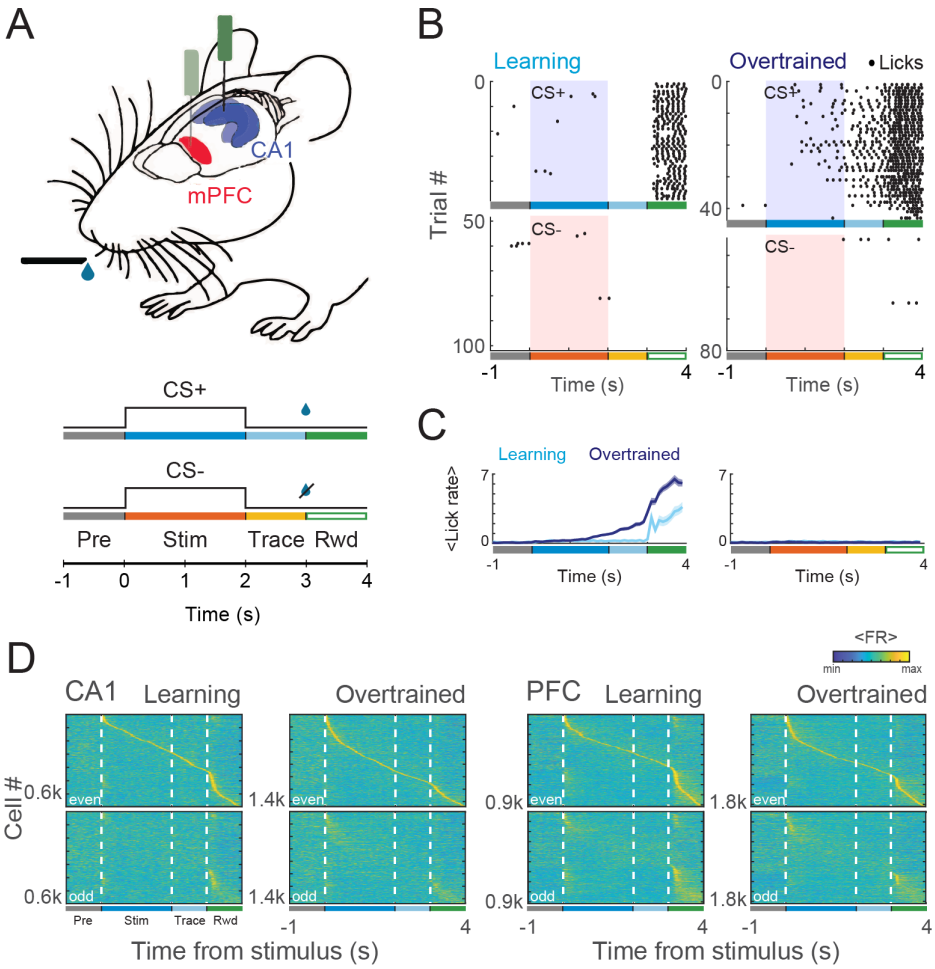

Figure 1: **Behavioral and firing rate profile during a trace conditioning. A.** Schematic representation of the task. Head-fixed animals with acute implants on CA1 and mPFC were presented a 2-sec auditory stimulus (either CS+ or CS-). After a 1-sec (trace period) a reward (US) was available on CS+ trials. **B.** Lick response of an example animal in CS+ and CS- trials before and after learning. **C.** Average lick response for CS+ and CS- trials for all animals. Overtrained sessions were identified by the presence of anticipatory licks in CS+, but not in CS- trials. **D.** Average firing rate profile of all recorded cells over the course of a CS+ trial. Peak firing rate time on even trials was used to sort cells on both groups (even and odd). Note the presence of temporally structured firing after CS and US presentation.

In this work, we aimed at elucidating the neural mechanisms leading to the emergence of a temporal association, linking events occurring at different times. To this end, we recorded neurons from the hippocampal area CA1 and mPFC of mice engaged in an appetitive auditory trace-conditioning task (Klee et al. (2021)), and examined how learning shaped the relationship between stimulus responses, temporal encoding, and behavior. We found that both conditioned (CS) and unconditioned stimuli (US) triggered the onset of temporal sequences in both CA1 and mPFC, which in turn were differentially modulated by learning. Moreover, we found, after

learning, that unexpected reward delivery evoked an echo of the CS+ temporal sequence in mPFC during the US. We uncovered strong and persistent stimulus representation in the mPFC, which started after the stimulus onset and persisted until the US, bridging stimulus and reward over the trace periods. Using single-trial decoding based on hidden Markov models, we found that the reactivation of CS+ coding states in mPFC predicted the onset of wrong licks in response to CS-stimuli, leading to error trials. Coordination between CA1 and PFC states decreased during CS+ trials, suggesting that CA1 may support internal memory representation during the initial stimulus sampling, while mPFC activity is dominated by network states encoding stimulus-reward mapping and driving behavior in overtrained animals.

Results

We previously developed an appetitive trace conditioning task (Klee et al, 2021) which trained mice to differentiate between two sound stimuli (conditioned stimulus: CS+ and CS-; Figure 1A). In each CS+ trial, the 2s stimulus was followed by a 1s delay (trace period), and a reward delivered in a water port. The same trial structure was present in CS- trials, except that there was no reward. In order to investigate the neural dynamics in CA1 and mPFC during the trace conditioning task, we acutely inserted high-density silicon probes in CA1 and/or mPFC and recorded spiking activity during the task, before and after learning. Crucially, the presence of a long trace period in this experimental paradigm rules out a simple Hebbian mechanism to link CS and US neural representations, because the two stimuli occur in different epochs of the trial without any temporal overlap Raybuck and Lattal (2014). Additionally, this paradigm allowed us to separate representations that might emerge during the sensory cue, from the internally generated representations that bridge the stimulus to the behavior during the trace period. Learning criteria was achieved when animals licked in anticipation of US in CS+ trials, but not in CS- trials (lick rates statistically different between CS+ and CS- trials - $p<0.05$, t-test; Figure 1A-C).

Differential encoding of temporal information in CA1 and mPFC ensembles

We first investigated whether CA1 and mPFC ensembles encoded the passing of time within the trial by generating sequential activity that consistently tiled the duration of a trial. To test this, we computed the average binned firing rate (80 ms running average) of the pooled neurons over a subset of the trials (even trials). The neurons were then sorted according to the time of their peak firing rate. We then applied the same order in the remaining, odd trials to unveil cells firing in a reliable temporal structure. If neurons generated consistent temporal sequences, we would expect the same sequential activation to be present both in the odd and in the even trials. We found a concentration of peak activation around the first 0.5 s after stimulus (CS), which seemed to be modulated by learning (mainly in CA1; Figure 1D). In addition, we also found a concentration of peaks after reward onset (US). Interestingly, some cells peaking in the CS had also a peak at the US and vice-versa (this effect was mainly seen in PFC). Those results unveiled the presence of temporal sequences coordinated by stimulus (CS+) and reward (US) onset at the population level (Figure 1D, Figure S1).

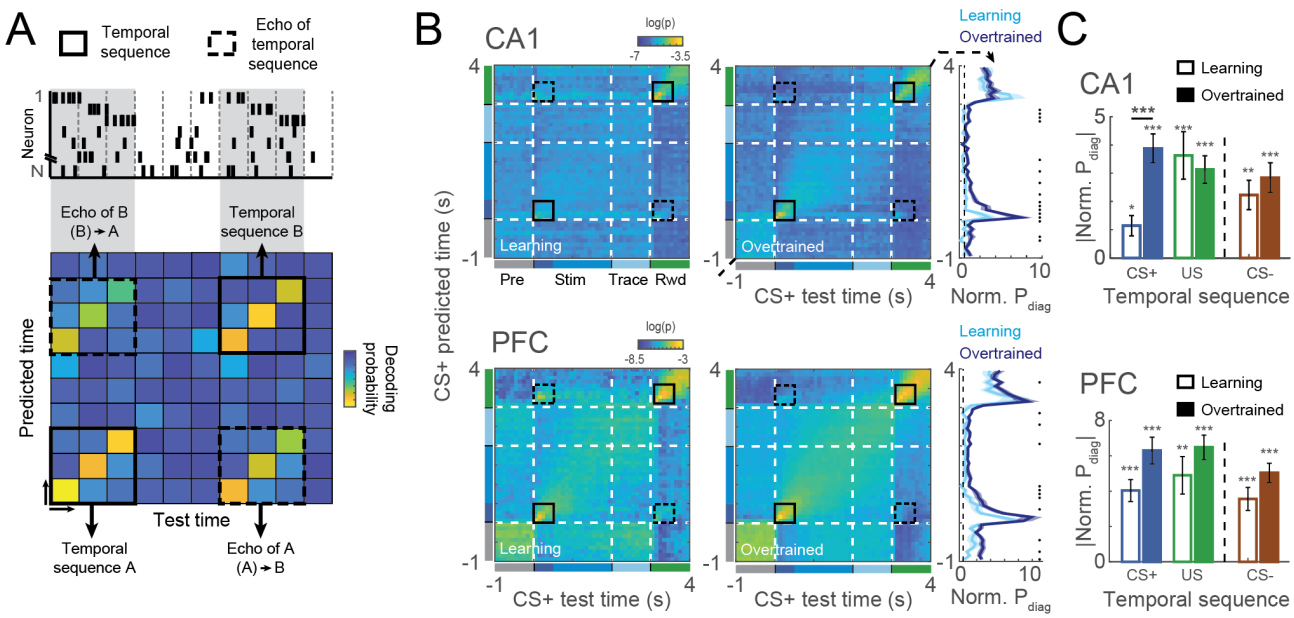

Figure 2: **Emergence of temporal sequences for CS+, CS- and US. A.** Schematic representation temporal sequences in the spike activity (top) and how they emerge in the time decoding matrices (bottom). Bayesian decoders were trained to predict the trial time of a given time bin, which was used to build a matrix with decoding probabilities for each time bin over the course of the trial. Temporal sequences were then identified as clusters of high performance in the main diagonal, where real time matches predicted time. Additionally, diagonal stripes of high temporal decoding performance outside the main diagonal revealed the similarity between two temporal sequences, and were defined as the temporal echo of one sequence (actual time) in the other (predicted time). **B.** Mean time decoding matrices for learning and overtrained sessions together with the average normalized probability on the main diagonal (rightmost panel). Decoders were trained and tested using CS+ activity from CA1 (top) and PFC (bottom). Black boxes show the regions in which temporal sequences (solid lines) and temporal echoes (dashed lines) were present. Black dots denote times in which learning and overtrained decoding are different (Wilcoxon ranksum, $p<0.05$). **C.** Average (normalized) decoding probability around CS+ and US temporal sequences, taken on main diagonal of the contiguous black boxes in B. Normalization was done by z-scoring probability values using the the mean and s.d. from surrogate decoding matrices. CS- sequences were computed similarly to CS+, but from a decoding matrix computed using CS- activity for training and testing (see also Figure S2). We used aligned rank transform to perform a nonparametric two-way ANOVA using recording area and learning condition on each temporal sequence as factors. We found no significant interaction effect for the two factors ($F(1,107)= 0.51$, $p=0.47$; $F(1,107)= 0.74$, $p=0.39$ and $F(1,107)= 0.64$, $p=0.43$ for CS+, US and CS-, respectively). Simple main effects analysis revealed that PFC had higher temporal sequences than CA1 ($F(1,107)= 20.08$, $p<0.001$; $F(1,107)= 18.02$, $p<0.001$ and $F(1,107)= 14.04$, $p<0.001$ for CS+, US and CS-, respectively). Also, for CS+ temporal sequences, we found an effect on learning condition ($F(1,107)= 18.08$, $p<0.001$; $F(1,107)= 1.30$, $p=0.26$ and $F(1,107)= 3.47$, $p=0.07$ for CS+, US and CS-, respectively). Notice the increase in CS+ temporal representation in CA1 after learning. Shaded areas and errorbars denote SEM. *$p<0.05$; **$p<0.01$; ***$p<0.001$.

We confirmed these results on a session-by-session basis, using a cross-validated Naive-Bayesian classifier to discriminate time points within a trial from population activity of each session (Figure 2A). Decoding probability was then evaluated for different time bins, building a time-decoding matrix with the corresponding testing bin time in the x-axis and the decoded time on y-axis. If temporal coding was present, we expected a high classification accuracy on the main diagonal where predicted time (y-axis in Figure 2A) matched real time (x-axis). The presence of temporal sequences was then assessed by measuring how much decoding probability successively concentrated in the main diagonal, where decoded time is equal to real sample time.

Similar to our previous population analysis, we found temporal sequences associated with CS+ and US (Figure 2B). This was evident in the mean time-decoding matrices, but also when looking exclusively to the

main diagonal (right-most panel). Interestingly, time-decoding matrices computed using CS- trials also showed a strong temporal sequence associated with the stimulus (Figure S2A). To quantify this effect, we computed the average of the main diagonal in the temporal sequence window (0-0.5s for CS and 3.1-3.6s for US) (Figure 2C). All the temporal sequences (CS+, CS- and US) were significantly different compared to a sham temporal sequence artificially detected before the stimulus onset ((-1)-(-0.5)s).

