

Patterns across multiple memories are identified over time

Blake A Richards¹⁻⁶, Frances Xia^{1,3}, Adam Santoro^{1,4}, Jana Husse¹⁻⁴, Melanie A Woodin⁶, Sheena A Josselyn¹⁻⁴ & Paul W Frankland¹⁻⁴

Memories are not static but continue to be processed after encoding. This is thought to allow the integration of related episodes via the identification of patterns. Although this idea lies at the heart of contemporary theories of systems consolidation, it has yet to be demonstrated experimentally. Using a modified water-maze paradigm in which platforms are drawn stochastically from a spatial distribution, we found that mice were better at matching platform distributions 30 d compared to 1 d after training. Post-training time-dependent improvements in pattern matching were associated with increased sensitivity to new platforms that conflicted with the pattern. Increased sensitivity to pattern conflict was reduced by pharmacogenetic inhibition of the medial prefrontal cortex (mPFC). These results indicate that pattern identification occurs over time, which can lead to conflicts between new information and existing knowledge that must be resolved, in part, by computations carried out in the mPFC.

Learning does not happen in a vacuum. Rather it occurs against the backdrop of an animal's cumulative experience. Related experiences form knowledge structures referred to as schemata¹⁻³ or categories^{4,5}, which then influence how new information is encoded. The degree to which new experiences match existing structures determines which brain networks are engaged and the speed of memory consolidation⁶⁻¹⁰. In particular, in both humans and rodents the mPFC appears to be important for mediating conflicts between new and old experiences and updating existing schemata^{6,7,9,11}.

Psychological and computational accounts suggest that the identification of commonalities (i.e., patterns) across experiences is necessary for the construction of schemata and categories^{3,12,13}. This process of identifying regularities in data may occur rapidly as patterns are encountered¹⁴ as well as in the immediate post-encoding rest period¹⁵⁻¹⁸. However, it has been hypothesized that a similar process persists over more extended time periods (e.g., weeks or months) after initial encoding is complete. This long-standing idea^{19,20} forms the basis of contemporary theories of systems consolidation²¹⁻²⁵, which propose that detailed episodic memories stored as individual patterns in the auto-associative networks of the hippocampus are transformed into semantic or schematized memories stored using distributed representations in the neocortex. Consistent with this, previous research has demonstrated that memories can generalize over time^{26,27} and that remote memories may be less detailed²⁸. However, there is no direct evidence that pattern identification persists over an extended time period after initial encoding. Moreover, how post-encoding pattern identification then influences new learning is not known.

To address these issues we developed a new water-maze paradigm in which mice are exposed to platform locations drawn stochastically

from a specific spatial distribution. This design allowed us to assess the extent to which mice could match the pattern defined by the spatial distribution 1 d or 30 d after training. Further, we examined how mice reacted to a new training platform location when it was either consistent or in conflict with the pattern. We found that after training was complete, the passage of time improved the match between mouse search paths and the pattern presented during training. This was associated with delay-dependent differences in the learning of new platforms, with increased sensitivity to conflicting platform locations 30 d after training. Increased sensitivity to a conflict between previous patterns and new information was reduced by pharmacogenetic inhibition of the mPFC. These results indicate that a continued process of pattern identification occurs over time, which can lead to conflicts between new information and existing knowledge structures that must be resolved, in part, by computations carried out in the mPFC.

RESULTS

The passage of time improves pattern matching

According to contemporary theories of systems consolidation, the passage of time allows for the identification of patterns across multiple experiences (Fig. 1a). Whereas one consequence of this process is that memory of details for specific experiences may fade with time, one benefit might be that the consolidated 'composite' memory is a better reflection of the cumulative experience^{22,23}. To assess this, we used a modified version of the delayed matching-to-place water-maze task²⁹. On each training day, we presented mice with four opportunities to find a platform in a fixed location. However, across days we moved the location of the platform, with each specific daily location selected

¹Program in Neurosciences and Mental Health, The Hospital for Sick Children, Toronto, Ontario, Canada. ²Department of Psychology, University of Toronto, Toronto, Ontario, Canada. ³Department of Physiology, University of Toronto, Toronto, Ontario, Canada. ⁴Institute of Medical Science, University of Toronto, Toronto, Ontario, Canada. ⁵Department of Biological Sciences, University of Toronto Scarborough, Toronto, Ontario, Canada. ⁶Department of Cell and Systems Biology, University of Toronto, Toronto, Ontario, Canada. Correspondence should be addressed to P.W.F. (paul.frankland@sickkids.ca).

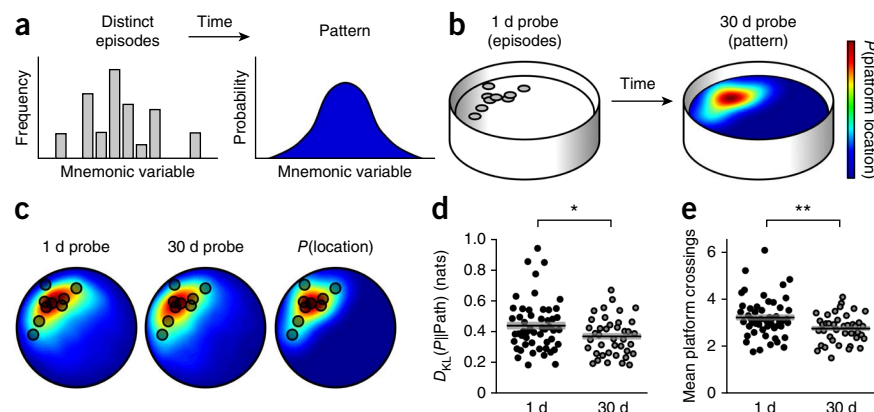
Received 29 January; accepted 8 May; published online 1 June 2014; doi:10.1038/nn.3736

Figure 1 Passage of time improves behavioral match to pattern. (a) Diagram illustrating hypothesized post-encoding transformation of memories from information about distinct episodes to information about the overall statistical patterns in the experiences.

(b) Schematic of the water-maze task we designed to test for post-encoding pattern identification, with gray circles indicating specific platform locations drawn from a probability distribution, and the heat map showing the distribution. (c) Averaged search paths for both groups (1 d and 30 d probe) and the platform distribution (right). Gray circles indicate specific platform locations.

(d) D_{KL} values 1 d and 30 d after training (t -test 1 d vs. 30 d: 1 d $n = 52$ mice, 30 d $n = 40$ mice, $t_{90} = 2.22$, $P = 0.03$).

(e) Mean platform location crossings 1 d and 30 d after training (Welch's t -test 1 d vs. 30 d: 1 d $n = 52$ mice, 30 d $n = 40$ mice, $t_{89.61} = 3.10$, $P = 0.002$). Data are for individual mice, with black line and shaded region indicating mean \pm s.e.m. (* $P < 0.05$, ** $P < 0.005$).



from a normal distribution in polar coordinates (Supplementary Fig. 1a). Accordingly, the pattern emerged across multiple training days but was not apparent on any single training day. If post-encoding processing facilitates pattern identification, then some delay between the completion of training and testing would allow mice to better match the overall distribution of platforms (perhaps at the expense of spending time at any single day's location; Fig. 1b). To evaluate the impact of a post-encoding delay, we presented mice with a probe test (with the platform removed from the pool) either 1 d or 30 d after the end of training (Supplementary Fig. 2a). To quantify the extent to which search paths on the probe test matched the pattern of platform locations experienced during training, we calculated the Kullback-Leibler divergence³⁰ (D_{KL}) between the search paths and the platform distribution. The D_{KL} is a measure of the difference between two probability distributions; low D_{KL} values indicate high pattern matching (Supplementary Fig. 3).

On any given training day, latency to locate the platform declined across trials and escape latencies generally declined over the course of training (Supplementary Fig. 4). During the probe test, the search patterns were sensitive to the delay between the completion of training and testing (Fig. 1c). Consistent with our prediction, there was a stronger correspondence between search path and overall platform distribution 30 d versus 1 d after the completion of training (indicated by a lower D_{KL} value; Fig. 1d). These results indicate that the passage of time, even in the absence of additional maze experience, allowed mice to better match the overall distribution of platforms.

