Title

Environmental complexity modulates information processing and the balance between decision-making systems

Authors

Ugurcan Mugan[1], Samantha L. Hoffman[2], A. David Redish[1]*

Affiliations

1 Department of Neuroscience, University of Minnesota, Minneapolis MN, United States
2 National Institute on Drug Abuse (NIDA), National Institute of Health, Bethesda MD, United States

*Correspondence: redish@umn.edu

Abstract

Decision-making relies on dynamic interplay between multiple neural circuits and decision-making systems integrating past experiences, current goals, and environmental demands. Past studies of rodent decision-making have largely occurred in simple environments and on isolated brain areas. This leaves unclear the impact of environmental complexity on the interaction between multiple brain areas to set adaptive strategies. To fill this gap, we either recorded neural activity across the key regions of hippocampus (HC), dorsolateral striatum (DLS), and medial prefrontal cortex (mPFC) or chemogenetically inactivated mPFC in a separate cohort while rats foraged for food under changing rules in differently complex environment. Environmental complexity increased behavioral variability, lengthened HC nonlocal sequences, and introduced intermediary subgoals. We found contrasting representations between DLS and HC, supporting a competition between decision systems. mPFC activity was indicative of setting this balance, in particular predicting nonlocal HC coding. Inactivating mPFC impaired short-term behavioral adaptation and produced long-term deficits balancing decision systems. Our findings reveal the dynamic nature of decision-making systems and how environmental complexity modulates their engagement with implications for behavior within naturalistic environments.

Acknowledgements

We thank G. Cannan, T. Lawrence, J. Williams for help with experiments, as well as C. Boldt, K. Seeland, and A. Sheehan for technical assistance. We thank E. Krook-Magnusson and members of the Redish lab for comments on earlier drafts and useful discussions. This work was supported by NIH grants R01-MH112688 and R01-MH080318.

Author contributions

U.M., S.L.H, and A.D.R. designed the study. U.M. collected the recording data and S.L.H. collected the DREADD manipulation data. U.M. analyzed the data recording data and S.L.H analyzed the DREADD manipulation data. U.M. and A.D.R. wrote the manuscript.

Declaration of Interests

The authors declare no competing interests.

Introduction

Naturalistic behaviors often occur in highly complex and uncertain environments, and thus necessitate the ability to create efficient representations and strategies that are context and situation specific. Neurophysiologically, adaptive decision-making is mediated by circuit dynamics that are distributed and coordinated across multiple brain areas that represent different components of these complex conditions.[1-8] Yet, the role that environmental complexity may play in adaptive decision making is often overlooked and not well understood[9,10] Traditionally, investigations into decision-making have focused on only a single brain region and employed highly-idealized simple environments. This has intrinsically limited our ability to understand interactions between the brain's multiple decision systems, the role of environmental demands on these interactions, and how information is represented in different brain areas.[11] To bridge this gap, we simultaneously recorded from multiple, interconnected brain regions on a foraging decision task with changing reward rules across a systematic variation of environment complexity.

Decision-making is thought to arise from multiple, parallel systems in the brain. Deliberation (or prospective, model-based, planning) relies on internally generated models of the environment and expectations of the future.[12-15] Conversely, the procedural (or habit, model-free) system uses learned and cached situation-action associations.[16-18] Therefore, while deliberative behaviors are often flexible and amenable to novel situations[19,20], procedural behaviors, while computationally efficient, are often inflexible.[2,6,7,21,22] Previous evolutionary analyses have found that environmental complexity and topology affect the usefulness of these strategies.[10,11] However, it remains unclear how environmental parameters and geometry influence the neural representations that are used to select between candidate actions, and how and where information is integrated to guide action selection.

Three regions---hippocampus (HC), dorsolateral striatum (DLS), and medial prefrontal cortex (mPFC)---are hypothesized to make distinct contributions to the decision-making process. HC is hypothesized to encode the state-space and thus enable flexible behaviors[2,6,7,23-26], while DLS is hypothesized to encode inflexible action chains and thus enable automated behaviors.[18,27-34] The theory that decision-making arises from at least two competing systems raises questions about how the deliberative (HC mediated) and procedural (DLS mediated) systems are balanced to create adaptive behaviors.[3,8,15,35,36] Dorsomedial prefrontal cortex (mPFC) has been shown to encode various task relevant features, such as the problem space and strategy, and as such has been put forward as a candidate region responsible for balancing between these parallel decision-systems.[37-44]

Research over the past decade suggests that HC is a crucial component of the network supporting memory function, cognitive, spatial and non-spatial information.[2,6,7,45-50] Furthermore, pyramidal cells in HC ('place cells') have been shown to be tuned to different environmental parameters such as geometry and space.[6,51-55] In addition to spatially specific firing at the animal's current position, decoding methods applied to the HC ensemble activity have elucidated time-compressed sequential extra-field firing at decision times---such as during single cycles of theta oscillations---that suggest imagination of future outcomes.[12,56-63] Thus, HC is thought of as a key component of prospective thinking, future planning, and deliberative decision-making. However, it is still an open question as to how environmental geometry changes HC network-level representations and how spatial complexity changes the features of nonlocal sequential activity, and thus how it influences prospective processes.

In contrast, dorsolateral striatum (DLS) is hypothesized to encode cached action sequences for procedural learning and decision making.[17,18,21,29,64] Prior research has shown punctate bursting activity by DLS ensembles that occur at the start and end of action sequences.[27,31,36] Furthermore, decoding of striatal activity during these bursts has shown that they reflect the current state and past experiences more than future options.[65] Together these results suggest that DLS is involved in initiating stereotyped behaviors associated with procedural decision-making. However, it is still an open question as to how environmental geometry changes these firing patterns, and thus how it influences procedural responses.

While there have been numerous studies looking at procedural and deliberative decision-making, and their neural instantiations, existing evidence has not addressed the role that task and environmental structure play in modifying the usefulness of procedural or prospective decision-making in differently-complex environments and the balancing between them. Many behavioral tasks commonly used to investigate these questions involve spatial tasks that are very simple (e.g., T-mazes). However, to understand behavior and neurobiology in naturalistic settings, we need to go beyond this simplicity to incorporate complexity into our studies.

To elucidate the arbitration between decision-making processes within complex environments, we developed a dynamic, complex spatial foraging tasks for rats. We used a variant of the left/right/alternation task, a spatial rule switching task in which rats were required to adjust their behavioral strategies after uncued changes to the reward contingency.[36,37,44,60,66-72] We quantitively varied the complexity to explicitly evaluate its impact on behavior and neural activity. Using simultaneous silicon probe recordings of HC, mPFC, and DLS ensembles in freely behaving rats, we found that environmental geometry impacted task-relevant neural representations and changed the timing of the transition between decision-making systems. Furthermore, chemogenetic inactivation of dorsal mPFC impaired the ability of the rats to recognize task regularity and stability, and led to both acute and long-term learning deficits, thus further highlighting mPFC's role in balancing between deliberative and procedural decision-systems.

Results

Environmental complexity differentially impacted the deliberative and procedural systems

To study the neural instantiations and representations of the two decision-making systems and the arbitration between them, we used a left/right/alternation (LRA) foraging task (Fig. 1A, Supplementary Fig. 1A)[36,37,44,60,66-72], and collected pairs of simultaneous recordings from dorsal HC CA1, dorsal mPFC, and DLS (Fig. 1B; $n_{HC} = 7$; $n_{DLS} = 4$; $n_{mPFC} = 7$). Using silicon probes, we recorded a total of 4292 single units from HC, 2293 from DLS, and 2609 from mPFC (Fig. 1B, Supplementary Fig. 1C, Supplementary Table 1). To study the role of mPFC in selecting the appropriate decision-system, in a separate cohort of rats, we chemogenetically inactivated dorsal mPFC (CaMKIIα-hM4Di DREADDs; see Methods) (Fig. 1C, Supplementary Fig. 6A, B; $n_{DREADDs\ Virus} = 8$, $n_{Control\ Virus} = 8$).

In the LRA task, rats had to traverse a central maze segment that included three low-cost decision points with dead ends on either side, followed by a high-cost choice point between the left and right return arms (Fig. 1A, Supplementary Fig. 1A).[36,37,44,60,66-72] There were three rules that determined which return arm would be rewarded as a function of the behavior of the rat. Under the left rule (L), food reward was presented on the left side; under the right rule (R), food reward was presented on the right side; and under the alternation rule (A), food reward was presented on the opposite side of the previous run ---

i.e., the rats had to alternate between left and right to receive a food reward. If the rats made the correct decision under the active rule, they received food at the reward arm and again at the start of the maze. Conversely, if the rats made an incorrect decision, they received no food reward at either location. Over the course of a 45 min session, there were three un-cued rule switches per session that occurred roughly at 15 and 30 mins. Each session included all three rules; thus, sessions were drawn from one of six possibilities: LRA, LAR, RLA, RAL, ALR, ARL (Fig. 1D).

In line with previous research[36,37,44,60,66-72], rats rapidly adjusted their behavior to the new rule (Fig. 1E, Supplementary Fig. 2A) and performed similarly under each rule with a slight decrease in performance under the alternation condition (Supplementary Fig. 2B, C; mean ± s.t.d. 45.7 ± 12.9 laps/block; mean ± s.t.d. 82.7 ± 10 %correct L & R and 70.3 ± 12 %correct Alt.).

This task is known to probe the interaction between deliberative and procedural decision-making and the neural correlates of their underlying systems (Fig. 1D).[38,44,71,73,74] Prior work has found that rats display vicarious trial and error (VTE) behaviors around rule switches at high-cost choice points. VTE behaviors entail pausing and orienting back and forth towards the two arm options before making a decision---quantified by IdPhi or when z-scored zIdPhi (Fig. 1F, Supplementary Fig. 2G; see Methods).[15,75,76] Neural recordings indicate that VTE is often accompanied by HC cell assemblies sweeping forward of the animal serially assessing each option, and thus is believed to reflect a behavioral correlate of deliberation.[12,56,60,67] In line with previous research, we found that VTE behaviors significantly increased after rule switches (Fig. 1G), which correspond to periods that are hypothesized to be driven under the deliberative system (Fig. 1D).[37,69,70]

In contrast, prior work has shown that rats develop procedural responses on this task---i.e., behavioral stereotypy---under stable contingencies and after the contingency has been learned.[28,36,65,74] Here, we quantified behavioral stereotypy by measuring the reciprocal of the Euclidean distance between pairs of path trajectories through the central maze segment on a lap-by-lap basis (Fig. 1H; see Methods). Notably, rats developed stereotyped path trajectories as they adopted the appropriate behavioral strategy (Fig. 1I). Both the pre-switch and post-switch periods were characterized by an increase in path similarity, with the rule changes disrupting the developed stereotypy.

Prior computational research suggests that environmental structure and complexity should impact the utility of using deliberative vs procedural strategies.[10,77] Here, we modulated environmental complexity by parametrically changing the central maze segment (central path) (Fig. 1A). Intuitively, high complexity environments are those that have highly variable navigational path lengths within the central segment. To quantify environmental complexity, we discretized the maze and created a graph of the discretized space. Structural complexity was defined as the total information content across all sets of equivalent trees that could be created from a given maze configuration (see Methods).[78-81] Walls used to manipulate the interior navigational sequence (A: open middle, B: open left, C: open right) created 27 unique mazes with different levels of complexity (Fig. 1J). These mazes roughly formed three clusters of complexity, which we partitioned into low, mid, and high complexity for subsequent analysis (Fig. 1J).

Theoretically, if rats are using different decision-making algorithms, they should show predictably different behavioral outcomes, neurophysiological representations, and behavioral-neurophysiological interactions based on the internal structure of the maze. Environments that feature multiple decision-points (low-cost decision-points aligned to walls B and C; Fig. 1A) are hypothesized to encourage decision-making that relies more on deliberation, thus resulting in more variable behaviors, longer

hippocampal nonlocal sequences, and differential engagement of dorsolateral striatum at these low-cost decision-points. In contrast, low complexity environments are hypothesized to encourage decision-making that relies more on habit, thus resulting in the development of stereotyped responses that span the entire central path, with shorter, less sequential hippocampal nonlocal sequences (Fig. 1K). Thus, this the environmental control should provide us with an experimentally tractable method for incorporating environmental complexity into a behavioral task that is well established for studying the engagement and competition between deliberative and procedural decision-systems.

To directly assess the role environmental complexity plays in modulating decision systems, we first analyzed the impact of changing complexity on overall task performance. We found that environmental complexity significantly affected performance: rats made more errors (poorer performance) in more complex environments (Supplementary Fig. 2D; One-way ANOVA F(18) = 2.35, p = 0.0033). Looking at the effect of environmental complexity on the behavioral markers of deliberative and procedural decision-making, VTE and stereotypy respectively, we found that there was a trend towards increased VTE on high complexity environments (Fig. 2A; One-way ANOVA F(18) = 1.59, p = 0.072).

Rats explored the central maze segment (Supplementary Fig. 1A) by occasionally wandering on unrepeated paths to the corners and edges of the dead-end spaces (Fig. 2B). These explorative events increased around rule switches and were correlated with the prominence of reorientation behaviors (Fig. 2C, D, Supplementary Figs. 2E, F, H). These data suggest that exploration also provides a behavioral correlate of prospection that, unlike VTE, is associated with the central path. Notably, the proportion of exploratory laps increased with environmental complexity, especially in highly complex environments (Fig. 2E; One-way ANOVA F(18) = 2.61, p = 0.0011), which is in line with the idea that environmental complexity modulates the competition between deliberative and procedural decision-making. Therefore, exploration, a behavior that only becomes observable in complex tasks that feature multiple decision points, also provides insight into different facets of deliberative decision-making.

In line with the modulation of deliberative behaviors with environmental complexity theory, we found that increased complexity resulted in overall lower behavioral stereotypy and slower development of stereotypy following rule switches (Fig. 2F, G).

Taken together these results indicate that modulating the environmental structure has important consequences on behavior and the brain's multiple decision-making systems.

Hippocampal sequences were modulated by environmental complexity

Place cells commonly found within the CA1 region of the HC have been shown to provide direct insight into how the brain's deliberative decision system represents important information.[12,15,20,82-85] Place fields appeared in both simple and complex environments (Fig. 3A), and the number of place cells and their distribution through the maze did not differ across the different environments (Supplementary Fig. 3B, C). As expected, the firing of putative pyramidal cells (89.3% recorded HC units) was largely local throughout the maze (Fig. 3B). However, place fields were significantly smaller in more complex environments, suggesting differences in network-level representations of space in differently-complex environments (Supplementary Fig. 3D; One-way ANOVA F(2) = 9.91, p < 0.001). These differences were not driven by differences in speed through the central path (Supplementary Fig. 3A; One-way ANOVA F(18) = 1.29, p = 0.20). To further investigate how these network level differences may affect nonlocal HC

representations, and putative physiological measures of deliberative thinking, we looked at theta timescale (hundreds of millisecond) sequences (Fig. 3C).

Theta oscillations (6-12 Hz) are prominent in CA1 and can contain fast time scale sequential activity of place cells.[12,56,63,86-91] To statistically identify significant theta timescale sequential activity in CA1, theta cycles were identified and limited to those that had ≥5 cells active. The spike trains within each candidate theta cycle were analyzed using Bayesian decoding[92] (20ms bins) to relate firing activity to underlying behavior, and the correlation between decoded position and time was identified as a metric of sequential activation (theta sequence score).[93-95] Intuitively, this metric indicates the ordered nature and clarity of the decoded representations in each theta cycle (see Fig. 3D, E for high and low score examples, respectively). Using this method, we found clear significant theta sequences during the entire navigational period.

While the proportion of significantly sequential theta cycles did not differ between environments (Fig. 3F; One-way ANOVA $F(12) = 1.23$, $p = 0.28$ and Supplementary Fig. 3E), we found that the average theta sequence score across all the theta cycles significantly increased with complexity (Fig. 3G; Kruskal-Wallis test $H(2) = 765.94$, $p < 0.001$). Moreover, the average theta sequence during traversals along the central path increased with complexity (Fig. 3H; One-way ANOVA $F(12) = 2.48$, $p = 0.01$). Notably, similar results were found when taking the slope of the decoded sequences (Supplementary Fig. 3F; Kruskal-Wallis test $H(2) = 66.67$, $p < 0.001$). Collectively these results suggest that increasing environmental complexity results in more robust nonlocal sequence representations in HC.

In conjunction with sequential activation of place cells, prior research suggests that theta cycles consist of two components: 1) local representation of the current location during the first-half of theta, and 2) a nonlocal sequential sweep ahead during the second-half of theta.[58] To quantify the prevalence of nonlocality during theta, we used Bayesian decoding to compute the difference in the posterior probability of the decoded location between the second- and first-half of theta aligned to the actual position of the rat (Fig. 3I). Consistent with past work, we found that the first half of theta contained more local information, and the second contained more nonlocal.[66,96-98] Importantly, environmental structure modulated this balance such that nonlocal decoding during the second half of theta was more prominent in more complex environments (Fig. 3I).

Together, these results suggest that environmental structure plays a pronounced role in modulating sequential hippocampal activity during theta cycles. Specifically, theta cycles showed more nonlocal activity with longer look-ahead sequences in more complex environments, which likely require increased cognitive and memory load. This network level modulation of HC activity implies important changes to the deliberative decision-making system based on environmental structure, with more complex environments giving rise to stronger neural representations of future options.

Theta sequence scores were correlated with deliberative behaviors

Given that we saw changes to CA1 theta sequences with increased environmental complexity, we asked whether there was a coherent relationship between HC nonlocal activity and deliberative behaviors. Briefly, is the prominence of nonlocal HC representations indicative of increased deliberative future thinking, as measured through exploration and VTE?