We then tested whether the strength of temporal encoding differed between areas and whether it was modulated by learning. We found a stronger temporal encoding in mPFC compared to CA1 following each stimulus (non-parametric two-way anova with aligned rank transform (Wobbrock et al. (2011); main effect (mPFC vs. CA1), $p<0.001$; no interaction effects; see Figure 2C). Moreover, we found stronger CS+ temporal sequence after learning in both areas (main effect (learning vs. overtrained), $p<0.001$; see Figure 2C). We could also see reduction on the relative dimensionality of the population activity during the temporal sequences in CA1 and (after learning) in PFC (Figure S2B). In general, those results suggest that encoding of time is stronger in PFC than in CA1 and, in particular, during CS+, which is the most task-relevant stimulus. They also imply a link between temporal sequences and learning (Figure 2C; see also Figure S2).

Temporal sequences of CS+, CS-, US 'echo' each other, in a learning-dependent fashion

The time-decoding analysis revealed a surprising 'echo' effect, whereby US and CS+ temporal sequences shared a similar representation of time following each stimulus (high classification accuracy in off-diagonal regions, Figure 2A-B). We refer to this effect as the temporal echo of one sequence (the one used to test the model) in another sequence (the predicted sequence).

Because (CS+)–evoked temporal sequences in CA1 were stronger after learning sessions, we hypothesized that temporal echos could be modulated by learning as well. We reasoned that the similarity between US and CS+ representation could be inferred by how well US sequences would generalize to CS+. To investigate this, we computed time-decoding matrices focusing on how US temporal sequences are decoded by a model restricted to the rest of the trial (pre-stimulus, stimulus and trace period; Figure 3A), from CS+ trials.

We found that, during learning sessions, the predicted time of the US temporal sequence strongly concentrated at the beginning of the stimulus presentation, where CS+ temporal sequences also happened (Figure 3A). This effect was strongly reduced in overtrained sessions in both CA1 and PFC, both in the general US average decoding performance (Figure 3A; rightmost panel) and in the temporal echo computed as the average performance on the main diagonal of the echo window (Figure 3B).

Because temporal echo strength decreased with learning, we hypothesized that echos could represent error signals arising during incorrect trials. Since after learning the animal committed less mistakes, that would explain weaker temporal echoes in overtrained sessions. To investigate this, we separated trials into correct (i.e., CS+ trials with anticipatory licks during trace or CS- trials with no licks) and incorrect (i.e., CS- trials with licks during trace period and CS+ trials with no licks). We then computed the time-decoding matrices

(trained on all CS+ trials) for correct and incorrect CS+ test trials (Figure 3C-D). We found that incorrect CS+ trials showed a stronger US echo (in CS) compared to correct trials. This was true for learning sessions, in which the temporal echo was evident (Figure 3C,E), but also in overtrained sessions (Figure 3D,E), partially recovering the effect we found before learning. Temporal echoes were stronger on incorrect trials (two-way ANOVA with factors outcome and learning condition, main effect (outcome) CA1: $p<0.05$; PFC: $p<0.001$) and during learning sessions (main effect (learning condition) CA1: $p<0.05$; PFC: $p<0.001$). Moreover, in PFC those two factors showed a significant interaction (outcome and learning condition; $p<0.05$) (Figure 3E).

Interestingly, the US echo in the CS+ was stronger in incorrect trials compared to correct ones. Computing the same analysis using CS- trials (US echo in CS-) showed no difference between correct and incorrect trials S3. Thus, CS and US representations of time are initially similar, but progressively decorrelate with learning. Because CS+ and US sequences recovered an increased similarity during incorrect trials (i.e., when the animal didn't expect the reward), we further hypothesized that the temporal sequences found were a general mechanism evoked by unexpected events. We reasoned that, although CS+ and US could initially be both perceived as unexpected events, after learning US occurrences would be predictable by CS+, leading to a change of the US-evoked temporal sequence and a decreases the similarity between the two sequences. To test this hypothesis we investigate the similarity between CS- temporal sequences and CS+ and US before and after learning.

To do so, we first used a Bayesian decoder trained to predict CS- time (as in Figure S2) to decode samples of CS+ activity. We reasoned that, if the same neural representation of time was shared between the CS+ and CS- conditions, then a classifier trained to discriminate time in CS+ trials would generalize well when tested on CS- trials; and vice versa. Alternatively, a failure of the classifiers to generalize between conditions would suggest the presence of different time encoding schema. This cross-stimulus decoder revealed that CS- and CS+ temporal sequences have a shared temporal representation, with the CS+ temporal sequence being correctly decoded during CS- in both CA1 and PFC (Figure 4A, see also Figure S4 for CS- sequences during CS+ trials). In CA1, this shared representation increased with learning, probably reflecting the strongest temporal CS+ representation in that area after learning. Moreover, the CS- temporal sequence also echoed on US temporal sequence in the test trial (Figure 4B). This effect could also be seen in the analogous case of testing CS- trials in a model trained with CS+ trials, where the CS- temporal sequence was echoed in the US epoch of the CS+ decoder (Figure S4A). Similarly to CS+, the CS- temporal echo decreased after learning of the task (Figure S4B). Together those results corroborate the idea of an initial unspecific representation of time evoked by unexpected events, which acquires stimulus specificity upon learning.

Learning induces persistent stimulus representations in PFC but not CA1

We next examined whether population activity encoded stimulus identity (i.e., CS+ or CS-). We trained an ensemble of SVM classifiers to discriminate between CS+ and CS- conditions in each bin. For each classifier trained on a particular time bin, we then predicted the trial identity (CS+ or CS-) of all other bins. We used different combinations of training and testing time bins to build a temporal generalization matrix (King and

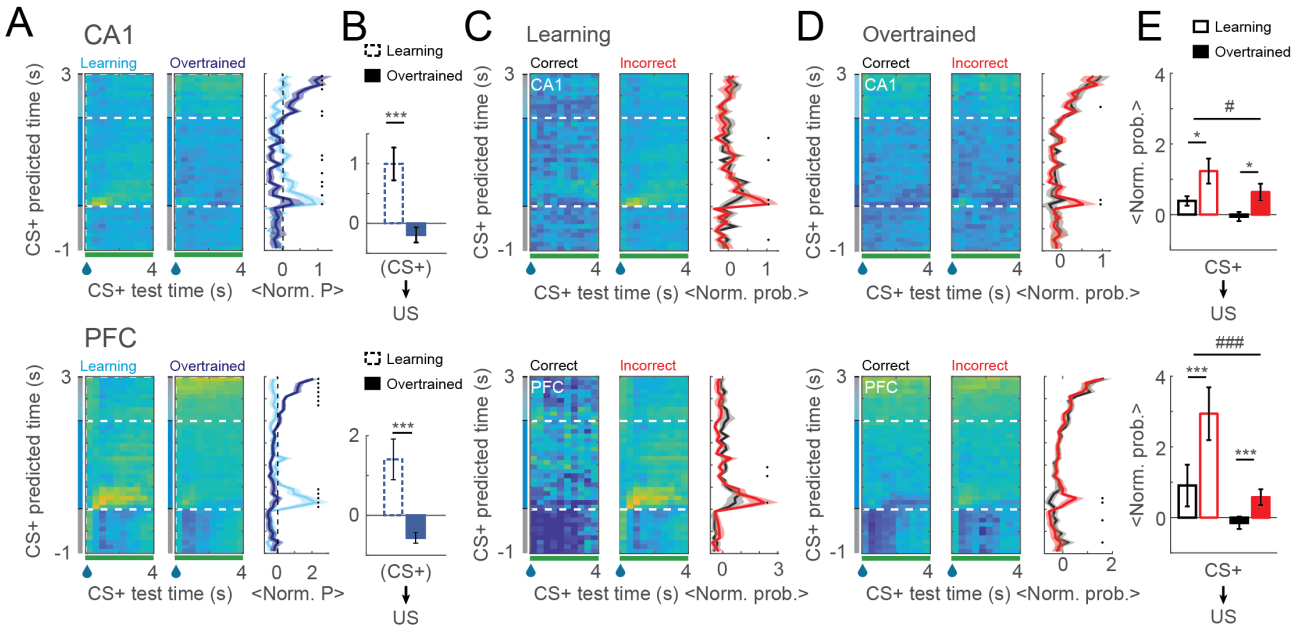

Figure 3: **Similarity between CS+ and US temporal representations is modulated by learning and trial outcome. A.** Time decoding matrices trained using CS+ activity before the US, and used to decode US time (left and middle). Average (normalized) probabilities for the US period are shown for learning and overtrained sessions (right). **B.** Mean decoding in the temporal echo window for CA1 and PFC. **C-D.** Time-decoding matrices as in A, but showing the temporal echo in correct and incorrect trials from learning (C) and overtrained (D) sessions. **E.** Average performance on the temporal echo window (main diagonal) for CA1 (top) and PFC (bottom). We used aligned rank transform to perform a nonparametric two-way ANOVA using trial outcome (i.e., correct vs. incorrect) and learning condition as factors for each area independently. We found no significant interaction effect for the two factors in CA1 and ($F(1,117)= 0.08$, $p=0.78$) and significant effect on learning condition ($F(1,117)=6.59$; $p<0.05$) and trial outcome independently ($F(1,117)=5.05$; $p<0.05$). In PFC both trial outcome ($F(1,92)= 12.05$, $p<0.001$) and learning ($F(1,92)= 18.13$, $p<0.001$) showed significant main effects, but there was also a significant interaction between factors ($F(1,92)= 4.93$, $p<0.05$). Shaded areas and errorbars denote SEM. */#: $p<0.05$; ***/###: $p<0.001$.

Dehaene (2014); Figure 5A), a framework which allowed us to examine the temporal nature of the stimulus representations in CA1 and PFC (Figure 5A). We tested three alternative scenarios: i) A high decoding accuracy restricted to a tight region around the main diagonal, suggesting a reliable yet time-varying stimulus representation (such as stimulus-specific sequential activity); ii) An extended area of high accuracy covering a large diagonal block, suggesting a persistent and stable stimulus representation; iii) No area of significant decoding accuracy, suggesting the absence of a reliable stimulus representation (for example, in a scenario of stochastic reactivation of stimulus-specific patterns at random times in each trial). We found very different encoding dynamics in CA1 and PFC. In CA1 decoding performance seemed to concentrate near the main diagonal, and we found a significant increase in decoding performance during trace period only (Figure 5C-D). On the other hand, a strong and persistent stimulus representation encompassing both stimulus and trace periods emerged in PFC after learning (Figure 5C-D). Both effects could also be seen by looking at stimulus representation at specific time points (i.e., rows of the decoding matrices in Figure 5C), as summarized in Figure 5E. Stimulus encoding increased strongly with learning in PFC during both the stimulus and trace periods. In CA1, stimulus encoding was weaker and mildly improved with learning only during the trace period.

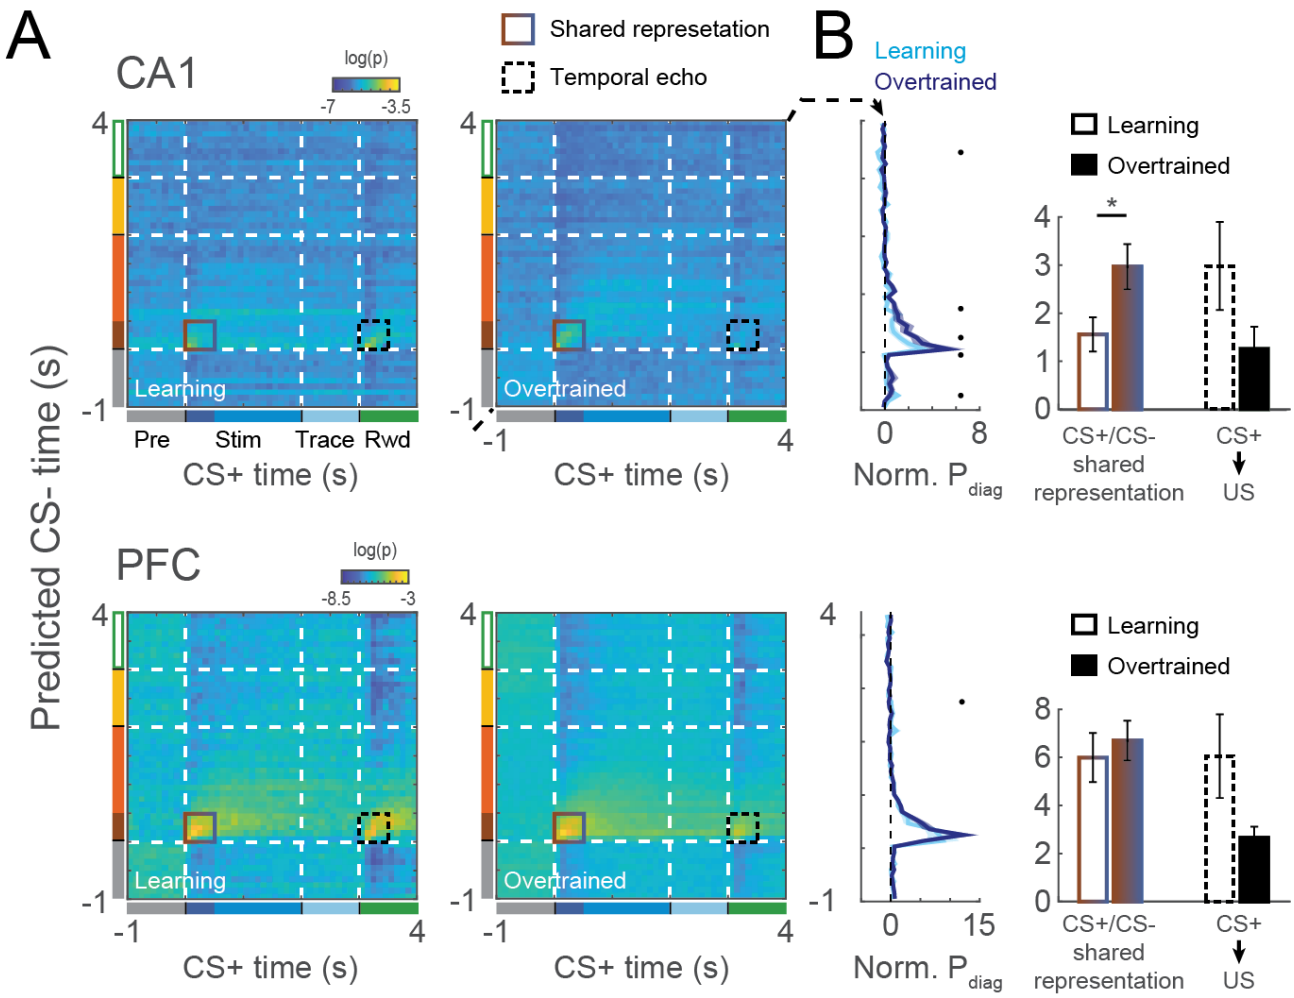

Figure 4: **CS- have a shared temporal representation with CS+. A.** Time-decoding matrices trained using CS- activity and used to decode time in CS+ trials in CA1 (top) and PFC (bottom). Solid boxes show areas in which a shared temporal representations can be seen. Dashed boxes show US temporal echoes on CS-. **B.** Normalized decoding performance in the main diagonal (left) and average decoding over the shared and echoed representations (right) in the matrices in A.

