

1 **Evolution of predictive memory in the hippocampus**

2 Adam M. P. Miller¹, Alex D. Jacob^{1,2}, Adam I. Ramsaran^{1,2}, Mitchell L. De Snoo^{1,4}, Sheena A. Josselyn¹⁻⁵,
3 Paul W. Frankland^{1-4,6*}

4 Affiliations:

5
6 1 Program in Neurosciences and Mental Health, The Hospital for Sick Children; Toronto, Canada
7 2 Department of Psychology, University of Toronto; Toronto, Canada
8 3 Department of Physiology, University of Toronto; Toronto, Canada
9 4 Institute of Medical Sciences, University of Toronto; Toronto, Canada
10 5 Brain, Mind, & Consciousness Program, Canadian Institute for Advanced Research; Toronto, Canada
11 6 Child & Brain Development Program, Canadian Institute for Advanced Research; Toronto, Canada

12
13 *Corresponding author. Email: paul.frankland@sickkids.ca
14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29 **Summary**

30 The brain organizes experiences into memories that can be used to guide future behavior. Hippocampal
31 CA1 population activity may reflect the retrieval of predictive models that contain information about
32 future events, but little is known about how these kinds of memories develop with experience. We
33 trained mice on a series of tone discrimination problems with or without a common statistical structure
34 to observe how memories are formed and updated during learning. Mice that learned structured
35 problems integrated their experiences into a predictive model that contained the solutions to upcoming
36 novel problems. Retrieving the model during learning improved discrimination accuracy and facilitated
37 learning by decreasing the amount of new information that needed to be acquired. Using calcium
38 imaging to track the activity of thousands of CA1 neurons during learning on this task, we observed the
39 emergence of a persistent hippocampal ensemble at the same time that mice formed a predictive model
40 of their environment. This ensemble was reactivated during training and incorporated new neuronal
41 activity patterns from each training problem. Interestingly, the degree to which mice reactivated the
42 ensemble was related to how well their model predicted the content of the current problem, ensuring
43 that the model was only updated with congruent information. In contrast, mice trained on unstructured
44 problems did not form a predictive model or engage a persistent ensemble. These results show how
45 hippocampal activity supports building predictive models by organizing newly learned information
46 according to its congruence with existing memories.

47 **Keywords:** hippocampus, learning, memory, retrieval, prediction, integration, context, schema, latent
48 state, inference, cognition

49

50

51

52

53

54

55

56

57 **Introduction**

58 All animals use information from their past to guide their present behavior. Memories derived
59 from multiple experiences may be particularly useful for supporting efficient behavior by providing
60 insight into the rules that govern environments. This type of memory, often described as a predictive
61 model, is considered a fundamental feature of high-level cognitive functions such as creativity and
62 intelligence (Kumaran et al., 2016; Tervo et al., 2016; Behrens et al., 2018; Momennejad, 2020; Morton
63 and Preston, 2021). However, little is understood about how predictive models are formed in the brain
64 and used to guide behavior.

65 The hippocampus has been implicated in supporting predictive models across several species
66 (Zeithamova et al., 2012; Pudhiyidath et al., 2022; Brunec and Momennejad, 2022; Vikbladh et al., 2019;
67 Schapiro et al., 2016; Knudsen and Wallis, 2021; Baraduc et al., 2019; Bulkin et al., 2016; McKenzie et al.,
68 2014; Nieh et al., 2021; Tse et al., 2007). One possibility is that the hippocampus supports the
69 development of predictive models by repeatedly retrieving and updating memories with new
70 information (Lee, 2009; Schlichting and Frankland, 2017; Gisquet-Verrier and Riccio, 2018; Mack et al.,
71 2018; Mau et al., 2020). In this framework, animals may infer a predictive model of their environment
72 from a subset of possible observations and then interpret subsequent experiences in terms of this
73 model (i.e. latent state theory; Fuhs and Touretzky, 2007; Gershman and Niv, 2010; Niv, 2019; Redish et
74 al., 2007). When an existing predictive model can explain the animal's current experience, it is retrieved
75 along with a corresponding hippocampal activity pattern (i.e. map or context code; Sanders et al., 2020;
76 Whittington et al., 2020). New learning that is consistent with the model may then be integrated into
77 the model, whereas new learning that is not consistent with any existing model may prompt the
78 formation of a new model.

79 Although it is challenging to study mental representations in the brain (Tervo et al., 2016),
80 several aspects of this idea have been examined. For instance, it has been shown that the hippocampus
81 encodes experience in state spaces (Wood et al., 2000; Nieh et al., 2021; Sun et al., 2020; Samborska et
82 al., 2021) and that pre-existing memories (McKenzie et al., 2014; Tse et al., 2007; Cai et al., 2016) and
83 neural structure (Epsztein et al., 2011; Dragoi and Tonegawa, 2011; Sadtler et al., 2014; McKenzie et al.,
84 2021) can affect learning. However, whether the interplay between new learning and existing memory
85 accrued over many experiences gives rise to a predictive model that guides behavior has not been
86 assessed directly, especially at the neural level. Here, we developed a new behavioral protocol that
87 repeatedly probes the task representation of a mouse as it learns a series of unique problems. At the

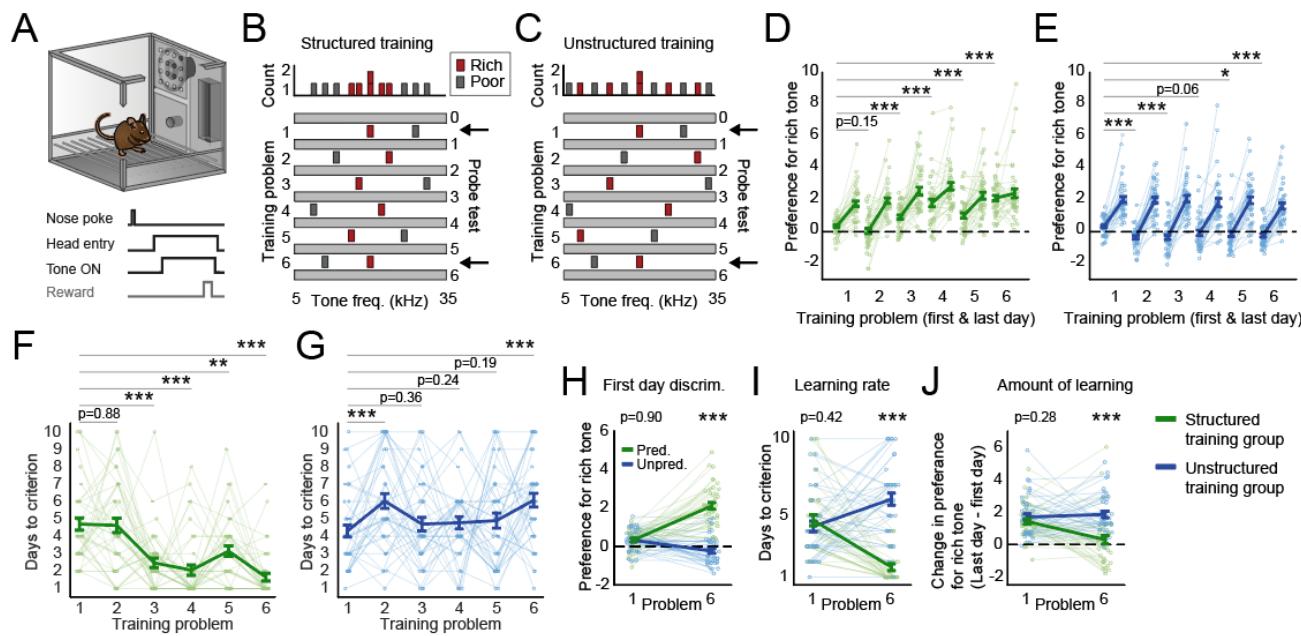
88 behavioral level, we found that mice formed a predictive model that accurately predicted the solution to
89 novel problems by integrating their memories of past problems, and that mice retrieved and updated
90 their model with new learning only when what they learned matched their predictions. At the neural
91 level, we used calcium imaging to track changes in hippocampal CA1 ensemble activity as mice formed
92 and retrieved a predictive model. We found that mice formed a persistent hippocampal ensemble by
93 incorporating neurons associated with prior training problems, and that they reactivated this ensemble
94 during new learning when they accurately predicted the solutions to the problems.

95 **Results**

96 Mice show superior learning of structured experiences

97 We trained 80 mice to perform a novel auditory discrimination and foraging task to examine
98 how memory-based predictions are formed and updated. In this task, mice hear tones and learn to wait
99 longer in response to certain frequencies in order to obtain chocolate milk rewards (**Figure 1A**;
100 **Supplemental figure 1**). Training was divided into a series of six unique training problems, and mice only
101 advanced to the next problem after successfully learning the current one. Half of the mice (n=40)
102 received a Structured training protocol where the highly rewarded (i.e., rich) tones from every problem
103 fell within a continuous band of frequencies (**Figure 1B**), while the other mice (n = 40) received an
104 Unstructured training protocol where the rich and poor tones were more evenly distributed (**Figure 1C**).
105 Consistent with previous work on memory transfer (Harlow, 1949), the Structured training group
106 showed increased preference for the rich tone on the first day of training of later problems (one-way
107 repeated measures ANOVA, $F(5, 195) = 23.99, p < 0.001$; **Figure 1D**), as well as more rapid learning over
108 the course of training ($F(5, 195) = 18.45, p < 0.001$; **Figure 1F**), even though every problem was novel,
109 indicating that early training experiences enhanced the ability of mice in the Structured training group to
110 learn novel problems. By contrast, the Unstructured training group reliably preferred the poor tone ($F(5,$
111 $195) = 3.76, p < 0.01$; **Figure 1E**) and often showed slowed learning ($F(5, 195) = 4.59, p < 0.001$; **Figure**
112 **1G**), indicating memory interference. Direct comparisons between the groups on the first and last
113 problems (which were the same in both training protocols; i.e., the common problems) confirmed that
114 only the Structured training group improved their performance on the first day of training problems
115 (two-way repeated measures ANOVA, group X problem interaction, $F(1, 78) = 93.89, p < 0.001$; between
116 groups on problem 1, $t(78) = 0.13, p = 0.90$; between groups on problem 6, $t(78) = 11.43, p < 0.001$;
117 **Figure 1H**) and their rate of learning (group X problem interaction, $F(1, 78) = 63.18, p < 0.001$; between
118 groups on problem 1, $t(78) = 0.81, p = 0.42$; between groups on problem 6, $t(78) = 9.80, p < 0.001$;

119 **Figure 1I).** Lastly, the Structured training group did not require as much training to reach criterion
 120 (defined as the difference between the preference for the rich tone on the first and last day; two-way
 121 repeated measures ANOVA, group X problem interaction, $F(1, 78) = 10.37$, $p < 0.01$; between groups on
 122 problem 1, $t(78) = 1.08$, $p = 0.28$; between groups on problem 6, $t(78) = 4.76$, $p < 0.001$; Structured
 123 training group compared to zero, $t(39) = 1.13$, $p = 0.26$; Unstructured training group compared to zero,
 124 $t(39) = 9.14$, $p < 0.001$; **Figure 1J**), indicating that the last problem contained less new information for
 125 the Structured training group despite neither group having been trained on that problem before.



126 **Figure 1.** Efficient learning of structured versus unstructured problems. **(A)** Mice were trained and tested in an
 127 operant chamber with a nose-poke port, reward-delivery hopper, and a speaker that delivered different tones.
 128 Each trial consisted of a nose-poke followed by head entry into the hopper. Upon head entry, a tone played
 129 followed by a potential reward delivery. **(B)** Two groups of mice were trained on a series of tone discrimination
 130 problems. In the Structured training group, all of the rich tones fell within a continuous band of frequencies. **(C)** In
 131 the unstructured group, the rich tones were more evenly distributed. The first and last problem were the same for
 132 both groups (black arrows). **(D)** Preference for the rich tone is shown in terms of the standardized mean difference
 133 between how long mice waited in response to the rich tone and the poor tone. Preferences are shown for the first
 134 and last day of training on every problem. Mice were trained to the same criterion (longer wait times on rich tone
 135 trials than poor tone trials as determined by a significant T-test, alpha level = 0.01) on every problem and then
 136 were given one additional day of training (i.e., the last day). The Structured training group showed better
 137 discrimination on the first day of problems later in learning compared to the first problem. **(E)** The Unstructured
 138 training group did not improve first-day performance and typically showed significantly worse performance. **(F)**
 139 The number of days required to reach criterion. The Structured training group learned new problems faster later in
 140 learning. **(G)** The learning rate of the Unstructured training group failed to improve. Direct comparisons between
 141 the two groups on the first and last problem revealed that the Structured training group **(H)** showed higher initial
 142 discrimination on the last training problem, **(I)** learned the last problem in fewer days, and **(J)** learned less new
 143 information on the last problem compared to the Unstructured training group.

144 Mice integrate structured experiences to form a stable, generalized model of the environment

145 We hypothesized that the Structured training group performed better on later discrimination
146 problems by integrating its early experiences into a predictive model of the environment that supported
147 new learning (Tolman, 1948; Tse et al., 2007). If this is true, then the Structured training group should
148 develop tone preferences (increased wait times) that match the statistics of the full Structured training
149 set (and not simply the most recent problem; **Figure 1B**). Furthermore, if this process depends on
150 updating a predictive model with new information, then tone preferences should change (be updated)
151 when mice encounter unpredicted problems but should remain stable when the mice encounter
152 problems that are already predicted by the model (and therefore do not require additional learning).

153 To observe how tone preferences changed during Structured or Unstructured training, we
154 examined the behavioral responses of both groups during probe tests administered before and after
155 each training problem (**Figure 1B, C**). During probe tests, 41 unique tones (min frequency = 5kHz, max
156 frequency = 35kHz; all tones, $p(\text{reward}) = 0.5$) were presented in random order to examine the length of
157 time mice would wait in response to a range of tones, including tones that had not been trained yet.
158 Early in training, both groups generalized beyond the two most recent training tones, preferring all
159 tones beyond the recent rich tone to all tones beyond the recent poor tone (e.g., probes 1 and 2; see
160 Purtle, 1973). Later in training, the Structured training group began to selectively prefer tones in the
161 center of the auditory range (including novel tones that were not previously trained), consistent with
162 the statistics of the training set (**Figure 2A, B**). This suggests that the Structured training group
163 integrated its training experiences into an accurate model of the relationship between tone and reward.

