nature neuroscience

Evaluation of ambiguous associations in the amygdala by learning the structure of the environment

Tamas J Madarasz[1,2], Lorenzo Diaz-Mataix[1], Omar Akhand[1], Edgar A Ycu[2], Joseph E LeDoux[1,3] & Joshua P Johansen[2,4]

Recognizing predictive relationships is critical for survival, but an understanding of the underlying neural mechanisms remains elusive. In particular, it is unclear how the brain distinguishes predictive relationships from spurious ones when evidence about a relationship is ambiguous, or how it computes predictions given such uncertainty. To better understand this process, we introduced ambiguity into an associative learning task by presenting aversive outcomes both in the presence and in the absence of a predictive cue. Electrophysiological and optogenetic approaches revealed that amygdala neurons directly regulated and tracked the effects of ambiguity on learning. Contrary to established accounts of associative learning, however, interference from competing associations was not required to assess an ambiguous cue–outcome contingency. Instead, animals' behavior was explained by a normative account that evaluates different models of the environment's statistical structure. These findings suggest an alternative view of amygdala circuits in resolving ambiguity during aversive learning.

To enhance their chance of survival, animals learn to make predictions based on sensory cues in their environment. However, it is not clear how they identify stimuli that are relevant for specific predictions and how they distinguish coincidences between environmental events from actual predictive relationships. If an outcome occurs both in the presence and the absence of a cue, for example, a contingent and therefore predictive relationship between the two is no longer obvious. An understanding of how the brain assesses such ambiguity in cue–outcome relationships is missing, and most accounts of animal learning confound ambiguity in the environment's statistical structure (that is, which relationships are predictive or causal in the environment) and uncertainty about the strength of established associations (for example, the probability with which an outcome follows a predictive cue).

We investigated how animals assess ambiguous predictive relationships using classical threat conditioning. In this model animals come to display defensive responses to stimuli predicting dangerous or aversive events after pairings of an initially neutral conditioned stimulus (CS), such as a tone, and a biologically salient unconditioned stimulus (US), such as a mild footshock[1–4]. Humans and non-human animals alike show graded contingency learning, depending on how well a given outcome is predicted by a sensory cue. In particular, rodents are known to exhibit reduced conditioning to a tone–CS if footshocks are presented both in the presence and absence of the tone, a phenomenon known as 'contingency degradation'[5].

The prevailing interpretation explains contingency degradation in terms of cue competition[6–8], where multiple cues compete for the ability to predict an outcome by partitioning a limited associative strength. For example, it is thought that during contingency degradation a strong association formed between the conditioning context and the shock reduces subsequent learning of the tone–shock association. This process is referred to as contextual blocking[5,9] and is thought to be implemented in the brain through attenuation of US processing during tone–shock pairings when the US is already predicted by the context[3,10]. Alternatively, a strong contextual association could be competing with the tone–CS at the time of memory expression[8]. Either type of cue competition, however, would rely on contextual learning, a hippocampus-dependent process.

Cue competition can be problematic under some circumstances because it assesses the ambiguity of predictive relationships only indirectly: instead of checking for dependencies between variables and learning statistical structure by evaluating different models of the environment, it sidesteps model selection and learns associations between any contiguous cue–outcome pair in a competitive manner.

Suggesting a different view, a previous *in vitro* study[11] found that the cellular-level process thought to underlie aversive memory storage in the lateral amygdala (LA) is itself sensitive to stimulus contingencies. Thus the brain might possess neural mechanisms at the level of the amygdala to evaluate contingencies between environmental stimuli, without relying on cue competition.

However, to make predictions in a statistically principled way from a small number of observations, the learning mechanism also needs to take into account the overall pattern of events in the environment and account for possibly complex interactions between the different cue–outcome associations. While there is strong evidence that sensory cues become associated with aversive (or rewarding) outcomes

[1]Center for Neural Science, New York University, New York, New York, USA. [2]RIKEN Brain Science Institute, Wako, Saitama, Japan. [3]The Emotional Brain Institute, Nathan Kline Institute for Psychiatric Research, Orangeburg, New York, USA. [4]Department of Life Sciences, Graduate School of Arts and Sciences, University of Tokyo, Tokyo, Japan. Correspondence should be addressed to J.P.J. (jjohans@brain.riken.jp) or T.J.M. (tamas.madarasz@unige.ch).

Received 24 February; accepted 26 April; published online 23 May 2016; doi:10.1038/nn.4308

through strengthening of sensory input synapses in the LA during associative learning[1–4], a principled learning strategy must go beyond evaluating single cue–outcome contingencies in isolation.

An understanding of both the learning strategy animals use on the computational level and of the neural circuitry involved is thus critical for identifying the circuit mechanisms and algorithmic level processes that could implement contingency evaluations in the face of ambiguity. Here we used a combination of behavioral and computational approaches together with optogenetics, electrophysiology and pharmacology to address these questions. We found that cue competition is not necessary for contingency degradation and does not give a satisfactory explanation of this phenomenon. Instead, animals' behavior is best explained by models that evaluate the overall statistical structure of the environment. Furthermore, we demonstrate that the amygdala tracks contingency changes in the environment, and we reveal that it is important in resolving ambiguity during learning.

RESULTS

Contingency evaluation independently of cue competition

We first determined whether predictions of the cue-competition models were supported when ambiguity in the ability of a given CS to predict the US was high. To test this, we first examined trial order sensitivity and the relationship between context and CS memory strength by varying the order of CS–US pairings and unsignaled USs (UUSs). Animals were given either three massed tone–shock pairings before, or three spaced tone–shock pairings intermixed with, 12 unsignaled shocks (both with 20% contingency) and were tested for aversive memories by measuring contextual and tone-evoked freezing 24 h later (**Fig. 1a**,**b**). Control I and II animals were given three CS–US pairings only (100% contingency). The control I group received three CS–US pairings spaced identically to those in the intermixed protocol but with all UUSs omitted. The control II group received massed CS–US pairings spaced identically to those in the pairings-first group, with the subsequent UUSs omitted, and conditioning terminated after the third CS–US pairing (**Fig. 1b**). Animals showed similar levels of tone-evoked freezing in both reduced-contingency conditions, and these freezing levels were significantly lower than for control animals (**Fig. 1c** and **Supplementary Fig. 1**). Animals were therefore sensitive to the ambiguity of the CS–US relationship and demonstrated the ability to integrate contingency information irrespective of the temporal order of training trials, contradicting a traditional cue-competition-based 'contextual blocking' account of contingency degradation.

Cue competition could also account for contingency degradation beyond such a forward blocking account. Some learning models suggest competition between associations at the time of memory retrieval[10] or trial-order-independent cue completion based on statistical learning principles, such as when learning strength parameters for predictive cues or causes in a predetermined generative model of the US[12]. However, we observed a reduction in CS memory strength between the pairings-first and control II groups without a corresponding change in context memory strength (**Fig. 1d**). This was also true upon time-binned analysis and when using a more salient conditioning context (**Supplementary Fig. 2a**,**b**). This suggests that competition, where a strong contextual association would suppress tone-evoked responding at the time of memory retrieval, also fails to account for contingency learning. Thus while under some circumstances there can be an apparent inverse relationship between the different cue–outcome associations (notably in the case of the spaced condition, where the low rate of shock delivery in the control I group results in low context freezing), this is not generally the case, and in particular is not necessary for the animals to learn a degraded tone–shock contingency. Looking at individual animals, we also observed that the correlation between tone and context freezing was positive in all four conditions (**Supplementary Fig. 1**).

To better understand the influence of contextual associations on learning the tone–shock contingency and to directly test for cue competition during learning and/or retrieval, we next infused the NMDA-receptor antagonist 2-amino-5-phosphonovalerate (APV) into the dorsal hippocampus before conditioning (**Fig. 2a**), a manipulation known to block the formation of contextual memories[13]. Consistent with previous results using a different procedure[14], this intervention had no effect on contingency degradation, despite significantly impairing contextual learning both in the spaced and the massed conditions (**Fig. 2b–e**). This provided further evidence that contingency degradation of auditory threat memories does not depend on competition between auditory and contextual cues, whether information

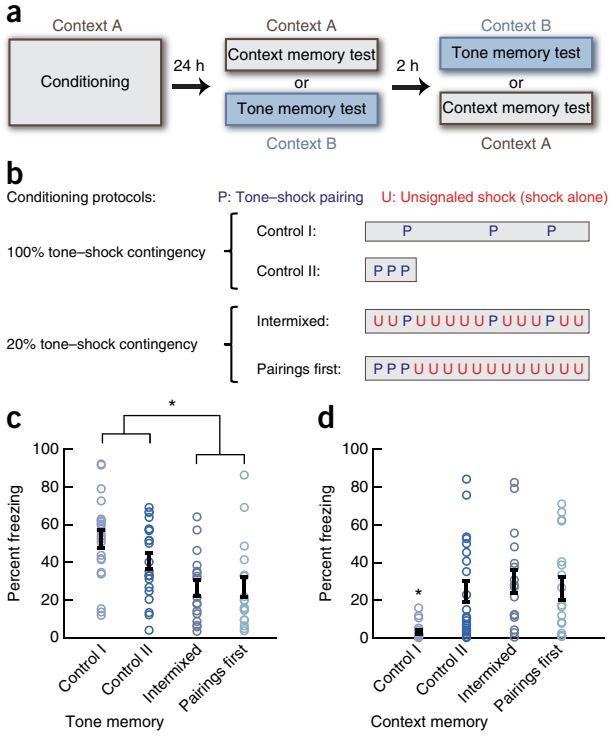

Figure 1 Reduced CS–US contingency results in reduced CS memory, irrespective of trial order and with or without changes in context memory. (**a**) Experimental design. Animals underwent threat conditioning on day 1 and were tested for contextual and tone (or tone and contextual) memories 24 h later. (**b**) Conditioning protocols. Rats received sequences of tone–shock pairings (P) and unsignaled shocks (U), or tone–shock pairings only. The boxes indicate time spent in the conditioning chamber: ~7 min for control II group and ~31 min for all other groups. (**c**) A 20% CS–US contingency during conditioning leads to significantly lower CS–induced freezing than a 100% contingency, whether unpaired shocks are given intermixed with or after tone–shock pairings ($n = 22, 22, 17, 18$, two-way ANOVA, no significant interaction $F_{1,75} = 1.63$, $P = 0.21$, main effect for contingency $F_{1,75} = 18.02$, $^*P = 0.00006$, simple effects for contingency $F_{1,75} = 15.0$, $P = 0.0002$; $F_{1,75} = 4.48$, $P = 0.038$ for control I vs. intermixed and control II vs. pairings first, respectively). (**d**) Context memory strengths for the same animals. Reduction of CS memory with degraded CS–US contingency is not explained by changes in context memory strengths, as there was no difference between context memories of control II and pairings-first groups (two-way ANOVA, significant interaction $F_{1,75} = 6.44$, $P = 0.013$, simple effect for contingency $F_{1,75} = 0.00008$, $P = 0.81$, not significant; $F_{1,75} = 10.65$, $^*P = 0.001$, for control II vs. pairings first and control I vs. intermixed, respectively). Error bars indicate s.e.m.

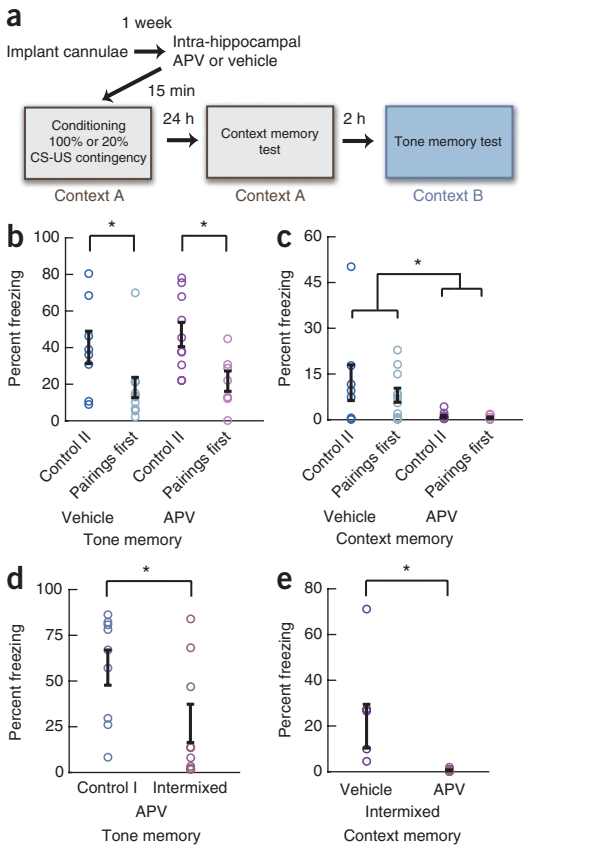

Figure 2 Conditioning to the context is not required for contingency degradation. (**a**) Experimental design depicting pharmacological inactivation of NMDA receptors in dorsal hippocampus before conditioning. (**b**) Hippocampal APV injections had no effect on learning the reduced auditory CS–US contingency ($n = 8, 11, 10, 7$, two-way ANOVA, no significant interaction $F_{1,32} = 0.07$, $P = 0.79$, main effect for contingency $F_{1,32} = 11.98$, $P = 0.0015$, simple effect for contingency $F_{1,32} = 8.60$, $^*P = 0.011$; $F_{1,32} = 5.45$, $^*P = 0.026$ for vehicle and APV groups, respectively). (**c**) NMDA receptor blockade impairs the acquisition of contextual aversive memories (two-way ANOVA, no significant interaction, $F_{1,32} = 0.38$, $P = 0.54$, main effect for drug treatment, $F_{1,32} = 9.47$, $^*P = 0.0043$). (**d**) Similarly to the pairings-first case, contingency degradation to the auditory stimulus is unaffected in the intermixed condition by APV infusion in dorsal hippocampus ($n = 9, 9$, unpaired sample t-test, $t_{16} = 2.14$, $^*P = 0.048$). (**e**) Impaired contextual aversive memory formation after NMDA receptor blockade in the intermixed condition ($n = 7, 9$, unpaired sample t-test, $t_{14} = 2.31$, $^*P = 0.037$). Error bars indicate s.e.m.

about the reduced tone–shock contingency is delivered after tone–shock pairings or the different types of shocks are intermixed (**Fig. 2d**,**e**). We further validated these results by comparing an alternative measure of threat response (defecation) in the spaced condition and found that it paralleled our results measuring freezing (**Supplementary Fig. 3**).

We also verified that the observed decrements in the tone-evoked responding were not due simply to delivering a larger number of shock USs (so-called 'reinforcer devaluation'). Groups of animals that received 15 or 21 tone–shock pairings (**Supplementary Fig. 4a**) displayed similarly high levels of tone freezing, indicating that learning the tone–shock association was at a stable asymptote and that the larger number of footshocks did not lead to a devaluation of this US. Further, if, instead of delivering UUSs, we signaled shocks following the three tone–shock pairings by a second discrete CS (a flashing light), contingency degradation did not occur (**Supplementary Fig. 4b**), consistent with the so-called cover-stimulus effect[15,16]. Thus, delivering a larger number of USs did not in itself cause contingency degradation; instead, the animals' learning reflected the precise environmental contingencies during learning.

LA neural activity controls and tracks contingency

As animals could learn a reduced tone–shock contingency without relying on hippocampal plasticity and contextual memory formation, we next explored the role of the amygdala in contingency degradation. Previous research suggests that the amygdala is important for contingency evaluations during reward learning[17,18]. It is also well established that synaptic enhancement of auditory inputs to LA pyramidal neurons occurs during, and is necessary for, auditory aversive learning, and that this enhancement is dependent on US-evoked activation of LA neurons coincident with the auditory CS[1–4,19]. A direct representation of the CS–US contingency needs to integrate information about the number of CS–US pairings versus UUSs, so the activation of LA pyramidal cells by the UUSs could be an important trigger for learning contingency degradation. To test whether this is the case, we expressed the outward proton-pump Arch-T[20] in these neurons, using intra-LA injection of a lentiviral vector (**Fig. 3a**,**b**). In previous work[21] we demonstrated pyramidal-cell-specific targeting of Arch-T expression using this viral targeting approach and laser-induced inhibition of shock-evoked responses in these cells, which we also validated here (**Fig. 3a**). We used this technique to test whether activity in LA pyramidal neurons during UUSs is necessary for the degraded contingency effects to occur. We found that inactivating these cells during UUSs, but not at other times in the conditioning session, rescued freezing to the tone on the long-term memory test (**Fig. 3b**,**c** and **Supplementary Fig. 5**), but caused no significant change in context memory (**Fig. 3d**). The US-evoked depolarization of LA pyramidal neurons can thus differentially modulate the strength of auditory aversive memories depending on its timing relative to the CS, and this can occur independently of changes in contextual memory strength.

As discussed above, the enhancement of auditory input synapses in the LA underlies the expression of auditory aversive memories. At the level of the LA, this representation corresponds to the association between the sensory features of the auditory stimulus and aversive outcome and is not correlated with the motor output directly[22]. Additionally, previous work in humans[23] and primates[18,24] has indicated that amygdala neurons can adapt their activity according to the higher order structure of the task environment. A reduction in the overall enhancement of auditory processing in the LA could therefore regulate behavioral responses during retrieval in the case of contingency degradation. To test whether this is the case, we next examined whether UUSs given after CS–US pairings reduced the learning-induced enhancement of the auditory-evoked local field potential (A-LFP) response, a measure of synaptic enhancement in the threat learning circuit. We recorded A-LFPs in the LA before and after three tone–shock pairings (control II protocol) or three tone–shock pairings followed by unpaired shocks (pairings-first protocol) (**Fig. 4a**). Consistent with previous findings, A-LFP was enhanced 24 h after conditioning with 100% tone–shock contingency, however this enhancement was significantly reduced in animals that were trained with a reduced contingency (**Fig. 4b**,**c** and **Supplementary Figs. 6** and **7**), paralleling a reduction in freezing behavior in the same animals (**Fig. 4d**). Thus, consistent with the behavioral results, the learning-induced changes in auditory processing in the LA reflect the broader environmental contingencies.