Transient co-activation of neural assemblies underlies stimulus-coding neural representations

The classification analysis discussed above showed that CA1 ensembles do not encode stimulus information in a time-ordered representation consistent across trials. Alternatively, stimulus-specific assemblies could stochastically reactivate for brief intervals, at random times in different trials. To investigate the presence of transient yet stochastic reactivation of neural assemblies, we fit a hidden Markov model to each recording session, separately in CA1 and mPFC. The HMM is a probabilistic generative model for neural population data, capturing both the network organization in a small set of stereotyped neural states, and also the dynamical transitions between these state forming stochastic sequences. Model selection for the number of neural states was obtained via the Bayesian information criterion (BIC) (S5A; Recanatesi et al. (2022)). Figure 6A shows an example of neural activity in a CS+ trial in PFC, where two states were detected, whose onset closely followed the onset of different trial epochs (state detection occurred whenever a state's posterior probability exceeded 0.8).

A rastergram showing state activation in CA1 and PFC for all trials in a representative session can be seen

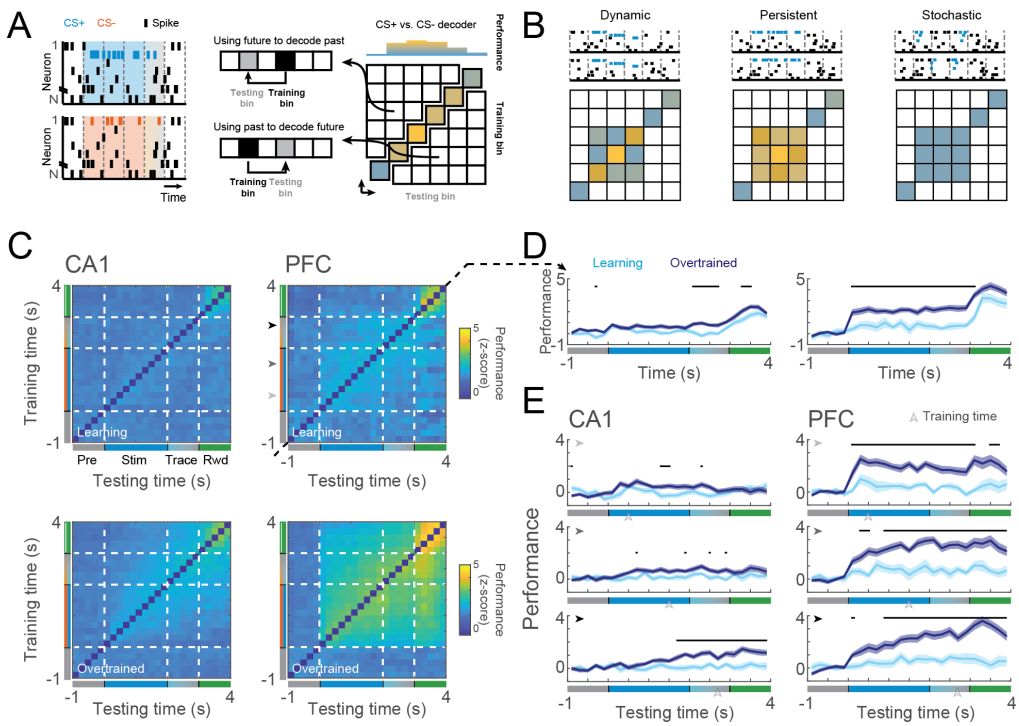

Figure 5: **Different trial type decoding dynamics in CA1 and PFC A.** Schematic representation of trial type decoding matrices. The activity in each time bin was used to train a CS+/CS- SVM classifier. Test bins were evaluated on decoders based on past, current and future time bins, enabling an evaluation of stability of neural representations over time. **B.** Possible modes of neural representations. Trial decoding matrix can reveal whether the nature of the stimulus representation is time-varying (left), persistent (middle) or undetectable (right). In case of a time-varying representation, population activity would encode stimulus identity in a stereotypical trajectory, with non-generalizable decoders over the trial time. Alternatively, in case of a persistent representations, population activity would reach a fixed stable attractor, with decoders being generalized over the trial time. At last, the stimulus decoding matrix could detect no reliable stimulus representation. **C.** Mean CS decoding matrices for learning (top) and overtrained (bottom) sessions of CA1 and PFC. **D.** Average performance of the main diagonal in the matrices in C. Black dots represent bins in which learning and overtrained performance are significantly different ($p<0.05$; Wilcoxon ranksum test). **E.** Rows of the CS decoding matrices in C. Black arrows denote the time bin used to train the models (y-axis in C). Notice the stability of PFC stimulus representation over the stimulus and trace period while CA1 representation mildly strengthens during trace. Black traces denote different learning and overtrained performances ($p<0.05$; Wilcoxon ranksum test).

in Figure 6B. In general, CA1 showed a higher number of states than PFC in both learning and overtrained sessions (Figure S5B) and state activation in CA1 was also shorter than in PFC (Figure S5C). It's worth to mention that the HMM latent states could also capture the presence of the temporal sequences in the form of sequential state activation. However, those state probabilities were often not strong enough to pass our activation criterion (state probability > 0.8; Figure S6).

To investigate whether HMM states encoded stimulus identity, we used the distribution of state activation during stimulus and trace periods in CS+ or CS- trials (Figure 6C). States that were significantly more active during CS+ or CS- were defined as 'CS+ coding states' or 'CS- coding states', respectively (Mazzucato et al. (2019)). Before learning, a small number of coding states were detected in PFC but not in CA1 (Figure 7A). Both areas showed a significant increase in the number of coding states after learning, with a much larger fraction of coding states in PFC compared to CA1 (despite the larger number of HMM states detected in CA1 overall, most CA1 sessions had no coding state), consistent with the larger stimulus-decoding accuracy found in PFC

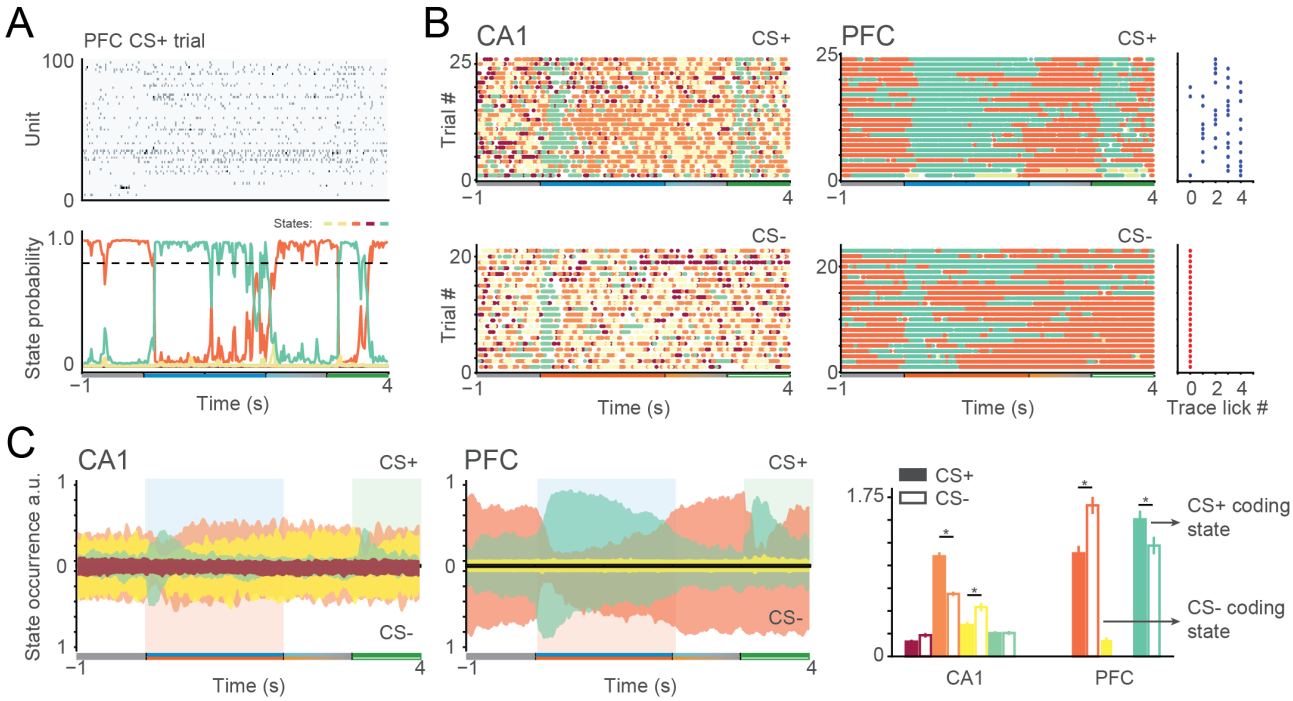

Figure 6: **Transient co-activation of neural assemblies captured by hidden Markov models. A.** Example of PFC neural activity during a CS+ trial (top) and its respective hidden states probabilities (bottom). State activation was defined by moments in which state probability was above 0.8 (color-coded curves represent posterior probabilities of each state). **B.** Raster plot of state activation in CA1 and PFC over different CS+ (top) nad CS- (bottom) trials for the session in A. Scatter plots show the number of licks during the trace period of each trial. **C.** Probability of state activation over CS+ (upper quadrant) and CS- (lower quadrant) trials. Bar plots show the average state activation probability during trace and stimulus period. Note that some states are significantly more active CS+ (CS+ coding states) or CS- (CS- coding states). HMM states were computed separately for CA1 and PFC. Errorbars denote SEM (*p<0.05; Wilcoxon signed rank test).

compared to CA1. Interestingly, occurrence of CA1 and PFC coding and non-coding states was uncorrelated (Figure S5 H), suggesting an independence of CA1 and PFC state onsets.

We then focused on PFC coding states in overtrained sessions, and looked at which moments those states were most active along the trial (Figure 7B). In CS+ trials, while CS- coding states were active during pre-stimulus period, CS+ states were silent. After stimulus onset, this relationship inverted, with CS+ becoming more active. In CS- trials, CS- states seemed to be active during both pre-stimulus, stimulus and trace, while CS+ states were only briefly more active after stimulus onset. These results are in accordance to the poor decoding previously found in CA1, ruling out the possibility of stochastic coding states (i.e., not time-locked with the stimulus) dominating CA1 activity. Also, we conclude that encoding stimulus identity by coding states in PFC is biased towards the relevant stimulus of the task (i.e., CS+).