164 To quantify the degree to which the model changed after each training experience, we assessed
165 the similarity of tone preferences before and after each problem. The Structured training group
166 developed stable tone preferences during the second half of training (linear mixed-effects (LME) model
167 with fixed effects for first probe, second probe, and group, and a random effect for subject, first probe X
168 second probe X group interaction, $t(3352) = 9.16$, $p < 0.001$; **Figure 2C**; the Structured training group,
169 one-way repeated measures ANOVA, $F(5, 195) = 27.79$, $p < 0.001$; **Figure 2D**). In contrast, the tone
170 preferences of the Unstructured training group were either uncorrelated or anticorrelated ($F(5, 195) =$
171 7.27, $p < 0.001$; **Figure 2E**), indicating that this group did not form a stable model. Between-group
172 comparisons confirmed that the Structured training group developed a stable model while the
173 Unstructured training group continued to show learning-related changes (two-way repeated measures
174 ANOVA, group X stage interaction, $F(1, 78) = 61.86$, $p < 0.001$; Structured training group compared to

175 zero, $t(39) = 13.78$, $p < 0.001$; Unstructured training group compared to zero, $t(39) = 0.67$, $p = 0.50$;
 176 **Figure 2F).**

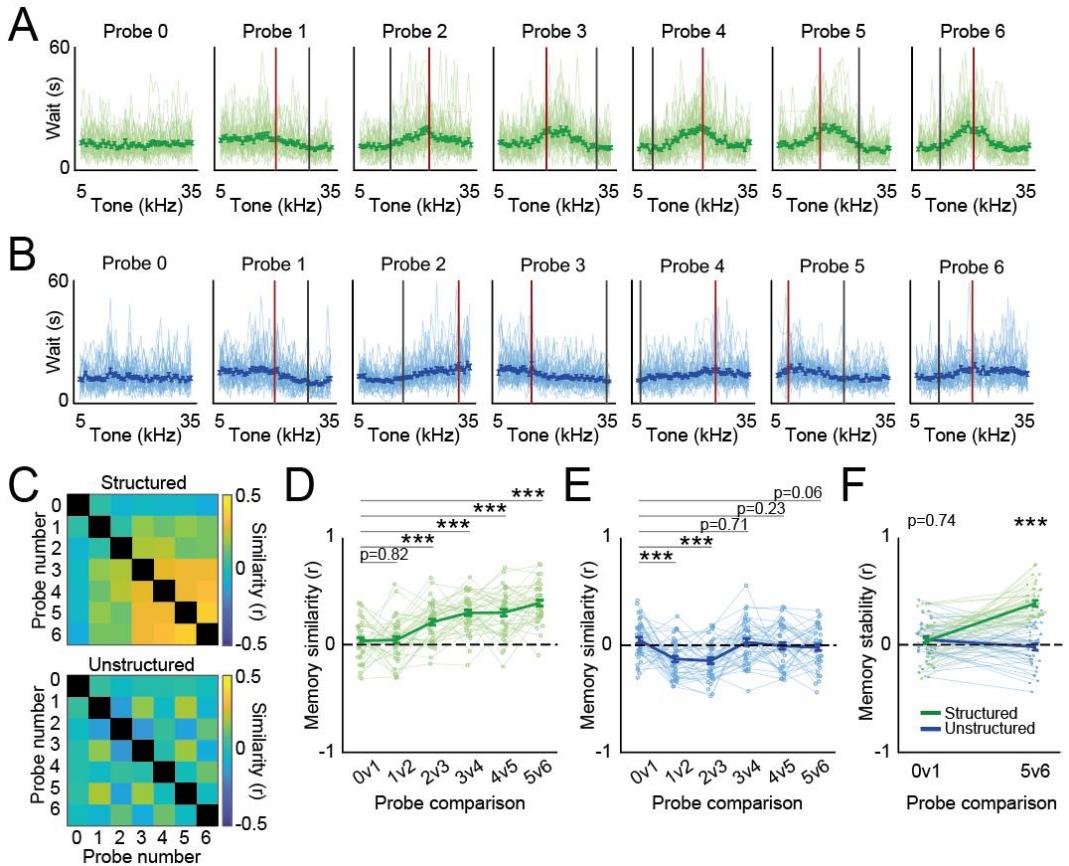


Figure 2.
 Integrating structured experiences into a generalized model. **(A)** The wait time responses to every tone are shown for every subject in every probe test of the Structured training group. Probe 0 is the pre-probe administered before any tone training. Subsequent probes were administered the day after completing the corresponding problem. The rich and poor tone

204 frequencies of the prior problem are shown as red and grey lines, respectively. **(B)** Same as A, but for the
 205 Unstructured training group showing no integration of new information into previous memories. **(C)** Correlation
 206 matrices showing the mean of all subject correlations between every pair of probe tests for the Unstructured
 207 training group (top) and Unstructured training group (bottom). **(D)** The correlation between every pair of
 208 sequential probe tests is shown for the Structured training group. **(E)** Same as D, but for the Unstructured training
 209 group. **(F)** The correlation between the probes before and after the two common problems (problem 1 and
 210 problem 6) are shown for both groups.

211 Generalized model predicts future discriminations and improves learning

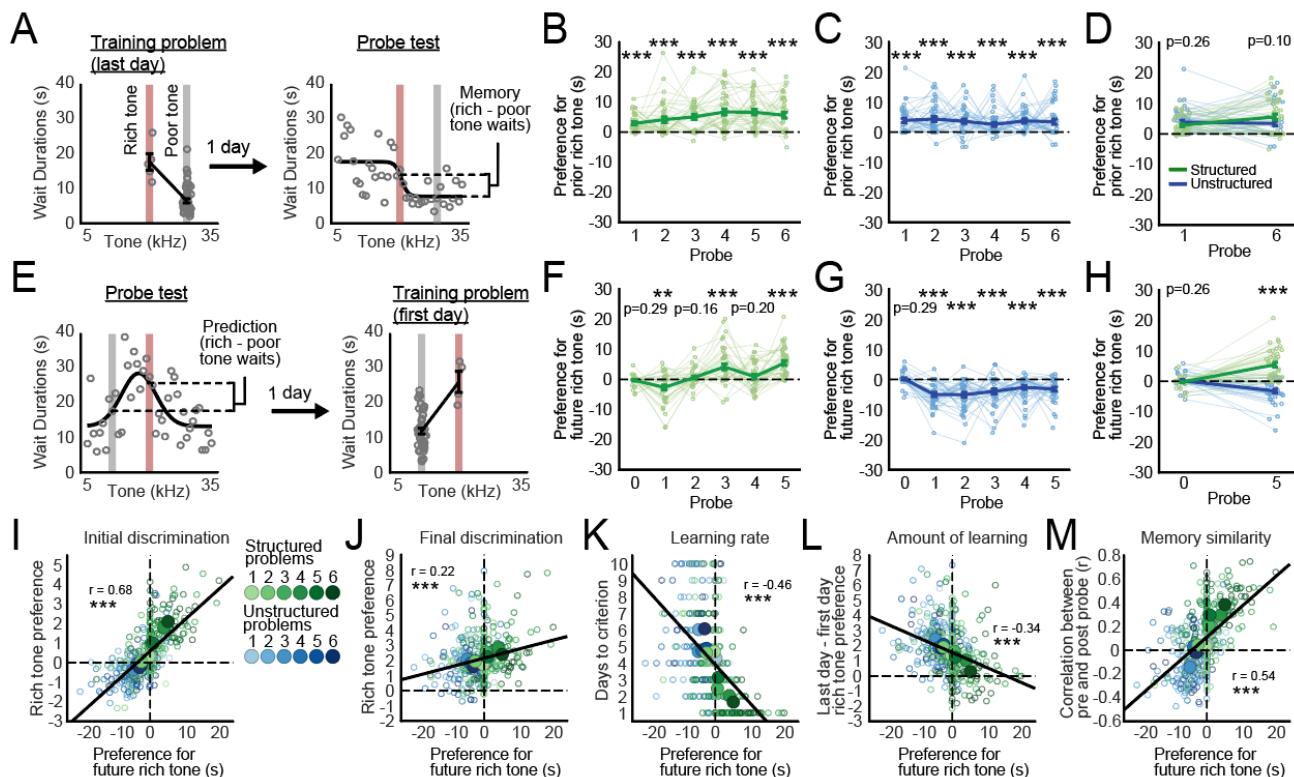
212 The above observations indicate that the Structured training group formed a stable model at the
 213 same time that they began to show efficient learning of novel discriminations. This raises the possibility
 214 that the model (integrated memories of past problems) contained information about future (not-yet
 215 trained) problems that facilitated learning on these problems. To quantify the information about
 216 discrimination problems contained in the model, we fit a curve to the tone preferences (wait times)
 217 observed during every probe test (see methods) and then measured the difference between how long
 218 the mouse waited in response to the rich and poor training tones used in the training problems. As a

219 proof of concept, we first asked whether the model expressed during probe tests correctly discriminated
220 between the rich and poor tones of the most recent training problem. Because both groups successfully
221 learned to prefer the rich tone to the poor tone over the course of every training problem (**Figure 1 D-E**),
222 we anticipated that the mice would similarly discriminate between these tones during the subsequent
223 probe test one day later. Consistent with this, mice showed a similar preference for the rich tone during
224 training problems and the probe tests the following day (**Figure 3A**; every Structured training group
225 session compared to zero, $p < 0.001$; **Figure 3B**; every Unstructured training group session compared to
226 zero, $p < 0.001$; **Figure 3C**; direct comparisons after the common problems, two-way repeated measures
227 ANOVA, group X probe interaction, $F(1, 78) = 4.35$, $p = 0.04$; at probe 1, $t(78) = 1.14$, $p = 0.26$; at probe 6,
228 $t(78) = 1.66$, $p = 0.10$; **Figure 3D**), indicating that the model expressed during the probe test contained
229 information about the discrimination learned during the preceding training problem.

230 To examine whether the model expressed during probe tests also contained information about
231 *future* problems, we asked whether mice waited longer in response to the upcoming rich tone than to
232 the upcoming poor tone. Similar to above, we found that the amount of time mice waited in response to
233 the training tones on the first day of a problem closely matched how long they waited in response to
234 those same tones on the probe test one day prior (**Figure 3E**), indicating that the mice were retrieving
235 the same model during training that they retrieved during the prior probe. However, only the Structured
236 training group developed a preference for future rich tones over future poor tones (all compared to
237 zero: probe 0, $t(39) = 1.07$, $p = 0.29$; probe 1, $t(39) = 3.32$, $p < 0.01$; probe 2, $t(39) = 1.42$, $p = 0.16$; probe
238 3, $t(39) = 4.21$, $p < 0.001$; probe 4, $t(39) = 1.31$, $p = 0.20$; probe 5, $t(39) = 6.87$, $p < 0.001$, **Figure 3F**). The
239 Unstructured training group never predicted upcoming problems accurately and was often inaccurate
240 (all compared to zero: probe 0, $t(39) = 0.62$, $p = 0.54$; probe 1, $t(39) = 7.24$, $p < 0.001$; probe 2, $t(39) =$
241 6.15 , $p < 0.001$; probe 3, $t(39) = 4.90$, $p < 0.001$; probe 4, $t(39) = 3.95$, $p < 0.001$; probe 5, $t(39) = 5.18$, p
242 < 0.001 , **Figure 3G**; direct comparisons on the probes preceding the common problems, two-way
243 repeated measures ANOVA, group X probe interaction, $F(1, 78) = 78.62$, $p < 0.001$; at probe 0, $t(78) =$
244 1.14 , $p = 0.26$; at probe 5, $t(78) = 8.61$, $p < 0.001$; **Figure 3H**). These findings indicate that the model
245 expressed during probe tests contained information that could be used to predict the solutions to
246 discrimination problems that had not yet been encountered (i.e., a predictive model).

247 If mice retrieve a predictive model during learning, then their learning might be affected by the
248 accuracy of the predictions. For example, an accurate model may reduce the amount of new
249 information that needs to be acquired and incorporated into memory, while inaccurate predictions

250 might interfere with performance. To test this, we asked how discrimination performance on the first
 251 day of training was related to the accuracy of the prediction measured during the preceding probe test.
 252 We found that accurate predictions were associated with greater preference for the rich tone on the
 253 first day of training ($r = 0.68$, $p < 0.001$; LME model with a fixed effect for prediction accuracy and a
 254 random effect for subject, $t(478) = 19.58$, $p < 0.001$; **Figure 3I**; last day, $r = 0.22$, $p < 0.001$, LME model,
 255 $t(478) = 5.00$, $p < 0.001$; **Figure 3J**) and faster learning (i.e. fewer days to criterion, $r = -0.46$, $p < 0.001$,
 256 LME, $t(478) = 11.22$, $p < 0.001$; **Figure 3K**), indicating that mice learned new problems more easily when
 257 they possessed a predictive model. Relatedly, accurate predictions were associated with *less* learning
 258 overall. Mice that retrieved accurate predictions showed a smaller change in their preference for the
 259 rich tone between the first and last day of training ($r = -0.34$, $p < 0.001$; LME, $t(478) = 7.61$, $p < 0.001$;
 260 **Figure 3L**), as well as a smaller change in their memory, as indicated by higher correlations between
 261 their responses on the probe tests before and after training ($r = 0.54$, $p < 0.001$; LME, $t(478) = 14.19$, $p <$
 262 0.001 ; **Figure 3M**). Together, these findings indicate that predictive models improve learning by reducing
 263 the amount of new information that must be learned.



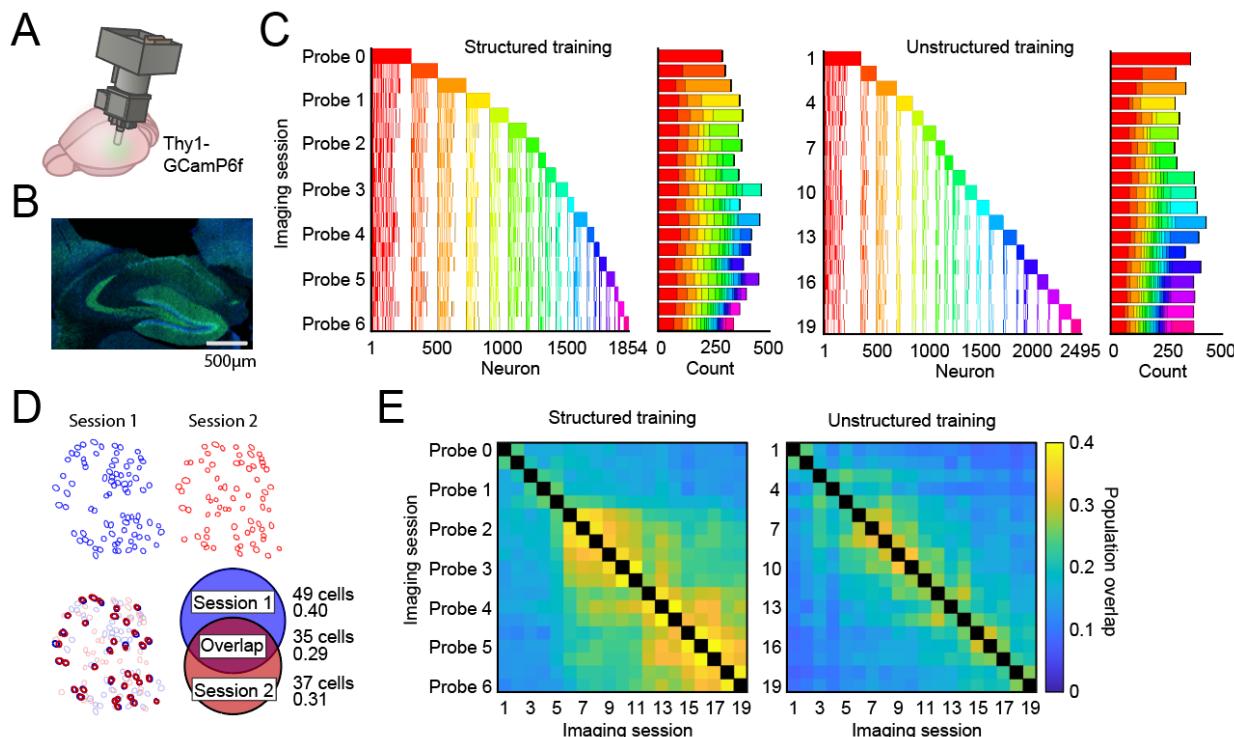
264 **Figure 3.** Predictive models guide new learning. **(A)** Memory for the reward values of the rich and poor tones used
 265 in the prior problem can be measured during the probe test one day after training. All wait times from two
 266 example sessions (the last day of a training problem and the probe test one day later) are shown from one mouse.
 267 Grey circles indicate the wait times on individual trials. We quantified the memory for the prior problem in terms

268 of the difference between the how long the mouse waited in response to the same tones during the subsequent
269 probe test. Wait times were determined by fitting a curve to the observed waits from every trial. **(B)** The
270 Structured training group and **(C)** the Unstructured training group showed significant memory for the prior
271 problem during every probe test. **(D)** The two groups also showed similar memory for each of the two common
272 training problems. **(E)** Predictions for the reward values of future rich and poor tones can be measured during the
273 probe test one day before the start of training. We measured predictions in the same way that we measured
274 memory, but with the rich and poor tones from the upcoming (and not the prior) problem. **(F)** The Structured
275 training group accurately predicted the solutions to some problems later in learning after initially making
276 inaccurate predictions. **(G)** The Unstructured training group showed inaccurate predictions throughout training.
277 **(H)** The Structured training group developed superior predictions with training compared to the Unstructured
278 training group. **(I)** The more accurate the prediction (measured during the preceding probe test), the stronger the
279 preference for the rich tone on the first day of training, defined as the standardized mean difference between the
280 wait times in response to the rich and poor tones. Large circles show condition means. **(J)** The more accurate the
281 prediction, the stronger the preference for the rich tone on the last day of training. **(K)** The more accurate the
282 prediction, the faster the subject learned the training problem. **(L)** The more accurate the prediction, the less new
283 information was learned during a training problem, defined as a smaller change in the preference for the rich tone
284 from the first to the last day of training. **(M)** More accurate predictions were associated with smaller learning-
285 related changes to the memory, defined in terms of the correlation between the behavioral responses on probe
286 tests before and after training (higher correlations indicate smaller changes).

