

An Emergent Neural Coactivity Code for Dynamic Memory

Mohamady El-Gaby^{1,2*}, Hayley M Reeve¹, Vítor Lopes-dos-Santos¹, Natalia Campo-Urriza¹, Pavel V Perestenko¹, Alexander Morley¹, Lauren A M Strickland¹, István P Lukács³, Ole Paulsen², David Dupret^{1*}

¹ Medical Research Council Brain Network Dynamics Unit, Nuffield Department of Clinical Neurosciences, University of Oxford, Oxford, OX1 3TH, United Kingdom

² Physiological Laboratory, Department of Physiology Development and Neuroscience, University of Cambridge, Cambridge, CB2 3EG, United Kingdom

³ Department of Pharmacology, University of Oxford, Oxford, OX1 3TH, United Kingdom

*Correspondence: mohamady.el-gaby@ndcn.ox.ac.uk and david.dupret@bndu.ox.ac.uk

Abstract

Neural correlates of external variables provide potential internal codes that guide an animal's behaviour. Notably, first-order features of neural activity, such as single-neuron firing rates, have been implicated in encoding information. However, the extent to which higher-order features, such as multi-neuron coactivity, play primary roles in encoding information or secondary roles in supporting single-neuron codes remains unclear. Here we show that millisecond-timescale coactivity amongst hippocampal CA1 neurons encodes short-lived behavioural contingencies. This contingency discrimination was unrelated to the tuning of individual neurons but instead an emergent property of their short-timescale coactivity. Contingency discriminating patterns were reactivated offline after learning and their reinstatement predicted trial-by-trial memory performance. Moreover, optogenetic suppression of inputs from the upstream CA3 region selectively during learning impaired coactivity-based contingency information in CA1 and subsequent dynamic memory retrieval. These findings identify short-timescale coactivity as a primary feature of neural firing that encodes behaviourally-relevant variables and supports memory retrieval.

Main

What features of neural activity does the brain use to encode information about the external world? Ample evidence suggests that the firing rates^{1,2} and temporal tuning properties^{3,4} of individual neurons show robust correlations with external variables. These first-order features of neural activity could serve as neural codes that are read by downstream structures to subsequently guide behaviour⁵. In addition, advances in *in vivo* multi-unit recordings have allowed further appreciation for the role of neuronal population dynamics in supporting internal representations^{6–8}. The timescale at which population activity is organized may be critical. In particular, coincidental spiking at the timescale of a neuron's membrane time constant (~10–30 ms for cortical neurons⁹) effectively drives downstream receiver neurons^{5,10}, can be parsed within network oscillations that pace firing of neuronal populations⁵, and can be rapidly stabilized through spike-timing-dependent plasticity (STDP)^{11,12}. Indeed, millisecond-timescale coactivity is a hallmark of some neural codes^{13–15}. Such short-timescale coactivity organises the firing of neurons with related tuning to external variables, giving rise to robust, population-based representations that are congruent with those of their participating neurons^{14,16,17}. Moreover, millisecond timescale coactivity could also play a primary role in encoding information. That is, groups of neurons may encode a variable as a function of their joint activity regardless of whether, individually, each neuron is tuned to this variable. While this type of emergent, coactivity-based coding has been described for physically well-defined

variables such as specific sensory inputs and actions^{18–20}, its possible cognitive function has not been explored.

Given the potential for rapid stabilization and retrieval of neural codes based on millisecond-timescale coactivity, such codes may support behavioural performance when animals must rapidly learn and flexibly retrieve salient information – a process we refer to here as “*dynamic memory*”. Converging evidence suggests a prominent role of the hippocampus for such rapid and flexible learning^{21–23}, supporting models that frame the hippocampus as a fast learning system²⁴. Moreover, neural activity in the hippocampus is organized into temporally precise coactivity patterns^{15,22,25}. We therefore hypothesized that millisecond-timescale coactivity patterns in the hippocampus serve a primary role in encoding behaviourally-relevant information supporting dynamic memory. To test this hypothesis, we developed a one-day, two-contingency discrimination task that we combined with multi-unit recording of hippocampal CA1 neurons and causal optogenetic manipulation of intra-hippocampal synapses. Our findings demonstrate a role for emergent coactivity-based representations in encoding contingency information and supporting dynamic memory retrieval.

Results

Mice learn and dynamically retrieve two new behavioural contingencies every day

We first established a one-day behavioural paradigm that recruits dynamic memory (Fig. 1). Mice were initially pre-trained to collect a transiently available (5-second) drop of sucrose from a liquid dispenser after the presentation of an auditory cue (pre-training phase 1; Extended Data Fig. 1). Subsequently, animals experienced a novel learning enclosure every day, which was defined by a new spatial topology, two new sets of wall-mounted LED-displays and two newly positioned dispensers (pre-training phase 2; Extended Data Fig. 1). In this learning enclosure, animals encountered the following rule: immediately after tone presentation, one dispenser delivers a drop of sucrose solution whereas the other simultaneously delivers a bitter (quinine) solution; both drops are transiently available. Importantly, the dispenser-solution pairing was contingent on which of the two sets of LED cues is illuminated concurrently with the tone (Fig. 1a,b). When animals reached an average of 80% performance in this pre-training phase, we then started the training phase, which included three stages every day (Fig. 1c). In the first stage, animals explored the new learning enclosure in two sessions, each with one of the two LED sets continuously illuminated but without tone presentation or drop delivery, as well as another exploration session in a control (task-unrelated) enclosure (“*Exploration*” stage; Fig. 1c and Extended Data Fig. 2a). In the second stage, animals learned to associate each LED set with the tone-triggered delivery of a selective drop outcome at each dispenser over four sessions alternating between active LEDs (“*Learning*” stage; Fig. 1c). We refer to these associations as LED-defined behavioural “contingencies” (*X* and *Y*; Fig. 1a), with animals learning two new contingencies every day (extended Data Fig. 2b). During learning, mice rapidly developed a successful approach response to the correct (sucrose) dispenser over the incorrect (quinine) dispenser in each contingency (Fig. 1d,e and Extended Data Fig. 2b). In the final stage conducted at the end of each day (one hour after the end of learning), memory for the newly-learned contingencies was tested in a probe session where the tone was presented without drop delivery while pseudo-randomly switching between the two LED sets (“*Probe*” stage; Fig. 1c). In these probe trials, mice continued to identify the correct dispenser (Fig. 1f and Extended Data Fig. 2c). Memory performance on a given day was unrelated to that on the previous day (Extended Data Fig. 3a) and held when averaging across all probe performances for each individual mouse (Extended Data Fig. 3b). Furthermore, while animals made more mistakes on the first probe trial following a switch in LEDs compared to the other trials (Extended Data Fig. 3c), there was no deterioration of performance as the probe session

progressed (Extended Data Fig. 3d). Thus, mice successfully learned to discriminate two new behavioural contingencies each day and flexibly retrieved a memory of this discrimination, providing a paradigm to study the neural substrates of dynamic memory.

Emergent millisecond-timescale coactivity discrimination of behavioural contingencies

To investigate whether an emergent coactivity code develops in our task, we monitored hippocampal CA1 neuronal ensembles during training days. We first trained a Bayesian classifier to decode the prevailing contingency on a trial-by-trial basis from both average firing rates of principal neurons and short-timescale (25-ms) pairwise temporal correlations between neuronal spike trains. Shuffling temporal correlations across trials, while preserving trial-by-trial average firing rates, significantly impaired decoding of ongoing contingency (Fig. 2a; Extended data Fig. 4a). Moreover, contingency information in temporal correlations alone was drastically impaired when shifting spikes to destroy short-timescale coactivity while maintaining correlations due to slow fluctuations of population firing rate in each trial (Fig. 2b). Short-timescale correlations also had significant explained variance for task contingencies (Fig. 2c). These results indicated the presence of contingency-related information in short-timescale coactivity beyond the information in single-neuron firing rates.

To investigate the task relevance of contingency-related coactivity, we isolated coactivity patterns nested within 25-ms time-windows²⁶ separately in each contingency within the learning enclosure. We represented each pattern by a weight vector containing the contribution of each neuron to the coactivity underpinning that pattern (Fig. 2d). These coactivity patterns differed from those extracted in the control enclosure (Fig. 2e,f), showing their spatial context-selective expression. In addition, some learning enclosure patterns discriminated the two contingencies, being selective to either X or Y (Fig. 2e,f; Extended data Fig. 4b; orange). To investigate the functional significance of such patterns, we compared them to a matched group of learning enclosure patterns with high between-contingency similarity (Fig. 2e,f Extended Data Fig. 4b; blue). We refer to these as contingency-discriminating and contingency-invariant coactivity patterns, respectively. Neurons that contributed the most to a given pattern are henceforth referred to as “members” of that pattern (see Methods).

We confirmed that members of the same contingency discriminating, but not invariant, pattern were more correlated in one contingency than the other (Fig. 3a). Importantly however, members of contingency-discriminating patterns were not individually contingency selective (Fig. 3b; Extended Data Fig. 4c), regardless of the membership threshold used (Extended Data Fig. 4d,e); and are hence separable from previously reported contextually modulated neurons^{27–29}. Moreover, such coactivity-based contingency discrimination was not explained by differences in temporal firing properties of individual member neurons between contingencies (Fig. 3c-e). Furthermore, contingency-discriminating pattern members were not tuned to goal locations (Fig. 3f,g) and hence did not report trajectories to goal locations³⁰. We also noted no differences in the participation of neurons along the transverse axis of the CA1 to contingency-discriminating and invariant patterns (52.1% and 48.2% of pattern member neurons found in proximal and distal CA1, respectively; Fisher’s exact test: odds ratio=1.17; $P=0.55$), no segregation by hemisphere (Extended Data Fig. 4f) and no differences in the participation of neurons from the deep or superficial CA1 pyramidal sublayer to contingency-discriminating compared to contingency-invariant patterns (39.9% and 28.4% of pattern member neurons, respectively; Fisher’s exact test: odds ratio=1.44; $P=0.20$). However, we observed a trend towards contingency-discriminating coactivity pattern members firing at earlier theta phases compared to members of contingency invariant patterns (Extended data Fig. 4g). Overall, these

findings identify an emergent, short-timescale neural coactivity-based discrimination of behavioural contingencies in the hippocampal CA1.

We next asked whether contingency-discriminating coactivity patterns relate to contingency learning. When we tracked the strength of each pattern (Fig. 2d) over time, we found that contingency-invariant patterns began increasing in strength during the initial exploration of the new learning enclosure on each day, before animals experienced task contingencies; their strength further increased and subsequently plateaued during learning (Fig. 4a; Extended data Fig. 5a,b). Conversely, contingency-discriminating coactivity was more stable during exploration but markedly increased during contingency learning (Fig. 4a; Extended data Fig. 5a,b). Pattern strengthening during learning reflected increased temporal correlations between members' activity rather than changes in their average firing rates (Extended data Fig. 5c,d). Furthermore, the cofiring of contingency-invariant pattern members increased during sharp-wave/ripples in post- compared to pre-exploration sleep, and increased again in the sleep session after learning, while contingency-discriminating pattern members only increased their sharp-wave/ripple cofiring after contingency learning (Fig. 4b,c). Thus, the distinction between contingency-invariant and contingency-discriminating pattern members was not equivalent to the difference between the previously described "rigid" and "plastic" cells³¹. Moreover, our findings did not simply reflect representations of rewarded/aversive locations²⁷ since pattern strength was calculated outside dispenser locations, nor were they a simple reflection of the animal's differential behaviour across the two contingencies (i.e. heading towards a given dispenser; Extended Data Fig 5e). Importantly, the reinstatement of contingency-discriminating patterns during memory retrieval predicted trial-by-trial performance; these patterns were stronger before correct, compared to incorrect, behavioural responses to tone presentation (Fig. 4d; Extended Data Fig. 5f). This contingency-selective and performance-related reinstatement of CA1 coactivity was not associated with a firing rate bias of member neurons nor animal running speed (Extended Data Fig. 5g,h), and was notably absent when animals performed at chance (i.e., when correct and incorrect trials were equivalent; Extended Data Fig. 5i). In contrast, the strength of contingency-invariant patterns was not related to trial-by-trial memory performance (Fig. 4d). Moreover, while there was contingency-related information in longer (1-second) timescale coactivity (Extended Data Fig. 5j), the reinstatement of second-timescale contingency-discriminating coactivity during the probe session did not predict memory performance (Extended Data Fig. 5k). These findings show that CA1 neuronal spiking is gradually organised during learning to form millisecond-timescale coactivity patterns representing newly-learned contingencies, which are subsequently reinstated on a trial-by-trial basis during dynamic memory retrieval.

Distinct spatial tuning of contingency-discriminating and invariant coactivity patterns

During exploration of a novel environment, CA1 neurons with overlapping place fields can form spatially tuned coactivity patterns^{17,32}. To investigate the spatial tuning of coactivity patterns during contingency learning, we computed for each detected pattern the spatial map corresponding to the time-course of its activation strength, as well as the individual firing rate maps of each of its member neurons. Contingency-discriminating coactivity was markedly less spatially coherent than contingency-invariant coactivity (Fig. 5a-d and Extended Data Fig. 6). This was concomitant with the less spatially coherent firing of contingency-discriminating pattern members relative to their contingency-invariant counterparts (Extended Data Fig. 7a,b), with contingency-discriminating members also exhibiting a trend towards more place fields within a given session compared to contingency-invariant members (Extended Data Fig. 7c). Moreover, while members of a given contingency-invariant pattern had overlapping firing fields, members of a given contingency-discriminating pattern were markedly less spatially

correlated (Fig. 5a,b,e; Extended Data Fig. 6). This weaker spatial overlap was observed regardless of the membership threshold used (Extended Data Fig. 7d,e) and was robust to differences in temporal correlation amongst members' spike trains (Extended Data Fig. 7f). In addition, while contingency-invariant coactivity was spatially biased towards the place fields of their member neurons, this bias was significantly weaker for contingency-discriminating patterns (Extended Data Fig. 7g,h). This finding was corroborated by a separate analysis showing lower place field similarity of neuron pairs with high explained variance for contingency compared to neuron pairs with low explained variance (Extended Data Fig. 7i). Finally, we found no evidence that contingency discrimination by a given coactivity pattern reflects contingency-gated spatial remapping of its member neurons. In fact, the spatial map of an individual member of either pattern type was on average more similar across sessions of different contingencies than sessions of the same contingency (Fig. 5f,g; Extended Data Fig. 7j), even when matching the spatial coherence of contingency-discriminating pattern members to that of contingency-invariant counterparts (Extended Data Fig. 7k). Moreover, members of the same contingency-discriminating pattern were as spatially correlated with each other across sessions of their preferred contingency as they were across sessions of opposite contingency (Extended Data Fig. 7l). Overall, these findings show that contingency-invariant coactivity provides robust place representations by binding spatially congruent neurons. In contrast, contingency-discriminating patterns stitch together neurons irrespective of their spatially correlated activity, giving rise to spatially discontinuous coactivity consistent with a specialization in representing ongoing behavioural contingency.

CA3^L→CA1 inputs are necessary for contingency discriminating coactivity and dynamic memory retrieval

Finally, to address the functional role of contingency-discriminating coactivity, we sought to identify and manipulate a neural pathway necessary for their formation. CA1 coactivity could rely on synaptic inputs from the recurrently-connected upstream hippocampal CA3 area^{33,34}, and recent work suggests a critical mnemonic role of left CA3 (CA3^L) inputs to CA1^{35,36}. Accordingly, we transduced CA3^L pyramidal neurons of Grik4-Cre mice with the yellow light-driven proton pump Archaeorhodopsin-3.0 (Fig. 6a,b); bilateral implantation of tetrodes and optic fibres further allowed simultaneous monitoring of, and light delivery to, CA1 ensembles. Light delivery targeting CA3^L axons in CA1 during learning markedly reduced the power of theta-nested slow-gamma, but not mid-gamma, oscillations in CA1 (Fig. 6c; Extended Data Fig. 8a-c), consistent with the suggestion that CA1 slow-gamma oscillations report incoming CA3 inputs^{37,38}. While suppressing CA3^L→CA1 inputs preserved both the organisation of CA1 neurons into coactivity patterns during learning, and the reinstatement of such patterns during memory retrieval (Extended Data Fig. 8d,e), this intervention altered the information content of CA1 coactivity. Firstly, the distribution of between-contingency pattern similarity and pattern strength ratio was shifted towards contingency-invariance (Fig. 6d,e; Extended Data Fig. 8f, Extended Data Fig. 9). Secondly, this manipulation reduced the explained variance for contingencies in short-timescale pairwise correlations (Extended Data Fig. 8g). Thirdly, Bayesian decoding of contingency using such short-timescale coactivity was markedly impaired with CA3^L→CA1 input suppression (Extended Data Fig. 8h). At the behavioural level, suppressing CA3^L→CA1 inputs selectively during learning had no effect on ongoing performance (Extended Data Fig. 8i) but reduced memory performance to chance levels in the subsequent probe test one hour after, during which there was no input suppression (Fig. 6f). This latent memory impairment was seen when mice had to flexibly retrieve two contingencies (Fig. 6f; Extended Data Fig. 8j-l), but not when retrieving only one contingency (Extended Data Fig. 8m). Moreover, flexible memory retrieval of the two contingencies was preserved after suppressing right CA3 inputs to CA1 (Extended Data Fig. 8n-u). Together, these findings

show that short-timescale CA1 coactivity-based contingency-information necessitates CA3^L inputs and is required for dynamic retrieval of two-contingency memory.

Discussion

In this study, we report a coactivity-based hippocampal code for dynamic memory retrieval of short-lived behavioural contingencies. Encoding information as an emergent property of coactivity amongst multiple neurons (Extended Data Figure 10a) allows effective discrimination of newly encountered contingencies every day, without committing individual neurons to represent such short-lived cognitive variables. The emergent nature of this code points to short-timescale coactivity as a primary feature of neural activity that is used to encode information and guide cognition, rather than only playing secondary roles, such as organizing or stabilizing single neuron rate-based codes. In particular, our findings show that millisecond-timescale coactivity is highly suited for mnemonic processing of short-lived information: it is rapidly formed and readily reinstated to support flexible memory retrieval. Millisecond-timescale neural coactivity may preferentially exhibit spike-timing dependent plasticity (STDP)^{11,12,39} to rapidly stabilize the code in memory. In contrast, while second-timescale coactivity contained contingency information in our task, its reinstatement during dynamic memory retrieval was not predictive of trial-by-trial performance. Second-timescale coactivity may exhibit slower plasticity, and hence be more suited for stable representation of long-lived contingencies.

Our findings also provide new insights into the role of correlated neural activity in guiding contextual behaviour. Spatial remapping, where patterns of spatial correlations between hippocampal principal cells disambiguate distinct spatial contexts, has been proposed as a neural basis for contextual learning²³. In this study we observe that contingency discriminating coactivity is not a reflection of spatial remapping. Instead, our findings are consistent with the view that spatial remapping may be a specific instance of a more general phenomenon of “temporal remapping”, in which the short-timescale temporal correlation structure of neurons differs across distinct contexts²³. Indeed, in tasks where animals must disambiguate different spatial reference frames, millisecond timescale coactivity is a robust correlate of moment-by-moment behavioural discrimination of different contexts, both in networks that show spatial remapping²² and those that do not²⁵. This is also consistent with a reader-centric view of neural codes⁵, since downstream reader/actuator neurons can detect temporal, but not spatial, correlations amongst their input neurons. Notably, one prediction from this coding scheme is that downstream receiver neurons “read” the incoming information, represented as an emergent property of the collective activity of multiple neurons, by disambiguating the relevant patterns of millisecond input coincidence from the myriad of other inputs they receive⁵. Such decoding may be implemented by a “reader” network⁴⁰, or even a single “reader” neuron⁴¹ (see also Extended Data Fig. 10b).

Our findings further establish that, to “write” a millisecond coactivity code for learned contingencies in memory, CA3^L→CA1 inputs are necessary. Whether this is related to lateralisation in information content, processing and/or plasticity^{35,42} of CA3→CA1 inputs remains to be investigated. Nevertheless, we show that distinct types of coactivity patterns show qualitatively distinct functional plasticity. While contingency-invariant patterns develop during exploration and are reactivated in sharp-wave/ripples during offline (sleep/rest) periods following spatial exploration, contingency-discriminating patterns show robust increases in strength during learning and are reactivated offline in sharp-wave/ripples after contingency learning. Thus, both invariant and discriminating patterns show a signature of previously described “plastic” cells³¹, albeit in different behavioural stages. This is consistent with a

division of labour amongst hippocampal coactivity patterns, with contingency-invariant patterns reflecting unsupervised learning about the spatial structure of the environment, and contingency-discriminating patterns supporting flexible memory-guided behaviour. Altogether, our findings open new perspectives for future empirical and modelling studies to elucidate mechanisms for writing and reading coactivity-based information and to relate coding schemes across multiple timescales of population activity.

How can the code be written and read?

The hippocampus is embedded in a wider network of cortical and subcortical structures that may mediate or modulate the formation of the emergent coactivity code we describe here (writing) and its subsequent use by downstream neurons (reading) to select contingency-specific behaviour. Below we outline hypotheses about possible mechanisms for both writing and reading processes.