Onset of stimulus-coding neural states predicts behavioral performance

We next investigated whether one could reliably predict the animal's behavior in single trials from the specific onset of PFC stimulus-coding states. We hypothesized that the activation of a CS+ neural representation could lead to the behavioral expression of the conditioned response. Specifically, we examined whether the onset of a CS+ coding state in a correct CS+ trial preceded the onset of the first lick. Likewise, we tested whether the

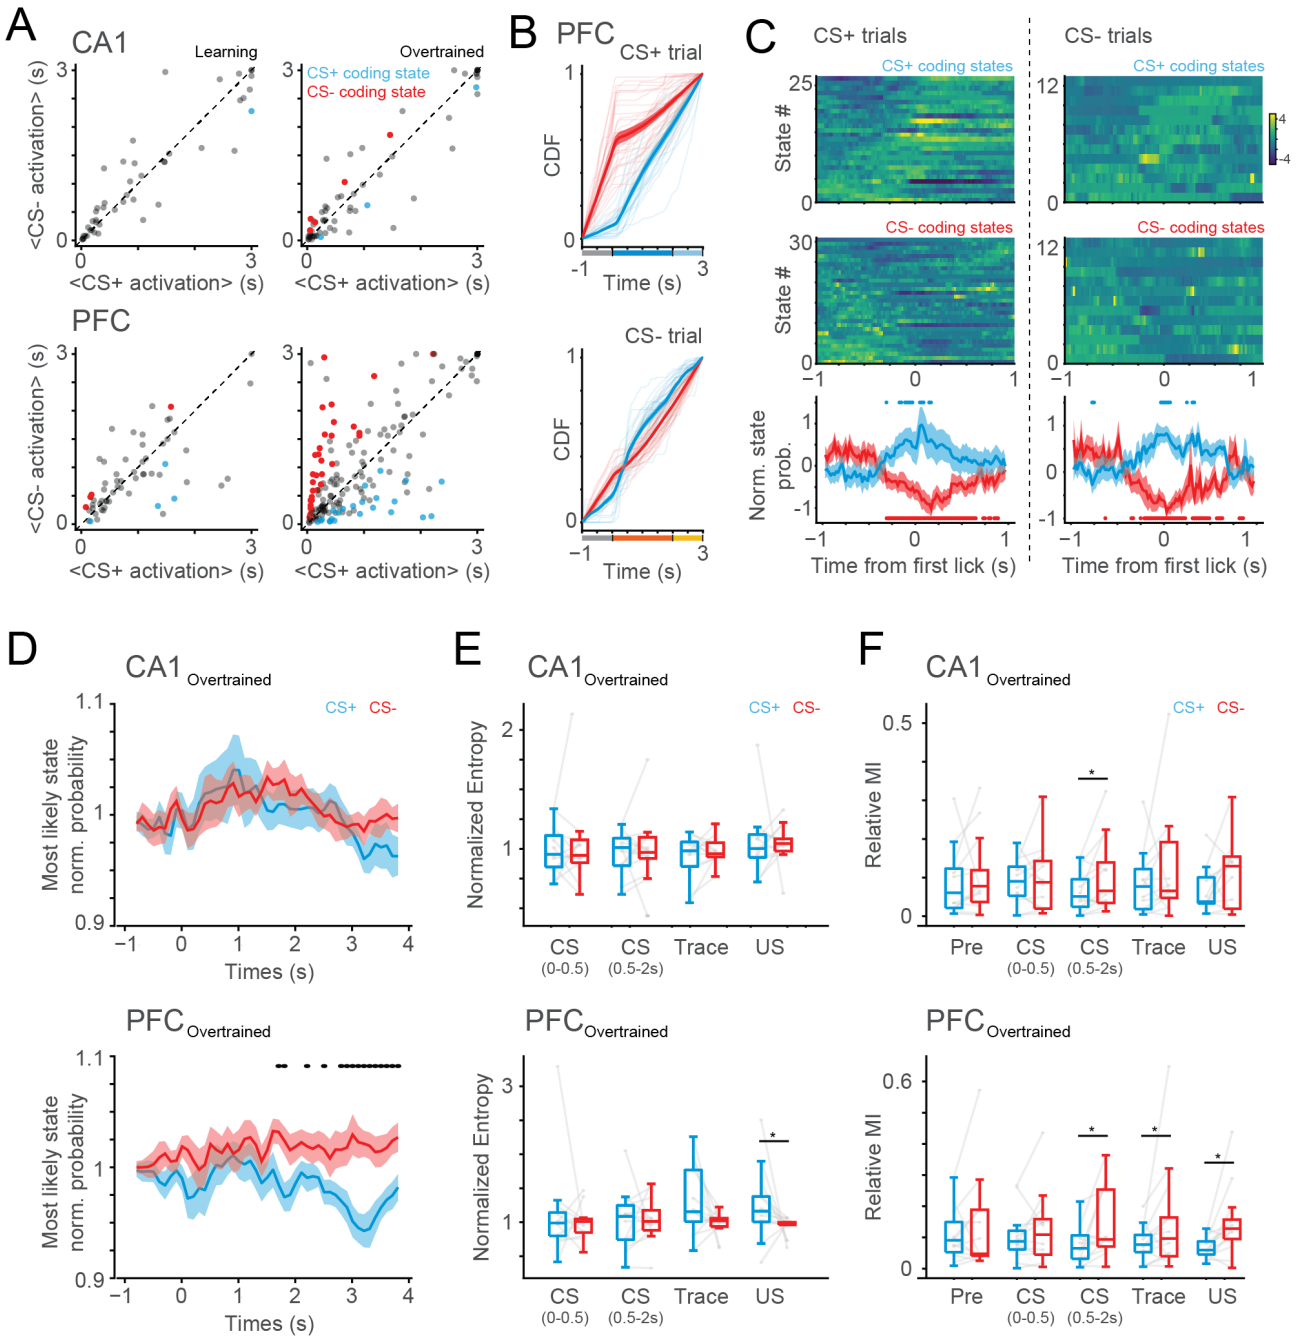

Figure 7: **Emergence of CS coding states in PFC after learning. A.** Average CS+ and CS- activation during trace and stimulus period for every state in learning (left) and overtrained (left) sessions. Color denotes CS+ (blue) and CS- (red) coding states. Note the strong increase in the number of coding states for PFC. **B.** Cumulative distribution function of coding states in PFC. CS- coding states show increased activation during pre-stimulus period while CS+ states are mostly active after stimulus presentation. **C.** Coding state's activation on overtrained sessions centered on the first lick of CS+ (left) and CS- (right) trials. Notice that activation and inactivation of, respectively, CS+ and CS- coding states precedes the first lick of the trial. This can not be explained by proximity with stimulus onset (see methods). Dots represent significantly different state activation probability than surrogate distribution ($p<0.05$; blue: CS+ coding states; red: CS- coding states). Shaded areas denote SEM. **D.** Most likely state probability over the trial for CA1 (top) and PFC (bottom) overtrained sessions during CS+ and CS- trials. Black dots represent significantly different CS+ and CS- state activation probability. Shaded areas denote SEM. **E.** Entropy of most likely state for CA1 and PFC during different trial periods and normalized by pre-stimulus entropy. Blue and red denote CS+ and CS- trials. **F.** Mutual information of most likely states in CA1 and PFC for different trial periods. Mutual information was expressed in relative terms, that is, normalized by the total entropy in CA1 (top) or PFC (bottom).

onset of erroneous licks in CS- incorrect trials was also preceded by the activation of a CS+ state. Figure 7C shows the (PFC) coding-state probability normalized by a surrogate distribution and centered at the time of the first lick. To rule out effects of stimulus onset, only first licks happening 1 s after stimulus were considered. We found that the probability of CS+ coding states increased before lick onset in both CS+ (correct) and CS- (incorrect) trials. Analogously, CS- coding state probability consistently decreased before lick onset. This effect could not be explained by delays in lick onset detection, since the minimum inter-lick interval happened on a shorter scale (Figure S5E). We thus concluded that the transient activation of stimulus-specific neural assemblies in PFC predicts the behavioral expression of the conditioned response in single trials.

CA1-PFC states are less coordinated during the task-relevant moments

Finally, we investigated whether CA1 and PFC states were more or less coordinated during the task. To do that, we focused our analyses on the hidden states found in the simultaneous recordings. We found that the probability of the most likely state, normalized by pre-stimulus period, evolved differently across CS+ and CS- trials in PFC. In CS+ trials, PFC activity would result in less certain states as the trial approached the US (Figure 7D). Interestingly, this effect could also be seen looking at the (normalized) entropy of the most likely state in PFC during the different parts of the trial (Figure 7E). There was a tendency to a greater entropy during the trace period, and a significant increase in the PFC entropy from the US.

We then computed the mutual information between the most likely state of CA1 and PFC for the different parts of the trial. We normalized the mutual information (MI) by the entropy of each area and found that the relative MI in PFC decreased in the late part of the stimulus (0.5-2 s) and trace period during CS+ trials (Figure 7F). Because this decrease was not reflected on the non-normalized MI (Figure S5G), it indicate an CA1-independent increase of PFC entropy. Together those results suggest that the two areas seem to be independently coordinated after the second part of the stimulus, with state changes in PFC happening independently from CA1.

Discussion

In this work we elucidated how the dynamics of stimulus and temporal coding in CA1 and PFC are shaped by learning during an auditory trace conditioning task. We found that both areas encode time via temporal sequences, evoked by stimulus and reward delivery. Interestingly, we found evidence that stimulus and reward sequences may represent surprise signals, whose representations are similar in naive animals, and progressively acquire specificity during the course of learning. We found that learning induces the emergence of a persistent and stable stimulus representation after learning in mPFC, bridging the gap from stimulus and to reward across the trace period. Looking for stimulus representation on neural states in single-trials showed that transient activation of stimulus-coding states in mPFC reliably predicted the onset of the first lick response, providing a link between neural coding and behavior.

Time encoding and surprise signals.

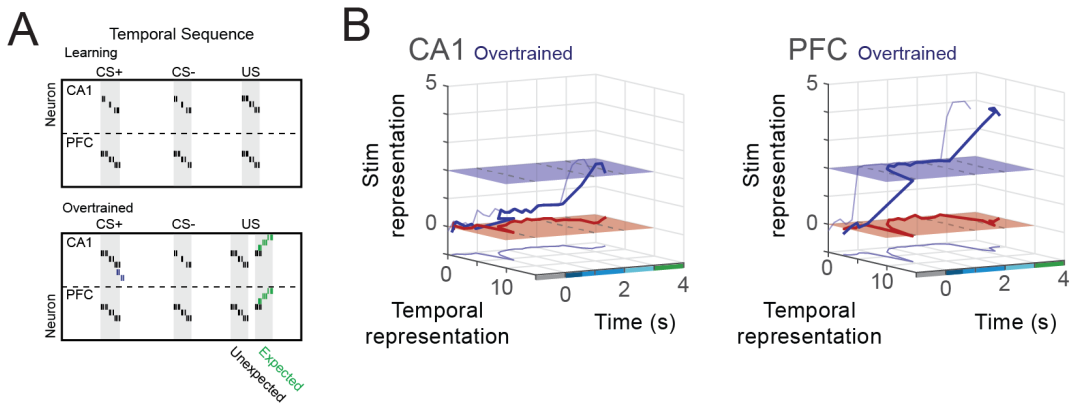

Figure 8: **Schematic representation of temporal and stimulus coding in CA1 and PFC A.** Joint temporal dynamics of time and stimulus representations in overtrained sessions. In CA1 overtrained sessions the stimulus representation of CS+ and CS- lie most of the time in the same plane, starting to increase during trace, and peaking at reward consumption, where CS+ and CS- are differentiable. In contrast, PFC stimulus representation of CS+ differentiates from baseline and CS- representation right after stimulus onset. This effect was also accompanied by a peak on temporal representation, present after CS+ and CS- onset on both areas. The temporal representation also peaked after US onset. **B.** Learning dynamics of temporal sequences. The temporal sequences of CS+, CS- and US are initially similar during learning sessions. After learning, a stronger temporal sequence appears in CA1 for CS+, and US temporal representation disentangle from CS+ and CS-, but only for expected events.

Since time is an important component of declarative memory, and given the role of the hippocampus in this type of memory, it's not surprising to find some form of temporal coding in this area. In fact, hippocampal neurons can have timely coordinated spikes in the form of phase precessing place cells and time cells (Eichenbaum (2014), O'Keefe and Recce (1993)). The increase in temporal decoding we found in CA1 upon learning indicates that the learning process sculpts the transient response to the CS in a behavior-dependent way, making it more temporally unique, and longer-lasting (Figure 2). It has been shown that the hippocampus (Solomon et al. (1986), Bangasser et al. (2006)), and specifically CA1 (Kesner et al. (2005)), is needed to create the CS-US associations in trace conditioning tasks. In accordance with this, previous works suggest that time cells may bridge stimulus representation during the trace period (Pastalkova et al. (2008b), MacDonald et al. (2011, 2013), Kitamura et al. (2015), Sellami et al. (2017)). Although we found temporal representations of the stimulus, differently from those other works, they faded after the first 0.5 s of stimulus presentation. This is similar to other CA1 studies showing no time encoding in calcium traces bridging CS and US during trace fear conditioning (Ahmed et al. (2020)). In our case, our shorter time-scale analyses allowed us to extend those findings and find the fast onset and offset of temporal sequences. This might indicate a lower hippocampal load in this particular task and an CA1-independent maintenance of stimulus representation during the trace period, as manifested in mPFC.

Interestingly, the increase in time encoding with learning was seen specifically after CS+ presentation - the relevant stimulus of the task - and not after CS- or US (which also presented high values of temporal decoding, but no changes after learning) (Figure 2C; see also Figure S2). It has been previously shown that CA1 is important to activate cortical memory representations during memory retrieval (Tanaka et al. (2014)). It might be that the emergence of a stronger temporal sequences is the mechanism used to process the sensory

information arriving in CA1 from the entorhinal cortex and propagate them to the PFC, where the stimulus-context associations are held (Euston et al. (2012)).

Moreover, before learning CS temporal sequences seemed to be echoed during the US presentation indicating that initially the same temporal pattern was elicited after all the three events (CS+, CS- and US) (Figure 3). After learning, CS and US similarity drastically decreased in both areas, except during putative incorrect trials where their pre-learning similarity was partially recovered (Figure 3). This suggests that, in error trials where the CS+ stimulus identity was likely not encoded and thus the animal did not expect a reward, US-evoked sequences may encode a surprise signal, more similar to the CS-evoked sequence, as it occurred before learning. Additionally, this effect seems to be supported by changes in US temporal sequence, while CS+ and CS- shared representation increases with learning (Figure 4).

Together those results suggest a common temporal component triggered by surprise. For example, it might be that initially unexpected events (CS or US) evoked a partially overlapping temporal sequence, whose shared element represents a 'surprise' signal. After learning, the surprise element wanes from the US since it then becomes expected (putative correct trials), leading to the emergence of stimulus-specific responses and enabling a decorrelation between CS+ and CS- (Figure 8) and between those and the US. It is worth to note that, despite the difference between correct and incorrect trials in US echo in CS+, there was no such difference for the US echo in CS- (Figure S3). This suggests that US temporal sequence are differently captured by CS+ and CS- decoders, and there might be an additional component beyond 'surprise' playing a role in the similarity between CS+ and US.