287 Structured training reactivates CA1 ensembles active during prior learning

288 Our behavioral data indicate that mice retrieve memories of prior problems to support
289 predictions and new learning. The hippocampus plays a crucial role in memory retrieval by reactivating
290 ensembles (e.g., engrams) that were active during prior experiences (Liu et al., 2012; Goshen et al.,
291 2011; Tanaka et al., 2014). To examine how ensemble reactivation in the hippocampus might support
292 the formation of a predictive model in this task, we implanted GRIN lenses above the CA1 layer of the
293 hippocampus in 13 Thy1-GCaMP6f mice, and imaged 4349 unique neurons with custom miniature
294 microscopes (Jacob et al., 2018; **Figure 4A, B**) while mice performed the discrimination and foraging test
295 described above (n = 8 Structured training mice, n = 5 Unstructured training mice; **Supplemental figure**
296 **2**). Imaging data was collected during the seven probe tests and the first and last day of each of the six
297 training problems (19 sessions). To measure ensemble reactivation in our task, we registered active
298 neurons across sessions (*CellReg*, Sheintuch et al., 2017; **Figure 4C; Supplemental figure 3**) and then
299 quantified excess population overlap for every pair of sessions, defined as the number of neurons that
300 were active in both sessions divided by the number that were active in either session (**Figure 4D**) minus
301 the overlap expected due to the amount of time between imaging sessions (**Supplemental figure 4**). We
302 found that excess overlap increased with Structured training, but not Unstructured training (LME model
303 with fixed effects for sessions and group, and a random effect for subject, session X session X group
304 interaction, $t(4438) = 3.78$, $p < 0.001$; post-hoc LME models with fixed effects for sessions, and a random
305 effect for subject, Structured training group session X session interaction, $t(2732) = 6.09$, $p < 0.001$,
306 Unstructured training group session X session interaction, $t(1706) = 0.23$, $p = 0.81$; **Figure 4E**;

307 **Supplemental figure 4D, E**), indicating that Structured training increased the degree to which
308 hippocampal ensembles were reactivated during later sessions. Interestingly, ensemble reactivation
309 increased in the Structured training group at approximately the same time that these mice began to
310 form a predictive model, suggesting that neuronal reactivation may be a mechanism for integrating
311 experiences into a model (**Figure 2C-D**).



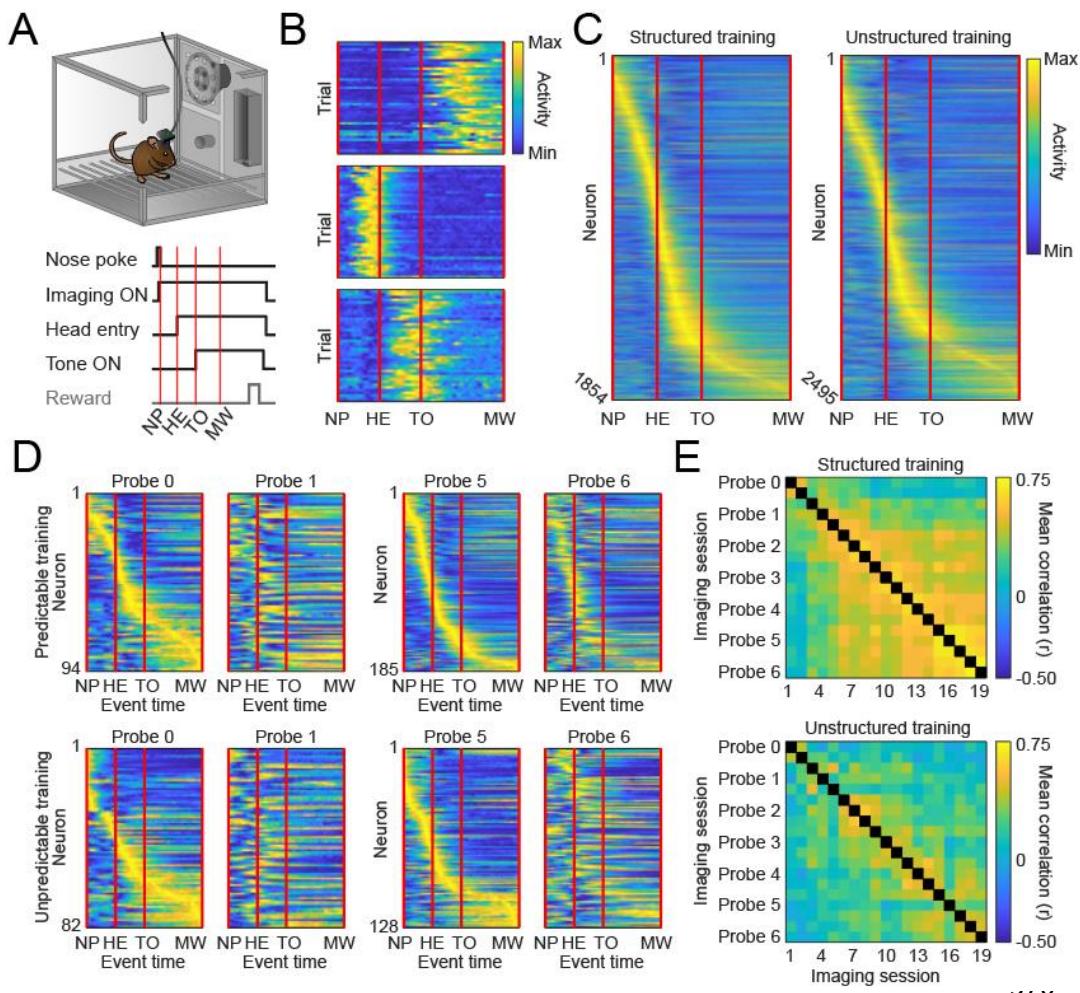
312 **Figure 4.** Mice reactivate CA1 populations during structured, but not unstructured, training. **(A)** Cartoon of the
313 custom miniature microscope implanted into the hippocampus. **(B)** Histological section used to identify the
314 location of a GRIN lens implanted on the surface of CA1. **(C)** We tracked the activity of every neuron throughout
315 training for both the Structured training and Unstructured training groups (see also **Supplemental figure 3**). The
316 columns show the activity of one neuron over all 19 training sessions (7 probes and the first and last day of each of
317 6 problems). Neurons are colored according to the session in which they were first active. Rows show the
318 composition of every session in terms of which cells were active. The composition is summarized in the outset bar
319 chart. **(D)** An explanation and example of the population overlap computation, defined as the number of cells
320 active in both sessions divided by the number active in either. **(E)** The mean observed overlap between every pair
321 of imaging sessions for the Structured training and Unstructured training groups. Matrices showing excess overlap
322 are shown in **Supplemental figures 4D-E**.

323 CA1 ensemble activity patterns stabilize during Structured training

324 In addition to activating unique ensembles of neurons, the hippocampus also guides memory
325 retrieval by reorganizing neuronal activity patterns along dimensions that are important for task
326 performance (i.e., remapping; Muller and Kubie, 1987; Markus et al., 1995; Smith and Mizumori, 2006;
327 Kelemen and Fenton, 2010; Bulkin et al., 2016). To assess how the reorganization of hippocampal

328 activity might support the formation of a predictive model in our task, we investigated time-warped
329 neuronal activity during events common to every trial (from nose poke until the mouse waited in the
330 food hopper for 2 seconds after tone-on; **Figure 5A**). Individual neurons reliably became active at
331 discrete periods during every trial (**Figure 5B**) such that the population of neurons encoded the entire
332 trial progression (**Figure 5C**). To evaluate how patterns of activity changed over training, we identified
333 neurons that were active in multiple sessions and measured the similarity of their activity patterns
334 between sessions (**Figure 5D**) while controlling for the effect of time between sessions as above (i.e.,
335 excess activity similarity, **Supplemental figure 4**). Comparing the excess activity similarity between every
336 pair of sessions for both groups showed that population activity became more similar with Structured
337 training only (LME model with fixed effects for sessions and group, and a random effect for subject,
338 session X session X group interaction, $t(3730) = 5.77$, $p < 0.001$; post-hoc LME models with fixed effects
339 for sessions and a random effect for subject, Structured training group session X session interaction,
340 $t(2334) = 7.23$, $p < 0.001$, Unstructured training group session X session interaction, $t(1396) = 1.57$, $p =$
341 0.12; **Figure 5E; Supplemental figure 4**). Therefore, Structured training increased the degree to which
342 hippocampal neurons from early sessions maintained their activity patterns when reactivated in later
343 sessions, consistent with the idea that this activity supported the retrieval of training memories late in
344 learning. As with the increase in ensemble overlap described above, this increase in activity similarity
345 occurred at approximately the same time that the Structured training group formed a predictive model
346 (**Figure 2C-D**). Together, this suggests that the hippocampus supported a predictive model by repeatedly
347 reactivating prior training memories that predicted the answers to future problems and enabled
348 efficient learning.

349 **Figure 5.** Mice reactivate CA1 activity patterns during Structured training. **(A)** We examined activity during the
350 window between nose-poke and then end of the minimum wait (MW) period, occurring 2s after tone-on. **(B)** Three
351 examples of neurons that were reliably active during discrete periods on every trial. **(C)** The mean activity is shown
352 for every neuron imaged from each group, arranged by when during the trial it was most active. **(D)** The activity
353 patterns of neurons that were active across any pair of probe sessions surrounding common problems (probes 0 &
354 1, 5 & 6) are shown. We examined the similarity of population activity across sessions by measuring the correlation
355 between the activity patterns of individual cells in both sessions. The top row shows all neurons imaged from mice
356 in the Structured training group that were active in both probe 0 and 1 (left side), and then neurons that were
357 active in both probe 5 and 6 (right side). The bottom row shows neurons imaged from mice in the Unstructured
358 training group. For each pair of probes, the neurons are sorted according to their activity in the left-side probe
359 (probe 0 or probe 5), such that corresponding rows show the activity of the same neuron. Note the highly similar
360 activity patterns between probes 5 and 6 in the Structured training group only. **(E)** The mean observed activity
361 pattern similarity between every pair of imaging sessions for the Structured training and Unstructured training
362 groups. Matrices showing excess activity similarity are shown in **Supplemental figure 4F-G**.



377 Learning-related neuronal activity is incorporated into a task memory ensemble

378 Because the hippocampal ensemble stabilized at the same time that mice began to form a predictive
 379 model, we hypothesized that this stability reflected the transition from learning new problems to
 380 predicting them. In particular, we hypothesized that mice reactivated and updated ensembles during
 381 learning, and that the ensembles stabilized when mice developed a predictive model. To test this, we
 382 quantified the degree to which the ensembles from earlier sessions were reactivated during later
 383 sessions in terms of a reactivation score that described both the ensemble overlap between a pair of
 384 sessions and the similarity of their activity patterns (see methods; time-corrected results reported in
 385 **Supplemental figure 5**). We then tested the relationship between ensemble reactivation and behavior.

386 We first investigated whether mice reactivated prior ensembles and associated memories during
 387 new learning by comparing the ensemble reactivation scores between probe sessions and the training
 388 sessions one day later. Consistent with the above findings, we observed that the Structured training

389 group increased ensemble reactivation during learning specifically during later training problems
390 (Structured training, $F(5, 35) = 6.35$, $p < 0.001$), while the Unstructured training group did not
391 (Unstructured training, $F(5, 20) = 1.67$, $p = 0.19$; **Figure 6A**; between groups comparisons, LME model
392 with fixed effects for session and group, and a random effect for subject, session X group interaction,
393 $t(22) = 2.44$, $p < 0.05$; **Figure 6B**). However, if ensemble reactivation is indicative of predictive memory
394 retrieval, then the mice should show greater reactivation specifically when they show behavioral
395 evidence of accurate memory retrieval. Consistent with this, we observed that ensemble reactivation
396 scores were correlated with accurate discrimination performance on the first day of training, our
397 behavioral measure of memory retrieval during learning ($r = 0.39$, $p < 0.001$; LME model with fixed effect
398 for first day discrimination and a random effect for subject, $t(75) = 3.52$, $p < 0.001$; **Figure 6C**).

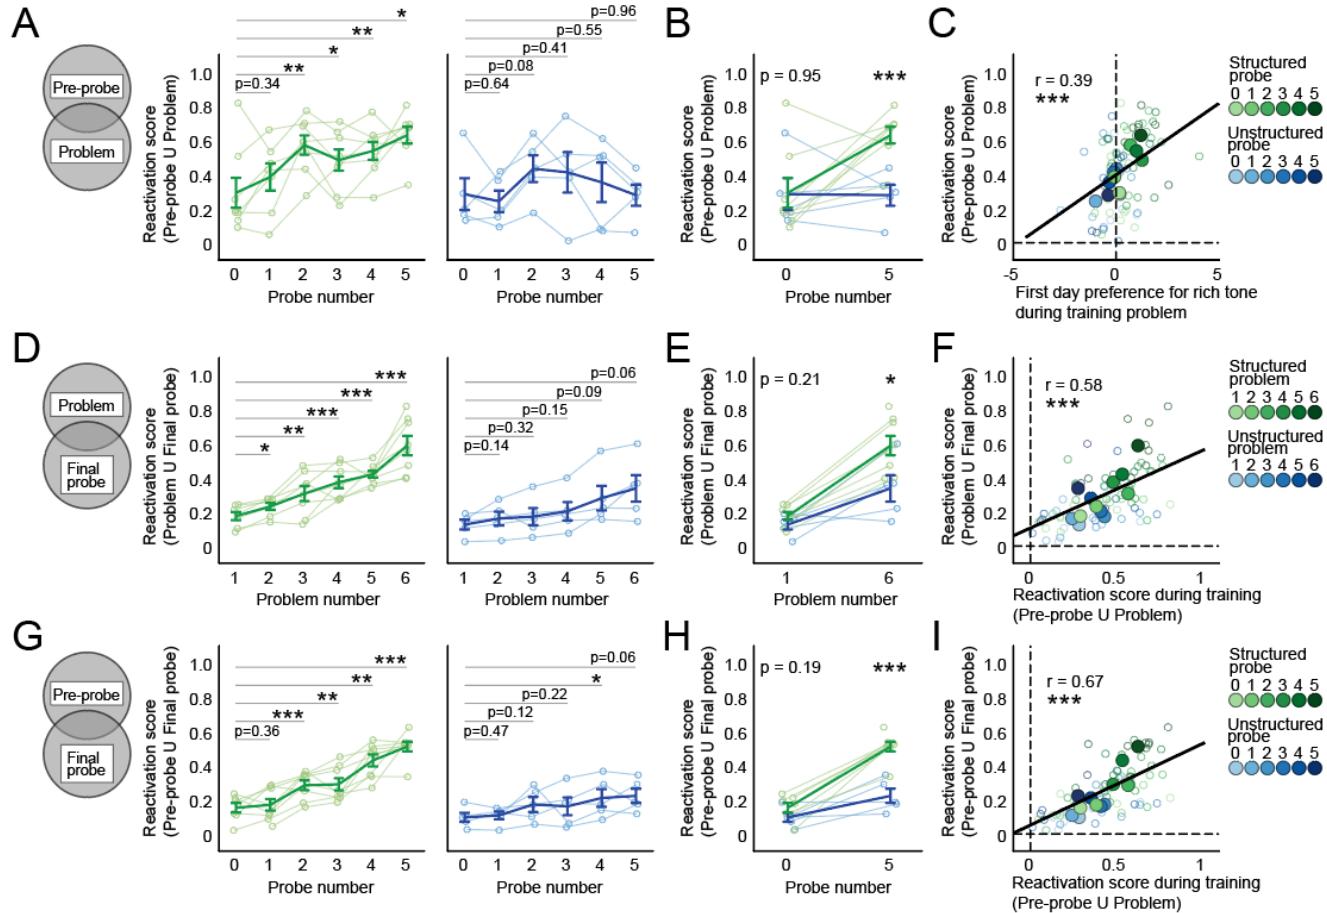
399 In addition to supporting memory retrieval, reactivating prior ensembles during learning may also
400 support the formation of predictive memories by promoting memory updating and integration
401 (Zeithamova et al., 2012; Cai et al., 2016; Rashid et al., 2016). To investigate the degree to which
402 reactivated ensembles were updated with new neurons and activity patterns during training sessions,
403 we compared the ensemble observed during the final probe test—when mice expressed their fully
404 integrated predictive model—to the ensembles observed during individual training sessions. If training-
405 session information is incorporated into the final predictive model, then there should be increased
406 reactivation of the training ensemble during the final probe. Consistent with the evidence for increased
407 reactivation in the Structured training group described above, we observed that the later Structured
408 training sessions were most strongly reactivated during the final probe session (Structured training
409 group, $F(5, 35) = 27.18$, $p < 0.001$; Unstructured training group, $F(5, 20) = 5.64$, $p < 0.01$; **Figure 6D**;
410 between-group comparison, LME model with fixed effects for session and group, and a random effect
411 for subject, session X group interaction, $t(21) = 2.30$, $p < 0.05$; **Figure 6E**). Analyses controlling for time
412 between imaging sessions confirmed that only Structured training sessions showed excess reactivation
413 scores that exceeded chance (i.e., what was expected due to the amount of time between imaging
414 sessions), beginning with the second training session (**Supplemental figure 5G-H**). This indicates that
415 Structured training promoted the development of a final probe ensemble that was comprised of
416 neurons and activity patterns seen during earlier training problems.

417 Importantly, if the incorporation of training-related information into the predictive model is due to a
418 memory updating process occurring during training, then the degree to which training ensembles are
419 reactivated during the final probe should be related to neural processes occurring during training.