Consistent with this hypothesis, average theta sequence scores at the choice point were significantly correlated to zIdPhi values (associated with VTE events) (Fig. 4A; t-test(84) = -3.3, $p = 0.0015$). In line

with previous findings, we found higher theta sequence scores on VTE laps (Fig. 4B; Kruskal-Wallis H(1) = 12.71, p = 0.0004).[44,66] Similarly, , we found that the average theta sequence scores and the amount of exploration in the central path were also significantly correlated (Fig. 4C; t-test t(84) = -3.3, p = 0.0015): theta sequence scores were significantly higher on exploratory laps (Fig 4D; Kruskal-Wallis H(1) = 37.64, p < 0.001). These results replicated when we instead examined sequence slopes (Supplementary Fig. 4A-D).

When the HC engages in prospective representations for use in deliberation, we would expect behavior to be less driven by the procedural system. To directly test this, we calculated the average stereotypy of each lap comparing the similarity of its trajectory to prior and subsequent laps (± 2 laps) (see Methods). Notably, trajectory stereotypy and theta sequence scores were inversely correlated (Fig. 4E; t-test t(84) = 5.79, p < 0.001 and Supplementary Fig. 4E), in line with theories that hypothesize a competition between the deliberative and procedural decision-making systems.

These results suggest that CA1 theta sequences accompanied deliberative behaviors, and importantly did not have the same level of sequential structure during procedural behaviors, suggesting that, theta sequences may be underlying a mechanism for fast timescale memory recall and deliberation.

Environmental structure gave rise to subgoals and modulated procedural behaviors

After adapting to the new task rule, rats automated their behaviors and showed increased behavioral stereotypy (Fig. 1I, Fig. 2G). This switch to behavioral automation is hypothesized to be driven by dorsolateral striatum (DLS) and is often associated with the development of characteristic firing patterns of DLS cell assemblies, which are hypothesized to reflect the boundaries of well-learned action sequences. These boundaries are believed to indicate the beginning and end of cached action sequences---referred to as "task-bracketing".[21,27,30-32,99] To examine the firing pattern of DLS cells, we quantified the average firing rates of putative medium spiny neurons (81.6% of recorded DLS cells) across the linearized maze (~3cm bins) per lap for 25 laps aligned to the start of a session and rule switches. Importantly, DLS firing patterns were highly influenced by environmental complexity (Fig. 5A). In simple environments, which are comparable to other experimental paradigms in which task-bracketing has been studied, such as the classical T-maze, we found action boundaries at the start of the maze and at the choice point (Fig. 5A left panel). However, increasing complexity changed DLS activity throughout the navigation sequence, especially around these walls.

Previous research has quantified task-bracketing---increased firing at the beginning of runs with behavioral automation---as the subtraction of firing rate in the central path from the firing at the start of the maze (task-bracketing index).[27,29,30] Positive task-bracketing values suggest the formation of meaningful action boundaries for automation, while negative task-bracketing values indicate that the ensemble firing increases while the action sequence is being implemented and suggest that the action sequence has not been chunked into initiation and termination. Interestingly, we found that task-bracketing index decreased with environmental complexity (Fig. 5B; One-way ANOVA F(8) = 3.64, p = 0.0025), as would be expected from slower behavioral automation and increased firing in the central path (Fig. 5B).

Given the firing patterns that we observed in mid and high complexity environments, we hypothesized that the walls might have constituted intermediary subgoals or action boundaries (Fig. 5A bottom row). Spatial cross correlation of ensemble activity along a linearized projection of the path elucidates

representational similarity between different portions of an environment and provides valuable information of how the environment might be chunked. For example, if ensemble activity is largely local and distinct, one would expect to see high correlation along the diagonal (for HC example see Supplementary Fig. 5A). We found that in simple environments, DLS represented the central path almost as a single entity, where activity along the central path was highly correlated throughout (Fig. 5C left panel; broad block structure), with a representational switch appearing at the choice point. In contrast, with increased environmental complexity, DLS representations became much more local, and, interestingly, in these environments, DLS represented each set of walls similarly (Fig. 5C middle and right panels, Supplementary Fig. 5B).

Given the representational similarity at the walls in complex environments, we hypothesized that walls may be constituting subgoals. This would suggest that DLS ensemble bursts should also be occurring at these locations to indicate subgoal to subgoal action chunking. To measure this, we quantified the proportion of bursting activity at the different maze locations (Supplementary Fig. 1A). We found a significant difference in the proportion of bursting activity at the start of maze, choice point, and at intermediary walls across differently-complex mazes, suggesting multiple action boundaries in complex environments (Fig. 5D; One-way ANOVA at each maze location) in line with the development of subgoals throughout the central path.

We found that HC theta sequences were correlated with deliberative behaviors and inversely correlated with procedural behaviors (Fig. 4E). Concurrently, previous work suggests that task-bracketing is related to the development of procedural behaviors and inversely correlated with VTE behaviors.[10,21,27,30-32,77,99] Theories hypothesizing a competition between decision systems predict that there should be an inverse relationship between theta sequence scores and task-bracketing, however no one has directly compared these two neural signatures. To compare them, we plotted HC theta sequence scores and DLS task-bracketing aligned to rule changes. We found that there was an increase in theta sequence score and corresponding decrease in task-bracketing aligned to rule switches predominantly in simple environment. While there was a slight dynamic change in DLS task-bracketing in mid complexity environments, there was no corresponding decrease in task bracketing in complex environments, this was because task bracketing remained low (or never developed) in complex environments (Fig. 5E).

While this analysis gives insight into the dynamics of DLS and HC activity around behaviorally relevant events, it does not provide direct insight into how DLS firing patterns are related to HC theta sequences at a finer, within-session scale. Two rats had simultaneous DLS and HC recordings allowing us to quantify the relationship between DLS task-bracketing and HC theta sequence score on a lap-by-lap basis. While there was a significant negative correlation between task-bracketing and theta sequence score in low complexity environments (r_{Low} = -0.13, p_{Low} < 0.001), there was no relationship in mid and high complexity environments (r_{Mid} = 0.03 , p_{Mid} = 0.34; r_{High} = 0.05, p_{High} = 0.31) (Fig. 5F, Supplementary Fig. 5C) . These findings demonstrate that the inverse relationship between HC and DLS held in simple environments, even when looking at fine scale dynamics, but broke down in more complex environments. This suggests that the potential arbitration between the two decision-systems, deliberative and procedural, may not be as distinct in environments that require a significant memory/cognitive load and further highlights the importance of employing varying environmental complexity when looking at the competition between decision-making systems.

mPFC encoded changes to strategy by engaging the deliberative system around rule switches

The prefrontal cortex is hypothesized to encode the 'problem task space', goals, subgoals, and strategy.[37-39,100,101] Accordingly, mPFC has been associated with balancing deliberative and procedural decision-making.[42,102,103] Previous research has found that mPFC changes its representations around rule switches and behavioral shifts between the two decision systems on the LRA task.[37,38] One analysis for identifying these changes in mPFC representations has been to correlate lap-by-lap ensemble firing rates.[37] In line with previous results, population activity in mPFC was highly consistent when behavior was automated (Fig. 6A left panel)---i.e., prior to the switch and after ~10 laps into the next rule. However, when the rule changed, so too did mPFC ensemble activity.[37,38] Interestingly we found similar ensemble dynamics in DLS and HC (Fig. 6A).

If mPFC is responsible for setting strategy, we would expect mPFC to change its firing patterns sooner after a rule switch as compared to either HC or DLS. To identify when ensemble changes occurred, we used a combination of hierarchical clustering to identify reliable ensemble representations and a change point analysis to identify moments that neural activity shifted (Fig. 6B; see Methods). Across all recordings, representational transitions across all three regions predominantly occurred after rule switches, however, changes to mPFC ensemble firing were earlier (Fig. 6C). This overall finding would suggest that mPFC representational changes leads those in HC and/or DLS, in line with a putative role of mPFC in setting a strategy and relaying this information to downstream regions.[38]

To directly evaluate the change point timing between regions, we leveraged our simultaneous recordings between pairs of brain areas (mPFC-HC and mPFC-DLS; n_{Rats} = 5, n_{Rats} = 2). We found that changes to mPFC ensemble firing patterns occurred prior to both HC (mean 1.4 laps; Wilcoxon signed-rank test z = -2.39, p = 0.017) and DLS (mean 9.3 laps; Wilcoxon signed-rank test z = -4.01, p < 0.001) (Fig. 6D). Together these results suggest that the representational changes in mPFC ensembles reflected a change in decision-making strategy (e.g., from procedural to deliberative) that needed to be transmitted to HC and DLS. To study the information transfer between mPFC and HC for strategy setting, we next looked at short timescale prediction of HC state by mPFC activity.

To determine whether mPFC activity carried information about what HC was representing, we used a 5-fold cross-validated Bayesian decoding to predict theta sequence scores from mPFC ensemble activity in rats with simultaneous recordings between the two structures (n_{Rats} = 5) (see Methods). Given the positive relationship between theta sequences and deliberative behaviors, mPFC's prediction of HC nonlocal representations and their sequential structure would elucidate mPFC's role in switching from procedural to deliberative decision-making.

Critically, we found that mPFC activity did in fact predict the nature of nonlocal representations in HC. mPFC activity was more predictive of the HC state compared to: 1) when mPFC representations were destroyed (circularly shifted spike times), and 2) when short-timescale coupling between mPFC and HC state was destroyed, even though internal representations of mPFC were preserved (Fig. 6E). Thus, mPFC representations encoded strategy and could be used to predict HC sequential activity (HC state). However, mPFC did not predict the HC state on a theta-cycle-by-theta-cycle basis. mPFC activity was not more predictive than shuffled theta sequence scores restricted to a lap (breaking short-timescale relationship between mPFC and HC state while preserving longer timescale relationships). Therefore,

there was a significant coupling between mPFC activity and HC state, but it was over longer behaviorally relevant timescales (Fig. 6E), suggesting a role in strategy setting on any given lap.[44]

Inactivating mPFC reduced deliberative behaviors and delayed behavioral adaptations to rule switches

While we found a relationship between mPFC population activity and initiation of strategies, it remained unclear to what extent mPFC causally influenced the balance between deliberative and procedural strategies. To test this question, we used inhibitory DREADDs (designer receptor activated by designer drugs[104]) (AAV5-CaMKIIα-hM4Di(Gi)-mCherry) (n_{Rats} = 8; DREADDs) compared to an across-subject - mCherry control (n_{Rats}= 8; Control) to investigate the behavioral effects of dorsal mPFC inactivation (DCZ[105] condition) (Fig. 7A, Supplementary Fig. 6A, B; see Methods) ($n_{Sessions}$ = 508). To investigate the long-term effects of mPFC inactivation and 2^{nd} day compensatory changes after either mPFC inactivation (hereafter referred to as DREADDs: DCZ) or within-subject control (hereafter referred to as DREADDs: VEH) conditions, we also included a post-manipulation saline repeat of the task (post DCZ or post VEH), which provided us with a washout day (Supplementary Fig. 6C).

Interestingly, we found that mPFC inactivation resulted in overall worse performance when compared to the across-subject control virus rats, pointing towards a fundamental involvement of mPFC functioning in proper execution of the task (Fig. 7B, Supplementary Fig. 6D, E). Critically, the central component of the task is the rats' ability to adapt their use of deliberative vs procedural decision-making following a change to the task rule. To quantify how mPFC manipulation changed the speed of behavioral adaptation, or the time it took for responses to stabilize into a new pattern, we computed the behavioral change point (see Methods). mPFC inactivation (DREADDs: DCZ) led to slower behavioral adaptation to rule switches when compared to post-manipulation saline day and VEH day, as well as control virus animals (Fig. 7C; One-way ANOVA F(3) = 3.51, p = 0.015; HSD Post-hoc DCZ vs all others p < 0.05, all other comparisons p > 0.05).

Given evidence that mPFC activity may be influencing deliberative strategies through HC engagement (Fig. 6D, E), we investigated the effects of mPFC inactivation on VTE and exploration, behaviors reflective of deliberation and putative HC engagement. Surprisingly, we found that the proportion of VTE events under inactivation (DREADDs: DCZ) was not different from the proportion of VTE events observed in mPFC intact animals (Control) (Fig. 7D; Wilcoxon signed-rank test z = 0.039, p = 0.97). Similarly, the amount of space explored after mPFC inactivation was comparable to control virus animals (Fig. 7F; Wilcoxon signed-rank test z = 0.95, p = 0.34). Thus, while mPFC inactivation impaired the rats' ability to adapt to new task rules, we paradoxically found no evidence for acute deficits in rats' ability to employ deliberative behaviors. One possibility is that while rats engaged in deliberative behaviors, inactivation of mPFC impaired their ability to use the information that they gained about the task and rules to perform efficiently.

Building on this hypothesis that mPFC serves a critical role in the use of gained information, we turned to a unique component of our experimental design. By repeating each behavioral condition on a second day in a saline washout condition between manipulations (DCZ rebound or VEH rebound), we were able to directly evaluate prolonged short-term impacts of mPFC inactivation on task performance and compensatory behavior rebounds following acute mPFC inactivation. Interestingly, the effects of mPFC manipulation were much more pronounced when we considered behavioral differences between these rebound effects. There was a significant compensatory increase in VTE behaviors following mPFC

inactivation on the previous day (Fig. 7E DREADDs: DCZ rebound; Wilcoxon signed-rank test z = 2.53, p = 0.0043); but a significant decrease in VTE if mPFC functioning was not manipulated (Fig. 7E DREADDs: VEH rebound; Wilcoxon signed-rank test z = -1.65, p = 0.049 and Supplementary Fig. 6F). Similarly, exploration increased on the second day following mPFC inactivation (Fig. 7G DREADDs: DCZ rebound; Wilcoxon signed-rank test z = 3.63, p < 0.001), but decreased on the second day if mPFC functioning was not manipulated (Fig. 7G DREADDs: VEH rebound; Wilcoxon signed-rank test z = -2.41, p = 0.0079 and Supplementary Fig. 6G). Neither of these compensatory effects after mPC manipulation were seen in control animals (Fig. 7E; Control: DCZ rebound Wilcoxon signed-rank test z = -0.13, p = 0.9, Control: VEH rebound Wilcoxon signed-rank test z = -0.38, p = 0.7 and Fig. 7G; Control: VEH rebound; Wilcoxon signed-rank test z = -1.63, p = 0.1, VEH condition: Wilcoxon signed-rank test z = -0.18, p = 0.89).

Given the impact that our mPFC manipulation had in eliciting compensatory behavioral rebounds following inactivation, we hypothesized that our manipulation may be causing long-term deficits in strategy setting and learning. To analyze the potential effects mPFC manipulation may have had on deliberative strategies, we compared the changes to VTE and exploration around rule switches in DREADDs (active virus) and control virus animals by pooling across all drug and rebound conditions. In line with our previous findings, we found that both the prominence of reorientation behaviors (zIdPhi) and the proportion of exploratory laps increased around rule switches in both DREADDs and control virus animals. Importantly, this increase happened earlier in control animals, suggesting that control animals switched decision-making strategies faster (Fig. 7H; Repeated measures ANOVA around switch interaction F(5) = 2.61, p = 0.023 and Fig. 7I; Repeated measures ANOVA around switch interaction F(6) = 2.12, p = 0.048, Post-hoc on each position bin via two-sided Wilcoxon rank sum corrected for multiple comparisons).

To determine if mPFC inactivation affected rats' ability to employ procedural strategies, we compared the development of behavioral stereotypy around rule switches in DREADDs and control virus animals. We found that DREADDs animals developed behavioral stereotypy slower and showed overall less stereotypy compared to control virus animals (Fig. 7J, Supplementary Fig. 7A). This is in line with previous findings that suggest dorsal mPFC disruption causes long-term deficits in habit formation.[106]

Finally, we evaluated the long-term progression of task performance over the course of our multi-week experiment. Notably, we found that DREADDs animals performed worse throughout the entire experimental sequence (Fig. 7K; ANOVA virus F(1) = 36.97, p < 0.001). In considering how rats' behavioral strategies may have evolved over the course of the experimental timeline, we hypothesized that as rats became more familiar with the general task structure (rule changes and adaptation), deliberation became less critical. In line with this hypothesis, we found that control rats exhibited progressively less exploratory behavior over the course of the four-week experiment (Supplementary Fig. 7C; ANOVA virus*week interaction F(3) = 2.65, p = 0.049), though VTE remained stable (Supplementary Fig. 7B; ANOVA virus*week interaction F(3) = 0.77, p = 0.51). In contrast, DREADDs rats did not exhibit this observed decrease in exploration or VTE.

Collectively, these data suggest that mPFC inactivation impaired long-term behavioral adaptation and recognition of task regularities. Furthermore, it impaired effective engagement and balancing between deliberative and procedural decision-making systems. These deficits resulted in both short-term and long-term performance decreases and highlight the importance of mPFC in this system.

Discussion

Prior research into decision-making suggests that adaptive decision-making necessitates the interaction between multiple competing decision-making systems, two of which are deliberative (or prospective) and procedural (or habit). Current theories suggest that these two systems depend on different neural circuits, are optimized for different situations, and are reflected in behavioral differences during decision-making.[1-3,6,7,107,108] However, the environment in which these decision-making processes are taking place is often overlooked. In the present study, we examined how: behavior, neural representations, and the interactions between the key brain regions of hippocampus (HC), dorsolateral striatum (DLS), and medial prefrontal cortex (mPFC) changed in uncertain and differently-complex spatial environments that better approximate naturalistic scenarios. The left/right/alternation foraging task with changing internal barriers, coupled with simultaneous recordings across dorsal HC, DLS, and dorsal mPFC, allowed us to directly and quantifiably study the effects of environmental complexity on neural representations and decision-making strategies.