Stimulus coding via mPFC persistent activity.

Even tough temporal sequences encoding the outcome signal of the task (i.e., the presence or not of the US) are present in both CA1 and PFC at the end of the trial, those two areas show different stimulus encoding profiles. While in overtrained animals CA1 decoders only had a mild increase in performance during trace period, mPFC showed a strong and stable stimulus decoding since CS presentation (Figure 5). Thus, beyond the outcome of the task, mPFC neurons also contained information about the stimulus, even in its absence, which is in accordance to previous works (Asaad et al. (2017)) and further support the idea that mPFC holds the stimulus-action maps.

mPFC stimulus representation could be generalized from stimulus to trace period, suggesting the neuronal dynamic encoding the stimulus in this area reach a stable attractor state. On the other hand, CA1 decoders could not be well-generalized for different time periods. Could the poor decoding performance in CA1 be explained by fast stimulus-specific patterns reactivated stochastically at random times (i.e., not time-locked with the stimulus)? A single-trial pattern analysis based on hidden Markov models failed to reveal stimulus-coding states, thus making the hypothesis of stochastic reactivations unlikely (7). It might be that in mPFC assemblies compete with each other in a winner-takes-all dynamic, with the strongest assembly dominating the state of the network, while in CA1 there are different independent patterns happening concomitantly, making it hard to identify independent states. This is in accordance with the lower and higher HMM posterior probabilities found in CA1 and PFC, respectively, (Figure S5D and F), corroborates previous works showing

neuronal representations encoding different context in mPFC are well separated while highly overlapping in CA1 (Hyman et al. (2012), Leutgeb et al. (2004)), and is consistent with previous findings showing very sparse stimulus selectivity in CA1 (Ahmed et al. (2020)).

PFC showed the presence of both CS+ and CS- coding states. Interestingly, we found that CS- coding states transiently occurred during pre-stimulus period as well (Figure 7 B)). On the other hand, the occurrence of CS+ coding states was limited to the CS+ stimulus epoch. Since the activation of a CS+ coding state was reliably preceding of the upcoming first lick of the trial, we can rule out that it is driven by backdrop of movements reflected in lick modulated neurons as previously reported in mPFC (Figure 7 C; see also the inter-lick interval in Figure S5E). Crucially, we found that the onset of (CS+)–coding states in single trials reliably predicted upcoming erroneous licks in CS- error trials. This suggests the presence of a causal relationship between transient activation of stimulus representations in mPFC and an animal decision to lick.

CA1-PFC coordination.

We did not find evidence for coordination between PFC and CA1 states (e.g., no increased correlation between coding states; (Figure S5 H). In fact, we found a decrease in coordination emerging after learning, during and after the late stimulus (decrease in the mutual information between the most likely state in CA1 and PFC (Figure 7 F). Thus, after the initial stimulation the entropy of PFC in CS+ trials increased without any change in the mutual information between PFC and CA1, suggesting a lesser role for CA1-PFC communication.

It seems that the presence of coordinated activity between those two areas may vary according to the type and period of the task (e.g., sleep, sharp-wave ripples, sensory encoding), which would explain previous work showing both dependent Jadhav et al. (2016), Peyrache et al. (2009), Shin et al. (2019) and independent (Kaefer et al. (2020), Klee et al. (2021)) co-activation of neuronal patterns in those areas. In our case, any interplay between those two areas during the late stimulus or trace period could not be captured by HMM states. This is consistent with a scenario where CA1 encodes representations of multiple task-variables in parallel or in a temporally continuous rather than discrete way. An HMM analysis would likely fail in these scenarios, and alternative methods such as factor analysis (Yu et al. (2008)) or matrix factorization (Mackevicius et al. (2019)) might help overcome these issues.

In summary, our findings unveiled rich and complex dynamics in CA1 and mPFC during an auditory trace conditioning, opening new interesting questions. Our results suggests that in our trace conditioning task CA1-mPFC interplay might happen mostly following the CS presentation, when time encoding is higher in both areas (Figure 8A). After learning, strong, sustained and stimulus-specific representations emerged in mPFC, bridging across the trace period, with neural states predicting the animal behavior in single trials (Figure 8B). Because those temporal representations also encoded the outcome of the task, the emergence of stronger and more specific temporal sequences in CA1 might support stimulus encoding and help to organize the following stable memory representation in PFC.

Materials and Methods

Animals

In this study we used a total of 17 male C57Bl/6 mice. From this total, seven animals were used for (acute) silicon probe recordings in the dorsal hippocampus, six for silicon probe recordings in the medial prefrontal cortex, and four additional animals were used for simultaneous silicon probe recordings. Animals were kept in a reversed light/dark cycle, and had access to food ad libitum until two days after the surgery, when food restriction was introduced (90% of initial body weight). A detailed description of experimental procedures can be found in Klee et al. (2021). This study was approved by the Dutch Central Commissie Dierproeven (CCD) and conducted in accordance with the Experiments on Animals Act (Dutch law) and the European Directive 2010/63/EU on animal research.

Surgical Procedures, behavioral training and data collection

Animals were kept under isoflurane (1-2%) anesthesia and placed in the stereotaxic frame. Carprofen (5 mg/kg) and lidocaine were subcutaneously injected in the area of the scalp before skull exposure. A custom-made circular head-fixation plate was placed on the skull and fixated with dental cement (Super-Bond C&B). A skull screw was placed over the cerebellum to serve as ground and reference for the later electrophysiological recordings. For animals used in CA1 recordings, a craniotomy was performed at the left hippocampus (-2.3 mm posterior and +1.5 lateral to Bregma). For animals used in PFC recordings the craniotomy was done over the left frontal cortex (+1.78 mm anterior and +0.4 lateral to bregma). A silicon elastomer (Body Double Fast, Smooth-on) was then used to covered the exposed skull until the first recording. After at least 2 days of recovery from surgery, animals began habituation to the head-fixed setup which consisted of two rods that could be screwed to either side of the implanted head-plate. In the head-fixed setup animals could move through a virtual linear track using a air-supported spherical treadmill. Feedback from treadmill movement was projected in a spherical screen surround the animal head. A plastic reward port delivering soy-milk was placed 0.5 cm aways from the posterior lip of the animal. Animals underwent at least 6 habituation sessions (3 sessions of 10 min per day) in which they received about 0.2 ml of soy milk before the beginning of trace electrophysiological recordings. The appetitive auditory trace conditioning task required the animals to associate a specific sound stimulus (the positive conditioned stimulus; CS+) to a later reward (unconditioned stimulus; US), in contrast to a second stimulus not associated with the reward (CS-). In our task, CS+ and CS- stimuli lasted for 2 s, and in the case of CS+ trials, the stimulus was followed by a silence period of 1 s (trace period) and then US delivery. Trials were then structured as 1 s of pre-stimulus, 2 s of stimulus, 1 s of trace, and 1 s of reward period. Learning of the task was accessed via anticipatory licking, detected using a infrared sensor beam in front of the reward spout. Sessions in which the amount of anticipatory licks during the trace period in CS+ trials was significantly different than in CS- trials (t-test) were classified as overtrained, while the others were classified as learning sessions. CS+ and CS- trials were presented in a pseudo-random order, and counterbalanced during the session.

Neural data recording and analysis

At the beginning of each recording session the mice were head-fixed and had the elastomer removed for skull exposure. One or two 128-channel silicon probes were then acutely inserted above the previously prepared craniotomies using a micromanipulator. For mPFC recordings, electrodes were slowly lowered up until -2.0 mm ventral to Bregma. In CA1 recordings the electrophysiological signal was monitored for high power over the ripple frequency band (150-300 Hz), and spiking activity (600-2000 Hz). Electrodes were placed so that the highest ripple power and spiking activity were localized in the middle of the probe. Spike sorting was automatically done using Kilosort and visually inspected using 'phy' gui (https://github.com/cortex-lab/phy). Following analysis, unless specified otherwise, were done using custom-made Matlab scripts.

Decoding of trial time

Temporal coding at the single session level was assessed through the ability of decoding the time of the trial using a Bayesian classifier. For each session, we first binned trials into 100 ms bins with no overlap. The number of trials for each stimulus was balanced between conditions. A naive-Bayesian classifier was trained with a leave-one-out cross validation scheme in for each iteration a pair of CS+/CS- trial was used to test. The classifier modeled neuronal firing as an inhomogeneous Poisson process whose rate was a function of trial time and neurons were assumed to be independent given trial time.

Using test samples of a given time bin, we estimated its probability of being decoded as each of the other time bins of the trial. We then stacked the probability distributions of all the time bins into a confusion matrix ('time-decoding' matrix), pseudo-colored with the probability of predicting a tested time bin in the x-axis, as one of the different time bins in the y-axis. Decoding probability values of each pair of time bins were z-scored using the mean and the standard deviation from a surrogate distribution (100 surrogates). Surrogates were computed by shuffling the time bin identity of each trial in the training set before computing the Bayesian decoder in each fold.

Temporal sequences could be seen in the time-decoding matrix as high decoding values concentrated near the main diagonal. We selected two windows to investigate the emergence of temporal sequences, one at the stimulus onset (0-0.5 s; 100-ms-bins) and another after reward onset (3.1-3.6 s; 100-ms-bins). Temporal sequence strength was then defined as the mean probability over the main diagonal of time-decoding matrices at those windows. Reward window was delayed 100 ms from reward presentation based on the echo of the stimulus temporal sequence, which started at this moment. This was probably due to the fact that the animal needs to lick the reward to perceive it, which can introduce a sampling delay.

To assess significance of the temporal sequences, their strength was tested against the strength of a null temporal sequence of equal size, extracted from pre-stimulus period (Wilcoxon sign-rank test). Comparison between temporal sequence strength across different areas and learning condition was done using aligned-rank transform in order to normalize non-parametric data. The equivalent of a two-way Anova was then performed (see ARTool toolbox, Wobbrock et al. (2011)).

US echo in CS+ (or CS-) temporal sequences was evaluated by analyzing the probabilities of decoding time

bins of reward period (after US presentation) as time bins up to the end of trace period. This was done to avoid influences on the strength of the US temporal sequence itself on the echo estimation. Echo strength was computed in a similar way as the temporal sequence strength.

Decoding of stimulus identity

We performed a cross-validated classification analysis to investigate the stability of stimulus representations. Neuronal activity was divided into bins of 200 ms with no overlap and used to train an ensemble of support-vector machines (SVM) binary classifiers. Each binary classifier was trained to classify CS+ and CS- trials using the activity of a given time bin. We then used test bins from previous, subsequent, or current times to see how good the prediction could be generalized. The number of CS+ and CS- samples was always balanced and paired prior to training. Because the size of sample pool was different when training and testing in the same time bin compared to training and testing with different time bins, we used different cross-validation schemes for those two situations. In the first case (training and testing in the same time bin), we used a leave-one-out cross-validation, applied to pairs of CS+ and CS- trials. In the second case (training and testing bins different), we used the entire pool of trials to train, testing in all trials from the other time bin.

For each of those two cases we also trained surrogate classifiers (100 surrogates) in which the identity of the training trials was randomized. We then used the surrogate distribution of performance to normalize (z-score) the performance values of each classifier. Similar to the time decoding matrices, classifiers with different training and testing bins were stacked together forming a trial-type confusion matrix with training time bin in the y-axis and testing (time) bin in the x-axis. Notice that in those matrices we omit the main diagonal, that has different sample statistics and cross-validation scheme. The main diagonal is shown separately.

Computing hidden Markov models

To investigate whether stimulus identity was encoded by fast co-activation of neuronal assemblies, we first binned spikes in 20-ms bins with no overlap. HMMs were computed using Linderman et al. (2020) and assuming that neuron's firing rate followed a Poisson distribution. Because during training of HMMs each trial is initialized in the same state, we included an additional second before the stimulus in those analyses to account for stabilization of states (i.e., training was done using 2 s before CS up to 1 s after US). For each session, CS+ and CS- trials were balanced and split in a 5-fold cross validation scheme. The Bayesian information criterion (BIC; Recanatesi et al. (2022)) was used to select the best number of hidden states of each model (ranging from 2 to 20). Using Akaike information criterion instead held more states, but qualitatively similar results (data not shown).

After selecting the number of hidden states, we used the best model to predict the posterior probability of each state across trials. We then defined a state as 'detected' whenever the probability of that state was above 80%, allowing us to have a mean state activation during stimulus and trace period (Mazzucato et al. (2019, 2015)). We looked for hidden states with significantly different activation in CS+ and CS- trials during the stimulus and trace periods ($p<0.05$, Wilcoxon ranksum test). For this analysis data was preprocessed defining valid trials for each state as trials with at least one activation. This was done to avoid permanent state transitions due to changes in recording stability occurring across the session. To account for that, for each

state only trials with at least one activation were taken into consideration. CS+ and CS- coding states were then defined as states that were significantly more active during CS+ and CS- ($p<0.05$; Wilcoxon signed-rank tests), respectively (we found similar results when restricting those analysis to one single block of stable trials per session).

We also computed the state activation triggered by the first lick of the trial focusing on CS+ and CS- coding states. To avoid effects driven by the stimulus, we restricted this analysis to trials in which the first lick happened at least 1 s after the stimulus onset. Also, we computed the inter-lick interval distribution of the licks pooled across all animals to have an upper bound estimation of the delay between the onset and detection of the licks (Figure S5 E).

