Structured navigation emerges from self-guided spatial learning in freely moving common marmosets

Nada El Mahmoudi[1]#*, Francesca Lanzarini[1]#*, Farzad Ziaie Nezhad[1], Deepak Surendran[1], Jean Laurens[1]*

1- Ernst Strüngmann Institute (ESI) for Neuroscience in Cooperation with Max Planck Society, Frankfurt am Main, 60528, Germany

Co-first authors

*Corresponding authors: nada.el-mahmoudi@univ-amu.fr; francesca.lanzarini@esi-frankfurt.de; jean.laurens@esi-frankfurt.de;

Abstract

How do structured, memory-guided behaviors emerge in freely moving primates? We addressed this question by training common marmosets to perform a foraging task where they retrieved food from eight dispensers in a semi-naturalistic environment, over multiple sessions. Animals didn't receive any instruction or shaping and were free to visit dispensers in any sequence. We found that animals engaged with the task in most trials, and, within engaged trials, transitioned from exploratory to efficient behavior. Learning was marked by performance gains and the emergence of reproducible movement patterns between specific dispensers. Through trajectory analyses and probabilistic modeling, we showed that these patterns reflected the consolidation of local transition motifs characterized by increased regularity and reduced spatial entropy. These patterns were flexibly recombined into variable global sequences. Thus, animals exhibited a hybrid strategy based on route segments flexibly embedded within a topological structure. These findings demonstrate how efficient, adaptive navigation can emerge through self-guided experience in complex environments. Our approach provides a naturalistic and longitudinal framework for studying the formation of structured spatial strategies in non-human primates, bridging ecological behavior with theoretical models of learning and memory.

Keywords

Common marmosets, Naturalistic Behavior, Spatial memory, Navigation strategies, Self-guided learning, Topological organization, Behavioral structure, Non-Human Primates, Neuroethology

Introduction

How structured behavior emerges from repeated experience in complex environments is a fundamental question in neuroscience, cognition, and behavioral ecology. Navigation provides a powerful model to address this question, as it requires animals to extract, organize, and flexibly use spatial information across multiple timescales and contexts.

Although the cognitive and neural foundations of navigation have been extensively studied in rodents, how these mechanisms operate in non-human primates (NHPs) remains poorly understood. Primates differ markedly from rodents in their sensorimotor capacities, brain architectures, and ecological demands, particularly under ethologically valid conditions (Kell et al., 2023; Krakauer et al., 2017; Martinez-Trujillo, 2025; Trapanese et al., 2019; Zhang et al., 2023).

Spatial memory supports an animal's ability to remember where it has been, and to use this information to guide future movement (Olton, 1977). It is classically divided into spatial working memory, supporting short-term retention during ongoing behavior, and spatial reference memory, which stores stable knowledge across repeated experience (Olton et al., 1979). Together, these systems support a range of navigation strategies, defined by how spatial information is encoded and used to guide behavior (Arleo and Rondi-Reig, 2007; O'Keefe and Nadel, 1978; Redish, 1999): route-based navigation relies on learned sequences of subgoals, topological navigation involves the representation of interconnected spatial nodes, and metric navigation enables flexible rerouting based on distances and directions (Parra-Barrero et al., 2023; Poucet, 1993).

NHPs display remarkable spatial flexibility and foraging efficiency (Trapanese et al., 2019), yet the mechanisms by which such strategies emerge and evolve over time remain elusive. Experimental approaches to primate navigation have traditionally relied on simplified laboratory tasks, which limit ecological validity, or on observational field studies, which preclude high-resolution behavioral tracking. As a result, we lack a detailed understanding of how spatial learning unfolds in freely moving primates.

The common marmoset (*Callithrix jacchus*) provides a powerful model to address this gap. These arboreal primates navigate three-dimensional environments to access both stable and unpredictable food sources, such as exudates, fruits, and insects (Abreu et al., 2020, 2016; Garber, 1992; Schiel and Souto, 2017; Xavier et al., 2024). In the wild, they follow repeatedly used foraging routes embedded within topological maps, enabling flexible yet structured navigation (Abreu et al., 2021). However, how such strategies are acquired and optimized over time, and whether similar processes can be observed under controlled yet ethologically relevant conditions, remains unknown. To address this, we developed a semi-naturalistic foraging task that combines ecological richness with longitudinal tracking of self-guided behavior.

Here, we show how freely moving marmosets progressively structured their navigation through repeated experience in a semi-naturalistic environment. By combining behavioral clustering, trajectory analysis, and probabilistic modeling, we uncover how structured yet flexible navigation strategies emerge through self-guided learning in primates. Without instruction or shaping, animals transitioned from exploratory to organized behavior via a set of discrete engagement states. With experience, their navigation became increasingly regular, shaped by the repeated use of familiar movement patterns. These patterns were not rigidly replayed but flexibly recombined, revealing a dynamic balance between behavioral consolidation and adaptive variability. Altogether, our findings provide a proof of concept for using naturalistic behavioral paradigms to uncover the emergence and consolidation of memory-guided navigation strategies in non-human primates. This framework bridges naturalistic observation and theoretical models of spatial learning, offering new insight into the cognitive foundations of primate navigation.

Results

We implemented a spatial memory task in an enriched, semi-naturalistic environment to study navigation behavior in common marmosets (**Fig. 1A**). Animals moved freely within the arena and collected rewards from eight food dispensers arranged around its perimeter (**Fig. 1E**). Each trial began when the central "Home Dispenser" was activated, signaled by auditory and visual cues along with reward delivery. Upon visiting this dispenser, all eight peripheral dispensers became simultaneously active and dispensed reward (**Fig. 1B**). A trial was completed when the animal had visited all eight unique dispensers, regardless of any revisits along the way. If this condition was not met within 10 minutes, the trial ended automatically. The next trial began with the reactivation of the Home Dispenser. Each subject engaged with the task over 30 sessions, performing on average 7.96 ± 0.4 trials per session for Animal E and 8 ± 0.4 for Animal F.

Marmosets show stable dispenser interactions and distinct learning dynamics across error types

To assess marmoset interactions with the dispensers, we analyzed the total number of visits and omissions (i.e., dispensers not visited within a trial) for each dispenser (**Fig. 1C-D**). A Chi-squared test revealed significant deviations from uniformity in the distribution of visits for both Animal E ($\chi^2 = 15.47$, df = 7, $p = 0.03$) and Animal F ($\chi^2 = 33.6$, df = 7, $p < 0.001$). Post hoc binomial tests with Bonferroni correction, however, indicated that Animal E visited all dispensers at comparable rates ($p > 0.05$), while Animal F visited Dispenser 1 significantly less often than expected ($p < 0.05$), with no significant differences among the remaining dispensers. For omissions, neither animal showed significant deviations from uniformity (Animal E: $\chi^2 = 8.9$, df = 7, $p = 0.26$, Animal F: $\chi^2 = 9.89$, df = 7, $p = 0.19$), suggesting no systematic neglect of specific locations. Together, these results indicate that dispenser interactions were globally stable and well-balanced across the arena, providing a robust basis for analyzing how spatial behavior and error patterns evolve over time.