At the longer delay, mice crossed specific platform locations less frequently (Fig. 1e). This suggests that behavior in the 30-d probe test was driven less by any single day's training experience and instead reflected the cumulative training experience. Improved pattern matching depended on exposure to the pattern and did not emerge spontaneously as a result of either decreased memory accuracy or increased search-path entropy (Supplementary Fig. 5). Similarly, time-dependent shifts in search strategy (e.g., swimming closer to the walls of the pool or reduced searching for more recent platforms) could not account for better pattern matching at the remote delay (Supplementary Figs. 6 and 7a,b).

Our data suggest that the passage of time allowed for improved pattern matching. We next asked whether similar time-dependent improvements would emerge after training on a more complex distribution. In this case, we trained mice on a series of platforms drawn from a weighted bimodal distribution, with uneven peaks centered in the north and east quadrants of the pool. Platforms were twice as likely to occur in the north quadrant (Fig. 2a, and Supplementary Figs. 1b and 2b), and therefore the best match to the pattern would be to spend twice as much time in the north versus east quadrant. At the 1-d delay, mice spent almost equal time in the north and east quadrants. In contrast, at the 30-d delay, mice spent roughly twice as much time in the north versus east quadrant, which indicated that pattern matching improved over time (Fig. 2b). Consistent with this, after a 30-d delay, but not a 1-d delay, the time that mice spent searching in the zone around any given platform

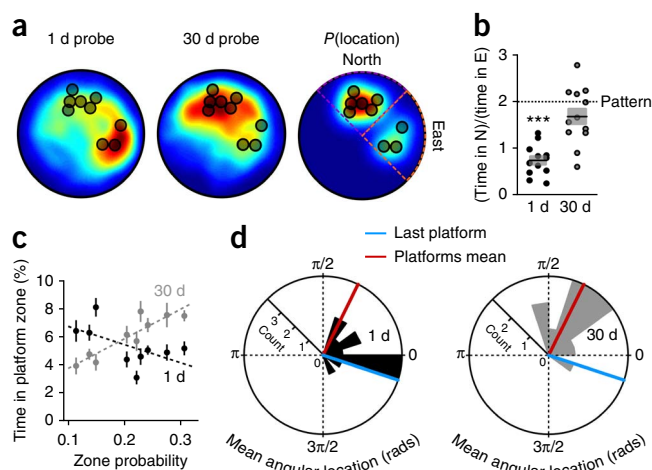


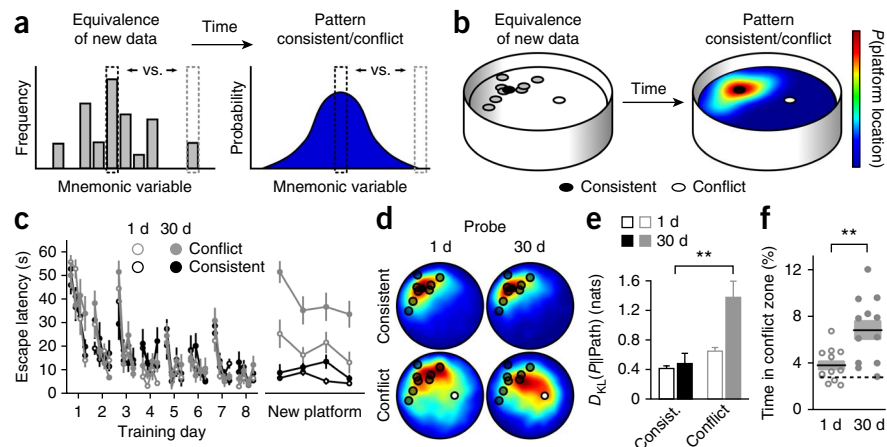
Figure 2 Pattern matching improves over time with alternate distribution and measures. (a) Averaged search paths for both groups (1 d and 30 d probe) and the platform distribution with 2:1 weighting in the north vs. the east quadrant (right). Gray circles indicate specific platform locations.

(b) Time spent searching in the north quadrant (N) vs. the east quadrant (E) 1 d and 30 d after training (t -test vs. mean of 2: 1 d $n = 12$ mice, $t_{11} = -13.05$, $P = 4.91 \times 10^{-8}$; t -test vs. mean of 2: 30 d $n = 12$ mice, $t_{11} = -1.73$, $P = 0.11$). Data are for individual mice, with black line and shaded region indicating mean \pm s.e.m. The dotted line indicates the ratio that corresponds to the pattern of the platform distribution.

(c) Time spent in 10-cm zones centered on each platform vs. platform probability in the distribution 1 d and 30 d after training (Pearson's: $n = 9$ platforms, 1 d $r = -0.58$, $P = 0.1$; 30 d $r = 0.91$, $P = 0.0006$; z -test 1 d vs. 30 d: $z = 2.19$, $P = 0.001$). Data are mean \pm s.e.m. across mice. Dotted lines indicate lines of best fit. (d) Mean angular positions of the mouse search paths 1 d and 30 d after training. Data are histograms, with last platform and mean platform locations highlighted (** $P < 0.0005$).

Figure 3 Sensitivity to difference between pattern consistency versus conflict increases over time. (a) Diagram illustrating hypothetical difference between assimilating new information into memory before and after post-encoding pattern identification has occurred. (b) Schematic of the water-maze task we designed to test for the effects of post-encoding pattern identification on learning new platforms.

(c) Escape latencies (time to find the platform) for all groups across training and on the day of exposure to the new platforms (two-way analysis of variance (ANOVA), 1 d vs. 30 d \times consistent vs. conflict, final trial latency: (1 d, consistent) $n = 24$ mice, (30 d, consistent) $n = 11$ mice, (1 d, conflict) $n = 12$ mice, (30 d, conflict) $n = 12$ mice; 1 d vs. 30 d $F_{1,55} = 11.42$, $P = 0.001$; consistent vs. conflict $F_{1,55} = 29.27$, $P = 1.41 \times 10^{-6}$; interaction $F_{1,55} = 7.54$, $P = 0.008$). (d) Averaged search paths for all groups from a probe test the next day; small circles indicate specific platform locations. (e) D_{KL} values 1 d and 30 d after training for mice exposed to consistent and conflict platforms (two-way ANOVA, 1 d vs. 30 d \times consistent vs. conflict: 1 d vs. 30 d $F_{1,55} = 11.74$, $P = 0.001$; consistent vs. conflict $F_{1,55} = 24.07$, $P = 8.65 \times 10^{-6}$; interaction $F_{1,55} = 8.16$, $P = 0.006$; *post hoc* Tukey's test, (30 d, conflict) vs. (1 d, consistent): $P < 0.001$, (30 d, conflict) vs. (30 d, consistent): $P < 0.001$). (f) Time spent searching in a 10-cm radius zone around the conflict platform 1 d and 30 d after training (*t*-test, (30 d, conflict) vs. (1 d, conflict): $t_{16,11} = -3.29$, $P = 0.005$). The dashed line indicates the chance level of searching in a 10 cm radius circular zone. Data are mean \pm s.e.m., or individual mice, with black line and shaded region indicating mean \pm s.e.m. (** $P < 0.005$).



correlated with the probability assigned to that zone by the distribution (Fig. 2c). The 'recency' of the last platform location (i.e., the shorter length of time that has passed since the last location was encountered relative to the other locations) may represent another factor that influences search paths. To address this, we determined whether order of platform presentation during training influenced search behavior in the probe test. Overall, platform order weakly influenced search behavior (Supplementary Fig. 7c). However, the last platform location did influence search behavior. At the 1-d delay, we found that the focus of the mouse search paths corresponded with the last platform location, but at the 30-d delay, the focus shifted to the mean of the distribution (Fig. 2d). Therefore, probe-test performance is more influenced by the specific memory for the most recent platform 1 d after training but more so by the overall distribution at the 30-d delay.

in pattern matching (Fig. 3a). To evaluate this idea, we modified the training protocol. As previously, we trained mice to find a new platform location on each day, with locations for the first 8 d matching those in the first experiment (Supplementary Fig. 1c). However, when that training was complete, we gave mice a set of test trials where they were presented with a new platform location. We placed this new platform at the mean of the former distribution ('consistent') or in an area of the pool that had a low probability according to the distribution ('conflicting') (Fig. 3b and Supplementary Fig. 1c). To assess the impact of the passage of time, these test trials occurred either 1 d or 30 d after training. We then gave all groups a probe test 1 d later (Supplementary Fig. 2c).