We used the simultaneous recordings to compute the correlation and mutual information between PFC and CA1 hidden states. To compute the mutual information across the trials we first computed the most likely state of each 20-ms bin (i.e., the state with higher probability). Then, the adjusted mutual information was computed in bins of 100 ms (5 20-ms bins), separately for CS+ and CS- trials.

Supplementary Figures and Methods

Measuring dimensionality of temporal representations

To calculate how the dimensionality of the stimulus representation varied over the trial. For that, we used principal component analysis to computed a relative dimensionality of the populational activity for each time bin. This measure was defined as ratio between the amount of principal components that accounted for 90% of the variance in a given time bin and the respective amount computed during the pre-stimulus period. We found that both during CS+ and CS- trials, CA1 neural activity showed a decrease in dimensionality after the stimulus onset (i.e., when temporal sequences appeared) and this effect was present both before and after learning. On the other hand, PFC showed a similar decrease, but restricted to overtrained sessions (Figure S2B).

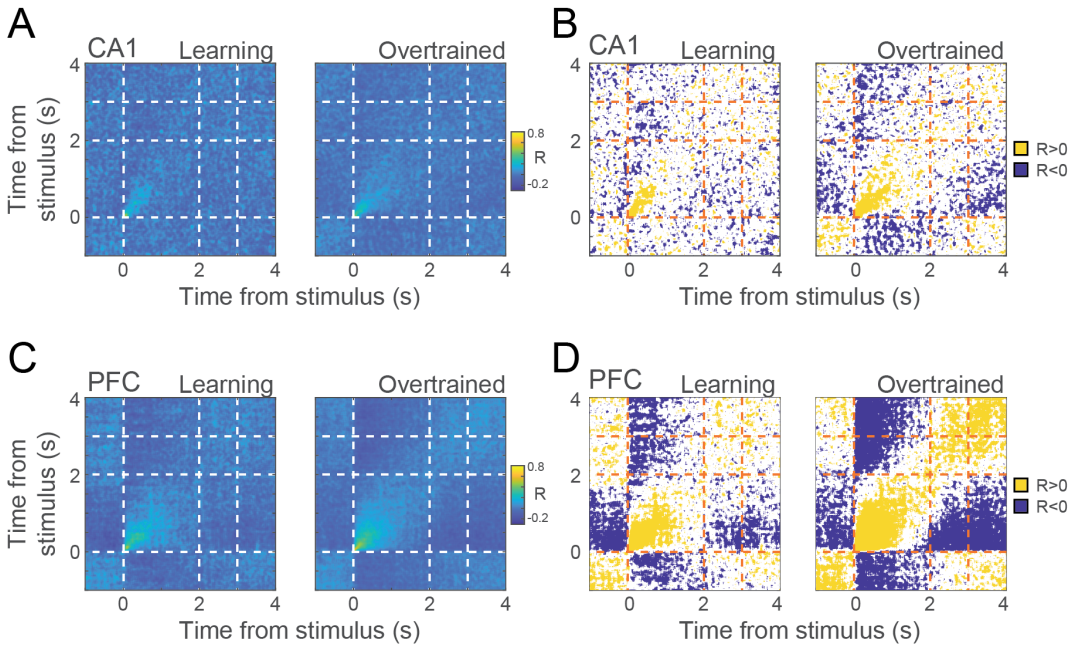

Figure S1: **Temporal structure in pooled activity.** **A.** Correlation between mean activity of all CA1 neurons on even and odd trials before and after learning (see Figure 1D). **B.** Significant correlation values ($p<0.05$ compared from a surrogate distribution) for A. **C** and **D** same as A and B, but for PFC.

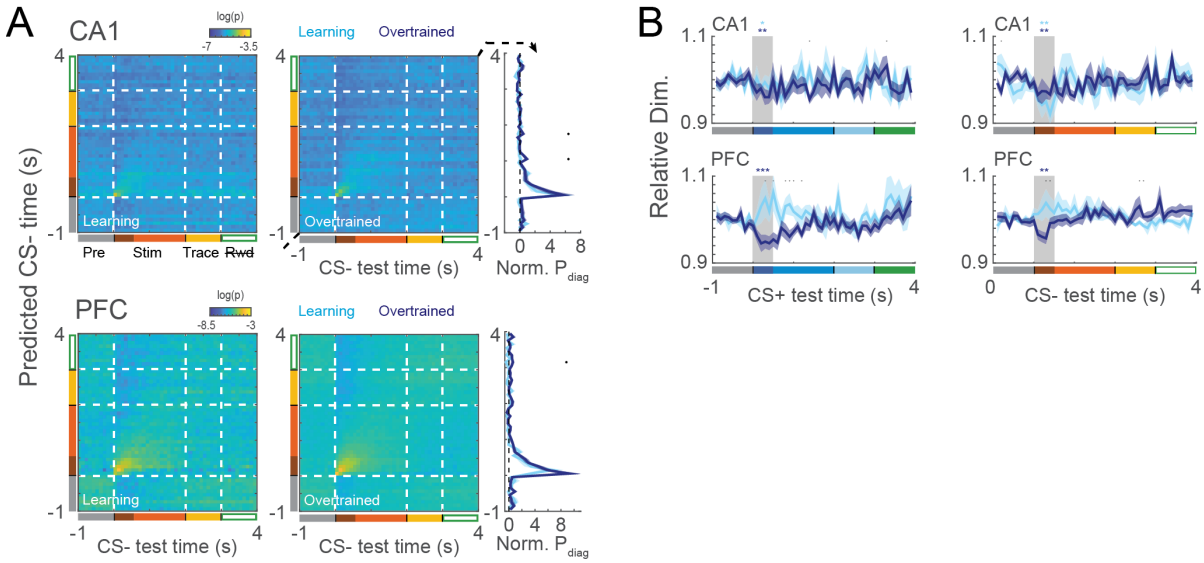

Figure S2: **Emergence of temporal sequences in CS-.** **A.** Mean time decoding matrices for learning and overtrained sessions together with the average probability on the main diagonal (rightmost panel). Decoders were trained and tested using CS- activity from CA1 (top) and PFC (bottom). Black dots denote times in which learning and overtrained decoding are different (Wilcoxon ranksum, $p<0.05$). **B.** Relative dimensionality of the population activity over time for CS+ (left) and CS- (right) trials. Relative dimensionality was measured as the fraction of principal components necessary to explain 90% of the data variance over each time bin. Note the lower dimensionality after stimulus presentation (in the same period of CS+ and CS- temporal representations) in CA1 before and after learning. PFC sessions showed a similar decrease only in overtrained sessions.

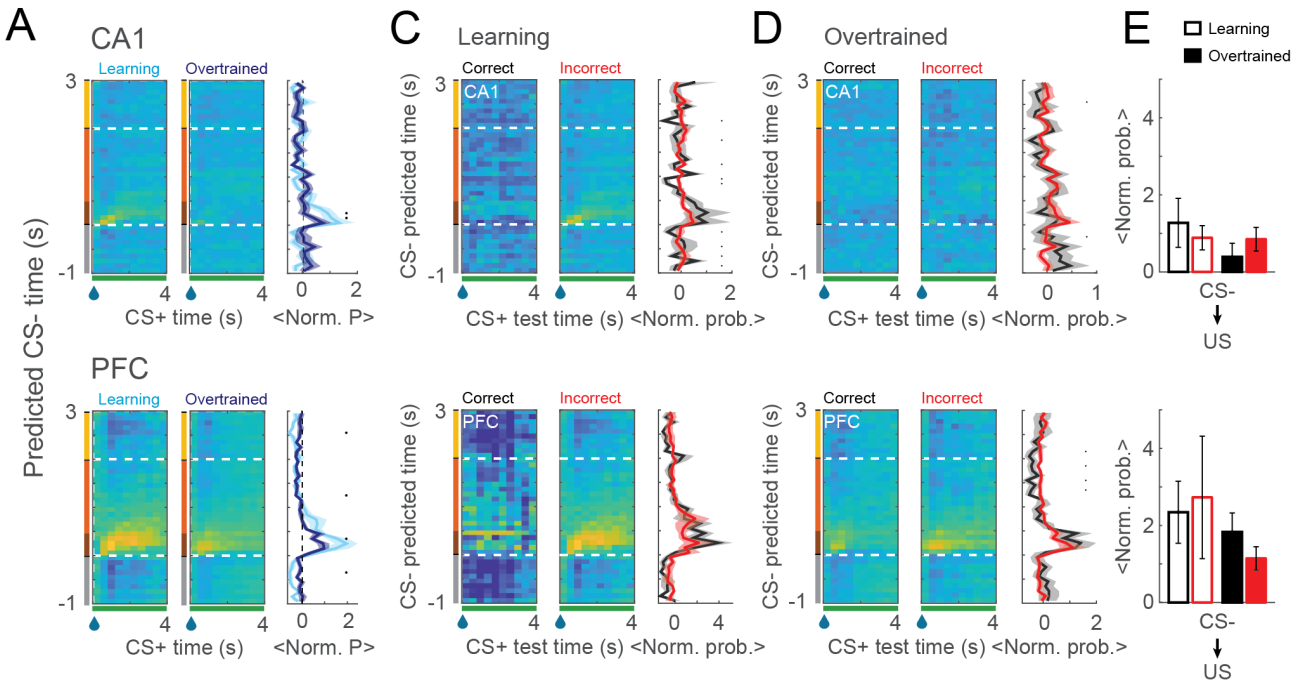

Figure S3: **Similarity between CS- and US temporal representations is not modulated by learning or error signal of trial outcome. A.** Similar to Figure 3, but using CS- trials for training. Time decoding matrices trained using CS- activity before the US, and used to decode US time (left and middle). Average (normalized) probabilities for the US period are shown for learning and overtrained sessions (right). **B.** Mean decoding in the temporal echo window for CA1 and PFC. **C-D.** Time-decoding matrices as in A, but showing the temporal echo in correct and incorrect trials from learning (C) and overtrained (D) sessions. **E.** Average performance on the temporal echo window (main diagonal) for CA1 (top) and PFC (bottom). Aligned rank transform followed by a two-way ANOVA using trial outcome (i.e., correct vs. incorrect) and learning condition as factors for each area independently revealed no significant interaction or main effect in either area. Shaded areas and errorbars denote SEM. Black dots denote $p<0.05$.

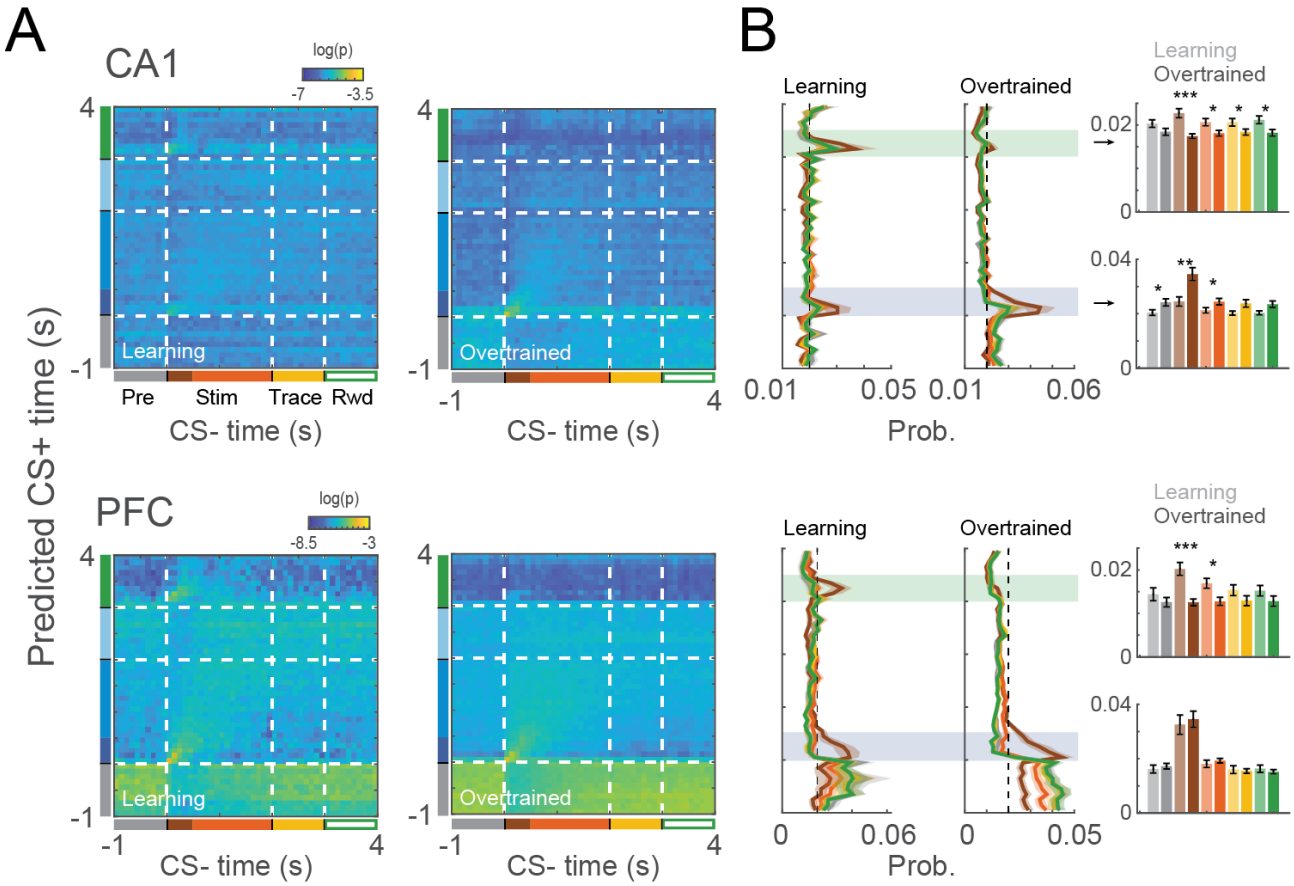

Figure S4: **US echo in CS- before and after learning. A.** Time decoding matrices trained using CS+ activity before the US, and used to decode US time in CS- trials during learning (left) and overtrained (right) sessions for CA1 and PFC. **B.** Average decoding probability for testing bins in different periods of the trial (right and middle), together with the mean probability in the CS and US temporal windows (shaded rectangles). Notice the decrease of similarity between CS- stimulus transient and US after learning for both areas (dark brown bars). Shaded areas and errorbars denote SEM. *: $p<0.05$; **: $p<0.01$; ***: $p<0.001$; Wilcoxon ranksum test.