We show a necessity of CA3^L-CA1 inputs during learning for the expression of an emergent coactivity code for short-lived behavioural contingencies, which opens a window into the generative mechanisms at play. Left hemisphere originating CA3 inputs in mice exhibit more robust long-term plasticity^{35,42}, including STDP⁴², and are preferentially involved in long-term memory compared to right CA3 inputs^{35,36}. Such a difference in plasticity may provide part of the mechanism by which contingency-discriminating patterns are strengthened during learning. The dynamic memory task we assess here necessitates the rapid acquisition and stabilization of contingency information (within 30 trials in each contingency across ~3 hours) as well as its rapid and flexible retrieval in the memory probe test (1 hour after learning, with frequent, pseudorandom switches in contingency). Such rapid mnemonic processing may be preferentially coded by short-timescale coactivity, since STDP mechanisms are more likely to rapidly stabilize neuronal co-firing within short (10s of ms) compared to that within longer (1s) windows^{11,12,39}. Indeed, we show that short (25ms) but not longer (1s) timescale coactivity is reinstated to predict performance (Fig. 4d; Extended data Fig. 5k). However, it is also possible that other plasticity mechanisms are at play (including non-synaptic ones). While STDP might stabilize millisecond timescale coactivity patterns, what processes generate such contingency-discriminating coactivity in the first place? Neurons in the dentate gyrus, two synapses upstream of the CA1, have been implicated in pattern separation processes that may be necessary for contextual behaviour⁴³ and can do so through differences in millisecond-timescale coactivity²⁵. Moreover, there is evidence for a left dominance in the expression of the activity marker cFos in the dentate gyrus during novel object exploration⁴⁴. Importantly, while CA3 neurons in one hemisphere send commissural projections to the contralateral hemisphere³³, the two hemispheres seem to retain functional differences in their projections to CA1^{35,42} which we target directly. Such lateralization could in part result from a developmental lateralization of factors involved in activity and plasticity⁴⁵, which may be robust to any potential synchronizing effects of commissural projections. It is plausible therefore that a combination of lateralized processing of contextual information, starting as early as the pattern separation circuits of the dentate gyrus, and lateralized plasticity at CA3-CA1 synapses contribute to the formation and stabilization of emergent contingency discriminating patterns in the CA1, respectively. Given that contingency discriminating patterns emerge during learning rather than spatial exploration (Fig. 4a), their formation is not simply a reflection of sensory differences between the two LED displays (which are also distinct during exploration sessions X0 and Y0) but instead relates to the different reward contingencies the animal must learn to discriminate. Indeed, recent evidence suggests that neural discrimination of distinct spatial contexts in the CA1, but not dentate gyrus, is related to behavioural discrimination of these contexts⁴⁶, suggesting an additional gating of behaviourally relevant environmental

differences between the dentate gyrus and CA1. How such behavioural contingency information is conveyed to the CA1 is currently unclear and may involve direct inputs from the prefrontal cortex⁴⁷. The finding described here, that an emergent coactivity code in hippocampal CA1 is necessary for dynamic retrieval of contingency-discrimination, will motivate subsequent empirical and modelling studies that elucidate the cross-circuit interactions involved in generating such functional coactivity.

How is the hippocampal coactivity code for contingencies decoded by downstream neurons in executive and motor areas to elicit appropriate behaviours in each contingency? Cortical neurons have membrane time constants in the range of 10-30ms⁹ meaning that convergent input from neurons coactive at the 25-ms timescale we investigate here can exhibit effective temporal summation in the downstream (“reader”) neuron’s dendrites and contribute to its spiking. Moreover, coincident synaptic activation within this time-window is consistent with the initiation of active, voltage-gated conductances in dendrites, which allows their supra-linear summation⁴⁸. This may also serve as a mechanism for disambiguating different patterns of coactivity by a single reader neuron, where inputs that are preferentially spatially clustered on individual dendrites will be more likely to elicit such non-linearities than more dispersed inputs⁴⁹, even when the mean synaptic weights of such inputs are indistinguishable (Extended Data Figure 10b). This would allow selective reading of emergent coactivity, as readers would not disambiguate the firing of individual members of contingency discriminating coactivity patterns, only their synchronous activity. Other single-neuron and network-based coactivity reading mechanisms have also been suggested^{40,41}. For all of these cases, the fast (10s of ms) nature of this code should allow rapid processing of contingency information supporting rapid behavioural responses in dynamically changing environments. These outlined candidate mechanisms by which emergent coactivity codes could be read by downstream circuits may be tested in future *ex vivo*, *in vivo* and *in silico* studies.

Acknowledgements We thank David Bannerman and Stephen McHugh for helpful advice when designing the behavioural protocol; Jane Westcott and Ben Micklem for technical support; William Podlaski and Tim Vogels for useful discussions on neural coding models; Andrew J. Quinn for developing the EMD toolbox; Stéphanie Trouche, Helen Barron and all the members of the Dupret laboratory for useful discussions. This work was supported by the Biotechnology and Biological Sciences Research Council UK (Awards BB/N00597X/1 to D.D. and BB/N006836/1 to O.P.) and the Medical Research Council UK (Programmes MC_UU_12024/3 and MC_UU_00003/4 to D.D.).

Author contributions M.E., O.P. and D.D. conceptualised the study; M.E. and D.D. designed the experiments. M.E., H.M.R., P.V.P., N.C-U., L.A.M.S. and I.P.L. performed the experiments and acquired the data; M.E., V.L. and A.M. pre-processed the data; M.E., and V.L. analysed the data; M.E. and D.D. wrote the manuscript; D.D. supervised the project. All authors discussed the results and commented on the manuscript.

Competing interests: The authors declare no competing interests.

References

1. Adrian, E. D. The impulses produced by sensory nerve endings: Part I. *J. Physiol.* **61**, 49–72 (1926).

2. O'Keefe, J. & Dostrovsky, J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* **34**, 171–175 (1971).
3. Optican, L. M. & Richmond, B. J. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *J. Neurophysiol.* **57**, 162–178 (1987).
4. O'Keefe, J. & Recce, M. L. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* **3**, 317–330 (1993).
5. Buzsáki, G. Neural Syntax: Cell Assemblies, Synapsembles, and Readers. *Neuron* **68**, 362–385 (2010).
6. Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).
7. Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).
8. Yuste, R. From the neuron doctrine to neural networks. *Nat. Rev. Neurosci.* **16**, 487–497 (2015).
9. Koch, C., Rapp, M. & Segev, I. A brief history of time (constants). *Cereb. Cortex N. Y. N* **1991** **6**, 93–101 (1996).
10. Song, S., Miller, K. D. & Abbott, L. F. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci.* **3**, 919–926 (2000).
11. Zhang, Z., Russell, L. E., Packer, A. M., Gauld, O. M. & Häusser, M. Closed-loop all-optical manipulation of neural circuits in vivo. *Nat. Methods* **15**, 1037–1040 (2018).
12. Cassenaer, S. & Laurent, G. Hebbian STDP in mushroom bodies facilitates the synchronous flow of olfactory information in locusts. *Nature* **448**, 709–713 (2007).
13. Stopfer, M., Bhagavan, S., Smith, B. H. & Laurent, G. Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature* **390**, 70–74 (1997).
14. Robbe, D. *et al.* Cannabinoids reveal importance of spike timing coordination in hippocampal function. *Nat. Neurosci.* **9**, 1526–1533 (2006).
15. Harris, K. D., Csicsvari, J., Hirase, H., Dragoi, G. & Buzsáki, G. Organization of cell assemblies in the hippocampus. *Nature* **424**, 552–556 (2003).
16. Wilson, M. & McNaughton, B. Dynamics of the hippocampal ensemble code for space. *Science* **261**, 1055–1058 (1993).
17. van de Ven, G. M., Trouche, S., McNamara, C. G., Allen, K. & Dupret, D. Hippocampal Offline Reactivation Consolidates Recently Formed Cell Assembly Patterns during Sharp Wave-Ripples. *Neuron* **92**, 968–974 (2016).
18. deCharms, R. C. & Merzenich, M. M. Primary cortical representation of sounds by the coordination of action-potential timing. *Nature* **381**, 610–613 (1996).
19. Friedrich, R. W., Habermann, C. J. & Laurent, G. Multiplexing using synchrony in the zebrafish olfactory bulb. *Nat. Neurosci.* **7**, 862–871 (2004).
20. Dejean, C. *et al.* Prefrontal neuronal assemblies temporally control fear behaviour. *Nature* **535**, 420–424 (2016).
21. Nakazawa, K. *et al.* Hippocampal CA3 NMDA receptors are crucial for memory acquisition of one-time experience. *Neuron* **38**, 305–315 (2003).
22. Kelemen, E. & Fenton, A. A. Dynamic grouping of hippocampal neural activity during cognitive control of two spatial frames. *PLoS Biol.* **8**, e1000403 (2010).
23. Kubie, J. L., Levy, E. R. J. & Fenton, A. A. Is hippocampal remapping the physiological basis for context? *Hippocampus* **30**, 851–864 (2020).
24. McClelland, J. L., McNaughton, B. L. & O'Reilly, R. C. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* **102**, 419–457 (1995).

25. van Dijk, M. T. & Fenton, A. A. On How the Dentate Gyrus Contributes to Memory Discrimination. *Neuron* **98**, 832–845.e5 (2018).
26. Lopes-dos-Santos, V., Ribeiro, S. & Tort, A. B. L. Detecting cell assemblies in large neuronal populations. *J. Neurosci. Methods* **220**, 149–166 (2013).
27. McKenzie, S. *et al.* Hippocampal Representation of Related and Opposing Memories Develop within Distinct, Hierarchically Organized Neural Schemas. *Neuron* **83**, 202–215 (2014).
28. Leutgeb, S. *et al.* Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* **309**, 619–623 (2005).
29. Frank, L. M., Brown, E. N. & Wilson, M. Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron* **27**, 169–178 (2000).
30. Pfeiffer, B. E. & Foster, D. J. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* **497**, 74–79 (2013).
31. Grosmark, A. D. & Buzsáki, G. Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. *Science* **351**, 1440–1443 (2016).
32. O'Neill, J., Senior, T. J., Allen, K., Huxter, J. R. & Csicsvari, J. Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nat. Neurosci.* **11**, 209–215 (2008).
33. Amaral, D. & Lavenex, P. Hippocampal Neuroanatomy. in *The Hippocampus Book* (eds. Andersen, P., Morris, R., Amaral, D., Bliss, T. & O'Keefe, J.) 37–114 (Oxford University Press, 2006).
34. Middleton, S. J. & McHugh, T. J. Silencing CA3 disrupts temporal coding in the CA1 ensemble. *Nat. Neurosci.* **19**, 945–951 (2016).
35. Shipton, O. A. *et al.* Left–right dissociation of hippocampal memory processes in mice. *Proc. Natl. Acad. Sci.* **111**, 15238–15243 (2014).
36. El-Gaby, M. *et al.* Archærhodopsin Selectively and Reversibly Silences Synaptic Transmission through Altered pH. *Cell Rep.* **16**, 2259–2268 (2016).
37. Bragin, A. *et al.* Gamma (40–100 Hz) oscillation in the hippocampus of the behaving rat. *J. Neurosci. Off. J. Soc. Neurosci.* **15**, 47–60 (1995).
38. Colgin, L. L. *et al.* Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature* **462**, 353–357 (2009).
39. Bi, G. Q. & Poo, M. M. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci. Off. J. Soc. Neurosci.* **18**, 10464–10472 (1998).
40. Zenke, F. & Ganguli, S. SuperSpike: Supervised Learning in Multilayer Spiking Neural Networks. *Neural Comput.* **30**, 1514–1541 (2018).
41. Gütig, R. & Sompolinsky, H. The tempotron: a neuron that learns spike timing-based decisions. *Nat. Neurosci.* **9**, 420–428 (2006).
42. Kohl, M. M. *et al.* Hemisphere-specific optogenetic stimulation reveals left-right asymmetry of hippocampal plasticity. *Nat. Neurosci.* **14**, 1413–1415 (2011).
43. Leutgeb, J. K., Leutgeb, S., Moser, M.-B. & Moser, E. I. Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science* **315**, 961–966 (2007).
44. Jordan, J. T., Shanley, M. R. & Pytte, C. L. Behavioral state-dependent lateralization of dorsal dentate gyrus c-Fos expression in mice. *Neuronal Signal.* **3**, (2019).
45. Kawahara, A. *et al.* Neuronal major histocompatibility complex class I molecules are implicated in the generation of asymmetries in hippocampal circuitry. *J. Physiol.* **591**, 4777–4791 (2013).
46. Allegra, M., Posani, L., Gómez-Ocádiz, R. & Schmidt-Hieber, C. Differential Relation between Neuronal and Behavioral Discrimination during Hippocampal Memory Encoding. *Neuron* **108**, 1103–1112.e6 (2020).

47. Vertes, R. P., Hoover, W. B., Szigeti-Buck, K. & Leranth, C. NUCLEUS REUNIENS OF THE MIDLINE THALAMUS: LINK BETWEEN THE MEDIAL PREFRONTAL CORTEX AND THE HIPPOCAMPUS. *Brain Res. Bull.* **71**, 601–609 (2007).
48. Stuart, G. J. & Spruston, N. Dendritic integration: 60 years of progress. *Nat. Neurosci.* **18**, 1713–1721 (2015).
49. Kastellakis, G., Cai, D. J., Mednick, S. C., Silva, A. J. & Poirazi, P. Synaptic clustering within dendrites: An emerging theory of memory formation. *Prog. Neurobiol.* **126**, 19–35 (2015).
50. Ho, J., Tumkaya, T., Aryal, S., Choi, H. & Claridge-Chang, A. Moving beyond P values: data analysis with estimation graphics. *Nat. Methods* **16**, 565–566 (2019).

Figures

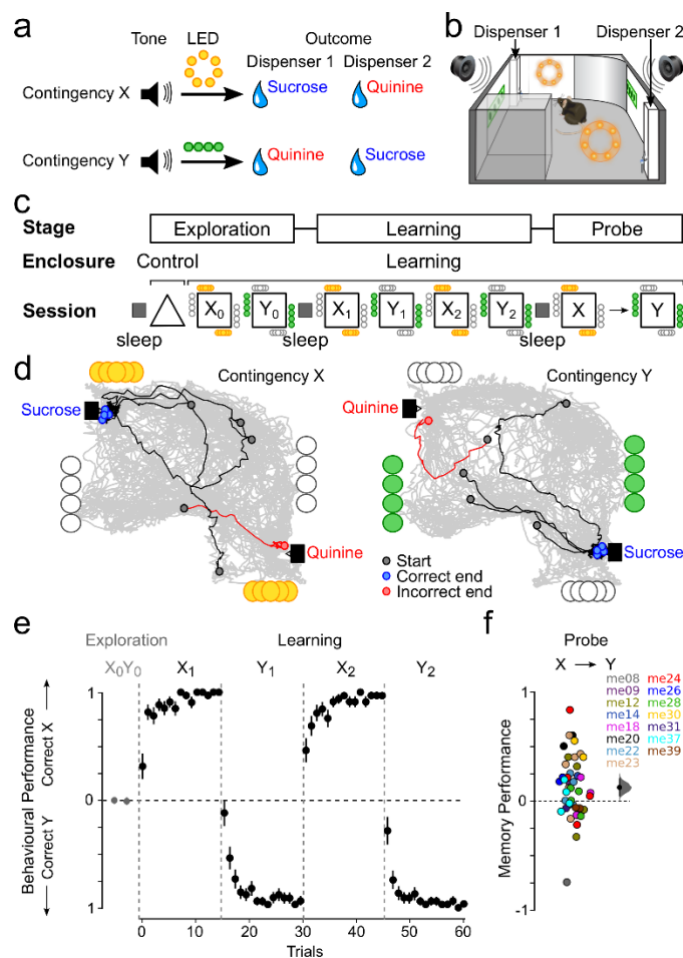


Figure 1. Mice rapidly acquire and flexibly retrieve a one-day-two-contingency memory

a) The two-learning-contingency layout. A tone signalled that both outcome dispensers deliver a liquid drop, the identity of which (sucrose versus quinine) depended on the active set of LEDs. **b)** Schematic of an example learning enclosure. **c)** The three-stage task structure. Tone-defined trials occurred in learning and probe sessions, with drop outcomes only delivered during learning. Sleep/rest sessions were recorded before and after exploration and after learning. **d)** Example animal paths during trials in contingency X and contingency Y (correct paths: black; incorrect paths: red), overlaid on the overall animal path (grey) for one learning session. Black and blue/red circles represent path starting and correct/incorrect ending points, respectively. **e)** Behavioural performance for contingency-defined correct dispensers during exploration and learning ($n=71$ days from 15 mice). For each learning trial, a score of 1 indicates that mice successfully identified the correct (sucrose-delivering) dispenser while a score of -1 indicates that mice opted for the incorrect (quinine-delivering) dispenser; behavioural performance is reported as the mean score per trial number across days across mice (see Methods). Since the correct dispenser in one contingency (e.g. X) was the incorrect dispenser in the other (e.g. Y), behavioural performance is shown with respect to the current contingency, with the y-axis ranging from 1 (correct Y) to 0 (chance) to 1 (correct X). **f)** Behavioural performance during memory probe test showing that animals identified the correct dispenser for a given contingency (mean performance: 0.12 ± 0.04). Left: color-coded raw data points represent individual mice. Right: effect size for the difference against zero, computed from 1,000 bootstrapped resamples, with black-dot representing mean; black-ticks

representing 95% confidence interval; and filled-curve representing resampled mean difference distribution.

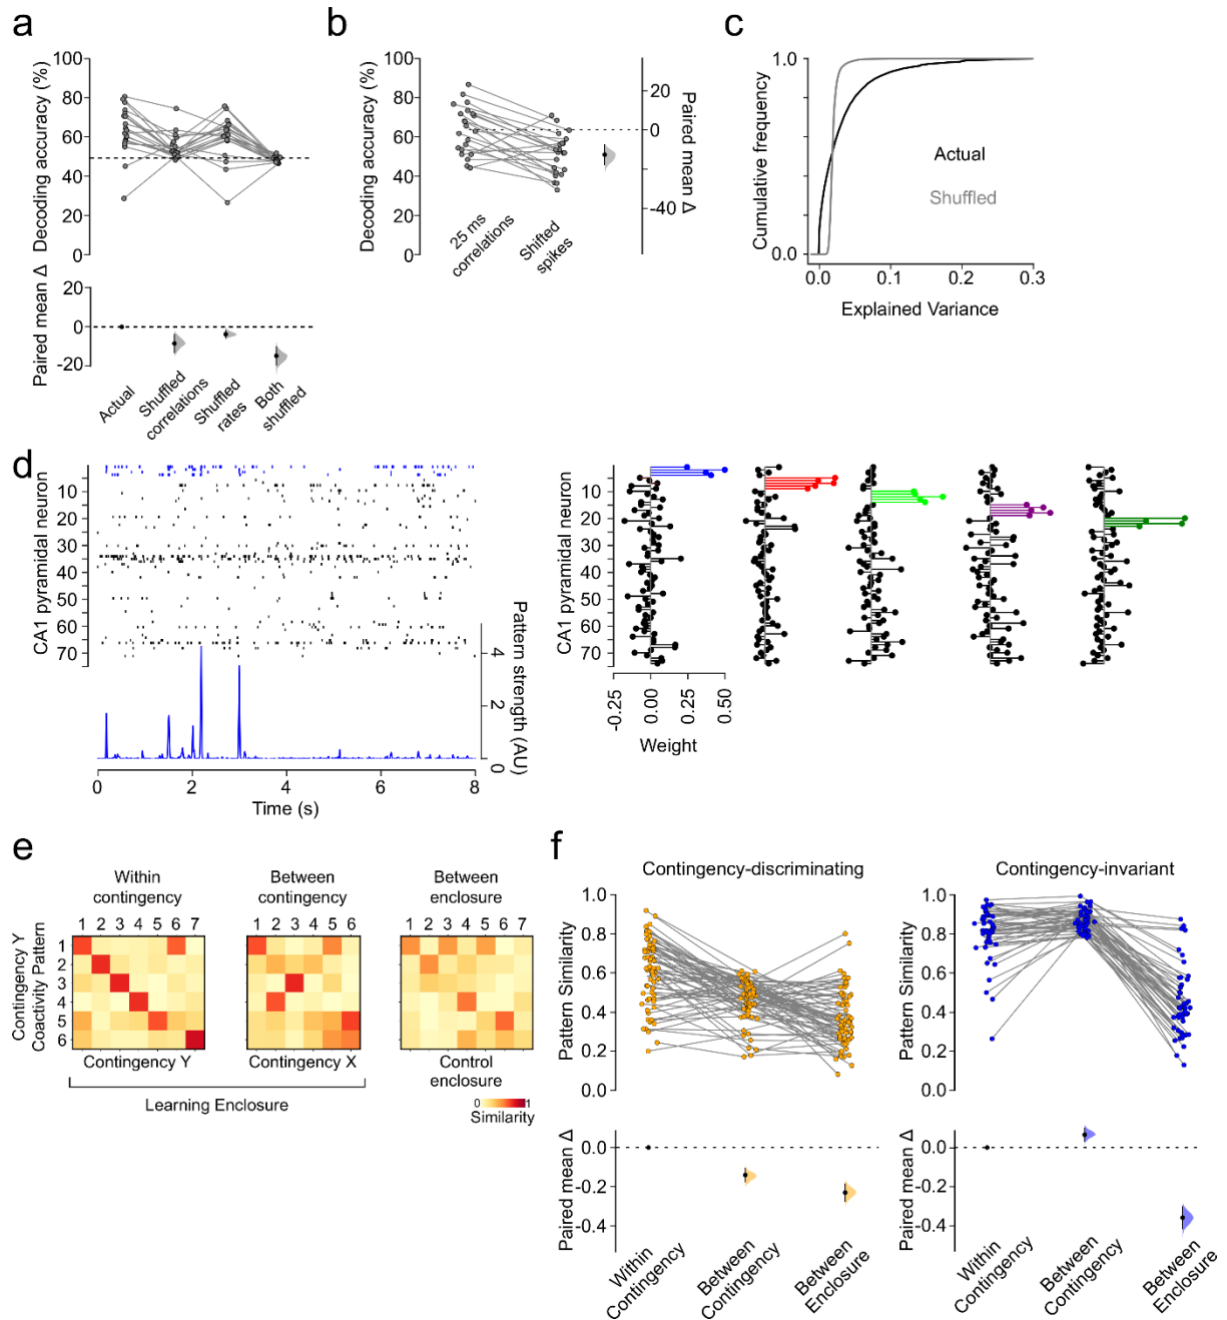


Figure 2. CA1 coactivity-based discrimination of task contingencies

a) DABEST (*Data Analysis with Bootstrap-coupled ESTimation*) plot⁵⁰ used to visualize the effect size of a Gaussian Naive Bayesian classifier decoding contingencies during learning using a combination of CA1 principal neuron firing rates and pairwise correlations. Top panel shows the raw data points for individual days (mean accuracy: actual=66.5±2.9%, shuffled correlations=55.1±1.7%, shuffled rates=62.0±2.5%, both shuffled=48.3±0.2%; N=23 days from 10 mice). Bottom panel shows the effect size for the difference with respect to the left most group (i.e., “Actual”), computed from 1,000 bootstrapped resamples: black-dot, paired mean difference; black-ticks, 95% confidence intervals; filled-curve, resampled paired mean difference distribution. DABEST plots used from here onwards (see Methods under “statistics”). **b)** Decoding accuracy using 25-ms correlations compared to that with spikes jittered to maintain correlations due to slow population dynamics but destroy short-timescale coactivity. Mean accuracy; Actual: 62.8±2.6%, Shifted spikes: 50.8±2.2%; N=23 recording