To evaluate task performance, we first classified the errors observed during the trials into three types: *repetition errors* (revisiting a dispenser already visited during the same trial), *omission errors* (failing to visit all eight dispensers), and *rule errors* (returning to the home dispenser mid-trial). We then computed the mean number of total errors per session and assessed how this measure evolved over time using linear regression (**Fig. S1A**). Both animals exhibited a significant decrease in error rate across sessions, reflecting a spontaneous improvement in performance (Animal E: $R^2 = 0.26$, $p < 0.01$, $\beta = -0.07$, Animal F: $R^2 = 0.62$, $p < 0.001$, $\beta = -0.28$). To assess potential contributors to this improvement, we used a multiple linear regression model (**Table S1**), which confirmed that the number of sessions was the main predictor of performance gains, with no detectable contribution from other variables. To further describe the evolution of different error types over time, we compared their session-wise dynamics using a Kruskal-Wallis test (**Fig. S1B**). This analysis revealed that the temporal profiles of repetition, omission, and rule errors differed significantly (Animal E: $H = 13.06$, $p < 0.01$, Animal F: $H = 39.05$, $p < 0.001$), suggesting that the different error types did not follow the same trajectory across training. Finally, to examine how error types related to each other and to overall performance, we conducted a correlation matrix analysis (**Fig. S1C-D**). Repetition and rule errors were positively correlated with each other and with the total number of visits, suggesting they tended to co-occur in more exploratory or

disorganized trials. In contrast, omission errors were uncorrelated with both repetition and rule errors, suggesting that they may reflect distinct behavioral processes.