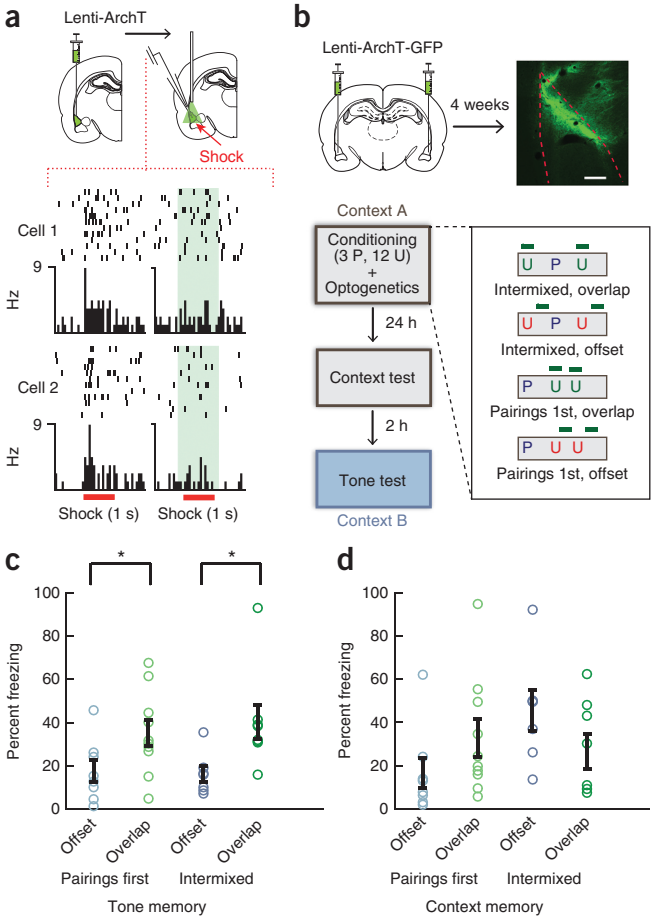

Figure 3 Activation of LA pyramidal cells during unsignaled USs is required to learn the degraded CS–US contingency. (**a**) Optogenetic inhibition of shock-evoked firing rate responses in single LA neurons. Peri-event time rasters (top) and histograms (bottom) show shock-evoked (red bars) responses in two example neurons without (left) and with optical inhibition (right, laser-on denoted by green bar). (**b**) Design of optogenetic behavioral experiments. Top, lentivirus injection and example of Arch-T expression in LA pyramidal neurons (scale bar, 160 µm). Virus expression in LA was verified for all experimental animals included in study. Bottom, protocols with laser illumination either coinciding with or offset from UUSs. Three tone–shock US pairings (P) were presented either intermixed with or before 12 unsignaled USs (U). (**c**) Inactivation of LA pyramidal neurons during, but not offset from, UUSs prevented the learning of the degraded auditory CS–US contingency ($n = 7, 8, 8, 10$, two-way ANOVA, no significant interaction, $F_{1,29} = 0.28$, $P = 0.60$, main effect for inactivation, $F_{1,29} = 11.41$, $P = 0.0021$, simple effects for inactivation $F_{1,29} = 7.02$, $^*P = 0.013$, $F_{1,29} = 4.45$; $^*P = 0.044$ for pairings-first and intermixed groups respectively). (**d**) Context memory strength was unaffected by optogenetic manipulation (two-way ANOVA, significant interaction, $F_{1,29} = 4.39$, $P = 0.045$, main effect for inactivation $F_{1,29} = 0.0297$, $P = 0.86$, not significant; simple effects for inactivation $F_{1,29} = 2.37$, $P = 0.13$; $F_{1,29} = 2.03$, $P = 0.17$ for pairings-first and intermixed groups respectively; not significant). Error bars indicate s.e.m.

Assessing ambiguity with structure learning

Amygdala processing thus plays a key role in the learning and retrieval of an ambiguous CS–US relationship, and this learning does not rely on cue competition, although it might incorporate more complex context-cue interactions. However, in the absence of cue competition, it is not clear what computational strategy animals use to resolve ambiguity during associative learning. Addressing this question requires the establishment of a computational framework that can quantitatively account for our behavioral findings, as well as predict the effects of the neural manipulations we performed and account for known conditioning phenomena that arise as a result of ambiguity in the predictive relationship between cues and outcomes.

We propose a structure learning model (SLM) that directly assesses uncertainty in the environment's statistical structure, determining which relationships are actually predictive by considering statistical dependencies (as well as temporal order and contiguity) between variables. Given events during conditioning, SLM learns a posterior probability distribution over the possible sets of predictive relationships in the environment, represented by different graph structures (**Fig. 5a**) using the formalism of Bayesian networks[25]. During retrieval, the strength of an association can be evaluated by calculating the posterior probability of a connection (a direct edge or a path in the graph) between the corresponding cue and outcome using a model-averaging procedure (Online Methods). Unlike simple cue competition, structure learning compares different configurations of interactions between variables and weighs these representations against each other. Such a model is able to learn a cue–outcome contingency even in the absence of a competing cue while incorporating flexibility in the range of possible interactions between cues.

We examined whether this type of model could simultaneously explain responses to discrete cues and the conditioning context. We built on previous work characterizing human causal judgments using a structure learning approach[26], extending it to the threat conditioning framework and to modeling neural interventions and more complex environments. We also analyzed the importance of the different components of the model in fitting a wide range of behavioral data.

To enable model fitting and comparison, we collected further behavioral data in a manner similar to experiment 1 (**Fig. 1a**), but using varied numbers of UUSs and CS–US pairings, allowing us to test which models can simultaneously explain learning under different conditions of ambiguity. In particular, if USs arrive only in the presence of the CS (that is, only CS–US pairings are given), the association between context and US is itself ambiguous, as it is not clear whether predictive power should be attributed to just the context, just the CS, or both[27]. We further included behavioral results for different degrees of contingency degradation by varying the number of UUSs after CS–US pairings.

We found that SLM successfully accounted for standard learning curves of context and tone memory strength and predicted how associative strength is attributed under ambiguity, including the effects of contingency degradation, the effects of partially reinforcing (or extinguishing) the context and the U-shaped learning curve of the context memory strength during overshadowing by the CS (**Fig. 5b**). SLM was also able to account for freezing levels in the control I group (resulting from a low rate of shock delivery) and successfully explained our first experiment (**Fig. 5b** and **Supplementary Fig. 8**). In addition, SLM successfully predicted the effects of hippocampal NMDA receptor blockade, using the best-fit parameters from the behavioral data set, and predicted the result of the amygdala inactivation experiments (**Fig. 5b**). In summary, SLM was able to capture how the different associations interact in driving behavior both in cases where these interactions appear competitive and in cases where there is an apparent dissociation or facilitation between associations.

A straightforward extension of SLM (**Supplementary Fig. 9**) that included a second CS (such as a light) but kept the best-fit parameters and scaling of the original model could also account for a range of previously documented conditioning phenomena involving the assessment of ambiguous stimuli. SLM could thus account for the effects

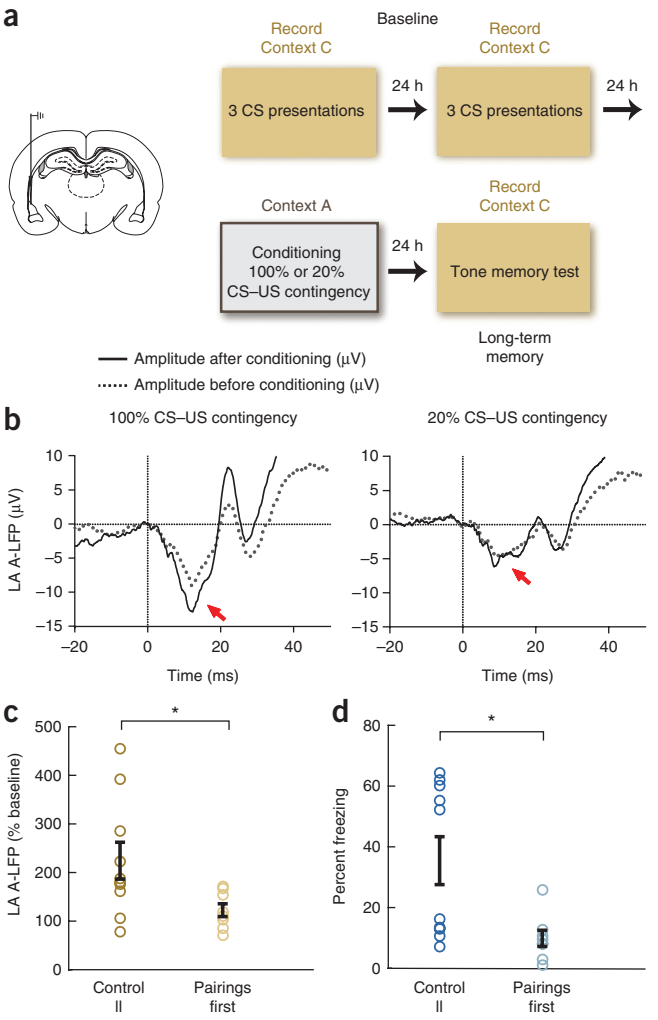

Figure 4 Degraded CS–US contingency leads to reduced enhancement of CS processing in the LA. (**a**) Design for the *in vivo* physiology experiment. (**b**) Population-averaged auditory-evoked local field potential (A-LFP) traces before and after conditioning in the control II group (left) and pairings-first group (right). Vertical line at $t = 0$ indicates CS onset; red arrow, peak depolarization. (**c**) Population-averaged post-training A-LFP as a percentage of pre-training baseline. Reduced CS–US contingency led to reduced potentiation of auditory CS processing ($n = 10, 8$, unpaired t-test, $t_{11,1} = 2.54$, $^*P = 0.028$). (**d**) Animals in the *in vivo* physiology experiment also showed reduced CS memory in the reduced contingency (pairings-first) condition (unpaired t-test, $t_{11,0} = 3.07$, $^*P = 0.011$). Error bars indicate s.e.m.

of signaling the UUSs with a second CS (the cover-stimulus effect described above) and gave a good fit both for our replication of this phenomenon (**Fig. 5b** and **Supplementary Fig. 10a**) and qualitatively similar predictions to the data from previous studies[28] using related experimental procedures (**Supplementary Fig. 10a**). The model's prediction for other phenomena (including blocking[10], overshadowing and recovery from overshadowing[29]) are further detailed in **Supplementary Figure 10b–d**.

Model comparison

We compared SLM to three models that assume a fixed structure and evaluate contingencies by learning strength parameters for associations through some form of cue competition (**Supplementary Fig. 11** and **Supplementary Modeling**). For the most direct comparison between a structure learning and a combined parameter learning/cue competition approach, we fit a parameter learning model[12], or PLM, that uses an identical Bayesian network representation but assumes the maximally connected structure (**Fig. 5a**, Graph 6). PLM learns a strength parameter for each edge starting from flexible, independent prior distributions over these edge parameters, fit to best explain behavioral data. Despite this flexibility, the parameter learning approach that implements cue competition in a statistical learning framework did not capture well how animals evaluated contingencies across the different conditions (**Table 1** and **Supplementary Fig. 11**).

Further, we included two advanced associative models that represent modern implementations of the cue competition idea formulated in the original Rescorla-Wagner model. These extend the Rescorla-Wagner model to allow for retrospective updating of associations and to capture the covariance information between cues and outcomes. Like the Rescorla-Wagner model, Van Hamme and Wasserman's extension[7] (**Supplementary Fig. 11**) implements cue competition during learning, but also updates associations when either the cue or the outcome (or both) are absent. Although this model utilizes the covariance information between a cue and an outcome, it evaluates these cue–outcome correlations in isolation for each cue, and as such did not give a good account of the behavior we observed (**Table 1**). A further shortcoming of this model is that it cannot account for the hippocampal interventions, since it does not predict contingency degradation in the absence of a competing variable (compare **Fig. 2b**,**c**). We therefore also evaluated a version of this model in which we added the background cue; however, this modification did not result in a better model fit (**Supplementary Table 1**). The sometimes-competing-retrieval model[8] (SOCR) considers the covariance information both between cues and outcomes and between different cues, in this sense approximating the principles of a Bayesian parameter learning model, and implements cue competition at the time of memory retrieval. We fit this model to the behavioral data both in its original form and with the added background variable, but it did not provide a fit comparable to that of SLM (**Table 1** and **Supplementary Table 1**).

SLM thus provided a better quantitative fit than PLM or associative cue-competition models, while also using fewer free parameters, and was robust to changes in specific components of the model (**Supplementary Table 2**), suggesting that it is the principle of evaluating different models of the environment that enables it to match observed behavior. A model implementing full Bayesian inference by learning both a distribution over structures and corresponding parameters (SPLM) provided a similar fit to SLM (**Supplementary Fig. 11**), but performed worse according to measures controlling for extra model parameters (**Table 1**) with the Bayesian information criterion, indicating that improvements from adding parameter learning did not justify adding even a single parameter to the model.

DISCUSSION

Here we examined the neural and computational processes through which ambiguity regulates aversive memory strength. First we identified key neural processes regulating contingency learning, revealing a new function of amygdala pyramidal neurons: in addition to their known role in storing associative aversive memories, they also actively participate in regulating a given association in response to signals (unsignaled aversive outcomes) that increase ambiguity in the cue–outcome association. Further, our results demonstrate that the degree of enhancement of auditory CS processing in amygdala neurons directly reflects a given CS–US contingency. Finally, we found

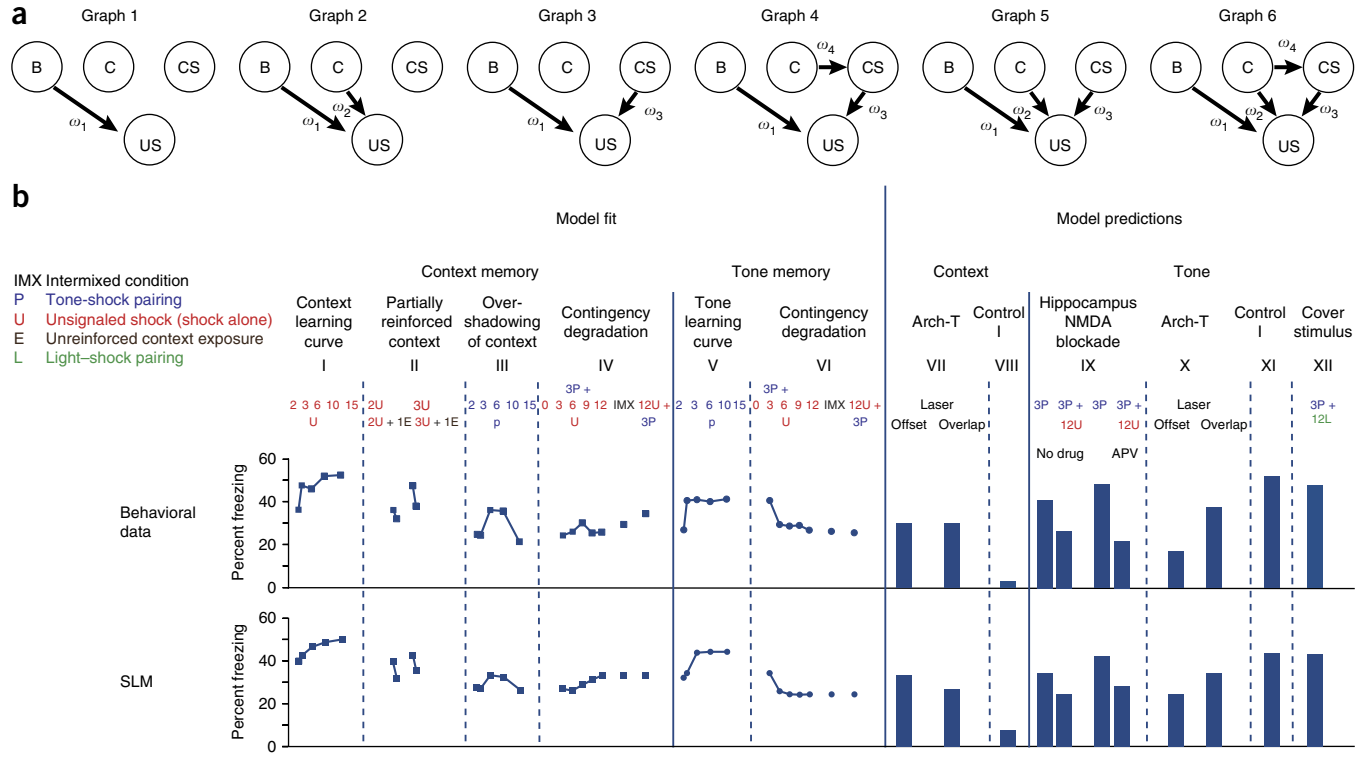

Figure 5 Comparison of the structure learning model (SLM) and behavioral data. (**a**) The six different graphical representations of statistical associations in the environment that are compared in SLM. B, background; C, context. (**b**) Context and tone CS memory strengths (behavioral data, top; SLM, bottom). Data points are percentage freezing, tested 24 h after the different conditioning protocols. Context memory: (I) 2, 3, 6, 10 or 15 unsignaled USs (UUSs). (II) 2 or 3 UUS followed by 0 or 1 unreinforced ITIs in the conditioning chamber. (III) 2, 3, 6, 10 or 15 CS–US pairings. (IV) 3 CS–US pairings followed by 0, 3, 6, 9, or 12 UUSs. Intermixed condition and pairings-last condition (3 CS–US pairings + 12 UUSs). Tone memory: (V) Same groups as in III. (VI) Same groups as in IV. (VII–XII) Bottom, memory strength as predicted by SLM. Model was fitted to data in I–VI. Best-fit parameters were then used to predict the effects of neural interventions and freezing scores for control I and cover stimulus groups (see Online Methods and **Supplementary Modeling**).

converging evidence on the computational and implementation levels against learning models that rely only on learning competing cue–outcome associations, supporting instead an account that directly assesses ambiguity in the environment's structure.

Structure and parameter learning as complementary strategies

When learning from sparse and ambiguous data, structure learning and model selection are important prerequisites for making successful predictions. Falsely assuming predictive relationships where they do not exist leads to a form of overfitting[30] and to poor generalization for future predictions. Quickly distinguishing spurious and predictive relationships is therefore important, and structure learning achieves this by also considering sparser structures that might lead to better predictions by identifying which variables actually interact.

While the exact contingencies between variables (for example, the strength of a generative causal process) often change over time, the existence or lack of a predictive relationship tends to be a stable property of an environment over time. This provides a strong rationale for separating the structure and parameter aspects of learning in certain domains and for engaging a structure learning mechanism when the brain is initially faced with a new environment or task.

Once enough information is gathered to evaluate different structures with a certain confidence, an important next step is to fine-tune the individual parameters of those models. We theorize that, as animals explore their environment, initial learning is geared toward structure learning, with a (potentially gradual) switch to parameter learning following, resulting in distributed representations of associations. Since continually updating a distribution over structures is computationally expensive and likely inadvisable, structure might be reengaged only if new environmental variables are encountered or the events in the environment strongly violate expectations based on the current model. Such a dual learning mechanism could in turn help explain the difficulty of persistently weakening aversive

Table 1 Comparison of model fits

Model	MSE	Standard error of the MSE	MSE for hippocampus APV injection	Free parameters	Bayesian information criterion
Structure learning	13.44	0.0054	34.03	2	88.82
Structure and parameter	12.61	0.054	18.29	9	110.54
Parameter learning	20.80	0.024	245.64	8	121.69
SOCR (extended comparator hypothesis)	32.45	–	292.53	6	127.85
Extended Rescorla-Wagner model	67.06	–	848.88	8	155.64

Mean squared error (MSE) for the best fit of each model, followed by the standard error of the MSE, and the error of the model in predicting the results of hippocampal interventions, with each value representing percentage freezing squared. The Bayesian information criterion provides a principled measure of model comparison, taking into account the number of model parameters.

memories and account for some important phenomena in memory updating and reconsolidation[31].