On each training day, latency to locate the platform declined across trials, and escape latencies generally declined over the 8 d of training (Fig. 3c). Although there were no group differences in latency on the final day of training, differences emerged on the test trials. In particular, mice took longer to locate the conflicting (compared to the consistent) platform on the final trial, and this difference was more pronounced 30 d after training (Fig. 3c). This suggests that over time mice became increasingly sensitive to the difference between a new instance that was consistent with the pattern versus an instance that conflicted with the pattern.

Improved pattern matching influences new learning

Only once a pattern has been identified is it possible to determine whether a new event is consistent or in conflict with that pattern. As such, the sensitivity to conflicts may increase as a function of time-dependent improvements

in pattern matching (Fig. 3a). To evaluate this idea, we modified the training protocol. As previously, we trained mice to find a new platform location on each day, with locations for the first 8 d matching those in the first experiment (Supplementary Fig. 1c). However, when that training was complete, we gave mice a set of test trials where they were presented with a new platform location. We placed this new platform at the mean of the former distribution ('consistent') or in an area of the pool that had a low probability according to the distribution ('conflicting') (Fig. 3b and Supplementary Fig. 1c). To assess the impact of the passage of time, these test trials occurred either 1 d or 30 d after training. We then gave all groups a probe test 1 d later (Supplementary Fig. 2c).

Improved pattern matching influences new learning

Only once a pattern has been identified is it possible to determine whether a new event is consistent or in conflict with that pattern. As such, the sensitivity to conflicts may increase as a function of time-dependent improvements

Figure 4 Time-dependent sensitivity to pattern conflict with alternate distribution. (a) Schematic of the experimental design to control for any possible effects of displacement from the last training platform or the pool wall; consistent and conflict platforms were equidistant from the last platform and equidistant from the wall.

(b) Escape latencies for all groups on the day of exposure to the new platforms (two-way ANOVA, 1 d vs. 30 d \times consistent vs. conflict: (1 d, consistent) $n = 12$ mice, (30 d, consistent) $n = 15$ mice, (1 d, conflict) $n = 22$ mice, (30 d, conflict) $n = 23$ mice; 1 d vs. 30 d $F_{1,68} = 0.46$, $P = 0.5$, consistent vs. conflict: $F_{1,68} = 0.77$, $P = 0.38$; interaction $F_{1,68} = 4.07$, $P = 0.047$). (c) Averaged search-path distributions from the probe test; small circles indicate specific platform locations. (d) D_{KL} values 1 d and 30 d after training for mice exposed to consistent and conflict platforms (two-way ANOVA 1 d vs. 30 d \times consistent vs. conflict: 1 d vs. 30 d $F_{1,68} = 3.96$, $P = 0.05$, consistent vs. conflict: $F_{1,68} = 16.01$, $P = 0.001$; interaction $F_{1,68} = 1.34$, $P = 0.25$; *post hoc* Tukey's test, high P , 1 d vs. low P , 30 d: $P < 0.05$; *post hoc* Tukey's test, (30 d, conflict) vs. (1 d, consistent): $P < 0.05$, (30 d, conflict) vs. (30 d, consistent): $P < 0.05$). (e) Time spent searching in a 10-cm radius zone around the conflict platform 1 d and 30 d after training (*t*-test (30 d, conflict) vs. (1 d, conflict): $t_{43} = -2.41$, $P = 0.02$). The dashed line indicates the chance level of searching in a 10 cm radius circular zone. Data are mean \pm s.e.m., or individual mice, with black line and shaded region indicating mean \pm s.e.m. (* $P < 0.05$).


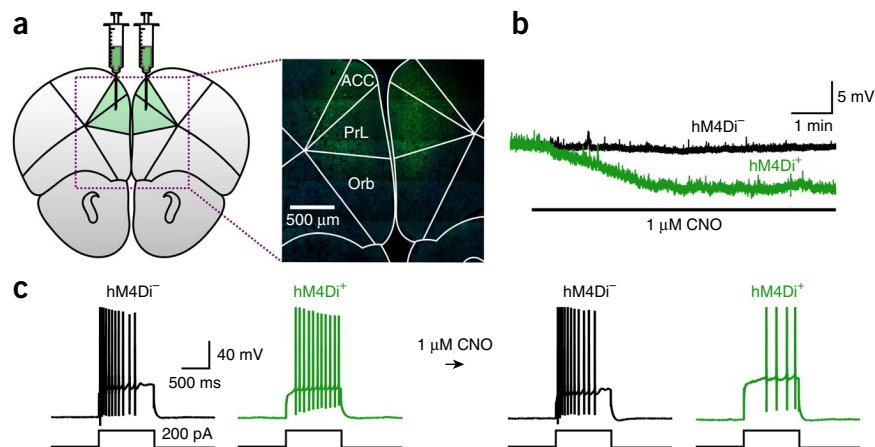


Figure 4: Time-dependent sensitivity to pattern conflict with alternate distribution. (a) Schematic of the experimental design to control for any possible effects of displacement from the last training platform or the pool wall; consistent and conflict platforms were equidistant from the last platform and equidistant from the wall. (b) Escape latencies for all groups on the day of exposure to the new platforms (two-way ANOVA, 1 d vs. 30 d \times consistent vs. conflict: (1 d, consistent) $n = 12$ mice, (30 d, consistent) $n = 15$ mice, (1 d, conflict) $n = 22$ mice, (30 d, conflict) $n = 23$ mice; 1 d vs. 30 d $F_{1,68} = 0.46$, $P = 0.5$, consistent vs. conflict: $F_{1,68} = 0.77$, $P = 0.38$; interaction $F_{1,68} = 4.07$, $P = 0.047$). (c) Averaged search-path distributions from the probe test; small circles indicate specific platform locations. (d) D_{KL} values 1 d and 30 d after training for mice exposed to consistent and conflict platforms (two-way ANOVA 1 d vs. 30 d \times consistent vs. conflict: 1 d vs. 30 d $F_{1,68} = 3.96$, $P = 0.05$, consistent vs. conflict: $F_{1,68} = 16.01$, $P = 0.001$; interaction $F_{1,68} = 1.34$, $P = 0.25$; *post hoc* Tukey's test, high P , 1 d vs. low P , 30 d: $P < 0.05$; *post hoc* Tukey's test, (30 d, conflict) vs. (1 d, consistent): $P < 0.05$, (30 d, conflict) vs. (30 d, consistent): $P < 0.05$). (e) Time spent searching in a 10-cm radius zone around the conflict platform 1 d and 30 d after training (*t*-test (30 d, conflict) vs. (1 d, conflict): $t_{43} = -2.41$, $P = 0.02$). The dashed line indicates the chance level of searching in a 10 cm radius circular zone. Data are mean \pm s.e.m., or individual mice, with black line and shaded region indicating mean \pm s.e.m. (* $P < 0.05$).

Figure 5 CNO inhibits spiking activity in hM4Di infected neurons in the mPFC. (a) Illustration of the regions targeted for AAV8-*Camk2a*-hM4Di-mCitrine microinfusion (left). Image of a representative infection in the anterior cingulate cortex (ACC), prelimbic cortex (PrL) and orbital cortex (Orb) (right). (b) Example traces showing responses to bath perfusion of 1 μ M CNO recorded via whole-cell current clamp in infected (hM4Di⁺) and uninfected (hM4Di⁻) neurons of *ex vivo* mPFC slices. (c) Example traces showing responses to injection of square pulses of 200 pA current before and after CNO bath application.



Learning is driven by differences between expected and observed outcome (i.e., prediction error³¹). Therefore, if the passage of time allows for refinement of an identified pattern, then presentation of the conflicting platform at the 30-d delay should produce a larger prediction error and lead to greater learning. The probe test provided an opportunity to evaluate learning after the presentation of the new platforms. When we presented mice with the conflicting platform after a 1-d delay, they reverted to searching the former distribution in the subsequent day's probe test (Fig. 3d). In contrast, when we presented mice with the conflicting platform after a 30-d delay, the mice altered their search strategy, as reflected by a significant increase in the D_{KL} value (Fig. 3e). These data suggest that after a 30-d delay, mice were more sensitive to instances that violated the pattern that they had experienced during training, and this produced a high prediction error and consequently supported new learning. In line with this interpretation, in the probe test mice searched for the conflicting platform more if it was presented 30 d after training (Fig. 3f). Presentation of the consistent platform did not alter search paths on the probe test, regardless of the post-training delay (Fig. 3d,e). We replicated these findings using a different distribution as well as conflicting and consistent platforms that were equidistant from the edge of the pool (Fig. 4a–e and Supplementary Fig. 1d). Therefore, the tendency for a conflicting platform to induce longer escape latencies, as well as stronger learning, does not depend on specific distributions or particular platform locations but only on the passage of time.