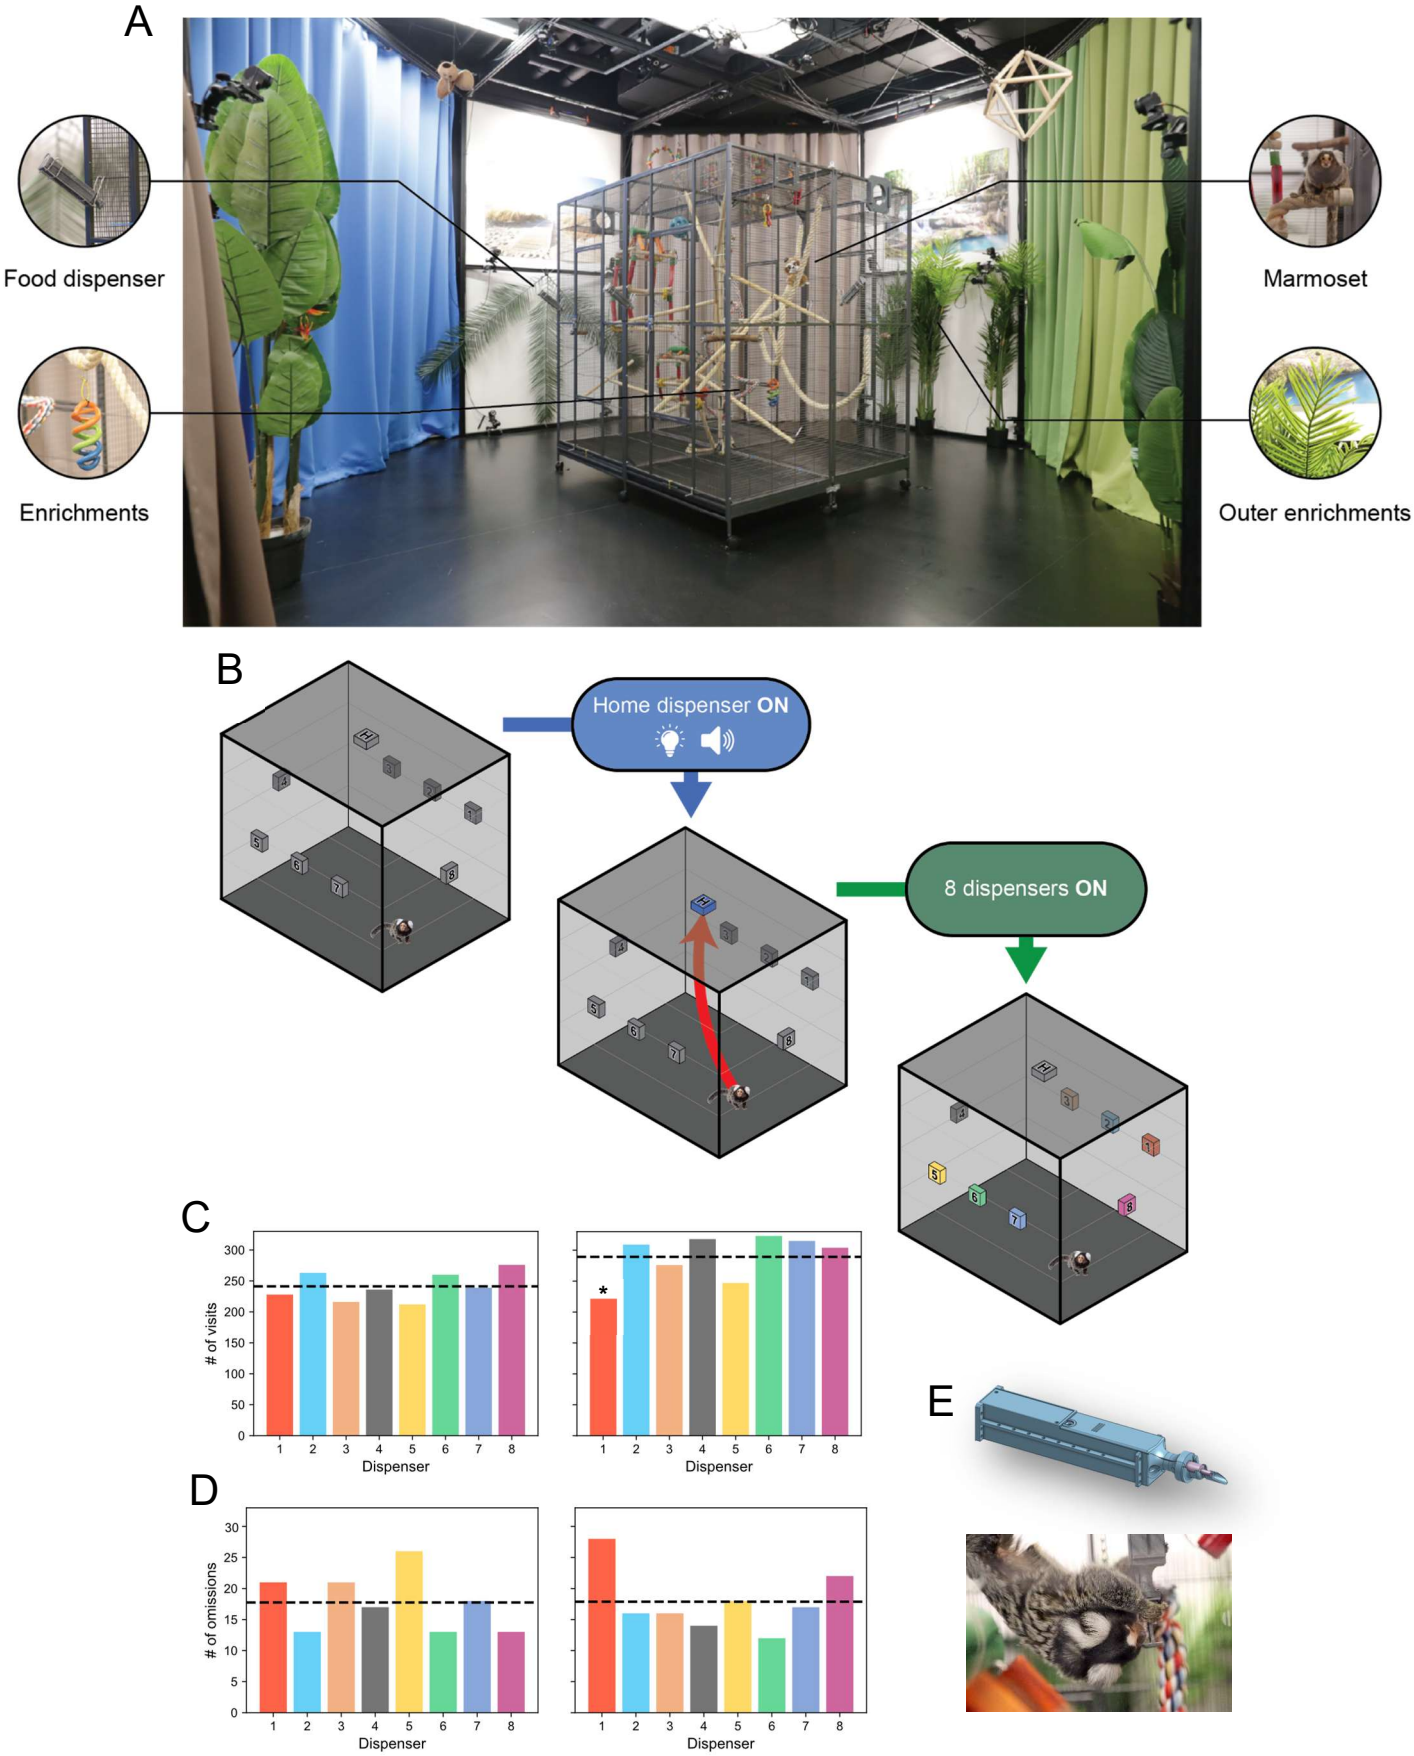

Figure 1: A freely moving setup for investigating spatial memory in marmosets
A. Experimental setup. The enclosure consisted of a wire mesh cage (1.8 m × 1.4 m × 1.6 m; 4 m^3) that allowed the marmosets to climb and move freely. Structural enrichments (wooden sticks, semi-rigid ropes, and various objects) were included to promote locomotion and exploration. External spatial cues were provided by four white walls decorated with distinct landscape posters, four colored curtains, and large plastic plants. An action camera was mounted above the setup to record behavior. A centrally mounted food dispenser, designated the Home Dispenser (HD), was placed on the ceiling, while eight peripheral food dispensers were evenly distributed around the perimeter of the cage to deliver rewards during the task.
B. Trial design. At the start of each trial, the HD was activated with an LED and a buzzer. Once the marmoset reached the HD, all eight peripheral dispensers were simultaneously activated and delivered gum arabic. To complete the trial, the animal had to visit each of the eight dispensers at least once. If not completed within 10 minutes, the trial was automatically terminated.
C. Total visits per dispenser. Bar plots showing the total number of visits to each of the eight dispensers for Animal E (*left*) and Animal F (*right*). The dashed horizontal line indicates the expected number of visits under a uniform distribution. A chi-squared test assessed the overall distribution, followed by binomial tests comparing each dispenser to the expected proportion. Significant p-values (after Bonferroni correction) are indicated with an asterisk (*$p < 0.05$).
D. Total omissions per dispenser. Bar plots showing the total number of omissions for each dispenser in Animal E (*left*) and Animal F (*right*). The dashed horizontal line indicates the expected number of omissions under a uniform distribution. A chi-squared test was used to assess deviation from uniformity.
E. Reward dispenser design. Design of the custom-made food dispenser used during the experiment (*top*). Eight dispensers were placed around the setup, while the HD was placed at the center of the ceiling upside-down (*bottom*, a marmoset eating from the HD).

Clustering analysis reveals different task execution profiles

Based on this dissociation, we applied an unsupervised k-means clustering analysis to capture trial-by-trial variability in task execution. We used omission errors and total visit counts as two complementary metrics, indexing disengagement and activity level, respectively. Although only moderately correlated (**Fig. S1C-D**), these measures captured distinct behavioral dimensions, enabling the identification of discrete behavioral states.

The optimal number of clusters ($k = 4$) was determined using silhouette score and inertia, which quantify within-cluster coherence and between-cluster separation (**Fig. 2A**). For both animals, clustering produced coherent and interpretable partitions (Animal E: silhouette = 0.73, inertia = 41.9, Animal F: silhouette = 0.68, inertia = 56.4), reflecting consistent behavioral variability across trials. Cluster centroids are reported in **Table S2**. To interpret the behavioral significance of each cluster, we compared key performance metrics across them (**Fig. S3**):

- **Cluster 0** exhibited the fewest errors and shortest trial durations ($p < 0.001$ for both animals vs. all other clusters), reflecting a state of optimal performance marked by efficient and precise behavior.
- **Cluster 1** was associated with higher visit counts and increased rule and repetition errors ($p < 0.001$), but few omissions, suggesting active but imprecise task engagement.
- **Cluster 2** showed the highest omission rates and no dispenser visits, consistent with disengagement or lack of task initiation.
- **Cluster 3** reflected partial engagement, with fewer visits and more omissions than Clusters 0 and 1 ($p < 0.001$), but slightly more visits and fewer omissions than Cluster 2 (non-significant), indicating incomplete task execution.

Cluster distribution was remarkably consistent between animals: Clusters 0 and 1 together accounted for the vast majority of trials (89% for Animal E, 90% for Animal F), while Clusters 2 and 3 comprised only 10% each, suggesting that most trials reflected engaged task performance (**Fig. 2A**).

We next assessed how cluster distributions evolved over time (**Fig. 2B-C**). Linear regressions showed that the proportion of Cluster 0 trials increased with session number (Animal E: $R^2 = 0.13$, $p = 0.05$, Animal F: $R^2 = 0.45$, $p < 0.001$), while Cluster 1 decreased (Animal E: $R^2 = 0.24$, $p < 0.01$, Animal F: $R^2 = 0.48$, $p < 0.001$), reflecting a gradual shift from error-prone to optimized task

execution. In contrast, the frequency of Clusters 2 and 3 remained stable over time ($p > 0.1$), suggesting that disengagement levels were not affected by session progression.

Altogether, these findings support a functional distinction between clusters reflecting task-related engagement (Clusters 0 and 1) and those associated with limited or absent task participation (Clusters 2 and 3), enabling the identification of distinct behavioral states during task performance.