Brain structures such as the medial prefrontal cortex or the anterior cingulate have been implicated in updating or representing internal models of the environment[32,33]. These brain regions, together with the amygdala[24,34], as well as certain neuromodulators[35], could help determine which type of learning is employed, depending on the level of ambiguity in the environmental contingencies and on how rapidly or drastically these contingencies appear to change. A more exact understanding of the circumstances that engage the different learning strategies would be important in understanding how aversive memories are updated, with possible clinical applications in the treatment of persistent and/or exaggerated responses.

Context as both cue and modulator for internal models

Here we examined the conditioning context's role as a CS; however, the context, by modulating memory retrieval, is also known to have important effects on learning that go beyond forming predictive associations. While first-learned associations often easily transfer between (physical or temporal) contexts for retrieval[36], if learning takes place over multiple epochs or in multiple contexts, behavior can be sensitive to the retrieval context as well. Thus it is possible that in complex situations different distributions over structures are associated with different environments or a change of context determines whether structure or parameter learning is preferentially engaged, which could explain the context's role as a modulator of memory.

A different formulation of structure learning using latent causes, on which our work also builds[37], proceeds by clustering similar events in the environment, with subsequent work successfully modeling phenomena related to extinction and to renewal[38,39]. The context-specific nature of these phenomena in particular could be an example of how structure learning results in context-specific behaviors. Though these different formulations of structure learning rely on different computational processes and explain different learning phenomena, they all give support to the idea that the brain could employ structure learning to deal with certain types of uncertainty.

Neural implementation of structure learning

The experimental findings and SLM together suggest an algorithmic-level view on how structure learning and structured representation of the environment could emerge in associative learning by implying a circuit architecture in which this learning could be implemented. The LA is known to be an important integrative site through which sensory information from different modalities is associated with aversive (or rewarding) outcomes. Current views suggest that plasticity of modality-specific sensory input synapses to LA neurons mediates this form of aversive learning. However, cells in the LA and in thalamic and cortical structures that provide sensory input to the LA show a diversity of response properties, with some cells responding to several sensory cues rather than a single one[40–42]. This representation parallels the diversity of graph structures seen in our statistical model, with different combination of cues associated with the US in different graphs.

Several models have been proposed for how neurons might compute inference in graphical models[43,44] Some in particular have suggested that simple learning rules can produce synaptic weights and firing rates that represent how well patterns of sensory stimuli in the environment agree with an internal generative model[45]. A synaptic learning rule tracking the likelihood of a generative model represented by input synapses, together with an appropriately learned normalization to translate these likelihoods into a probability distribution across the structures, could then implement structure learning in SLM. In such an implementation the priors of the model correspond to initial distributions over synaptic weights and over the ratio of cells with different combinations of sensory input. Such a neural representation could provide a simple and efficient probabilistic code for structure learning[46]. Unlike traditional models of associative learning where a single weight and corresponding synaptic connection(s) control an association, here information about each association is represented by, and distributed over, multiple weights. An important characteristic of such a distributed representation is that computations can proceed in parallel over different microcircuits representing different models of the environment, but with all of them affecting each other at the time of behavioral readout. Updating multiple graphical structures representing the different features of a given learning environment at amygdala neuron synapses and/or at synapses in upstream areas could be accomplished through well-established heterosynaptic plasticity mechanisms[47,48] that allow synaptic weight changes even at synapses which are not directly recruited during plasticity induction.

Explicitly representing the many possible structures of a complex environment can be a challenge, even though calculating the posterior probability over specific features (such as edge probability) can be done efficiently even for a large number of variables under reasonable constraints[49] (such as are imposed by temporal relationships between cues and the complexity of models considered). However, a synaptic sampling mechanism where the inherent variability of synapses represents a distribution of synaptic strengths might provide a more efficient alternative to an exact enumeration of graph structures and, in particular, might implement the integration over many different parameter values through sampling over stochastic synaptic features and spine motility[50].

Our electrophysiology data demonstrate that averaged neural activity (as reflected by the local field potential) in LA can track contingencies over broad timescales and that activation of LA neurons is important in regulating contingency evaluations during learning. This supports the idea that LA neural activity reflects and can causally modulate inferential processes. While these data suggest that LA (or other) neuronal ensembles can encode sensory information as probabilistic graphical structures, an ideal test of this model would be to examine more closely whether neuronal ensembles in these circuits encode information in this way and how learning affects these representations. However, this requires the ability to chronically monitor large-scale neuronal population dynamics. Until recently this has not been possible, but recent advances in neuronal recording and imaging techniques[4] may allow researchers to examine when and how these types of representations are encoded and altered with learning. The SLM along with the experimental data described here provide a framework for guiding future research in this area. This approach could provide insights into how environmental stimuli are selected to become associated with biological threats and could be a key step in understanding anxiety disorders that are characterized by maladaptive and inappropriate responses to stimuli.

METHODS

Methods and any associated references are available in the online version of the paper.

Note: Any Supplementary Information and Source Data files are available in the online version of the paper.

ACKNOWLEDGMENTS

We thank J. Gardner, W.J. Ma and C. Yokoyama for advice and comments on the manuscript and N. Daw for discussions during the course of this work. We thank

the UNC vector core and E. Boyden (MIT) for the lentiviral vectors. This study was funded by National Institute of Mental Health (NIMH) grants R01-MH046516 and R01-MH38774, National Institute on Drug Abuse (NIDA) grant DA029053 (J.E.L.), MEXT Strategic Research Program for Brain Sciences (11041047) and Grants-in-Aid for Scientific Research (25710003, 25116531, 15H04264, 16H01291, 15H01301) (J.P.J.), and an NYU Williamson Fellowship (T.J.M.).

AUTHOR CONTRIBUTIONS

T.J.M. designed the experiments, collected and analyzed data, developed the computational models and wrote the manuscript. J.P.J. designed the experiments, collected data and wrote the manuscript. L.D.-M. collected and analyzed electrophysiology data. O.A. collected data. E.A.Y. collected and analyzed single-cell electrophysiology data. J.E.L. contributed to data interpretation and the final version of the manuscript.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

1. Maren, S. & Quirk, G.J. Neuronal signalling of fear memory. *Nat. Rev. Neurosci.* **5**, 844–852 (2004).
2. LeDoux, J.E. Emotion circuits in the brain. *Annu. Rev. Neurosci.* **23**, 155–184 (2000).
3. Herry, C. & Johansen, J.P. Encoding of fear learning and memory in distributed neuronal circuits. *Nat. Neurosci.* **17**, 1644–1654 (2014).
4. Gründemann, J. & Lüthi, A. Ensemble coding in amygdala circuits for associative learning. *Curr. Opin. Neurobiol.* **35**, 200–206 (2015).
5. Rescorla, R.A. Probability of shock in the presence and absence of CS in fear conditioning. *J. Comp. Physiol. Psychol.* **66**, 1–5 (1968).
6. Rescorla, R.A. & Wagner, A.R. A theory of pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. in *Classical Conditioning II Current Research and Theory* (eds. Black, A.H. & Prokasy, W.F.) 64–99 (Appleton-Century-Crofts, New York, 1972).
7. Van Hamme, L.J. & Wasserman, E.A. Cue competition in causality judgments: the role of nonpresentation of compound stimulus elements. *Learn. Motiv.* **25**, 127–151 (1994).
8. Stout, S.C. & Miller, R.R. Sometimes-competing retrieval (SOCR): a formalization of the comparator hypothesis. *Psychol. Rev.* **114**, 759–783 (2007).
9. Pearce, J.M. & Bouton, M.E. Theories of associative learning in animals. *Annu. Rev. Psychol.* **52**, 111–139 (2001).
10. Kamin, L.J. Attention-like processes in classical conditioning. in *Miami Symposium On The Production Of Behavior Aversive Stimulation* (ed. Jones, M.R.) 9–33 (University of Miami Press, 1968).
11. Bauer, E.P., LeDoux, J.E. & Nader, K. Fear conditioning and LTP in the lateral amygdala are sensitive to the same stimulus contingencies. *Nat. Neurosci.* **4**, 687–688 (2001).
12. Holyoak, K.J. & Cheng, P.W. Causal learning and inference as a rational process: the new synthesis. *Annu. Rev. Psychol.* **62**, 135–163 (2011).
13. Kim, J.J., DeCola, J.P., Landeira-Fernandez, J. & Fanselow, M.S. *N*-Methyl-D-aspartate receptor antagonist APV blocks acquisition but not expression of fear conditioning. *Behav. Neurosci.* **105**, 126–133 (1991).
14. McNish, K.A., Gewirtz, J.C. & Davis, M. Disruption of contextual freezing, but not contextual blocking of fear-potentiated startle, after lesions of the dorsal hippocampus. *Behav. Neurosci.* **114**, 64–76 (2000).
15. Durlach, P.J. Effect of signaling intertrial unconditioned stimuli in autoshaping. *J. Exp. Psychol. Anim. Behav. Process.* **9**, 374–389 (1983).
16. Gunther, L.M. & Miller, R.R. Prevention of the degraded-contingency effect by signalling training trials. *Q. J. Exp. Psychol. B* **53**, 97–119 (2000).
17. Balleine, B.W., Killcross, A.S. & Dickinson, A. The effect of lesions of the basolateral amygdala on instrumental conditioning. *J. Neurosci.* **23**, 666–675 (2003).
18. Bermudez, M.A. & Schultz, W. Responses of amygdala neurons to positive reward-predicting stimuli depend on background reward (contingency) rather than stimulus-reward pairing (contiguity). *J. Neurophysiol.* **103**, 1158–1170 (2010).
19. Rogan, M.T., Stäubli, U.V. & LeDoux, J.E. Fear conditioning induces associative long-term potentiation in the amygdala. *Nature* **390**, 604–607 (1997).
20. Chow, B.Y. *et al.* High-performance genetically targetable optical neural silencing by light-driven proton pumps. *Nature* **463**, 98–102 (2010).
21. Johansen, J.P. *et al.* Hebbian and neuromodulatory mechanisms interact to trigger associative memory formation. *Proc. Natl. Acad. Sci. USA* **111**, E5584–E5592 (2014).
22. Goosens, K.A., Hobin, J.A. & Maren, S. Auditory-evoked spike firing in the lateral amygdala and Pavlovian fear conditioning: mnemonic code or fear bias? *Neuron* **40**, 1013–1022 (2003).
23. Prévost, C., McNamee, D., Jessup, R.K., Bossaerts, P. & O'Doherty, J.P. Evidence for model-based computations in the human amygdala during Pavlovian conditioning. *PLoS Comput. Biol.* **9**, e1002918 (2013).
24. Saez, A., Rigotti, M., Ostojic, S., Fusi, S. & Salzman, C.D. Abstract context representations in primate amygdala and prefrontal cortex. *Neuron* **87**, 869–881 (2015).
25. Pearl, J. *Probabilistic Reasoning in Intelligent Systems* (Morgan Kauffmann, San Mateo, California, USA, 1988).
26. Griffiths, T.L. & Tenenbaum, J.B. Structure and strength in causal induction. *Cognit. Psychol.* **51**, 334–384 (2005).
27. Hall, G. & Symonds, M. Overshadowing and latent inhibition of context aversion conditioning in the rat. *Auton. Neurosci.* **129**, 42–49 (2006).
28. Witnauer, J.E. & Miller, R.R. Degraded contingency revisited: posttraining extinction of a cover stimulus attenuates a target cue's behavioral control. *J. Exp. Psychol. Anim. Behav. Process.* **33**, 440–450 (2007).
29. Matzel, L.D., Schachtman, T.R. & Miller, R.R. Recovery of an overshadowed association achieved by extinction of the overshadowing stimulus. *Learn. Motiv.* **16**, 398–412 (1985).
30. Koller, D. & Friedman, N. *Probabilistic Graphical Models: Principles and Techniques*, Chs. XVII and XVIII, pp. 726, 784 (MIT Press, 2009).
31. Quirk, G.J. *et al.* Erasing fear memories with extinction training. *J. Neurosci.* **30**, 14993–14997 (2010).
32. Karlsson, M.P., Tervo, D.G.R. & Karpova, A.Y. Network resets in medial prefrontal cortex mark the onset of behavioral uncertainty. *Science* **338**, 135–139 (2012).
33. Behrens, T.E.J., Woolrich, M.W., Walton, M.E. & Rushworth, M.F.S. Learning the value of information in an uncertain world. *Nat. Neurosci.* **10**, 1214–1221 (2007).
34. Likhtik, E. & Paz, R. Amygdala-prefrontal interactions in (mal)adaptive learning. *Trends Neurosci.* **38**, 158–166 (2015).
35. Yu, A.J. & Dayan, P. Uncertainty, neuromodulation, and attention. *Neuron* **46**, 681–692 (2005).
36. Rosas, J.M., Todd, T.P. & Bouton, M.E. Context change and associative learning. *Wiley Interdiscip. Rev. Cogn. Sci.* **4**, 237–244 (2013).
37. Courville, A., Daw, N.D., Gordon, G. & Touretzky, D.S. Model uncertainty in classical conditioning. *Adv. Neural Inf. Process. Syst.* **16**, 977–984 (2004).
38. Gershman, S.J., Blei, D.M. & Niv, Y. Context, learning, and extinction. *Psychol. Rev.* **117**, 197–209 (2010).
39. Soto, F.A., Gershman, S.J. & Niv, Y. Explaining compound generalization in associative and causal learning through rational principles of dimensional generalization. *Psychol. Rev.* **121**, 526–558 (2014).
40. Uwano, T., Nishijo, H., Ono, T. & Tamura, R. Neuronal responsiveness to various sensory stimuli, and associative learning in the rat amygdala. *Neuroscience* **68**, 339–361 (1995).
41. Romanski, L.M., Clugnet, M.C., Bordi, F. & LeDoux, J.E. Somatosensory and auditory convergence in the lateral nucleus of the amygdala. *Behav. Neurosci.* **107**, 444–450 (1993).
42. Weinberger, N.M. Auditory associative memory and representational plasticity in the primary auditory cortex. *Hear. Res.* **229**, 54–68 (2007).
43. Pecevski, D., Buesing, L. & Maass, W. Probabilistic inference in general graphical models through sampling in stochastic networks of spiking neurons. *PLoS Comput. Biol.* **7**, e1002294 (2011).
44. Lochmann, T. & Deneve, S. Neural processing as causal inference. *Curr. Opin. Neurobiol.* **21**, 774–781 (2011).
45. Nessler, B., Pfeiffer, M., Buesing, L. & Maass, W. Bayesian computation emerges in generic cortical microcircuits through spike-timing-dependent plasticity. *PLoS Comput. Biol.* **9**, e1003037 (2013).
46. Pouget, A., Beck, J.M., Ma, W.J. & Latham, P.E. Probabilistic brains: knowns and unknowns. *Nat. Neurosci.* **16**, 1170–1178 (2013).
47. Chistiakova, M., Bannon, N.M., Bazhenov, M. & Volgushev, M. Heterosynaptic plasticity: multiple mechanisms and multiple roles. *Neuroscientist* **20**, 483–498 (2014).
48. Abraham, W.C. How long will long-term potentiation last? *Phil. Trans. R. Soc. Lond. B* **358**, 735–744 (2003).
49. Friedman, N. & Koller, D. Being Bayesian about network structure. A Bayesian approach to structure discovery in Bayesian networks. *Mach. Learn.* **50**, 95–125 (2003).
50. Kappel, D., Habenschuss, S., Legenstein, R. & Maass, W. Network plasticity as Bayesian inference. *PLoS Comput. Biol.* **11**, e1004485 (2015).

ONLINE METHODS

Subjects. Male Sprague-Dawley rats (Hilltop) approximately 8 weeks old and weighing 275–300 g (225–250 g for the electrophysiology experiments) on arrival were individually housed on a 12-h light/dark cycle and given food and water *ad libitum*. All animals were naive and had no previous history before the conditioning experiment or surgery appropriate to their group. All procedures were approved by the New York University Animal Care and Use Committee or the Animal Care and Use Committees of the RIKEN Brain Science Institute, and conducted in accordance with the National Institutes of Health Guide for the Care and Use of Experimental Animals.

Viral vectors. Lentiviral vectors (lentivirus-CaMKII-ArchT-GFP) were produced by, and purchased from, the University of North Carolina Vector Core. Previous work[21] has demonstrated specific expression in LA pyramidal neurons using these vectors.

Behavioral conditioning experiments. Animals were placed into a custom-modified Med Associates sound-isolating chamber with Plexiglas walls, illuminated only by infrared light, and underwent one of several conditioning protocols consisting of sequences of CS–US pairings and/or unsignaled USs (UUSs). The CS for all experiments was a series of 5-kHz tone pips (pips at 1 Hz with 250 ms on and 750 ms off) for 30 s. US onset occurred and coterminated with the final pip. The US was a 1-s, 1-mA footshock. Inter-trial intervals (ITIs) between the USs were randomized around 120 s. For the cover stimulus experiment, the light CS was a 30-s flashing white light. Animals were removed from the training context 60 s after the final US of the conditioning protocol (except for animals that received unreinforced context exposure, which were removed 60 s after the end of the last ITI), and spent around 120 s in total outside both the conditioning chamber and the behavioral colony (in the room used for conditioning while the conditioning chamber was cleaned, and in transit to and from the behavioral colony). During the long-term context memory testing phase 24 h later, animals were placed back in the original conditioning context for 330 s. During long-term CS memory testing, animals were placed in a novel, peppermint-scented testing chamber (context B, Coulbourn Instruments), that was different from the conditioning chamber in shape and size, was illuminated by a visible houselight, and had a smooth plastic floor. After a 150-s acclimation period, animals were presented with the identical CS five times, with a randomized ITI of around 120 s. During the training and testing phases the animals' behaviors were recorded on DVD or on a digital storage unit. A rater who was blind with respect to the treatment group scored the animals' behavioral freezing during the first 5 min of the context test and during the 5 CSs, as well as the 2 min before the first CS in the CS test. Scoring was done offline using a digital stopwatch, and freezing was defined as the cessation of all bodily movement with the exception of respiration-related movement. Percentages were calculated as the ratio of time spent freezing to the total time of 300 s for the context memory test, and to the combined 150 s duration of the 5 CSs for the CS memory test. Animals that froze for more than 18 s (15%) of the 2 min before the onset of the first test CS in the novel testing environment of the CS test were excluded from the study, as this freezing interfered with our ability to evaluate the level of the CS memory. The remaining animals showed very low levels of pre-CS freezing (with a mean < 1%). Sample sizes for the different conditioning protocols used for the modeling study are summarized in **Supplementary Table 3**. Eleven animals only received the CS test, but no context test, as noted in **Supplementary Table 3**. These animals were included in the modeling study, but not in the analysis of experiment 1. As the order of the context and CS tests had no statistically significant effect on freezing (**Supplementary Table 4**), context testing was always done first for behavioral experiments with animals that had undergone surgery. All conditioning and testing was done during the light cycle.