Past research has emphasized the competition between deliberative and procedural control of behavior and has suggested that they arise based on the animal's experience and the consistency of the environment.[4,7,16,64,109] Deliberative control is favored when the optimal actions remain variable.[2,6,7,15,110-113] In contrast, procedural control is favored when the optimal behavior is relatively consistent across time.[2,6,7,27,99,107,114] In line with this idea, decision-making strategy in our task changed from initially deliberative on entry into the environment to procedural as the rats settled into a behavioral flow, but then transitioned back to deliberative around rule switches before becoming automatic again under the procedural system after adaptation to the new rule. Importantly, environmental structure affected this process, slowing behavioral adaptation in complex environments (Fig. 2A, E, G).

Competition between deliberative and procedural decision-making in differently-complex environments

The LRA task engages the brain regions hypothesized to play prominent and distinct roles in deliberative and procedural decision-making (HC, DLS, and mPFC) at distinct moments in the task that correspond to hypothesized periods of deliberative and procedural decision-making (Fig. 1).[38,44,71,73,74]

The role of HC is often thought of as providing an adequate model of the environment (or cognitive map) to support future thinking of potential action sequences and their outcomes to aid deliberative decision-making.[2,7,12,15,20,57,60,82,115,116] We found that time-compressed HC nonlocal sequences predominantly occurred during deliberative behaviors such as VTE (Fig. 4A, B). Importantly, we observed another deliberative behavior, which we identified as exploration, that was also affected by maze complexity and was correlated with the prominence of HC nonlocal sequences, significantly more than VTE, thus further buttressing the role of HC in model-building (Fig. 2B-E, Fig. 4A-D, Supplementary Fig. 4A-D). Furthermore, complex environments featured more theta-timescale nonlocal activity that were longer and more prominent (Fig. 3). Together these data suggest that prolonged deliberation in complex environments may indicate increased cognitive/memory load due to the need to create fine-scale schema of the environment to learn the consistencies therein.

Prior research has shown that there is a parallel between deliberative decision-making and curiosity, which is often associated with the intrinsic motivation and natural tendency to proactively explore the environment to gather information about its structure.[117] Thus, this structural knowledge can be flexibly

used when planning actions in new environments or when reward locations change.[2,6,19,20,62,84,112,118] Here, the complexity of the environment---which quantifies the diversity and length of available paths within the central segment of the maze---may be changing the utility of exploration (Fig. 2E, Supplementary Fig. 2E, F). In low complexity environments the ability of the animal to visually observe the high-cost choice point at the end of the central path may be enough for it to infer the structure of the environment, decreasing the overall usefulness of physical exploration. Paradoxically, we found that propensity to explore increased around rule changes, with more complex environment exhibiting more exploration (Supplementary Fig. 2I). Considering that exploration of the central path does not confer any advantages for task performance, this observation suggests any task changes may induce a general rise in uncertainty, promoting re-exploration of the entire physical space, where the degree of this sensitivity scales with the underlying environmental complexity.[119-121] Another possibility is that the diversity of the number of paths increases the perceived state-space of the problem thus favoring increased physical exploration that implicitly organizes the mental search and planning occurring over the state-space.

Procedural decision-making is often thought to be encoded by DLS bursts aligned to meaningful boundaries of well-learned action sequences.[18,21,27,28] Consistent with previous research[36,65,99], we found increased firing at the beginning and end of action sequences in simple mazes (Fig. 5A, B, D). Task-bracketing and DLS representations of cached action-sequences have primarily been studied in simple environments, such as T-mazes. In more complex environments, DLS ensemble firings were skewed toward the turn sequence of the task, as well as reward locations and the start (Fig. 5A, D). Notably, even in simple environments, we found elevated DLS firing at the choice point that developed over the course of a stable contingency, suggesting that there may be multiple cached action sequences, e.g., from start to choice point, choice point to feeder, and feeder to start (Fig. 5A, D). Furthermore, DLS representations of space changed with environmental complexity, from highly similar throughout the navigation sequence in simple environments, to distinct and only similar at wall locations in complex environments (Fig. 5C). If elevated firing of DLS ensembles at action boundaries are indicative of action caching, one possibility is that environmental complexity is favoring the creation of intermediary subgoals aligned to the low-cost decision-points (Fig. 5A, C).

The dorsolateral striatum receives most of its excitatory inputs from sensorimotor cortex and thalamus, and can influence movements by modulating, for example, motor cortical activity.[122-124] This raises questions as to cognitive vs motor components in DLS firing patterns. While it is possible that the increased firing and shared representations at wall locations (Fig. 5C; subgoals) may be a byproduct of distinct motor actions, these observations are also consistent with theories that suggest these bursts are indicative of action selection and initiation through the integration of sensorimotor, cognitive, and motivational information. It is worth noting that these hypotheses are not necessarily in conflict with each other.

Our findings suggest that the potential arbitration between the deliberative and procedural decision-making systems may not be as distinct in environments that require significant cognitive/memory load, and that striatum may not be encoding classical task-bracketing representations in complex environments aligned to task start and end points. The hypothesis that environmental complexity increases memory load is in line with computational research that suggests that planning and deliberation grow quadratically with the state-space.[11] In complex environments striatum may be encoding the state-space based on the regularity of their features, which may be manifesting as increased representations in the central segment of the maze (Supplementary Fig. 5D). Therefore, task-

bracketing in complex environments may not be the right metric to understand how DLS may be representing the boundaries of cached action-sequences.

The role of dorsal prefrontal cortex in arbitrating between decision systems

The hypothesis that behavior arises from an interaction between multiple algorithms leaves open the critical question about how disagreements between decision-making systems are resolved. Current theories suggest that these are resolved through information processing in the medial prefrontal cortex.[37-43,125,126] Consistent with this hypothesis, dorsal mPFC ensemble firing patterns changed following rule switches, and preceded representational changes in both HC and DLS (Fig. 6A, C, D). At the level of single HC theta cycles, mPFC activity predicted the prominence of HC nonlocal sequences, which further emphasizes mPFC's role in setting HC state to retrieve contextual information (Fig. 6E). However, this prediction was not on a short theta-cycle-by-theta-cycle timescale but rather at much longer lap-by-lap behaviorally relevant strategy timescales. Consistent with the idea that mPFC and HC coordinate to produce adaptive behaviors, prior research has shown that they exhibit oscillatory synchronization during decision-making tasks, which is interpreted as information transfer and integration between mPFC and HC.[13,127-130] mPFC leading both HC and DLS in changing its activity patterns following a rule switch (Fig. 6C, D), suggests that mPFC may be integrating contextual information to balance between deliberative and procedural decision-making, by updating the current task rule and inhibiting prepotent responses to enable flexible behaviors.

To examine the role of dorsal mPFC on behavioral adaptation and flexibility, we chemogenetically inactivated dorsal mPFC. We found that mPFC inactivation led to acute performance deficits and slower adaptation to rule changes. This is in line with previous research that has found disruptions in working memory and preservative errors with dorsal mPFC lesions.[131-133] Consistent with prior research that has shown that neural activity in dorsal anterior cingulate cortex shifts between exploration and choice repetition[134,135], we found that dorsal mPFC inactivation caused preservation of left or right choices after rule changes leading to maladaptive behaviors (Fig. 7C). Surprisingly, mPFC inactivation did not decrease the prevalence of deliberative behaviors when compared to control animals but resulted in long-term changes to strategy setting, suggesting inadequate learning from deliberative behaviors. Importantly, we saw compensatory mechanisms that increased deliberative behaviors following mPFC inactivation (Fig. 7E, G). This suggests that mPFC manipulations may be having long-term strategy effects, where continuous inactivation may be leading to global impairments in building context associations that are necessary for connecting a situation with a decision strategy that is appropriate for the degree of behavioral flexibility perceived to be required.

References

1   Balleine, B. W. & Dickinson, A. Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology* **37**, 407-419, doi:10.1016/s0028-3908(98)00033-1 (1998).

2   O'Keefe, J. & Nadel, L. *The hippocampus as a cognitive map*. (Clarendon Press ;

Oxford University Press, 1978).

Daw, N. D., Niv, Y. & Dayan, P. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci* **8**, 1704-1711, doi:10.1038/nn1560 (2005).

Dezfouli, A. & Balleine, B. W. Actions, action sequences and habits: evidence that goal-directed and habitual action control are hierarchically organized. *PLoS Comput Biol* **9**, e1003364, doi:10.1371/journal.pcbi.1003364 (2013).

Balleine, B. W., Doya, K., O'Doherty, J. & Sakagami, M. Current trends in decision making. *Ann N Y Acad Sci* **1104**, xi-xv, doi:10.1196/annals.1390.2226 (2007).

Redish, A. D. *Beyond the cognitive map: from place cells to episodic memory*. (MIT press, 1999).

Redish, A. D. *The mind within the brain : how we make decisions and how those decisions go wrong*. (Oxford University Press, 2013).

van der Meer, M., Kurth-Nelson, Z. & Redish, A. D. Information processing in decision-making systems. *Neuroscientist* **18**, 342-359, doi:10.1177/1073858411435128 (2012).

MacIver, M. A., Schmitz, L., Mugan, U., Murphey, T. D. & Mobley, C. D. Massive increase in visual range preceded the origin of terrestrial vertebrates. *Proc Natl Acad Sci U S A* **114**, E2375-E2384, doi:10.1073/pnas.1615563114 (2017).

Mugan, U. & MacIver, M. A. Spatial planning with long visual range benefits escape from visual predators in complex naturalistic environments. *Nat Commun* **11**, 3057, doi:10.1038/s41467-020-16102-1 (2020).

Hunt, L. *et al.* Formalizing planning and information search in naturalistic decision-making. *Nature neuroscience* **24**, 1051-1064 (2021).

Johnson, A. & Redish, A. D. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J Neurosci* **27**, 12176-12189, doi:10.1523/JNEUROSCI.3761-07.2007 (2007).

Ito, H. T., Zhang, S.-J., Witter, M. P., Moser, E. I. & Moser, M.-B. A prefrontal–thalamo–hippocampal circuit for goal-directed spatial navigation. *Nature* **522**, 50-55 (2015).

Pastalkova, E., Itskov, V., Amarasingham, A. & Buzsaki, G. Internally generated cell assembly sequences in the rat hippocampus. *Science* **321**, 1322-1327, doi:10.1126/science.1159775 (2008).

Redish, A. D. Vicarious trial and error. *Nat Rev Neurosci* **17**, 147-159, doi:10.1038/nrn.2015.30 (2016).

Balleine, B. W. & O'Doherty, J. P. Human and rodent homologies in action control: corticostriatal determinants of goal-directed and habitual action. *Neuropsychopharmacology* **35**, 48-69, doi:10.1038/npp.2009.131 (2010).

Graybiel, A. M. Habits, rituals, and the evaluative brain. *Annu Rev Neurosci* **31**, 359-387, doi:10.1146/annurev.neuro.29.051605.112851 (2008).

Balleine, B. W., Delgado, M. R. & Hikosaka, O. The role of the dorsal striatum in reward and decision-making. *J Neurosci* **27**, 8161-8165, doi:10.1523/JNEUROSCI.1554-07.2007 (2007).

Gupta, A. S., van der Meer, M. A., Touretzky, D. S. & Redish, A. D. Hippocampal replay is not a simple function of experience. *Neuron* **65**, 695-705, doi:10.1016/j.neuron.2010.01.034 (2010).

Pfeiffer, B. E. & Foster, D. J. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* **497**, 74-79, doi:10.1038/nature12112 (2013).

Jog, M. S., Kubota, Y., Connolly, C. I., Hillegaart, V. & Graybiel, A. M. Building neural representations of habits. *Science* **286**, 1745-1749, doi:10.1126/science.286.5445.1745 (1999).

Yin, H. H., Knowlton, B. J. & Balleine, B. W. Inactivation of dorsolateral striatum enhances sensitivity to changes in the action-outcome contingency in instrumental conditioning. *Behav Brain Res* **166**, 189-196, doi:10.1016/j.bbr.2005.07.012 (2006).

Eichenbaum, H. *Memory, amnesia, and the hippocampal system*. (MIT press, 1993).

Knierim, J. J., Kudrimoti, H. S. & McNaughton, B. L. Place cells, head direction cells, and the learning of landmark stability. *J Neurosci* **15**, 1648-1659, doi:10.1523/JNEUROSCI.15-03-01648.1995 (1995).

Lee, J. Q., LeDuke, D. O., Chua, K., McDonald, R. J. & Sutherland, R. J. Relocating cued goals induces population remapping in CA1 related to memory performance in a two-platform water task in rats. *Hippocampus* **28**, 431-440, doi:10.1002/hipo.22843 (2018).

Bradfield, L. A., Leung, B. K., Boldt, S., Liang, S. & Balleine, B. W. Goal-directed actions transiently depend on dorsal hippocampus. *Nature Neuroscience* **23**, 1194-1197, doi:10.1038/s41593-020-0693-8 (2020).

Barnes, T. D., Kubota, Y., Hu, D., Jin, D. Z. & Graybiel, A. M. Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature* **437**, 1158-1161, doi:10.1038/nature04053 (2005).

Graybiel, A. M. The basal ganglia and chunking of action repertoires. *Neurobiol Learn Mem* **70**, 119-136, doi:10.1006/nlme.1998.3843 (1998).

Smith, K. S. & Graybiel, A. M. A dual operator view of habitual behavior reflecting cortical and striatal dynamics. *Neuron* **79**, 361-374, doi:10.1016/j.neuron.2013.05.038 (2013).

Thorn, C. A., Atallah, H., Howe, M. & Graybiel, A. M. Differential dynamics of activity changes in dorsolateral and dorsomedial striatal loops during learning. *Neuron* **66**, 781-795, doi:10.1016/j.neuron.2010.04.036 (2010).

Jin, X. & Costa, R. M. Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature* **466**, 457-462 (2010).

Jin, X., Tecuapetla, F. & Costa, R. M. Basal ganglia subcircuits distinctively encode the parsing and concatenation of action sequences. *Nat Neurosci* **17**, 423-430, doi:10.1038/nn.3632 (2014).

Yin, H. H. & Knowlton, B. J. The role of the basal ganglia in habit formation. *Nat Rev Neurosci* **7**, 464-476, doi:10.1038/nrn1919 (2006).

Schmitzer-Torbert, N. C. & Redish, A. D. Task-dependent encoding of space and events by striatal neurons is dependent on neural subtype. *Neuroscience* **153**, 349-360, doi:10.1016/j.neuroscience.2008.01.081 (2008).

Gahnstrom, C. J. & Spiers, H. J. Striatal and hippocampal contributions to flexible navigation in rats and humans. *Brain Neurosci Adv* **4**, 2398212820979772, doi:10.1177/2398212820979772 (2020).

Regier, P. S., Amemiya, S. & Redish, A. D. Hippocampus and subregions of the dorsal striatum respond differently to a behavioral strategy change on a spatial navigation task. *J Neurophysiol* **114**, 1399-1416, doi:10.1152/jn.00189.2015 (2015).

Powell, N. J. & Redish, A. D. Representational changes of latent strategies in rat medial prefrontal cortex precede changes in behaviour. *Nat Commun* **7**, 12830, doi:10.1038/ncomms12830 (2016).

Hasz, B. M. & Redish, A. D. Dorsomedial prefrontal cortex and hippocampus represent strategic context even while simultaneously changing representation throughout a task session. *Neurobiology of Learning and Memory* **171**, 107215, doi:https://doi.org/10.1016/j.nlm.2020.107215 (2020).

Durstewitz, D., Vittoz, N. M., Floresco, S. B. & Seamans, J. K. Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron* **66**, 438-448, doi:10.1016/j.neuron.2010.03.029 (2010).

Euston, D. R., Gruber, A. J. & McNaughton, B. L. The role of medial prefrontal cortex in memory and decision making. *Neuron* **76**, 1057-1070, doi:10.1016/j.neuron.2012.12.002 (2012).

McLaughlin, A. E., Diehl, G. W. & Redish, A. D. Potential roles of the rodent medial prefrontal cortex in conflict resolution between multiple decision-making systems. *Int Rev Neurobiol* **158**, 249-281, doi:10.1016/bs.irn.2020.11.009 (2021).

Wirt, R. A. & Hyman, J. M. Integrating Spatial Working Memory and Remote Memory: Interactions between the Medial Prefrontal Cortex and Hippocampus. *Brain Sci* **7**, doi:10.3390/brainsci7040043 (2017).

Kidder, K., Gillis, R., Miles, J. & Mizumori, S. J. Y. The medial prefrontal cortex during flexible decisions: Evidence for its role in distinct working memory processes. *Hippocampus* **34**, 141-155, doi:10.1002/hipo.23594 (2024).

Hasz, B. M. & Redish, A. D. Spatial encoding in dorsomedial prefrontal cortex and hippocampus is related during deliberation. *Hippocampus* **30**, 1194-1208, doi:10.1002/hipo.23250 (2020).

O'Keefe, J. & Dostrovsky, J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res* **34**, 171-175, doi:10.1016/0006-8993(71)90358-1 (1971).

Morris, R. G., Garrud, P., Rawlins, J. N. & O'Keefe, J. Place navigation impaired in rats with hippocampal lesions. *Nature* **297**, 681-683, doi:10.1038/297681a0 (1982).

Fortin, N. J., Agster, K. L. & Eichenbaum, H. B. Critical role of the hippocampus in memory for sequences of events. *Nat Neurosci* **5**, 458-462, doi:10.1038/nn834 (2002).

Behrens, T. E. J. *et al.* What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron* **100**, 490-509, doi:10.1016/j.neuron.2018.10.002 (2018).

Aronov, D., Nevers, R. & Tank, D. W. Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit. *Nature* **543**, 719-722, doi:10.1038/nature21692 (2017).