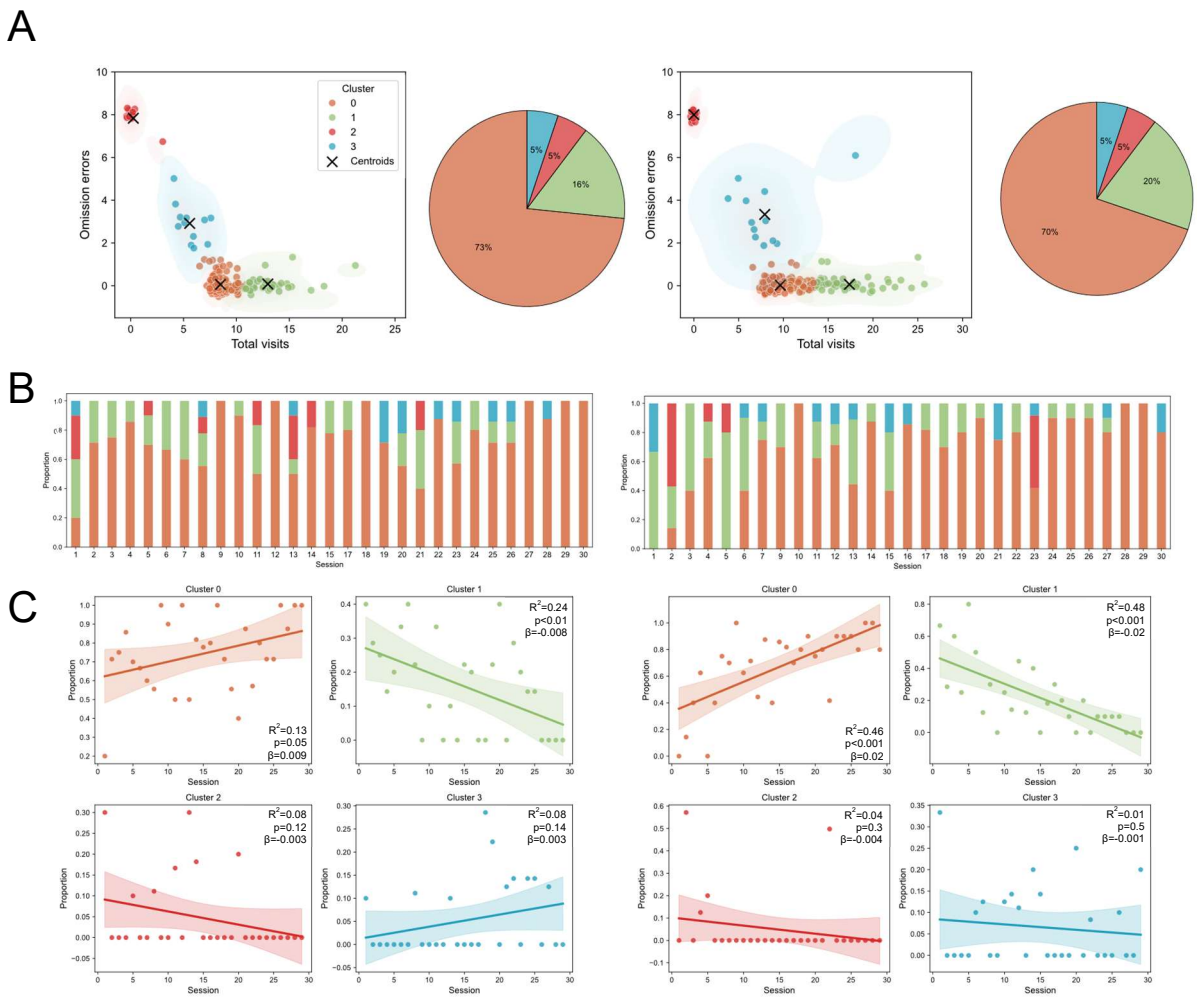

Figure 2: Clustering analysis reveals different task execution profiles

A. Unsupervised clustering analysis. Scatter plots showing the outcome of the k-means clustering analysis for Animal E (*left*) and Animal F (*right*), with each trial assigned to one of four clusters (color-coded). Cluster centroids, defined by omission errors and total number of visits, are marked with crosses. Data points are jittered for visibility. Pie charts indicate the proportion of trials assigned to each cluster. Colors are consistent across panels: Cluster 0 (orange), Cluster 1 (green), Cluster 2 (red), and Cluster 3 (blue).

B. Cluster distributions across sessions. Stacked bar plots showing the normalized proportion of trials in each cluster across training sessions for Animal E (*left*) and Animal F (*right*). Each bar represents one session, and segment height indicates the proportion of trials falling into each cluster.

C. Temporal evolution of cluster proportions. Scatter plots showing the proportion of trials in each cluster across sessions. Dashed lines represent regression fits; shaded areas indicate 95% confidence intervals. R^2, p-values, and regression slope (β) quantify the strength and direction of the association.

Clusters 0 and 1 reflect variations in efficiency within a shared behavioral framework

Since Clusters 0 and 1 were associated with task engagement, we next asked whether they reflected distinct navigation strategies or simply differences in execution efficiency. We focused on the first eight choices within each trial, as successful task completion required visiting eight unique dispensers without repetition.

We quantified the performance rate as the percentage of correct responses within the first eight choices of each trial (**Fig. 3A, H**), where a response was considered correct if the animal visited a dispenser that had not yet been visited during that trial. To determine whether performance deviated from random behavior, we compared performance rates to an expected chance level of 50%, based on uniform random choice among all dispensers at each step (see Methods). Wilcoxon signed-rank test with Bonferroni correction confirmed that both clusters exhibited significantly greater performance rates than chance ($p < 0.001$, for both animals), confirming that even in Cluster 1, despite lower precision, marmosets showed goal-directed behavior.

To further explore how they navigated the environment, we analyzed their transitions between dispensers within a trial, i.e., how they moved from one dispenser to another. By reconstructing sequences of dispenser visits, we captured the order in which animals visited dispensers. We quantified the global organization of transitions using two complementary measures: global entropy, capturing the uncertainty in transition patterns and transition diversity, reflecting the evenness of transition use. In our framework, high entropy reflects unpredictable, wide-ranging transitions, while low entropy indicates a reliance on a smaller set of recurrent transitions. Similarly, high diversity corresponds to evenly distributed transitions, whereas low diversity signifies a preference for specific paths. For both animals, entropy and transition diversity were significantly lower than chance ($p < 0.001$, bootstrap with 1,000 random transitions, **Fig. 3B-C, I-J),** indicating structured and non-random transitions in both clusters.

To visualize these patterns, we built transition matrices and transition maps. Global transition maps showed all transitions, with arrow thickness indicating transition probabilities (**Fig. 3D-E, K-L**). Significant transition maps (based on binomial tests with Bonferroni correction) highlighted transitions occurring more frequently than expected by chance (binomial tests with Bonferroni correction, **Fig. 3F-G, 3M-N**). Correlation analyses revealed a strong alignment in transition structure between clusters (global matrices: Animal E: $r=0.89$ $p<0.001$, Animal F: $r=0.91$ $p<0.001$, significant matrices: Animal E: $r=0.8$ $p<0.001$, Animal F: $r=0.96$ $p<0.001$). Yet, Euclidean distances indicated differences in absolute transition probabilities (global matrices: Animal E: 0.90, Animal F: 0.84, significant matrices: Animal E: 1.92, Animal F: 0.72), suggesting that while the structure of transitions was conserved, the specific weights assigned to each path varied between clusters.