The mPFC mediates learning of conflict with prior patterns

The mPFC may be important for resolving conflicts between new information and background knowledge¹¹. Our data indicate that sensitivity to conflicts increases as a function of time and therefore predict a more pronounced role for the mPFC in conflict resolution at remote time points after training. To test this, we microinfused an adeno-associated virus (AAV) carrying a *Camk2a*-driven modified sequence encoding the human muscarinic M4 receptor, known as hM4Di, tagged with the fluorescent reporter mCitrine, into the mPFC and then trained mice in our water-maze paradigm. hM4Di is a $G_{i/o}$ -coupled designer receptor exclusively activated by designer drugs (DREADDs) that is insensitive to endogenous ligands but activated by a synthetic ligand (clozapine-N-oxide (CNO)³²). When bound to CNO, this $G_{i/o}$ -coupled DREADD induces membrane hyperpolarization and inhibition of spiking activity³³.

Histological examination revealed robust, bilateral expression of hM4Di in neurons in the mPFC, including the anterior cingulate cortex and prelimbic cortex regions (with some limited infection in the orbital cortex) (Fig. 5a). In a subset of mice, we performed *ex vivo*

whole-cell patch clamp recordings to characterize the effects of CNO in AAV-infected neurons. Application of 1 μ M CNO hyperpolarized infected neurons (Fig. 5b) by -7.07 ± 2.09 mV (mean \pm s.e.m) but not uninfected neurons, which showed a change of -0.12 ± 0.81 mV after treatment with CNO (*t*-test for hM4Di⁻ (no hM4Di expression) versus hM4Di⁺ (hM4Di expression): hM4Di⁻ $n = 3$ cells, hM4Di⁺ $n = 4$, $t_4 = 3.09$, $P = 0.04$). This was associated with inhibition of spiking responses to 200 pA square current pulses (Fig. 5c), with a $38.59 \pm 9.32\%$ reduction in the spike rate of infected neurons after application of CNO, compared to a $3.03 \pm 8.02\%$ increase in uninfected neurons (*t*-test versus mean of 1: hM4Di⁻ $n = 3$, $t_2 = 0.38$, $P = 0.74$; *t*-test versus mean of 1: hM4Di⁺ $n = 4$, $t_3 = -4.14$, $P = 0.03$).

After mice recovered from surgery, we trained them in the same modified water-maze protocol as in the previous experiment (Fig. 4 and Supplementary Fig. 1d). On any given training day, latency to locate the platform declined across trials, similar to previous experiments (Supplementary Fig. 8). Either 1 d or 30 d after training, we presented mice with a platform that was either consistent or in conflict with the pattern. To pharmacogenetically inhibit mPFC neurons, we injected the mice with CNO or vehicle 30 min before the mice learned the new platform (Supplementary Fig. 9). We then probed the mice the next day in a drug-free state in order to evaluate how mPFC inhibition had altered learning. After a 1-d post-training delay, CNO had no effect on learning either the consistent or conflicting platforms. During the subsequent probe test, search paths in CNO and vehicle groups did not differ in D_{KL} or the time spent in the new platform zone (Fig. 6a–c).

However, mPFC inhibition significantly impacted learning of a conflicting platform when there was a 30-d delay after training. As expected, vehicle-treated mice presented with the conflicting platform after 30 d altered their search behavior, shifting their focus to the new platform in the subsequent probe test (Fig. 6d). In contrast, CNO-treated mice presented with the conflicting platform divided their search between the training distribution and the new platform in the subsequent probe test (Fig. 6d). This division between training distribution and the new platform was indicated by a significant reduction in the D_{KL} value and time spent searching for the new platform in the CNO injection group compared to the vehicle group (Fig. 6e,f). Injection of CNO during presentation of the consistent platform did not alter search paths on the subsequent probe test (Fig. 6d–f). Altogether, these results indicate that the mPFC has an important role in updating search behavior in response to new platforms that conflict with patterns encountered during training, but only after a 30-d delay.

Figure 6 mPFC inhibition impairs learning of conflicting platform after 30-d delay.

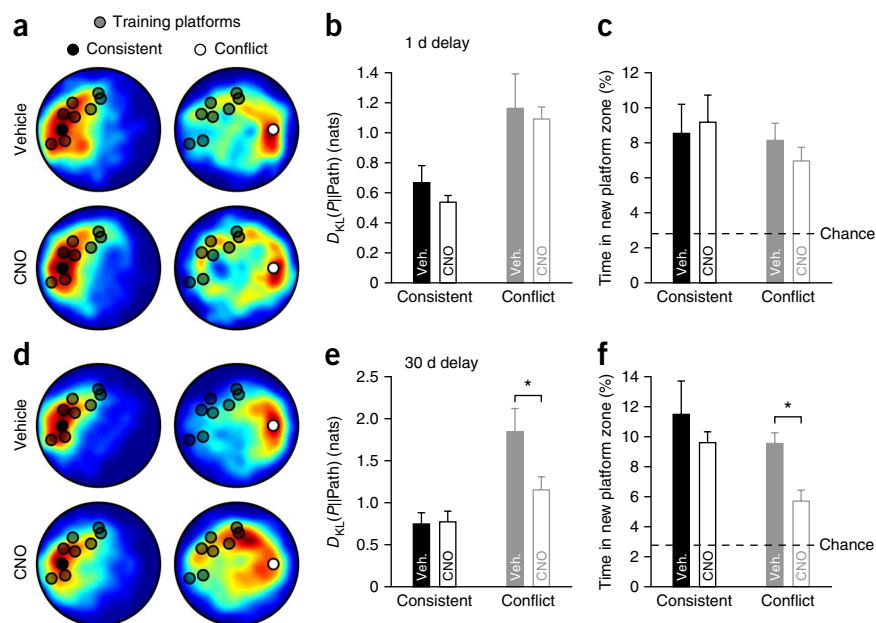
(a) Averaged search-path distributions from the probe tests after a 1-d delay; small circles indicate specific platform locations.

(b) D_{KL} values from probe tests for mice exposed to new platform (consistent or conflict) and given either vehicle or CNO injections 1 d after training (two-way ANOVA, consistent vs. conflict \times CNO vs. vehicle: (consistent, vehicle) $n = 7$ mice, (consistent, CNO) $n = 8$, (conflict, vehicle) $n = 5$, (conflict, CNO) $n = 4$; consistent vs. conflict $F_{1,20} = 17.02$, $P = 0.0005$; vehicle vs. CNO $F_{1,20} = 0.6$, $P = 0.45$; interaction $F_{1,20} = 0.05$, $P = 0.82$).

(c) Time spent searching in a 10-cm radius zone around the new platform for mice exposed to consistent and conflict platforms and given either vehicle or CNO injections 1 d after training (two-way ANOVA, consistent vs. conflict \times CNO vs. vehicle: consistent vs. conflict $F_{1,20} = 0.67$, $P = 0.42$; vehicle vs. CNO $F_{1,20} = 0.03$, $P = 0.87$; interaction $F_{1,20} = 0.33$, $P = 0.57$).

(d) Averaged search-path distributions from the probe tests after a 30-d delay.

(e) D_{KL} values as in **b**, 30 d after training (two-way ANOVA, consistent vs. conflict \times CNO vs. vehicle: (consistent, vehicle) $n = 6$ mice, (consistent, CNO) $n = 8$, (conflict, vehicle) $n = 8$, (conflict, CNO) $n = 13$; consistent vs. conflict $F_{1,31} = 14.54$, $P = 0.0006$; vehicle vs. CNO $F_{1,31} = 2.93$, $P = 0.1$; interaction $F_{1,31} = 3.43$, $P = 0.07$; *post hoc* Tukey's test, (conflict, vehicle) vs. (conflict, CNO): $P < 0.05$). (f) Analysis as in **c**, 30 d after training (two-way ANOVA, consistent vs. conflict \times CNO vs. vehicle: consistent vs. conflict $F_{1,31} = 7.36$, $P = 0.01$; vehicle vs. CNO $F_{1,31} = 7.04$, $P = 0.01$; interaction $F_{1,31} = 0.83$, $P = 0.37$; *post hoc* Tukey's test, (conflict, vehicle) vs. (conflict, CNO): $P < 0.05$). Data are mean \pm s.e.m. (* $P < 0.05$).