Randomization. Animals were randomly assigned to experimental groups before the start of each experiment. Experiments were blocked so that groups alternated and the first group for each day was randomly selected. ITIs were pseudorandom around 2 min.

Stereotaxic cannula implantation, virus injection, and electrode surgery. Animals were anesthetized with a mixture of ketamine/xylazine and implanted with bilateral chronic guide cannulae (22 gauge, Plastics One) above the dorsal hippocampus (stereotaxic coordinates from bregma anterior–posterior –3.8 mm, dorsal–ventral –2.6 mm, medial–lateral 1.5 mm) or the LA (21 gauge, stereotaxic coordinates from bregma anterior–posterior –3.0 mm, dorsal–ventral –6.6 mm, medial–lateral 5.4 mm). For optogenetic experiments simultaneous bilateral injections of 0.5 µl lentivirus were made following cannula placement, through an injector cannula on each side (26 gauge, Plastics One) that protruded 1.4 mm beyond the tip of the guide cannula and was attached to a 1-µl Hamilton syringe (gauge 25s) by polyethylene tubing. Injections were controlled by an automatic pump (PHD 2000, Harvard Apparatus) and were made at a rate of 0.07 µl/min. Injector cannulae were left in place for 20 min after injection and then replaced with clean dummy cannulae.

For awake, behaving electrophysiological experiments, animals were anesthetized as above, and an insulated stainless steel recording wire (1–2 MΩ) (FHC, Inc) attached to a circuit board (Pentalogix) was lowered such that the tip of the electrode targeted the left LA (stereotaxic coordinates from bregma, anterior–posterior –3.0 mm, dorsal–ventral –8.0 mm, medial–lateral 5.4 mm). Additionally, two silver wires, one placed contralaterally and one ipsilaterally above the neocortex, served as a reference and ground respectively. For all experiments, guides and electrode boards were affixed to the skull using surgical screws and dental cement.

Awake-behaving psychopharmacology experiments. Approximately 1 week after dorsal hippocampus cannula surgery, the competitive NMDA receptor antagonist APV (Sigma-Aldrich) was dissolved in saline at a concentration of 10 µg/µl. Animals were taken one by one and injection guides (28 gauge) connected to 1-µl Hamilton syringes (gauge 25s) mounted on an automatic pump (PHD 2000, Harvard Apparatus) were inserted through the implanted cannulae, such that they extended 1 mm below tip of the cannulae. After the injectors were in place, rats received bilateral infusions of ether 0.5 µl saline or 0.5 µl of the 10 µg/µl APV-saline solution (5 µg APV per hemisphere), at a rate of 0.1 µl/min, for 5 min. The injectors were left in place for 4 min after the infusion was completed and then replaced with clean dummy cannulae. Animals were returned to the animal colony for 6 min, after which time the conditioning session began. Conditioning and testing were conducted as described under "Behavioral conditioning experiments."

Awake-behaving optogenetic experiments. Approximately 4 weeks after virus infusion, a fiber optic cable attached to a 532-nm diode-pumped solid state laser (Shanghai Laser and Optics Century Co, Ltd.) was inserted through and screwed onto each of the bilateral cannulae targeting the LA, such that the tip of the fiber optic cable extended 1 mm beyond the tip of the cannula. The tubing surrounding the fiber optic cables was painted black so that the laser illumination caused no perceptible illumination of the conditioning chamber. Rats with the fiber optic cables attached then underwent conditioning as described under "Behavioral conditioning experiments," except that they received laser illumination either occurring 250 ms before UUS onset and lasting 50 ms after UUS termination ('Overlap' group) or an identical laser illumination delayed after the UUS by a random time interval of around 30 s ('Offset' group). The fiber optic cables were also attached to the cannulae before the context test, but no laser illumination was given.

Awake-behaving local field potential physiology. During the first 2 consecutive days of the awake-behaving physiology experiments, animals were taken one by one, attached to the electrophysiological setup, and placed in a novel, peppermint-scented testing chamber conditioning chamber (context C) that was distinct from context A (and context B) in shape and size, illuminated by a visible houselight, with metal bar walls and a plastic floor. After a 5-min acclimation period, animals were habituated with three presentations of the CS (with the same CS as described under "Behavioral conditioning experiments") with a randomized ITI of between 1 and 5 min. LA local field potentials were recorded during these two sessions. The third day all rats were conditioned as previously described in the behavioral conditioning experiments method section. 24 h after conditioning rats were placed back in context C, and after 5 min acclimation 5 CSs were delivered with a random ITI of between 90 s and 150 s, while LA local field potentials and freezing behavior were recorded. CS presentation was automated using Spike2 software (CED, Cambridge, UK). Electrical signals were recorded and analyzed as described previously[21]. Latencies of the A-LFP and the average waveform

amplitudes during habituation for the two groups are given in **Supplementary Table 5**. Statistical comparisons were made using two-tailed unpaired *t*-tests, and two-way ANOVA for the latencies.

Awake-behaving extracellular single-unit physiology. For single-unit electrophysiological studies, rats received chronically implanted microdrives consisting of 16 stereotrode bundles (0.001-inch insulated tungsten wire (diameter 25 μm), California Fine Wire Company) and eyelid wires for shock delivery[51]. Following recovery from surgery, daily screening sessions were conducted until single, shock responsive units were isolated (testing was done using mild, single-pulse (2 ms) 1-mA eyelid shocks). Animals then received intermixed shocks (12 trials of each condition) alone (2 ms, 2 mA at 7 Hz for 1 s) or shocks with laser illumination (589 nm, Shanghai Laser Company). Laser onset occurred 400 ms before shocks and was turned off 50 ms after US termination. Spike data were acquired using a Neuralynx data acquisition system. Spike clustering and single unit isolation were performed using Neuralynx SpikeSort 3D software and spiking data. Single unit isolation was assumed if spike trains had a refractory period of greater than 1 ms and a mean spike amplitude of at least 70 μV.

Histology. After behavioral testing was completed, animals were anesthetized with an overdose of chloral hydrate and perfused with paraformaldehyde (for optogenetic experiments) or with either 10% buffered formalin or Prefer (Anatech, Ltd.). For animals with electrode implants, the location of the electrode was marked by passing a small current (4 μA; 5 s) through the electrode tips before perfusion. Following perfusions, brains were sectioned into 40-μm coronal slices and stained with Nissl (Sigma-Aldrich, C5042, staining only for animals with electrode implants or hippocampal cannulation). An experimenter blind to the identity of the animal and treatment assessed the placement of the cannulae, electrodes and virus expression. For animals to be included in the analysis of the optogenetic experiment, Arch-T had to be expressed in LA neurons, with the tip of the each guide cannula dorsal and proximal to the LA (**Supplementary Fig. 5**).

Statistical analysis. Experiments 1–3 had a two-way design and were analyzed accordingly with a two-way ANOVA model with interaction. CS and context scores were analyzed separately. Experiment 4 was analyzed using unpaired *t*-tests. We tested for normality using a Lilliefors test with a critical value of 0.01, and for equality of variances in experiments 1–3 using Levene's test. The groups compared were found to be normally distributed with equal variances, with two exceptions. The Lilliefors test was significant for the context test scores for control II group in experiment 1. However, given the large sample size ($n = 22$) in this experiment and the strong negative result ($P > 0.8$) for a difference between control II and pairings-first groups, the result of the ANOVA test can be expected to be robust to this violation. Levene's test found unequal variances among the context test scores in experiment 2, since the scores from the APV groups tended to lie very close to 0, resulting in a small variance. We used the Keppel correction to correct for this violation by substituting $\alpha/2$ for the original critical value $\alpha = 0.05$. Since our *P* value was very small ($P = 0.0043$), changing the critical value had no effect on the test's conclusion, and our result is expected to be robust against this violation. We also found unequal variances using the two-sampled *F* test for both freezing scores and amplitude changes in experiment 4, and accordingly used an unpaired two sample *t*-test with unequal variances. Since repeated-measures ANOVAs can be especially susceptible to violations in sphericity, we used a lower bound correction when sphericity was violated (**Supplementary Figs. 2** and **7**).

F and *P* values for interaction and main effects, as well as for simple effects, are summarized in **Supplementary Table 6**. For simple effects we report the individual *P* values, as adjusting for multiple comparison by the Holm-Sidak procedure did not affect statistical significance. We also used a two-way ANOVA to evaluate the effect of the order of the context and CS tests for data from experiment 1, as well as using data from all the behavioral experiments where the order of testing was varied (**Supplementary Table 4**). We measured the effect size of contingency on CS memory in experiment 1 and performed power analysis to determine an appropriate range of sample sizes for the subsequent experiments. The effect size of $f = 0.31$ fell in the medium (0.25) to high (0.40) range for this type of test, with a power of 0.78. We set the target sample size for experiments 2 and 3 to detect a strong effect ($f = 0.4$) with a power of at least 0.6, requiring a total *n* of at least 33. The *t*-test comparing changes in A-LFP amplitudes in experiment 3 had an effect size of 1.15.

To compare means of discrete measures, such as defecation (**Supplementary Fig. 3**), we used the Mann-Whitney *U* test. All tests used in this study were two-tailed. Mean and standard error values for our data are listed in **Supplementary Table 7**.

Bayesian network models. The Bayesian network models represented the environment with graphs over four binary variables, the background, context, CS (tone) and US. For notational simplicity we will also refer to these as X_1, X_2, X_3 and X_4 respectively, or as the vector of variables $\mathbf{X}$, with each taking either the value 0 (absent) or 1 (present). For each training protocol, a series of observations $\mathbf{X}^t$ was summarized into counts of the eight different configurations of the four binary variables (eight rather than sixteen, since the background, by definition, will always be 'present' during the experiment). We adopted the use of the background variable from causal learning models, to represent the sum of all unobservable or unspecified influences on our system (in particular, on the US occurrence). As such, the background will always be present during learning but absent for predictions during recall, and an edge from the background to the US ($X_1 \rightarrow X_4$) present in all graphs. An alternative to having the background variable is to specify a prior distribution (e.g., β) for the probability of US occurrence for the case when the US has no parent variables or when all of its parent variables are absent, allowing one to calculate likelihoods of observations. This can yield to a similar fit as the original SLM, but the background variable from the PLM is highly detrimental to its fit. See **Supplementary Modeling** and **Supplementary Tables 1** and **2**.

We considered potential edges that conform with the ordering $X_i \prec X_j$ iff $i < j$: between the context and the US $X_2 \rightarrow X_4$, the tone and the US $X_3 \rightarrow$ US and between the context and the tone $X_2 \rightarrow X_3$, with corresponding parameters $\omega_{1,4}$, $\omega_{2,4}$, $\omega_{3,4}$ and $\omega_{2,3}$, respectively. We assumed that the animals learn this ordering because of the temporal order and duration of the stimuli. Edges between variables represented noisy-or generating functions, corresponding to the assumption that different parent variables predict a child variable independently (analogous to independent generative causes, but without making assumptions about causality). For edges with parameters $0 \leq \omega_{i,j} \leq 1$, the relevant probabilities are then given by

$$P\left(X_4 = 1 \mid \mathrm{Pa}(X_4)\right) = 1 - \left(\prod_{X_i \in \mathrm{Pa}(X_4)} (1 - \omega_{i,4})^{x_i} \right) \quad (1)$$

and

$$P(X_3 = 1 \mid X_2 = 1) = \omega_{2,3}, \text{ when } X_2 \rightarrow X_3 \quad (2)$$

with uniform $U[0,1]$ priors for

$$P(X_3 = 1 \mid X_2 = 0) \text{ when } X_2 \rightarrow X_3, \text{ and } P(X_3 = 1) \text{ when } X_2 \not\rightarrow X_3 \quad (3)$$

$\mathrm{Pa}(X)$ is the parent set of X (all the variables sending edges to X). More details about the temporal representation of trials and about the assumptions of the model about the stationary nature of the environment are available in the **Supplementary Modeling**.

Structure learning (SLM). For SLM we calculated the posterior distribution over different Bayesian network structures, without assuming or learning specific parameter values ω_i for the edges. We considered the six possible graph structures $G_i \in G$ that can lead to different predictions about the US (**Fig. 5a**). In graphs 1 and 2, leaving out, or adding the edge $X_2 \rightarrow X_3$ is irrelevant when making predictions about the US, we therefore considered only one of each of these pairs of functionally equivalent graphs (the one with no $X_2 \rightarrow X_3$ edge).

By Bayes' rule

$$P(G_i \mid D) \propto \int_{\omega_{|G_i}} P(G_i) \cdot P(\omega \mid G_i) \cdot P(D \mid G_i, \omega_{|G_i}) \mathrm{d}\omega_{|G_i} \quad (4)$$

doi:10.1038/nn.4308

so to calculate the posterior probability of a graph structure G, we integrated out the parameters in the graph, assuming that each comes, independently, from the uniform distribution $U[0,1]$. We fitted a prior $P(G_1) = \rho$ for the minimally connected graph G_1, to account for the fact that the CS and the context are initially largely neutral stimuli that do not predict threats. The other graphs had equal priors

$$P(G_i) = \frac{1-\rho}{5}$$

Unlike parameter priors, which strongly influence structure learning no matter the amount of data, the effect of these structure priors on the predictions of the model becomes less important as the number of training trials increases (i.e., as the data overwhelmed the priors).

The likelihood term for a graph G_i is the probability of observing a particular combination of stimuli during a complete training protocol, given a graph structure G and parameters $\omega_{|Gi}$ (for the edges present in G_i). To calculate this probability, we took the product over the sequence of observations $\mathbf{X}^t$ so that

$$P(D \mid G_i, \omega_{|G_i}) \propto \prod_{t=1:T} P(X_4^t = x_4^t \mid \{x_j^t, \omega_{j,4}^t\}_{j \in Pa(X_4)}) \cdot P(X_3^t = x_3^t \mid \{x_j^t, \omega_{j,3}^t\}_{j \in \mathrm{Pa}(X_3)}) \tag{5}$$

T is the total number of time bins during the experiment, including the time outside the training context (see **Supplementary Modeling**). For notational simplicity we chose to write the integral as integrating over a sequence of trials, rather than counts of a specific trial type, but the two approaches are of course equivalent.

To calculate the posterior probability of a feature f, such as particular edge, or a path, we used model averaging over the graph structures

$$P(f \mid D) = \sum_{G_i \in \mathbf{G}} P(G_i \mid D) f(G_i) \tag{6}$$

where $f(G)$ is 0 or 1, depending on whether the feature f is in graph G or not. Such model averaging is a popular tool for prediction problems when limited data means that the posterior distribution over graphs is not peaked at a single structure (i.e., the choice of a single structure for predictions is inappropriate).

The behavioral response to the tone CS is then predicted to be proportionate to the posterior probability of the edge $X_3 \rightarrow X_4$:

$$\text{CS elicited response} \propto P(X_3 \rightarrow X_4 \mid D)$$

The context can be connected to the US both by a direct edge $X_2 \rightarrow X_4$ and indirectly through the path $X_2 \rightarrow X_3 \rightarrow X_4$. In cases where a direct connection does not exist, an indirect connection still signifies statistical dependency in cases when the intermediate variable(s) cannot be observed. Such a connection can therefore serve as a basis for a (possibly weaker) behavioral response. Such a weaker response has been observed in various studies in the form of second-order conditioning, or facilitation. Such a relationship could be represented in the brain by disynaptic or polysynaptic connections, resulting in a weaker feedforward response. We therefore introduced a second model parameter α, $0 \le \alpha \le 1$, that reflects a discounting factor for such secondary relationships, as well as weighing this indirect context-US relationship by a simple estimate of the context–CS association, depending on the frequency with which the CS appeared in the context,

$$\frac{\text{trials with CS}}{\text{total trials}}$$

$$\text{Context-elicited response} \propto P(X_2 \rightarrow X_4 \mid D) + \gamma \cdot P(X_2 \rightarrow X_3 \rightarrow X_4, X_2 \nrightarrow X_4 \mid D)$$

where

$$\gamma = \alpha \cdot \frac{\text{trials with CS}}{\text{total trials}}$$

for some constant α. This parameter could also potentially account for influences of temporal discounting, as well as substituting for the need to specify a nonuniform prior distribution for the CS-on probability.

For training protocols where no tone was played, the posterior is calculated only over two structures with the variables X_1, X_2 and X_4. Here

$$P(G_i \mid D) \propto \int_{\omega_{|G_i}} \left(P(G_i) \cdot \prod_{t=1:T} P(X_4^t = x_4^t \mid \{x_j^t, \omega_{j,4}^t\}_{j \in Pa(X_4)}) \right) d\omega_{|G_i} \tag{7}$$

where $P(G)$ is a uniform prior. In this case we have

$$\text{Context-elicited response} \propto P(X_2 \rightarrow X_4 \mid D)$$

Since the integrals in (5) and (6) cannot be evaluated analytically when any variable has two or more parents, we used a Monte Carlo stimulation to approximate their value. For each calculation, 2.5×10^5 samples of the parameter vector ω were drawn from a uniform distribution, and the resulting likelihoods were averaged over. For given parameters ρ and α, this gave predicted behavioral responses for all the different conditioning protocols.

Parameter learning (PLM). PLM predicts behavioral responses based on learning the posterior mean of the parameter values in the maximally connected graph, Graph 6 (**Fig. 5a**). For parameter $\omega_{j,k}$ (for the edge $X_j \rightarrow X_k$) using the joint prior over ω, we have

$$\hat{\omega}_{j,k} = E(\omega_{j,k} \mid D, G_6) = \int_{\omega|G_6} \omega_{j,k} \cdot P(\omega) \cdot P(D \mid G_6, \omega) \mathrm{d}\omega \tag{8}$$

During retrieval, predictions are based on standard inference in the network with parameters $\hat{\omega}$ such that

$$\text{CS-elicited response} \propto P(X_4 = 1 \mid X_3 = 1, X_1 = X_2 = 0) = \hat{\omega}_{3,4}$$

$$\begin{aligned} \text{Context-elicited response} &\propto P\left(X_4 = 1 \mid X_2 = 1, X_1 = 0\right) \\ &= P(X_4 = 1 \mid X_2 = 1, X_1 = 0, X_3 = 0) \cdot P(X_3 = 0 \mid X_2 = 1) \\ &\quad + P(X_4 = 1 \mid X_2 = 1, X_1 = 0, X_3 = 1) \cdot P(X_3 = 1 \mid X_2 = 1) \\ &= (1 - \hat{\omega}_{2,3}) \cdot \hat{\omega}_{2,4} + \hat{\omega}_{2,3} \cdot (\hat{\omega}_{2,4} + \hat{\omega}_{3,4} - \hat{\omega}_{2,4} \cdot \hat{\omega}_{3,4}) \end{aligned}$$

Each of the four Bayesian network parameters $\omega_{j,k}$ has an independent β prior distribution. Fitting the model thus includes finding a pair of parameters for each of these four prior β distributions (eight parameters in total), such that they best explain the behavioral data across all training protocols. We carried out this optimization using a genetic algorithm separately for different discretization parameters t that determined the temporal subdivision of the 2-min trials. We allowed some flexibility toward the discretization of the CS in form of a binary choice when t does not uniquely determine the discretization of the CS (e.g., when $t = 5$ the CS could be both length 2 or length 1). We found that the default discretization of $t = 1$ provided a considerably better fit than all other values of t.