Altogether, these findings indicate that Clusters 0 and 1 share a common navigational structure, differing primarily in execution efficiency rather than in strategy type. Based on this similarity, we pooled Clusters 0 and 1 for subsequent analyses of spatial learning.

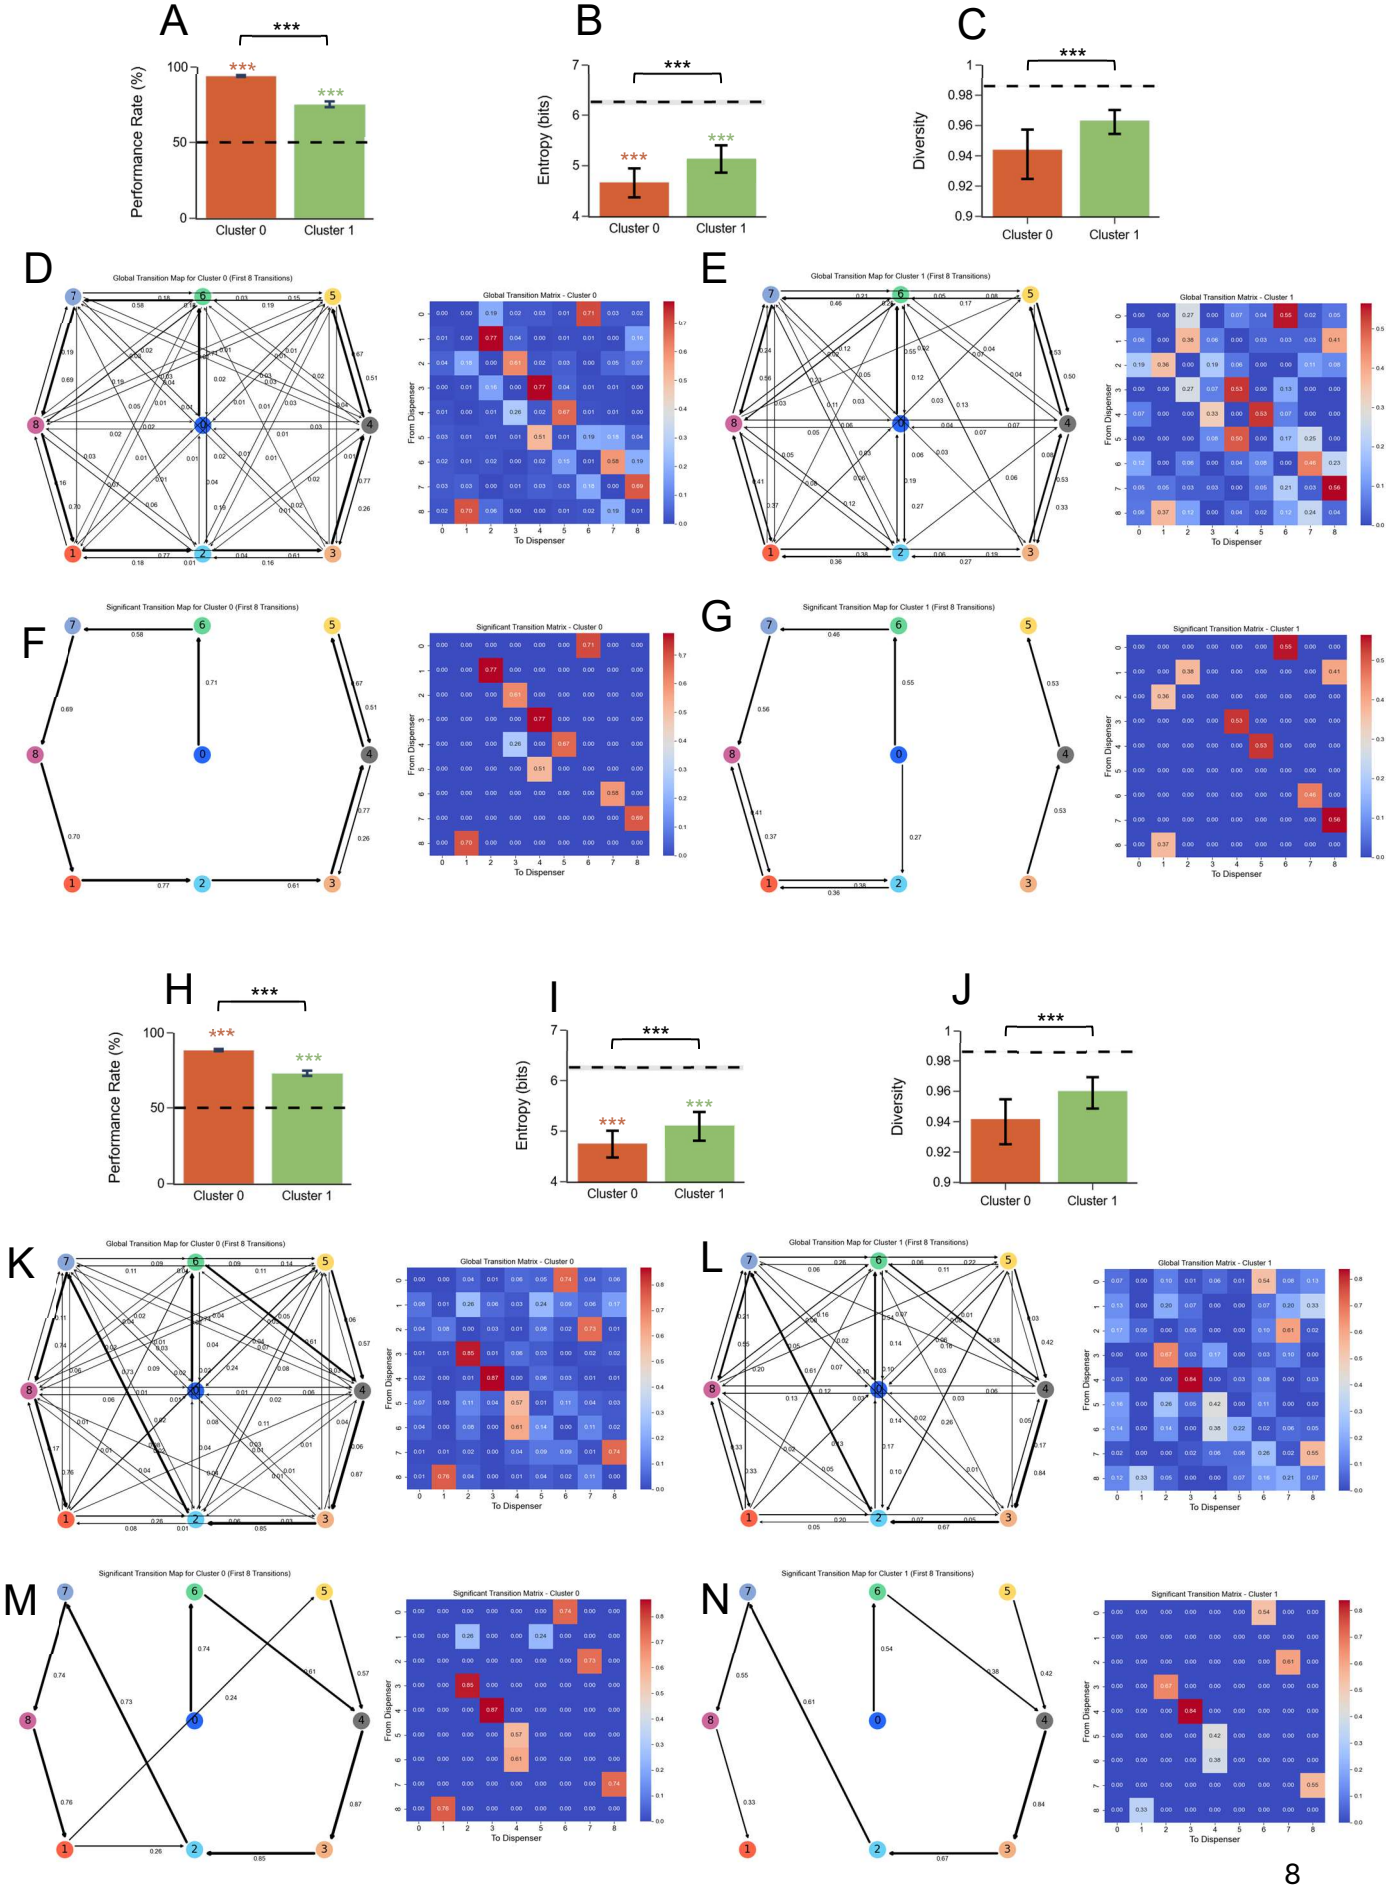
A
B
C
D
E
F
G
H
I
J
K
L
M
N
Performance Rate (%)
Entropy (bits)
Diversity
Cluster 0
Cluster 1
Global Transition Map for Cluster 0 (First 8 Transitions)
Global Transition Matrix - Cluster 0
Global Transition Map for Cluster 1 (First 8 Transitions)
Global Transition Matrix - Cluster 1
Significant Transition Map for Cluster 0 (First 8 Transitions)
Significant Transition Matrix - Cluster 0
Significant Transition Map for Cluster 1 (First 8 Transitions)
Significant Transition Matrix - Cluster 1
From Dispenser
To Dispenser