420 Specifically, we would expect to see greater updating during training sessions with stronger reactivation
421 of the to-be-updated ensemble. Consistent with this, we found that reactivation during the final probe
422 was correlated with the amount of reactivation observed during training: the greater the reactivation of
423 the preceding probe ensemble during training, the greater the reactivation of that training ensemble
424 during the final probe test ($r = 0.58$; $p < 0.001$; LME model with fixed effect for first day discrimination
425 and a random effect for subject, $t(74) = 5.52$, $p < 0.001$; **Figure 6F, Supplemental figure 5I**). This suggests
426 that the ensemble observed during the final probe was repeatedly reactivated during—and updated
427 with new neuronal activity patterns from—each training problem.

428 If so, then the ensembles observed during each probe test should become progressively more
429 similar to the final probe test ensemble as mice accumulate information from each training problem into
430 a predictive model. Consistent with this, probe reactivation scores increased over training for both
431 groups (Structured training group, $F(5, 35) = 27.62$, $p < 0.001$; Unstructured training group, $F(5, 20) =$
432 3.05 , $p < 0.05$; **Figure 6G**), but this increase was greater in the Structured training group (LME model
433 with fixed effects for session and group, and a random effect for subject, session X group interaction,
434 $t(21) = 4.21$, $p < 0.001$; **Figure 6H**), and only the Structured training group showed excess reactivation
435 over that expected by chance (**Supplemental figure 5J-K**). Similar to above, we also found a strong
436 correlation between probe ensemble reactivation during training and probe ensemble reactivation
437 during the final probe ($r = 0.67$, $p < 0.001$; linear mixed effects model with fixed effect for first day
438 discrimination and a random effect for subject, $t(74) = 7.51$, $p < 0.001$; **Figure 6I, Supplemental figure**
439 **5L**), consistent with the idea that the incorporation of new information required ensemble reactivation
440 during learning. Together, these data suggest that the increased ensemble stability observed in the
441 Structured training group may be attributed to the repeated reactivation of a predictive model and
442 corresponding ensemble, both formed by incorporating information from each of the training problems.

443



444 **Figure 6.** Learning-related neuronal activity is incorporated into a reactivated ensemble. **(A)** We computed a
445 reactivation score that combined ensemble overlap and activity similarity measures. The reactivation probe
446 session population during subsequent training problems is shown for the Structured training group (left) and the
447 Unstructured training group (right). **(B)** The reactivation of probe ensembles during subsequent common problems
448 is shown for both groups. **(C)** The reactivation of probe ensembles during subsequent problems was predicted by
449 first day discrimination accuracy, a behavioral measure of predictive memory retrieval. Large circles show
450 condition means. **(D)** The reactivation of training session ensembles during the final probe session is shown for the
451 Structured training group (left) and the Unstructured training group (right). **(E)** The reactivation of the common
452 training problem ensembles during the final probe session is shown for both groups. **(F)** The reactivation of training
453 ensembles during the final probe was predicted by the reactivation of probe ensembles during training. **(G)** Like D,
454 but with the reactivation of probe ensembles during the final probe session. **(H)** Like E, but with probe ensembles.
455 **(I)** The reactivation of probe ensembles during the final probe was predicted by the reactivation of probe
456 ensembles during training problems.

457

458

459

460

461

462 **Discussion**

463 Hippocampal activity patterns are hypothesized to reflect predictive models of the environment
464 (Sanders et al., 2020). Here, we tested this hypothesis by investigating the relationship between
465 memory, prediction, and CA1 ensemble activity as mice learned the rules that governed their
466 environment. By using a new behavioral method for observing how memories change after discrete
467 learning experiences, we found that Structured learning experiences promoted the integration of new
468 learning into a memory of the task environment that accurately predicted future problems (i.e., a
469 predictive model). We showed that this model included predictions about untrained stimuli, and that
470 the accuracy of these predictions had a large effect on the efficiency of new learning. Then, by observing
471 hippocampal activity throughout learning, we found that mice learning Structured training problems
472 reactivated prior ensembles, thereby supporting the integration of new information into the reactivated
473 ensemble. These data indicate that hippocampal ensemble reactivation supports the formation of
474 predictive models that explain the underlying structure of an environment. Moreover, these data shed
475 light on how memory systems organize experiences to maximize prediction accuracy and minimize
476 memory interference.

477 The hippocampus has long been implicated in memory processes that support predictive models
478 (Tolman, 1948) and considerable evidence suggests that hippocampal activity develops as subjects learn
479 about their environment (Smith and Mizumori, 2006; Gill et al., 2011; Plitt and Giocomo, 2021).
480 However, a fundamental feature of a predictive model is the ability to infer information beyond what
481 has been directly experienced—to generalize what has been learned to explain and predict aspects of
482 the environment that have not been learned (Tervo et al., 2016; Whittington et al., 2018, 2022).
483 Consistent with this, we found that mice readily generalized what they learned from individual
484 discrimination problems into predictions about novel contingencies during subsequent probe tests.
485 Predictions expressed during probe tests, in turn, affected how mice learned new problems, such that
486 accurate predictions accelerated learning and inaccurate predictions impeded learning. These findings
487 are consistent with longstanding ideas that it is easier to learn things that are related to what is already
488 known (Bartlett, 1932; Bransford and Johnson, 1972; Craik and Tulving, 1975; Harlow, 1949; Chase and
489 Simon, 1973; Tse et al., 2007). These findings also highlight the reciprocal relationship between learning
490 and memory whereby learning is stored in memory and then retrieved during new learning.

491 We additionally found that the Structured training group successfully integrated new
492 experiences into a predictive model, whereas the Unstructured training group did not. This indicates

493 that the ability to integrate experiences was influenced by whether they shared an underlying structure.
494 It has been previously noted that the relationship between new experiences and existing memories
495 plays an important role in memory formation (Fernández et al., 2016; van Kesteren et al., 2012).
496 However, it has not always been clear what makes a given experience sufficiently related to another to
497 promote memory integration. One possibility is that mice selectively integrate information that is
498 predicted by the memories retrieved during learning. Such a rule could support the formation of
499 predictive models by organizing experiences in memory according to shared hidden causes (i.e. whether
500 both experiences were caused by the same underlying rule or structure). It is also likely that memory
501 integration in the Structured training group was facilitated by the learning-related benefits of retrieving
502 accurate predictions: predicted experiences are learned more quickly, thereby increasing the likelihood
503 that learning occurs in the presence of the retrieved memory. Inversely, if a memory interferes with
504 learning (as was observed in both groups early in training and in the Unstructured training group
505 throughout training), then the retrieved memory may be suppressed, thereby preventing the new
506 information from being integrated into it. Over many iterations, these two processes could organize a
507 series of experiences into a predictive model that closely matches the underlying structure of the
508 environment and can therefore support predictions about future experiences in that environment.

509 Integrating experiences requires a system for retrieving relevant memories and then updating
510 these memories with new information. The hippocampus plays essential roles in memory retrieval (Liu
511 et al., 2012; Robinson et al., 2020) and consolidating new learning (Davis and Squire, 1984; Girardeau et
512 al., 2009). One way these functions cooperate is through the process of integrating memories that are
513 learned under the same hippocampal activity pattern, such as when two experiences occur close in time
514 (Cai et al., 2016; Chowdhury et al., 2022) or when a second experience reminds one of the first
515 (Zeithamova et al., 2012; Lee et al., 2017). In the current study, we found increased reactivation of
516 hippocampal activity patterns associated with prior problems exclusively in the Structured training
517 group and specifically late in learning, coinciding with the time that memory integration became
518 apparent behaviorally. This suggests that memory integration occurs when predictive memories and
519 their corresponding hippocampal activity states are retrieved during new learning, leading to integration
520 of newly learned information into the reactivated state. If so, then the memory integration might result
521 from a two-part process whereby memory retrieval increases the activity of the subset of neurons
522 associated with the retrieved memory (i.e., the engram), and then a neuronal allocation process
523 preferentially assigns newly learned information to those active neurons (Rashid et al., 2016). This
524 process would result in the first and second experience becoming linked in memory due to both

525 experiences sharing a single neuronal population that drives the retrieval of both memories when
526 reactivated.

527 A complementary possibility is that the hippocampus changes its activity state in response to
528 unpredicted problems. Although retrieving related memories is important for memory integration and
529 the formation of predictive memories, it is arguably even more important to avoid retrieving and
530 integrating unrelated information, which can cause memory interference (Smith and Vela, 2001) and
531 possibly amnesia (McCloskey and Cohen, 1989). The hippocampus is well known for its ability to
532 reorganize its firing activity in response to new environments (Muller and Kubie, 1987). Interestingly,
533 this type of reorganization can also free subjects from memory interference (Bulkin et al., 2016; Butterly
534 et al., 2012), likely by discouraging the retrieval of the interfering information (Park et al., 2021).
535 Therefore, one explanation for our finding that the Unstructured training group showed little
536 reactivation of prior hippocampal activity patterns during learning is that the hippocampus actively
537 generated new patterns to facilitate the learning of unstructured problems. This would also help to
538 avoid erroneously integrating experiences that do not share an underlying hidden cause and whose
539 integrated memory would therefore have little predictive value.

540 Memory organization can also occur 'offline' during sleep or otherwise outside of task
541 performance. In the days and weeks after learning, newly learned information goes through a systems
542 consolidation process that transfers information into neocortical regions (Squire, 2004), and there are
543 changes to memory content such as the loss of detailed information (Wiltgen and Silva, 2007; Richards
544 et al., 2014) and, in some cases, additional learning or insight (Barron et al., 2020; Payne et al., 2009).
545 Mice in the current study learned training problems over the course of weeks (**Figure 1F-G**), providing
546 sufficient opportunity for long-term memory generalization into the cortex (Makino and Komiyama,
547 2015; Miller et al., 2019; Takehara-Nishiuchi and McNaughton, 2008; Tse et al., 2007, 2011). However,
548 our data show that the process also occurs locally in the hippocampus. One possibility is that
549 hippocampus supports the selective retrieval of predictive memories stored in the neocortex in order to
550 update them with relevant new information (Debiec et al., 2002).

551 Together, these data suggest a concise description of prediction-based memory organization in
552 the hippocampal system. After any learning experience, the resulting memory trace generalizes to form
553 predictions about novel stimuli that can be retrieved when those stimuli are encountered in the future.
554 When the generalized memories are retrieved during new learning, they are accompanied by the
555 reactivation of hippocampal activity patterns that were present when the original memory was learned.

556 If the retrieved memory improves learning, then the prior hippocampal activity patterns remain
557 activated while the new memory is formed, thereby integrating the two memories. Alternatively, if the
558 retrieved memory impedes learning, then the hippocampus changes its activity pattern to discourage
559 the retrieval of the interfering memories and to avoid erroneous integration. Over many learning
560 experiences, a predictive model is formed that matches the statistics of the experiences from which it
561 was formed, and that can be used to predict other similar future experiences. Memories that are not
562 incorporated into the model may be used to seed new models, or they may become isolated and
563 forgotten.

564

565

566

567

568

569

570

571

572

573

574

575

576

577

578

579

580

581 **Acknowledgements:** A.M.P.M. and P.W.F. designed the experiments and wrote the first draft of the
582 paper. A.M.P.M., A.I.R. and M.L.D. performed the surgeries and did the histology. A.M.P.M. and A.D.J.
583 designed and updated the behavioral and imaging hardware. A.M.P.M. conducted the experiments and
584 performed the analyses. A.M.P.M., P.W.F. and S.A.J. supervised the project. All authors helped edit the
585 paper. We thank Tao Zhang, Shuo Huang, Sandy Ma, and Lulu Liu for assistance with pilot versions of
586 this project, Chen Yan, Lina Tran, and Andrew Mocle for assistance adapting the microscopes and
587 imaging analysis pipeline, and Hendrick Steenland for designing components of the imaging acquisition
588 hardware. We thank Nathan Insel for helpful comments on the manuscript. This work was supported by
589 a Brain Canada platform grant and NIMH grant (RO1MH119421) to S.A.J. and P.W.F. A.M.P.M. was
590 supported by a Restracomp award and a Research Institute Exceptional Trainee Award Fund Bursary
591 from The Hospital for Sick Children. A.D.J. was supported by an Ontario Graduate Scholarship. A.I.R was
592 supported by an NSERC CGS-D award and an NIH 1 F31 MH120920-01 award. M.L.D. was supported by a
593 CIHR Vanier Canada Graduate Scholarship.

594

595 **Declaration of Interests:** The authors declare no competing interests.

596

597

598

599

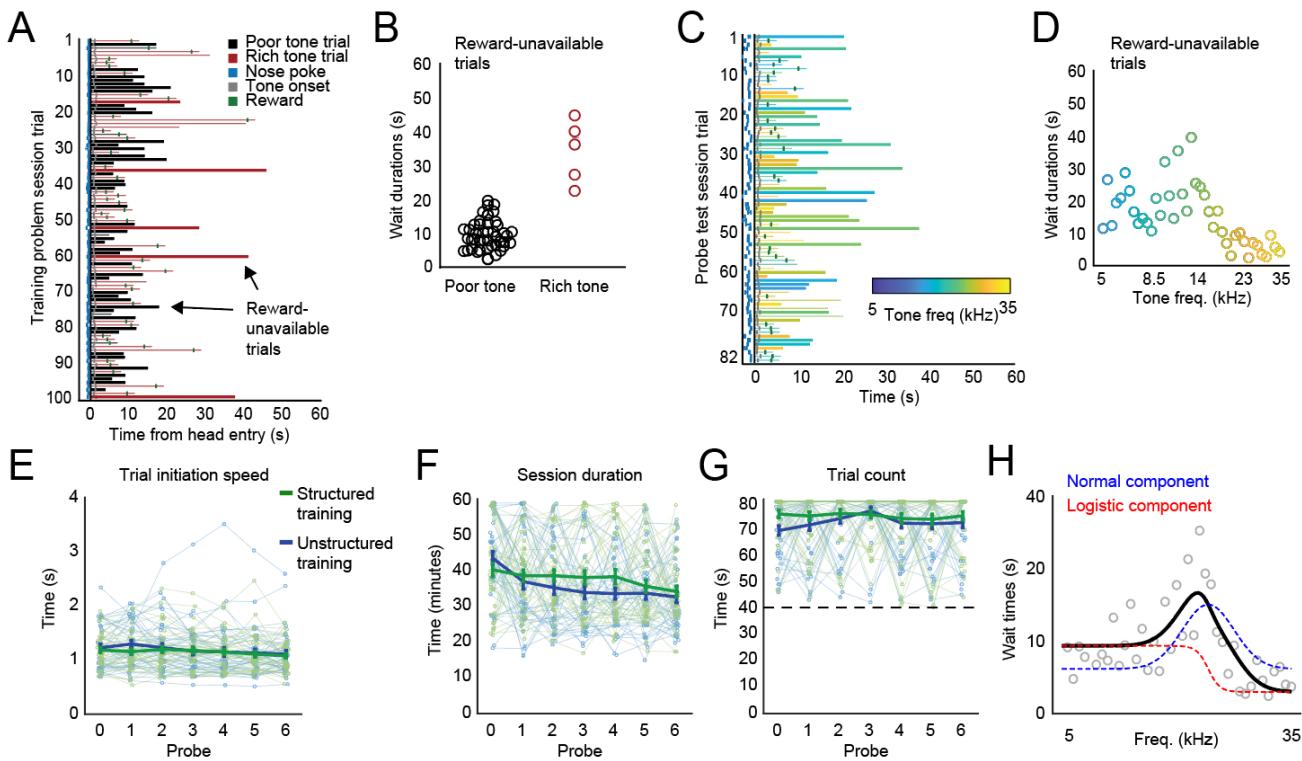
600

601

602

603

604



605 **Supplemental figure 1.** Operant behavior in unoperated mice. **(A)** Every measured behavior during an exemplary
606 training session. Mice had the potential to earn 50 rewards over 100 trials in an hour-long session. On each trial,
607 one of two unique tones would play: the rich tone ($p(\text{reward})=0.9$) or the poor tone ($p(\text{reward}) = 0.1$) and the
608 mouse responded by waiting for a potential reward. The wait durations on trials when no reward was available
609 (thick lines) indicated the mouse's preference for the tone that played on that trial. **(B)** All wait durations from
610 reward-unavailable trials in A. **(C)** Every measured behavior during an exemplary probe test session. Mice had the
611 potential to earn 41 rewards over 82 trials in an hour-long session. On each trial, one of 41 unique tones would
612 play. No reward was available on half of all trials (all tones $p(\text{reward}) = 0.5$) and the wait durations on these trials
613 (thick lines) indicated the mouse's preference for each tone. In this example of behavior on probe 1 (after training
614 on problem 1), the mouse waits longer in response to lower-frequency tones **(D)** All wait durations from reward-
615 unavailable trials in C. **(E)** The speed of task behavior was similar between the groups throughout learning. The
616 amount of time between nose poke entry and head entry is shown for every subject at every probe session. **(F)** The
617 amount of time required to complete a session is shown for every probe session. Probe sessions were completed
618 after 82 trials with a maximum time of 60 minutes. **(G)** The number of trials completed in every probe session is
619 shown. Sessions with fewer than 41 trials (dotted line) were repeated. **(H)** For some analyses (Figure 3;
620 Supplementary figure 2), we estimated the preference for every tone by fitting a curve composed of a normal
621 component and a logistic component to the observed wait times on every trial (probe sessions only; see methods).