Cohen, N. J. & Eichenbaum, H. *Memory, amnesia, and the hippocampal system*. (MIT Press, 1993).

Leutgeb, S. *et al.* Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* **309**, 619-623, doi:10.1126/science.1114037 (2005).

Lever, C., Wills, T., Cacucci, F., Burgess, N. & O'Keefe, J. Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature* **416**, 90-94, doi:10.1038/416090a (2002).

Jeffery, K. J. & Anderson, M. I. Dissociation of the geometric and contextual influences on place cells. *Hippocampus* **13**, 868-872, doi:10.1002/hipo.10162 (2003).

Fenton, A. A. *et al.* Unmasking the CA1 ensemble place code by exposures to small and large environments: more place cells and multiple, irregularly arranged, and expanded place fields in the larger space. *J Neurosci* **28**, 11250-11262, doi:10.1523/JNEUROSCI.2862-08.2008 (2008).

Duvelle, E. *et al.* Hippocampal place cells encode global location but not connectivity in a complex space. *Curr Biol* **31**, 1221-1233 e1229, doi:10.1016/j.cub.2021.01.005 (2021).

Kay, K. *et al.* Constant Sub-second Cycling between Representations of Possible Futures in the Hippocampus. *Cell* **180**, 552-567 e525, doi:10.1016/j.cell.2020.01.014 (2020).

Foster, D. J. & Wilson, M. A. Hippocampal theta sequences. *Hippocampus* **17**, 1093-1099, doi:10.1002/hipo.20345 (2007).

Lisman, J. & Redish, A. D. Prediction, sequences and the hippocampus. *Philosophical Transactions of the Royal Society B: Biological Sciences* **364**, 1193-1201 (2009).

Dragoi, G. & Buzsaki, G. Temporal encoding of place sequences by hippocampal cell assemblies. *Neuron* **50**, 145-157, doi:10.1016/j.neuron.2006.02.023 (2006).

Gupta, A. S., van der Meer, M. A., Touretzky, D. S. & Redish, A. D. Segmentation of spatial experience by hippocampal theta sequences. *Nat Neurosci* **15**, 1032-1039, doi:10.1038/nn.3138 (2012).

Ambrose, R. E., Pfeiffer, B. E. & Foster, D. J. Reverse Replay of Hippocampal Place Cells Is Uniquely Modulated by Changing Reward. *Neuron* **91**, 1124-1136, doi:10.1016/j.neuron.2016.07.047 (2016).

Wu, X. & Foster, D. J. Hippocampal replay captures the unique topological structure of a novel environment. *J Neurosci* **34**, 6459-6469, doi:10.1523/JNEUROSCI.3414-13.2014 (2014).

Buzsaki, G. Theta oscillations in the hippocampus. *Neuron* **33**, 325-340, doi:10.1016/s0896-6273(02)00586-x (2002).

Dezfouli, A., Lingawi, N. W. & Balleine, B. W. Habits as action sequences: hierarchical action control and changes in outcome value. *Philos Trans R Soc Lond B Biol Sci* **369**, doi:10.1098/rstb.2013.0482 (2014).

Cunningham, P. J., Regier, P. S. & Redish, A. D. Dorsolateral Striatal Task-initiation Bursts Represent Past Experiences More than Future Action Plans. *J Neurosci* **41**, 8051-8064, doi:10.1523/JNEUROSCI.3080-20.2021 (2021).

Amemiya, S. & Redish, A. D. Hippocampal Theta-Gamma Coupling Reflects State-Dependent Information Processing in Decision Making. *Cell Rep* **22**, 3328-3338, doi:10.1016/j.celrep.2018.02.091 (2018).

Amemiya, S. & Redish, A. D. Manipulating Decisiveness in Decision Making: Effects of Clonidine on Hippocampal Search Strategies. *J Neurosci* **36**, 814-827, doi:10.1523/JNEUROSCI.2595-15.2016 (2016).

Hasz, B. M. & Redish, A. D. Dorsomedial prefrontal cortex and hippocampus represent strategic context even while simultaneously changing representation throughout a task session. *Neurobiol Learn Mem* **171**, 107215, doi:10.1016/j.nlm.2020.107215 (2020).

Powell, N. J. & Redish, A. D. Complex neural codes in rat prelimbic cortex are stable across days on a spatial decision task. *Front Behav Neurosci* **8**, 120, doi:10.3389/fnbeh.2014.00120 (2014).

Steiner, A. P. & Redish, A. D. The road not taken: neural correlates of decision making in orbitofrontal cortex. *Front Neurosci* **6**, 131, doi:10.3389/fnins.2012.00131 (2012).

van der Meer, M. A., Johnson, A., Schmitzer-Torbert, N. C. & Redish, A. D. Triple dissociation of information processing in dorsal striatum, ventral striatum, and hippocampus on a learned spatial decision task. *Neuron* **67**, 25-32, doi:10.1016/j.neuron.2010.06.023 (2010).

Blumenthal, A., Steiner, A., Seeland, K. & Redish, A. D. Effects of pharmacological manipulations of NMDA-receptors on deliberation in the Multiple-T task. *Neurobiol Learn Mem* **95**, 376-384, doi:10.1016/j.nlm.2011.01.011 (2011).

Schmitzer-Torbert, N. & Redish, A. D. Neuronal activity in the rodent dorsal striatum in sequential navigation: separation of spatial and reward responses on the multiple T task. *J Neurophysiol* **91**, 2259-2272, doi:10.1152/jn.00687.2003 (2004).

Schmitzer-Torbert, N. & Redish, A. D. Development of path stereotypy in a single day in rats on a multiple-T maze. *Arch Ital Biol* **140**, 295-301 (2002).

Papale, A. E., Stott, J. J., Powell, N. J., Regier, P. S. & Redish, A. D. Interactions between deliberation and delay-discounting in rats. *Cogn Affect Behav Neurosci* **12**, 513-526, doi:10.3758/s13415-012-0097-7 (2012).

George, A. E., Stout, J. J. & Griffin, A. L. Pausing and reorienting behaviors enhance the performance of a spatial working memory task. *Behav Brain Res* **446**, 114410, doi:10.1016/j.bbr.2023.114410 (2023).

Mugan, U., Amemiya, S., Regier, P. S. & Redish, A. D. Navigation through the complex world--the neurophysiology of decision-making processes. *arXiv preprint arXiv:2306.03162* (2023).

Trucco, E. A note on the information content of graphs. *The bulletin of mathematical biophysics* **18**, 129-135 (1956).

Newman, M. E. J., Barabási, A.-L. s. & Watts, D. J. *The structure and dynamics of networks*. (Princeton University Press, 2006).

Dehmer, M. & Mowshowitz, A. A history of graph entropy measures. *Information Sciences* **181**, 57-78 (2011).

Bonchev, D. G. Information theoretic complexity measures. *Encyclopedia of Complexity and Systems Science* **5**, 4820-4838 (2009).

Wikenheiser, A. M. & Redish, A. D. Hippocampal theta sequences reflect current goals. *Nat Neurosci* **18**, 289-294, doi:10.1038/nn.3909 (2015).

Buzsaki, G. Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus* **25**, 1073-1188, doi:10.1002/hipo.22488 (2015).

Widloski, J. & Foster, D. J. Flexible rerouting of hippocampal replay sequences around changing barriers in the absence of global place field remapping. *Neuron* **110**, 1547-1558 e1548, doi:10.1016/j.neuron.2022.02.002 (2022).

Olafsdottir, H. F., Barry, C., Saleem, A. B., Hassabis, D. & Spiers, H. J. Hippocampal place cells construct reward related sequences through unexplored space. *Elife* **4**, e06063, doi:10.7554/eLife.06063 (2015).

Zielinski, M. C., Shin, J. D. & Jadhav, S. P. Coherent Coding of Spatial Position Mediated by Theta Oscillations in the Hippocampus and Prefrontal Cortex. *J Neurosci* **39**, 4550-4565, doi:10.1523/JNEUROSCI.0106-19.2019 (2019).

Mizuseki, K., Sirota, A., Pastalkova, E. & Buzsaki, G. Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron* **64**, 267-280, doi:10.1016/j.neuron.2009.08.037 (2009).

Maurer, A. P., Cowen, S. L., Burke, S. N., Barnes, C. A. & McNaughton, B. L. Organization of hippocampal cell assemblies based on theta phase precession. *Hippocampus* **16**, 785-794, doi:10.1002/hipo.20202 (2006).

Maurer, A. P. & McNaughton, B. L. Network and intrinsic cellular mechanisms underlying theta phase precession of hippocampal neurons. *Trends Neurosci* **30**, 325-333, doi:10.1016/j.tins.2007.05.002 (2007).

Skaggs, W. E., McNaughton, B. L., Wilson, M. A. & Barnes, C. A. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* **6**, 149-172, doi:10.1002/(SICI)1098-1063(1996)6:2<149::AID-HIPO6>3.0.CO;2-K (1996).

O'Keefe, J. & Recce, M. L. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* **3**, 317-330, doi:10.1002/hipo.450030307 (1993).

Zhang, K., Ginzburg, I., McNaughton, B. L. & Sejnowski, T. J. Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *J Neurophysiol* **79**, 1017-1044, doi:10.1152/jn.1998.79.2.1017 (1998).

Farooq, U. & Dragoi, G. Emergence of preconfigured and plastic time-compressed sequences in early postnatal development. *Science* **363**, 168-173, doi:10.1126/science.aav0502 (2019).

Feng, T., Silva, D. & Foster, D. J. Dissociation between the experience-dependent development of hippocampal theta sequences and single-trial phase precession. *J Neurosci* **35**, 4890-4902, doi:10.1523/JNEUROSCI.2614-14.2015 (2015).

Drieu, C., Todorova, R. & Zugaro, M. Nested sequences of hippocampal assemblies during behavior support subsequent sleep replay. *Science* **362**, 675-679, doi:10.1126/science.aat2952 (2018).

Schmidt, B., Duin, A. A. & Redish, A. D. Disrupting the medial prefrontal cortex alters hippocampal sequences during deliberative decision making. *J Neurophysiol* **121**, 1981-2000, doi:10.1152/jn.00793.2018 (2019).

Zheng, C., Hwaun, E., Loza, C. A. & Colgin, L. L. Hippocampal place cell sequences differ during correct and error trials in a spatial memory task. *Nat Commun* **12**, 3373, doi:10.1038/s41467-021-23765-x (2021).

Wang, M., Foster, D. J. & Pfeiffer, B. E. Alternating sequences of future and past behavior encoded within hippocampal theta oscillations. *Science* **370**, 247-250 (2020).

Smith, K. S. & Graybiel, A. M. Habit formation coincides with shifts in reinforcement representations in the sensorimotor striatum. *J Neurophysiol* **115**, 1487-1498, doi:10.1152/jn.00925.2015 (2016).

Hok, V., Save, E., Lenck-Santini, P. P. & Poucet, B. Coding for spatial goals in the prelimbic/infralimbic area of the rat frontal cortex. *Proc Natl Acad Sci U S A* **102**, 4602-4607, doi:10.1073/pnas.0407332102 (2005).

Botvinick, M. M. Hierarchical models of behavior and prefrontal function. *Trends Cogn Sci* **12**, 201-208, doi:10.1016/j.tics.2008.02.009 (2008).

Belluscio, M. A., Mizuseki, K., Schmidt, R., Kempter, R. & Buzsaki, G. Cross-frequency phase-phase coupling between theta and gamma oscillations in the hippocampus. *J Neurosci* **32**, 423-435, doi:10.1523/JNEUROSCI.4122-11.2012 (2012).

Zielinski, M. C., Tang, W. & Jadhav, S. P. The role of replay and theta sequences in mediating hippocampal-prefrontal interactions for memory and cognition. *Hippocampus* **30**, 60-72, doi:10.1002/hipo.22821 (2020).

Roth, B. L. DREADDs for neuroscientists. *Neuron* **89**, 683-694 (2016).

Nagai, Y. *et al.* Deschloroclozapine, a potent and selective chemogenetic actuator enables rapid neuronal and behavioral modulations in mice and monkeys. *Nat Neurosci* **23**, 1157-1167, doi:10.1038/s41593-020-0661-3 (2020).

McLaughlin, A. E. & Redish, A. D. Optogenetic disruption of the prelimbic cortex alters long-term decision strategy but not valuation on a spatial delay discounting task. *Neurobiol Learn Mem* **200**, 107734, doi:10.1016/j.nlm.2023.107734 (2023).

Hull, C. L. Principles of behavior: an introduction to behavior theory. (1943).

Botvinick, M. M., Niv, Y. & Barto, A. G. Hierarchically organized behavior and its neural foundations: a reinforcement learning perspective. *Cognition* **113**, 262-280, doi:10.1016/j.cognition.2008.08.011 (2009).

Niv, Y. Reinforcement learning in the brain. *Journal of Mathematical Psychology* **53**, 139-154, doi:10.1016/j.jmp.2008.12.005 (2009).

Averbeck, B. & O'Doherty, J. P. Reinforcement-learning in fronto-striatal circuits. *Neuropsychopharmacology* **47**, 147-162 (2022).

Buckner, R. L. & Carroll, D. C. Self-projection and the brain. *Trends Cogn Sci* **11**, 49-57, doi:10.1016/j.tics.2006.11.004 (2007).

Tolman, E. C. Cognitive maps in rats and men. *Psychol Rev* **55**, 189-208, doi:10.1037/h0061626 (1948).

Gardner, R. S. *et al.* A secondary working memory challenge preserves primary place strategies despite overtraining. *Learn Mem* **20**, 648-656, doi:10.1101/lm.031336.113 (2013).

Packard, M. G. & McGaugh, J. L. Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiol Learn Mem* **65**, 65-72, doi:10.1006/nlme.1996.0007 (1996).

Wikenheiser, A. M. & Redish, A. D. Decoding the cognitive map: ensemble hippocampal sequences and decision making. *Curr Opin Neurobiol* **32**, 8-15, doi:10.1016/j.conb.2014.10.002 (2015).

Miller, K. J., Botvinick, M. M. & Brody, C. D. Dorsal hippocampus contributes to model-based planning. *Nat Neurosci* **20**, 1269-1276, doi:10.1038/nn.4613 (2017).

Gottlieb, J. & Oudeyer, P. Y. Towards a neuroscience of active sampling and curiosity. *Nat Rev Neurosci* **19**, 758-770, doi:10.1038/s41583-018-0078-0 (2018).

Ormond, J. & O'Keefe, J. Hippocampal place cells have goal-oriented vector fields during navigation. *Nature* **607**, 741-746, doi:10.1038/s41586-022-04913-9 (2022).

Berylne, D. E. Curiosity and Exploration. *Science* **153**, 25-33, doi:DOI: 10.1126/science.153.3731.25 (1966).

Kidd, C. & Hayden, B. Y. The Psychology and Neuroscience of Curiosity. *Neuron* **88**, 449-460, doi:10.1016/j.neuron.2015.09.010 (2015).

Gruber, M. J. & Ranganath, C. How Curiosity Enhances Hippocampus-Dependent Memory: The Prediction, Appraisal, Curiosity, and Exploration (PACE) Framework. *Trends Cogn Sci* **23**, 1014-1025, doi:10.1016/j.tics.2019.10.003 (2019).

Hunnicutt, B. J. *et al.* A comprehensive excitatory input map of the striatum reveals novel functional organization. *Elife* **5**, doi:10.7554/eLife.19103 (2016).

Wolff, S. B. E., Ko, R. & Olveczky, B. P. Distinct roles for motor cortical and thalamic inputs to striatum during motor skill learning and execution. *Sci Adv* **8**, eabk0231, doi:10.1126/sciadv.abk0231 (2022).

Alexander, G. E., DeLong, M. R. & Strick, P. L. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu Rev Neurosci* **9**, 357-381, doi:10.1146/annurev.ne.09.030186.002041 (1986).

Tang, W. & Jadhav, S. P. Multiple-Timescale Representations of Space: Linking Memory to Navigation. *Annu Rev Neurosci*, doi:10.1146/annurev-neuro-111020-084824 (2021).

Diehl, G. W. & Redish, A. D. Differential processing of decision information in subregions of rodent medial prefrontal cortex. *Elife* **12**, doi:10.7554/eLife.82833 (2023).

Benchenane, K. *et al.* Coherent theta oscillations and reorganization of spike timing in the hippocampal- prefrontal network upon learning. *Neuron* **66**, 921-936, doi:10.1016/j.neuron.2010.05.013 (2010).

Colgin, L. L. Oscillations and hippocampal-prefrontal synchrony. *Curr Opin Neurobiol* **21**, 467-474, doi:10.1016/j.conb.2011.04.006 (2011).

Spellman, T. *et al.* Hippocampal-prefrontal input supports spatial encoding in working memory. *Nature* **522**, 309-314, doi:10.1038/nature14445 (2015).

Stout, J. J., Hallock, H. L., George, A. E., Adiraju, S. S. & Griffin, A. L. The ventral midline thalamus coordinates prefrontal-hippocampal neural synchrony during vicarious trial and error. *Sci Rep* **12**, 10940, doi:10.1038/s41598-022-14707-8 (2022).

Guise, K. G. & Shapiro, M. L. Medial Prefrontal Cortex Reduces Memory Interference by Modifying Hippocampal Encoding. *Neuron* **94**, 183-192 e188, doi:10.1016/j.neuron.2017.03.011 (2017).

Rich, E. L. & Shapiro, M. Rat prefrontal cortical neurons selectively code strategy switches. *Journal of Neuroscience* **29**, 7208-7219 (2009).

Rich, E. L. & Shapiro, M. L. Prelimbic/infralimbic inactivation impairs memory for multiple task switches, but not flexible selection of familiar tasks. *J Neurosci* **27**, 4747-4755, doi:10.1523/JNEUROSCI.0369-07.2007 (2007).