days. **c)** Explained variance for contingency using trial-by-trial pairwise correlations among all CA1 principal neurons compared to shuffled pairwise correlations. Mean explained variance; Actual: $0.028 \pm 0.000\%$, shuffled: $0.020 \pm 0.000\%$; $N=19,852$ neuron pairs. Wilcoxon test (two-sided): $Z = 96651960.0$, $P=0.02$. **d)** Example CA1 coactivity patterns detected in one learning session. Each pattern is represented as a vector containing the contribution (weight) of each neuron's spiking to the coactivity defining that pattern²⁶. For each pattern, neurons with a weight above 2SD of the mean were referred to as members (color-coded). Shown is an example raster plot of the spike trains (top-left; one neuron per row) along with the coactivity strength of one (dark-blue) pattern over time (bottom-left) and the vectors (right) of other coexisting patterns. Projecting such vectors onto neuron spike trains allowed tracking the time-course of each pattern's strength (e.g., bottom-left dark-blue time course of coactivity peaks for the left-most vector, with the member spiking shown in dark-blue on the raster plot above). **e)** Example similarity matrices of patterns detected in the learning enclosure with contingency Y , compared to patterns detected in sessions with the same (within-contingency; *left*) or the other (between-contingency; *middle*) contingency, or to patterns detected in the control enclosure (between-enclosure; *right*). **f)** Cosine similarity for contingency-discriminating and contingency-invariant patterns across conditions. Contingency-discriminating: within-contingency= 0.60 ± 0.02 , between-contingency= 0.46 ± 0.01 , between-enclosure= 0.37 ± 0.02 ; Contingency-invariant: within-contingency= 0.80 ± 0.02 , between-contingency= 0.87 ± 0.01 , between-enclosure= 0.44 ± 0.02 .

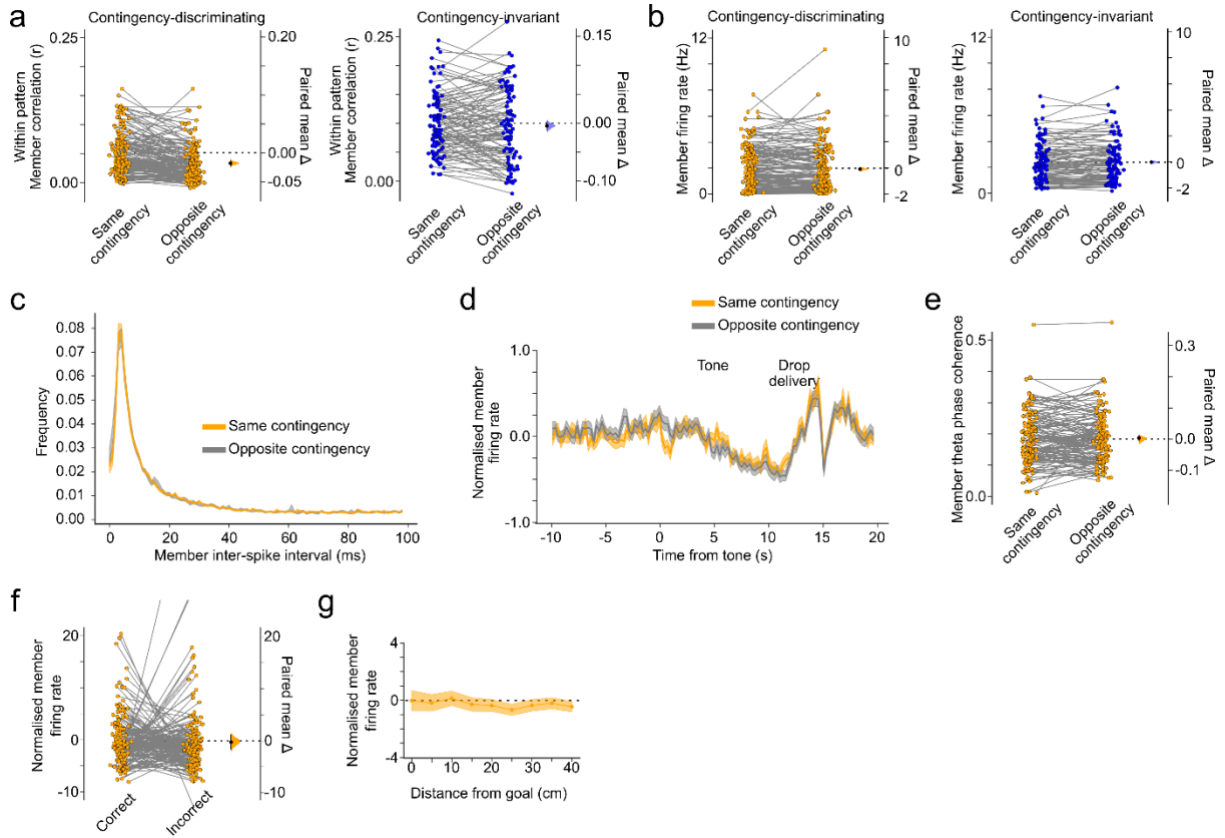


Figure 3. Discrimination of task contingencies is an emergent property of neural coactivity

a) Mean temporal correlations amongst members of a given contingency-invariant and contingency-discriminating patterns in the same contingency (i.e. the contingency in which the patterns were detected) and the opposite contingency. Mean Pearson correlation: Contingency-invariant members: same contingency: 0.100 ± 0.005 , opposite contingency: 0.096 ± 0.006 ; contingency-discriminating members: same contingency: 0.052 ± 0.003 , opposite contingency: 0.033 ± 0.002 . **b)** Mean firing rates of contingency-discriminating and contingency-invariant members. Contingency-discriminating: same-contingency = 2.06 ± 0.14 Hz, opposite-contingency = 1.93 ± 0.15 Hz; Contingency-invariant: same-contingency = 2.28 ± 0.15 Hz, opposite-contingency = 2.31 ± 0.15 Hz. **c)** Inter-spike intervals of contingency-discriminating pattern member neurons in the same contingency and the opposite contingency (Two way repeated measures ANOVA: No main effect of contingency: $F(1)=0.549$, $P=0.458$, $\eta^2=8 \times 10^{-6}$, Main effect of interval: $F(98)=373.0$, $P=0.000$, $\eta^2=0.520$, No contingency:interval interaction: $F(98)=0.738$, $P=0.976$, $\eta^2=0.001$), **d)** Z-scored firing rates during tone and drop delivery contingency-discriminating pattern members pattern member neurons in the same contingency and the opposite contingency (Two way repeated measures ANOVA: No main effect of contingency: $F(1)=4.97 \times 10^{-25}$, $P=1.00$, $\eta^2=1.19 \times 10^{-29}$, Main effect of time: $F(117)=14.4$, $P=2.52 \times 10^{-269}$, $\eta^2=0.040$, No contingency:time interaction: $F(117)=0.92$, $P=0.716$, $\eta^2=0.003$) **e)** spike-phase coherence to theta oscillations of contingency-discriminating pattern member neurons is indistinguishable across contingencies (Mean coherence: same contingency: 0.186 ± 0.007 , opposite contingency: 0.188 ± 0.007). Similarly, the theta-phase preference of contingency discriminating neuron firing is indistinguishable across contingencies (Mean theta-phase preference, with respect to theta peak; same contingency: $156 \pm 6^\circ$; opposite contingency: $146 \pm 6^\circ$; Watson-Wheeler test: $W(2)=0.44$, $P=0.801$). **f)** Contingency-discriminating pattern

members firing rates (z-scored) are indistinguishable at correct and incorrect dispensers (Normalized mean: correct: -0.2 ± 0.4 , incorrect: -0.4 ± 0.5) and **g**) not modulated by distance from goal location. Linear regression of rate against distance from sucrose dispenser: $r = -0.02$, $P = 0.41$.

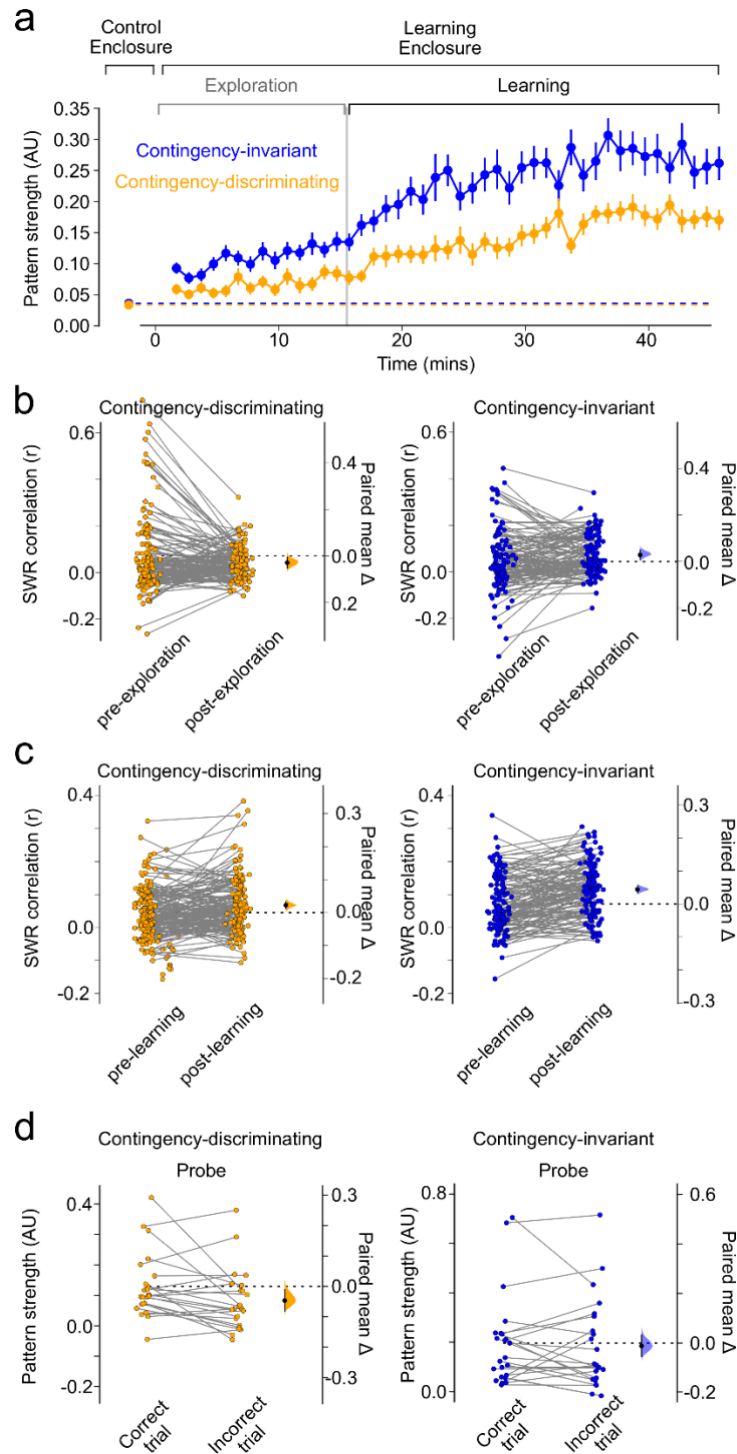


Figure 4. Contingency discriminating patterns develop with learning and predict probe performance

a) Time-course of pattern strength. Contingency X-discriminating and Y-discriminating patterns were pooled, and the coactivity strength of each pattern was quantified in exploration/learning sessions of its preferred contingency. Dashed-lines represent mean pattern strength in control enclosure. Linear regression of strength against time during Exploration (contingency-invariant: $r=0.17$, $P=4.22 \times 10^{-6}$; contingency-discriminating: $r=0.09$, $p=0.003$) and Learning (contingency-invariant: $r=0.15$, $P=2.17 \times 10^{-8}$; contingency-discriminating: $r=0.19$, $P=1.80 \times 10^{-17}$). Slopes of contingency-invariant patterns were steeper than those of contingency-discriminating patterns during exploration (slope= 0.0036 ± 0.0007

and 0.0019 ± 0.0006 units/minute, respectively; Mann-Whitney U test (two-sided): $U=1279.0$, $P=0.03$) but not during learning (slope= 0.0032 ± 0.0009 and 0.0032 ± 0.0005 units/minute, respectively; Mann-Whitney U test (two-sided): $U=1441.0$, $P=0.21$). Error bars represent standard error of the mean across patterns. **b)** Contingency-discriminating and contingency-invariant pattern member correlations during SWRs in the sleep before and after exploration of the novel learning and **c)** before and after contingency learning. (Mean SWR correlation: contingency-invariant members: pre-exploration: 0.045 ± 0.012 , post-exploration (pre-learning): 0.071 ± 0.008 , post-learning: 0.116 ± 0.007 ; contingency-discriminating members: pre-exploration: 0.070 ± 0.015 , post-exploration (pre-learning): 0.045 ± 0.006 , post-learning: 0.068 ± 0.007) **d)** Mean pattern strength before animal's choice during probe trials in sessions where animals performed above chance. Contingency-discriminating: correct= 0.13 ± 0.02 , incorrect= 0.08 ± 0.02 ; contingency-invariant: correct= 0.20 ± 0.04 , incorrect= 0.19 ± 0.04 .

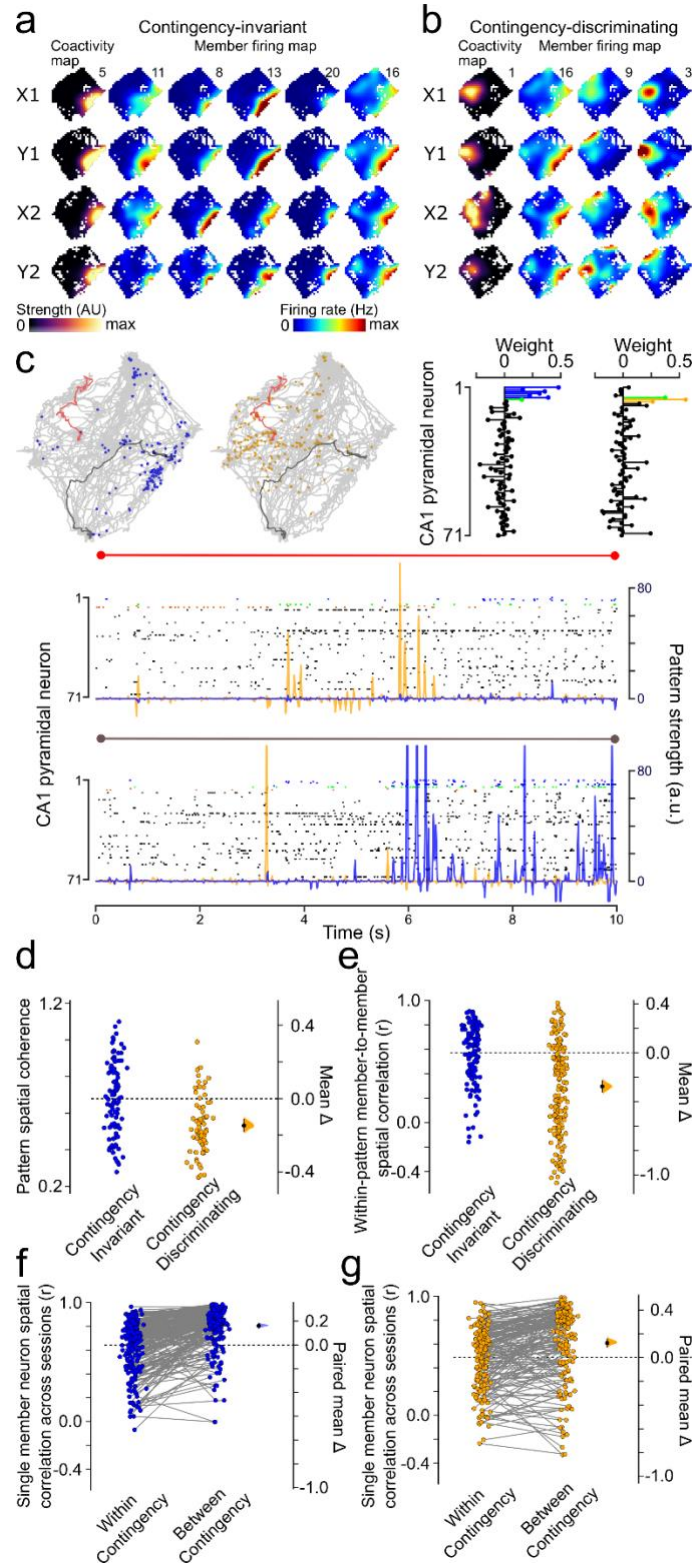


Figure 5. Contingency-discriminating CA1 coactivity is spatially discontinuous
a,b Example coactivity strength maps and corresponding firing rate maps of individual members for a contingency-invariant (a) and a concomitantly recorded contingency-discriminating (b) pattern across all sessions. Note that the right-most member of the contingency-invariant pattern is also a member of the contingency-discriminating pattern. Maximum firing rate (in Hz) or maximum coactivity strength (AU) are shown above each firing rate map or pattern strength map, respectively. **c** Principal neuron spike trains and coactivations along two animal's paths (red and charcoal) for the contingency-invariant

(blue) and contingency-discriminating (orange) patterns shown in (a-b). Members are indicated by the colour-coded assembly weight-vectors and spike times in the raster plots (Note shared member is indicated in green). Line above each raster color-coded as in the path plots. **d)** Contingency-discriminating coactivity is less spatially coherent than that of contingency-invariant patterns (mean spatial coherence: contingency-invariant= 0.64 ± 0.02 , contingency-discriminating= 0.51 ± 0.02). **e)** Contingency-discriminating member neuron firing fields are less spatially overlapping (mean spatial correlation: contingency-invariant= 0.57 ± 0.02 , contingency-discriminating= 0.30 ± 0.03). **f,g)** Spatial correlation of individual contingency-invariant (f) and contingency-discriminating (g) pattern members across sessions of the same (“within”) contingency or of opposite (“between”) contingencies (mean spatial correlation; contingency-invariant: within-contingency= 0.64 ± 0.01 , between-contingency= 0.80 ± 0.01 ; contingency-discriminating: within-contingency= 0.49 ± 0.02 , between-contingency= 0.61 ± 0.02).