622

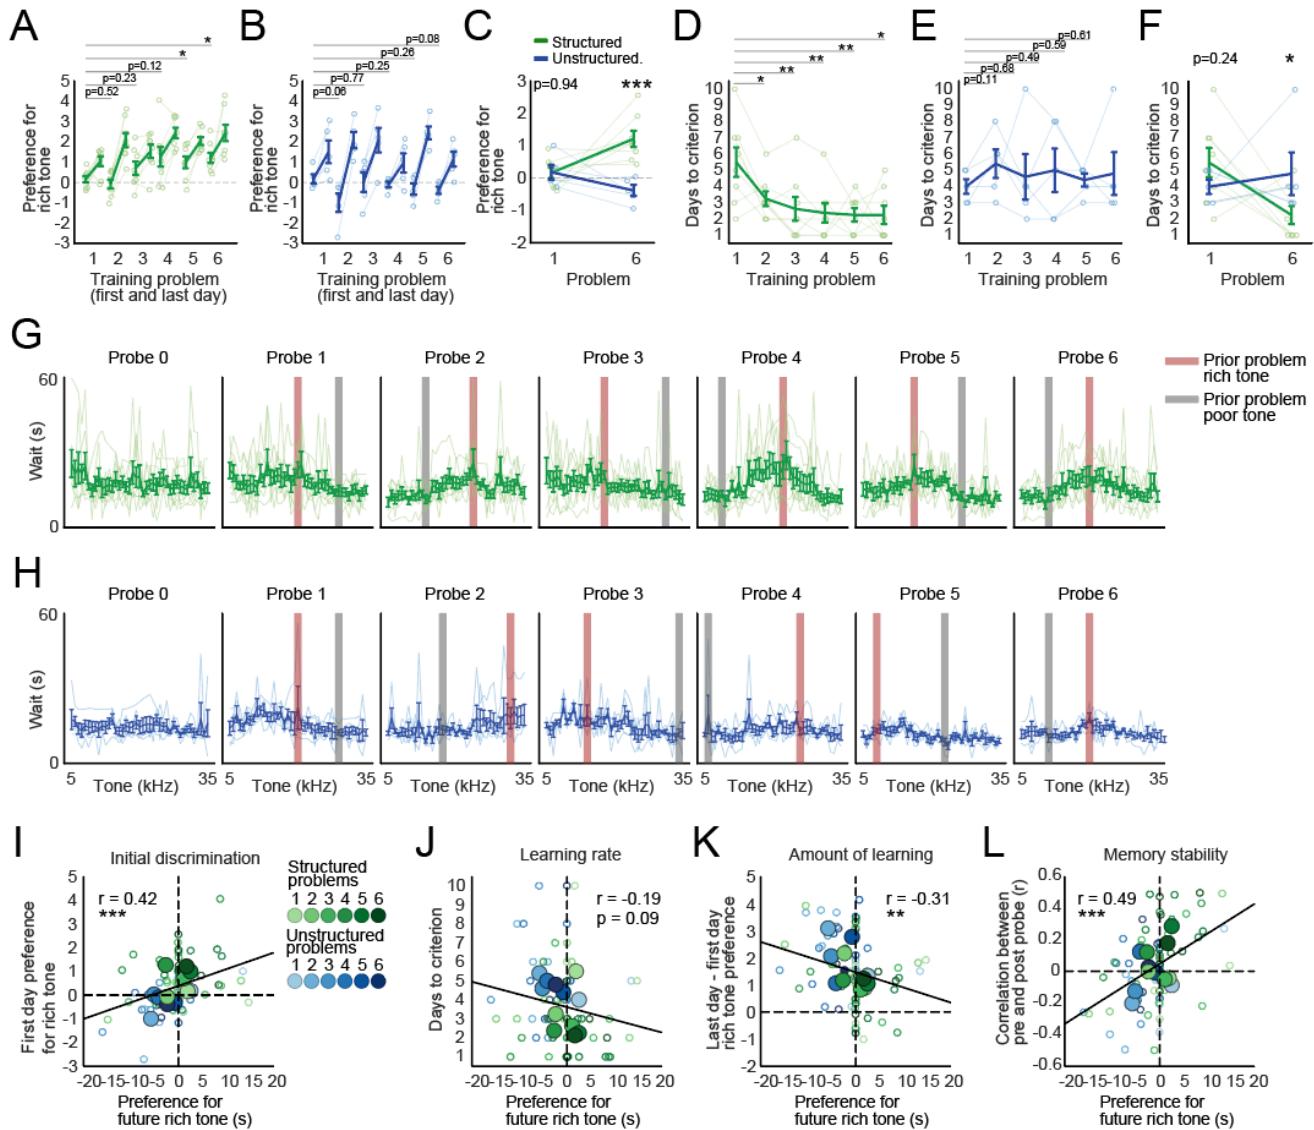
623

624

625

626

627



628 **Supplemental figure 2.** Operant behavior during calcium imaging. **(A)** Thirteen mice (8 Structured training, 5
629 Unstructured training) were imaged performing the discrimination and foraging task in an operant chamber.
630 Preference for the rich tone (standardized mean difference between wait times in response to rich and poor tones)
631 is shown for the first and last day of training on every training problem. The Structured training group showed a
632 greater preference for the rich tone on the first day of training on problems later in learning (one-way repeated
633 measures ANOVA, $F(5,35) = 3.35$, $p < 0.05$). **(B)** The Unstructured training group never improved its first-day
634 preference for the rich tone (one-way repeated measures ANOVA, $F(5,20) = 1.59$, $p = 0.21$). **(C)** Direct comparisons
635 between the two groups on the first and last problem revealed that the structured group showed greater
636 preference for the rich tone on the first day of the last problem (mixed model with fixed effects for intercept,
637 training group, problem number, and group X problem and a random effect for subject, group X problem
638 coefficient, $t(22) = 3.97$, $p < 0.001$). **(D)** The number of days required to reach criterion is shown. The Structured
639 training group learned new problems faster later in learning (one-way repeated measures ANOVA, $F(5,35) = 6.16$, p
640 < 0.001). **(E)** The Unstructured training group never improved its learning rate (one-way repeated measures
641 ANOVA, $F(5,20) = 0.28$, $p = 0.92$). **(F)** Direct comparisons between the two groups on days required to reach
642 criterion on the common problems (group X problem coefficient, $t(22) = 2.64$, $p < 0.05$). **(G)** The wait time
643 responses to every tone are shown for every subject in every probe test of the Structured training group. **(H)** The
644 same as G, but for the Unstructured training group. **(I)** We measured predictions from behavior during the probes
645 preceding each problem (see **Figure 3E**). The more accurate the prediction, the better the discrimination

646 performance on the first day of training (linear effects mixed model with fixed effects for intercept, prediction
647 accuracy, and a random effect for subject: prediction accuracy term, $t(76) = 3.33$, $p < 0.01$). Large dots indicate
648 condition means. **(J)** The relationship between prediction and learning rate on the training problem (prediction
649 accuracy term, $t(76) = 1.92$, $p = 0.06$). **(K)** The more accurate the prediction, the less new information was learned
650 during training (prediction accuracy term, $t(76) = 2.35$, $p < 0.05$). **(L)** The more accurate the prediction, the more
651 stable the memory as measured during pre and post probe tests (prediction accuracy term, $t(76) = 4.53$, $p < 0.001$).

652

653

654

655

656

657

658

659

660

661

662

663

664

665

666

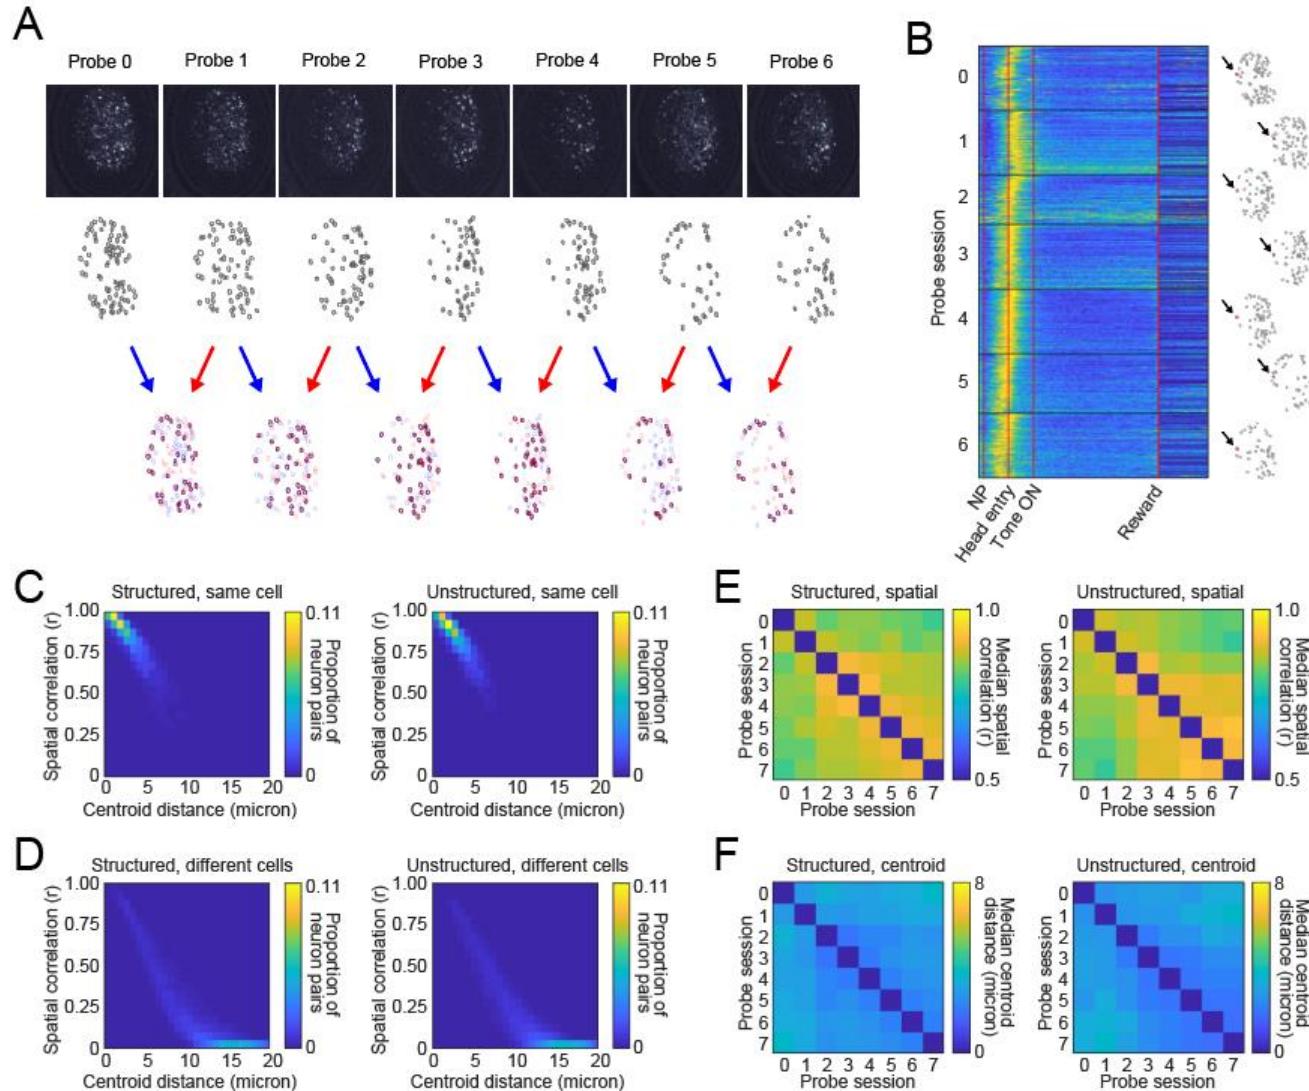
667

668

669

670

671

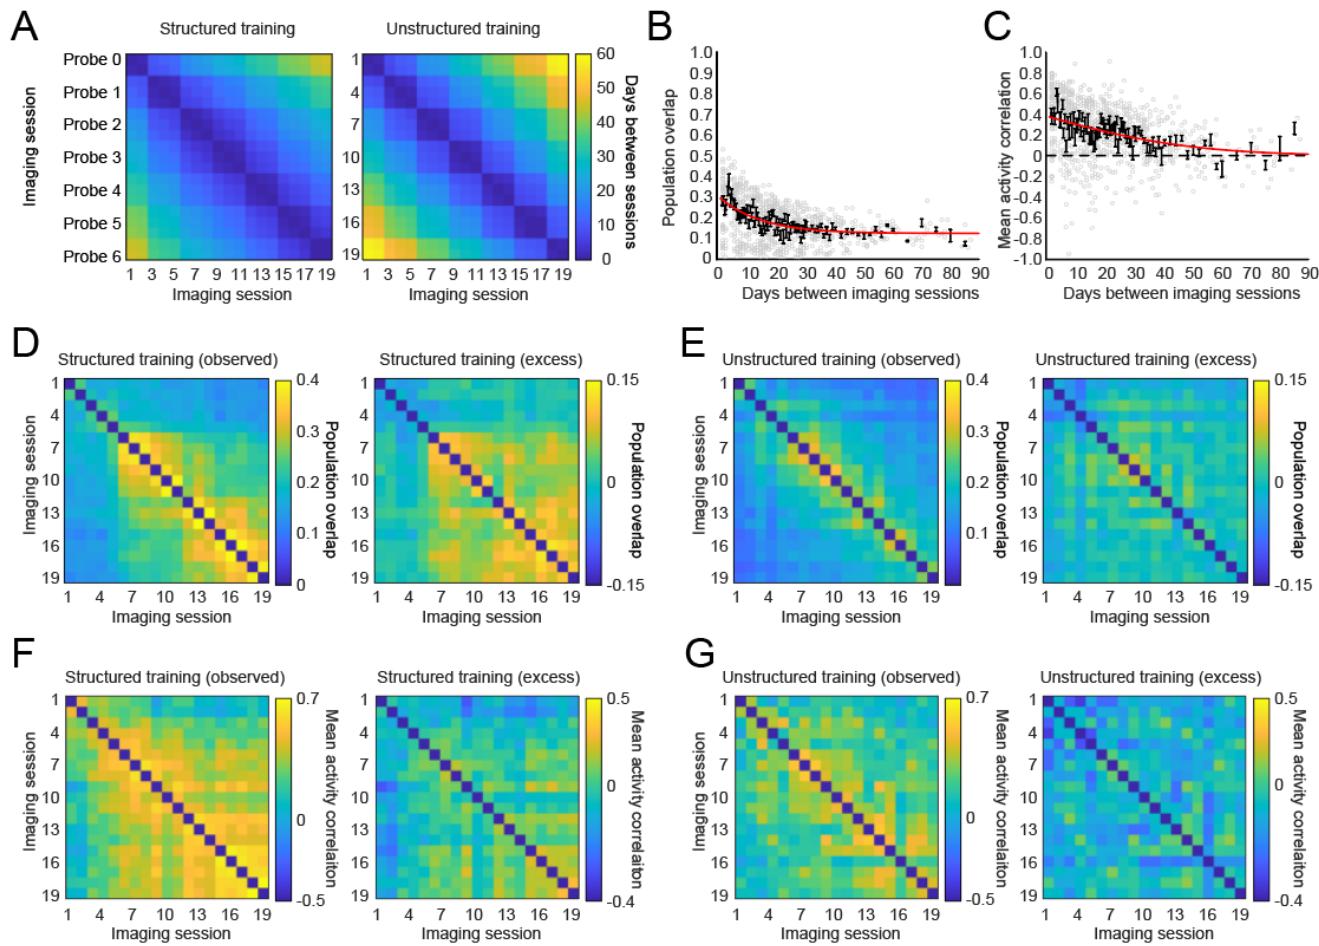


672 **Supplemental figure 3.** Registering CA1 neurons. **(A)** Example showing the registration of imaging data collected
673 over weeks. Top row, fields of view (maximal projection images) from the same mouse from every probe session.
674 Second row, spatial footprints are identified from the video file and aligned (translated and rotated) to a common
675 reference frame. Bottom row, session pairs are then overlaid to identify overlapping neurons. **(B)** Example of a
676 neuron that was active during every probe 0 – 6 and showed a stable activity pattern. Rows show normalized
677 activity on every trial from every probe session. The neuron's spatial footprint is highlighted to the right of each
678 probe session. **(D)** Anatomical distance vs spatial footprint similarity for every registered (same cell) cell pair in the
679 dataset. For each cell pair, we plotted the distance between the two cells versus the correlation coefficient of their
680 cell mask spatial correlations. **(E)** Same as B, but for unregistered (different cells) cell pair. **(E)** Median spatial
681 correlations for every registered cell pair from each probe session pair for each group. **(F)** Same as E, but with
682 anatomical distance between centers of each cell.