Boorman, E. D., Behrens, T. E., Woolrich, M. W. & Rushworth, M. F. How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. *Neuron* **62**, 733-743, doi:10.1016/j.neuron.2009.05.014 (2009).

Procyk, E., Tanaka, Y. L. & Joseph, J. P. Anterior cingulate activity during routine and non-routine sequential behaviors in macaques. *Nat Neurosci* **3**, 502-508, doi:10.1038/74880 (2000).

Methods

Subjects

Eight adult Fisher-Brown Norway F1 hybrid rats (FBNF-1) (5M, 3F) and one adult Brown Norway (BN) female were used in the recording study (n = 9 rats). FBNF-1 rats were bred in-house, and the BN rat was purchased from Envigo. All rats were aged between 6-11 months old during experiments.

Ten adult BN rats (5M, 5F) and six FBNF-1 (3M, 3F) were used in the manipulation (DREADDs) study (n = 16 rats). All rats were aged between 7-15 months old during experiments. While two strains of rats were used in the reported study, we found no evidence for differences in behavior, physiology, or response to manipulations between FBNF-1 and BN rats, so data are combined.

Rats were single housed in a temperature-controlled colony room with a 12 hr light/dark cycle (lights on at 8:00 AM). Throughout the experiments the rats were food restricted and had to maintain a weight of at least 80% of their free feeding weight. Rats were given full access to food one-week prior to silicon probe implantation or DREADDs (designer receptors activated by designer drugs) viral infusion. Surgeries were conducted when the rats were at least at 100% of their prior free feeding weight. All training and experimental sessions were conducted at around the same time and during the rats' light phase. All procedures were approved by the University of Minnesota Institutional Animal Care and Use Committee (IACUC) and were in accordance with the NIH guidelines.

Surgical procedures

Of the 25 rats used in the study, 9 rats were used for *in vivo* electrophysiology recordings, and 16 rats underwent viral infusions. Of the 16 virus rats, 8 rats were bilaterally injected with DREADDs virus AAV5-CaMKIIα-hM4Di-mCherry (titer, 2.3×10^{13}; Addgene http://addgene.org/50477) and 8 rats were bilaterally injected with AAV5-CaMKIIα-mCherry control virus (titer, 2.3×10^{13}; Addgene http://addgene.org/114469).

For the recording study (n = 9), rats were anesthetized with isoflurane (1-2% in O_2) and mounted in a stereotaxic frame (Kopf, Tujunga CA). Prior to surgery Carprofen (5 mg/kg, SC) and sterile saline (3 mL, SC) were administered for analgesia and hydration, respectively. A pair of ground screws were implanted over cerebellum and a pair of anchor screws were implanted to ensure proper bonding to the 3D printed ring that was attached to the skull with C&B Metabond adhesive. Craniotomies were drilled over either HC and DLS contralaterally, HC and mPFC ipsilaterally, or DLS and mPFC ipsilaterally (Male coordinates HC: ± 3 ML -4 AP; DLS: ± 3.3 ML 0.7 AP; mPFC: ± 0.7 ML 2.8 AP; Female coordinates were adjusted by 93%). One 64-channel silicon probe was chronically implanted targeting each region, for a total of two probes in each rat. HC and DLS were always implanted with four shank 64-channel Cambridge Neurotech P-1 probes, while four mPFC rats were implanted with linear 64-channel Cambridge Neurotech H3 probes, one mPFC rat was implanted with two shank 64-channel Cambridge Neurotech H6 probe (R786 – simultaneous recordings from mPFC & HC), and one mPFC rat was implanted with four shank 64-channel Cambridge Neurotech P-1 probe (R787 – simultaneous recordings from mPFC & HC). Targeting for the left and right hemisphere was counterbalanced across rats accounting only for sex (Supplementary Fig. 1C, D, Supplementary Table 1).

The probes were mounted on custom-made 3D printed micro-drives for precise and independent adjustment of the vertical position. The probes were lowered to ~1.5mm and ~2.5 mm from skull surface

for P-1 and H3/H6 probes respectively to ensure all recording sites were in the brain. The craniotomies were then filled with a sterile mixture of bone wax and mineral oil (3/1 by weight) and the micro-drives were adhered to the skull with Metabond. A 3D printed shell that was painted with conductive paint, head cap, and LED mount were attached to the skull ring, and the ground screw was attached to the inside of the shell via conductive epoxy.

For viral vector injections, rats went through the same pre-surgery protocol. Craniotomies were made bilaterally over mPFC (±0.7 ML 2.8 AP), and bilateral cannulae were lowered into depth -3.8 DV for females and -4 DV for males (Supplementary Fig. 6A). Rats were randomly assigned to receive either the active virus (CaMKIIα-hM4Di(Gi)-mCherry; n = 8) or control virus (CaMKIIα-mCherry; n = 8). 1μL of virus was delivered at a rate of 200 nL/min. The cannulae were slowly raised following a 15-minute wait period after infusion. The craniotomies were then filled with sterile bonewax. Finally, an LED attachment was adhered to the skull using C&B Metabond to facilitate video tracking during behavior.

At the completion of the surgery rats were given Baytril (10 mg/kg, SC), additional saline (20 ml/kg, SC), and Children's Tylenol (0.8 mL, orally) for antibiotic, hydration, and analgesia, respectively.

Left/Right/Alternation Task

The left/right/alternation task has been detailed in prior works.[1-5] Briefly, in the left/right/alternation task (LRA), rats started from the start of maze (SOM) region at the bottom of the track, and ran through a navigation sequence, that consisted of three low-cost choice points (Wall 1, Wall 2, Wall 3) to a high-cost choice point (CP) at the top of the track (Supplementary Fig. 1A). At the high-cost choice point the rats had to choose whether to proceed to the left or to the right, where each return track consisted of one feeder that automatically delivered one food pellet (45 mg, Test Diet, Richmond IN) if the choice was correct. The rats would then proceed back to the SOM where one food pellet would be delivered if the current choice was correct. We refer to the region between SOM and CP as the central track. The walls of the maze were constructed using DUPLO LEGO bricks with vinyl flooring under the maze. Permanent walls had unique patterns but did not change throughout the experiment. The interior walls of the central path could change daily, but they remained constant within the day. For the recordings experiment, the maze was 108 cm × 163 cm with 25 cm wide tracks, and for the DREADDs virus study the maze was 94 cm × 100 cm with 13 cm wide tracks.

There were three different reward criteria that specified the correct choice on each lap, left reward (L) where food was delivered on the left return rail and SOM, right reward (R) where food was delivered on the right return rail and SOM, and alternation (A) where the rats had to alternate between the left and right return rails to receive reward at the chosen side and SOM. For the left and right reward criteria a choice was considered correct if on each lap rats ran to the left and right sides of the track, respectively. For alternation a choice was considered correct if on each lap rats alternated which side of the track they went (e.g., if prior choice was left the next choice had to be right and vice versa). If the starting rule was alternation, the rats always received a reward (both at side feeder and SOM) on their first lap of the session.

At the beginning of the session rats were placed on the track at the SOM, and then were allowed to run continuously for 45 minutes. Rats were pre-trained to run in a loop. Each correct lap resulted in the automated delivery of two reward pellets, one at the chosen side and one at SOM, thus with each

correct lap, rats earned 90 mg of food. All pellets were unflavored and performance on the task constituted the sole opportunity for rats to earn food each day.

All behavioral events were operated through custom built Matlab software (Matlab v2015b, Mathworks) to track rat position, control task state, and trigger reward delivery. Video was collected via an overhead camera sampled at 30 Hz, tracking either the backpack LED (pre-op) or head-mounted LED (post-op). Food delivery was via a Matlab-Arduino interface (Arduino UNO) sending TTL pulses to Med-Associate feeders (Fairfax VT).

Initial task training

Prior to training all rats were handled by the experimenter for 3 to 7 days until they were acclimated. The rats were then food restricted for 5 to 7 days and were offered 30 min of access to unflavored pellets (45 mg/pellet, TestDiet, Richmond IN). After the rats were familiar and ate the available food, they were placed back on free feeding for 7 days to return them to their original free-feeding weight. Training started after food acclimation and weight rebound. The training took place before probe implantation for the recording study, and after the virus surgery for the manipulation study to ensure viral expression prior to the experimental phase. Behavioral tracking was done using backpacks with an LED for the recording rats, and a head-mounted LED for the virus rats. Rats were trained over three phases that ranged from 7 to 12 days. During the first two phases there were no walls (low-cost choice points) within the central track, i.e., it was an open arena.

In the first training phase, rats were trained to run in a loop starting from the SOM and ending at SOM. L and R rules were alternated each day. The incorrect (opposite to the session rule) arm was blocked such that rats were forced to turn congruent with the rule, and thus make the correct choice (e.g., if the reward contingency was L, the right arm was blocked). During this initial phase each feeder (side feeder and SOM) delivered to 2 unflavored pellets/feeder (total of 4 pellets/lap, 180mg of food). Training continued until the rats were running continuously and were completing at least 50 laps/session for each reward contingency (left/right).

In the second phase of training, blocks were removed, and alternation was introduced as a rule. During the rest of the training days rats experienced one pseudo-randomly chosen reward criterion per day. Rats remained in this phase until they reached 80% correct for each reward criterion (left, right, or alternation) and ran ~100 laps continuously (10-15 days).

In the third phase of training, interior walls were introduced. All possible maze configurations within a given complexity (Fig. 1A, J) were randomly assigned to either be part of the training or the test set. Each day, a maze configuration from the training set was pseudo-randomly chosen without replacement, such that everyday had a unique maze configuration. LEGO bricks were used to create the maze structure. Similar to training phase two, animals experienced a pseudo-randomly chosen single reward criterion. Rats remained in this phase until they reached 80% correct for each reward criterion (left, right, alternation) and ran ~100 laps (5-7 days).

For the recording animals, surgery was conducted after training phase two. After the rats underwent probe implantation, they were reacclimated with the task starting from training phase two, during which the probes were lowered to their appropriate targets (7-12 days). After rats reached phase two criterion, they moved onto phase three, where along with the introduction of the maze walls, rats ran the task

with the recording tether attached, familiarizing them with the process of performing the task under full task and recording conditions.

Experimental sequence

Once rats completed their training sequence and the probes were positioned at their appropriate targets, the rats immediately entered the experimental sequence, where all three reward rules were presented to the animal within a single session. Each 45 min session had three un-cued rule changes that occurred roughly 15 and 30 minutes into the session. A random amount of time between ± 2.2 mins was added or subtracted from the target switch times to randomize their precise occurrence. The rule sequence for a given session was pseudo-randomly drawn from one of six possibilities: LRA, LAR, RLA, RAL, ALR, ARL and maze configurations were pseudo-randomly chosen from the test set.

Recording animals (n = 9 rats) underwent a 16-day experimental sequence (1 session/day). Recordings were performed over these 16 days, totaling 144 recording sessions. Before the task start and after the task end rats sat on a flowerpot for 5 min to record baseline neural activity during quiet resting. The pre/post session recording data were processed but were otherwise not included in analysis.

The 16-day sequence was further broken up into two experimental sequences spanning 8 days. Each 8-day sequence consisted of 4 unique mazes that were randomly chosen without replacement along the complexity axis, however, each 8-day sequence included the simplest maze (AAA). Rats were run on the same maze configuration under the same task rules for 2 consecutive days (Supplementary Fig. 1B).

The same experimental protocol was used for the DREADDs rats; however, the experimental sequence was extended to 32-days ($n_{Sessions}$ = 508), such that every maze configuration within the test set was used. Deschloroclozapine (DCZ; Hello Bio. Inc., Princeton NJ)[6] was used to activate the DREADDs, or VEH (saline) was used as a control (Supplementary Fig. 6B). The active and control virus animals were paired such that each pair had the same exact same experimental sequence (mazes and switches). DCZ and VEH conditions were presented in a pseudorandomized order, controlling for first-order effects. Importantly, the 2nd day-repeat of a given maze/rule configuration was always a saline washout day (Supplementary Fig. 6C). Every day the rats received subcutaneous (SC) injections of either DCZ, VEH, or saline 30 min prior to the start of the experiment. DCZ was dissolved in saline to a concentration of 0.05mg/mL and a single dosage of 0.1mg/kg of DCZ was used throughout the study. For VEH and saline a comparable volume was injected SC. All experimenters were blind to the identity of the solution (DCZ or VEH) injected on a given day.

Recordings

All recordings were collected at 30kHz sampling rate using a 256 channel Intan RHD recording system connected to two 64-channel Intan RHD headstages. An Arduino interface between Matlab and the Intan system facilitated time synchronization of neural data to ongoing behavior. In DLS and mPFC, throughout the 16-day recording sequence probes were periodically lowered (~30μm intervals) to increase the likelihood of recording a diverse array of units. All movement of probes occurred after completion of day's recordings leaving at least 23 hrs for probes to restabilize before the next recording session. After reaching the CA1 layer in dorsal HC during the training phase, the probes were not subsequently moved, unless there was visible spontaneous movement of the probe from one day to the next. In those cases, HC probes were either lowered or raised to ensure recordings from the CA1 cell layer. While we took appropriate efforts to record diverse ensemble et of units across days, we are unable to claim that all

recorded cells are unique. No attempts were made to identify recordings of the same cells across successive days.

Spikes were extracted and classified using Kilosort v1.0 for HC, and Kilosort v2.0 for DLS and mPFC using a custom pipeline. Kilosort v1.0 clustering for HC used previously published parameters tuned for HC (https://github.com/brendonw1/KilosortWrapper), and Kilosort v2.0 clustering for mPFC and DLS used default parameters.[7] Automated sorting was followed by manual curation of the waveform clusters using Phy (https://github.com/cortex-lab/phy).[8] Each silicon probe was analyzed independently and included the entire 45 min task data along with the 5 min pre/post recordings (~55 min total recording time). Prior to Kilosort processing, a median reference subtraction was performed across the 64 channels and a 600 Hz high-pass filter was used for spike identification. Collectively we identified 4292 single units from HC, 2293 single units from DLS, and 2609 single units from mPFC. All unit isolation and manual curation was performed blind to the session's corresponding behavioral data.

HC units were separated into putative pyramidal cells and interneurons using their waveform characteristics and firing rates. DLS units were separated as phasic and nonphasic based on the distribution of inter-spike-intervals (ISIs).[9] Phasic cells were those with at least 40% of all ISIs greater than 2 seconds, which we take as putative medium spiny neurons (MSNs). Otherwise, the DLS cell was classified as a nonphasic firing cell. Following previous studies, if a cell that had a post-spike suppression index longer than 100 ms it was classified as tonic firing (putative tonically active cholinergic interneurons), otherwise, it was classified as high firing (putative high firing interneurons). Nonphasic cells were not analyzed further here.

DREADD validation

In one additional rat (Male FBNF-1) we bilaterally transfused the dorsal mPFC (± 0.7 ML 2.8 AP -3 DV) with the same active inhibitory DREADDs used in the main study (AAV5-CaMKIIα-hM4Di(Gi)-mCherry) (Supplementary Fig. 6B histology). Roughly six weeks after the viral injection, a single 64-channel Cambridge Neurotech H3 probe was chronically implanted into the infusion craniotomy on the right hemisphere and was lowered to 2.3 mm from skull surface. These surgeries used the same procedures outlined above. Over four days the probe was lowered to 4.13 mm (2.86 mm top of probe) from skull surface to span the volume around the DREADD infusion site (Supplementary Fig. 6B histology).

We performed two recording sessions, one with a DCZ[6] injection and one with a VEH (saline) injection. The recordings were taken for ~3hrs during quiet rest of the rat on a flowerpot. We first recorded a 30 min baseline period before drug injection. At 30 mins the rat was injected with either DCZ or VEH (1 session/drug), while the recordings were ongoing. The rat was then placed back on to the pot for another 2.5 hrs to establish a time course of drug effect. Multiunit spiking activity (18 isolated units taken collectively for VEH injection Supplementary Fig. 6B left panel, and 5 units for DCZ injection Supplementary Fig. 6B right panel) was binned into 30s bins to calculate the average firing rate over the 3-hour experiment.

Histology

Following the conclusion of all experiments, rats were overdosed with pentobarbital. Rats were perfused transcardially with phosphate-buffered saline and paraformaldehyde. For recording rats, probes were left in place for 3-5 hrs and were then carefully removed from the brain, leaving a visible track of their final position (Fig. 1B, Supplementary Fig. 1D). Brains were then extracted and post fixed in paraformaldehyde

for 24 hrs before being transferred to a mixture of paraformaldehyde and sucrose for cryoprotection. Coronal sections (40μm sections) surrounding the probe location were cut on a cryostat, mounted to slides. For recording rats, mounted sections were stained with cresyl violet for visualization. In two rats with DLS recordings we performed immunohistochemical staining for MOR1 on a subset of the tissue (anti-μ-Opioid Receptor gifted by Maureen Riedl). For virus rats, we used immunohistochemical procedures to amplify the endogenous -mCherry fluorophore. Mounted tissue was labeled with an anti-RFP antibody to visualize viral expression and with an anti-NeuN antibody to counter stain cell bodies. In both cases fluorescent secondary antibodies were used for visualization.

Data Analysis

Most data analyses were performed using custom Matlab software (Matlab v2021a, Mathworks). Quantification of environment complexities were performed using the Networkx package for Python. Each subject was treated independently. No formal replication was performed, but our results agree with previously published work. For all analyses data were pooled across all subjects (physiology or behavioral responses), but in all cases we confirmed that reported findings were qualitatively consistent across rats and not driven by a single subject.