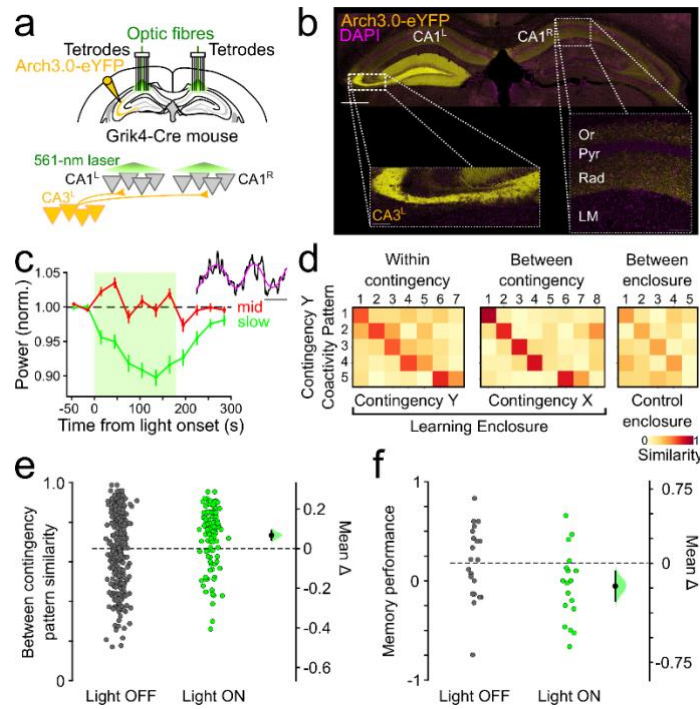


Figure 6. Contingency-discriminating CA1 coactivity requires CA3^L inputs and supports dynamic memory retrieval

a) CA3^L→CA1 optogenetic suppression protocol. CA3^L neurons were transduced with Arch3.0-eYFP in Grik4-Cre mice (n=5) and their axonal projections in the CA1 targeted bilaterally during learning with yellow 561nm-light delivery from implanted optic fibres; 12 tetrodes monitored CA1 neurons. **b)** Expression of Arch3.0-eYFP in the somata of CA3^L neurons and their axons in CA1 bilaterally (top); DAPI-stained nuclei. Higher-magnification images of eYFP-expressing CA3 neurons (bottom-left) and their axons in the contralateral CA1 (bottom-right). Representative of images from 5 animals. Scale bars (top=100μm, bottom=10μm). *Stratum*: Or, *Oriens*; Pyr, *Pyramidale*; Rad, *Radiatum*; LM, *Lacunosum Moleculare*. **c)** Light delivery to Arch3.0-expressing CA3^L axons reduced the power of theta-nested slow, but not mid, gamma oscillations in CA1 (Wilcoxon test (two-sided): slow gamma: $Z=8.0$, $P=9.78 \times 10^{-8}$; mid gamma: $Z=291.0$, $P=0.167$). Inset: example raw trace showing two theta cycles nesting strong mid (~50-90Hz) and slow (~25-40Hz) gamma oscillations respectively; raw trace and theta component in black and magenta, respectively; scale bar=100ms. **d)** Example similarity matrices of patterns detected in the learning enclosure during CA3^L→CA1 input suppression with contingency Y, compared to patterns detected in subsequent sessions with the same (within-contingency; *left*) or the other (between-contingency; *middle*) contingency, or to patterns detected in the control enclosure (between-enclosure; *right*). **e)** CA3^L→CA1 input suppression shifted the between-contingency similarity of CA1 patterns towards contingency invariance; n=54 and 57 patterns detected in contingency X and contingency Y respectively on days with CA3^L→CA1 input suppression (light ON days). Mean between-contingency cosine similarity: light OFF=0.65±0.01, light ON=0.74±0.01. **f)** CA3^L→CA1 input suppression during learning impaired subsequent probe trial performance (mean performance: light OFF days=0.18±0.08, light ON days=-0.05±0.09). All error bars, mean±S.E.M except when used with Gardner-Altman DABEST plots, where they represent mean difference (or paired mean; as indicated) ±95% confidence intervals.

Methods

Animals. These experiments used adult male C57BL/6J mice ($n = 4$; Charles River Laboratories) and transgenic hemizygous Grik4-Cre mice⁵¹ ($n = 11$; The Jackson Laboratories; C57BL/6-Tg(Grik4-cre)G32-4Stl/J, stock number 006474, RRID: IMSR_JAX:006474). Animals were pre-selected based on their propensity to cover a novel open field and to approach a sucrose-baited dispenser within this open field. Animals were housed with their littermates up until the start of the experiment, with free access to water in a dedicated housing facility with a 12/12 h light/dark cycle (lights on at 07:00h), 19–23°C ambient temperature and 40–70% humidity. All mice held in IVC's, with wooden chew stick and nestlets. Food was available *ad libitum* before the experiments (see below), and water available *ad libitum* throughout. Mice were 4–7 months old at the time of testing. Experimental procedures performed on mice in accordance with the Animals (Scientific Procedures) Act, 1986 (United Kingdom), with final ethical review by the Animals in Science Regulation Unit of the UK Home Office.

Surgical procedures. All surgical procedures were performed under deep anaesthesia using isoflurane (0.5–2%) and oxygen (2 l/min), with analgesia provided before (0.1 mg/kg vetergesic) and after (5 mg/kg metacam) surgery. For optogenetic manipulations, AAV5-EF1a-DIO-Arch3.0-eYFP viral vector injections (2×500nl) were performed unilaterally in the dorsal CA3 on either the left or right hemispheres (CA3^L: 5 animals or CA3^R: 6 animals) of Grik4-cre mice using stereotaxic coordinates (site 1: −1.7 mm anteroposterior, ±1.5 mm lateral and −2.1 mm ventral from bregma; site 2: −2.3 mm anteroposterior, ±2.3 mm lateral and −2.3 mm ventral from bregma). The viral vector was delivered at a rate of 100 nl.min^{−1} using a glass micropipette. For electrophysiological recordings, mice were subsequently implanted, 4–6 weeks later, with a microdrive with 12–14 independently movable tetrodes (combined with two optic fibres for optogenetic manipulations; Doric Lenses) targeting the dorsal CA1 bilaterally⁵².

Behaviour. After the recovery period of at least one week following surgical implantation, mice were familiarised daily to the experimental paradigm, including handling, connection to the recording system and exploration of various open fields. Mice were maintained at 90–95% of their free-feeding bodyweight. Animals explored every day the same (triangular) open-field (the “control enclosure”; equilateral triangle; 45-cm side) and a new open-field (45-cm outer width) wherein they were trained in the following three task phases.

Pre-training phase 1 involved conditioning mice to collect transiently available drops of 15% sucrose solution from a single liquid dispenser following a ten-second tone (Extended Data Fig. 1). Sucrose was initially available for 20 seconds before the drop was automatically aspirated by the dispenser. Over multiple pre-training sessions, the drop availability was gradually reduced in 5-second intervals every time the mouse successfully collected sucrose three times consecutively, until a 5-second availability period was reached. To encourage full coverage across the open field, and discourage behavioural persistence at the sucrose dispenser, tones were only delivered after the mouse had moved away from the dispenser to explore the open field. This pre-training phase 1 continued until mice successfully obtained reward on more than 80% of trials, while uniformly exploring the open field; this typically required 5–7 days. All mice actively covered the open-field enclosure and approached the dispenser upon tone presentation.

Next, for pre-training phase 2, mice experienced two pairs of wall-mounted LED displays and two identical liquid dispensers in a novel spatial configuration of the learning enclosure each day (Extended Data Fig. 1). One dispenser delivered sucrose and the other quinine (0.02mM), with both drops simultaneously available for 5 seconds following a 10-second tone. The identity of the dispenser delivering sucrose versus quinine solution could be inferred from the

currently illuminated set of LEDs, but was not directly indicated by the LEDs e.g., Fig. 1b). The LEDs therefore defined a given task contingency (X or Y; Fig. 1a). On a given day of this phase 2 of pre-training, mice initially explored the control enclosure for one session (~15-20 minutes), followed by the exploration of the new learning enclosure for two sessions with only one of the two different sets of LEDs being continuously illuminated in each exploration session. Mice were allowed to rest in a sleep box before and after having explored the enclosure (~20-minute each sleep/rest session). Subsequently, a total of 6 learning sessions (3 of each contingency; ~15-25 minutes per session) were conducted in a pseudorandom order (e.g., X-Y-Y-X-Y-X), with 15 tone presentations (thus 15 trials; inter-trial interval of ~1-2 minutes) in each session. Sessions of the same contingency were never presented 3 times in a row. Sucrose and quinine were delivered simultaneously after 80% of tone presentations in each session, with the remaining 20% of tone presentations being non-reinforced (no sucrose nor quinine delivered). After at least 3 days (and up to 7 days) of pre-training phase 2, animals reached an average performance of at least 80% correct choices on a given day and thus were ready for the third phase: i.e. the training phase.

All behavioural and electrophysiological data quantified in this study are from the training phase (Fig. 1c). Here, the procedure was identical to pre-training phase 2 except that: (i) only two learning sessions of each contingency were presented in alternation (Fig. 1c; X-Y-X-Y) in a novel configuration of the learning enclosure each day; and (ii) a memory probe session was carried out one hour after the final learning session of the day, with an intervening sleep session in between the last learning session and the probe. In this probe session, a total of 24 trials were presented under extinction (i.e., non-reinforced trials where neither sucrose nor quinine was delivered after the tone); 12 trials were presented in each LED-defined contingency, with pseudorandom transitions between the two sets of LEDs defining contingency X and Y while the animal was in the learning enclosure, and with the restriction that either 2 or 5 trials were delivered in succession while a given set of LEDs was active, before the LEDs were switched. The first probe trials for a given recording day were equally likely to be of contingency X or Y. Probe sessions lasted 30 minutes in total, with trials within a given contingency occurring at a rate of 1 trial per minute (with an additional minute delay between the last trial in one contingency and the first trial of another). Only probe sessions where animals covered at least 50% of the enclosure and completed at least 4 trials (i.e. visited at least one dispenser for at least 4 trials) were included in probe analyses (48 out of 71 days satisfied this criteria). Note that animals were allowed to rest/sleep in the sleep box after every session although only three sessions were recorded: (i) Pre-exploration: sleep prior to first exploration session, (ii) Post-exploration/pre-learning: sleep after last exploration session and (iii) Post-learning: sleep after last learning session (Figure 1c).

In addition to these two-choice discrimination training days, mice also performed “One-contingency training days” where we tested learning and memory retrieval of a single behavioural contingency (as opposed to two behavioural contingencies as described above). Here, the task structure was identical to that of training days but one dispenser always delivered sucrose and the other always delivered quinine regardless of the currently illuminated set of LED displays.

To quantify behavioural performance during the learning stage in each training day, we first identified for each tone trial which dispenser the animal approached within the 5-second period of reward availability. To quantify behavioural performance during the probe stage, we identified for each tone trial which dispenser the animal preferred to visit (i.e., spent more time within 5-cm vicinity of the dispenser) during the period from tone onset to 10 seconds after tone offset (i.e. 20-second period). For both learning and probe stages, we next classified each tone trial as being correct or incorrect depending on whether the animal had opted for the sucrose-delivering or the quinine-delivering dispenser, respectively. We finally scored

behavioural performance during learning (e.g. Fig. 1e) and memory performance during probe (e.g. Fig. 1f) sessions by calculating the difference between the number of correct trials minus the number of incorrect trials, divided by the total number of completed trials. A score of 1 thus indicates that mice always visited the correct (sucrose-delivering) dispenser while a score of -1 indicates that mice always visited the incorrect (quinine-delivering) dispenser. Note that during learning trials, since the correct dispenser in one contingency (e.g. X) was the incorrect dispenser in the other (e.g. Y), we display behavioural performance (e.g. Fig. 1e) with respect to the current contingency, with the y-axis ranging from 1 (correct Y) to 0 (chance) to 1 (correct X). A similar behavioural score was calculated for the exploration session (Fig. 1e) on the basis of the total number of visits to each dispenser (since there were no trials during exploration). In order to give equal weighting to trials in each contingency during the probe, the overall probe score for a given day was the mean of means for contingencies X and contingency Y (i.e. (mean score in Y trials + mean score in X trials)/2). Data collection could not be performed blind to the conditions of the experiments since the experimenter had to be aware as to which condition they had to expose each mouse on a given day (Light ON vs Light OFF) and on a given session (which open-field arena/session type).

In vivo ensemble recordings and light delivery. On the morning of each recording day, optimal positioning within the CA1 pyramidal layer was carried out using the local field potential (LFP) signals obtained from each tetrode⁵² in search of multi-unit spiking activity. Tetrodes were then left in position for ~1.5h before commencing recordings. Tetrodes were raised at the end of each recording day to avoid possible mechanical damage overnight. Optical interrogation was performed during learning using a diode-pumped solid-state laser (Laser 2000, Ringstead) that delivers yellow light (561nm; ~18mW output power) to the optic fibres implanted bilaterally above the CA1 pyramidal cell layer in order to suppress CA3→CA1 inputs in Arch3.0-expressing Grik4-Cre mice. Mice were accustomed to light delivery before training. During training, light was delivered for 3-minute periods, 5 times per learning session, with a 2-minute light OFF gap between each light delivery. Trials occurred during the light ON epochs, and at least 1 minute after the onset of each light pulse to allow sufficient time for axonal suppression³⁶. Note that, for the quantification of behavioural effects of input suppression, we compared light ON days to light OFF days from the same animals to provide a within-subject control.

Multichannel data acquisition. Amplification, multiplexing and digitisation of the signals from the electrodes was carried out using a single integrated circuit located on the head of the animal (RHD2164, Intan Technologies; gain x1000; http://intantech.com/products_RHD2000.html). The amplified and filtered (0.09Hz to 7.60kHz) electrophysiological signals were digitised at 20kHz and saved to disk along with the synchronisation signals (transistor-transistor logic digital pulses) reporting the animal's position tracking, laser activation, tone presentation, sucrose and quinine drop delivery, drop removal and LED display illumination. To track the location of the animal, three LED clusters were attached to the electrode casing and captured at 25 frames per second by an overhead colour camera.

Spike detection and unit isolation. The electrophysiological signal was band-pass filtered (800Hz to 5kHz) and single extracellular discharges were detected through thresholding the RMS power spectrum using a 0.2ms sliding window. Detected spikes of the individual electrodes were combined for each tetrode. To isolate spikes which putatively belong to the same neuron, spike waveforms were first up-sampled to 40kHz and aligned to their maximal trough. Principal component analysis was applied to these waveforms ± 0.5 ms from the trough

to extract the first 3-4 principal components per channel, such that each individual spike was represented by 12 waveform parameters. An automatic clustering program (KlustaKwik 2.0, <https://github.com/kwikteam/klustakwik2/>) was run on this principal component space and the resulting clusters were manually recombined and further isolated based on cloud shape in the principal component space, cross-channel spike waveforms, auto-correlation and cross-correlation histograms⁵². An automated clustering was further performed using Kilosort 1.0⁵³ (<https://github.com/cortex-lab/KiloSort>) via the SpikeForest sorting framework⁵⁴ (<https://github.com/flatironinstitute/spikeforest>), with units then automatically curated using metrics derived from the waveforms and spike times, and verified by the operator. All sessions recorded on the same day were concatenated and clustered together. A cluster was only used for further analysis if it showed: stable cross-channel spike waveforms, a clear refractory period in its auto-correlation histogram, well-defined cluster boundaries and an absence of refractory period in its cross-correlation histograms with the other clusters. This study includes a total of 1,124 CA1 principal neurons (853 in light OFF days and 271 in light ON days). Only principal neurons isolated on light OFF days were used in Figures 2 and 3 (and Extended Data figures 4-7) while Figure 4 (and Extended Data figures 8 and 9) included analyses on all principal neurons. All data analysis (with the exception of sharp-wave/ripple analysis) was conducted during active locomotion periods (when the animal was running at a speed of at least 2 cm.s⁻¹).

Neuronal pattern isolation and tracking. Firing patterns of co-active CA1 principal cells were detected using a statistical framework based on independent component analysis²⁶. Spikes discharged by each neuron were counted in 25-ms (or 1000-ms where indicated) time bins and standardised (z-scored, i.e., the activity of each neuron was set to have null mean and unitary variance), to avoid an analytical bias toward neurons with higher firing rates. The neuronal population activity was represented by a matrix in which each element represents the z-scored spike count of a given neuron within a given time bin. We extracted coactivity patterns from this matrix in a two-step procedure. First, the number of significant co-activation patterns embedded within the neuronal population was estimated as the number of principal components of the activity matrix with variances above a threshold derived from an analytical probability function for uncorrelated data. Second, we applied independent component analysis to extract the coactivity patterns from projection of the data into the subspace spanned by the significant principal components (i.e., each coactivity pattern was captured by an independent component). Pattern detection was performed using active periods (speed > 2 cm.s⁻¹) separately during the entire last session of contingency X, contingency Y (i.e. X2 and Y2) or the exploration session of the control enclosure as appropriate. On average, we detected one coactivity pattern for every 5.5±0.3 recorded neurons. To assess the enclosure- or contingency-specificity of coactivity patterns, we compared all patterns detected across enclosures or contingencies, respectively. This was carried out as follows:

(1) We computed the cosine similarity between the weight vector representing a given pattern detected in one session (e.g., X2) and the weight vectors representing each individual pattern detected in another task session (e.g., Y2). By considering a pair of recording sessions this way, this procedure gives a matrix containing the cosine similarity values between each individual pattern detected in one session with each individual pattern detected in the other session. For each pattern, we thus identified its “maximum similarity” value with a pattern (i.e., “the best match”) of another session.

(2) Using this procedure, we obtained three distributions of maximum similarity values for the patterns detected: (i) between two sessions of the same contingency in the learning enclosure (e.g., the “within-contingency” *left* matrix in Fig. 2e), (ii) between two sessions of opposing contingencies in the learning enclosure (e.g., the “between-contingency” *middle* matrix in Fig.

2e), and (iii) between a given task contingency in the learning enclosure versus the patterns detected in the control enclosure (e.g., the “between-enclosure” *right* matrix in Fig. 2e).

(3) We next defined contingency-discriminating patterns as patterns with between-contingency similarity values below the 90th percentile of the between-enclosure similarity distribution.

(4) To investigate the properties of such contingency-discriminating patterns we compared them to a matched group of patterns detected in the same session but characterized with the highest between-contingency similarity values. In other words, we selected the n patterns with the highest between-contingency similarity scores (with n being the number of contingency-discriminating patterns across all recordings). We therefore isolated n contingency-discriminating patterns (with the lowest between-contingency similarity values) and n contingency-invariant patterns (with the highest within-contingency discriminating values) in each contingency. Note that the total number of contingency discriminating patterns before exclusions is $2n$ since there are two contingencies. Subsequently, all patterns that had a within-session maximum similarity below the 90th percentile of the between-enclosure similarity distribution were excluded from further analysis.

Since detected weight vectors were typically asymmetrical (Fig. 2d), the direction where weights were highest was assigned positive weights, and principal CA1 neurons whose weight was positive and exceeded 2 standard deviations from the mean were defined as pattern ‘members’ (mean of $6.0 \pm 0.2\%$ member neurons for each pattern from an average of 37.1 ± 4.2 neurons recorded per day). To assess the robustness of findings regarding member neurons isolated using this membership threshold, we further used a threshold of either 1 or 3 standard deviations (Extended Data Fig. 4d,e; Extended Data Fig. 7d,e). In total, the analyses shown in Figures 2 and 3 included 67 contingency-discriminating patterns (32 in contingency X and 35 in contingency Y) and 152 member neurons (79 in contingency X and 73 in contingency Y), 49 contingency-invariant patterns (104 member neurons); all patterns detected in light OFF (151 in contingency X and 155 in contingency Y) and light ON (56 in contingency X and 52 in contingency Y) days were used in Figure 4.

The activation strength A of each coactivity pattern at time t (e.g. Fig. 2d) was computed as:

$$A_t = Z_t^T P Z_t$$

Where Z_t is a population vector carrying the z-scored rate of each neuron at time t , P is the projection matrix (outer product) of the corresponding independent component, and T is the transpose operator. A_t is therefore the squared projection of Z_t onto the component that represents the coactivity pattern. This projection represents the similarity between the independent component (representing all neurons recorded on that day) and the population rate at a given time bin of 25 ms (or 1000 ms for patterns tracked at this window). The main diagonal of P was set to zero before calculating A_t , in order to eliminate the contribution of single neurons to the coactivity pattern strength. The resulting value of A_t reflected expression strength of a particular coactivity pattern and was used in subsequent calculations of coactivity pattern emergence and spatial tuning. Therefore, the strength of a given pattern at any time point does not reflect only the small number of (“member”) neurons with the highest contribution to that pattern, but rather the entire weight vector representing all neurons. To determine whether pattern expression strength predicted memory probe performance, we calculated each pattern’s strength during the period of tone presentation but before the animal approached either dispenser, and averaged these values during theta cycles across epochs preceding correct or incorrect choices. The same calculation was performed for member neuron firing rates. Strength change across contingencies (Extended Data Fig. 8f) was calculated for each pattern as: (the difference between mean strength (same contingency) and mean strength (opposite contingency) normalised (divided) by mean strength (same contingency); where “same” contingency is the contingency in which a pattern was detected.

Significant co-activation events were defined as time points when co-activation strength was more than 2 standard deviations above the mean for the learning session in which the patterns were detected. Using this threshold, the mean coactivation rate was 0.54 ± 0.02 Hz for contingency X and 0.52 ± 0.02 Hz for contingency Y. This quantification was used in extended data figure 7g,h, and in example traces in figure 3 and extended data figure 9. For all other pattern analyses, we used the raw coactivity strength.