Learning both structure and parameters (PSLM). Learning a full posterior over the Bayesian network representations includes first learning a distribution over the graph structures as in SLM, and then learning a posterior distribution for the parameters present for each structure as in PLM, but separately for each graph. Predictions are then made by averaging over predictions from the different graph structures weighed by the posterior probability of each graph.

$$\text{CS-elicited response} \propto \sum_{G_i \in \mathbf{G}} P(G_i \mid D) \cdot \hat{\omega}_{3,4}^i$$

$$\text{Context-elicited response} \propto \sum_{G_i \in \mathbf{G}} P(G_i \mid D) \cdot \left[(1 - \hat{\omega}_{2,3}^i) \cdot \hat{\omega}_{2,4}^i + \hat{\omega}_{2,3}^i \cdot (\hat{\omega}_{2,4}^i + \hat{\omega}_{3,4}^i - \hat{\omega}_{2,4}^i \cdot \hat{\omega}_{3,4}^i) \right]$$

where for notational simplicity we take $\hat{\omega}_{jk}^i = 0$ if the corresponding edge is not present in graph G_i. We assumed a prior over the relevant graph structures as in

SLM, with ρ a free parameter and uniform priors for parameters for structure learning, while fitting β-distributed edge-parameter priors for the parameter learning part as in PLM, such that these priors were shared across all graph structures where the respective parameters were present. Accordingly, this model fits nine parameters.

Model fitting. Since our data has two distinct measures (CS freezing and context freezing), we used mean squared error (MSE) rather than R^2 as a measure of model fit. To evaluate the fit of the model, predicted responses were scaled to freezing scores by multiplication with a scaling factor found by linear regression. This was done separately for the context and for the CS freezing scores, since the different behavioral testing procedures are likely to result in different scaling factors. Mean squared error (MSE) was then calculated by summing these squared error terms and dividing by 29, the number of different conditioning protocols. The best-fitting parameters were found using a genetic algorithm, using MATLAB's ga function. Parameters for the β priors were constrained to lie between 0.01 and 30. For each model, we repeated the optimization process at least ten times, with each run giving approximately the same minimum error values. Values for α and $P(G_1)$ for SLM (and SPLM) were also consistent across runs, but the best-fit β parameters varied, since different pairs of β parameters can determine very similar distributions. For each model, we then averaged over twenty runs with each of the 50 best-fitting set of parameters found during the optimization process and chose the set of parameters that gave the smallest average error (**Supplementary Table 1**). For each model, we checked the feasibility of this optimization by generating a data set from the model using four sets of randomly generated parameters. These four data sets per model were scaled so that the means of the CS and context scores matched the means from the behavioral data (so that MSEs could be appropriately compared). Fitting these generated data sets by the procedure outlined above (but running the genetic algorithm only once rather than ten times), we obtained MSEs that were all below $0.05\%^2$. Best-fit parameters for the models are listed in **Supplementary Table 8**.

Modeling neural interventions. The inactivation of the hippocampus during learning was modeled using the best-fit parameters from the behavioral data, and removing the variable X_2 and the corresponding edges from the model (or equivalently, by setting the prior for all these edges to the delta function), and calculating the predicted CS-elicited freezing scores. Amygdala inactivation during the US was modeled by excluding trials with inactivation from the trial counts (such that they counted neither toward the reinforced nor the unreinforced trials).

A **Supplementary Methods Checklist** is available.

Code availability. All MATLAB (R2014b) scripts used to fit and compare the computational models are available upon request.

Data availability. The data that support the findings of this study are available from the corresponding author upon request.

51. Johansen, J.P., Tarpley, J.W., LeDoux, J.E. & Blair, H.T. Neural substrates for expectation-modulated fear learning in the amygdala and periaqueductal gray. *Nat. Neurosci.* **13**, 979–986 (2010).

doi:10.1038/nn.4308

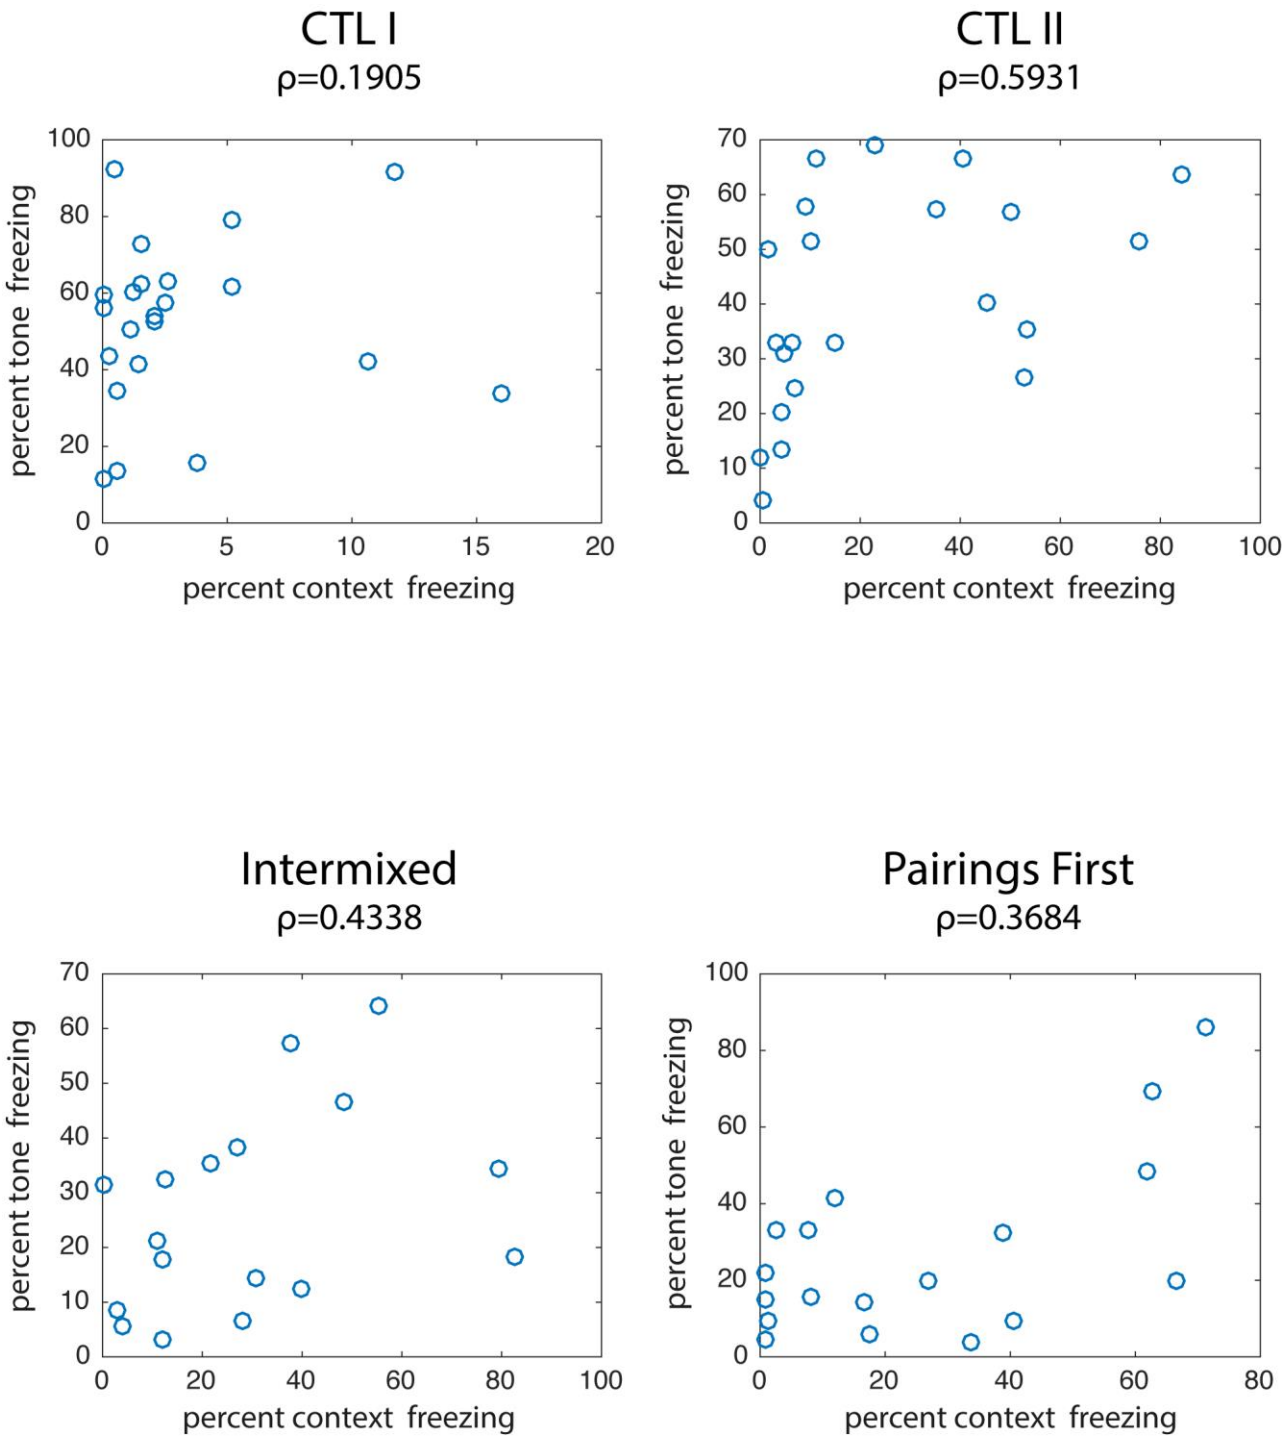

Supplementary Figure 1

Correlations between tone and context freezing by animal in each of the four groups in experiment 1.

Every animal is represented by a blue circle. Correlation was measured by Spearman's rank correlation coefficient (ρ).

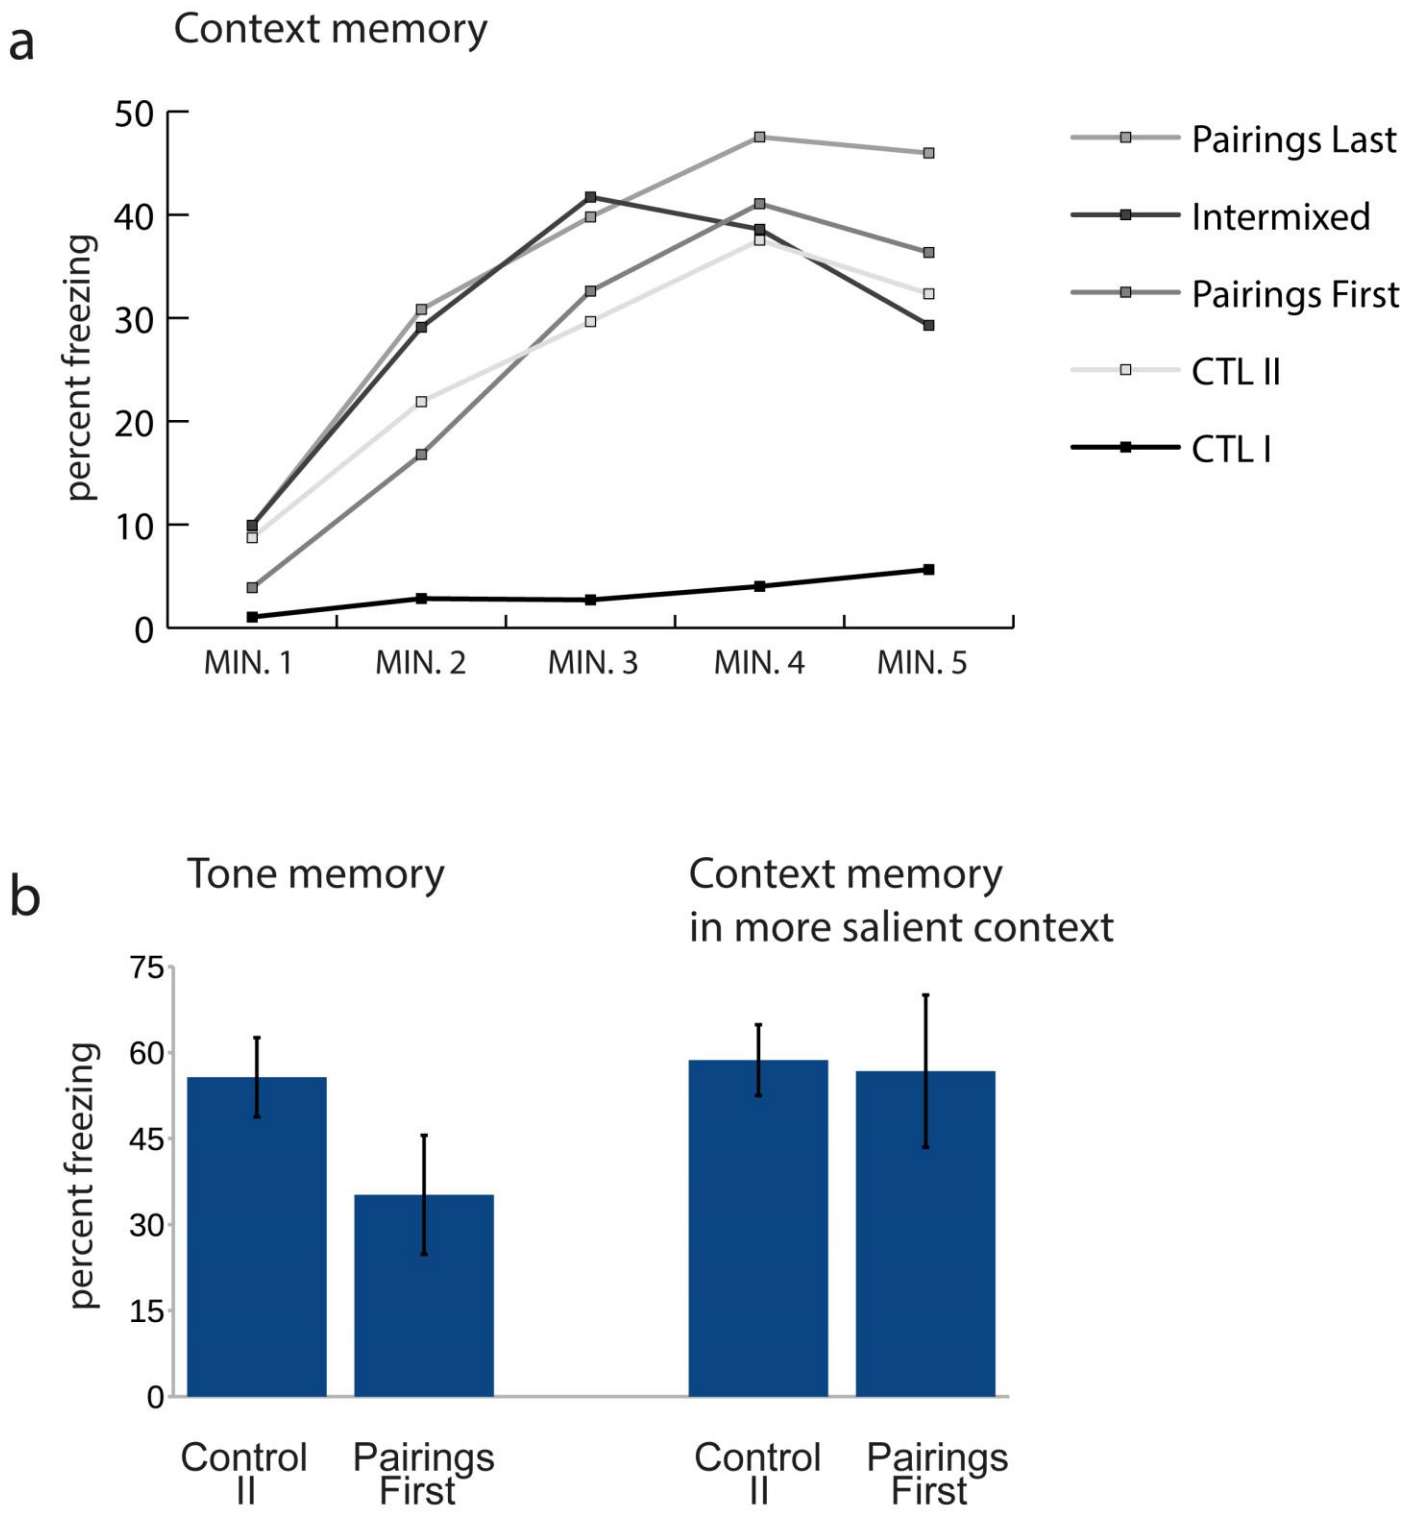

Supplementary Figure 2

Comparisons of context memory between groups are stable over time and invariant to the salience of the conditioning context.

(a) Minute-by-minute analysis of freezing during the context test for the groups in experiment 1 (and Pairings Last, included for illustration but not in the statistical analysis). A repeated measures ANOVA showed no Time*Contingency*Spacing interaction (n = 16,17,18,20,21, $F_{1,288} = 1.96$, $P = 0.17$). A comparison restricted to the massed condition (between CTL II and Pairings First) also showed no Time*Contingency interaction ($F_{1,144}) = 0.74$, $P = 0.40$). (b) Comparison of CTL II and Pairings First groups with conditioning and context test performed in a more salient context (lit by a visible light and with citrus odor). Reduction in Tone memory matched previous result (ratio between CTLII and Pairings First 0.63 vs. 0.66 originally), whereas Context memory was similar between the groups, as before. Error bars indicate s.e.m.