Quantification of maze complexity

Three different wall configurations were used to change the maze structure: *A*---two walls with open space in the middle, *B*---one wall with open space on the left, and *C*---one wall with open space on the right. Combinatorial combination with repetition of the three different wall types (A, B, C) created 27 unique mazes (3^3 = 27) (Fig. 1J).

Graph theoretic measures have been previously used to quantify the complexity of an environment.[10-14] To determine the structural complexity of a maze, we first discretized the inner navigation sequence into 15 bins. These bins corresponded to graph nodes, and each node was connected to another node if it could be reached by a cardinal action---i.e., an edge existed between the two nodes if there was no wall between the two bins. Starting from each node we created 15 generalized trees, with each root node being one of the 15 nodes of the discretized maze. The structural complexity of a tree graph is defined as the Shannon entropy of the distribution of the number of nodes at each level of the tree---e.g., the distribution of the number of nodes across all possible trees one step from the root node. The total structural complexity of the space was then defined as the total information content across all equivalent generalized trees that could be created from a given maze configuration.[10,11] The distribution of the total information across mazes was then min-max normalized. This distribution had 3 distinct clusters denoted by the colored circles in Fig. 1J that were labeled as low (light green), mid (teal), and high complexity (dark blue) environments, respectively.

Behavior

Each lap (SOM to SOM) was identified as correct or error dependent on whether the rat chose the action congruent with the current contingency/rule and was subsequently rewarded. Percent correct for a given session was calculated based on the ratio between the number of correct laps and the total number of laps.

The percent correct for a given lap was calculated based on the ratio between the number of sessions in which that lap was correct and the total number of sessions, analyzed for each rat separately. Given that

the sessions did not have equal number of laps and the switch times did not occur on the same lap, analyses were restricted to 10 laps pre- and 20 laps post-switch. Everything was computed for each switch independently and then averaged across the switches.

Behavioral change point

To identify laps where rats stabilized their behavior after rule switches, we used a modified version of the change point analysis developed by Gallistel et. al. (2004) [15]. Briefly, behavior in each lap was identified as a left or a right choice and was subsequently coded as either a 1 (left turn) or a 2 (right turn). Note that this differs slightly from prior implementations of this method which consider the correct/error nature of each choice as these methods can generate artificial change points when rats make an error on the first lap following a rule change (before they have any information about the change). This created a num. of Laps x 1 length vector. We then calculated the cumulative sum along this choice vector. Following[15], we identified behavioral change points as those where the slope of the cumulative sum changed most abruptly. Computationally, this was done using the *findchangepts* function in Matlab with linear change point detection (statistic = linear) and minimum improvement in total residual set to 2 (minthreshold = 2) so as to not limit the number of change points that could be identified. After detection, change points in every session were visually confirmed by an experimenter blind to all information about maze configuration, task rules, and drug conditions.

Running speed

We computed the running speed as the change in $< x, y >$ position $< \mathrm{d}x, \mathrm{d}y >$ using an adaptive windowing of best-fit velocity vectors.[16]

Vicarious trial and error (VTE)

To quantify vicarious trial and error (VTE) behaviors at the choice point, we calculated the time integrated change in angular movement of the head (IdPhi), which captures both pause duration and the prominence of head re-orientation behaviors (Fig. 1F, G).[17,18] IdPhi is calculated by the integral of the absolute value of the arctangent of the velocity over the duration that the rat is in the choice point. Velocity $(\mathrm{d}x, \mathrm{d}y)$ was calculated as the first derivative of the rat's head position $< x, y >$ using the Janabi-Sharifi and Marey (2010)[16] method.

$$\mathrm{IdPhi} = \int\limits_{t_{start}}^{t_{end}} \left| \frac{\delta}{\delta t} \tan^{-1}\left(\frac{\delta y}{\delta t}, \frac{\delta x}{\delta t} \right) \right| dt$$

zIdPhi was then computed as the z-scored IdPhi on a lap-by-lap basis for each session. In line with previous work, the distribution of the calculated zIdPhi created a positively skewed bimodal distribution, where the prominent distribution (larger proportion of the overall data) captured non-VTE events, and the smaller distribution captured VTE events. Based on this distribution, and similar to prior work, we identified a cutoff threshold between VTE or non-VTE laps as 0.25 zIdPhi (Supplementary Fig. 2G).[4,18]

Exploration

To calculate the amount of space explored, we binned the 2D space into ~2cm bins. We then calculated the per session aggregate occupancy in the central path to create a 2D histogram (Fig. 2B, Supplementary Fig. 2F). The percentage of space explored was quantified as the total number of counts divided by the number of bins.

To classify a lap as explorative vs. non-explorative, we first found the median path through the central track. For each lap the trajectory through the central maze segment was resampled to 500 positional points to create a [num. of laps × 500] matrix for each *x* and *y* position. The median path was taken to be the median $< x, y >$ location across laps for each interpolated sample. We then calculated the point-by-point distance between the resampled lap trajectory and the median path using the *dsearchn* function in Matlab. If more than 2.5% of the points in a given lap had distances that were above the 95% percentile of the overall distances across all laps, the lap in question was labeled as exploration.

Behavioral trajectory stereotypy

We calculated behavioral trajectory stereotypy from times that the rat was navigating through the central maze segment (central path) (Fig. 1H, I, Fig. 2F, G). Each central path trajectory was resampled to 500 positional points for each *x* and *y* position. For lap-by-lap stereotypy quantification, we found the reciprocal of the average Euclidean distance between lap *i* and lap *j* for each interpolated point.[19] The stereotypy of a given lap *i* was calculated as the average of the trajectory similarity between lap *i* and ±2 laps around the diagonal of the stereotypy matrix (i.e., laps *i-2*, *i-1*, *i+1*, *i+2*-).

Maze linearization

The 2D maze was broken into 2 chunks: central path that spanned SOM to CP and return rails that spanned CP to SOM for each left and right turn (Supplementary Fig. 1A). To avoid any potential overlaps, the zone entry times for each lap were aligned such that the entry of one zone corresponded to the exit of another zone. We subsequently reflected the positional data that corresponded to the time when the rat was in the return rail zone to one side. Similar to above, we found both the median central path and the median rail path and resampled them to 200 equidistant points. We then combined the resampled median central and rail paths to form a single path across the whole track. This was done using the *polyshape* function in Matlab which created a fully enclosed polygon while respecting the overall structure of the input trajectory. This median track path was then resampled in 2D using the *interparc*[20] function, which interpolates points along a 2D perimeter. All laps were then compared against this median template. The 2D $< x, y >$ location of the actual position tracking was then reflected onto this median path using the *dsearchn* function, where the linearized 1D coordinate corresponded to the closest 2D position index, and thus binning the positional data in to 200 bins, which correspond to ~2cm bins. All laps that were further than 20 cm away from the idealized path were excluded from further analysis.

To account for any inconsistencies in linearized maze space due to different maze configurations, we identified the particular bins corresponding to relevant maze locations (e.g., SOM, Wall, CP) in each individual sessions. These identified bins were used for task bracketing analysis that depended on these designations.

Place field analysis

Tuning curves of HC pyramidal cells were calculated based on the linearized maze by first binning the space into ~3cm bins, and then generating the maps of bin spike counts over bin occupancy. Bins with occupancies less than 200 ms were omitted. A smoothed tuning curve (smoothing size: 2 bins) was constructed by convolving a gaussian with the generated tuning curve. The smoothed tuning curves were then thresholded to 25% of their maximum firing rate, and place fields were defined as regions within the thresholded maps with an area greater than 15 cm and minimum firing rate of 1Hz. We then extracted the field centers and sizes using *regionprops* Matlab function on the binarized rate map.[21]

Hippocampal Theta cycle identification

The raw recordings taken from HC were decimated to 2 kHz to extract LFPs. Recordings were bandpass filtered with a Hilbert transform between 6 – 12 Hz to obtain the theta signal.[22-31] The instantaneous phase and amplitude were estimated via the Hilbert transform. The channel to identify theta cycles from was chosen based on the probe channel that had maximal theta power. Given that Hilbert transform assumes the data is sinusoidal and because theta is asymmetric, using previous approaches[32], we first found the peaks and troughs using the filtered signal. We then took these estimated peak and trough times in the raw signal and found the local minima (trough) and maxima (peak) around the estimated times. Theta cycles were limited to 80 to 200 ms. A single theta cycle was defined as peak-to-peak.

Bayesian decoding and sequence analysis

We used Bayesian decoding to calculate ensemble neural representations at both behavioral time-scales and theta-timescales.[33] Briefly, this method computes the posterior probability of the ensemble representing each position in space given spike counts from each single unit in the ensemble. Posterior distributions were normalized to sum to one. Probability of an animal's position x given all observed spikes from recorded units can be operationalized by:

$$P(x, |\text{spikes}) = \frac{P(\text{spikes}|x)P(x)}{P(\text{spikes})}$$

where x is the set of all linear positions on the track, binned in to ~3cm bins. This method assumes that every cell is independent from each other and follows a Poisson process. For decoding of behavioral sequences, as in Fig. 3C, we decoded spiking in 120 ms time windows using uniform spatial prior. This allowed us to validate both our decoding procedure and maze linearization.

For all theta-timescale decoding analyses we used two criteria: 1) minimum of 5 cells had to be spiking in the given theta cycle, and 2) the running speed of the rat had to be faster than 2 cm/s. To examine the spatial decoding in theta cycles (Fig. 3I), we decoded spiking in each half of the theta cycle---peak to trough and trough to subsequent peak---to obtain a probability distribution across the linearized maze (~3cm bins). The obtained posterior probability distribution was subsequently aligned to the rat's position (± 50 cm around the rat). Posterior probabilities obtained from the two halves of theta were subtracted from each other (2nd half – 1st half) to obtain the local vs non-local decoding.[34,35] The null distribution was obtained by shuffling (n = 1000) the identity of the decoding, 1st vs 2nd half (red line in Fig. 3I).

To identify sequential structure of theta cycles, we adapted two widely used measures: 1) theta sequence score[36,37], and 2) theta sequence slope[37-39]. In both cases we used Bayesian decoding methods with time bins defined by 1/6 of a theta cycle measure (measure 1) or 20ms bins (measure 2). In the first method, we used weighted correlation coefficient between time and the decoding posterior $r(x, t|\boldsymbol{P})$. The posterior probability matrix $\boldsymbol{P}$ [num. of spatial bins × num. of time points] was used as weights. In a general sense, correlation coefficient is simply the ratio of the covariances of the two variables, normalized by the square root of their variances. Thus, in this case, the weighted means for location x and time t can be written as:

$$\mu_x(x|\boldsymbol{P}) = \frac{\sum_{i=1}^{T} \sum_{j=1}^{D} P_{ij} x_j}{\sum_{i=1}^{T} \sum_{j=1}^{D} P_{ij}}$$

$$\mu_T(t|\boldsymbol{P}) = \frac{\sum_{i=1}^{T} \sum_{j=1}^{D} P_{tj} t_i}{\sum_{i=1}^{T} \sum_{j=1}^{D} P_{ij}}$$

where P_{ij} is the probability at time i and spatial bin j. The weighted covariance therefore can be written as:

$$covar(x, t|\boldsymbol{P}) = \frac{\sum_{i=1}^{T} \sum_{j=1}^{D} P_{ij} * (x_j - \mu_x)(t_i - \mu_T)}{\sum_{i=1}^{T} \sum_{j=1}^{D} P_{ij}}$$

Finally, the weighted correlation can be calculated by:

$$r(x, t|\boldsymbol{P}) = \frac{covar(x, t|\boldsymbol{P})}{\sqrt{covar(t, t|\boldsymbol{P}) covar(x, x|\boldsymbol{P})}}$$

Weighted correlation for each reconstructed theta sequence was calculated from the decoded probabilities of positions 50 cm ahead and 20 cm behind the animal's current position, and decoding was performed for time windows corresponding to $\pi/6$ phase intervals of each theta cycle. Intuitively, a sequential structure sweeping along the animal's running direction would yield a large positive correlation, whereas lack of structure would be close to 0.

In the second method we computed theta sequence slope by fitting a line that yielded maximum decoded probability within a ± 20 cm window around the rat. For a given line with slope v and intercept ρ, the average likelihood $R(v, \rho)$, that the decoded position is located within a distance d of that line can be written as:

$$R(v, \rho) = \frac{1}{n} \sum_{k=0}^{n-1} P(|pos - vk * \Delta t| \le d)$$

where k is the temporal bin of the posterior probability matrix and Δt is the moving step of the decoding window, in this case 20 ms. The value of d was set to be 10 cm around the rat. In line with prior research, we densely sampled the parameter space for v to find the value that maximized R.

To calculate the significance of either theta sequences (both sequence score and sequence slope), we circularly shifted the space-bins of the posterior probability distribution (n = 1000) for each time point. As previously described this shuffle conserves the structure of decoded probabilities within a bin but randomizes association between bins. A theta sequence was significant if it's score was above the 95[th] or below the 5[th] percentile of the shuffled distribution to account for both forward and reverse sequences.[36,37,39]

Correlation between theta metrics and behavior

Three behavioral metrics were correlated to theta sequence score and theta sequence slope (Fig. 4, Fig. Supplementary Fig. 4). For exploration and stereotypy correlations, only theta sequences within the central path were considered. For zIdPhi correlations, only theta sequences at the choice point were considered. For each session we performed a lap-by-lap correlation between the median theta sequence metric and the behavioral metric (Fig. 4A, C, E & Supplementary Fig. 4A, C, E). For binarized exploration vs. no exploration or VTE vs. no VTE, for each session we calculated the median theta score of

explorative laps, VTE laps, non-explorative laps, and non-VTE laps, thus the distributions reflect session means (Fig. 4B, D, Supplementary Fig. 4B, D).

Linearized firing rate by lap

To construct the firing rate maps in Fig. 5A, we created rate maps for every lap and every DLS cell as a function of the linear location on the maze (binned into ~3cm bins) resulting in a [laps × num. of spatial bins × num. of cells] matrix (35.8 ± 2.8 S.E.M number of single units/session). For each session we took the 25 laps from the start of the session and 25 laps following a rule switch (overall rats ran 45.7 ± 12.9 number of laps/rule block), and averaged firing rates across sessions for each maze complexity (low, mid, and high).

Task-bracketing index

Following previous research, we calculated task bracketing index by subtracting the mean normalized firing rate in the central path from the mean normalized firing rate at the SOM for each cell.[19,40-43] Zones were defined based on entry and exit times for each session. Positive task bracketing indicates relatively strong firing in the SOM compared to the subsequent navigation through the central path, while values near 0 indicate that there is little to no difference in firing rate at the start compared to the rest of the maze. Given the observed firing rates of DLS cells across the linearized maze, we also defined the start and end positions and SOM and Choice point, and this analysis also yielded the same results.

Bursts of ensemble activity and firing rate percentiles

A DLS cell was considered to be bursting at a given time point if its firing rate exceeded 97.5% of its firing rate throughout the session (binned into 500ms bins). The burst time was then mapped onto the closest behavioral time point. Thus, a cell was considered to burst at a given zone if the behavioral time point resided within the zone entry and exit times (Supplementary Fig. 1A)

We used the HC place field analysis to calculate the number of DLS cells that restrict their firing to the central segment of the maze (Supplementary Fig. 5D). Briefly, we created tuning curves of each DLS cell on the linearized maze (~3cm bins). A smoothed tuning curve (smoothing size: 2 bins) was constructed by convolving a gaussian with the generated tuning curve. The smoothed tuning curves were then thresholded to 20% of their maximum firing rate, and place fields were defined as regions within the thresholded maps with an area greater than 15 cm and minimum firing rate of 0.1Hz. We then extracted the field centers and sizes using *regionprops* Matlab function on the binarized rate map and counted the number of "field" centers that were within the central segment of the maze.

Spatial cross correlation of ensemble activity

To investigate how spatial representations were similar to each other across the linearized maze we calculated the spatial cross correlation of ensemble activity (Fig. 5C, Supplementary Fig. 5A). We correlated neural rate maps between pairs of spatial bins. For each cell we calculated the firing rate in each of the spatial bins throughout the entire maze. For each pair of spatial bins, we computed the correlation between the normalized firing rate vectors for those bins. Therefore, if representations are local (e.g., as is in HC Supplementary Fig. 5A) the spatial cross correlation would only be high restricted around the diagonal of the matrix.

Population correlation of ensemble activity

Following previous research, to measure the similarity between population activity across laps, we correlated neural rate maps between pairs of laps.[1,4] For each cell we calculated the firing rate in each of the 5 zone bins (Supplementary Fig. 1A) along the central stem (SOM, Wall 1, Wall 2, Wall 3, Choice point). For each pair of laps, we computed the correlation between the normalized firing rate vectors for the lap (Fig. 5B). We averaged windows of 10 laps prior and 20 laps post rule switch across sessions to generate switch-aligned average correlation matrices for mPFC, HC, and DLS (Fig. 5A).

Ensemble population activity change point

To calculate the population activity change point we used the individual session population correlations (Fig. 6B). To determine where there are abrupt changes to the population correlation matrix, we used hierarchical clustering to identify reliable neural representations. Briefly, this method seeks to cluster data in an input matrix by constructing a hierarchical cluster tree. Unlike k-means clustering, because the linking tree is fixed the starting point is not random.

We first created the linking tree for the population correlation matrix using *linkage* function in Matlab and set the method to minimize the inner Euclidean distance between elements (method = ward). Subsequently, we performed agglomerative clustering of the created hierarchical cluster tree (linkage tree) by using *cluster* function in Matlab. Because the clustering depends on the number of allowed clusters, we followed previous work and repeated our analysis multiple times allowing the maximum clusters to range from 3 to 7.[1,4] We then applied a change point analysis similar to that for detecting behavioral changes[15], but here based on cluster identity number as opposed to choice identity. For each cluster size and session, we cumulatively summed the output cluster IDs, and using the *findchangepts* function in Matlab, found the laps in which there was a significant change in the cumulative sum of cluster IDs based on slope and mean (statistic = linear) (Fig. 6B). Intuitively, this change point reflects the moment in which neural ensemble data shifted from one consistent representation to another. For each cluster size and session, to calculate the null distribution for lap change points, we shuffled (n = 75) the cluster IDs across laps and found the change point of the shuffled data.