Spatial maps. The recording arena was divided into bins of 1.5×1.5 cm to generate spike count maps (number of spikes fired in each bin) for each unit, or pattern strength map for each co-activation pattern, and an occupancy map (time spent by the animal in each bin). Rates and occupancy were calculated only during active periods (i.e. speed > 2 cm.s⁻¹) and bins visited less than a total of 5 times per session were excluded from subsequent analysis. All maps were then smoothed by convolution with a two-dimensional Gaussian kernel of s.d. equal to two bin widths. Finally, spatial rate maps were generated for each session individually by normalising the smoothed spike count maps by the smoothed occupancy map. Spatial coherence reflects the similarity of the firing rate in adjacent bins, and is the z-transform of the Pearson correlation (across all bins) between the rate in a bin and the smoothed rate of the same bin⁵⁵. The same calculation was used on coactivity pattern strength to calculate pattern spatial coherence. Spatial correlation between maps of member neurons, or co-activation patterns, was calculated as the Pearson correlation coefficient from the direct comparison of the spatial bins between the smoothed place rate maps. This comparison was made between spatial maps of member neurons of the same patterns within the same session to assess the spatial similarity of members of the same coactivity patterns (e.g. Fig. 3e) or between maps of the same member neuron across sessions to assess any possible member neuron contingency-dependent remapping (e.g. Fig. 3f,g). To determine the degree to which pattern coactivations were biased by member firing fields, we calculated an infield coactivation score for each member as the spatial density of coactivations inside the member neuron's firing field (spatial bins within 70% of the peak firing rate bin) minus the outside-the-field coactivation density divided by the sum of those two values. To match the spatial tuning of contingency-discriminating and contingency-invariant pattern members we used the 90th percentile of the contingency-invariant pattern member spatial coherence distribution as a threshold and only included contingency-discriminating pattern members with spatial coherence values above this threshold (Note that this was only done for Extended Data Fig. 7k).

Decoding. To quantify information in spike time correlations and firing rates of CA1 principal neurons we used a Gaussian naïve Bayesian classifier to decode contingency on a trial-by-trial basis (across all learning trials) from pair-wise Pearson correlations of two neurons' spike trains and/or individual neuron firing rates. Uniform priors were used throughout. Since LED-defined contingencies were signalled to the animal throughout learning, we used activity across the 30-second epochs preceding tone onset, the 10-second tone and the 5-second reward availability period to ensure sufficient spikes are used for decoding. We only considered neuronal activity when the animal was at least 10 cm away from both dispensers and during active locomotion (speed > 2 cm.s⁻¹). Trials with less than 10 seconds satisfying these criteria were excluded from analysis. Decoding was carried out using either a combination of individual firing rates and 25-ms pairwise temporal correlations (e.g. Fig. 2a), or pairwise correlations alone (using 25-ms bins, e.g. Fig. 2b; or 1000-ms bins, e.g. Extended Data Fig. 5j). Decoding accuracy was then compared to the mean accuracy of a null distribution, generated by randomly shuffling the contingency label across trials (i.e., X versus Y) 100 times. In addition, we assessed whether contingency information was present in temporal spike correlations beyond population firing rate in two ways. First, we shuffled 25-ms correlations relative to trial labels while keeping

features representing average firing rates aligned with trial labels and re-quantified decoding accuracy (Fig. 2a). Second, we compared decoding accuracy using only 25-ms pairwise correlations to that of using pairwise correlations between the same neurons but with spikes shifted randomly by a value between -1000 and 1000 ms before binning, to remove correlations due to short-timescale coincident activity but keep coactivity associated with slower changes in population activity (Fig. 2b).

Explained Variance. For each pair of neurons, the explained variance for task contingency was calculated using a trial-by-trial Pearson correlations on activity binned using 25-ms bins as follows:

$$EV = \frac{SS_{between\ contingency}}{SS_{Total}} = \frac{nX(\bar{rX} - \bar{rXY})^2 + nY(\bar{rY} - \bar{rXY})^2}{\sum_{i=0}^n (rX_n - \bar{rXY})^2 + \sum_{i=0}^n (rY_n - \bar{rXY})^2}$$

where SS is the sum of squares, nX and nY are number of trials in contingency X and contingency Y respectively, rX , rY and rXY are temporal correlation values in X , Y and all trials respectively.

We compared this EV value with the mean of a null distribution generated by shuffling the contingency labels (i.e., X versus Y) of trials. As for decoding, we used activity across the 30-second epochs preceding tone onset, the 10-second tone and the 5-second reward availability window, only considering neuronal activity when the animal was at least 10 cm away from both dispensers and during active locomotion (speed > 2 cm.s $^{-1}$). Trials with less than 10 seconds satisfying these criteria were excluded from analysis. To assess the spatial congruence of high explained variance pairs, we calculated the spatial map correlation (see ‘‘Spatial maps’’ section above) of each pair in the upper 5th percentile of the EV distribution that was temporally positively correlated in at least one contingency. We used the maximum spatial correlation across both contingencies (e.g., if neurons had more spatially correlated place fields in contingency X than in contingency Y , the value for contingency X was used). We compared such spatial correlation values to those from pairs of neurons in the lower 5th percentile (low explained variance pairs) that were positively temporally correlated in both contingencies (Extended Data Fig. 7i).

LFP analyses. Raw local field potentials (LFPs) were down-sampled from 20kHz to 1250 Hz (order 8 Chebyshev type I filter was applied prior to decimation to avoid aliasing) and then decomposed using Empirical Mode Decomposition (EMD⁵⁶; <https://pypi.org/project/emd/>). In order to avoid mode mixing, we used the mask sift EMD procedure⁵⁷, with sinusoidal masks with the following frequencies: 350, 200, 70, 40, 30 and 7 Hz, which captured mid gamma, slow gamma and theta oscillations as isolated components. To determine individual theta cycles and theta phase, we first detected peaks and troughs of theta with absolute values higher than low-frequency component (sum of all components with main frequencies below the theta signal) envelope, then a theta cycle was defined by pairs of supra-threshold troughs separated at least by 71ms (~ 14 Hz) and no more than 200ms (5 Hz) that surrounded a supra-threshold peak⁵⁸. Theta phase was calculated by linear interpolating neighbouring theta troughs, zero crossings and peaks. For nested-gamma analysis (Fig. 4c, Extended Data Fig. 8a,b,r), instantaneous envelopes and frequencies were calculated by means of the normalised-Hilbert transform⁵⁹. For the time course analysis shown in Fig. 4c and Extended Data Fig. 8r, we adopted a bootstrap procedure to keep the speed distribution of each time bin virtually equal⁵⁸. For each experiment, we used 60 to 30-second pre-laser stimulus windows as a reference for speed distribution. More specifically, we calculated the histogram (linearly-spaced speed bins from 2 to 30 cm.s $^{-1}$) of instantaneous speed values for each theta cycle within that reference window; then a bootstrap consisted of (I) subsampling theta cycles from that reference time window by randomly choosing 75% of the cycles in each speed bin (i.e., maintaining the

original speed histogram proportions); and (2) from all remaining time windows, for each theta cycle in the reference window we randomly chose a cycle with matched speed (no more than 2.5% away from the reference cycle). One hundred such bootstraps were computed for each tetrode, then all tetrodes of each experiment were averaged. Figures show the mean across recording days.

Sharp-wave/ripple (SWR) detection was carried out as follows. First, LFPs of each pyramidal CA1 channel were subtracted by the mean across all channels (common average reference). These re-referenced signals were then filtered for the ripple band (110 to 250 Hz; 4th order Butterworth filter) and their envelopes (instantaneous amplitudes) were computed by means of the Hilbert transform. The peaks (local maxima) of the ripple band envelope signals above a threshold (5 times the median of the envelope values of that channel) were regarded as candidate events. Further, the onset and offset of each event were determined as the time points at which the ripple envelope decayed below half of the detection threshold. Candidate events passing the following criteria were determined as SWR events: (1) ripple band power in the event channel was at least 2 times the ripple band power in the common average reference (to eliminate common high frequency noise); (2) an event had at least four ripple cycles (to eliminate events that were too brief); (3) ripple band power was at least 2 times higher than the supra-ripple band defined as 200-500 Hz (to eliminate high frequency noise, not spectrally compact at the ripple band, such as spike leakage artefacts). We classified tetrodes as either being in the deep or superficial sublayer of the CA1 pyramidal cell layer based on the mean peak amplitude of the SWR events across all sleep sessions. Positive values indicated that the tetrode was in the deep sublayer (i.e. closest to *stratum oriens*) while negative values indicated tetrode was located in the superficial sublayer (i.e. closest to *stratum radiatum*)^{60–62}. To calculate pattern member firing correlations during sleep/rest SWRs, we used SWR events as activity bins and calculated the Pearson correlations between pairs of pattern members separately across each sleep session (i.e. separately for pre-exploration, pre-learning and post-learning sleep sessions; only sessions where at least 200 SWR events were detected were included in this analysis).

Anatomical and histological analysis. All mice were anaesthetised with pentobarbital following completion of the experiments and transcardially perfused with PBS followed by 4% PFA / 0.1% glutaraldehyde in PBS solution. Brains were extracted and kept in 4% PFA for at least 24 h before slicing. Coronal sections (50 μ m thick) were then made and stored in PBS-azide combined with DAPI to stain neuronal somata. All sections were mounted in Vectashield (Vector Laboratories, Cat. No. H-1000) and images of native eYFP fluorescence and DAPI fluorescence were captured with a LSM 880 (Zeiss) confocal microscope using ZEN software (Zeiss Black 2.3).

Statistical analyses. Data were analyzed in Python 3.6 (<https://www.python.org/downloads/release/python-363/>) and using the packages scikit-learn 0.23.2, statsmodels 0.12.1, Numpy 1.18.1, Scipy 1.4.1, Matplotlib 3.1.2, Pandas 0.25.3 and Seaborn 0.11.0. Error bars, mean \pm S.E.M unless otherwise stated. *Ns* refer to recording days for behavioural preference figures and LFP analysis. For unit data, *Ns* refer to coactivity patterns, coactivity pattern members or all principal neurons as indicated. Where indicated, we also used mice as *Ns* for behavioural and unit data, with values averaged across days for a given mouse for behavioural data and across neurons/patterns for a given mouse for unit data. Data Analysis with Bootstrap-coupled ESTimation (DABEST) plots⁵⁰ are used throughout the manuscript to visualise the effect size by plotting the data against a mean (or paired mean) difference between the left-most condition and one or more conditions on the right (right y-axis) and compare this difference against zero using 1,000 bootstrapped resamples: black-dot indicates mean

difference or mean paired difference (as indicated in the right y-axis label) and black ticks depict error bars representing 95% confidence intervals, shaded area represents kernel density estimate for visualizing the resampled mean (or paired mean) difference distribution. All *P* values were calculated as specified in figure legends. For ANOVAs, type II sums of squares were used throughout, with degrees of freedom reported in parentheses after “F” (e.g. $F(1)=...$). Data distributions were assumed to be normal for ANOVAs but this was not formally tested. No statistical methods were used to pre-determine sample sizes but our sample sizes are similar to those reported in previous publications (e.g.,^{25,27,31,34,63}). Neural and behavioural data analyses were conducted in an identical way regardless of the experimental condition from which the data were collected. See also the corresponding Life Sciences Reporting Summary.

A total of 15 mice were used in this study: 5 animals injected with an Archaelrhodopsin expressing construct in the left CA3 and 6 animals injected with the same construct in the right CA3. A further 4 mice did not receive CA3 injections.

The behavioural data in figure 1 (and extended data figures 2 and 3) are from light OFF training days (termed light OFF days throughout) from all 15 animals used in this study (71 light OFF days). For behavioural quantification in figure 6 (and extended data figure 8) we used both light OFF and light ON days (separately as indicated) from 5 animals injected with an Archaelrhodopsin expressing construct in the left CA3 (33 light OFF days and 20 light ON days) and 6 animals injected in the right CA3 for behavioural experiments (18 light OFF days and 23 light ON days). Mice were randomly assigned to left versus right CA3 injected groups. Light OFF and light ON days were pseudorandomized and interleaved for each animal (ensuring the first day of training was always Light OFF). A further 16 days of one-contingency learning were recorded (10 light OFF days and 6 light ON (CA3L-suppression) days; extended data figure 8m).

Electrophysiological unit data in figures 2-5 (and extended Data figures 4-7) are from light OFF days from 10 animals (23 recording days); 4 (of the 5) left CA3 injected animals, 2 (of the 6) right CA3 injected animals and all 4 mice without CA3 injections. Electrophysiological unit data from light ON days from 4 (of the 5) left CA3 injected animals are used (8 recording days), and compared to all light off electrophysiological unit data, in figures 4 (and extended data figure 8 and 9).

Code availability: The software used for data acquisition and analysis are available using the web links mentioned in the methods.

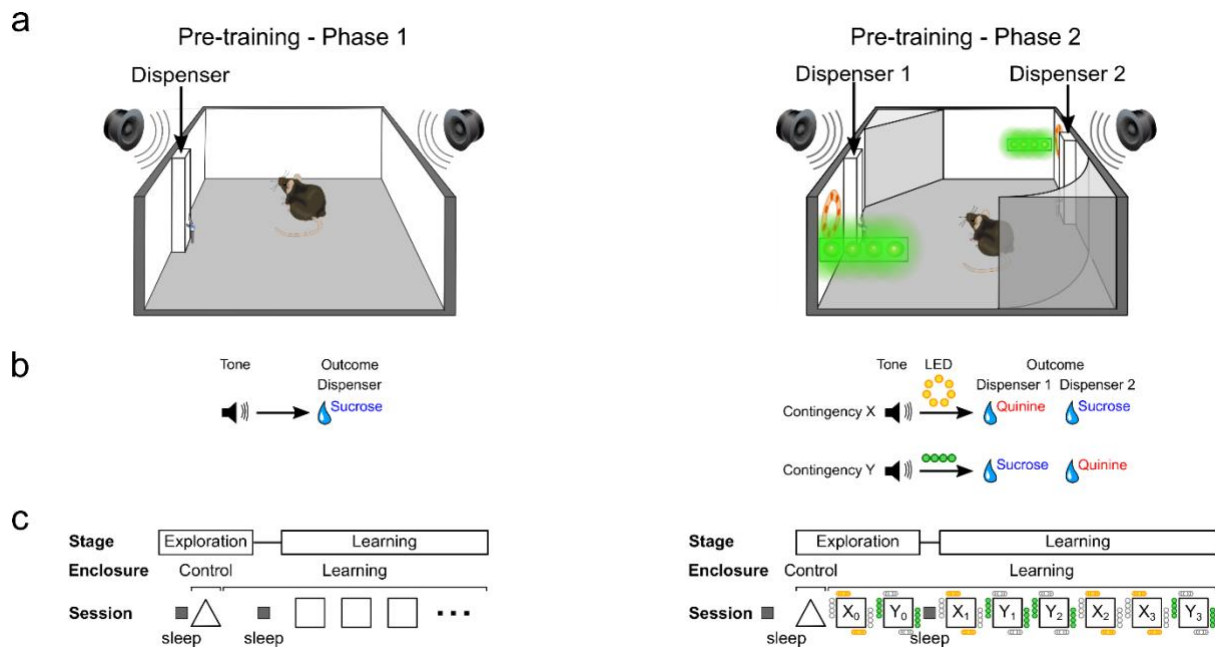
Data availability: The data that support the findings of this study are available from the corresponding author upon request.

Methods-only References

51. Nakazawa, K. *et al.* Requirement for hippocampal CA3 NMDA receptors in associative memory recall. *Science* **297**, 211–218 (2002).
52. McNamara, C. G., Tejero-Cantero, Á., Trouche, S., Campo-Urriza, N. & Dupret, D. Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. *Nat. Neurosci.* **17**, 1658–1660 (2014).
53. Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M. & D, H. K. Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. *bioRxiv* 061481 (2016) doi:10.1101/061481.

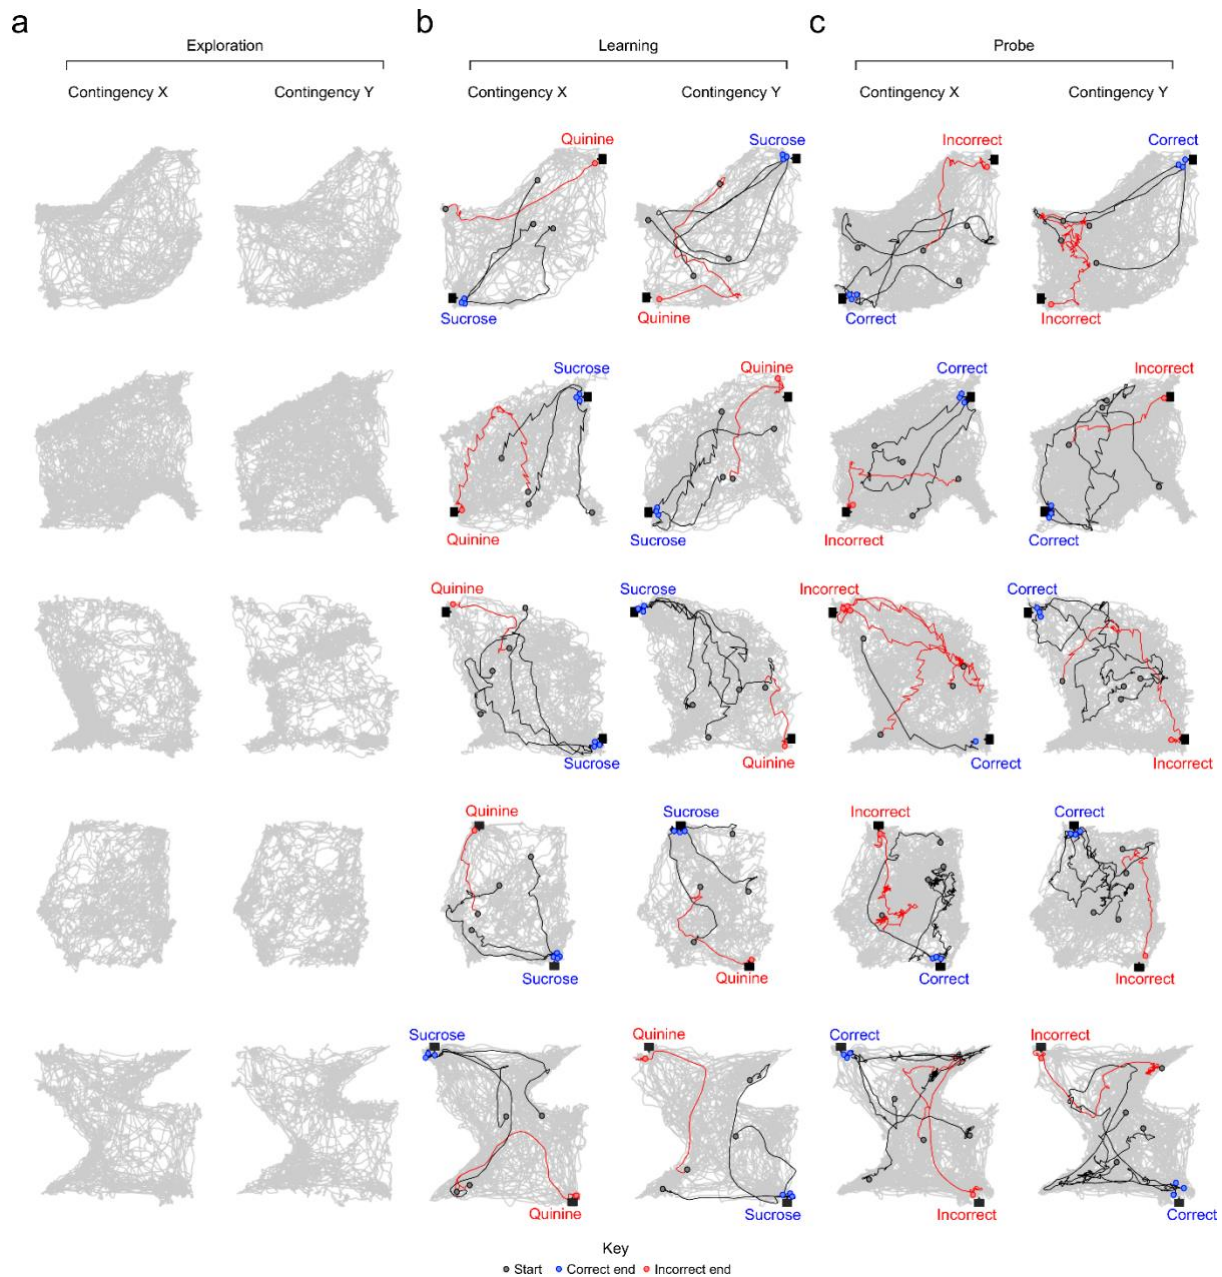
54. Magland, J. *et al.* SpikeForest, reproducible web-facing ground-truth validation of automated neural spike sorters. *eLife* **9**, e55167 (2020).
55. Zhang, S., Schönfeld, F., Wiskott, L. & Manahan-Vaughan, D. Spatial representations of place cells in darkness are supported by path integration and border information. *Front. Behav. Neurosci.* **8**, (2014).
56. Huang, N. E. *et al.* The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proc. R. Soc. Lond. Ser. Math. Phys. Eng. Sci.* **454**, 903–995 (1998).
57. Deering, R. & Kaiser, J. F. The Use of a Masking Signal to Improve Empirical Mode Decomposition. in *Proceedings. (ICASSP '05). IEEE International Conference on Acoustics, Speech, and Signal Processing, 2005.* vol. 4 485–488 (IEEE, 2005).
58. Lopes-Dos-Santos, V. *et al.* Parsing Hippocampal Theta Oscillations by Nested Spectral Components during Spatial Exploration and Memory-Guided Behavior. *Neuron* **100**, 940-952.e7 (2018).
59. Huang, N. E. *et al.* On Instantaneous Frequency. *Adv. Adapt. Data Anal.* **01**, 177–229 (2009).
60. Mizuseki, K., Diba, K., Pastalkova, E. & Buzsáki, G. Hippocampal CA1 pyramidal cells form functionally distinct sublayers. *Nat. Neurosci.* **14**, 1174–1181 (2011).
61. Oliva, A., Fernández-Ruiz, A., Buzsáki, G. & Berényi, A. Spatial coding and physiological properties of hippocampal neurons in the Cornu Ammonis subregions. *Hippocampus* **26**, 1593–1607 (2016).
62. Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A. & Buzsáki, G. Fast network oscillations in the hippocampal CA1 region of the behaving rat. *J. Neurosci.* **19**, RC20 (1999).
63. Davoudi, H. & Foster, D. J. Acute silencing of hippocampal CA3 reveals a dominant role in place field responses. *Nat. Neurosci.* **22**, 337–342 (2019).