Figure 3: Clusters 0 and 1 reflect variations in efficiency within a shared behavioral framework
(A–G) Analyses for Animal E; (H–N) Corresponding analyses for Animal F.
A and H. Mean performance in the first eight choices for Clusters 0 and 1. Bar plots show the percentage of correct responses (i.e., visits to non-repeated dispensers) within the first eight choices of each trial. Error bars indicate SEM. The dashed line marks the 50% chance level (random choice among nine options). Wilcoxon signed-rank tests (vs. chance) and Mann-Whitney U tests (between clusters) were performed, with Bonferroni correction. ***p < 0.001.
B and I. Global entropy across Clusters 0 and 1. Bar plots display mean global entropy (in bits), with 95% bootstrap confidence intervals (CI). Dashed lines represent mean entropy expected under random transitions; shaded areas indicate 95% CI of random baseline. Permutation tests with Bonferroni correction assessed significance. ***p < 0.001.
C and J. Global transition diversity across Clusters 0 and 1. Same conventions as in B and I. Bars represent mean transition diversity; dashed lines and shaded areas reflect the random baseline. ***p < 0.001.
D and K. Global transition maps and matrices for Cluster 0. *Left:* Directed graphs represent all observed transitions within the first eight choices. Node positions correspond to dispenser locations and arrow thickness is proportional to transition probability. *Right*: transition matrices show the probability of moving between all dispenser pairs, with color intensity reflecting transition strength (dark red: high, dark blue: low).
E and L. Global transition maps and matrices for Cluster 1. Same conventions as in D and K.
F and M. Statistically significant transition maps and matrices for Cluster 0. *Left:* Directed graphs show transitions that occurred significantly more often than expected under a uniform random model (p = 1/9; binomial test with Bonferroni correction). *Right*: corresponding transition matrices highlight these statistically significant transitions. Arrow thickness and color intensity reflect transition frequency, with thicker arrows and warmer colors indicating higher occurrence rates.
G and N. Statistically significant transition maps and matrices for Cluster 1. Same conventions as in F and M.

Evolution of performance and global transition pattern across training phases

We next examined whether improvements in task performance were accompanied by changes in the global organization of navigation behavior. To do so, we divided the 30 training sessions into six consecutive phases of five sessions each: early training (Beg 1: sessions 1-5, Beg 2: sessions 6-10), intermediate training (Int 1: sessions 11-15, Int 2: sessions 16-20), and late training (Last 1: sessions 21-25, Last 2: sessions 26-30).

As expected, performance, quantified as the percentage of correct responses within the first eight choices, significantly increased across training phases (**Fig. 4A, F**, Kruskal-Wallis test, Animal E: H = 11.66, p = 0.03, Animal F: H = 64.6, p < 0.001).

In parallel, global transition entropy within the first eight choices significantly decreased over training (**Fig. 4B, G**, permutation test, p < 0.001 for all comparisons), indicating a progressive emergence of more structured behavioral patterns. Moreover, this reduction in entropy was strongly associated with improved performance (**Fig. 4C, H**): Pearson correlation analyses revealed a robust negative relationship between mean global entropy and performance across training phases (Animal E: r = -0.88, p = 0.04, Animal F: r = -0.93, p = 0.01).

To further characterize this evolution, we calculated the overlap index, a measure of similarity between entropy distributions across training phases (0 = no overlap, 1 = complete overlap, **Fig. S4A, F**). Early phases exhibited low overlap (Beg 1 vs. Int 2: Animal E = 0.15, Animal F = 0.004), reflecting a marked reorganization of transition patterns early in training. In contrast, overlap increased between later phases (Int 2 vs. Last 2: Animal E = 0.7, Animal F = 0.5), suggesting stabilization and refinement of navigation strategies over time.

Global (**Fig. 4D, I**) and significant (**Fig. 4E, J**) transition maps further illustrated this refinement. Permutation-based correlation analyses (Bonferroni-corrected) confirmed that the overall structure of transition matrices remained consistent across training phases (**Fig. S4B, G**). However, Euclidean distances between matrices revealed substantial changes in individual transition probabilities, which progressively decreased with training, indicating reduced variability in transition patterns.