DISCUSSION

The hypothesis that there is a persistent process of pattern identification after encoding to combine and schematize memories is central to contemporary theories of systems consolidation^{21–25}. According to these theories, statistically common elements across multiple memories constitute patterns that can be used during consolidation to construct more general knowledge about the world^{21–25}. Here we provided evidence supporting this hypothesis. Our data indicate that mice trained in a water-maze task to find platforms drawn from a statistical distribution exhibited better pattern matching after a 30-d delay compared to a 1-d delay. We observed this improvement in pattern matching for two distinct distributions, using different measures of pattern matching. This time-dependent improvement in pattern matching was associated with increased sensitivity to conflicting platforms. Given that conflicts between new information and background knowledge engage the mPFC, this suggested that the mPFC may have a more pronounced role in learning about conflicting platforms at remote time points. In agreement with this, pharmacogenetic inhibition of the mPFC only affected the learning of conflicting platforms after a 30-d delay. Altogether, our results support the idea that there is an extended process of pattern identification after encoding and that identified patterns can influence new learning as a result of computations performed in the mPFC.

An important question that our studies do not directly address is: what is the nature of the process that continues after encoding, i.e., what leads to better pattern matching and greater sensitivity to pattern conflicts at remote time points? One component of the transformation may be degradation of the original memories. However, the time-dependent improvement in pattern matching to the weighted bimodal distribution in **Figure 2** is inconsistent with a memory-degradation process (reduced memory precision would not alter the ratio of time spent in north versus east quadrants). Another factor that is very likely to be involved is recency³⁴, where memory strength is related

to how recently an event has been experienced. We found that after a 1-d delay search-path behavior was often centered at the position of the last platform encountered, whereas after a 30-d delay search paths were centered at the mean of the platform distribution (**Fig. 2d**). This suggests that one contribution to better pattern matching after a 30-d delay was reduced dominance of the memory for the last platform. However, recency effects cannot explain all of our results. The increased sensitivity of mice to the difference between a consistent and conflicting platform after a 30-d delay, measured by their escape latencies (**Fig. 3c**), is the opposite of what a recency account would predict: a reduction in the strength of the memories for the most recent platforms should have made it easier to learn a new platform in a new location of the pool, not harder. Therefore, although a reduction in the focus on recent platforms is likely to contribute to improved pattern matching over time, it is unlikely to be the sole factor driving the behavior we observed.

What other processes might be involved in time-dependent pattern matching and sensitivity to pattern conflicts? It has been proposed that statistical regularities across memories are extracted via hippocampal replay, which leads to the categorization of experiences in the neocortex^{19,20} (an idea that subsequently has been elaborated in computational models of systems consolidation²⁴). With this in mind, it is plausible that time-dependent improvements in pattern matching could reflect an active process of statistical learning. Given that the patterns in platform locations that our mice were exposed to were explicitly statistical in nature, our data are consistent with this interpretation. In line with this, there is evidence that the brain can learn statistical models and estimate the probability of different values or events based on experience. Early psychological studies indicated that people can rapidly identify statistical patterns in visual stimuli³, and computational models of sensory learning in the neocortex have long proposed that it is statistical in nature^{35–37}. These ideas have been supported more recently by neuroimaging work in humans^{14,38} and

in vivo recordings in primary sensory cortices of ferrets³⁹. Likewise, learning in sensorimotor integration tasks is predicted by Bayesian statistical models^{40,41}. Studies in rodents using multiple-choice behavioral tasks have also demonstrated that the neocortex can compute estimates of expectation and confidence^{42,43}. Finally, studies in humans have shown that sleep facilitates statistical learning, possibly via memory replay^{15–18}. Therefore, the types of learning mechanisms known to operate in the brain are well suited to identify statistical patterns across multiple memories during consolidation.

Our finding that inhibition of the mPFC affects learning of conflicting platforms after a 30-d delay has interesting implications for rapid, schema-based consolidation^{6,7,10}. First, these data suggest that the role of the mPFC in rapid consolidation may be limited to the learning of new information that breaks established patterns¹¹. It is possible that consolidation of information that closely matches (or only slightly deviates from) previous experiences could be achieved without the functions performed by the mPFC. Second, our findings leave open the possibility that rapid, schema-based consolidation would not be possible in tasks where initial learning is too rapid to permit time-dependent pattern identification. Recent memories may not provide sufficient scaffolding for quickly encoding information outside of the hippocampus. The framework provided by the task we used here provides an approach for investigating these issues in the future.

METHODS

Methods and any associated references are available in the [online version of the paper](#).

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

ACKNOWLEDGMENTS

This work was supported by grants from the Canadian Institutes of Health Research to P.W.F. (MOP-77561), S.A.J. (MOP-74650) and M.A.W. (MOP-123466). B.A.R. received financial support from a Banting Postdoctoral Fellowship via the Natural Sciences and Engineering Research Council of Canada. B.A.R., F.X. and J.H. received financial support from the Canadian Institutes of Health Research Sleep and Biological Rhythms Training Program. A.S. received financial support from the Hospital for Sick Children. We thank G. Hinton and M. Moscovitch for comments on earlier drafts of this manuscript.

AUTHOR CONTRIBUTIONS

B.A.R., S.A.J. and P.W.F. conceived and designed the experiments. B.A.R., F.X., A.S. and J.H. performed the behavioral experiments. B.A.R. and M.A.W. performed the patch clamp experiments. B.A.R. conducted the analyses. B.A.R., S.A.J. and P.W.F. wrote the paper.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

Reprints and permissions information is available online at <http://www.nature.com/reprints/index.html>.

- Bartlett, F.C. *Remembering: A Study in Experimental and Social Psychology* (Cambridge University Press, 1932).
- Piaget, J. & Cook, M. *The Origins of Intelligence in Children* (International Universities Press, 1952).
- Posner, M.I. & Keele, S.W. On the genesis of abstract ideas. *J. Exp. Psychol.* **77**, 353–363 (1968).
- Knowlton, B. & Squire, L. The learning of categories: parallel brain systems for item memory and category knowledge. *Science* **262**, 1747–1749 (1993).
- Squire, L.R. & Knowlton, B.J. Learning about categories in the absence of memory. *Proc. Natl. Acad. Sci. USA* **92**, 12470–12474 (1995).
- Tse, D. *et al.* Schema-dependent gene activation and memory encoding in neocortex. *Science* **333**, 891–895 (2011).
- Wang, S.-H., Tse, D. & Morris, R.G.M. Anterior cingulate cortex in schema assimilation and expression. *Learn. Mem.* **19**, 315–318 (2012).
- Winocur, G., Moscovitch, M., Fogel, S., Rosenbaum, R.S. & Sekeres, M. Preserved spatial memory after hippocampal lesions: effects of extensive experience in a complex environment. *Nat. Neurosci.* **8**, 273–275 (2005).