Extended Data Figures



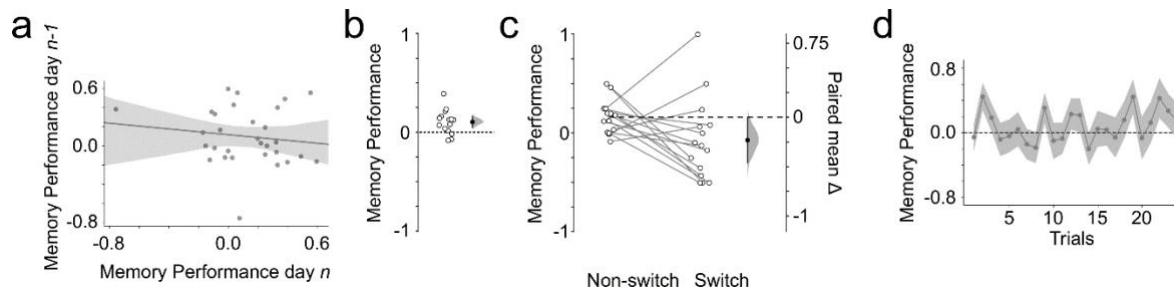
Extended Data Fig. 1. Schematics of pretraining protocol.

Schematics of pretraining phases 1 (left) and 2 (right). a) Schematic of example learning enclosures. b) Learning in pretraining phase 1 involved associating a tone with delivery of sucrose from one dispenser. In pretraining phase 2, animals learned two new LED-tone-outcome associations each day. c) In pretraining phase 1 animals initially explored the control enclosure and then experienced between 2-6 sessions of tone-defined trials. In pretraining phase 2, after exploring the control enclosure and the learning enclosure (with each LED set active in turn), tone-defined trials were presented in 6 learning sessions (3 in contingency X and 3 in contingency Y) that were pseudo-randomly ordered each day. No probe tests were carried out in either pretraining phase.



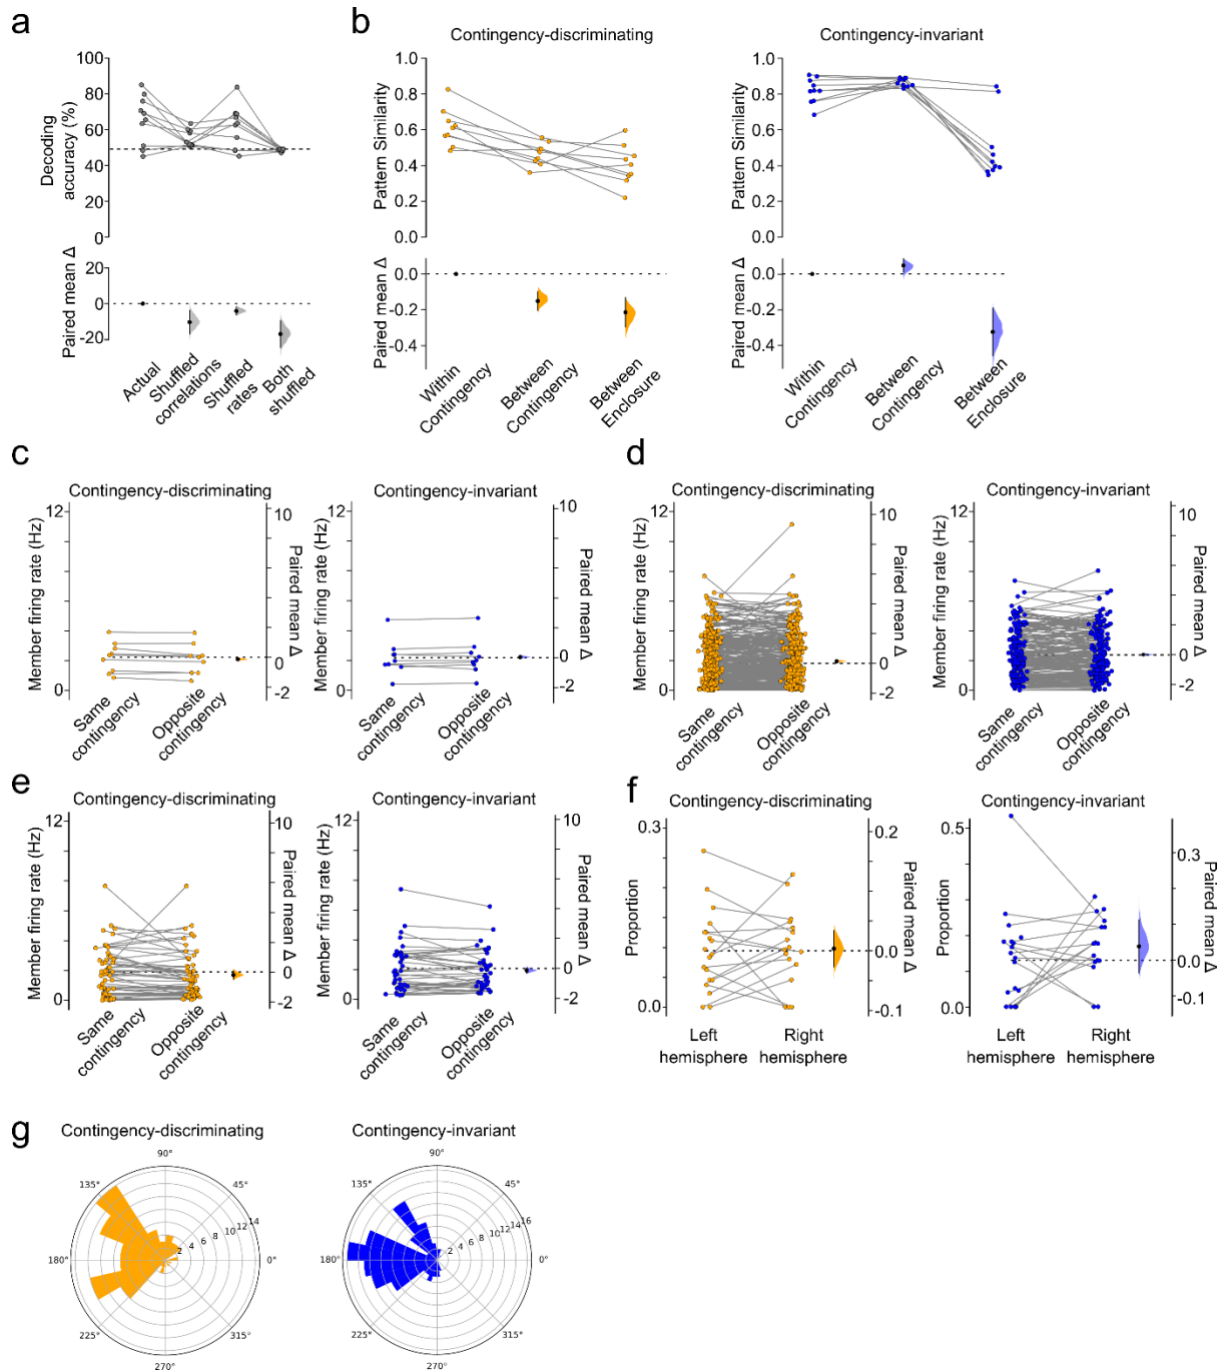
Extended Data Fig. 2. Examples of enclosure set ups and animal paths across task stages and recording days.

Enclosure set ups across distinct behavioural days. Animal coverage represented in grey. **a)** Example coverage paths for pre-learning exploration of learning enclosures. **b)** Example animal paths during learning trials in contingency X and contingency Y. **c)** Example animal paths during probe trials in contingency X and contingency Y. Paths of the animal during trials (correct path: black; incorrect path: red) are overlaid onto overall coverage (grey) for a single learning session. Black circles represent path starting points; blue and red circles represent correct and incorrect end points, respectively.



Extended Data Fig. 3. Dynamics of memory performance.

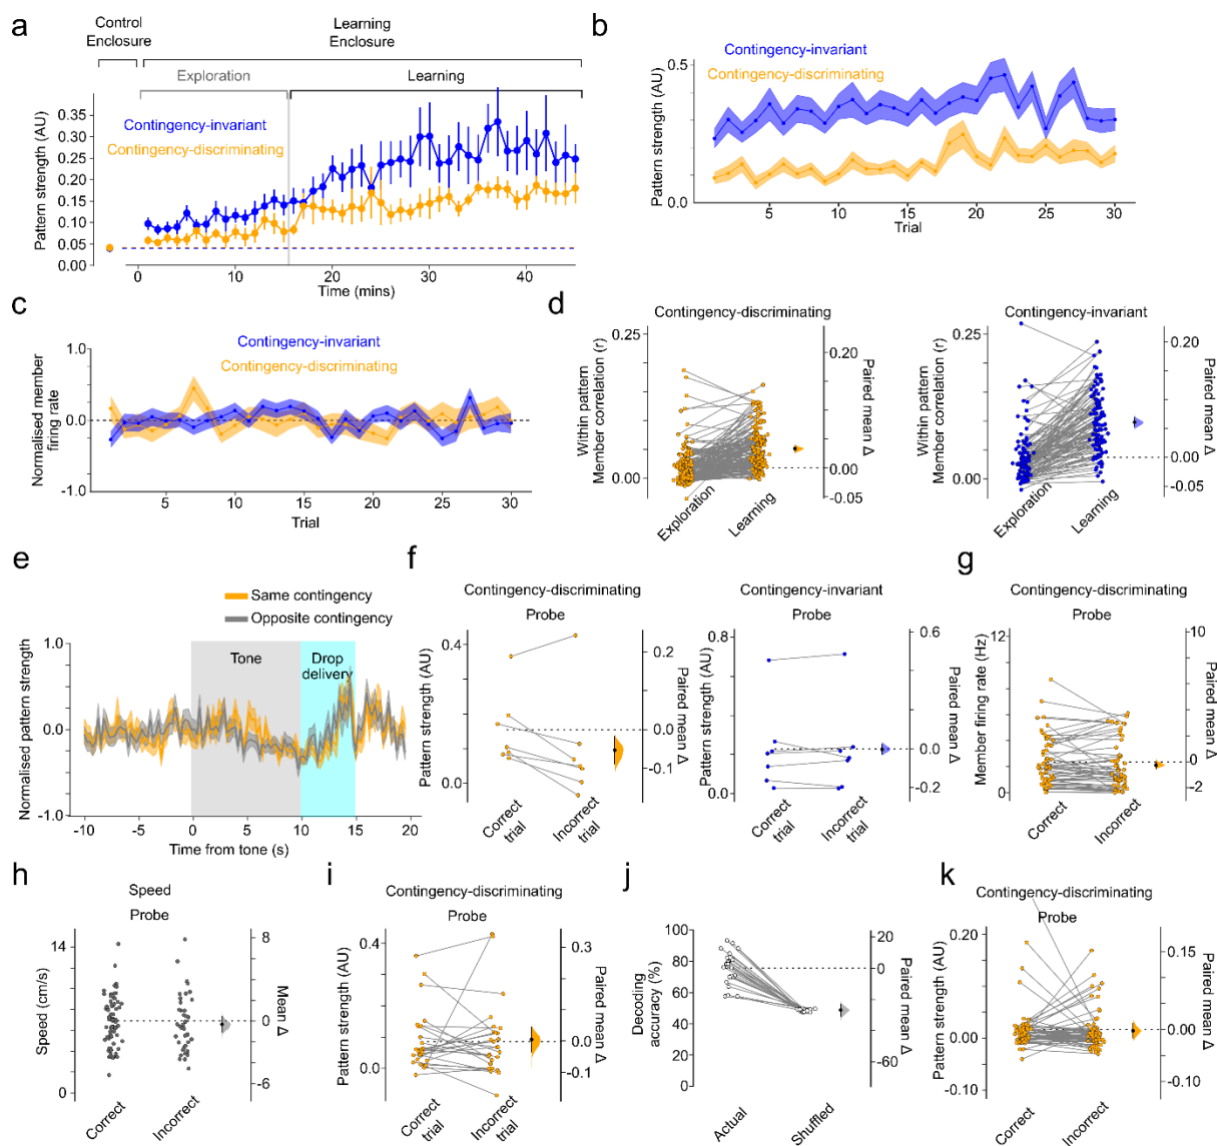
a) Lack of a relationship between performance on probe trials of a recording day and those of the previous day (Regression line shown in dark grey; light grey shaded area represents 95% confidence intervals). Linear regression of probe performance on day n against probe performance on day " $n-1$ ": $r=-0.155$, $P=0.413$. b) Behavioural performance during memory probe test, per mouse. Here the memory performance for each individual mouse is averaged across days, with each data point showing average performance for a single mouse (mean performance= 0.10 ± 0.03). c) Probe performance (per mouse) is weaker during the first trial following a switch in LED displays (switch trials; Mean performance: -0.07 ± 0.11) compared to subsequent trials (non-switch trials; Mean performance: 0.16 ± 0.04). d) Probe performance does not change systematically across probe trials and hence no further learning is observed during memory retrieval. Linear regression of performance against trial number $r=0.030$, $P=0.442$.



Extended Data Fig. 4. Within- and between-contingency properties of coactivity patterns and their member neurons.

a) A Gaussian Naive Bayesian classifier was trained to decode contingencies using a combination of CA1 principal neuron firing rates and pairwise correlations. Shuffling rates, correlations or both impairs classifier performance. Data points represent individual mice. Mean accuracy; actual: $65.3 \pm 4.3\%$, shuffled correlations: $55.0 \pm 1.5\%$, shuffled rates: $61.2 \pm 3.7\%$, both shuffled: $48.3 \pm 0.2\%$; $N=10$ mice. b) Cosine similarity of contingency-discriminating and contingency-invariant patterns across conditions per mouse. Contingency-discriminating: Within-contingency: 0.62 ± 0.03 , between-contingency: 0.47 ± 0.02 , between-enclosure: 0.41 ± 0.04 ; Contingency-invariant: Within-contingency: 0.82 ± 0.02 , between-contingency: 0.87 ± 0.01 , between-enclosure: 0.49 ± 0.06 . Note that $N=10$ animals for contingency-invariant patterns but 9 animals for contingency-discriminating patterns as no

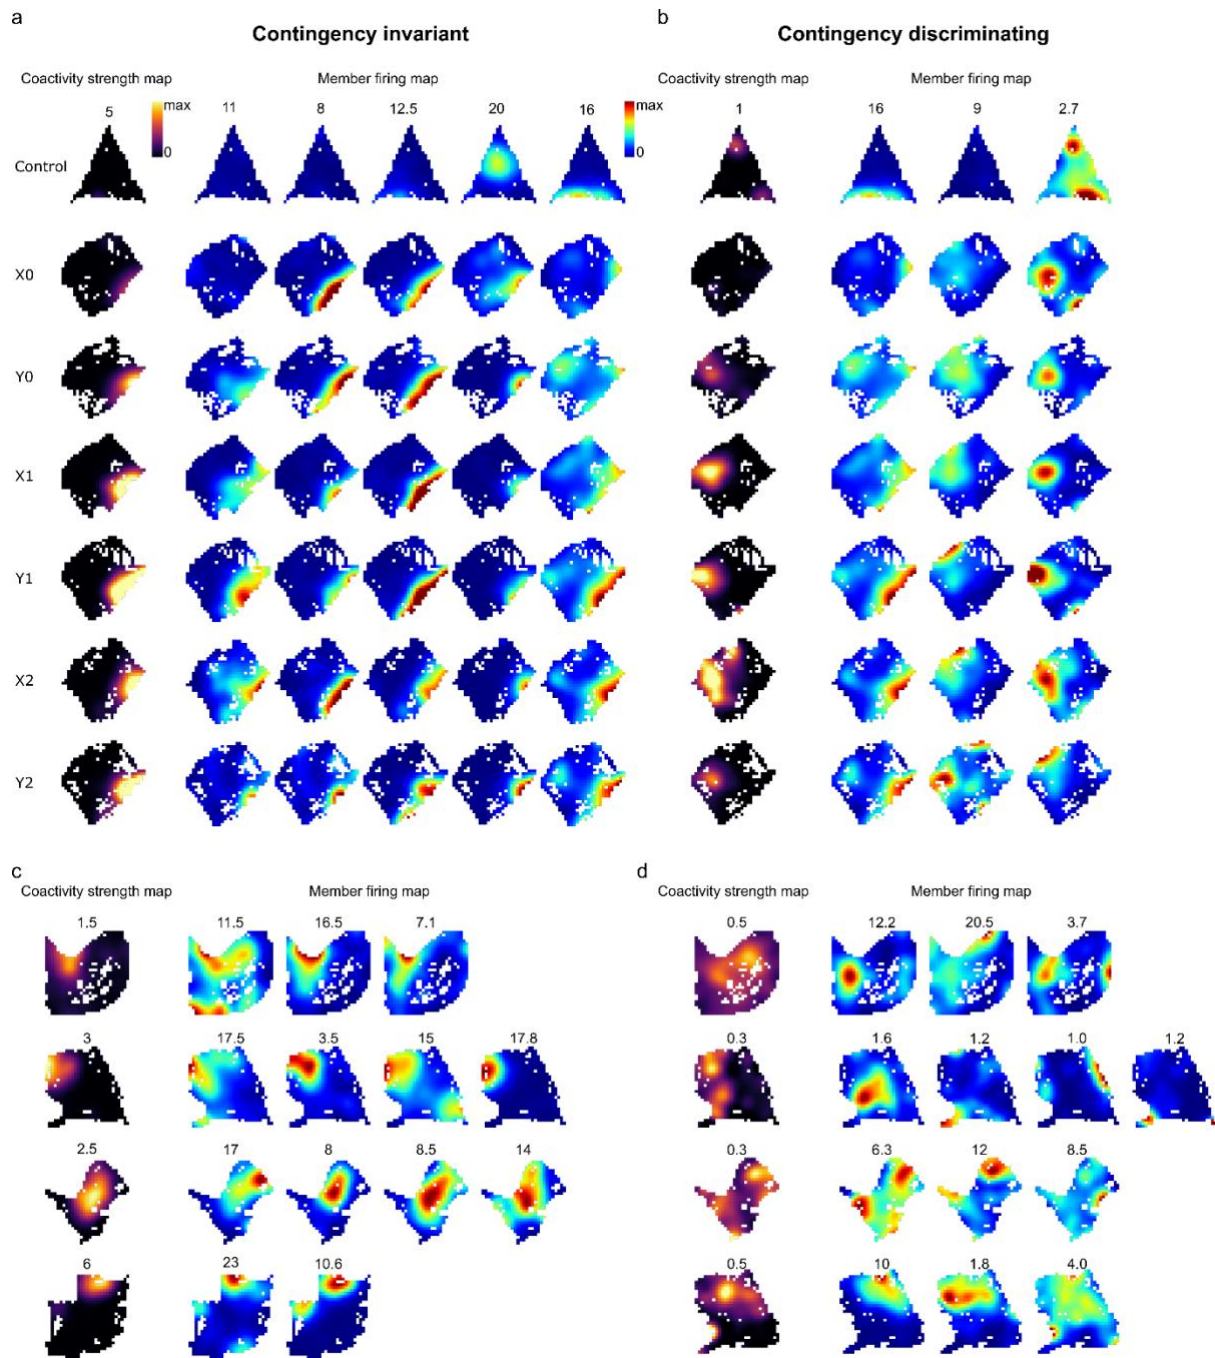
such patterns could be detected in one animal. c) Average firing rate of contingency-discriminating and contingency-invariant member neurons per mouse. Contingency-discriminating: Same contingency: 2.22 ± 0.33 Hz, opposite contingency: 2.09 ± 0.33 Hz; Contingency-invariant: Same contingency: 2.15 ± 0.41 Hz, opposite contingency: 2.20 ± 0.42 Hz. N=9 animals for contingency discriminating and contingency invariant patterns as for one animal, none of the detected contingency-invariant patterns had “members” (i.e. neurons with a weight of more than 2 standard deviations above the pattern weight vector mean). Average firing rate of contingency-discriminating and contingency-invariant member neurons using d) 1 standard deviation (Contingency-discriminating: Same contingency: 2.00 ± 0.09 Hz, opposite contingency: 1.95 ± 0.09 Hz; Contingency-invariant: Same contingency: 2.38 ± 0.11 Hz, opposite contingency: 2.42 ± 0.11 Hz) or e) 3 standard deviations as weight thresholds for defining pattern members (Contingency-discriminating: Same contingency: 1.93 ± 0.23 Hz, opposite contingency: 1.70 ± 0.22 Hz; Contingency-invariant: Same contingency: 2.06 ± 0.25 Hz, opposite contingency: 1.93 ± 0.21 Hz). f) Proportion of principal neurons recorded from the CA1 on the left or the right hemisphere that are members of contingency-discriminating patterns (Mean proportion: left hemisphere: 0.104 ± 0.017 , right hemisphere: 0.087 ± 0.015) or contingency-invariant patterns (Mean proportion: left hemisphere: 0.129 ± 0.022 , right hemisphere: 0.179 ± 0.033). g) Contingency discriminating pattern members showed a trend towards a preference for earlier phases of theta cycles compared to contingency invariant pattern members (Mean theta-phase preference, with respect to theta peak; contingency-discriminating pattern members: $156 \pm 6^\circ$; contingency-invariant pattern members: $174 \pm 5^\circ$; Watson-Wheeler test: $W(2)=5.23$, $P=0.073$).