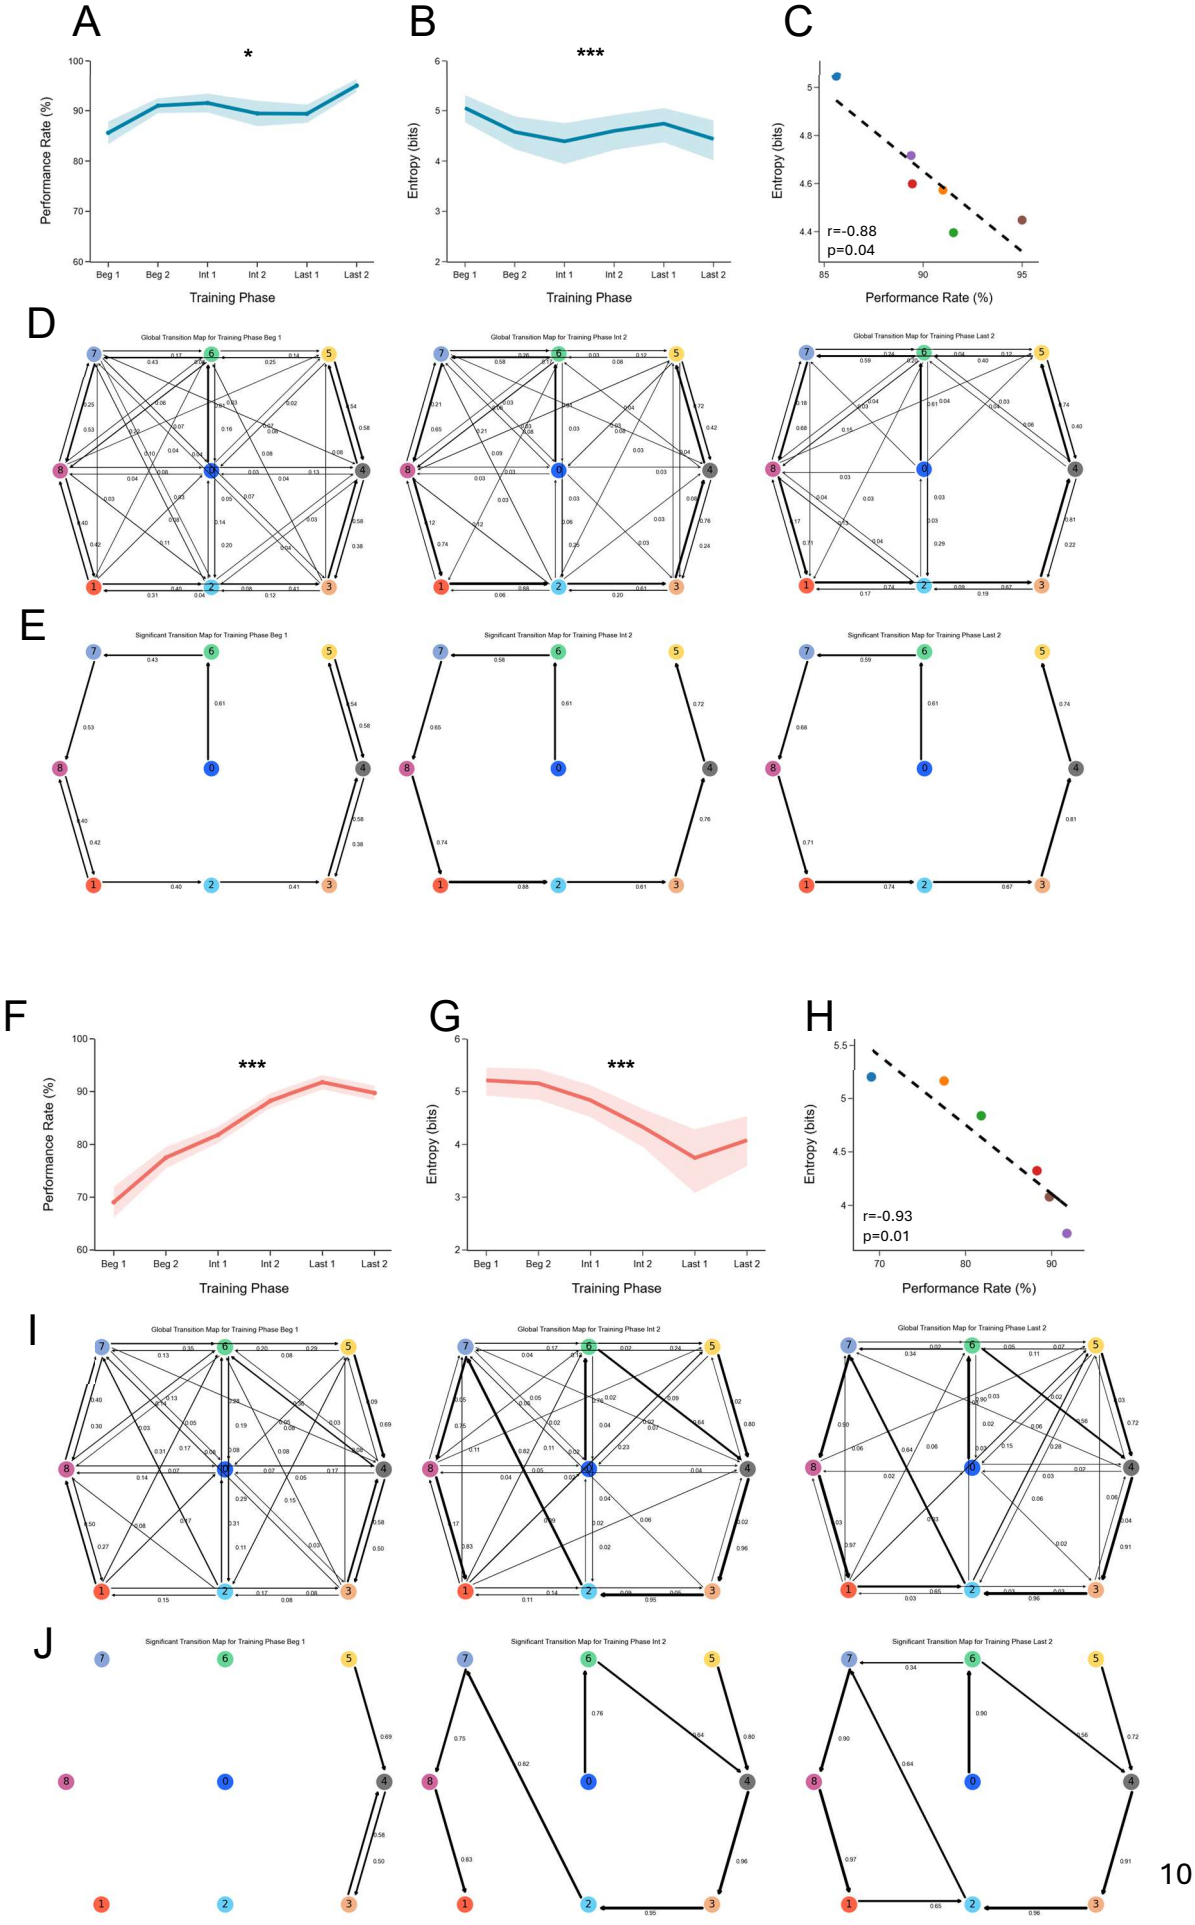
A
B
C
D
E
F
G
H
I
J
*
***
Performance Rate (%)
Entropy (bits)
Training Phase
Beg 1
Beg 2
Int 1
Int 2
Last 1
Last 2
r=-0.88
p=0.04
r=-0.93
p=0.01
Global Transition Map for Training Phase Beg 1
Global Transition Map for Training Phase Int 2
Global Transition Map for Training Phase Last 2
Significant Transition Map for Training Phase Beg 1
Significant Transition Map for Training Phase Int 2
Significant Transition Map for Training Phase Last 2