- Van Kesteren, M.T.R., Fernández, G., Norris, D.G. & Hermans, E.J. Persistent schema-dependent hippocampal-neocortical connectivity during memory encoding and postencoding rest in humans. *Proc. Natl. Acad. Sci. USA* **107**, 7550–7555 (2010).
- Tse, D. *et al.* Schemas and memory consolidation. *Science* **316**, 76–82 (2007).
- Preston, A.R. & Eichenbaum, H. Interplay of hippocampus and prefrontal cortex in memory. *Curr. Biol.* **23**, R764–R773 (2013).
- Rumelhart, D.E., Smolensky, P., McClelland, J.L. & Hinton, G.E. Schemata and sequential thought processes in PDP models. *Parallel Distrib. Process. Explor. Microstruct. Cogn.* **2**, 7–57 (1986).
- Smolensky, P. Information processing in dynamical systems: foundations of harmony theory. *Parallel Distrib. Process. Explor. Microstruct. Cogn.* **1**, 194–281 (1986).
- Turk-Browne, N.B., Scholl, B.J., Chun, M.M. & Johnson, M.K. Neural evidence of statistical learning: efficient detection of visual regularities without awareness. *J. Cogn. Neurosci.* **21**, 1934–1945 (2009).
- Durrant, S.J., Taylor, C., Cairney, S. & Lewis, P.A. Sleep-dependent consolidation of statistical learning. *Neuropsychologia* **49**, 1322–1331 (2011).
- Lewis, P.A. & Durrant, S.J. Overlapping memory replay during sleep builds cognitive schemata. *Trends Cogn. Sci.* **15**, 343–351 (2011).
- Durrant, S.J., Cairney, S.A. & Lewis, P.A. Overnight consolidation aids the transfer of statistical knowledge from the medial temporal lobe to the striatum. *Cereb. Cortex* **23**, 2467–2478 (2013).
- Djonlagic, I. *et al.* Sleep enhances category learning. *Learn. Mem.* **16**, 751–755 (2009).
- Marr, D. A theory for cerebral neocortex. *Proc. R. Soc. Lond. B Biol. Sci.* **176**, 161–234 (1970).
- Marr, D. Simple memory: a theory for archicortex. *Phil. Trans. R. Soc. Lond. B* **262**, 23–81 (1971).
- Nadel, L. & Moscovitch, M. Memory consolidation, retrograde amnesia and the hippocampal complex. *Curr. Opin. Neurobiol.* **7**, 217–227 (1997).
- Moscovitch, M., Nadel, L., Winocur, G., Gilboa, A. & Rosenbaum, R.S. The cognitive neuroscience of remote episodic, semantic and spatial memory. *Curr. Opin. Neurobiol.* **16**, 179–190 (2006).
- Winocur, G., Moscovitch, M. & Bontempi, B. Memory formation and long-term retention in humans and animals: Convergence towards a transformation account of hippocampal-neocortical interactions. *Neuropsychologia* **48**, 2339–2356 (2010).
- McClelland, J.L., McNaughton, B.L. & O'Reilly, R.C. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* **102**, 419–457 (1995).
- McClelland, J.L. Incorporating rapid neocortical learning of new schema-consistent information into complementary learning systems theory. *J. Exp. Psychol. Gen.* **142**, 1190–1210 (2013).
- Winocur, G., Moscovitch, M. & Sekeres, M. Memory consolidation or transformation: context manipulation and hippocampal representations of memory. *Nat. Neurosci.* **10**, 555–557 (2007).
- Wiltgen, B.J. & Silva, A.J. Memory for context becomes less specific with time. *Learn. Mem.* **14**, 313–317 (2007).
- Rosenbaum, R.S. *et al.* Remote spatial memory in an amnesic person with extensive bilateral hippocampal lesions. *Nat. Neurosci.* **3**, 1044–1048 (2000).
- Steele, R.J. & Morris, R.G.M. Delay-dependent impairment of a matching-to-place task with chronic and intrahippocampal infusion of the NMDA-antagonist D-AP5. *Hippocampus* **9**, 118–136 (1999).
- Kullback, S. & Leibler, R.A. On information and sufficiency. *Ann. Math. Stat.* **22**, 79–86 (1951).
- Schultz, W., Dayan, P. & Montague, P.R. A neural substrate of prediction and reward. *Science* **275**, 1593–1599 (1997).
- Farrell, M.S. & Roth, B.L. Pharmacosynthetics: reimagining the pharmacogenetic approach. *Brain Res.* **1511**, 6–20 (2013).
- Ferguson, S.M. *et al.* Transient neuronal inhibition reveals opposing roles of indirect and direct pathways in sensitization. *Nat. Neurosci.* **14**, 22–24 (2011).
- Baddeley, A. & Hitch, G. The recency effect: Implicit learning with explicit retrieval? *Mem. Cognit.* **21**, 146–155 (1993).
- Hinton, G.E. & Sejnowski, T.J. Learning and relearning in Boltzmann machines. *Parallel Distrib. Process. Explor. Microstruct. Cogn.* **1**, 282–317 (1986).
- Hinton, G.E., Dayan, P., Frey, B.J. & Neal, R.M. The 'wake-sleep' algorithm for unsupervised neural networks. *Science* **268**, 1158–1161 (1995).
- Dayan, P. & Hinton, G.E. Varieties of Helmholtz machine. *Neural Netw.* **9**, 1385–1403 (1996).
- Turk-Browne, N.B., Scholl, B.J., Johnson, M.K. & Chun, M.M. Implicit perceptual anticipation triggered by statistical learning. *J. Neurosci.* **30**, 11177–11187 (2010).
- Berkes, P., Orbán, G., Lengyel, M. & Fiser, J. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* **331**, 83–87 (2011).
- Makin, J.G., Fellows, M.R. & Sabes, P.N. Learning multisensory integration and coordinate transformation via density estimation. *PLoS Comput. Biol.* **9**, e1003035 (2013).
- Kording, K.P. & Wolpert, D.M. Bayesian integration in sensorimotor learning. *Nature* **427**, 244–247 (2004).
- Jaramillo, S. & Zador, A.M. The auditory cortex mediates the perceptual effects of acoustic temporal expectation. *Nat. Neurosci.* **14**, 246–251 (2011).
- Kepecs, A., Uchida, N., Zariwala, H.A. & Mainen, Z.F. Neural correlates, computation and behavioural impact of decision confidence. *Nature* **455**, 227–231 (2008).

ONLINE METHODS

Animals. All experiments conformed to the guidelines set forth by the Canadian Council for Animal Care and the Animal Care Committee at the Hospital for Sick Children. 8–9-week-old male, wild-type mice were used in all experiments. These mice were the F₁ generation derived from crossing C57BL/6n and 129SvEv mice (breeding mice were sourced directly from Taconic). Mice were weaned at 21 d, and then housed 4–5 mice per cage. The housing room was maintained on a 12 h–12 h light–dark cycle, with lights on during the day. All experiments were performed in the morning between 9 a.m. and 12 p.m. One week before the commencement of experiments mice were briefly anesthetized using isoflurane, and their ears were punched with individual markings to differentiate cage mates. Cages of mice were assigned to different experimental groups randomly. The experimenter was not blinded regarding group assignments because this would have required using different experimenters for training and probe tests, which could introduce unwanted behavioral effects. In the 4 d before training, mice were handled by the experimenter, two at a time, for 3 min each day. No mice were excluded from any analyses in purely behavioral experiments (Figs. 1–4), with the exception of one mouse that escaped the pool during the second day of training and did not search for the platform afterward. Mice were excluded from the pharmacogenetic experiments (Figs. 5 and 6) based on post-experiment histology: only mice for whom viral microinfusion resulted in robust bilateral expression of hM4Di-mCitrine in the appropriate location of the mPFC were included (see below for more details on the exclusion criteria).

Behavioral apparatus. Experiments were performed in a water-maze pool, which was 120 cm in diameter and 50 cm in depth. The pool was surrounded by white curtains with large red and black shapes to provide visual cues for navigation. Ambient light in the room was kept at a constant, moderately low level across all days of the experiments.

Each day, before mice were brought into the room, the pool was filled to 40 cm with water and the temperature of the water was adjusted to 22 °C. The water was rendered opaque using nontoxic, white children's paint. Video tracking of the mice was conducted using an Actimetrics WaterMaze system (Actimetrics) at a rate of 8 frames/s. The tracking camera was mounted above the center of the pool and fixed to the ceiling. A custom-built platform with a double-jointed rotating arm was used to ensure that the platform position could be adjusted to a specific position within the pool. The platform was circular (10 cm diameter). The platform positions were preselected using Matlab 7.12 (MathWorks) and entered into the WaterMaze software. The video tracking device was then used to confirm that the actual platform was set to the correct position on each day. The platform was submerged ~5–8 mm below the water level.