Nature Neuroscience: doi:10.1038/nn.4308

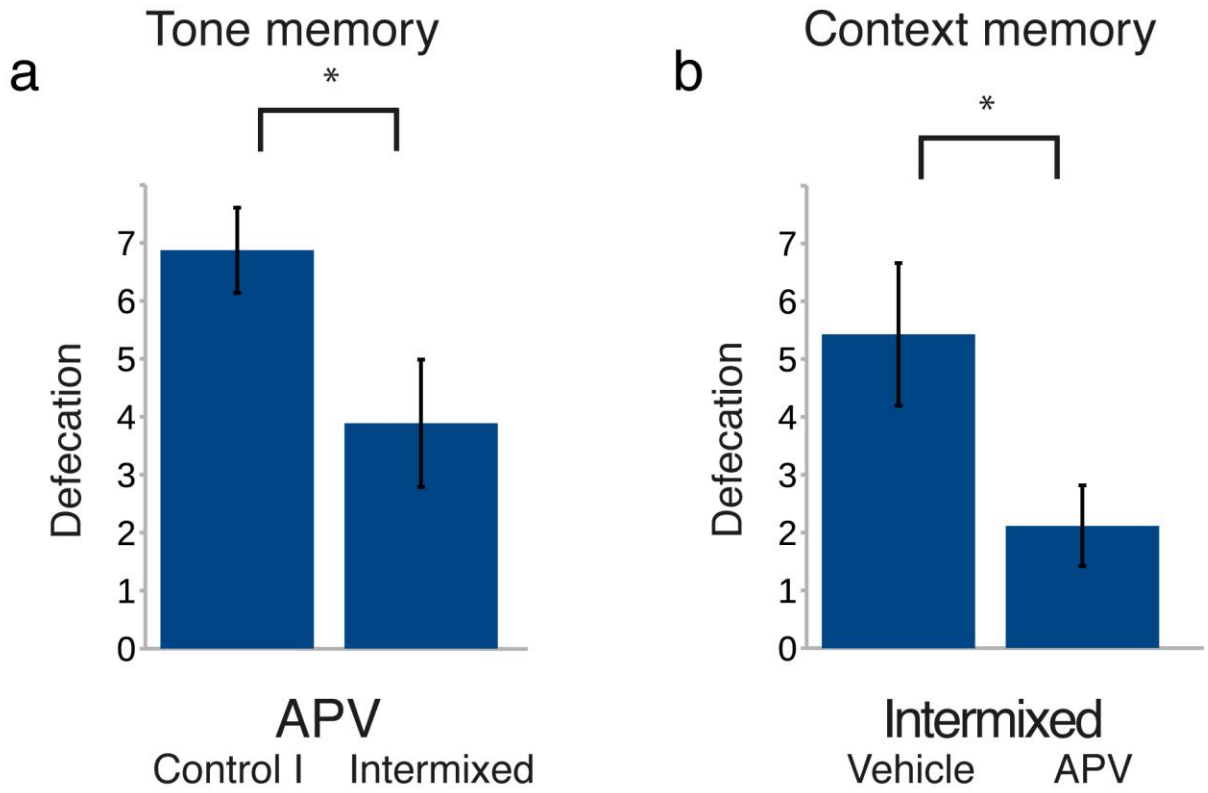

Supplementary Figure 3

Defecation gives similar results to freezing for experiments described in Figure 2.

(a) Contingency degradation in the Intermixed condition with APV infusion in dorsal hippocampus (DH) prior to conditioning, as measured by defecation during tone test. (n = 9,9, Mann-Whitney U test, U = 17.5, $P = 0.04$). (b) Impaired contextual aversive memory, as measured by defecation, following APV infusion in DH prior to conditioning (n = 9,7, Mann-Whitney U test, U = 9.5, $P = 0.018$). Error bars indicate s.e.m.

T = Tone-shock pairing
L = Light-shock pairing
U = Unsignalled shock (shock alone)
c = Unreinforced Context exposure

a

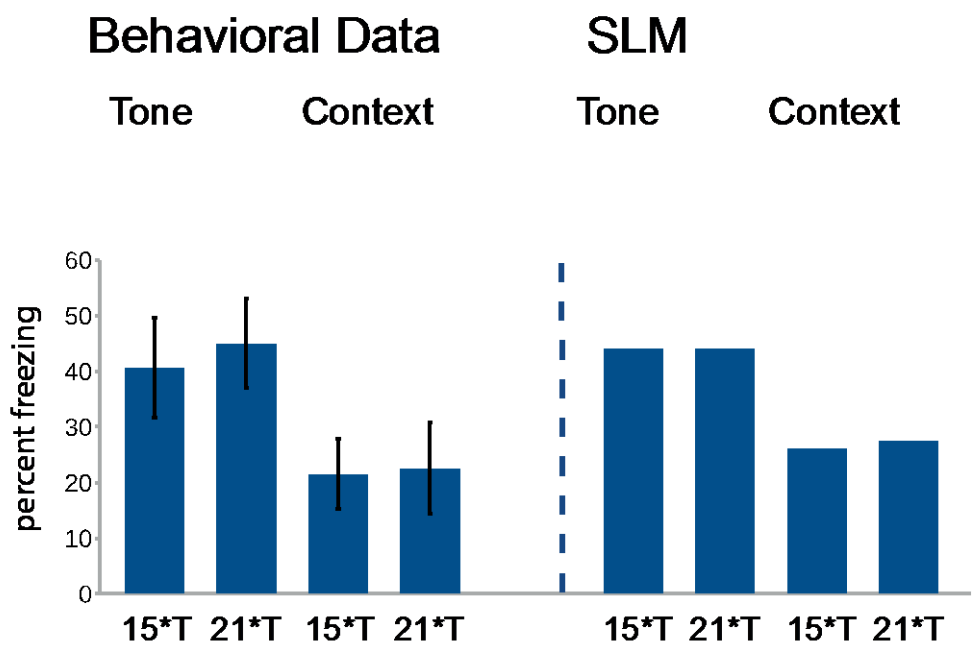

b

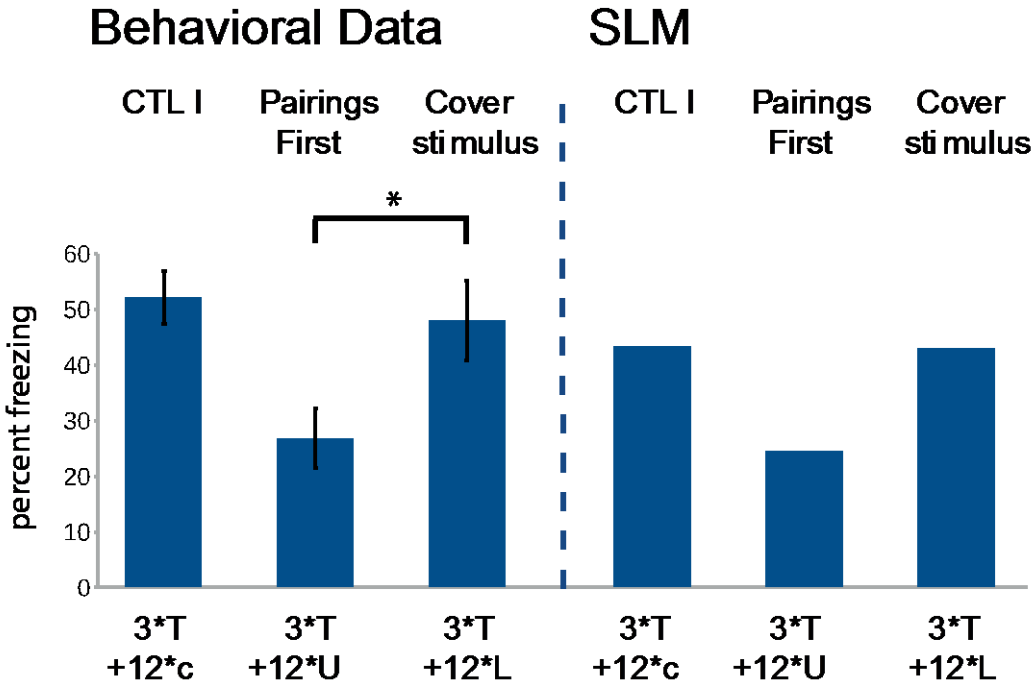

Nature Neuroscience: doi:10.1038/nn.4308

Supplementary Figure 4

Effect of repeated USs depends on contingencies.

(a) Behavioral data (left) and model simulation (right) for conditioning with 15, and 21 CS-US pairings (n=9, 7). Adding further shocks paired with the same tone CS (21 pairings in total) did not reduce tone memory strength. (b) Behavioral data (left) and model simulation (right) for the cover stimulus effect. Signaling shocks with a second CS (in this case a flashing light), instead of giving unsignaled shocks attenuates contingency degradation (comparison between Pairings First and Cover Stimulus groups, n = 18, 12, unpaired sample t-test, $t_{28} = 2.42$, * $P = 0.022$) Error bars indicate s.e.m.

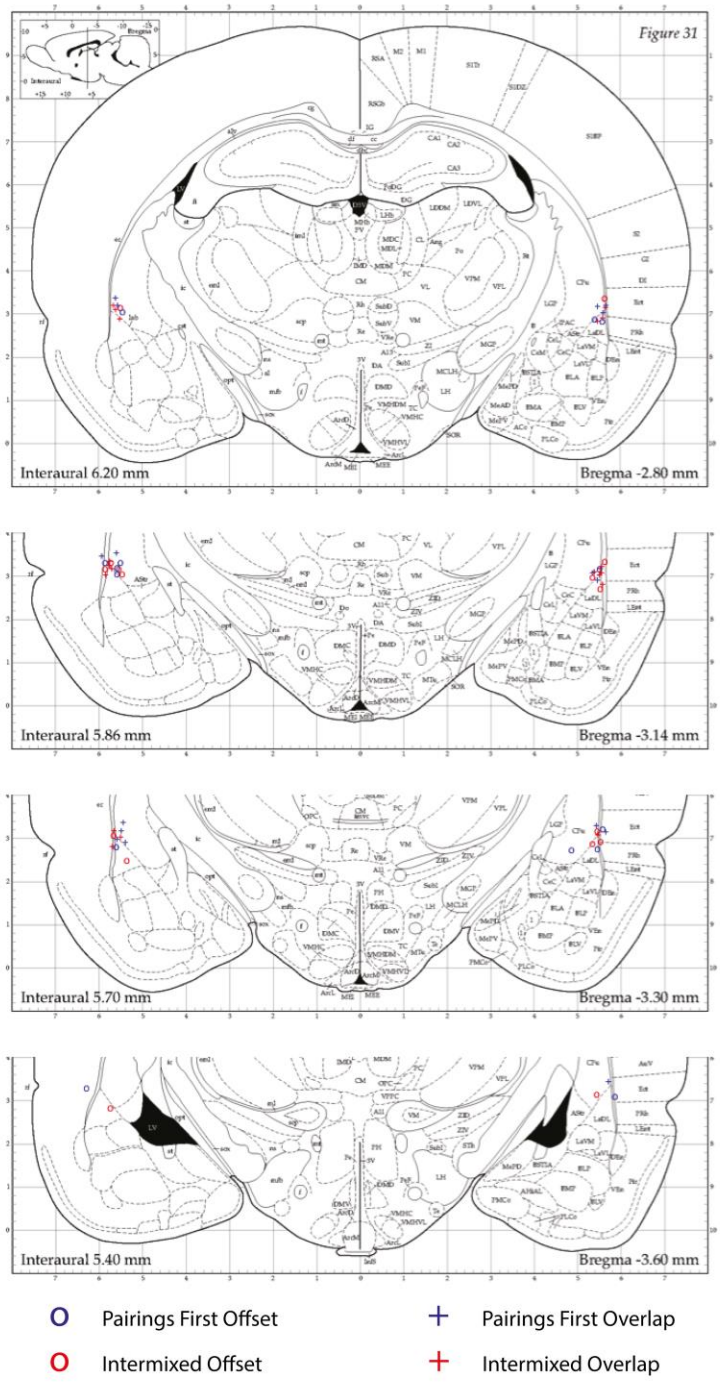

Supplementary Figure 5

Location of the optical fiber tips for optogenetic experiments.

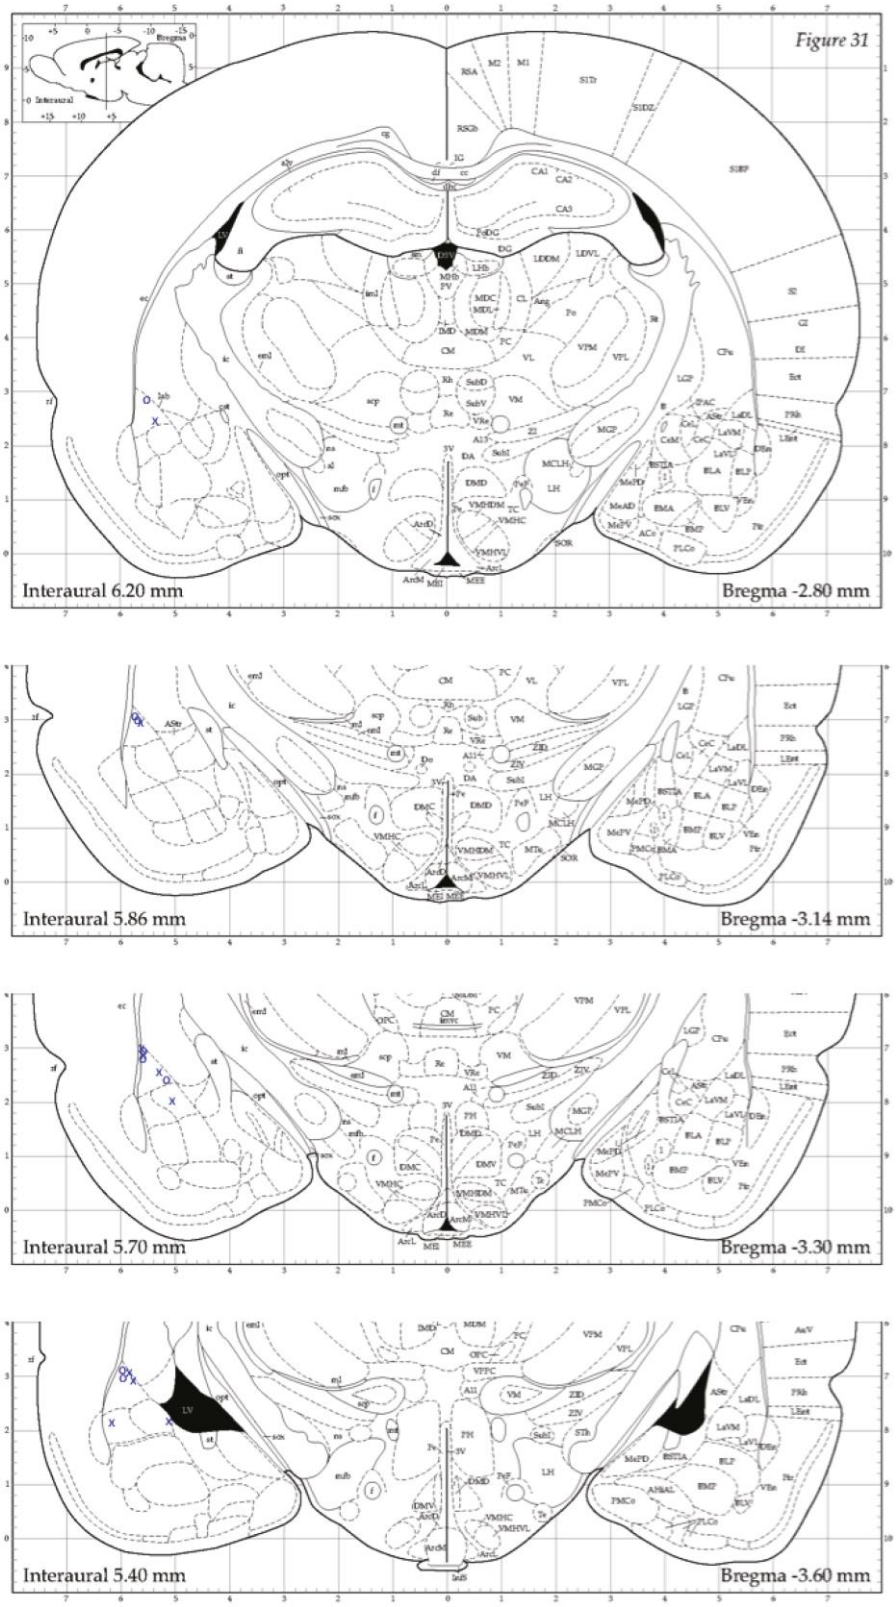

Nature Neuroscience: doi:10.1038/nn.4308

Supplementary Figure 6

Location of the electrode tips for electrophysiological experiments.

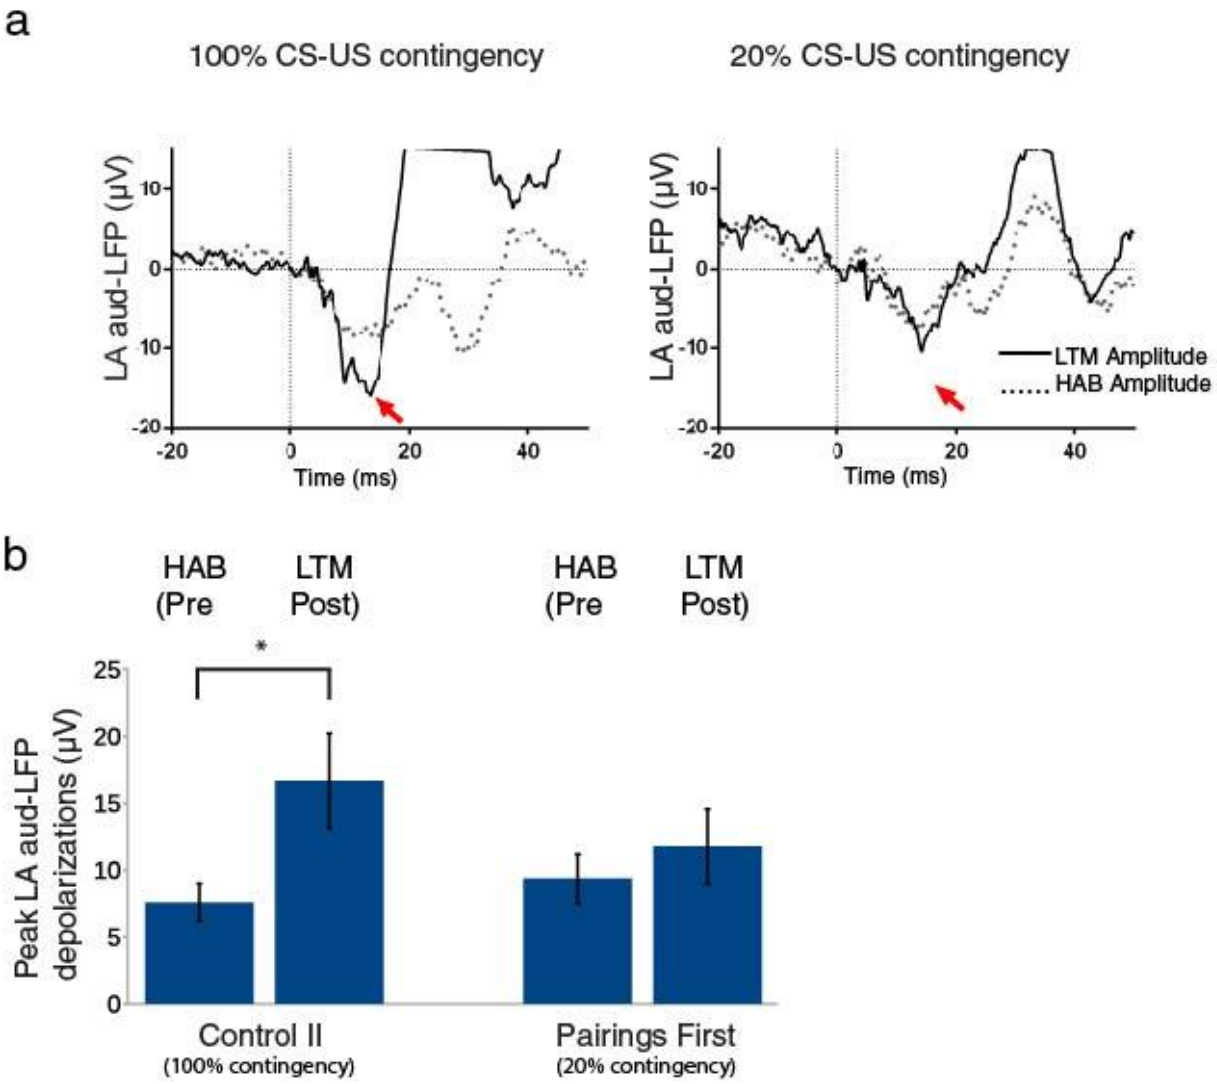

Supplementary Figure 7

Effects of contingency on amygdala LFP potentiation.