683

684

685



686

687 **Supplemental figure 4.** The relationship between ensemble similarity and time. **(A)** The average time between
 688 every imaging session is shown for both the structured and Unstructured training groups. **(B)** We modelled the
 689 effect of time on reactivation scores by fitting a truncated normal distribution to the ensemble overlap and mean
 690 activity similarity values obtained by comparing all pairwise combinations of imaging sessions from the
 691 Unstructured training group. Mean overlap in bins spanning time is shown in black. The fit model curve is shown in
 692 red. We obtained the time-corrected (excess) overlap values by subtracting the effect characterized by the fit
 693 curve from the observed overlap values. **(C)** Like B, but with the mean activity similarity across cells that were
 694 active any two sessions. **(D)** The ensemble overlap between every pair of sessions during Structured training is
 695 shown in terms of the observed values and the excess values. **(E)** Same as C, but for Unstructured training. **(F-G)**
 696 Like D and E, but with mean activity similarity.

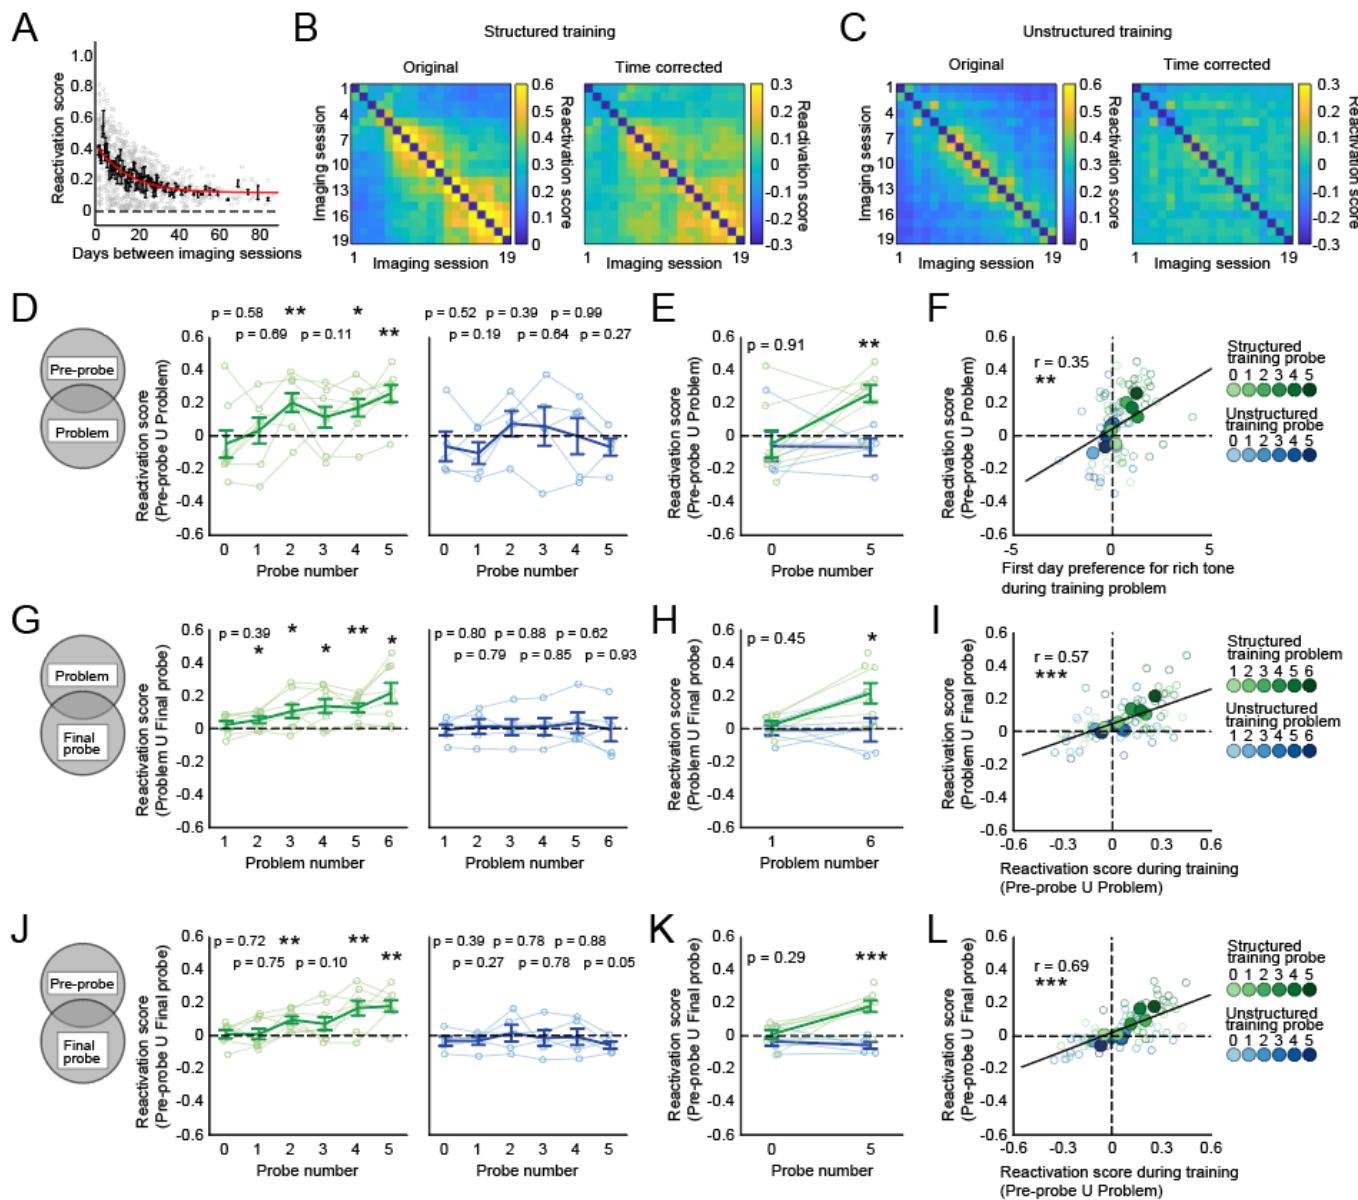
697

698

699

700

701



702 **Supplemental figure 5.** Learning related neuronal activity is incorporated into a reactivated ensemble, time
 703 corrected. **(A)** We modelled the effect of time on reactivation scores by fitting a truncated normal distribution to
 704 the reactivation scores obtained by comparing all pairwise combinations of imaging sessions from the
 705 Unstructured training group. **(B)** The observed and time-corrected mean reactivation score values for all pairwise
 706 Structured training group imaging session comparisons. **(C)** The observed and time-corrected mean reactivation
 707 score values for all pairwise Structured training group imaging session comparisons. There was a significant group
 708 (structured vs Unstructured training) X session 1 X session 2 interaction for the time-corrected reactivation scores
 709 (linear mixed model with fixed effects for session 1, session 2, and group, and a random effect for subject, $t(4344) = 7.34$, $p < 0.001$). **(D)** The reactivation probe ensembles during subsequent training problems increased with
 710 training for the Structured training group ($F(5,35) = 5.39$, $p < 0.001$; left) and but not the Unstructured training
 711 group ($F(5, 20) = 1.55$, $p = 0.22$; right). **(E)** The reactivation of probe ensembles during subsequent common
 712 problems are shown for both groups. Only the structured group showed an increased reactivation score late in
 713 learning (fixed effects for group and session and a random effect for subject, group X session interaction, $t(22) = 2.37$, $p < 0.05$). **(F)** The reactivation of probe ensembles during subsequent problems was predicted by first day
 714 discrimination accuracy (fixed effect for discrimination accuracy and a random effect for subject, $t(75) = 3.15$, $p <$
 715 0.01). **(G)** The Structured training group showed increased reactivation of later training session ensembles during
 716

718 the final probe session ($F(5,35) = 6.12$, $p < 0.001$; left), while the Unstructured training group did not ($F(5,20) =$
719 0.24 , $p = 0.93$, right). **(H)** Only the Structured training group showed an increase at the common training problems
720 (fixed effects for group and session and a random effect for subject, group X session interaction, $t(21) = 2.18$, $p <$
721 0.05). **(I)** The reactivation of training ensembles during the final probe was predicted by the reactivation of probe
722 ensembles during training (fixed effect for training reactivation and a random effect for subject, $t(21) = 2.18$, $p <$
723 0.05). **(J)** Like G, but with the reactivation of probe ensembles during the final probe session (Structured training,
724 $F(5,35) = 6.94$, $p < 0.001$; Unstructured training, $F(5,20) = 0.92$, $p = 0.49$). **(K)** Like H, but with probe ensembles
725 (fixed effects for group and session, group X session interaction, $t(21) = 3.75$, $p < 0.01$). **(L)** The reactivation of
726 probe ensembles during the final probe was predicted by the reactivation of probe ensembles during subsequent
727 training problems (fixed effect for training reactivation and a random effect for subject, $t(74) = 7.40$, $p < 0.001$).

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747 **Methods**

748 Animals

749 All experiments conformed to the guidelines set forth by the Canadian Council for Animal Care and the
750 Animal Care Committee at the Hospital for Sick Children. For the behavior-only study, we used 80 wild
751 type mice aged 8-10 weeks at the start of training. These mice were derived from crossing C57BL6n and
752 129SvEv mice (breeding mice were sourced directly from Taconic). For the imaging study, we used 13
753 transgenic mice expressing the fluorescence calcium indicator GCaMP6f under the Thy1 promoter
754 (GP5.17 mice, Jackson Laboratories, stock no. 025393; Dana et al., 2014). All mice were weaned at 21d,
755 housed 4 to a cage, and maintained on food restriction at 80-85% of their free-feeding weight for the
756 duration of the experiment. The housing room was maintained on a 12h-12h light-dark cycle with the
757 lights coming on at 8am. All experiments were performed between 9am and 6pm. Cages of mice were
758 randomly assigned to experimental groups.

759 GRIN lens implantation surgery for calcium imaging

760 Mice were pretreated with atropine sulfate (0.1mg/kg, ip), anesthetized with chloral hydrate (400mg/kg,
761 ip), injected with dexamethasone (5mg/kg, ip), and placed in a stereotaxic frame. An incision was made
762 to expose and clean the skull, and then a craniotomy was drilled above the right dorsal hippocampus (AP
763 = -2.0mm, ML = 1.5mm from bregma). To gain access to CA1, the cortex and corpus callosum above the
764 hippocampus were gently aspirated while the craniotomy was constantly irrigated with artificial
765 cerebrospinal fluid. A baseplate with an attached 2.0mm GRIN objective lens was slowly lowered into
766 the craniotomy to a depth of 1.5mm below the surface of the skull and secured with jeweler screws and
767 dental cement. The mice were then treated subcutaneously with meloxicam (5mg/kg) for analgesia.
768 Imaging and behavior began 3-6 weeks later when calcium activity was observed.

769 Behavioral apparatus

770 Operant conditioning chambers (Med Associates, VT, USA, ENV-307W-CT) were enclosed in sound
771 attenuating cubicles (ENV-016MD). A tall liquid delivery cup (Med Associates, ENV-303LPHDW-4.25) was
772 recessed in the center of the front wall and a nose poke (Med Associates, ENV-313W) was placed to the
773 left of the liquid delivery cup. Both the liquid delivery cup and the nose poke were outfitted with
774 infrared beams to detect head entries. A tweeter speaker (Med Associate, ENV-324D; pure tone range
775 from 5 - 35 kHz) was mounted above the nose poke. We 3D printed custom face panels to improve
776 functionality for mice with head-mounted microscopes. During imaging, the ceiling of the operant

777 chamber was removed, and a small hole was opened at the top of the sound attenuating chamber to
778 run imaging cables between the microscope and a commutator (NeuroTek Inc., ON, Canada). Operant
779 boxes were controlled with custom code (Med Associates, Med-PC) running on a desktop computer in
780 an adjacent room.

781 Auditory stimuli

782 Mice were exposed to 41 logarithmically spaced tone frequencies that spanned their auditory range.
783 The frequencies for all subjects were (1) 5000, (2) 5249, (3) 5511, (4) 5786, (5) 6074, (6) 6377, (7) 6695,
784 (8) 7028, (9) 7379, (10) 7747, (11) 8133, (12) 8538, (13) 8964, (14) 9411, (15) 9880, (16) 10372, (17)
785 10890, (18) 11432, (19) 12002, (20) 12601, (21) 13229, (22) 13888, (23) 14581, (24) 15307, (25) 16070,
786 (26) 16872, (27) 17713, (28) 18596, (29) 19523, (30) 20496, (31) 21518, (32) 22590, (33) 23716, (34)
787 24899, (35) 26140, (36) 27443, (37) 28811, (38) 30247, (39) 31755, (40) 33338, (41) 35000 kHz. To avoid
788 mice using volume information when learning to discriminate between tones, we applied a volume
789 correction based on R-weighting (Björk et al., 2000) to account for how hearing sensitivity varies across
790 frequencies. The resulting volumes for each tone in decibels were (1) 83, (2) 83, (3) 83, (4) 83, (5) 83, (6)
791 83, (7) 83, (8) 83, (9) 83, (10) 83, (11) 83, (12) 83, (13) 83, (14) 82, (15) 81, (16) 81, (17) 80, (18) 79, (19)
792 79, (20) 78, (21) 77, (22) 76, (23) 76, (24) 75, (25) 74, (26) 74, (27) 73, (28) 73, (29) 73, (30) 72, (31) 72,
793 (32) 73, (33) 74, (34) 74, (35) 75, (36) 76, (37) 77, (38) 78, (39) 81, (40) 83, (41) 83. Tone volume was also
794 randomly modified +/- 5% on each trial. All frequencies and volumes were controlled by a dedicated
795 programmable audio generator (Med Associates, ANL-926).

796 Behavioral procedures

797 *Pretraining*

798 Pretraining involved a series of training days that first shaped the mice to nose poke for 0.04 mL
799 chocolate milk rewards delivered in the adjacent food cup, and then shaped them to wait longer and
800 longer for rarer and rarer rewards. The first day of nose poke training involved placing hungry mice into
801 the operant box, with chocolate milk smeared on the nose poke port and in the liquid delivery cup.
802 Chocolate milk was then automatically delivered at 1 min +/- 30s intervals, while nose pokes were
803 reinforced with chocolate milk delivery in the liquid delivery cup. On subsequent days, chocolate milk
804 delivery was contingent on a nose poke followed by a head entry into the liquid delivery cup within 5s.
805 We shaped the mice to steadily increase the amount of time they were willing to wait for a reward after
806 head entry from 0s to approximately 12s over days by increasing the time between head entry and

807 reward delivery. By the end of pretraining, wait times on every trial were randomly drawn from a
808 distribution based on a negative power distribution (min = 0 s, max = 56.80 s, mean = 6.35 s, median =
809 2.69 s). At the same time that wait times were increasing, the probability of receiving a reward given a
810 sufficient wait response was steadily decreased from 100% to 50%. Pretraining typically lasted between
811 11 and 20 days.

812 *Session order and progression criteria*

813 After pretraining, mice began the task protocol. Mice were trained on six training problems (problems 1-
814 6) interleaved with 7 probe tests (probes 0-6). Training began with a pre-probe (probe 0) to test for any
815 pre-existing tone preferences. This was followed by the first problem (problem 1) and then probe 1,
816 followed by problem 2 and then probe 2, and so on until the last problem (problem 6), which was
817 followed by the final probe (probe 6). Probe sessions were administered on a single day, whereas
818 problem sessions were administered daily until the mice achieved a behavioral criterion (described
819 below). After achieving criterion, mice were given one additional day of training (i.e., the last day of
820 training), after which they progressed to the probe test.