We calculated the distribution of change points aligned to change points as the difference between all the change points identified by the clustering in a given switch block to the block's switch lap. The minimum distance was set as the minimum of the above distribution.

mPFC activity prediction of theta score

We used 5-fold cross-validated Bayesian decoding to decode whether ensemble activity in mPFC predicted hippocampal theta sequence score. For each theta cycle we calculated the spiking activity of mPFC ensembles. We only took theta cycles in which there were at least 5 mPFC and HC cells active and thus had computed corresponding theta sequence scores for the given cycles. Bayesian decoding requires the expected firing of a cell given the variable of interest to create the likelihood function. We defined our variable of interest as the theta sequence score (n = 15 bins), such that our likelihood function was the expected firing of mPFC cells given binned theta sequence score. Similar to above, we assumed a uniform prior over theta sequence scores. Subsequently, for each theta cycle we identified the discretized bin of its respective theta sequence score and found the decoded posterior probability for that particular bin---the probability of identifying the correct theta sequence score bin for the particular theta cycles, given a set of mPFC activity.

We used three different shuffle methods to create null distributions[2]: 1) shuffled mPFC, which breaks any relationships within mPFC and/or between mPFC and HC, 2) across session shuffled theta sequence scores, which preserves the mPFC representations, but breaks any moment-by-moment relationship between mPFC and HC, 3) within-lap shuffled theta sequence scores, which breaks the relationship between mPFC and HC within a given lap, but preserves longer behaviorally relevant timescale relationships.

In the first method, we shuffled mPFC spike times by circularly shifting the full spike train of each cell by a random amount ranging from 1 min to 45 mins to disassociate neural spiking from variable of interest. In the second method, we shuffled all theta sequence scores within a session by randomly permuting the vector of theta sequence scores to disassociate theta sequence scores from neural spiking in mPFC. Finally, in the third method, we isolated theta sequence scores for each lap, and similar, randomly permuted theta sequence scores within a given lap. For each shuffle we used 5-fold Bayesian decoding to create null distributions. All shuffles were performed 5 times and the mean value was taken across the shuffles as the respective null value.

Statistical analysis

Data analyses were performed using custom MATLAB (MATLAB 2021a, Mathworks) functions. For comparison across environmental complexity we used One-way ANOVAs when evaluating behavioral data and a Kruskal-Wallis test when evaluating physiology. For comparisons between pairs of groups we used a Wilcoxon rank-sum test. When relating exploratory and procedural behavioral parameters to HC theta sequence score we used one sample t-tests. To assess the differences in correlation between lap-by-lap DLS bursting activity and theta sequence score we used ANCOVA with post-hoc comparisons using Pearson correlations. For comparison of behavioral responses in DREADD experiments across days we used a Wilcoxon signed-rank test. All assessments of behavior and physiologic responses across laps around a rule switch was done with a One-way repeated measures ANOVA with pairwise post-hoc via rank-sum with Bonferroni corrections. All significance was assessed at α = 0.05. Where appropriate we performed ANOVA's with a random effect of rat.

References

Hasz, B. M. & Redish, A. D. Dorsomedial prefrontal cortex and hippocampus represent strategic context even while simultaneously changing representation throughout a task session. *Neurobiology of Learning and Memory* **171**, 107215, doi:https://doi.org/10.1016/j.nlm.2020.107215 (2020).

Hasz, B. M. & Redish, A. D. Spatial encoding in dorsomedial prefrontal cortex and hippocampus is related during deliberation. *Hippocampus* **30**, 1194-1208, doi:10.1002/hipo.23250 (2020).

Blumenthal, A., Steiner, A., Seeland, K. & Redish, A. D. Effects of pharmacological manipulations of NMDA-receptors on deliberation in the Multiple-T task. *Neurobiol Learn Mem* **95**, 376-384, doi:10.1016/j.nlm.2011.01.011 (2011).

Powell, N. J. & Redish, A. D. Representational changes of latent strategies in rat medial prefrontal cortex precede changes in behaviour. *Nat Commun* **7**, 12830, doi:10.1038/ncomms12830 (2016).

Gupta, A. S., van der Meer, M. A., Touretzky, D. S. & Redish, A. D. Segmentation of spatial experience by hippocampal theta sequences. *Nat Neurosci* **15**, 1032-1039, doi:10.1038/nn.3138 (2012).

Nagai, Y. *et al.* Deschloroclozapine, a potent and selective chemogenetic actuator enables rapid neuronal and behavioral modulations in mice and monkeys. *Nat Neurosci* **23**, 1157-1167, doi:10.1038/s41593-020-0661-3 (2020).

Pachitariu, M., Sridhar, S. & Stringer, C. Solving the spike sorting problem with Kilosort. *bioRxiv*, 2023.2001. 2007.523036 (2023).

Rossant, C., Kadir, S., Goodman, D., Hunter, M. & Kenneth, H. *Phy. Version 1.0.9* (2016).

Schmitzer-Torbert, N. C. & Redish, A. D. Task-dependent encoding of space and events by striatal neurons is dependent on neural subtype. *Neuroscience* **153**, 349-360, doi:10.1016/j.neuroscience.2008.01.081 (2008).

Dehmer, M. & Mowshowitz, A. A history of graph entropy measures. *Information Sciences* **181**, 57-78 (2011).

Bonchev, D. G. Information theoretic complexity measures. *Encyclopedia of Complexity and Systems Science* **5**, 4820-4838 (2009).

Newman, M. E. J., Barabási, A.-L. s. & Watts, D. J. *The structure and dynamics of networks*. (Princeton University Press, 2006).

Trucco, E. A note on the information content of graphs. *The bulletin of mathematical biophysics* **18**, 129-135 (1956).

Mugan, U. & MacIver, M. A. Spatial planning with long visual range benefits escape from visual predators in complex naturalistic environments. *Nat Commun* **11**, 3057, doi:10.1038/s41467-020-16102-1 (2020).

Gallistel, C. R., Fairhurst, S. & Balsam, P. The learning curve: implications of a quantitative analysis. *Proc Natl Acad Sci U S A* **101**, 13124-13131, doi:10.1073/pnas.0404965101 (2004).

Janabi-Sharifi, F. & Marey, M. A kalman-filter-based method for pose estimation in visual servoing. *IEEE transactions on Robotics* **26**, 939-947 (2010).

Papale, A. E., Stott, J. J., Powell, N. J., Regier, P. S. & Redish, A. D. Interactions between deliberation and delay-discounting in rats. *Cogn Affect Behav Neurosci* **12**, 513-526, doi:10.3758/s13415-012-0097-7 (2012).

Steiner, A. P. & Redish, A. D. The road not taken: neural correlates of decision making in orbitofrontal cortex. *Front Neurosci* **6**, 131, doi:10.3389/fnins.2012.00131 (2012).

Cunningham, P. J., Regier, P. S. & Redish, A. D. Dorsolateral Striatal Task-initiation Bursts Represent Past Experiences More than Future Action Plans. *J Neurosci* **41**, 8051-8064, doi:10.1523/JNEUROSCI.3080-20.2021 (2021).

D'Errico, J. interparc (https://www.mathworks.com/matlabcentral/fileexchange/34874-interparc). *MATLAB Central File Exchange* (2010).

Duvelle, E. *et al.* Hippocampal place cells encode global location but not connectivity in a complex space. *Curr Biol* **31**, 1221-1233 e1229, doi:10.1016/j.cub.2021.01.005 (2021).

Johnson, A. & Redish, A. D. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J Neurosci* **27**, 12176-12189, doi:10.1523/JNEUROSCI.3761-07.2007 (2007).

Kay, K. *et al.* Constant Sub-second Cycling between Representations of Possible Futures in the Hippocampus. *Cell* **180**, 552-567 e525, doi:10.1016/j.cell.2020.01.014 (2020).

Zielinski, M. C., Shin, J. D. & Jadhav, S. P. Coherent Coding of Spatial Position Mediated by Theta Oscillations in the Hippocampus and Prefrontal Cortex. *J Neurosci* **39**, 4550-4565, doi:10.1523/JNEUROSCI.0106-19.2019 (2019).

Mizuseki, K., Sirota, A., Pastalkova, E. & Buzsaki, G. Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron* **64**, 267-280, doi:10.1016/j.neuron.2009.08.037 (2009).

Buzsaki, G. Theta oscillations in the hippocampus. *Neuron* **33**, 325-340, doi:10.1016/s0896-6273(02)00586-x (2002).

Maurer, A. P., Cowen, S. L., Burke, S. N., Barnes, C. A. & McNaughton, B. L. Organization of hippocampal cell assemblies based on theta phase precession. *Hippocampus* **16**, 785-794, doi:10.1002/hipo.20202 (2006).

Maurer, A. P. & McNaughton, B. L. Network and intrinsic cellular mechanisms underlying theta phase precession of hippocampal neurons. *Trends Neurosci* **30**, 325-333, doi:10.1016/j.tins.2007.05.002 (2007).

Skaggs, W. E., McNaughton, B. L., Wilson, M. A. & Barnes, C. A. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* **6**, 149-172, doi:10.1002/(SICI)1098-1063(1996)6:2<149::AID-HIPO6>3.0.CO;2-K (1996).

O'Keefe, J. & Recce, M. L. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* **3**, 317-330, doi:10.1002/hipo.450030307 (1993).

Lubenov, E. V. & Siapas, A. G. Hippocampal theta oscillations are travelling waves. *Nature* **459**, 534-539, doi:10.1038/nature08010 (2009).

Cole, S. & Voytek, B. Cycle-by-cycle analysis of neural oscillations. *J Neurophysiol* **122**, 849-861, doi:10.1152/jn.00273.2019 (2019).

Zhang, K., Ginzburg, I., McNaughton, B. L. & Sejnowski, T. J. Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *J Neurophysiol* **79**, 1017-1044, doi:10.1152/jn.1998.79.2.1017 (1998).

Schmidt, B., Duin, A. A. & Redish, A. D. Disrupting the medial prefrontal cortex alters hippocampal sequences during deliberative decision making. *J Neurophysiol* **121**, 1981-2000, doi:10.1152/jn.00793.2018 (2019).

Amemiya, S. & Redish, A. D. Hippocampal Theta-Gamma Coupling Reflects State-Dependent Information Processing in Decision Making. *Cell Rep* **22**, 3328-3338, doi:10.1016/j.celrep.2018.02.091 (2018).

Farooq, U. & Dragoi, G. Emergence of preconfigured and plastic time-compressed sequences in early postnatal development. *Science* **363**, 168-173, doi:10.1126/science.aav0502 (2019).

Feng, T., Silva, D. & Foster, D. J. Dissociation between the experience-dependent development of hippocampal theta sequences and single-trial phase precession. *J Neurosci* **35**, 4890-4902, doi:10.1523/JNEUROSCI.2614-14.2015 (2015).

Davidson, T. J., Kloosterman, F. & Wilson, M. A. Hippocampal replay of extended experience. *Neuron* **63**, 497-507, doi:10.1016/j.neuron.2009.07.027 (2009).

Drieu, C., Todorova, R. & Zugaro, M. Nested sequences of hippocampal assemblies during behavior support subsequent sleep replay. *Science* **362**, 675-679, doi:10.1126/science.aat2952 (2018).

Barnes, T. D., Kubota, Y., Hu, D., Jin, D. Z. & Graybiel, A. M. Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature* **437**, 1158-1161, doi:10.1038/nature04053 (2005).

Smith, K. S. & Graybiel, A. M. A dual operator view of habitual behavior reflecting cortical and striatal dynamics. *Neuron* **79**, 361-374, doi:10.1016/j.neuron.2013.05.038 (2013).

Thorn, C. A., Atallah, H., Howe, M. & Graybiel, A. M. Differential dynamics of activity changes in dorsolateral and dorsomedial striatal loops during learning. *Neuron* **66**, 781-795, doi:10.1016/j.neuron.2010.04.036 (2010).

Regier, P. S., Amemiya, S. & Redish, A. D. Hippocampus and subregions of the dorsal striatum respond differently to a behavioral strategy change on a spatial navigation task. *J Neurophysiol* **114**, 1399-1416, doi:10.1152/jn.00189.2015 (2015).

Figure 1: Behavior in left/right/alternation task and environment complexity

(A) The left/right/alternation (LRA) foraging task. Three different wall types: open middle (Wall A), open left side (Wall B), and open right side (Wall C) were used to create intermediary decision-points. **(B)** Recording sites from dorsal mPFC, DLS, and dorsal HC. **(C)** Example histology (right) and illustration of viral spread (left) for DREADDs virus animals (red) and control virus animals (blue). **(D)** Example behavioral data from a single session and the hypothesized arbitration between deliberative and procedural decision-making systems. **(E)** Rat performance aligned to rule switches. **(F)** Example vicarious trial and error (VTE) and non-VTE paths. **(G)** z-Scored IdPhi aligned to rule switches. **(H)** Two example paths (orange and green) through the central maze segment showing low (top) and high (bottom) path stereotypy. **(I)** Trajectory stereotypy (reciprocal of Euclidian distance) between pairs of paths through the central maze segment around rule switches. **(J)** Environmental complexity of different maze structures created by rearranging inserts shown in (A). The colored circles indicate the low, mid, and high complexity designations. **(K)** Illustration of hypothesized behavior and neurophysiology in differently-complex environments. For panels E-G, shaded error bars show mean ± SEM. The vertical dotted lines correspond to the last lap of the previous rule block.

Figure 2: Behavioral effects of environmental complexity on the LRA task

(A) Proportion of VTE laps (zIdPhi > 0.25) in low, mid, and high complexity environments. **(B)** Two example sessions showing the occupancy of the central maze segment. *Top*: example of a low exploration session, and *Bottom*: an example of a high exploration session. **(C)** Percentage of exploratory laps aligned to rule switches. Black line indicates mean, and the shaded gray region ± SEM. The vertical dotted line corresponds to the last lap of the previous rule block. **(D)** Distribution zIdPhi values at the choice point for non-exploratory and exploratory laps as measured on the central track. Horizontal line denotes the boundary for identifying VTE events (Kruskal-Wallis test H(1) = 46.41, p < 0.001). **(E)** Percentage of exploratory laps per session broken down by maze complexity group. **(F)** Examples of trajectory stereotypy in different environments (as in Fig. 1H). **(G)** Development of stereotypy around rule switches (as in Fig. 1I).

Deliberative Behaviors

Procedural Behavior

Figure 3: Theta sequences in hippocampus in differently-complex environments

(A) Two example HC cell tuning curves. (B) Spatially binned activity of putative pyramidal cells on the linearized maze (~3cm bins). (C) An example of HC ensemble representation of space. *First row*: The maze and an example trajectory (dark blue) along with raw (gray) and theta frequency filtered LFP (green). *Second row*: spike raster of simultaneously recorded HC cells on an example lap. Cells are ordered and colored by their respective tuning curve centers. Gray line shows the actual trajectory of the rat. *Third row*: Decoded activity from the HC ensemble. Black line shows the actual trajectory of the rat in linearized space. Dashed horizontal lines indicate the high-cost choice point and feeder locations. (D, E) Examples of (D) high and (E) low score theta sequences. *Left*: HC cell spiking during a theta cycle ordered and colored by field location. Raw LFP (gray) and theta filtered LFP (green) are shown below with vertical lines indicating the peaks and trough of the theta cycle. *Right*: Bayesian decoding during the theta cycle. Theta sequence scores are indicated on the plots. The arrow and cyan line show the actual position of the rat during the theta time window. (F) Average percent of significant theta sequences per session. (G) Distribution of theta sequence scores for the different maze complexity groups. Dashed lines indicate medians. (H) Session average theta sequence scores along the central path for the different maze complexity groups. (I) Differences in posterior probability of the decoding in the 1st half (peak to trough) and 2nd half (trough to peak) of each identified theta cycle (1st half –2nd half) aligned to the actual position of the rat. Positive values indicate more local and negative values indicate more non-local decoding. The red line indicates the shuffle distribution. Black lines indicate significant differences between pairs of environment complexity (top line low vs mid, middle line mid vs high; bottom line low vs high). (Repeated measures ANOVA for complexities across positions; Interaction: $F_{(34)} = 51$, $p < 0.001$. Post-hoc on each position bin via two-sided Wilcoxon rank sum corrected for multiple comparisons;).

Figure 4: Theta sequence score and its relation to deliberative behaviors

(A) Distribution of correlations between zIdPhi and mean theta sequence score at the choice point. **(B)** Average theta sequence score at the choice point for laps identified as VTE and non-VTE. **(C)** Distribution of correlations between exploration amount and mean theta sequence score in the central path. Exploration amount was calculated as the average per-lap deviation from the median path in the central maze segment. **(D)** Average theta sequence score through the central path for laps that were identified as exploratory and non-exploratory. **(E)** Distribution of correlations between central path stereotypy and average theta sequence score during central path traversal. (A, C, E) Dashed lines indicate distribution medians.