(a) Example traces before, and after conditioning for a representative animal each in the Control II group (left) and Pairings First group (right). Red arrows indicate the peak depolarization. (b) Averaged peak depolarizations in the CTL II and Pairings First groups before (Habituation), and after (LTM) conditioning. There was a marginally significant interaction between time and contingency (n=10, 8, repeated measures ANOVA, $F_{1,16}$ = 4.30, P=0.055), and a simple effects analysis showed significant potentiation of the LFP response in the CTL II, but not the Pairings First condition ($F_{1,16}$ = 18.0, P = 0.001 and $F_{1,16}$ = 1.022, P = 0.33), further indicating that conditioning differentially effects synaptic processing depending on contingency. Error bars indicate s.e.m.

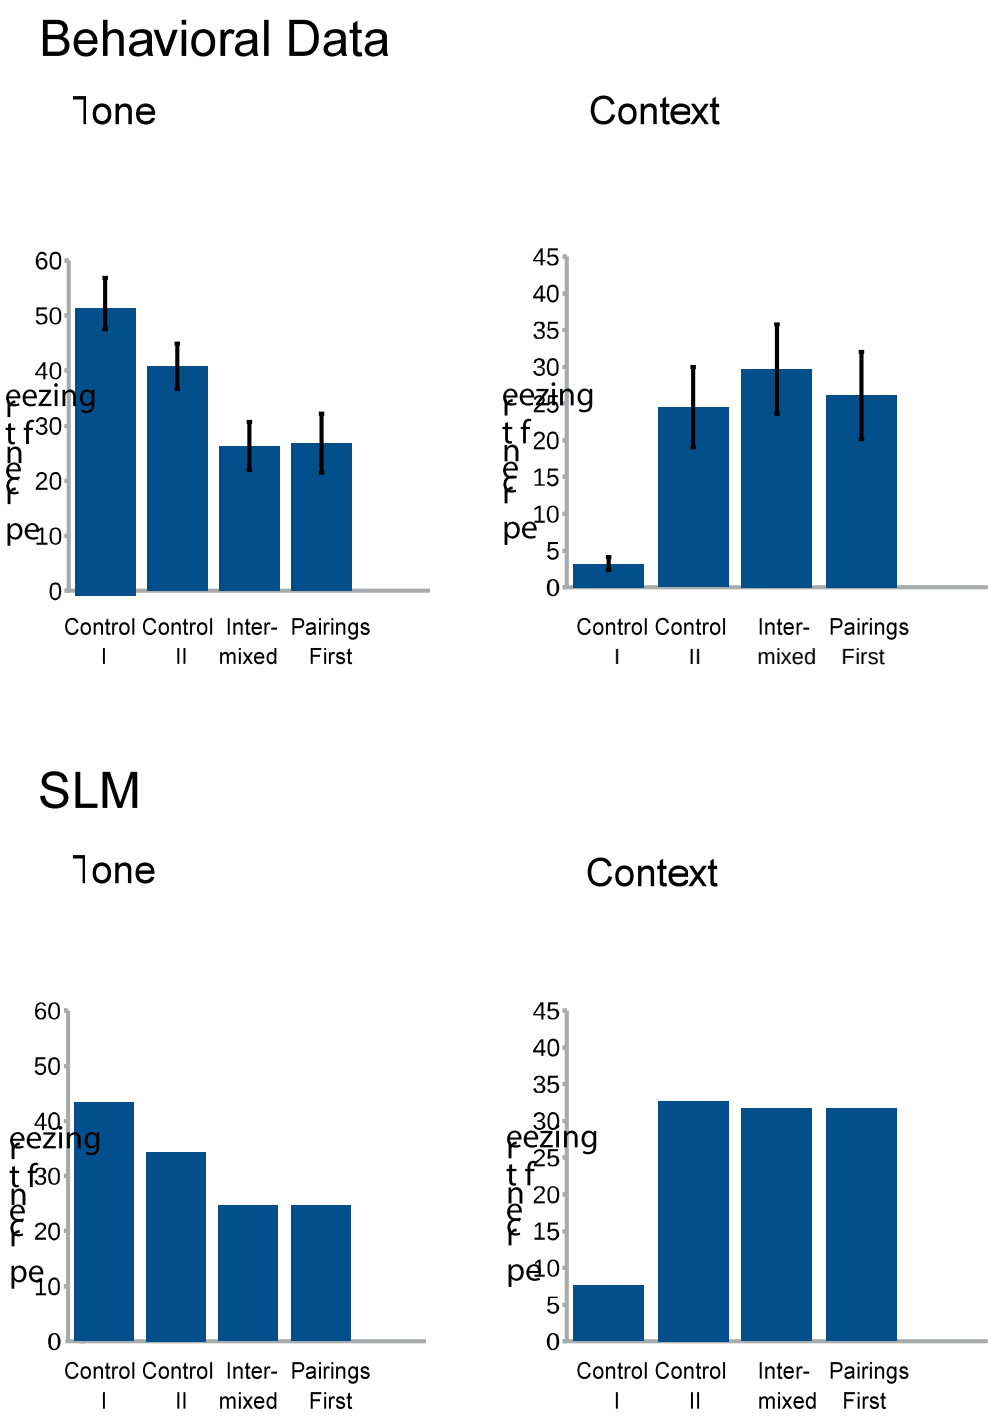

Supplementary Figure 8

Comparison of SLM to behavioral results for experiments described in Figure 1.

Direct comparison of behavioral data (top panel) and SLM (bottom panel) for experiment 1. Error bars indicate s.e.m.

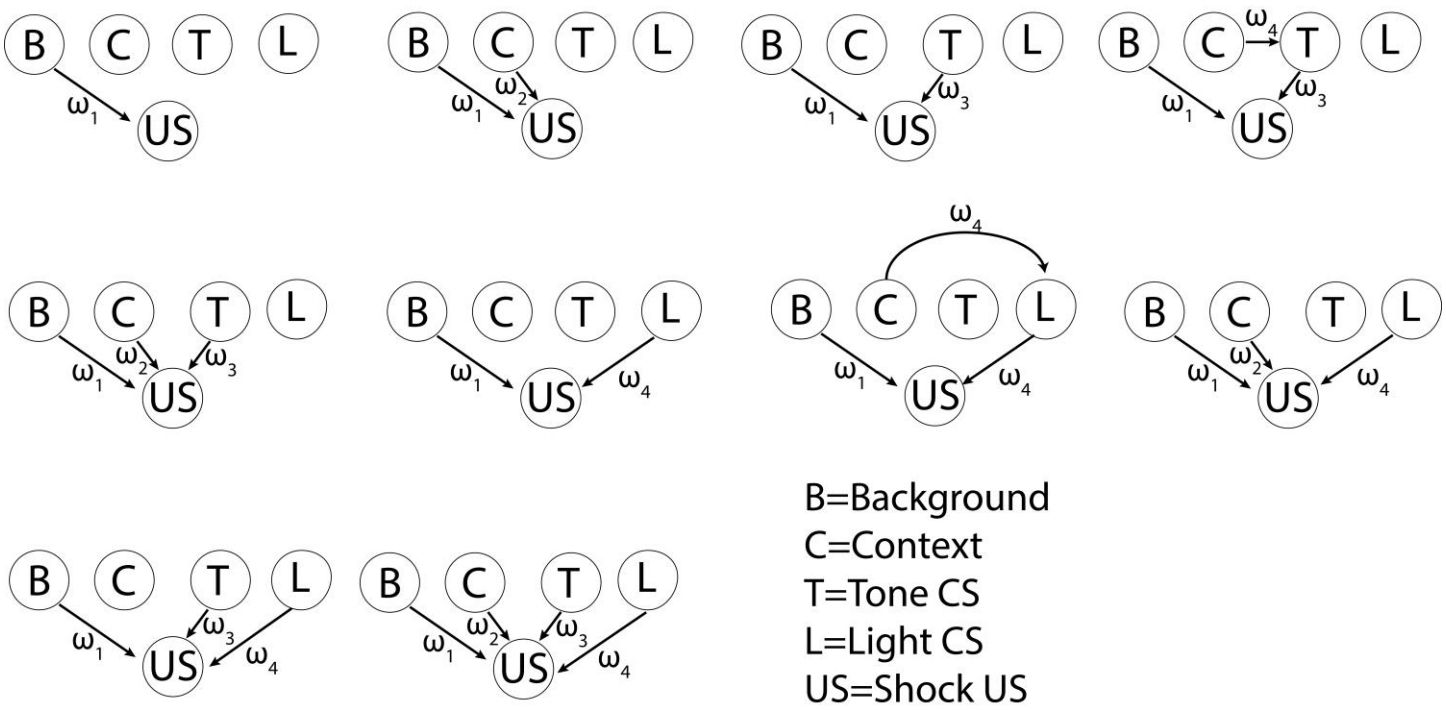

Supplementary Figure 9

The graph structures used for SLM, extended to include a second discrete variable.

T = Tone-shock pairing
L = Light-shock pairing
U = Unsignalled shock (shock alone)
CD=contingency degradation

t = Unreinforced Tone presentation
l = Unreinforced Light presentation
c = Unreinforced Context exposure

a

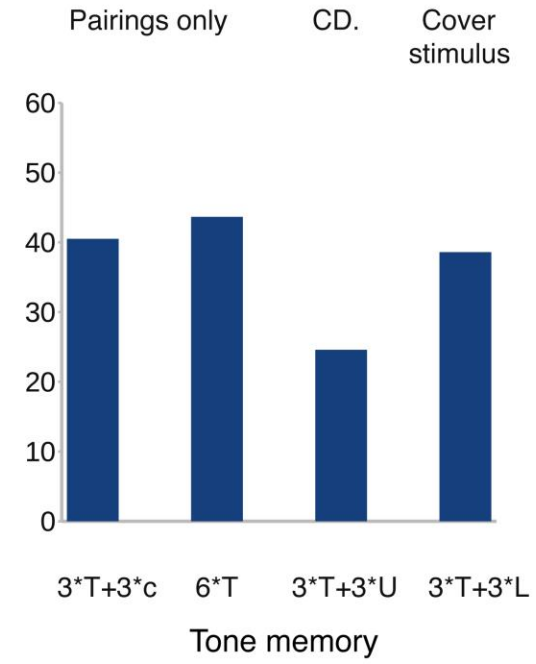

c

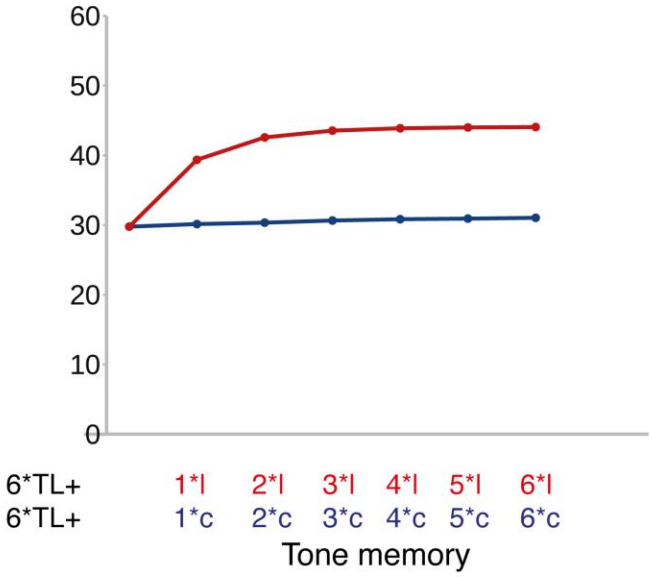

b

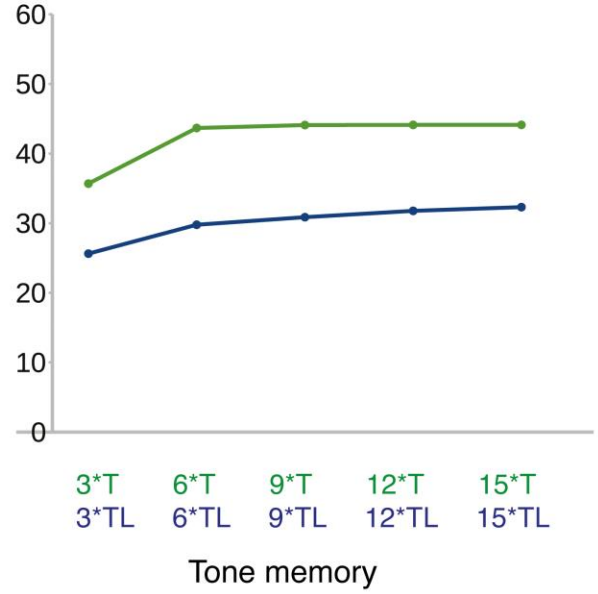

d

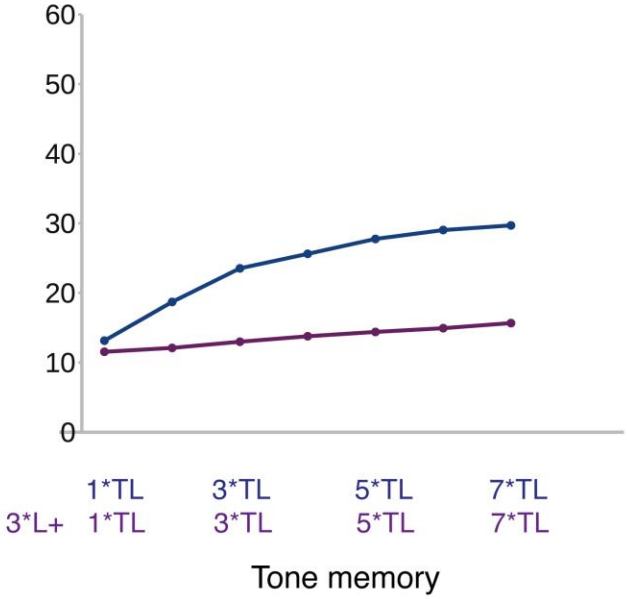

Nature Neuroscience: doi:10.1038/nn.4308

Supplementary Figure 10

SLM's predictions for further conditioning phenomena.

(a) Cover stimulus effect: Replacing unsignaled USs by USs signaled by a second discrete cue (e.g. a light) reverses the effects of contingency degradation. (b) Overshadowing: Conditioning to a single cue (Tone) is reduced if it is trained in compound with a second cue (Light). (c) Recovery from overshadowing: Unreinforced presentations of the overshadowing second cue (Light) restores the level of responding to the first cue. (d) Blocking: Initial conditioning to a Light reduces subsequent conditioning to the Tone when the Tone is conditioned in compound with the Light.

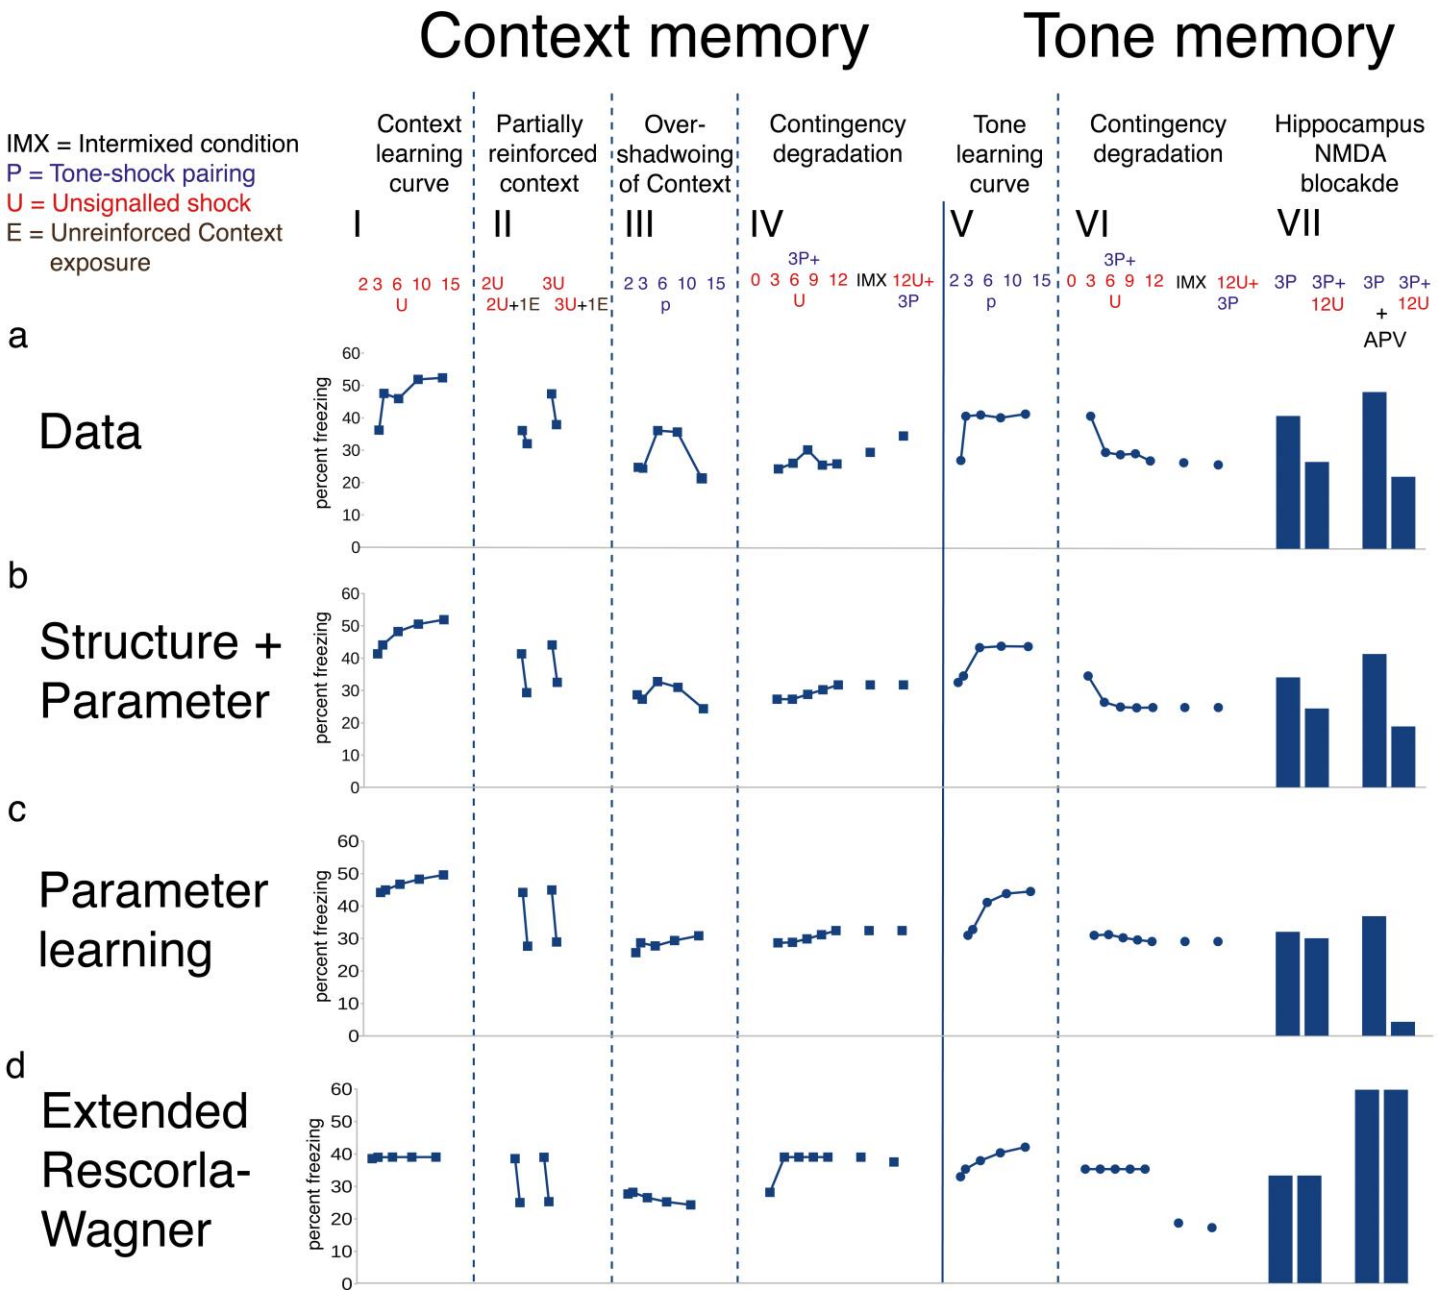

Supplementary Figure 11

Graphical illustration of the fits of some of the different models compared.