Extended Data Fig. 5. Contingency discriminating and invariant coactivity patterns show distinct relationships to task phase and performance.

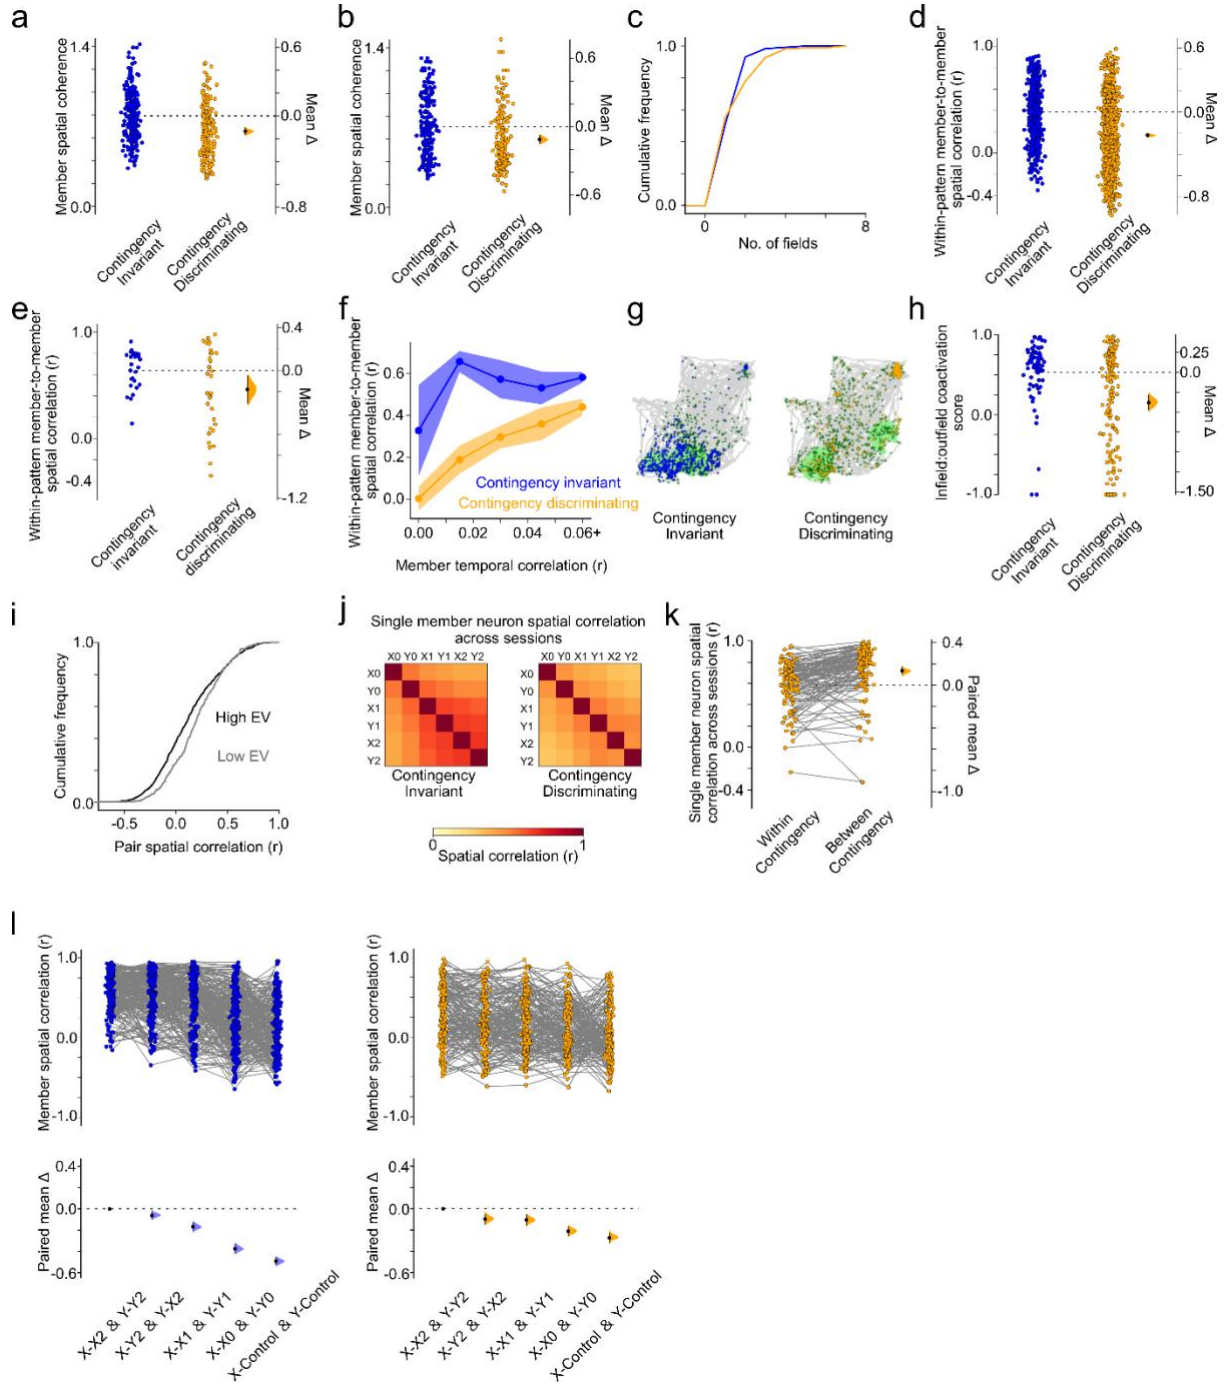
a) Time-course of pattern strength changes with mice as Ns. Contingency X-discriminating and contingency Y-discriminating patterns were pooled and the strength of all patterns of a given type in its preferred contingency were averaged per mouse and the mean value quantified in exploration/learning sessions. Dashed lines represent mean pattern strength in Control enclosure. Linear regression of strength against time during Exploration (contingency-invariant: $r=0.27$, $P=8.15 \times 10^{-4}$; contingency-discriminating: $r=0.20$, $P=0.02$) and Learning (contingency-invariant: $r=0.19$, $P=9.30 \times 10^{-4}$; contingency-discriminating: $r=0.21$, $P=4.92 \times 10^{-4}$). Slopes of contingency invariant patterns showed a trend towards being higher than those of contingency discriminating patterns during exploration (slope = 0.0041 ± 0.0012 and 0.0025 ± 0.0015 units/minute respectively; Mann Whitney U test (two-sided): $U=25.0$, $P=0.06$) but not during learning (slope = 0.0038 ± 0.0015 and 0.0021 ± 0.0007 units/minute respectively; Mann Whitney U test (two-sided): $U=36.0$, $P=0.24$). $N=10$ animals for contingency invariant patterns but 9 animals for contingency discriminating patterns as no such patterns were detected in one animal. b) Increases in contingency-invariant and contingency-discriminating pattern strengths plotted as a function of learning trials. Contingency X-discriminating and Y-discriminating patterns were pooled, and the coactivity

strength of each pattern was quantified in learning trials of its preferred contingency. Linear regression of strength against trials (contingency-invariant: $r=0.04$, $P=0.020$; contingency-discriminating: $r=0.13$, $P=1.05 \times 10^{-8}$). Shaded area represents variability (Standard error of the mean) across coactivity patterns. c) No changes in member neuron firing rates (z-scored) across learning trials. Linear regression of firing rate against trials (contingency-invariant: $r=0.0016$, $P=0.94$; contingency-discriminating: $r=0.0022$, $P=0.94$). Shaded area represents variability (Standard error of the mean) across coactivity pattern members. d) Temporal correlations (Pearson r values) amongst each member neuron of a pattern and other members in the same pattern between exploration and learning (Mean Pearson correlation: Contingency-invariant members: exploration: 0.037 ± 0.004 , learning: 0.098 ± 0.005 ; contingency-discriminating members: exploration: 0.019 ± 0.003 , learning: 0.052 ± 0.003). e) Z-scored contingency discriminating pattern strength in the same contingency and the opposite contingency during tone and drop delivery. This is the point when animals' behaviour is maximally different between contingencies, as animals head towards opposite dispensers (Figure 1e and Extended Data figure 2b). Despite this, the normalized time course of coactivity pattern strength was indistinguishable across contingencies (Two way repeated measures ANOVA: No main effect of contingency: $F(1)=1.5 \times 10^{-26}$, $P=1.00$, $\eta^2=9.39 \times 10^{-31}$, Main effect of time: $F(117)=3.41$, $P=7.61 \times 10^{-32}$, $\eta^2=0.025$, No contingency:time interaction: $F(117)=0.76$, $P=0.98$, $\eta^2=0.006$). f) Pattern strength before animal's choice during probe trials, on days where overall probe performance was above chance, averaged per mouse. Contingency discriminating mean strength: correct: 0.14 ± 0.04 , incorrect: 0.09 ± 0.05 ; contingency-invariant: correct: 0.23 ± 0.08 , incorrect: 0.23 ± 0.09 . $N=7$ animals for both contingency-discriminating and contingency-invariant patterns reflecting the number of animals with recording days in which: 1) units were recoded and isolated; 2) animals performed above chance in the probe; 3) coactivity patterns of the indicated type were detected. g) Contingency-discriminating pattern member firing rate is indistinguishable before correct vs incorrect probe trials on days where overall probe performance was above chance. Mean member rate: correct: 2.32 ± 0.26 Hz, incorrect: 2.15 ± 0.26 Hz. h) Mouse running speed before correct and incorrect trials. Mean speed: correct: 6.90 ± 0.28 cm.s⁻¹, incorrect: 6.58 ± 0.43 cm.s⁻¹. i) Contingency-discriminating coactivity patterns are indistinguishable before correct trials compared to incorrect trials when animal's overall probe performance is not above chance level. Mean strength: correct: 0.086 ± 0.019 , incorrect: 0.090 ± 0.023 . j) Decoding accuracy using 1000 ms pairwise correlations compared to shuffled controls. (Mean accuracy; Actual: $75.7 \pm 2.1\%$, shuffled: $48.8 \pm 0.2\%$; $N=23$ recording days). k) Contingency-discriminating coactivity patterns, detected across 1000ms windows, are not stronger before correct compared to incorrect choices on memory probe trials, on days where overall probe performance was above chance. Mean strength: correct: 0.017 ± 0.006 , incorrect: 0.016 ± 0.004 .



Extended Data Fig. 6. Example pattern activation and member firing rate maps.

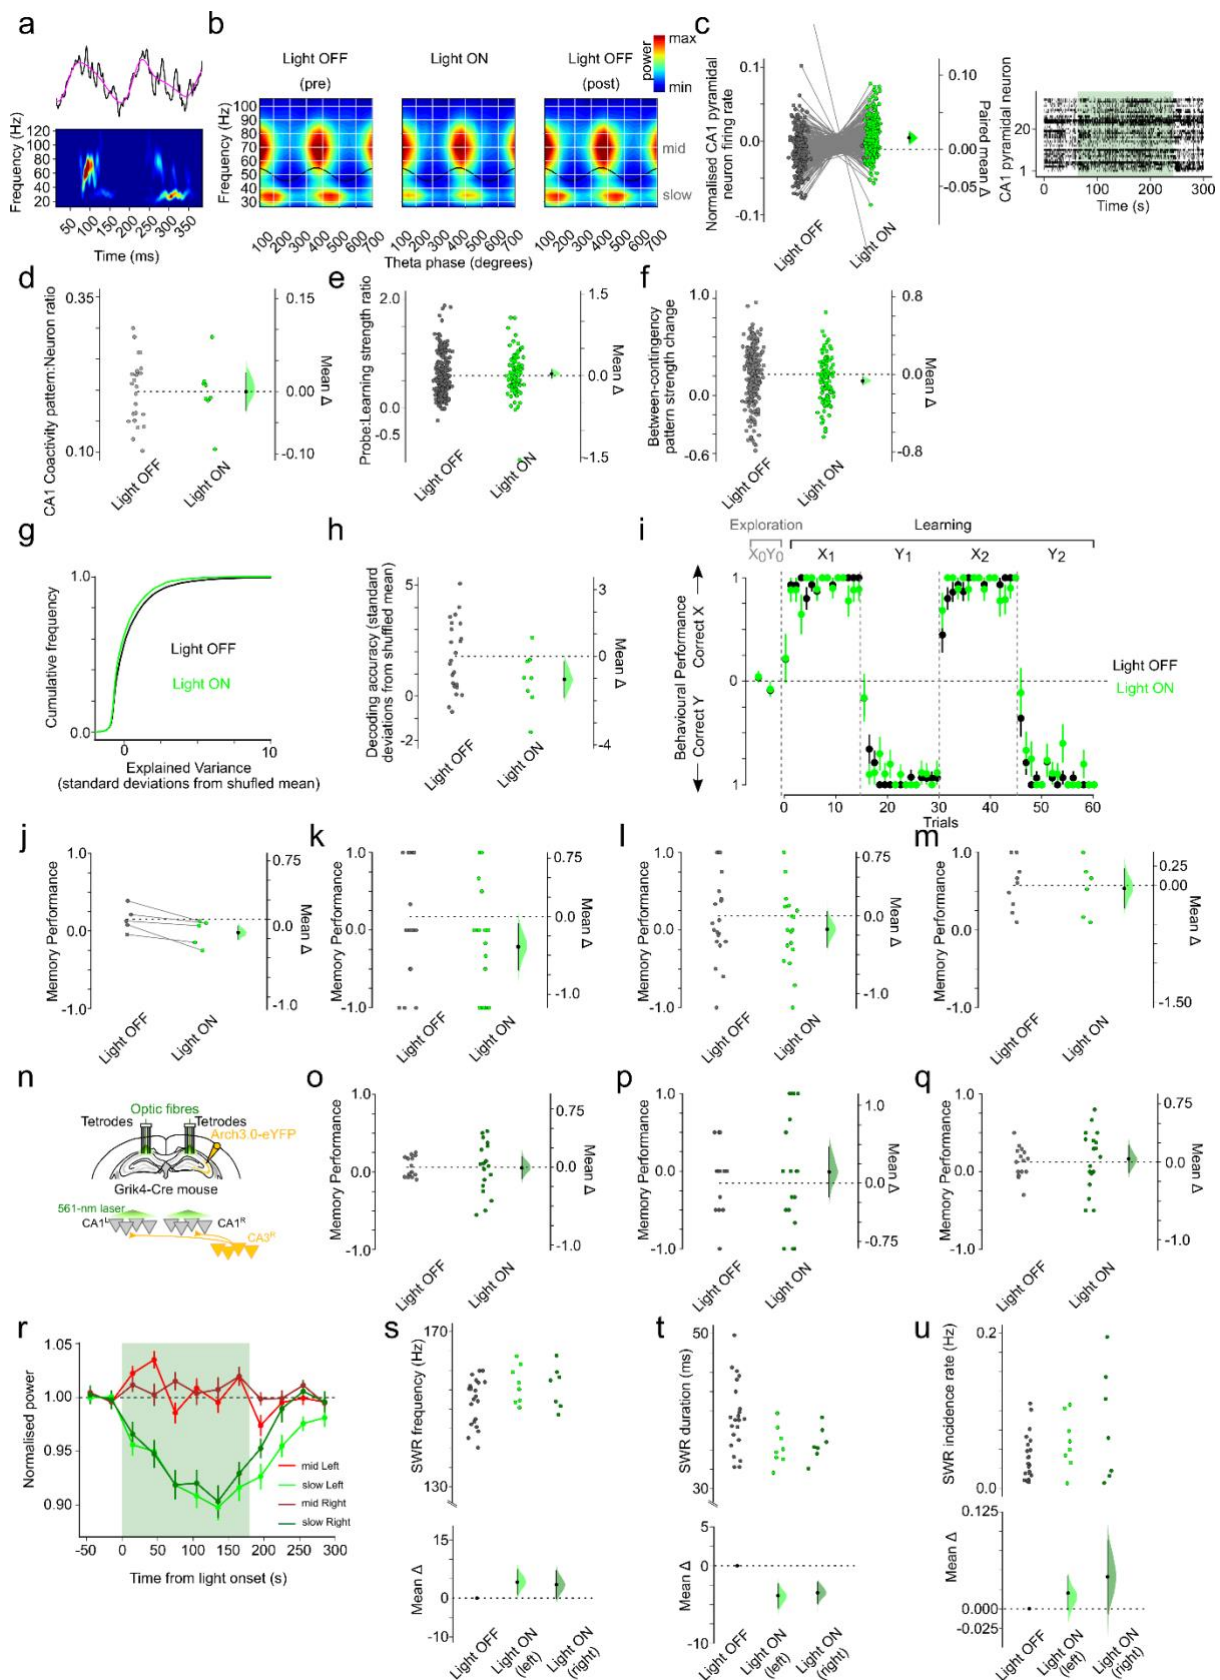
Example pattern activation maps and corresponding place maps of pattern member neurons for **a)** a contingency-invariant and a concomitantly recorded **b)** contingency-discriminating coactivity pattern across all sessions. Note that the right most member of the contingency-invariant pattern is also a member of the contingency-discriminating pattern. Further examples of coactivity pattern and strength maps and member rate maps for **c)** contingency invariant and **d)** contingency discriminating patterns. Maps are shown for the session in which these patterns were detected. Maximum firing rate (in Hz) or maximum coactivity strength (AU) are shown above each firing rate map or pattern strength map, respectively.



Extended Data Fig. 7. Spatial firing properties of coactivity pattern members.

Spatial coherence of contingency-invariant pattern members is higher than that of contingency-discriminating pattern members **a**) in the learning (Mean spatial coherence: contingency-invariant: 0.79 ± 0.01 , contingency-discriminating: 0.66 ± 0.02) and **b**) in the control enclosures (Mean spatial coherence: contingency-invariant: 0.71 ± 0.02 , contingency-discriminating: 0.60 ± 0.02). **c**) Cumulative distribution of spatial firing field numbers for contingency-discriminating and contingency-invariant pattern members. (Mean field number: contingency-invariant: 1.59 ± 0.07 , contingency-discriminating: 1.79 ± 0.09 ; Kolmogorov-Smirnov test (two-sided): $D=0.15$, $P=0.08$). Member neuron firing fields are less spatially overlapping for contingency-discriminating than contingency-invariant patterns using **d**) 1 standard deviation (Mean spatial correlation: contingency-invariant: 0.39 ± 0.01 , contingency-discriminating: 0.17 ± 0.01) or **e**) 3 standard deviations (Mean spatial correlation: contingency-

invariant: 0.64 ± 0.04 , contingency-discriminating: 0.46 ± 0.07) as weight thresholds for defining pattern members. f) Pairwise spatial correlations of contingency-discriminating pattern members are lower than those of contingency-invariant pattern members regardless of the degree of temporal correlation between the member neurons. Two-way ANOVA: main effect of pattern type ($F(1)=27.0$, $P=3.87 \times 10^{-7}$, $\eta^2=0.073$) and temporal correlation ($F(4)=9.3$, $P=4.12 \times 10^{-7}$, $\eta^2=0.102$). No pattern type: temporal correlation interaction ($F(4)=1.7$, $P=0.14$, $\eta^2=0.019$). g) Example coverage traces (grey) with overlaid spiking activity (dots) of a member of a contingency-invariant (left) and a member of a contingency-discriminating (right) coactivity pattern. Spikes during a co-activation event of a given pattern are marked in blue (contingency-invariant) or orange (contingency-discriminating), while the remaining spikes are marked in dark green. Spatial firing field of the member neuron is indicated by light green shading. h) Infield versus outfield co-activation score for member neurons of contingency-invariant and contingency-discriminating patterns (Mean score: contingency-invariant: 0.56 ± 0.04 , contingency-discriminating: 0.18 ± 0.05). i) Pairwise spatial correlations of high explained variance and low explained variance principal cell pairs. Mean spatial correlation: High explained variance pairs ($N=993$): 0.134 ± 0.010 , low explained variance pairs ($N=369$): 0.204 ± 0.014 ; Mann Whitney U test (two-sided): $U=155648.0$, $P=9.69 \times 10^{-6}$. j) Matrices showing mean spatial correlations of members of contingency invariant (left) and contingency-discriminating (right) patterns across all sessions. k) Spatial correlation of each contingency discriminating pattern member neuron across sessions of the same contingency or of opposite contingencies showing only member neurons with spatial coherence matching that of contingency-invariant pattern members (Mean spatial correlation: within-contingency 0.58 ± 0.02 , between-contingency: 0.72 ± 0.02). l) Spatial correlations between members of the same contingency-invariant (left) or contingency-discriminating (right) patterns across sessions. For both pattern types spatial correlations amongst pairs of neurons of the same coactivity patterns were higher during the learning stage than during the exploration stage further reflecting the development of these patterns with learning. Spatial correlations amongst members of the same contingency discriminating or those of contingency-invariant patterns were lowest in the control enclosure and the highest in the last learning sessions, confirming the enclosure-selectivity of hippocampal maps. Key to x-axis labels: first letter denotes contingency in which pattern was detected, subsequent letters denote session in which spatial maps of members were computed (e.g. X-Y2 are the spatial maps of members of coactivity patterns detected in contingency X, plotted in session Y2; i.e. second learning session of contingency Y). Mean spatial correlation: contingency-invariant: X-X2 & Y-Y2 (pooled): 0.605 ± 0.015 , X-Y2 & Y-X2: 0.543 ± 0.018 , X-X1 & Y-Y1: 0.438 ± 0.021 , X-X0 & Y-Y0: 0.230 ± 0.023 , X-Control & Y-Control: 0.114 ± 0.021 ; contingency-discriminating: X-X2 & Y-Y2: 0.297 ± 0.027 , X-Y2 & Y-X2: 0.200 ± 0.025 , X-X1 & Y-Y1: 0.191 ± 0.026 , X-X0 & Y-Y0: 0.086 ± 0.023 , X-Control & Y-Control: 0.023 ± 0.024 .