Deliberative Behaviors   ## Procedural Behavior

Figure 5: DLS ensemble activity across the linearized maze and the arbitration between deliberative and procedural decision-making

(A) *Top row*: Average firing rate of all putative dorsolateral striatum medium spiny neurons (MSNs) over the linearized track (~3cm bins). Firing rates are plotted as a function of laps aligned to the start of rule blocks. *Bottom row*: Illustration of maze locations where we found increased firing. **(B)** Average task-bracketing score for each session broken down by maze complexity group. **(C)** Spatial cross correlation of DLS ensemble activity broken down by maze complexity group. The dashed lines indicate the choice point and feeder sites. **(D)** Proportion of cells that were identified as bursting at each maze location (firing rate > 97.5% of its session average), broken down by maze complexity. The dots indicate mean and the error bars ± SEM. **(E)** *Blue:* Mean HC theta sequence score across the central path aligned to rule switches. *Orange:* DLS Task-bracketing aligned to rule switches. For both plots, solid lines indicate mean, and the shading ± SEM. The vertical dotted line corresponds to the last lap of the previous rule block. **(F)** Scatter plot of lap-by-lap pairs of simultaneously recorded DLS task-bracketing scores and mean HC theta sequence score on the central path for each lap as a function of maze complexity (ANCOVA $F(2) = 9.76$, $p < 0.001$). Lines indicate correlations.

Figure 6: Representational transitions after rule switches in mPFC, HC and DLS and predictions of prospective representations in HC from mPFC activity

(A) Lap-by-lap correlation of mPFC, DLS, and HC population activity aligned to rule switches. Correlation matrices are averaged across the rule changes. The vertical and horizontal dashed lines correspond to the last lap of the previous rule block. **(B)** *Top row:* Example session lap-by-lap correlations of mPFC, DLS and HC population activity. *Bottom row:* Cluster IDs across laps derived from hierarchical clustering performed on the population correlation matrix. The dashed black lines indicate task rule switches, and the magenta dashed lines indicate detected cluster change points. **(C)** Distribution of physiology-based change points aligned to the task rule switch laps. **(D)** Distribution of the difference in physiology-based change points for simultaneously recorded rats. **(E)** Distribution of the posterior probabilities of mPFC activity predicting binned theta sequence score. The red line indicates actual data. The blue lines indicate the different shuffles. The dashed line indicates uniform probability.

Figure 7: Behavioral effects of mPFC inactivation with DREADDs

(A) Example histology (right) and illustration of viral spread (left) for DREADDs virus animals (red) and control virus animals (blue). **(B)** Distribution of rat performance for mPFC inactivated rats (DREADDs: DCZ dark pink) and control virus animals (light pink). **(C)** Distribution of minimum distance between detected behavioral change points and task rule switch laps for DREADDs virus rats under DCZ, post DCZ (2nd day saline repeat of the task), and VEH, as well as all control virus animals. **(D)** Distribution of proportion of VTE events per session for prefrontal inactivated animals (DREADDs: DCZ, red), for the day after prefrontal inactivation (DREADDs: post DCZ, blue), and for control virus animals (Control, light pink). **(E)** Difference in the proportion of VTE events between 1st day manipulation and 2nd day rebound (Day 2 – Day 1) for DREADDs (dark pink) and control (light pink) virus animals with respect to drug condition. There was a significant virus and DCZ interaction: ANOVA virus*DCZ $F(1) = 3.86$, $p = 0.05$. **(F)** Per session distribution of the percentage occupied space in the central segment of the maze for prefrontal inactivated animals (DREADDs: DCZ, red), for the day after prefrontal inactivation(DREADDs: post DCZ, blue), and for control virus animals (Control, light pink). **(G)** Difference in the percentage of the central segment of the maze occupied between 1st day manipulation and 2nd day rebound (Day 2 – Day 1) for DREADDs (dark pink) and control (light pink) virus animals with respect to drug condition. There was a significant virus and DCZ interaction: ANOVA virus*DCZ $F(1) = 25.41$, $p < 0.001$. **(H)** zIdPhi aligned to rule switches for DREADDs and control virus animals. The horizontal dashed line indicates the threshold for which events were identified as binary VTE and non-VTE events. The vertical dotted line corresponds to the last lap of the previous rule block. The black line above indicates laps in which there was a significant difference between Active and Control groups. Data are shown as mean ± SEM. **(I)** Proportion of exploratory laps aligned to rule switches for DREADDs and control animals. Same display as (H). **(J)** Trajectory stereotypy (reciprocal of Euclidian distance) between pairs of paths through the central maze segment around rule switches for DREADDs and control virus animals. **(K)** Rat performance across the experimental weeks for DREADDs and control virus animals. Each 'week' is an 8-day sequence (see Supplementary Fig. 6C).

Supplementary Figure 1: Task design and recordings, Related to Figure 1

(A) The left/right/alternation maze. The colored blocks indicate the 11 zones the maze was broken into. **(B)** An example 8-day sequence of the experimental paradigm. A set of 4 mazes were selected across the complexity space. Rats were run on the same maze (i.e., same central path configuration) with the same contingencies for 2 subsequent days. This 8-day sequence was repeated 2 times with randomized mazes and rule contingencies. **(C)** Breakdown of the sex and probe targeting across rats. **(D)** *Left:* Silicon probe recording locations in prefrontal cortex (top row), hippocampus (middle row) and dorsolateral striatum (bottom row). Circles indicate the termination of the recording sites, and the extending line indicates the length of the recording sites. For 2 rats (outlined), no histology was available and so identified locations are estimates based on targeting coordinates, turning notes, and physiologic landmarks. *Right:* Histological examples of recording sites from prefrontal cortex, hippocampus, and dorsolateral striatum.

A

B

Example Experimental Phase				
Day	**Maze**	**Complexity**	**Contingency**	**Switch Times (sec)**
1 2	ABC	0.75	ALR	881, 1890
3 4	BBB	0.44	LRA	1018, 1782
5 6	AAA	0.00	RLA	852, 1837
7 8	CAB	0.70	LAR	914, 1855

Repeat x 2 — shuffled contingencies — shuffled mazes

C

Rat IDs	Sex	Recording region	Hemisphere
R668	Male	Hippocampus	Left
		dorsal mPFC	
R677	Female	Hippocampus	Left
		DLS	Right
R679	Female	DLS	Left
		dorsal mPFC	
R699	Female	Hippocampus	Right
		dorsal mPFC	
R710	Male	Hippocampus	Right
		DLS	Left
R721	Female	Hippocampus	Right
		dorsal mPFC	
R786	Male	Hippocampus	Left
		dorsal mPFC	
R787	Male	Hippocampus	Right
		dorsal mPFC	
R788	Male	DLS	Right
		dorsal mPFC	

D

R668 ■ R679 ■ R699 ■ R788
R721 ■ R787 ■ R786

A/P: 2.8 M/L: 0.7

R668 ■ R677 ■ R699 ■ R787
R710 ■ R721 ■ R786

A/P: -4 M/L: 3

R677 ■ R679
R710 ■ R788

A/P: 0.7 M/L: 3.5

R786

R721

R786

R721

R710

R677

Stained with anti-μ-Opioid Receptor

Supplementary Figure 2: Behavior on the left/right/alternation task, Related to Figure 1 and Figure 2
(A) Distribution of minimum distance between detected behavioral change points to the corresponding task rule switch across all sessions. **(B)** Number of laps run per contingency/rule identity (One-way ANOVA $F(18) = 3.61$, $p < 0.001$). **(C)** Rat performance for each contingency/rule (One-way ANOVA $F(18) = 13.47$, $p < 0.001$). **(D)** Rat performance in low, mid, and high complexity environments (One-way ANOVA $F(18) = 2.35$, $p = 0.0033$). **(E)** Proportion of exploratory laps aligned to the start of the session (initial rule block). The black line indicates mean, and the shaded gray region SEM. **(F)** *Left*: An example of high exploration of the start of maze (SOM) and central maze segment (central path). *Right*: Quantification across all sessions of the percent space occupied in each maze region (see Supplementary Fig. 1A for maze zone designations). The dots indicate mean and the error bars ± SEM. **(G)** Distribution of all zIdPhi values. zIdPhi was a positively skewed bimodal distribution, and in line with prior research, we set the threshold of VTE to be around the cutoff between the two peaks (0.25). **(H)** Distribution of correlations between zIdPhi and exploration amount (t-test $t(143) = 6.86$, $p < 0.001$). Exploration amount was calculated as the lap's average deviation from the median path through the central maze segment. The vertical dashed line indicates the median. **(I)** Distribution of average deviation from the median central path for each rule block (start, switch 1 and switch 2) separated by complexity (ANOVA complexity and switch block interaction $F(4) = 2.61$, $p = 0.0352$).

A

B

C

D

E

F

G

H

I

Supplementary Figure 3: Basic hippocampal physiology and theta sequence slopes in differently-complex environments, Related to Figure 3

(A) Distribution of speeds through the central path for different maze complexities (One-way ANOVA $F(18) = 1.29$, $p = 0.20$). **(B)** Total number of HC place fields broken down by complexity. There was no significant difference in the number of fields detected (One-way ANOVA $F(2) = 0.53$, $p = 0.59$). **(C)** Distribution of HC place field centers across the linearized maze (~3 cm bins). The distribution of HC place field centers is largely uniform with a slight overrepresentation at the start and choice point. However, the complexity and structure of the maze did not affect the distribution of the place fields. **(D)** HC place field sizes broken down by complexity. In more complex environments fields were smaller (One-way ANOVA $F(2) = 9.91$, $p = 0.0001$). **(E, F)** Same as Fig. 3F, G. but using theta sequence slopes instead of theta sequence scores. (E) Average percent of significant theta sequence slopes per session. Similar to Fig. 3F we did not find a significant difference in the proportion of theta sequence slopes that were identified as significant (see Methods) (One-way ANOVA $F(12) = 1.12$, $p = 0.36$). (F) Distribution of theta sequence slopes of the different maze complexity groups. Similar to Fig. 3G w found that complex environments favored steeper theta sequences (Kruskal-Wallis $H(2) = 66.67$, $p < 0.001$).

Supplementary Figure 4: Exploration as a deliberative behavior and the relation between theta sequence slopes and deliberative behaviors, Related to Figure 4

(A-E) Same analysis as Fig. 4A-E but for theta sequence slopes. (A) Distribution of correlation between theta sequence slopes and zIdPhi (t-test(81) = -3.68, p < 0.001). (B) Theta sequence slopes for VTE and non-VTE laps (Kruskal-Wallis Test H(1) = 5.72, p = 0.0168). (C) Distribution of correlation between theta sequence slopes and exploration amount (t-test(84) = -3.29, p = 0.0015). (D) Theta sequence slopes for exploratory and non-exploratory laps (Kruskal-Wallis Test H(1) = 12, p = 0.0005). (E) Distribution of correlation between theta sequence slope and path stereotypy through the central maze segment (t-test (84) = 2.56, p = 0.0122). For panels A, C, E the vertical line indicates the median.

Deliberative Behaviors   Procedural Behavior

Supplementary Figure 5: Spatial cross correlation of HC and DLS, Related to Figure 5

(A) Display as in Fig. 5C. Spatial cross correlation of HC ensemble activity broken down by maze complexity group. The dashed lines indicate the choice point and feeder sites. HC shows largely local representations throughout the linearized maze (~3cm bins). **(B)** Standard deviation of the spread off the diagonal of the spatial cross correlation matrix, restricted to the central path (Fig. 5C). Intuitively, greater standard deviation indicates rapid fall-off and a tighter distribution around the diagonal, thus showing that complex environments have significantly more local and restricted representations (One-way ANOVA $F(2) = 5.83$, $p = 0.0051$). **(C)** Same as Fig. 5F but broken down by rat. For each maze complexity, lap-by-lap correlation of simultaneously recorded DLS task-bracketing and mean HC theta sequence score on the central path. Solid lines indicate the line of best fit. *R677*: $r_{low} = -0.2$, $p_{low} < 0.001$, $r_{mid} = 0.05$, $p_{mid} = 0.23$, $r_{high} = 0.06$, $p_{High} = 0.88$ (ANCOVA $F(2) = 9.62$, $p < 0.001$). *R710*: $r_{low} = -0.12$, $p_{low} = 0.11$, $r_{mid} = -0.05$, $p_{mid} = 0.69$, $r_{high} = -0.11$, $p_{High} = 0.41$ (ANCOVA $F(2) = 0.01$, $p = 0.99$). **(D)** Distribution of the number of spatially restricted 'place-like' DLS cells in the central segment of the maze broken down by complexity (One-way ANOVA $F(2) = 3.38$, $p = 0.04$).

A

Hippocampus

Low complexity   Mid complexity   High complexity

B

C

R677   R710

score*bracket F(2) = 9.62, p<0.001   score*bracket F(2) = 0.01, p<0.99

D

Supplementary Figure 6: Histology and behavioral effects of mPFC inactivation with DREADDs, Related to Figure 7

(A) All histological sections from virus transfected rats included in the study. Top row shows the histological sections from rats transfected with the active DREADDs virus (AAV5-CaMKIIα-hM4Di(Gi)-mCherry), and the bottom row shows the histological sections from rats transfected with the control virus (AAV5-CAMKIIα-mCherry). Note that in most cases expression is bilateral and localized to dorsal mPFC (anterior cingulate and prelimbic cortices). **(B)** Proof of concept for DREADDs and DCZ inactivation (n_{Rats} = 1, $n_{Sessions}$ = 2). Average multiunit firing over 3hrs while the rat was on a flowerpot to measure the changes to baseline neural activity under the VEH (Left) and DCZ (Right) conditions. The vertical lines indicate the time of injection and the putative experiment start time (~30 mins after injection). Inset shows the histology for the DREADDs expression and probe placement. **(C)** An example 8-day sequence of the experimental paradigm for prefrontal manipulations. A total of 14 different mazes were used and selected from different variations of the maze inserts to span the complexity across the entire experimental phase. 4 unique mazes that spanned the complexity space were used across the 8-day sequence. Rats were run on the same maze (i.e., same central path configuration) with the same contingencies for 2 consecutive days, in which the second day was always a saline washout day. This 8-day sequence was repeated 4 times with randomized mazes, contingencies, and drug conditions. **(D)** Distribution of performance on the task (% correct) for DREADDs virus rats under DCZ, post DCZ Saline, and VEH, as well as all control virus animals (One-way ANOVA F(3) = 18.31, p < 0.001; HSD Post-hoc DCZ vs all others p < 0.05, post DCZ Saline vs VEH p = 0.43, control vs all others p < 0.05). **(E)** Difference in performance between 1st day manipulation and 2nd day rebound (Day 2 – Day 1) for DREADDs and control virus animals with respect to drug condition. There was no significant effect of virus, and all the rats significantly increased their performance on the repeat day (2nd day) (ANOVA virus F(1) = 0.35, p = 0.55; Wilcoxon signed rank DREADDs: rebound DCZ z = 2.38, p = 0.0088; DREADDs: rebound VEH z = 1.65, p = 0.049; Control: rebound DCZ z = 2.37, p = 0.0088; Control: rebound VEH z = 1.91, p = 0.028). **(F & G)** Same as Fig. 7D, F. (F) Distribution of proportion of VTE events/session for different drug conditions (G) Distribution of the occupancy of the central segment of the maze/session for different drug conditions.

Supplementary Fig. 7: Long-term behavioral effects of DREADD inactivation of dorsal mPFC, Related to Figure 7

(A) Same as Fig. 7J. Lap stereotypy aligned to task rule switches broken down by drug condition. Overall, DREADDs virus animals showed less stereotypy when compared to control virus animals. mPFC inactivation (DCZ) in DREADDs virus rats increased stereotypy; however, on the 2^{nd} day after both mPFC inactivation (post DCZ Saline and VEH, DREADDs rats displayed less behavioral stereotypy. **(B)** Distribution of the proportion of VTE events across the experimental weeks for DREADDs and control virus animals. Similar representation to Fig. 7K (ANOVA virus $F(1) = 7.3$, $p = 0.0069$; week $F(3) = 1.265$, $p = 0.29$; virus and week interaction $F(3) = 0.77$, $p = 0.51$). **(C)** Similar to (B). Distribution of the percentage of exploratory laps across the experimental weeks for DREADDs and control virus animals (ANOVA virus $F(1) = 11.8$, $p < 0.001$; week $F(3) = 3.95$, $p = 0.0084$; virus and week interaction $F(3) = 2.65$, $p = 0.049$). Control virus animals changed their propensity to explore across weeks (One-way ANOVA $F(3) = 3.18$, $p = 0.0245$).

Supplementary Table 1: Cell counts by rat and recorded region

Total, average, minimum, maximum, and standard deviation of the number of recorded units for each rat and brain region summed across all the sessions. Also showing the grand total of the number of recorded units for each brain region summed across all sessions and rats, as well as average, global minimum, global maximum, and average standard deviation.

Rat ID	Total Sessions	Region	Total	Avg	Min	Max	Std
R668	16	HC	588	36.75	20	53	9.38
		mPFC	390	24.38	6	39	11.45
R677	16	HC	1145	71.56	49	98	17.50
		DLS	258	16.13	5	27	5.24
R679	16	DLS	1037	64.81	37	90	14.87
		mPFC	1111	69.43	48	93	15.86
R699	16	HC	7	0.44	0	4	1.09
		mPFC	38	2.38	0	7	2.16
R710	16	HC	123	7.69	1	15	4.06
		DLS	298	18.63	10	24	4.27
R721	16	HC	999	62.44	36	93	16.41
		mPFC	111	6.94	4	12	2.05
R786	16	HC	1038	64.88	23	100	20.31
		mPFC	583	36.44	31	43	3.97
R787	16	HC	392	24.50	9	38	9.04
		mPFC	254	15.88	8	31	7.49
R788	16	DLS	700	43.75	32	59	8.20
		mPFC	122	7.63	0	24	7.34

Regions	Total	Avg	Min	Max	Std
HC	4292	38.32	19.71	57.29	11.11
DLS	2293	35.83	21	50	8.15
mPFC	2609	23.30	13.86	35.57	7.19