Behavioral training and testing. The training procedure was based on the delayed matching-to-place task in ref. 29. As in that task, the platform location changed on each day of training, and all mice were trained for four trials on this new location. However, in this task the locations were selected from a statistical distribution in polar coordinates. Specifically, we selected the distance of a given platform from the center of the pool, d , and the angle of the arc connecting the platform and the eastern cardinal direction in the pool, θ , using a variety of distributions (Supplementary Fig. 1). In the first set of groups (Fig. 1) the platforms for the first 8 d of training were generated by selecting d from a normal distribution with a mean of 35 cm and an s.d. of 10 cm, and selecting θ from a Von Mises distribution (a circular normal distribution) with a mean of 0.75π radians and a κ parameter of 5. The 9th platform for these groups was presented at the mean of the selection distribution (Supplementary Fig. 1a). In the weighted bimodal distribution groups (Fig. 2), d was again selected from a normal distribution with a mean of 35 cm and a s.d. of 10 cm, but now, θ was selected by first randomly selecting one of two Von Mises distributions, one with a mean of 0.5π radians, the other with a mean of 0 radians, and then selecting from the chosen distribution. The selection of the two distributions was weighted 2:1 toward the first (0.5π radians mean), and both used a κ parameter of 8 (Supplementary Fig. 1b). For the first set of groups presented with new platforms (Fig. 3), training was conducted with the same eight platforms used in the first experiments, but during the test trials either the same platform at the mean of the distribution was presented ('consistent') or a different platform was presented that was low probability in the distribution ('conflict') (Supplementary Fig. 1c). For the second set of groups presented with new consistent or conflict platforms (Fig. 4),

a new set of seven platforms was selected from the distribution used for the first experiments. The 8th platform was then placed along the northern axis of the pool, 35 cm from the center. The consistent and conflict platforms were then placed 35 cm from the center along either the western or eastern axis, respectively, in order to preserve the displacement between the final day of training and the two different new platforms (Supplementary Fig. 1d). In the mPFC inhibition experiments (Fig. 6), the same platforms used in the second set of new platform learning experiments were used. In the single platform groups (Supplementary Fig. 5), the same platform at the mean of the first distribution was presented each day. The platforms used and their specific order over the days of training for all groups are shown in Supplementary Figures 1 and 2.

For each training trial, mice were taken from their cage and carried to the pool area on the hand of the experimenter. Mice were then lowered into the pool by their tail at one of the 4 cardinal locations (north, south, east and west) facing the wall of the pool. The start locations were selected randomly and were identical for all mice in any given experiment (Supplementary Fig. 2). The experimenter started the trial manually once he or she left the water-maze area. Training trials ended automatically when tracking software determined that the mice had stayed on the platform for 1 s. Mice were then left on the platform for 15 s, at which point they were retrieved by the experimenter. If a mouse jumped off the platform before the end of 15 s, it was directed back to the platform by the experimenter. Trials lasted for a maximum of 60 s. On the first day, mice often did not get on the platform and were guided there by hand of the experimenter at the end of the trial. On the first trial of each day, when mice sometimes did not find the new platform in the allotted 60 s, the experimenter would place his or her hand on the platform until the mouse came to the platform and climbed onto it. At the start of each trial a short 320 Hz tone played, at the end of each trial a short 440 Hz tone played and during the intertrial interval a 220 Hz tone repeatedly played once every 5 s. The tones marked the trials for the experimenter, but they were unlikely to have affected behavioral results as they were below the hearing range of mice⁴⁴. During training, when mice were not in the water maze they were kept in their cages in an area behind the curtains.

On the new platform trials, when the consistent or conflict platforms were introduced, the procedure followed the same protocol as on the training days. On probe tests, the same procedure as on the training days was followed except that the platform was absent from the pool. The probe trials lasted for 60 s. In experiments 1, 2 and 3 (Figs. 1–3), mice underwent two probe trials and the search paths from both trials were used for data analysis (see below). At the end of the first probe trial, mice were removed from the pool by the experimenter at the 60 s mark at whichever location they happened to be at in the pool. In experiments 4 and 5 (Figs. 4 and 6), mice underwent a single probe trial. A subset of data (24 mice) from the first groups given a 1-d delay (Fig. 1) was reused for the first new platform analyses (Fig. 3) as these training procedures were identical and these mice were trained by the same experimenter (Supplementary Fig. 2). Further, because the experimental procedures were identical in experiments 4 and 5, a subset of the mice that were trained by the same experimenter included in the analysis in Figure 4 (15 of 72 mice) were from the vehicle condition in Figure 6. The specific training protocols and testing time lines for each experimental group are displayed in Supplementary Figure 2.

In the pharmacogenetic experiments (Fig. 6, and Supplementary Figs. 8b,c and 9), mice received intraperitoneal injections of either clozapine-N-oxide (CNO) or a vehicle solution 30 min before the first trial of the new platform learning. Mice in the CNO group received an injection of 5 mg/kg CNO, which was first dissolved in 20 μ l DMSO, then mixed into 380 μ l 1 \times PBS. Mice in the vehicle group received an injection of 20 μ l DMSO mixed with 380 μ l 1 \times PBS. All drugs and ingredients for these experiments were obtained from Sigma.

Viral microinfusions. 45 d before training with a new platform (Supplementary Fig. 9), mice in the pharmacogenetic experiments underwent stereotaxic surgery for viral microinfusion of AAV into the mPFC. AAV8-*Camk2a-hM4Di-mCitrine* was purchased from the University of North Carolina Vector Core. Mice were pretreated with atropine sulfate (0.1 mg/kg, intraperitoneal), anesthetized with chloral hydrate (400 mg/kg, intraperitoneal) and placed in a stereotaxic frame. Skin on top of the head was retracted, and holes were drilled in the skull bilaterally above the mPFC (anteroposterior = +1.9, mediolateral = \pm 0.3, ventral = –1.8 mm from bregma) according to ref. 45. AAV (2.0 μ l/side) was microinjected through glass micropipettes connected via polyethylene tubing to a microsyringe

(Hamilton) at a rate of 0.1 $\mu\text{l}/\text{min}$. Micropipettes were left in place an additional 5 min after microinfusion to ensure diffusion of vector. Micropipettes were slowly retracted, the incision closed, and mice treated with analgesic (ketoprofen, 5 mg/kg, subcutaneous) and 1 ml saline.

After the pharmacogenetic experiments mice were transcardially perfused with 1 \times PBS followed by 4% paraformaldehyde. Brains were fixed overnight (4 $^{\circ}\text{C}$) and transferred to a 30% sucrose solution. Coronal brain slices (30 μm) across the entire extent of the mPFC were collected using a cryostat (Leica CM1850). Every second section was mounted on a gel-coated glass slide and covered with a coverslip with Vectashield fluorescence mounting medium (Vector Laboratories). Mosaic images of the mPFC in the slices were then obtained using a confocal laser scanning microscope (LSM 710; Zeiss).

Images of AAV infection were scored via two judges blinded to the groups. Each mouse received an infection quality score from each judge ranging from 0 to 3, where 0 indicated no infection or severe damage to the tissue, 1 indicated infection, but only unilateral, weak or in the wrong location, 2 indicated robust bilateral infection centered on the prelimbic and anterior cingulate cortices, and 3 indicated ideal infection. Only data for animals whose score averaged across both judges was 2 or greater were included in subsequent analyses. As a result, of the 112 animals included in the original experiments, only 59 were included in the data analyses shown in **Figure 6**.

Electrophysiology. A subset of mice (4 mice) microinfused with AAV were used to test the effects of CNO on hM4Di-infected neurons. 4 weeks after microinfusion of AAV into the mPFC, these mice were perfused with cold modified artificial cerebrospinal fluid (mACSF) containing the following (in mM): 180 sucrose, 25 sodium bicarbonate, 25 glucose, 2.5 KCl, 1.25 sodium phosphate, 2 MgCl_2 , 1 CaCl_2 , 0.4 sodium ascorbate and 3 sodium pyruvate, saturated with 95% $\text{O}_2/5\% \text{CO}_2$ and with pH and osmolarity adjusted to 7.4 and 290–295 mOsm, respectively. Brains were quickly removed and placed for 30 s in a chilled mACSF slurry. The posterior portions of the brain were removed, and the posterior edge of the brain was glued to a slicing stage with an agarose block placed behind it. The stage was then placed in a slicing chamber filled with mACSF chilled to a slurry and continuously oxygenated with 95% $\text{O}_2/5\% \text{CO}_2$. Brain slices (350 μm) were prepared on a VT1000S vibratome (Leica) and slices recovered for 1 h at room temperature in a continuously oxygenated mixture of 25 ml of mACSF mixed with 25 ml ACSF (see recipe below).