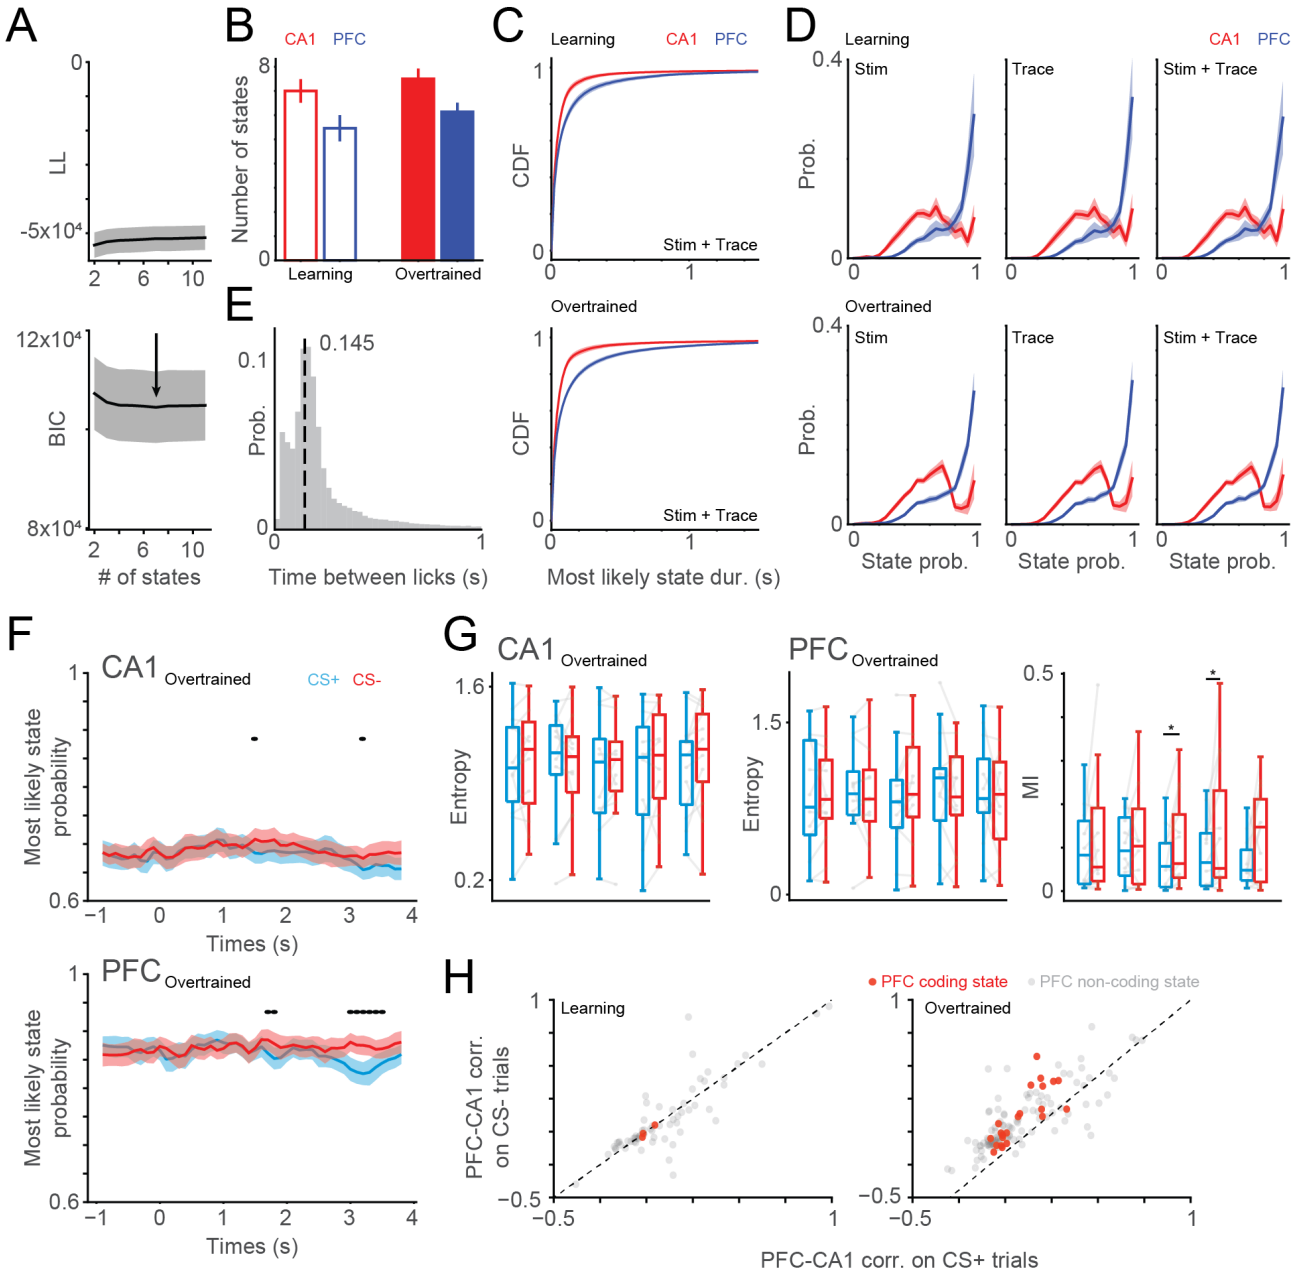

Figure S5: **HMM states characterization.** **A.** HMM model selection was done using the mean Bayesian information criterion (BIC) over the different cross-validation folds. The model with lowest BIC was used. **B.** Number of states in CA1 and PFC before and after learning. **C.** Cumulative distribution of most likely state duration during CS presentation and trace period. **D.** Most likely state probability distribution during CS, trace and both periods together. Note the higher values in PFC states. **E.** Distribution of the time between consecutive licks. Compare with PFC coding state activation before the first lick **F.** Most likely state probability over the trial for CA1 (top) and PFC (bottom) overtrained sessions during CS+ and CS- trials. **G.** Entropy (left, middle) and mutual information (right) of most likely state for CA1 and PFC during CS+ and CS- trials. **H.** Correlation between CA1 states with coding (red dots) and non-coding PFC states.

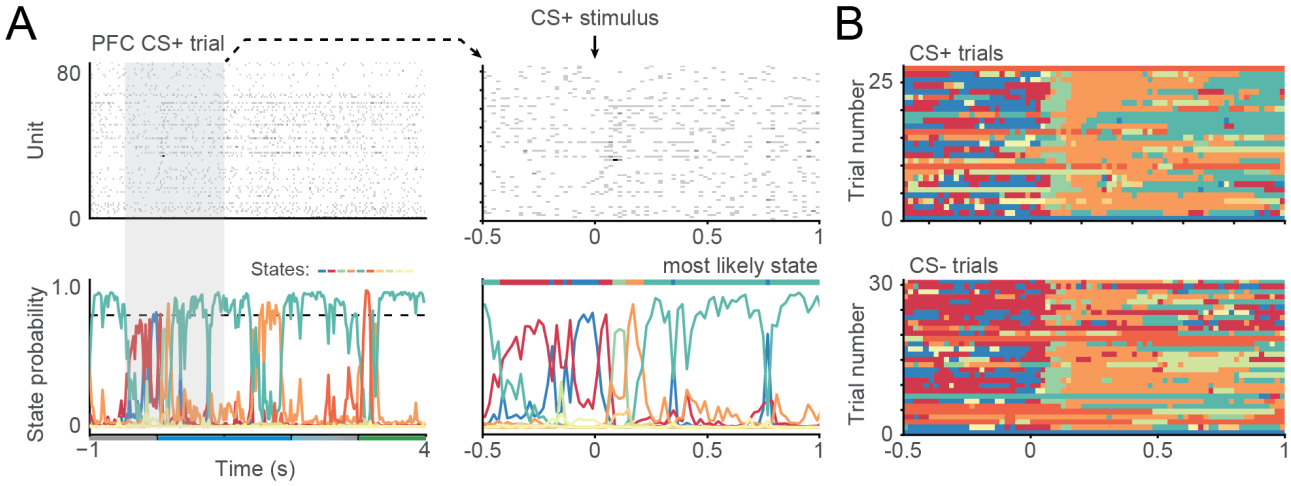

Figure S6: **Temporal sequences in HMM states. A.** (Left) Example of neuronal activity (top) and HMM state probabilities (bottom) over a trial showing the occurrence of temporal sequences after stimulus onset. (Right) The same as in left, but zooming over stimulus onset. Coloured bar in bottom codes for the most-likely state. **B.** Most-likely states around the stimulus for CS+ (top) and CS- (bottom) trials. Noticed that even though the probability of some states don't reach the 0.8 threshold (dashed line in A, bottom), the sequential activation of states is seen at the most-likely state across trials.

References

Mohsin S Ahmed, James B Priestley, Angel Castro, Fabio Stefanini, Ana Sofia Solis Canales, Elizabeth M Balough, Erin Lavoie, Luca Mazzucato, Stefano Fusi, and Attila Losonczy. Hippocampal network reorganization underlies the formation of a temporal association memory. *Neuron*, 107(2):283–291, 2020. Publisher: Elsevier.

Wael F Asaad, Peter M Lauro, János A Perge, and Emad N Eskandar. Prefrontal neurons encode a solution to the credit-assignment problem. *Journal of Neuroscience*, 37(29):6995–7007, 2017. Publisher: Soc Neuroscience.

E. H. Baeg. Fast Spiking and Regular Spiking Neural Correlates of Fear Conditioning in the Medial Prefrontal Cortex of the Rat. *Cerebral Cortex*, 11(5):441–451, May 2001. ISSN 14602199. doi: 10.1093/cercor/11.5.441. URL `https://academic.oup.com/cercor/article-lookup/doi/10.1093/cercor/11.5.441`.

Debra A Bangasser, David E Waxler, Jessica Santollo, and Tracey J Shors. Trace conditioning and the hippocampus: the importance of contiguity. *Journal of Neuroscience*, 26(34):8702–8706, 2006. Publisher: Soc Neuroscience.

Francesco P Battaglia, Karim Benchenane, Anton Sirota, Cyriel MA Pennartz, and Sidney I Wiener. The hippocampus: hub of brain network communication for memory. *Trends in cognitive sciences*, 15(7):310–318, 2011. Publisher: Elsevier.

Alberto Bernacchia, Hyojung Seo, Daeyeol Lee, and Xiao-Jing Wang. A reservoir of time constants for memory traces in cortical neurons. *Nature Neuroscience*, 14(3):366–372, March 2011. ISSN 1546-1726. doi: 10.1038/nn.2752. URL `https://www.nature.com/articles/nn.2752`.

Brian G. Burton, Vincent Hok, Etienne Save, and Bruno Poucet. Lesion of the ventral and intermediate hippocampus abolishes anticipatory activity in the medial prefrontal cortex of the rat. *Behavioural Brain Research*, 199(2):222–234, May 2009. ISSN 0166-4328. doi: 10.1016/j.bbr.2008.11.045. URL `https://www.sciencedirect.com/science/article/pii/S0166432808006621`.

Howard Eichenbaum. Time cells in the hippocampus: a new dimension for mapping memories. *Nature Reviews Neuroscience*, 15(11):732–744, 2014. Publisher: Nature Publishing Group.

Howard Eichenbaum. Memory: organization and control. *Annual review of psychology*, 68:19–45, 2017a. Publisher: Annual Reviews.

Howard Eichenbaum. Prefrontal–hippocampal interactions in episodic memory. *Nature Reviews Neuroscience*, 18(9):547–558, 2017b. Publisher: Nature Publishing Group.

David R Euston, Aaron J Gruber, and Bruce L McNaughton. The role of medial prefrontal cortex in memory and decision making. *Neuron*, 76(6):1057–1070, 2012. Publisher: Elsevier.

Shintaro Funahashi, Matthew V Chafee, and Patricia S Goldman-Rakic. Prefrontal neuronal activity in rhesus monkeys performing a delayed anti-saccade task. *Nature*, 365(6448):753–756, 1993. Publisher: Nature Publishing Group.

Marieke R. Gilmartin and Matthew D. McEchron. Single neurons in the medial prefrontal cortex of the rat exhibit tonic and phasic coding during trace fear conditioning. *Behavioral Neuroscience*, 119(6):1496–1510, December 2005. ISSN 1939-0084, 0735-7044. doi: 10.1037/0735-7044.119.6.1496. URL `http://doi.apa.org/getdoi.cfm?doi=10.1037/0735-7044.119.6.1496`.