821 *Training problem sessions*

822 During a training problem session, mice learned to wait longer in response to a rich tone than to a poor
823 tone. A trial began when the mouse nose-poked and then entered their head into the food cup within
824 4s, waited until the tone started playing (1 +/- 0.5s after head entry) and then waited for an additional
825 2s (i.e., the minimum wait). After achieving the minimum wait, a mouse could abandon the trial by
826 exiting the food cup or it could wait for a delay of 0-60s for a possible reward. However, successfully
827 waiting for the duration of the delay did not ensure a reward, since rewards were only available on 50%
828 of trials on average. Therefore, mice learned to use the tones to determine how long to wait (i.e., how
829 much time to invest on that trial) for a potential reward. On each trial, the tone was randomly selected
830 from two possible tones: the “rich” tone, which indicated that the trial had a 90% chance of having an
831 available reward, and the “poor” tone, which indicated a 10% chance. The tone frequencies for every
832 problem are shown in **Table 1** (also see **Figures 1B-C**). We measured the animal’s preference for a tone
833 in terms of how long it was willing to wait on trials when that tone was playing but no reward was
834 available. We defined successful discrimination learning (i.e., the behavioral criterion) as significantly
835 longer waits in response to the rich tone than to the poor tone as indicated by a significant t-test with an
836 alpha level of 0.01 (see **Supplemental figure 1A-B**). Individual sessions lasted for 100 trials or 1 hour,

837 whichever came first. Sessions with fewer than 41 trials were repeated. The number of days required to
838 learn a problem varied with the prior experience of the animal and can be seen in **Figure 1F**.

| Problem | Structured, rich | Structured, poor | Unstructured, rich | Unstructured, poor |
|---------|------------------|------------------|--------------------|--------------------|
| 1 | 13229 | 23716 | 13229 | 23716 |
| 2 | 16872 | 8538 | 28811 | 10890 |
| 3 | 11432 | 27443 | 8964 | 33338 |
| 4 | 15307 | 6377 | 19523 | 5249 |
| 5 | 10372 | 7379 | 6074 | 16070 |
| 6 | 13229 | 7379 | 13229 | 7379 |

839

840 **Table 1.** The tone frequencies used for each training problem. All frequencies are shown in kHz. The
841 tones used in problem 1 were common between both groups, as were the tones used in problem 6
842 (both rows are shaded). To control for presentation order, half of the subjects were trained on the
843 problems in reverse order.

844 *Probe test sessions*

845 Probe test sessions were identical to training problem sessions except that during a probe test session,
846 mice were exposed to all 41 tones used in the study to test their preferences. Tones were presented in
847 random order, twice each, such that a reward was available on exactly one trial for each tone, resulting
848 in an average reward probability of 50% (the same as in the problem training sessions). There was no
849 behavioral criterion during probe test sessions. Individual sessions lasted for 82 trials or 1 hour,
850 whichever came first. Each probe test (0-6) was only administered once unless the mice performed
851 fewer than 41 trials, at which point the session was repeated.

852 Calcium imaging

853 Imaging mice were acclimated to the weight of the imaging scope and cables during pretraining by
854 slowly introducing heavier and heavier dummy scopes as the mice learned to nose poke for a reward.
855 During the training protocol, we only collected video data during probe tests and during the first and last
856 session of each training problem. To minimize the effect of bleaching over the course of a session, the
857 LED light was kept off except for the period on each trial between nose poke and two seconds after
858 reward delivery. If a mouse abandoned a trial before reward delivery, then the LED was turned off as

859 soon as they left the food cup. The microscope LED was controlled by the MED-PC code so that the ON
860 and OFF times were synchronized with task performance.

861 To encourage the imaging mice to perform a sufficient number of trials on every session, the
862 probe test and training problem sessions were modified so that rewards were available on two thirds of
863 trials (instead of half) by doubling the number of reward available trials (on probe tests) and by
864 increasing the number of rich tone trials (on training sessions). Additionally, the delay times between
865 minimum wait and reward delivery were shortened by approximately 2s on average (distribution min =
866 0s, max = 54.63s, mean = 4.22s, median = 1.49s).

867 Histology

868 Lens placement was determined post-mortem. Mice were transcardially perfused with 4%
869 paraformaldehyde in phosphate buffered saline. Brains were removed and stored for at least 24 h in 4%
870 paraformaldehyde before being transferred to a 30% sucrose solution for storage until they could be
871 sliced at 50 μ m, DAPI stained, and examined.

872 Data analysis

873 *Subject-level comparisons*

874 We compared behavioral responses and reactivation scores using standard approaches. We used t-tests
875 to evaluate differences between two groups at a single time point, between one group at two time
876 points, or to compare one group at a single time point to zero. We used one-way repeated measures
877 ANOVA tests to evaluate the inequality of one group at multiple time points, and mixed ANOVA tests to
878 evaluate the interaction between two groups at multiple time points. In cases where the groups were
879 unbalanced, such as in the imaging conditions, we used linear mixed effects models in place of mixed
880 ANOVA tests. All statistics were performed in MATLAB (MathWorks, Natick, MA).

881 *Modeling probe test responses*

882 We fit a model to probe wait times for every subject to estimate their preferences for every tone. This
883 allowed us to estimate preferences similarly for all subjects even if a subject did not provide data about
884 their preference for a particular tone of interest (for example, if they did not run enough trials to sample
885 all tones). To do this we used the *nlinfit* function in MATLAB to fit a curve to the wait times of each
886 subject on each probe. We chose a linear combination of a normal and a logistic function (**Supplemental**
887 **figure 1H**) as these described the various curve shapes observed in this study (**Figure 2A-B**). To ensure

888 that we fit the most parsimonious model to each session, we first fit four separate models (intercept-
889 only, intercept + logistic, intercept + normal, and intercept + logistic + normal) and then compared them
890 by iteratively withholding a single data point from a session, fitting each model to the remaining wait
891 times, and then measuring the error between each model and the withheld data point. After computing
892 the error rates of each model, we selected the model with the lowest mean error rate as follows: we
893 first selected a partial model by comparing the intercept + normal vs intercept + logistic models and kept
894 the one with the lower error rate; we then compared the partial model vs the intercept-only model, and
895 if the intercept-only model had the lower error rate we selected the intercept-only model; however, if
896 the partial model had the lower error rate, then we compared the partial model vs the full (intercept +
897 logistic + normal) model and selected the model with the lower error rate. In this way, we only selected
898 the full model if it was more effective at predicting withheld data than any of the simpler models.

899 The intercept model was simply the mean wait time response in seconds. The intercept + logistic
900 model was defined as

$$901 \quad \alpha + \frac{\beta}{1 + e^{\lambda - x}}$$

902 where α is the intercept, β is the height of the curve (in seconds), λ is the midpoint of the curve (the
903 tone number, any value from 1 to 41). The intercept + normal model was written to constrain the width
904 of the curve (in number of tones) so that it was proportionate to the height, defined as

$$905 \quad \alpha + \theta e^{-\frac{(x-\psi)^2}{\theta^3/2}}$$

906 where α is the intercept, θ is the height of the curve and is proportional to the width of the curve, and ψ
907 is the midpoint of the curve. The intercept + logistic + normal model was defined as

$$908 \quad \alpha + \frac{\beta}{1 + e^{\lambda - x}} + \theta e^{-\frac{(x-\psi)^2}{\theta^3/2}}$$

909 where α is the intercept, β is the height of the logistic curve, λ is the midpoint of the logistic curve, θ is
910 the height and determines the width of the normal curve, and ψ is the midpoint of the normal curve.

911 *Preprocessing calcium imaging data*

912 Each imaging session produced as many videos as there were trials, along with the timestamps for each
913 frame. We concatenated the imaging videos into a single session video using custom Python software,
914 and then processed the resulting session video using CNMFe (Zhou et al., 2018) to identify the spatial

915 footprints of, and to extract calcium signals from, individual neurons. The resulting calcium signals were
916 then linearly up sampled to 100 Hz to match the sampling rate of the behavioral hardware.

917 *Cell registration across sessions*

918 In order to determine which neurons were active in each session, we used CellReg (Sheintuch et al.,
919 2017) to align and register neuronal spatial footprints across all 19 imaging sessions for each mouse.
920 Gross rotational differences between session videos, identified by comparing vasculature observable in
921 maximum projection images, were manually corrected using custom rotation software written in
922 MATLAB. We then registered the resulting footprints using rigid rotation and transformation in CellReg
923 (**Supplemental figure 3**).

924 *Time warping fluorescence signals*

925 We obtained calcium signals from each trial for the series of events starting with nose poke and ending
926 either when the mouse left the food cup or 2 s after reward delivery. To correct for slight differences in
927 the time between each event on different trials, we linearly interpolated the observed calcium signals
928 from each trial into a universal trial timeline with fixed durations between each event. The fixed
929 durations (as seen in **Figure 5**) were approximately the median observed time between each event.

930 *Mean activity correlations*

931 We determined the similarity of neuronal activity patterns across pairs of sessions by first finding the
932 mean time-warped activity pattern for each cell across all trials within a session and then computing the
933 correlation between the mean activity patterns of a neuron in one session and the mean activity pattern
934 of that same neuron in a second session. This analysis only included the portion of each trial that was
935 common to both rewarded and unrewarded trials. The included events were the nose poke exit, the
936 head entry into the food cup, the moment the tone on came on, and the end of the minimum wait time
937 (2 s after tone on). Lastly, we computed the mean activity correlation between two sessions as the
938 average correlation among all of the neurons that were active in both sessions.

939 *Reactivation scores*

940 We measured the reactivation of hippocampal ensemble activity in terms of the proportion of neurons
941 that were active in two different sessions (proportion overlap) and how similar the activity of each
942 neuron was across the two sessions (mean activity correlation). To do this, we used a reactivation score
943 computed as

944

$$R = O + (O * A)$$

945 where R is the reactivation score, O is the proportion overlap, and A is the mean activity correlation. A
946 score of 0 will occur if there is no overlap between the two sessions (no neurons are active in both
947 sessions) or if all of the overlapping neurons have perfectly anticorrelated activity patterns between the
948 two sessions. A score of 2 will only occur if there is perfect overlap (the exact same population of
949 neurons is active in both sessions) and all of the neurons have perfectly correlated activity patterns.

950 *Correcting for time between imaging sessions*

951 Because the Structured training group learned new problems faster than the Unstructured training
952 group (**Supplemental figure 2D-F**), there were slight differences in the number of days between the
953 corresponding imaging sessions of each group (**Supplemental figure 4A**). We controlled for this in all of
954 our measures of population similarity (population overlap, mean activity correlations, and reactivation
955 scores) with the same general approach. We first computed the population similarity measure between
956 every pair of imaging sessions within each Unstructured training group mouse and then pooled all of
957 these values together. We only used Unstructured training sessions because these sessions showed a
958 clear temporal trend that was unadulterated by learning related effects, and because the Structured
959 training group did not have any long (> 50 days) durations between imaging sessions. However, similar
960 results were obtained if we used sessions from both groups. We then estimated the effect of days
961 between imaging sessions by fitting a truncated normal curve to the scatterplot of population similarity
962 measurements using the *nlinfit* function in MATLAB with parameters for the intercept, height, width,
963 and mid-point of the curve. For the mean activity correlation model, we assumed an intercept of 0 (no
964 correlation between activity states as the time between imaging sessions approaches infinity) and
965 therefore only fit parameters for height, width, and mid-point. Fit curves can be seen for population
966 overlap and mean activity correlations in **Supplemental figures 4B-C** and for reactivation scores in
967 **Supplemental figure 5A**. We then computed the new (time-corrected) population similarity
968 measurements as the original measurement minus the value of the fit normal curve at the
969 corresponding x-value (days between imaging sessions). This resulted in both positive and negative
970 values, with negative values corresponding to population similarity below what was expected by the
971 number of days between the two imaging sessions.

972

973

974 **References**

975 Baraduc, P., Duhamel, J.-R., and Wirth, S. (2019). Schema cells in the macaque hippocampus. *Science*
976 363, 635–639. <https://doi.org/10.1126/science.aav5404>.

977 Barron, H.C., Reeve, H.M., Koolschijn, R.S., Perestenko, P.V., Shpektor, A., Nili, H., Rothaermel, R.,
978 Campo-Urriza, N., O'Reilly, J.X., Bannerman, D.M., et al. (2020). Neuronal Computation Underlying
979 Inferential Reasoning in Humans and Mice. *Cell* 183, 228-243.e21.
980 <https://doi.org/10.1016/j.cell.2020.08.035>.

981 Bartlett, F.C. (1932). Remembering: A study in experimental and social psychology (New York, NY, US:
982 Cambridge University Press).

983 Behrens, T.E.J., Muller, T.H., Whittington, J.C.R., Mark, S., Baram, A.B., Stachenfeld, K.L., and Kurth-
984 Nelson, Z. (2018). What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron* 100,
985 490–509. <https://doi.org/10.1016/j.neuron.2018.10.002>.

986 Björk, E., Nevalainen, T., Hakumäki, M., and Voipio, H.M. (2000). R-weighting provides better estimation
987 for rat hearing sensitivity. *Lab Anim* 34, 136–144. <https://doi.org/10.1258/002367700780457518>.

988 Bransford, J.D., and Johnson, M.K. (1972). Contextual prerequisites for understanding: Some
989 investigations of comprehension and recall. *Journal of Verbal Learning and Verbal Behavior* 11, 717–726.
990 .

991 Brunec, I.K., and Momennejad, I. (2022). Predictive Representations in Hippocampal and Prefrontal
992 Hierarchies. *J Neurosci* 42, 299–312. <https://doi.org/10.1523/JNEUROSCI.1327-21.2021>.

993 Bulkin, D.A., Law, L.M., and Smith, D.M. (2016). Placing memories in context: Hippocampal
994 representations promote retrieval of appropriate memories. *Hippocampus* 26, 958–971.
995 <https://doi.org/10.1002/hipo.22579>.

996 Butterly, D.A., Petroccione, M.A., and Smith, D.M. (2012). Hippocampal context processing is critical for
997 interference free recall of odor memories in rats. *Hippocampus* 22, 906–913.
998 <https://doi.org/10.1002/hipo.20953>.

999 Cai, D.J., Aharoni, D., Shuman, T., Shobe, J., Biane, J., Song, W., Wei, B., Veshkini, M., La-Vu, M., Lou, J.,
1000 et al. (2016). A shared neural ensemble links distinct contextual memories encoded close in time. *Nature*
1001 534, 115–118. <https://doi.org/10.1038/nature17955>.

1002 Chase, W.G., and Simon, H.A. (1973). Perception in chess. *Cognitive Psychology* 4, 55–81.
1003 [https://doi.org/10.1016/0010-0285\(73\)90004-2](https://doi.org/10.1016/0010-0285(73)90004-2).

1004 Chowdhury, A., Luchetti, A., Fernandes, G., Filho, D.A., Kastellakis, G., Tzilivaki, A., Ramirez, E.M., Tran,
1005 M.Y., Poirazi, P., and Silva, A.J. (2022). A locus coeruleus-dorsal CA1 dopaminergic circuit modulates
1006 memory linking. *Neuron* S0896-6273(22)00707-3. <https://doi.org/10.1016/j.neuron.2022.08.001>.

1007 Craik, F.I.M., and Tulving, E. (1975). Depth of processing and the retention of words in episodic memory.
1008 *Journal of Experimental Psychology: General* 104, 268–294. [3445.104.3.268](https://doi.org/10.1037/0096-
1009 3445.104.3.268).

1010 Dana, H., Chen, T.-W., Hu, A., Shields, B.C., Guo, C., Looger, L.L., Kim, D.S., and Svoboda, K. (2014). Thy1-
1011 GCaMP6 Transgenic Mice for Neuronal Population Imaging In Vivo. *PLOS ONE* 9, e108697.
1012 <https://doi.org/10.1371/journal.pone.0108697>.

1013 Davis, H.P., and Squire, L.R. (1984). Protein synthesis and memory: a review. *Psychol Bull* 96, 518–559. .

1014 Debiec, J., LeDoux, J.E., and Nader, K. (2002). Cellular and systems reconsolidation in the hippocampus.
1015 *Neuron* 36, 527–538. [https://doi.org/10.1016/s0896-6273\(02\)01001-2](https://doi.org/10.1016/s0896-6273(02)01001-2).

1016 Dragoi, G., and Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular
1017 assemblies. *Nature* 469, 397–401. <https://doi.org/10.1038/nature09633>.

1018 Epsztein, J., Brecht, M., and Lee, A.K. (2011). Intracellular determinants of hippocampal CA1 place and
1019 silent cell activity in a novel environment. *Neuron* 70, 109–120.
1020 <https://doi.org/10.1016/j.neuron.2011.03.006>.

1021 Fernández, R.S., Boccia, M.M., and Pedreira, M.E. (2016). The fate of memory: Reconsolidation and the
1022 case of Prediction Error. *Neurosci Biobehav Rev* 68, 423–441.
1023 <https://doi.org/10.1016/j.neubiorev.2016.06.004>.

1024 Fuhs, M.C., and Touretzky, D.S. (2007). Context learning in the rodent hippocampus. *Neural Comput* 19,
1025 3173–3215. <https://doi.org/10.1162/neco.2007.19.12.3173>.

1026 Gershman, S.J., and Niv, Y. (2010). Learning latent structure: carving nature at its joints. *Curr Opin
1027 Neurobiol* 20, 251–256. <https://doi.org/10.1016/j.conb.2010.02.008>.

1028 Gill, P.R., Mizumori, S.J.Y., and Smith, D.M. (2011). Hippocampal episode fields develop with learning.
1029 *Hippocampus* 21, 1240–1249. <https://doi.org/10.1002/hipo.20832>.

1030 Girardeau, G., Benchenane, K., Wiener, S.I., Buzsáki, G., and Zugaro, M.B. (2009). Selective suppression
1031 of hippocampal ripples impairs spatial memory. *Nat Neurosci* 12, 1222–1223.
1032 <https://doi.org/10.1038/nn.2384>.