(a) Behavioral Data. (b) Bayesian model that learned both structure and parameters (SPLM). (c) Bayesian model that learned parameters using Graph 6 (from Fig. 5a) and the best Beta priors for edge parameters. (d) Van Hamme and Wasserman's extension of the Rescorla-Wagner model.

Model	MSE	s.e. of the MSE	MSE for hippoampus APV injection (% freezing squared)	Parameters
Structure Learning (SLM)	13.44	0.0054	34.03	2
Parameter Learning (PLM)	20.80	0.024	245.64	8
SLM no Background	17.28	0.057	51.31	3
PLM no Background	34.69	0.023	201.87	6
SLM Linear	21.66	0.086	192.29	2
SOCR with Background	30.89	-	440.66	7
HW-RW with Background	47.66	-	234.61	11
SOCR	32.45	-	292.53	6
HW-RW	67.06	-	848.88	8

Supplementary Table 1: Comparison of model fits. Mean squared error (MSE) for the best fit of each model, followed by the standard error of the MSE, and the error of the model in predicting the results of hippocampal interventions, with each value representing percentage freezing squared. The final column lists the number of free parameters (excluding the two scaling parameters).

Model	BIC	Adjusted R^2 Context	Adjusted R^2 Tone
SLM	88.82	0.87	0.79
SLM no Background	99.47	0.74	0.73
SLM linear	102.66	0.68	0.93
SPLM	110.54	0.79	0.71
PLM	121.69	0.53	0.65
SOCR	127.85	0.63	0.78
HW-RW	155.64	-0.32	0.17

Supplementary Table 2: Bayesian Information Criterion (BIC) and adjusted R-squared values for the different models. For BIC the number of data points was 29, and the number of parameters was the number of model parameters plus the two scaling parameters. R-squared values were adjusted by the number of model parameters plus 1 scaling parameter.

Nature Neuroscience: doi:10.1038/nn.4308

Training protocol	2*U	3*U	6*U	10*U	15*U	2*U +1*E	3*U+1*E
n=	7	7	7	5	10	6	7
Training protocol	2*P	3*P	6*P	10*P	15*P	21*P	
n=	6	25	8	8	9	7	
including with CS test only	-	3	-	-	-		
Training protocol	3*P+3*U	3*P+6*U	3*P+9*U	3*P+12*U	Intermixed	12*U+3*P	
n=	12	10	12	18	17	16	
including with CS test only	3	2	3	-	-	-	

Supplementary Table 3: Group sizes for the different conditioning protocols.

Experiment 1

	Tone		Context	
	F statistic	p value	F statistic	p value
Interaction	$F_{1,71} = 0.4$	0.75	$F_{1,71} = 0.5$	0.68
Main effect for test order	$F_{1,71} = 0.53$	0.47	$F_{1,71} = 0.16$	0.60

All experiments with varied test order

	Tone		Context	
	F statistic	p value	F statistic	p value
Interaction	$F_{1,116} = 0.35$	0.93	$F_{1,116} = 0.76$	0.62
Main effect for test order	$F_{1,116} = 0.24$	0.62	$F_{1,116} = 2.83$	0.096

Supplementary Table 4: The order of the CS and Context tests didn't significantly affect freezing scores either for the groups in Experiment 1, or across all the different conditioning protocols where the order of the testing was varied. Testing order was varied for all protocols except when only CS-US pairings were given. For those, the context test was always given first to get the best possible measure of the overshadowing effect. Analysis using Two-way ANOVA, with the levels of the first factor being the different conditioning protocols.

Nature Neuroscience: doi:10.1038/nn.4308

Group	Latency HAB (ms)	Latency LTM	Amplitude HAB (μV)	Amplitude LTM as % from HAB (μV)
Control II	14.06 ± 1.23	14.44 ± 1.41	7.57 ± 1.40	224.34 ± 37.94
Pairings First	13.83 ± 1.09	13.86 ± 1.05	9.35 ± 1.83	122.46. ± 13.25

Supplementary Table 5: There was no statistically significant effect of group or conditioning on A-LFP latencies (p values for main effects and interaction > 0.7), or between average amplitudes during habituation between groups ($p > 0.4$). However, the increase (as percentage baseline) in A-LFP amplitude amplitude following conditioning was significantly higher in the 100% contingency group Control II ($p = 0.028$).

Nature Neuroscience: doi:10.1038/nn.4308

Experiment 1-no manipulation

	Tone		Context	
	F statistic	p value	F statistic	p value
Interaction	$F_{1,75} = 1.63$	0.21	$F_{1,75} = 6.44$	0.013
Main effect for contingency	$F_{1,75} = 18.02$	0.00006	$F_{1,75} = 3.26$	0.074
Main effect for spacing of CS-US pairings	$F_{1,75} = 1.35$	0.25	$F_{1,75} = 8.22$	0.005
Simple main effect for contingency with pairings spaced	$F_{1,75} = 15.00$	0.0002	$F_{1,75} = 14.37$	0.0003
Simple main effect for contingency with pairings massed	$F_{1,75} = 4.48$	0.038	$F_{1,75} = 0.00008$	0.82
Simple main effect for spacing with 100% contingency	$F_{1,75} = 3.35$	0.071	$F_{1,75} = 10.65$	0.001
Simple main effect for spacing with 20% contingency	$F_{1,75} = 0.006$	0.94	$F_{1,75} = 0.0004$	0.63

Experiment 2-APV infusions (Fig. 2b,c)

	Tone		Context	
	F statistic	p value	F statistic	p value
Interaction	$F_{1,32} = 0.07$	0.79	$F_{1,32} = 0.38$	0.54
Main effect for contingency	$F_{1,32} = 11.98$	0.0015	$F_{1,32} = 0.55$	0.46
Main effect for drug	$F_{1,32} = 0.59$	0.44	$F_{1,32} = 9.47$	0.0043

Experiment 2-APV infusions (Fig. 2d,e)

	Tone		Context	
	t statistic	p value	t statistic	p value
	$t_{16} = 2.14$	0.048	$t_{14} = 2.31$	0.037

Experiment 3-optogenetic inactivation

	Tone		Context	
	F statistic	p value	F statistic	p value
Interaction	$F_{1,29} = 0.28$	0.6	$F_{1,29} = 4.3903$	0.045
Main effect for laser treatment	$F_{1,29} = 11.41$	0.0021	$F_{1,29} = 0.0297$	0.86
Main effect for spacing of CS-US pairings	$F_{1,29} = 0.092$	0.76	$F_{1,29} = 1.79$	0.19
Simple main effect for laser with pairings spaced	$F_{1,29} = 4.4537$	0.044	$F_{1,29} = 2.03$	0.17
Simple main effect for laser with pairings massed	$F_{1,29} = 7.0165$	0.013	$F_{1,29} = 2.37$	0.13

Experiment 4

	LFP potentiation		Tone	
	t statistic	p value	t statistic	p value
	$t_{11.1} = 2.54$	0.028	$t_{11.0} = 3.07$	0.011

Supplementary Table 6: F and p values from Two-way ANOVA comparisons for Experiments 1-3 and t statistics for Experiments 2 and 4.

Nature Neuroscience: doi:10.1038/nn.4308

Experiment 1-no manipulation

	Tone	Context
	% freezing	
CTL I	52.22 ± 4.69	3.22 ± 0.91
CTLII	40.79 ± 4.14	24.53 ± 5.50
Intermixed	26.33 ± 4.41	29.73 ± 6.06
Pairings First	26.86 ± 5.34	26.14 ± 5.94

Experiment 2-APV infusions

	Tone	Context
	% freezing	
Vehicle CTLII	40.09 ± 8.88	12.12 ± 5.87
Vehicle Pairings First	18.2 ± 5.55	8.04 ± 2.31
APV CTL II	47.17 ± 8.64	1.00 ± 0.422
APV Pairings First	21.66 ± 5.56	0.63 ± 0.29
Vehicle Intermixed		19.92 ± 9.57 (Defecation 5.423 ± 1.23)
APV CTL I	57.33 ± 9.62 (Defecation 6.88 ± 0.73)	
APV Intermixed	26.87 ± 10.46 (Defecation 3.89 ± 1.10)	0.62 ± 0.22 (2.11 ± 0.70)

Experiment 3-optogenetic inactivation

	Tone	Context
	% freezing	
Pairings First/Laser Offset	16.13 ± 3.64	45.46 ± 9.38
Pairings Firs/Laser Overlap	40.23 ± 8.04	26.34 ± 8.05
Intermixed/LaserOffset	17.51 ± 5.06	16.51 ± 6.99
Intermixed/Laser Overlap	35.1 ± 6.10	32.73 ± 8.61

Experiment 4-Electrophysiology

	Tone
	% freezing
Control II	35.45 ± 7.894
Pairings Firs	9.87 ± 2.66
Cover Stimulus	48.04 ± 7.15

Supplementary Table 7: Mean ± standard error freezing (and defecation) scores.

Nature Neuroscience: doi:10.1038/nn.4308

Structure learning model	α = 0.809	ρ = 0.646							
Structure & parameter learning model	a_1 = 28.620	b_1 = 0.239	a_2 = 5.407	b_2 = 0.173	a_3 = 28.034	b_3 = 2.416	a_4 = 23.422	b_4 = 3.108	ρ = 0.782
Parameter learning model	a_1 = 5.968	b_1 = 0.766	a_2 = 1.058	b_2 = 0.372	a_3 = 0.214	b_3 = 0.140	a_4 = 0.018	b_4 = 3.267	
Extended Comparator Hypothesis	s_1 = 0.865	s_2 = 0.862	s_3 = 0.679	k_1 = 0.05	k_2 = 0.551	k_3 = 1.00			
Extended Rescorla-Wagner model	Ta_1 = 0.980	Ta_2 = 0.000	Tb_1 = 0.468	Tb_2 = 0.005	Ca_1 = 0.972	Ca_2 = -0.158	Cb_1 = 0.857	Cb_2 = 0.004	

Supplementary Table 8: Best fit parameters for Bayesian and Associative models.

Nature Neuroscience: doi:10.1038/nn.4308

Supplementary Modeling

Stationarity

Our Bayesian models assume that trials come from a stationary process and therefore ignore trial order effects. This is a good match for our data where we only see weak ordering effects, and is a suitable strategy for learning the structure of the environment, since the presence or absence of a predictive relationship between two variables would tend to be a stable property over time. On the other hand the exact strength of a relationship might change over time, and learning the parameters that represent this strength by assuming that the underlying state of the world is dynamic could give rise to strong trial ordering effects *(52)*. An interplay between structure and parameter learning at different time points could yield an optimal strategy for exploring the environment, and give rise to weaker or stronger trial order effects, depending on the particulars of the learning environment.

Temporal representation

The most straight-forward temporal discretization of the experiments assigns one time bin to each trial, such that the time length for each temporal unit is 2 mins. This corresponds to counting each US as 1 event where the CS is either present or absent. The approximately two minutes each animal spent outside the conditioning chamber, but before being returned to the home cage was also counted as 1 event. when only the Background, but not the other stimuli were present. Since the 30s tone CS was only present for a fourth of the two minute duration of tone-shock pairing trials, the remaining fractional time intervals, as well as the one minute 'half' trial after the last US, were added up, and the integer value of this sum counted towards the number of trials with context present and CS and US absent (i.e. added to the count of the [1 1 0 0] vectors).Including these fractional counts improved all Bayesian model fits slightly, but didn't affect the final order of the fits.

Since USs always arrived at the end of the CS in our experiments, we didn't consider potential effects of ambiguity arising by the timing of the US during the CS, which is a characteristic of the totally random control procedure, and has been suggested to slow conditioning to a CS *(53)*

Alternative formulations of the Bayesian models

An alternative for the Bayesian models without a Background variable can be implemented with the three variables X_2, X_3, and X_4, s.t.

$$P(X_4 = 1|Pa(X_4)) = 1 - \left(\prod_{X_i \in Pa(X_4)} (1 - \omega_{i,4})^{x_i} \right) \quad (1)$$

Nature Neuroscience: doi:10.1038/nn.4308

or if X_4 has no parents, or when all its parents are absent, then

$$P(X_4 = 1) = \omega_{p_{X_4}} \tag{2}$$

where $\omega_{p_{X_4}}$ has a prior Beta distribution. This formulation gave a higher MSE then SLM and SPLM, but a better fit than the other models. Removing the Background variable from PLM turns it into a model that is similar to simple cue competition models, and substantially decreases the model's ability to fit the data (Supplementary Table 1).

A further alternative is to replace the noisy-OR generating function with a thresholded linear function (see SLM linear in Supplementary Table 1) , such that

$$P(X_4 = 1|Pa(X_4)) = \max(1, \sum_{i|X_i \in Pa(X_4)} \omega_{i,4} \cdot X_i) \tag{3}$$

Extending SLM to model other behavioral phenomena

We tested if SLM was compatible with previously documented conditioning phenomena involving ambiguous cue-outcome associations, by extending the model to include a further variable representing an additional environmental stimulus (such as a light). We extended the number of graph structures to include the relevant configurations (Supplementary Fig. 9) , but otherwise left the model unchanged, including the hyperparameter $P(G_1)$, and the scaling factors previously fitted to our behavioral data. SLM made accurate predictions, qualitatively matching previous data for the cover stimulus effect, overshadowing, recovery from overshadowing and blocking (Supplementary Fig. 10). The original formulation of SLM can also explain partial reinforcement and latent inhibition effects for CS-US pairings, as well as for context-US pairings.

Associative models

Representing experiments

To allow maximum flexibility for the associative models, we used two temporal discretization parameters, one for the CS duration, t_{CS} and one for when the CS was not present, t_C. We restricted the relationship between these two parameters such that the discretized representation remained faithful to the original temporal structure. In particular we imposed

$$3 * t_C \leq t_{CS} \leq 3 * (t_C + 2)$$

These discretization parameters were not included in the parameter count in the model comparison, Models were implemented as described in references *(7)* and *(8)*, and parameters fitted

Nature Neuroscience: doi:10.1038/nn.4308

by minimizing MSE with an interior-point algorithm using MATLAB's fmincon function. Associative strengths given by the model were converted to freezing scores in the same way as for the Bayesian models (multiplication by a scalar through simple linear regression, separately for the Context and the CS).

Van Hamme and Wasserman's extension of the Rescorla-Wagner model required fitting 4 learning rate parameters each for the context and the tone (8 in total) that were constrained to lie in the interval $[0, 1]$ or $[-1, 0]$, as specified by the model. λ was taken to be one, since the multiplicative scaling when converting to freezing scores ensured that the value of λ didn't influence the model's fit. We fit 6 parameters for the SOCR model: 3 learning rate parameters (s_1, s_2 and s_3), the extinction parameter k_1, and the comparator parameters k_2 and k_3. The learning rates and comparator parameters were constrained to lie in $[0, 1]$, while k_1 was required to lie in $[0.05, 1]$ to ensure a realistic model that can account for extinction/partial reinforcement effects. We ran the optimization process at least 10 times from different random starting points, separately for each permitted combination of the integer-valued temporal discretization parameters with $C_t \leq 30$. We found that minimums for each such combination were consistent across most of the the runs, with the optimization terminating at one of the few different observed values. We also extended these models by adding a 'Background' cue, to enable factorial model comparison and a better understanding of the importance of such a variable. For HW-RW this meant adding further 4 learning rate parameters, and for SOCR a single extra learning rate parameter.

Modeling neural interventions

To model hippocampal inactivations, the learning rates for the context were set to zero. Amygdala inactivation during the US were modeled by excluding trials with inactivation from the trial counts (such that they counted neither towards the reinforced, nor the unreinforced trials).

Alternative formulations

In the extended RW model, we tried replacing the linear sum in the prediction error terms

$$\lambda - (V_{Context} + V_{CS})$$

and

$$0 - (V_{Context} + V_{CS})$$

by the or function

$$\lambda - (V_{context} + V_{CS} - V_{Context} \cdot V_{CS})$$

and

$$0 - (V_{Context} + V_{CS} - V_{Context} \cdot V_{CS})$$

Nature Neuroscience: doi:10.1038/nn.4308

to see if the linear vs or formulation were important in the differences we found between models. However, this change didn't significantly improve the fit of the model.

References

52. P. Dayan, S. Kakade, P. R. Montague, Learning and selective attention. Nat. Neurosci. 3 Suppl, 1218-1223 (2000).

53. P. D. Balsam, S. Fairhurst, C. R. Gallistel, Pavlovian contingencies and temporal information. J. Exp. Psychol. Anim. Behav. Process. 32, 284-294 (2006).