Patricia S Goldman-Rakic. Cellular basis of working memory. *Neuron*, 14(3):477–485, 1995. Publisher: Elsevier.

Aaron J. Gruber, Gwendolyn G. Calhoon, Igor Shusterman, Geoffrey Schoenbaum, Matthew R. Roesch, and Patricio O'Donnell. More Is Less: A Disinhibited Prefrontal Cortex Impairs Cognitive Flexibility. *Journal of Neuroscience*, 30(50):17102–17110, December 2010. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.4623-10.2010. URL `https://www.jneurosci.org/content/30/50/17102`.

Shoai Hattori, Lillian Chen, Craig Weiss, and John F. Disterhoft. Robust hippocampal responsivity during retrieval of consolidated associative memory. *Hippocampus*, 25(5):655–669, May 2015. ISSN 1050-9631, 1098-1063. doi: 10.1002/hipo.22401. URL `https://onlinelibrary.wiley.com/doi/10.1002/hipo.22401`.

James M Hyman, Liya Ma, Emili Balaguer-Ballester, Daniel Durstewitz, and Jeremy K Seamans. Contextual encoding by ensembles of medial prefrontal cortex neurons. *Proceedings of the National Academy of Sciences*, 109(13):5086–5091, 2012. Publisher: National Acad Sciences.

Shantanu P Jadhav, Gideon Rothschild, Demetris K Roumis, and Loren M Frank. Coordinated excitation and inhibition of prefrontal ensembles during awake hippocampal sharp-wave ripple events. *Neuron*, 90(1):113–127, 2016. Publisher: Elsevier.

Karola Kaefer, Michele Nardin, Karel Blahna, and Jozsef Csicsvari. Replay of behavioral sequences in the medial prefrontal cortex during rule switching. *Neuron*, 106(1):154–165, 2020. Publisher: Elsevier.

Raymond P Kesner, Michael R Hunsaker, and Paul E Gilbert. The role of CA1 in the acquisition of an object-trace-odor paired associate task. *Behavioral neuroscience*, 119(3):781, 2005. Publisher: American Psychological Association.

J-R. King and S. Dehaene. Characterizing the dynamics of mental representations: the temporal generalization method. *Trends in Cognitive Sciences*, 18(4):203–210, April 2014. ISSN 1364-6613. doi: 10.1016/j.tics.2014.01.002. URL `https://www.sciencedirect.com/science/article/pii/S1364661314000199`.

Takashi Kitamura, Christopher J Macdonald, and Susumu Tonegawa. Entorhinal–hippocampal neuronal circuits bridge temporally discontiguous events. *Learning & Memory*, 22(9):438–443, 2015. Publisher: Cold Spring Harbor Lab.

Jan L Klee, Bryan C Souza, and Francesco P Battaglia. Learning differentially shapes prefrontal and hippocampal activity during classical conditioning. *eLife*, 10:e65456, October 2021. ISSN 2050-084X. doi: 10.7554/eLife.65456. URL https://elifesciences.org/articles/65456.

Antonio H Lara and Jonathan D Wallis. The role of prefrontal cortex in working memory: a mini review. *Frontiers in systems neuroscience*, 9:173, 2015. Publisher: Frontiers.

Stefan Leutgeb, Jill K Leutgeb, Alessandro Treves, May-Britt Moser, and Edvard I Moser. Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science*, 305(5688):1295–1298, 2004. Publisher: American Association for the Advancement of Science.

Scott Linderman, Benjamin Antin, David Zoltowski, and Joshua Glaser. SSM: Bayesian Learning and Inference for State Space Models. October 2020. URL https://github.com/lindermanlab/ssm.

Christopher J MacDonald, Kyle Q Lepage, Uri T Eden, and Howard Eichenbaum. Hippocampal "time cells" bridge the gap in memory for discontiguous events. *Neuron*, 71(4):737–749, 2011. Publisher: Elsevier.

Christopher J MacDonald, Stephen Carrow, Ryan Place, and Howard Eichenbaum. Distinct hippocampal time cell sequences represent odor memories in immobilized rats. *Journal of Neuroscience*, 33(36):14607–14616, 2013. Publisher: Soc Neuroscience.

Emily L Mackevicius, Andrew H Bahle, Alex H Williams, Shijie Gu, Natalia I Denisenko, Mark S Goldman, and Michale S Fee. Unsupervised discovery of temporal sequences in high-dimensional datasets, with applications to neuroscience. *eLife*, 8:e38471, February 2019. ISSN 2050-084X. doi: 10.7554/eLife.38471. URL https://doi.org/10.7554/eLife.38471.

L. Mazzucato, G. La Camera, and A. Fontanini. Expectation-induced modulation of metastable activity underlies faster coding of sensory stimuli. *Nature Neuroscience*, 22(5):787–796, May 2019. ISSN 1546-1726. doi: 10.1038/s41593-019-0364-9. URL https://www.nature.com/articles/s41593-019-0364-9.

Luca Mazzucato, Alfredo Fontanini, and Giancarlo La Camera. Dynamics of Multistable States during Ongoing and Evoked Cortical Activity. *Journal of Neuroscience*, 35(21):8214–8231, May 2015. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.4819-14.2015. URL https://www.jneurosci.org/content/35/21/8214.

Matthew D. McEchron and John F. Disterhoft. Sequence of Single Neuron Changes in CA1 Hippocampus of Rabbits During Acquisition of Trace Eyeblink Conditioned Responses. *Journal of Neurophysiology*, 78 (2):1030–1044, August 1997. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.1997.78.2.1030. URL https://www.physiology.org/doi/10.1152/jn.1997.78.2.1030.

Matthew D. McEchron and John F. Disterhoft. Hippocampal encoding of non-spatial trace conditioning. *Hippocampus*, 9(4):385–396, 1999. ISSN 1050-9631, 1098-1063. doi: 10.1002/(SICI)1098-1063(1999)9:4<385::AID-HIPO5>3.0.CO;2-K. URL https://onlinelibrary.wiley.com/doi/10.1002/(SICI)1098-1063(1999)9:4<385::AID-HIPO5>3.0.CO;2-K.

Earl K. Miller. The prefontral cortex and cognitive control. *Nature Reviews Neuroscience*, 1(1):59–65, October 2000. ISSN 1471-0048. doi: 10.1038/35036228. URL https://www.nature.com/articles/35036228.

Earl K. Miller and Jonathan D. Cohen. An Integrative Theory of Prefrontal Cortex Function. *Annual Review of Neuroscience*, 24(1):167–202, March 2001. ISSN 0147-006X, 1545-4126. doi: 10.1146/annurev.neuro.24.1.167. URL https://www.annualreviews.org/doi/10.1146/annurev.neuro.24.1.167.

Edvard I. Moser, Emilio Kropff, and May-Britt Moser. Place Cells, Grid Cells, and the Brain's Spatial Representation System. *Annual Review of Neuroscience*, 31(1):69–89, July 2008. ISSN 0147-006X, 1545-4126. doi: 10.1146/annurev.neuro.31.061307.090723. URL https://www.annualreviews.org/doi/10.1146/annurev.neuro.31.061307.090723.

May-Britt Moser, David C. Rowland, and Edvard I. Moser. Place Cells, Grid Cells, and Memory. *Cold Spring Harbor Perspectives in Biology*, 7(2):a021808, February 2015. ISSN , 1943-0264. doi: 10.1101/cshperspect.a021808. URL http://cshperspectives.cshlp.org/content/7/2/a021808.

Masayoshi Murakami, Hanan Shteingart, Yonatan Loewenstein, and Zachary F. Mainen. Distinct Sources of Deterministic and Stochastic Components of Action Timing Decisions in Rodent Frontal Cortex. *Neuron*, 94(4):908–919.e7, May 2017. ISSN 0896-6273. doi: 10.1016/j.neuron.2017.04.040. URL https://www.sciencedirect.com/science/article/pii/S0896627317303975.

Yuji Naya and Wendy A. Suzuki. Integrating What and When Across the Primate Medial Temporal Lobe. *Science*, 333(6043):773–776, 2011. doi: 10.1126/science.1206773. URL https://www.science.org/doi/abs/10.1126/science.1206773. _eprint: https://www.science.org/doi/pdf/10.1126/science.1206773.

John O'Keefe and Michael L Recce. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3(3):317–330, 1993. Publisher: Wiley Online Library.

James M Otis, Vijay MK Namboodiri, Ana M Matan, Elisa S Voets, Emily P Mohorn, Oksana Kosyk, Jenna A McHenry, J Elliott Robinson, Shanna L Resendez, Mark A Rossi, and others. Prefrontal cortex output circuits guide reward seeking through divergent cue encoding. *Nature*, 543(7643):103–107, 2017. Publisher: Nature Publishing Group.

Eva Pastalkova, Vladimir Itskov, Asohan Amarasingham, and Gyorgy Buzsaki. Internally Generated Cell Assembly Sequences in the Rat Hippocampus. *Science*, 321(5894):1322–1327, September 2008a. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1159775. URL https://www.science.org/doi/10.1126/science.1159775.

Eva Pastalkova, Vladimir Itskov, Asohan Amarasingham, and Gyorgy Buzsaki. Internally generated cell assembly sequences in the rat hippocampus. *Science*, 321(5894):1322–1327, 2008b. Publisher: American Association for the Advancement of Science.

Adrien Peyrache, Mehdi Khamassi, Karim Benchenane, Sidney I Wiener, and Francesco P Battaglia. Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nature neuroscience*, 12(7): 919–926, 2009. Publisher: Nature Publishing Group.

Bradley R Postle. Working memory as an emergent property of the mind and brain. *Neuroscience*, 139(1): 23–38, 2006. Publisher: Elsevier.

Wayne E Pratt and Sheri J. Y Mizumori. Neurons in rat medial prefrontal cortex show anticipatory rate changes to predictable differential rewards in a spatial memory task. *Behavioural Brain Research*, 123(2):165–183, September 2001. ISSN 0166-4328. doi: 10.1016/S0166-4328(01)00204-2. URL `https://www.sciencedirect.com/science/article/pii/S0166432801002042`.

Alison R Preston and Howard Eichenbaum. Interplay of hippocampus and prefrontal cortex in memory. *Current Biology*, 23(17):R764–R773, 2013. Publisher: Elsevier.

Rene Quilodran, Marie Rothe, and Emmanuel Procyk. Behavioral Shifts and Action Valuation in the Anterior Cingulate Cortex. *Neuron*, 57(2):314–325, January 2008. ISSN 0896-6273. doi: 10.1016/j.neuron.2007.11.031. URL `https://www.sciencedirect.com/science/article/pii/S0896627307010215`.

Jonathan D. Raybuck and K. Matthew Lattal. Bridging the interval: Theory and neurobiology of trace conditioning. *Behavioural Processes*, 101:103–111, January 2014. ISSN 0376-6357. doi: 10.1016/j.beproc.2013.08.016. URL `https://www.sciencedirect.com/science/article/pii/S0376635713001915`.

Stefano Recanatesi, Ulises Pereira-Obilinovic, Masayoshi Murakami, Zachary Mainen, and Luca Mazzucato. Metastable attractors explain the variable timing of stable behavioral action sequences. *Neuron*, 110(1):139–153.e9, January 2022. ISSN 0896-6273. doi: 10.1016/j.neuron.2021.10.011. URL `https://www.sciencedirect.com/science/article/pii/S0896627321007790`.

Azza Sellami, Alice Shaam Al Abed, Laurent Brayda-Bruno, Nicole Etchamendy, Stephane Valerio, Marie Oule, Laura Pantaleon, Valerie Lamothe, Mylene Potier, Katy Bernard, and others. Temporal binding function of dorsal CA1 is critical for declarative memory formation. *Proceedings of the National Academy of Sciences*, 114(38):10262–10267, 2017. Publisher: National Acad Sciences.

Justin D Shin and Shantanu P Jadhav. Multiple modes of hippocampal–prefrontal interactions in memory-guided behavior. *Current opinion in neurobiology*, 40:161–169, 2016. Publisher: Elsevier.

Justin D Shin, Wenbo Tang, and Shantanu P Jadhav. Dynamics of awake hippocampal-prefrontal replay for spatial learning and memory-guided decision making. *Neuron*, 104(6):1110–1125, 2019. Publisher: Elsevier.

Paul R Solomon, Ellen R Vander Schaaf, Richard F Thompson, and Donald J Weisz. Hippocampus and trace conditioning of the rabbit's classically conditioned nictitating membrane response. *Behavioral neuroscience*, 100(5):729, 1986. Publisher: American Psychological Association.

Clara Kwon Starkweather, Samuel J Gershman, and Naoshige Uchida. The medial prefrontal cortex shapes dopamine reward prediction errors under state uncertainty. *Neuron*, 98(3):616–629, 2018. Publisher: Elsevier.

Kazumasa Z Tanaka, Aleksandr Pevzner, Anahita B Hamidi, Yuki Nakazawa, Jalina Graham, and Brian J Wiltgen. Cortical representations are reinstated by the hippocampus during memory retrieval. *Neuron*, 84 (2):347–354, 2014. Publisher: Elsevier.

Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 143–146, Vancouver BC Canada, May 2011. ACM. ISBN 9781450302289. doi: 10.1145/1978942.1978963. URL `https://dl.acm.org/doi/10.1145/1978942.1978963`.

Byron M Yu, John P Cunningham, Gopal Santhanam, Stephen Ryu, Krishna V Shenoy, and Maneesh Sahani. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. In *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc., 2008. URL `https://proceedings.neurips.cc/paper/2008/hash/ad972f10e0800b49d76fed33a21f6698-Abstract.html`.