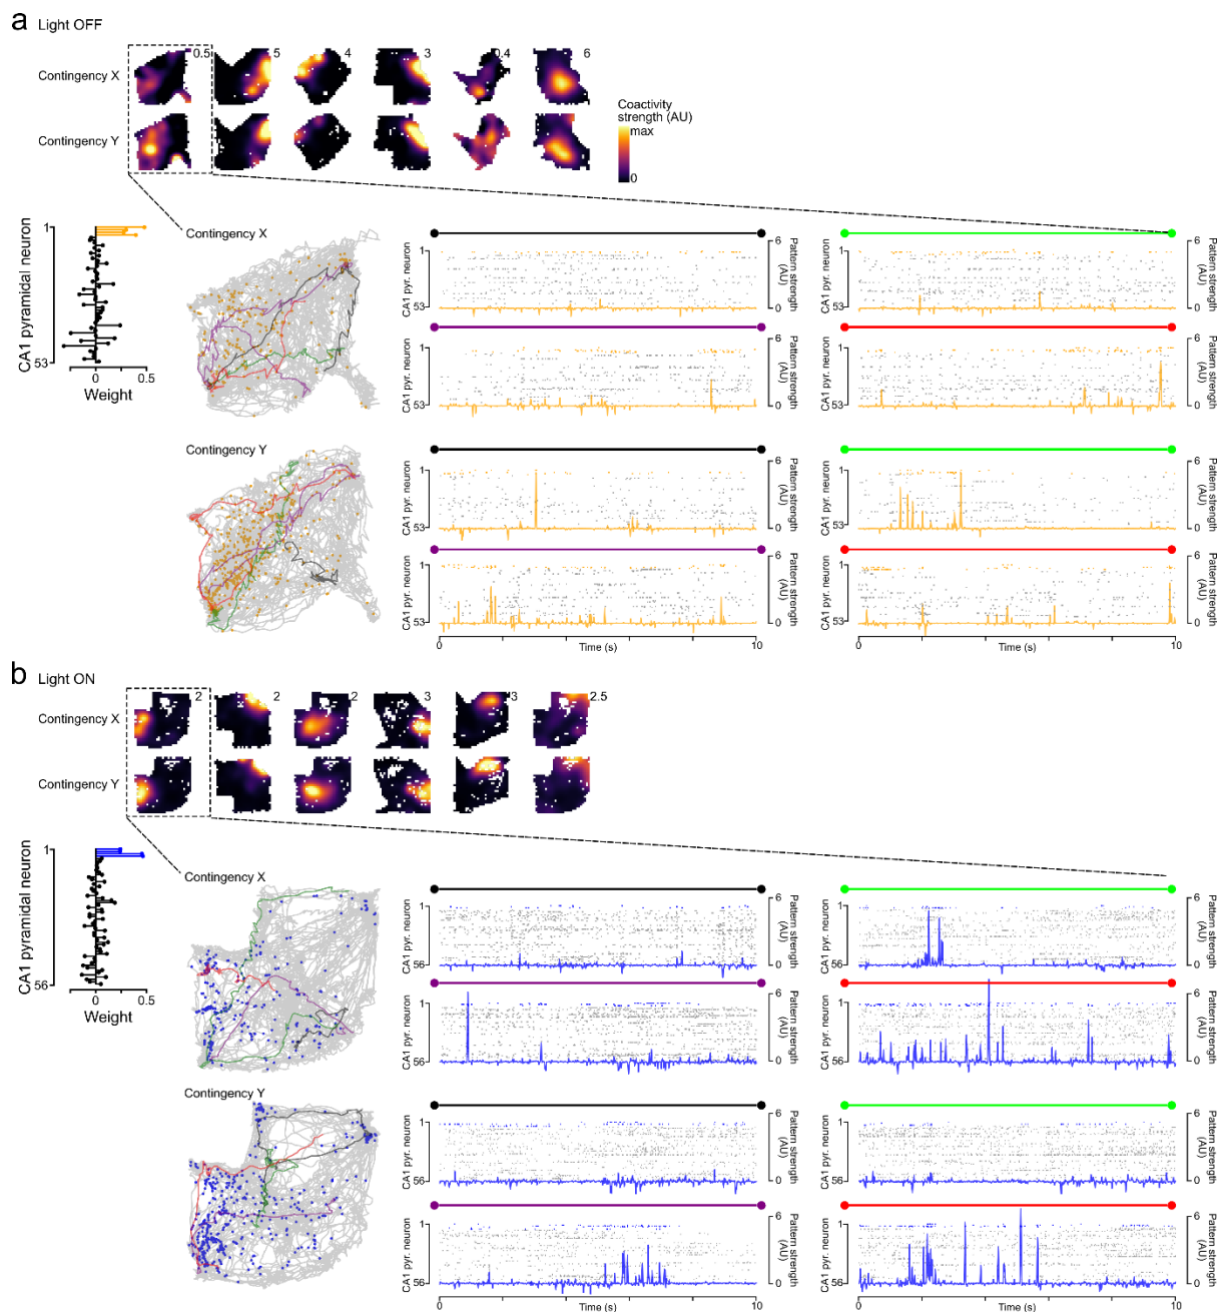


Extended Data Fig. 8. Behavioural and neural effects of silencing left or right hemisphere originating CA3-CA1 inputs.

a) Example LFP trace containing theta-nested mid and slow gamma oscillations (top; raw trace and theta component as black and magenta traces, respectively) along with its time-

frequency representation (bottom) b) Example of the selective effect of CA3^L→CA1 input suppression on the slow but not the mid gamma oscillations. Shown are Hilbert-spectra as a function of ongoing theta phase for pre, during and post light delivery in a representative session (colours represent the same scale in all three panels). Theta cycles were subsampled to maintain instantaneous speed distributions across panels. c) Firing rate of CA1 principal neurons is increased by light delivery (Mean normalized (z-scored) firing rate: light OFF epochs: -0.011 ± 0.002 , light ON epochs (1 minute after light onset): 0.004 ± 0.002 ; *right*, example raster plot during light ON period for one light ON epoch in a single recording day). d) The ratio of detected coactivity patterns to CA1 principal neurons is unaltered by CA3^L→CA1 input suppression (Mean pattern-to-neuron ratio: Light OFF days: 0.20 ± 0.01 , Light ON days: 0.20 ± 0.02). e) Reinstatement strength of all coactivity patterns is unaltered by CA3^L→CA1 input suppression (Mean probe:learning pattern strength ratio: Light OFF: 0.59 ± 0.02 , Light ON: 0.63 ± 0.04). Results in panels e and d show that input suppression does not alter the overall organisation of CA1 neurons into coactivity patterns nor the cross-session stability of such coactivity. f) The strength of coactivity patterns detected in the CA1 under CA3^L→CA1 input suppression is less sensitive to contingencies compared to light OFF days (Mean pattern strength change across contingencies: Light OFF days: 0.22 ± 0.01 , Light ON days: 0.15 ± 0.02). g) Explained variance for contingency is higher in light OFF days compared to days with CA3^L→CA1 input suppression. Mean normalised explained variance (standard deviations from mean): Light OFF days: 0.36 ± 0.01 (N=19852 neuron pairs), Light ON days (N= 5696 neuron pairs): 0.10 ± 0.02 ; Mann Whitney U test (two-sided): $U=51962023.0$, $P=9.36 \times 10^{-21}$. h) CA3^L→CA1 input suppression impairs Gaussian naïve Bayesian decoding accuracy using short-timescale (25ms) correlations (Mean normalised decoding accuracy (standard deviations from mean): Light OFF days (N=23 days): 1.80 ± 0.33 , Light ON days (N=8 days): 0.75 ± 0.45). i) CA3^L→CA1 input suppression does not impair performance during learning trials. Mean performance: Light OFF: 0.90 ± 0.02 (n=31 days), Light ON: 0.86 ± 0.02 (n=20 days); Mann Whitney U test (two-sided): $U=240.0$, $P=0.09$. j) Comparison of mean probe performance on light OFF and light ON (CA3^L→CA1 input suppression) days averaged per animal. (Mean performance: Light OFF days: 0.15 ± 0.07 , Light ON days: -0.02 ± 0.08). Effect of CA3^L→CA1 input suppression on performance on k) the first trial following a switch in LED displays (“switch” trials; Mean performance: Light OFF days: 0.18 ± 0.17 , Light ON days: -0.21 ± 0.16) and on l) subsequent trials (“non-switch” trials: Light OFF days: 0.18 ± 0.12 , Light ON days: 0.01 ± 0.11). m) Suppressing CA3^L inputs to CA1 during learning does not impair memory performance in probe trials when each LED set signals the same contingency (same dispenser-sucrose and dispenser-quinine pairing) throughout all learning sessions (“One-contingency training days”; Mean performance: Light OFF days: 0.57 ± 0.11 , Light ON days: 0.54 ± 0.14). n) Schematic of CA3^R→CA1 optogenetic suppression protocol. CA3^R neurons were transduced with Archaelrhodopsin 3.0 in Grik4-Cre mice (n=6) and their axonal projections in the CA1 targeted bilaterally during learning trials with yellow 561nm-light delivery from implanted optic fibres. CA3^R→CA1 input suppression during learning of the two-contingency task does not impair performance in probe trials, when taking o) all (Mean performance: Light OFF: 0.06 ± 0.03 , Light ON: 0.05 ± 0.08), p) switch (Mean performance: Light OFF: -0.16 ± 0.12 , Light ON: -0.01 ± 0.17) or q) non-switch trials (Mean performance: Light OFF: 0.12 ± 0.05 , Light ON: 0.17 ± 0.09). r) Suppression of CA3^R→CA1 input reduces the power of theta-nested slow gamma oscillations to a similar extent as with CA3^L→CA1 input suppression without affecting mid gamma oscillations. Two-way repeated measures ANOVA: Slow gamma: Main effect of light ($F(1)=64.2$, $P>0.001$, $\eta^2=0.592$), no main effect of CA3 hemisphere ($F(1)=0.571$, $P=0.457$, $\eta^2=0.005$) on normalised gamma power; Mid gamma: No main effect of light ($F(1)=1.22$, $P=0.281$, $\eta^2=0.029$), no main effect of CA3 hemisphere

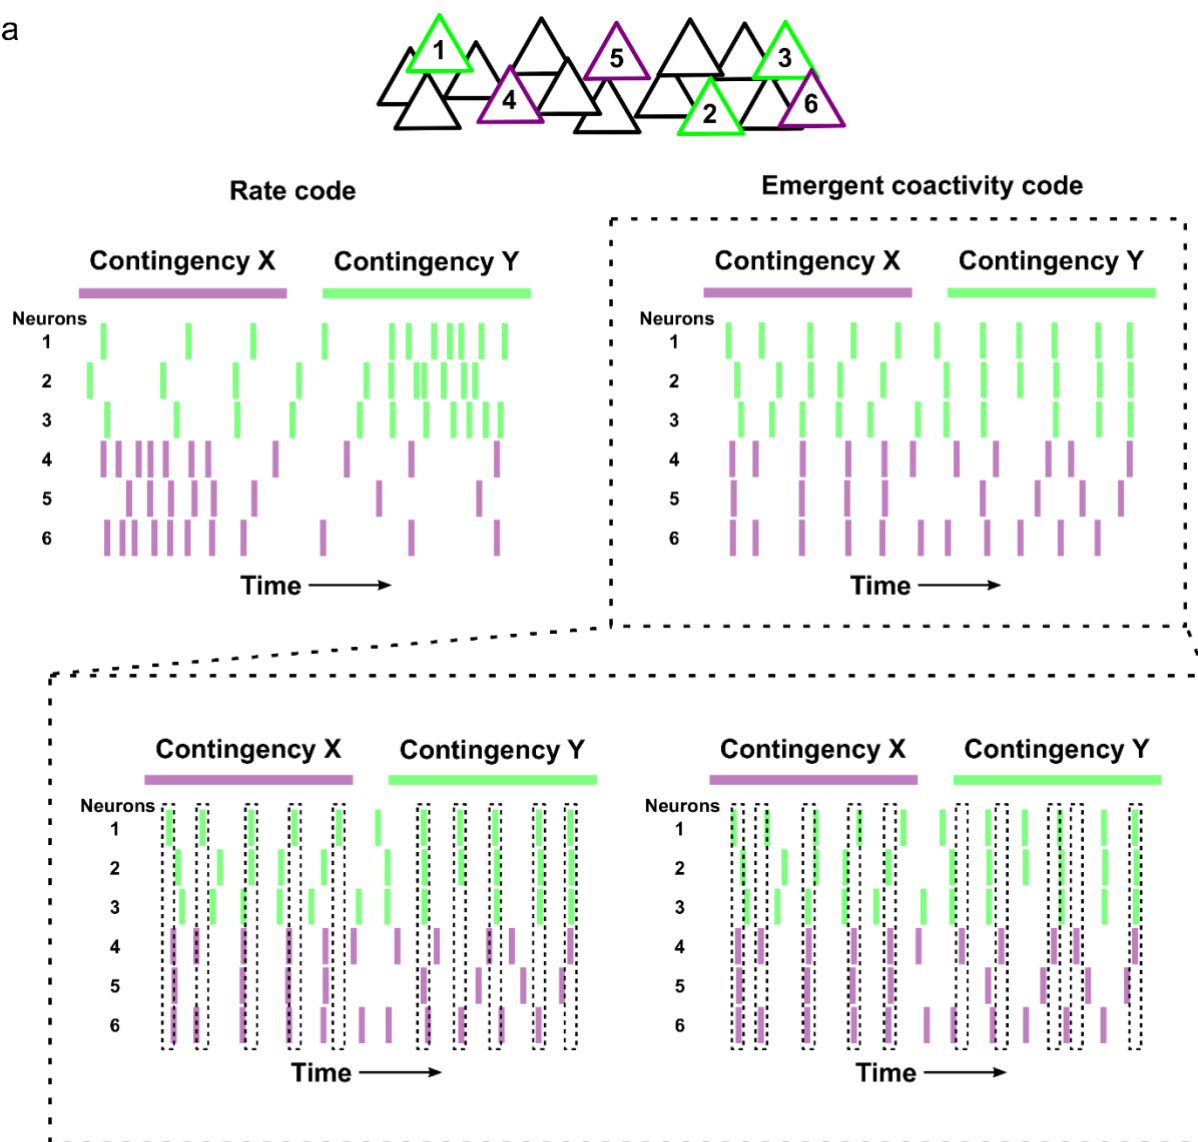
($F=0.226$, $P=0.639$, $\eta^2=0.005$) on normalised gamma power. s) SWR frequency increases with suppression of either $CA3^L \rightarrow CA1$ or $CA3^R \rightarrow CA1$ inputs (mean frequency: Light OFF days: 152 ± 1 Hz, Light ON days (left): 156 ± 2 Hz, Light ON days (right): 155 ± 2 Hz. t) Awake sharp-wave ripple (SWR) duration is reduced by suppression of either $CA3^L \rightarrow CA1$ or $CA3^R \rightarrow CA1$ inputs (mean duration: Light OFF days: 39 ± 1 ms, Light ON days (left): 35 ± 1 ms, Light ON days (right): 36 ± 1 ms). This reduction is therefore not sufficient to explain the selective impairment of memory performance after suppressing $CA3^L \rightarrow CA1$ but not $CA3^R \rightarrow CA1$ inputs. u) Incidence rates of awake SWRs during suppression of either $CA3^L \rightarrow CA1$ or $CA3^R \rightarrow CA1$ inputs (mean incidence rate: Light OFF days: 0.039 ± 0.006 Hz, Light ON (left) days: 0.059 ± 0.012 Hz, Light ON (right) days: 0.080 ± 0.027 Hz). We did not observe a reduction in awake SWR incidence rates, unlike a previous study using bilateral silencing of CA3 neurons in rats⁶³. Possible co-occurrence of SWR generating processes in the CA3 across hemispheres may explain why unilateral silencing does not impair the incidence rate of CA1 SWRs. Nevertheless, the reduction in SWR duration seen when silencing unilateral CA3 inputs from either hemisphere suggests that input from the CA3 on both hemispheres is needed for the full expression of a given CA1 SWR.



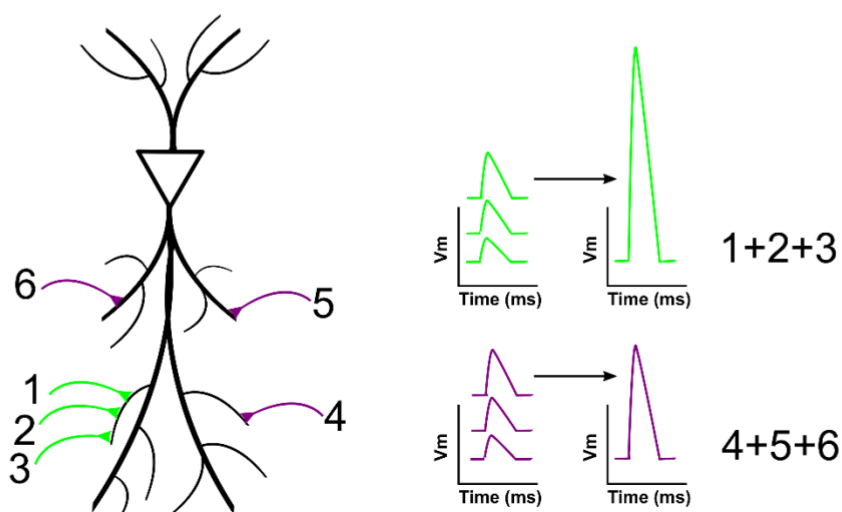
Extended Data Fig. 9. Example coactivation maps and raster plots of principal neuron activity with or without left CA3-CA1 silencing.

Example coactivity patterns during a) light OFF and b) CA3^L→CA1 input suppression (light ON) days across both learning contingencies (sessions X2 and Y2). Top of each panel depicts strength maps for several representative coactivity patterns, while below example coactivations for the left most pattern are shown in more detail. All coactivations (defined as coactivity strength above 2 standard deviation of mean strength in preferred contingency; displayed as coloured dots) are superimposed on coverage maps in each contingency (bottom left). Raster plots show the time-courses of neuronal firing (members color-coded orange or blue to denote contingency discriminating or contingency invariant pattern) and coactivation strengths for four separate paths (color-coded) across each contingency (bottom right) are plotted.

a



b



Extended Data Fig. 10. Schematic representation of emergent coactivity coding and a potential single-neuron reading mechanism.

a) *Top*: Schematics contrasting a hypothetical rate code (Adrian, 1926) (left) and emergent coactivity code (right) for disambiguating contingencies. *Bottom*: Emergent coactivity code for contingencies with temporal windows indicated by dashed rectangles aligned to spikes from neuron 1 (left) or neuron 4 (right) showing that neurons 1, 2 and 3 are more coactive for contingency *Y* while neurons 4, 5 and 6 are more coactive for contingency *X*. **b)** A hypothetical “reader” neuron can disambiguate distinct patterns of coactivity, for example by supralinear summation of one set of coactive inputs (e.g. from upstream neurons 1, 2 and 3), but only linear/sublinear summation of another (from upstream neurons 4, 5 and 6). Such non-linearities could result from the preferential activation of voltage-gated dendritic conductances, for example through clustering of synaptic inputs on dendritic branches (Stuart and Spruston, 2015). The membrane time-constant (~10-30 ms in forebrain pyramidal neurons (Koch et al., 1996)) constrains a reader neuron’s integration time-window, and hence this mechanism is particularly suited for short-timescale coactivity. Note that the converse may be true for another reader neuron, with inputs from neurons 4, 5, and 6 preferentially exhibiting supralinear summation and hence preferentially driving activity in this alternative neuron. Vm: membrane potential.