1033 Gisquet-Verrier, P., and Riccio, D.C. (2018). Memory integration: An alternative to the
1034 consolidation/reconsolidation hypothesis. *Prog Neurobiol* 171, 15–31.
1035 <https://doi.org/10.1016/j.pneurobio.2018.10.002>.

1036 Goshen, I., Brodsky, M., Prakash, R., Wallace, J., Gradinaru, V., Ramakrishnan, C., and Deisseroth, K.
1037 (2011). Dynamics of retrieval strategies for remote memories. *Cell* 147, 678–689.
1038 <https://doi.org/10.1016/j.cell.2011.09.033>.

1039 Harlow, H.F. (1949). The formation of learning sets. *Psychological Review* 56, 51–65.
1040 <https://doi.org/10.1037/h0062474>.

1041 Jacob, A.D., Ramsaran, A.I., Mocle, A.J., Tran, L.M., Yan, C., Frankland, P.W., and Josselyn, S.A. (2018). A
1042 Compact Head-Mounted Endoscope for In Vivo Calcium Imaging in Freely Behaving Mice. *Curr Protoc
1043 Neurosci* 84, e51. <https://doi.org/10.1002/cpns.51>.

1044 Kelemen, E., and Fenton, A.A. (2010). Dynamic grouping of hippocampal neural activity during cognitive
1045 control of two spatial frames. *PLoS Biol* 8, e1000403. <https://doi.org/10.1371/journal.pbio.1000403>.

1046 van Kesteren, M.T.R., Ruiter, D.J., Fernández, G., and Henson, R.N. (2012). How schema and novelty
1047 augment memory formation. *Trends Neurosci* 35, 211–219. <https://doi.org/10.1016/j.tins.2012.02.001>.

1048 Knudsen, E.B., and Wallis, J.D. (2021). Hippocampal neurons construct a map of an abstract value space.
1049 *Cell* 184, 4640-4650.e10. <https://doi.org/10.1016/j.cell.2021.07.010>.

1050 Kumaran, D., Hassabis, D., and McClelland, J.L. (2016). What Learning Systems do Intelligent Agents
1051 Need? Complementary Learning Systems Theory Updated. *Trends Cogn Sci* 20, 512–534.
1052 <https://doi.org/10.1016/j.tics.2016.05.004>.

1053 Lee, J.L.C. (2009). Reconsolidation: maintaining memory relevance. *Trends Neurosci* 32, 413–420.
1054 <https://doi.org/10.1016/j.tins.2009.05.002>.

1055 Lee, J.L.C., Nader, K., and Schiller, D. (2017). An Update on Memory Reconsolidation Updating. *Trends
1056 Cogn Sci* 21, 531–545. <https://doi.org/10.1016/j.tics.2017.04.006>.

1057 Liu, X., Ramirez, S., Pang, P.T., Puryear, C.B., Govindarajan, A., Deisseroth, K., and Tonegawa, S. (2012).
1058 Optogenetic stimulation of a hippocampal engram activates fear memory recall. *Nature* 484, 381–385.
1059 <https://doi.org/10.1038/nature11028>.

1060 Mack, M.L., Love, B.C., and Preston, A.R. (2018). Building concepts one episode at a time: The
1061 hippocampus and concept formation. *Neurosci Lett* 680, 31–38.
1062 <https://doi.org/10.1016/j.neulet.2017.07.061>.

1063 Makino, H., and Komiyama, T. (2015). Learning enhances the relative impact of top-down processing in
1064 the visual cortex. *Nat Neurosci* 18, 1116–1122. <https://doi.org/10.1038/nn.4061>.

1065 Markus, E.J., Qin, Y.L., Leonard, B., Skaggs, W.E., McNaughton, B.L., and Barnes, C.A. (1995). Interactions
1066 between location and task affect the spatial and directional firing of hippocampal neurons. *J. Neurosci.*
1067 15, 7079–7094. <https://doi.org/10.1523/JNEUROSCI.15-11-07079.1995>.

1068 Mau, W., Hasselmo, M.E., and Cai, D.J. (2020). The brain in motion: How ensemble fluidity drives
1069 memory-updating and flexibility. *eLife* 9, e63550. <https://doi.org/10.7554/eLife.63550>.

1070 McCloskey, M., and Cohen, N.J. (1989). Catastrophic Interference in Connectionist Networks: The
1071 Sequential Learning Problem. In *Psychology of Learning and Motivation*, G.H. Bower, ed. (Academic
1072 Press), pp. 109–165.

1073 McKenzie, S., Frank, A.J., Kinsky, N.R., Porter, B., Rivière, P.D., and Eichenbaum, H. (2014). Hippocampal
1074 representation of related and opposing memories develop within distinct, hierarchically organized
1075 neural schemas. *Neuron* 83, 202–215. <https://doi.org/10.1016/j.neuron.2014.05.019>.

1076 McKenzie, S., Huszár, R., English, D.F., Kim, K., Christensen, F., Yoon, E., and Buzsáki, G. (2021).
1077 Preexisting hippocampal network dynamics constrain optogenetically induced place fields. *Neuron* 109,
1078 1040-1054.e7. <https://doi.org/10.1016/j.neuron.2021.01.011>.

1079 Miller, A.M.P., Mau, W., and Smith, D.M. (2019). Retrosplenial Cortical Representations of Space and
1080 Future Goal Locations Develop with Learning. *Curr Biol* 29, 2083-2090.e4.
1081 <https://doi.org/10.1016/j.cub.2019.05.034>.

1082 Momennejad, I. (2020). Learning Structures: Predictive Representations, Replay, and Generalization.
1083 *Curr Opin Behav Sci* 32, 155–166. <https://doi.org/10.1016/j.cobeha.2020.02.017>.

1084 Morton, N.W., and Preston, A.R. (2021). Concept formation as a computational cognitive process. *Curr
1085 Opin Behav Sci* 38, 83–89. <https://doi.org/10.1016/j.cobeha.2020.12.005>.

1086 Muller, R.U., and Kubie, J.L. (1987). The effects of changes in the environment on the spatial firing of
1087 hippocampal complex-spike cells. *J. Neurosci.* 7, 1951–1968. <https://doi.org/10.1523/JNEUROSCI.07-07-01951.1987>.

1089 Nieh, E.H., Schottdorf, M., Freeman, N.W., Low, R.J., Lewallen, S., Koay, S.A., Pinto, L., Gauthier, J.L.,
1090 Brody, C.D., and Tank, D.W. (2021). Geometry of abstract learned knowledge in the hippocampus.
1091 *Nature* 595, 80–84. <https://doi.org/10.1038/s41586-021-03652-7>.

1092 Niv, Y. (2019). Learning task-state representations. *Nat Neurosci* 22, 1544–1553.
1093 <https://doi.org/10.1038/s41593-019-0470-8>.

1094 Park, A.J., Harris, A.Z., Martyniuk, K.M., Chang, C.-Y., Abbas, A.I., Lowes, D.C., Kellendonk, C., Gogos, J.A.,
1095 and Gordon, J.A. (2021). Reset of hippocampal–prefrontal circuitry facilitates learning. *Nature* 591, 615–
1096 619. <https://doi.org/10.1038/s41586-021-03272-1>.

1097 Payne, J.D., Schacter, D.L., Propper, R.E., Huang, L.-W., Wamsley, E.J., Tucker, M.A., Walker, M.P., and
1098 Stickgold, R. (2009). The role of sleep in false memory formation. *Neurobiology of Learning and Memory*
1099 92, 327–334. <https://doi.org/10.1016/j.nlm.2009.03.007>.

1100 Plitt, M.H., and Giocomo, L.M. (2021). Experience-dependent contextual codes in the hippocampus. *Nat
1101 Neurosci* 24, 705–714. <https://doi.org/10.1038/s41593-021-00816-6>.

1102 Pudhiyidath, A., Morton, N.W., Viveros Duran, R., Schapiro, A.C., Momennejad, I., Hinojosa-Rowland,
1103 D.M., Molitor, R.J., and Preston, A.R. (2022). Representations of Temporal Community Structure in
1104 Hippocampus and Precuneus Predict Inductive Reasoning Decisions. *Journal of Cognitive Neuroscience*
1105 34, 1736–1760. https://doi.org/10.1162/jocn_a_01864.

1106 Purtle, R.B. (1973). Peak shift: A review. *Psychological Bulletin* 80, 408–421.
1107 <https://doi.org/10.1037/h0035233>.

1108 Rashid, A.J., Yan, C., Mercaldo, V., Hsiang, H.-L.L., Park, S., Cole, C.J., De Cristofaro, A., Yu, J.,
1109 Ramakrishnan, C., Lee, S.Y., et al. (2016). Competition between engrams influences fear memory
1110 formation and recall. *Science* 353, 383–387. <https://doi.org/10.1126/science.aaf0594>.

1111 Redish, A.D., Jensen, S., Johnson, A., and Kurth-Nelson, Z. (2007). Reconciling reinforcement learning
1112 models with behavioral extinction and renewal: implications for addiction, relapse, and problem
1113 gambling. *Psychol Rev* 114, 784–805. <https://doi.org/10.1037/0033-295X.114.3.784>.

1114 Richards, B.A., Xia, F., Santoro, A., Husse, J., Woodin, M.A., Josselyn, S.A., and Frankland, P.W. (2014).
1115 Patterns across multiple memories are identified over time. *Nat Neurosci* 17, 981–986.
1116 <https://doi.org/10.1038/nn.3736>.

1117 Robinson, N.T.M., Descamps, L.A.L., Russell, L.E., Buchholz, M.O., Bicknell, B.A., Antonov, G.K., Lau,
1118 J.Y.N., Nutbrown, R., Schmidt-Hieber, C., and Häusser, M. (2020). Targeted Activation of Hippocampal
1119 Place Cells Drives Memory-Guided Spatial Behavior. *Cell* 183, 1586–1599.e10.
1120 <https://doi.org/10.1016/j.cell.2020.09.061>.

1121 Sadtler, P.T., Quick, K.M., Golub, M.D., Chase, S.M., Ryu, S.I., Tyler-Kabara, E.C., Yu, B.M., and Batista,
1122 A.P. (2014). Neural constraints on learning. *Nature* 512, 423–426. <https://doi.org/10.1038/nature13665>.

1123 Samborska, V., Butler, J., Walton, M., Behrens, T.E.J., and Akam, T. (2021). Complementary task
1124 representations in hippocampus and prefrontal cortex for generalising the structure of problems.
1125 2021.03.05.433967. <https://doi.org/10.1101/2021.03.05.433967>.

1126 Sanders, H., Wilson, M.A., and Gershman, S.J. (2020). Hippocampal remapping as hidden state inference.
1127 *eLife* 9, e51140. <https://doi.org/10.7554/eLife.51140>.

1128 Schapiro, A.C., Turk-Browne, N.B., Norman, K.A., and Botvinick, M.M. (2016). Statistical learning of
1129 temporal community structure in the hippocampus. *Hippocampus* 26, 3–8.
1130 <https://doi.org/10.1002/hipo.22523>.

1131 Schlichting, M.L., and Frankland, P.W. (2017). Memory allocation and integration in rodents and
1132 humans. *Current Opinion in Behavioral Sciences* 17, 90–98.
1133 <https://doi.org/10.1016/j.cobeha.2017.07.013>.

1134 Sheintuch, L., Rubin, A., Brande-Eilat, N., Geva, N., Sadeh, N., Pinchasof, O., and Ziv, Y. (2017). Tracking
1135 the Same Neurons across Multiple Days in Ca2+ Imaging Data. *Cell Rep* 21, 1102–1115.
1136 <https://doi.org/10.1016/j.celrep.2017.10.013>.

1137 Smith, D.M., and Mizumori, S.J.Y. (2006). Learning-related development of context-specific neuronal
1138 responses to places and events: the hippocampal role in context processing. *J Neurosci* 26, 3154–3163.
1139 <https://doi.org/10.1523/JNEUROSCI.3234-05.2006>.

1140 Smith, S.M., and Vela, E. (2001). Environmental context-dependent memory: A review and meta-
1141 analysis. *Psychonomic Bulletin & Review* 8, 203–220. <https://doi.org/10.3758/BF03196157>.

1142 Squire, L.R. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology
1143 of Learning and Memory* 82, 171–177. <https://doi.org/10.1016/j.nlm.2004.06.005>.

1144 Sun, C., Yang, W., Martin, J., and Tonegawa, S. (2020). Hippocampal neurons represent events as
1145 transferable units of experience. *Nat Neurosci* 23, 651–663. <https://doi.org/10.1038/s41593-020-0614-x>.

1147 Takehara-Nishiuchi, K., and McNaughton, B.L. (2008). Spontaneous changes of neocortical code for
1148 associative memory during consolidation. *Science* 322, 960–963.
1149 <https://doi.org/10.1126/science.1161299>.

1150 Tanaka, K.Z., Pevzner, A., Hamidi, A.B., Nakazawa, Y., Graham, J., and Wiltgen, B.J. (2014). Cortical
1151 Representations Are Reinstated by the Hippocampus during Memory Retrieval. *Neuron* **84**, 347–354.
1152 <https://doi.org/10.1016/j.neuron.2014.09.037>.

1153 Tervo, D.G.R., Tenenbaum, J.B., and Gershman, S.J. (2016). Toward the neural implementation of
1154 structure learning. *Curr Opin Neurobiol* **37**, 99–105. <https://doi.org/10.1016/j.conb.2016.01.014>.

1155 Tolman, E.C. (1948). Cognitive maps in rats and men. *Psychol Rev* **55**, 189–208.
1156 <https://doi.org/10.1037/h0061626>.

1157 Tse, D., Langston, R.F., Kakeyama, M., Bethus, I., Spooner, P.A., Wood, E.R., Witter, M.P., and Morris,
1158 R.G.M. (2007). Schemas and memory consolidation. *Science* **316**, 76–82.
1159 <https://doi.org/10.1126/science.1135935>.

1160 Tse, D., Takeuchi, T., Kakeyama, M., Kajii, Y., Okuno, H., Tohyama, C., Bito, H., and Morris, R.G.M. (2011).
1161 Schema-Dependent Gene Activation and Memory Encoding in Neocortex. *Science* **333**, 891–895.
1162 <https://doi.org/10.1126/science.1205274>.

1163 Vikbladh, O.M., Meager, M.R., King, J., Blackmon, K., Devinsky, O., Shohamy, D., Burgess, N., and Daw,
1164 N.D. (2019). Hippocampal Contributions to Model-Based Planning and Spatial Memory. *Neuron* **102**,
1165 683-693.e4. <https://doi.org/10.1016/j.neuron.2019.02.014>.

1166 Whittington, J.C.R., Muller, T.H., Mark, S., Barry, C., and Behrens, T.E.J. (2018). Generalisation of
1167 structural knowledge in the hippocampal-entorhinal system.
1168 <https://doi.org/10.48550/arXiv.1805.09042>.

1169 Whittington, J.C.R., Muller, T.H., Mark, S., Chen, G., Barry, C., Burgess, N., and Behrens, T.E.J. (2020). The
1170 Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the
1171 Hippocampal Formation. *Cell* **183**, 1249-1263.e23. <https://doi.org/10.1016/j.cell.2020.10.024>.

1172 Whittington, J.C.R., McCaffary, D., Bakermans, J.J.W., and Behrens, T.E.J. (2022). How to build a
1173 cognitive map: insights from models of the hippocampal formation.
1174 <https://doi.org/10.48550/arXiv.2202.01682>.

1175 Wiltgen, B.J., and Silva, A.J. (2007). Memory for context becomes less specific with time. *Learn. Mem.*
1176 **14**, 313–317. <https://doi.org/10.1101/lm.430907>.

1177 Wood, E.R., Dudchenko, P.A., Robitsek, R.J., and Eichenbaum, H. (2000). Hippocampal neurons encode
1178 information about different types of memory episodes occurring in the same location. *Neuron* **27**, 623–
1179 633. [https://doi.org/10.1016/s0896-6273\(00\)00071-4](https://doi.org/10.1016/s0896-6273(00)00071-4).

1180 Zeithamova, D., Dominick, A.L., and Preston, A.R. (2012). Hippocampal and ventral medial prefrontal
1181 activation during retrieval-mediated learning supports novel inference. *Neuron* **75**, 168–179.
1182 <https://doi.org/10.1016/j.neuron.2012.05.010>.

1183 Zhou, P., Resendez, S.L., Rodriguez-Romaguera, J., Jimenez, J.C., Neufeld, S.Q., Giovannucci, A., Friedrich,
1184 J., Pnevmatikakis, E.A., Stuber, G.D., Hen, R., et al. (2018). Efficient and accurate extraction of in vivo
1185 calcium signals from microendoscopic video data. *eLife* **7**, e28728. <https://doi.org/10.7554/eLife.28728>.