During recording, slices were placed in a recording chamber perfused with continuously oxygenated ACSF containing the following (in mM): 125 NaCl, 25 sodium bicarbonate, 25 glucose, 2.5 KCl, 1.25 sodium phosphate, 1 MgCl_2 , 2 CaCl_2 , 0.4 sodium ascorbate and 3 sodium pyruvate, saturated with 95% $\text{O}_2/5\% \text{CO}_2$ and with pH and osmolarity adjusted to 7.4 and 290–295 mOsm, respectively. Solution temperature was maintained at 36 $^{\circ}\text{C}$ with a TC-344B temperature controller and SH-27B in-line solution heater (Warner Instruments). Whole-cell recording pipettes with tip resistances of 4–7 $\text{M}\Omega$ were pulled from thin-walled borosilicate glass (World Precision Instruments, TW-150F) using a Sutter Instruments P-87. Pipettes were filled with a potassium gluconate-based internal solution containing the following (in mM): 130 potassium gluconate, 10 KCl, 10 HEPES, 0.2 EGTA, 4 ATP, 0.3 GTP and 10 phosphocreatine, with pH and osmolarity adjusted to 7.4 and 290–295 mOsm, respectively. Whole-cell recordings were performed using a Multiclamp 700B amplifier and digitized using an Axon Digidata 1440A (Molecular Devices). Recordings were made from both fluorescent and nonfluorescent neurons in the mPFC, visualized with an Olympus BX51WI equipped with infrared differential interference contrast and GFP epifluorescence. Recordings were performed in current-clamp with bridge balancing. After 10 min to ensure patches had settled, a 1-s, 200-pA current pulse was injected to induce spiking. After this, 10-min continuous recordings were performed. 1 min into the continuous recordings, ACSF containing 1 μM CNO was switched into the perfusion system, and the effects on resting membrane potential were observed. After this, a second 200-pA current pulse was injected to induce spiking in order to compare spiking before and after CNO application. All reagents used for electrophysiology were obtained from Sigma.

Behavioral analysis. All data analysis was fully automated. Data were analyzed in Matlab using a custom-designed toolbox called the Matlab water-maze (MWM) toolbox. This toolbox is freely available for download under a GNU Lesser Public License agreement at http://www.franklandlab.com/?page_id=306/. Additionally,

the analysis used the CircStat Toolbox⁴⁶ and the kernel density estimation (KDE) Toolbox (see below), which is also available under a GNU Lesser Public License agreement at <http://www.ics.uci.edu/~ihler/code/kde.html>.

Path and platform data from each training, testing and probe trial were loaded into Matlab directly from the Actimetrics software data files and used to estimate escape latency, platform crossings, time in quadrants, time in different platform zones, thigmotaxis, path entropy, mean angular search position and D_{KL} .

Escape latency (**Figs. 3c and 4b**, and **Supplementary Figs. 4 and 8**) was defined simply as the length of the trial. Platform crossings (**Fig. 1e** and **Supplementary Fig. 5c**) were defined as the number of times during the probe test when the mouse entered a region of the pool that had been occupied by a platform at any time during training. Time in platform zone (**Figs. 2c, 3f, 4e, 6c,f** and **Supplementary Fig. 7a,c**) was defined as the percentage of time during the probe trials that the mice spent in a circular zone with 10 cm radius centered on the platform. In experiments with two probe trials both platform crossings and time in platform zone measurements were averaged across both probe trials. Mean angular location (**Fig. 2d** and **Supplementary Fig. 7b**) was calculated by converting the search paths into polar coordinates and taking the circular mean via the function provided in the CircStat toolbox. Thigmotaxis was measured by examining the percentage of time that the animals spent in the outer 10-cm ring of the pool (**Supplementary Fig. 6**).

Path entropy (**Fig. 1f** and **Supplementary Fig. 5d**) and D_{KL} (**Figs. 1d, 3e, 4d** and **6b,e**) were estimated by calculating a probability density function from each search path during the probe using a kernel density estimation technique. Search paths from both trials were used as the centers for circular two-dimensional (2D) Gaussian kernels with bandwidths of 5 cm. These kernels were then used to generate a probability density function estimate, P_{search} (normalized across the space of the pool), at each grid point in a 1 cm \times 1 cm grid covering the pool. The path entropy, H_{search} , was then calculated as:

$$H_{\text{search}} = - \sum P_{\text{search}}(x) \ln(P_{\text{search}}(x))$$

where x ranges over space in the pool.

D_{KL} provides an estimate of the difference between two statistical distributions³⁰. It does so by providing a measure of the average number of bits that are lost when information from a target distribution is encoded using an alternative distribution. We used the D_{KL} to measure the difference between the distribution of individual mouse search paths on the probe test and the distribution of the platforms experienced during training. To estimate the D_{KL} , we defined the target distribution as the probability density function, P_{platform} (**Figs. 1c, 2a and 4a**), which was estimated using kernel density estimation. This was done using each training platform location as the center for circular 2D Gaussian kernels. The bandwidth of the Gaussian kernels was optimized using a likelihood cross-validation search technique (implemented in the KDE Toolbox). The resulting bandwidth was 11.35 cm for the unimodal distribution (**Supplementary Fig. 1a**), 14.84 cm for the first new platform test distribution (**Supplementary Fig. 1c**) and 10.02 cm for the second new platform test distribution (**Supplementary Fig. 1d**). It should be noted that these specific bandwidths were not critical as analyses using different bandwidths produced similar results. Using P_{platform} and P_{search} , D_{KL} was calculated as:

$$D_{\text{KL}}(P_{\text{platform}} || P_{\text{search}}) = \sum P_{\text{platform}}(x) \ln \left(\frac{P_{\text{platform}}(x)}{P_{\text{search}}(x)} \right)$$

An example illustrating the calculation of D_{KL} for two different search paths is given in **Supplementary Figure 3**. The D_{KL} measure is particularly high when the search path is absent from a nonzero region of the P_{platform} distribution. Therefore, a higher mean D_{KL} indicates that mice did not provide sufficient coverage in their search of areas which had a nonzero probability according to the P_{platform} distribution.

Electrophysiology analysis. Electrophysiology data was analyzed in iPython using NumPy, SciPy and custom scripts. The change induced by CNO in the resting membrane potential was measured as the difference between the median membrane potential in the 1 min of continuous recording before CNO application, and the median membrane potential in the final 1 min window of the 10 min recording after CNO application. Spike rate was estimated by first detecting

spikes via a thresholding of the derivative of the spike trace, then calculating the spike rate as the number of spikes generated over the 1 s pulse of current. The effect of CNO on spike rate was measured by taking the spike rate observed after CNO application and dividing it by the pre-CNO spike rate.

Statistical analysis. Statistical analyses were conducted in Matlab (MathWorks). No statistical methods were used to predetermine sample sizes and exact group numbers were determined by animal availability. However, we did ensure that our sample sizes were similar to those generally employed in the field. For the comparison of averages between two groups (e.g., **Fig. 1d**), data were tested for normality using Lilliefors test. If data for either group were significantly nonnormal (with $\alpha = 0.05$), then a nonparametric Mann-Whitney *U* test was used (e.g., **Supplementary Fig. 5b**). If data from neither group were significantly nonnormal, then the groups were tested for equality of variance using an *F* test. If the two groups had significantly different variances (with $\alpha = 0.05$), then Welch's *t*-test was used; otherwise a standard two-sample *t*-test was used. For the comparison of the average of a single group to a specified mean (e.g., **Fig. 2b**) the data were tested for normality using Lilliefors test. If the data were

significantly nonnormal (with $\alpha = 0.05$), then a sign-rank test was used; otherwise a *t*-test was used. For the comparison of averages between groups presented with consistent or conflict platforms at different delays, two-way, fixed-effect ANOVAs with interaction terms were used (e.g., **Fig. 2d**). Lilliefors test of the residuals from the models confirmed that they were not significantly nonnormal. *Post hoc* comparisons between individual groups were done using Tukey's test to control for multiple comparisons (e.g., **Fig. 3e**). Correlations were performed using Pearson's correlation coefficient. *P* values for individual correlations were calculated using transformation to a Student's *t* distribution, and *z*-tests were used to calculate the significance of differences between correlations. All statistical tests were two-tailed.

A **Supplementary Methods** checklist is available.

44. Heffner, H.E. & Heffner, R.S. Hearing ranges of laboratory animals. *J. Am. Assoc. Lab. Anim. Sci.* **46**, 20–22 (2007).
45. Paxinos, G. & Franklin, K.B.J. *The Mouse Brain in Stereotaxic Coordinates* Ed. 2nd edn. (Academic, 2001).
46. Berens, P. CircStat: a matlab toolbox for circular statistics. *J. Stat. Softw.* **31**, 1–21 (2009).