Figure 4: Evolution of performance and transition pattern across training phases.
(A–E) Analyses for Animal E; (F–J) Corresponding analyses for Animal F.
A and F. Performance across training phases.
Line plots showing the mean performance (% correct choices within the first 8 visits per trial) across six training phases. Solid lines represent the mean; shaded areas indicate the standard error of the mean (SEM). Kruskal-Wallis test followed by Dunn's post hoc test with Bonferroni correction was used for statistical comparisons. Significant results are marked: ***p < 0.001, **p < 0.01, *p < 0.05.
B and G. Global entropy across training phases.
Line plots showing the evolution of mean global entropy (in bits) over training. Points represent phase means; shaded areas represent 95% confidence intervals. Statistical analysis used permutation tests with Bonferroni correction. Asterisks indicate significant phase-wise differences.
C and H. Correlation between performance and entropy.
Scatter plots showing the relationship between performance rate and global entropy across phases. Each point represents a training phase, color-coded accordingly. Dashed black lines indicate linear regression fits. Pearson correlation coefficients are reported, with Bonferroni correction for multiple comparisons.
D and I. Global transition maps.
Directed graphs showing transition patterns across training phases (Beg 1: sessions 1–5; Int 2: sessions 16–20; Last 2: sessions 26–30). Node positions match dispenser locations; arrow thickness represents transition probabilities. Thicker arrows indicate more frequently used transitions.
E and J. Statistically significant transition maps.
Directed graphs showing transitions that occurred significantly more often than expected by chance (uniform model, p = 1/9). Significance assessed via binomial tests with Bonferroni correction. Maps are shown for Beg 1, Int 2, and Last 2. Thicker arrows highlight dominant, structured transitions.

Altogether, these results show that performance improvements were accompanied by a gradual reduction in transition entropy and variability, reflecting the emergence and consolidation of consistent, structured navigation patterns across training.

Local transition analysis uncovers dispenser-specific variability within a globally structured transition behavior

To determine whether the global structuring of transitions also emerged at finer spatial scales, we examined local entropy (e.g. uncertainty in the pattern of transitions leaving a given dispenser) patterns at the level of individual dispensers (**Fig. S4C, H**). For both animals, local entropy significantly decreased over training (Kruskal-Wallis, Animal E: H = 12.6, p = 0.02, Animal F: H = 25.67, p < 0.001), indicating increasing transition regularity at the dispenser level. Post hoc comparisons confirmed significant reductions between early (Beg 1) and late (Last 2) training phases (Animal E: p = 0.04, Animal F: p = 0.04).

To evaluate the behavioral relevance of this local structuring, we computed the probability of making an error after visiting each dispenser (**Fig. S4D, I**). For both animals, error probability decreased significantly between early (Beg 1) and late (Last 2) training phases (Kruskal-Wallis, Animal E: H = 6.5, p = 0.01, Animal F: H = 7.6, p < 0.01). Furthermore, we found a strong positive correlation between mean local entropy and mean error probability across phases (Animal E: r = 0.90, p < 0.01, Animal F: r = 1.0, p < 0.001, **Fig. S4E, J**), indicating that lower local entropy was closely associated with improved performance.

We next assessed spatial heterogeneity in local structuring by comparing entropy values across dispensers to a baseline generated from 1,000 bootstrap-sampled random transitions (mean = 2.65 bits, 95% CI: 2.15-2.98). All dispensers showed entropy values below this baseline (**Fig S4C,H**), confirming structured, non-random transitions. However, the degree of structuring varied across dispensers, as reflected by the standard deviation of local entropy values (Animal E: 0.29 bits, Animal F: 0.35 bits), indicating spatial heterogeneity in transition predictability.

Altogether, these results show that transition structuring develops not only globally but also at finer spatial scales. While overall entropy declines with training, individual dispensers vary in their

contribution to this structuring, suggesting that some locations support more stereotyped transitions, whereas others retain greater flexibility.

Marmoset transition patterns show local predictability but global flexibility

We next asked whether the increasing structuring of transition patterns over time reflected the emergence of rigid, repetitive strategies, or instead more flexible, probabilistic preferences. Specifically, we examined whether successful performance in error-free trials consistently relied on the same transition patterns, or whether multiple strategies could lead to success, suggesting variability in execution despite overall structure.

To test this, we applied a first-order Markov chain model to assess the predictability of transitions at both local and global levels. Markov models estimate the probability of moving from one state (dispenser) to another based solely on the current state, and are well suited to capture sequential dependencies. This allowed us to evaluate the model's ability to predict individual transitions (local level) as well as full sequences of dispenser visits (global level). The model was trained exclusively on error-free trials to isolate structured behaviors associated with successful task execution, and its performance was compared to a random baseline model assuming uniform transition probabilities.

The Markov model showed strong performance at the local level (**Fig. 5A-D**), with average prediction accuracy of 79.6% ± 7.8% (95% CI) for Animal E and 84.75% ± 2.7% for Animal F, well above the random baseline (Animal E: 10.5% ± 1.8%, Animal F: 9.0% ± 3.5%). Wilcoxon signed-rank tests confirmed that these differences were statistically significant ($p < 0.01$), indicating that local transitions were highly structured and predictable. To assess global predictability, we used the model to generate full action sequences by iteratively selecting the most probable next transition, starting from the first dispenser visited in each trial (**Fig. 5E-F**). As expected, accuracy declined relative to local prediction, but remained substantially above chance. Average full-sequence accuracy reached 57.1% ± 13.8% (Animal E) and 54.5% ± 6.5% (Animal F), while the random baseline achieved 0% for both animals. These differences were also statistically significant ($p < 0.01$).

Together, these findings indicate that while individual transitions during successful trials are highly predictable, full sequences retain flexibility. The Markov model accurately predicted local transitions but showed reduced accuracy for complete sequences, suggesting that while dominant transition patterns existed, sequence execution was variable across error-free trials.

Predictability of local and global transition patterns evolves dynamically across training phases

While the Markov model revealed both local predictability and global flexibility in error-free trials, how these properties evolved over training remained to be determined. We examined the evolution of transition predictability across training phases by evaluating the Markov model's performance in predicting local transitions and complete sequences within error-free trials. The model's accuracy was compared to a random baseline model in each phase. Kruskal-Wallis tests assessed phase-wise changes, while Wilcoxon signed-rank tests compared model accuracy against baseline within phases (Bonferroni-corrected).

Error-free trials first emerged at different phases in each animal (Animal E: Beg 1, Animal F: Int 1), but once present, both animals showed a similar trajectory in terms of local transition predictability (**Fig. 5G,K, I,M**).

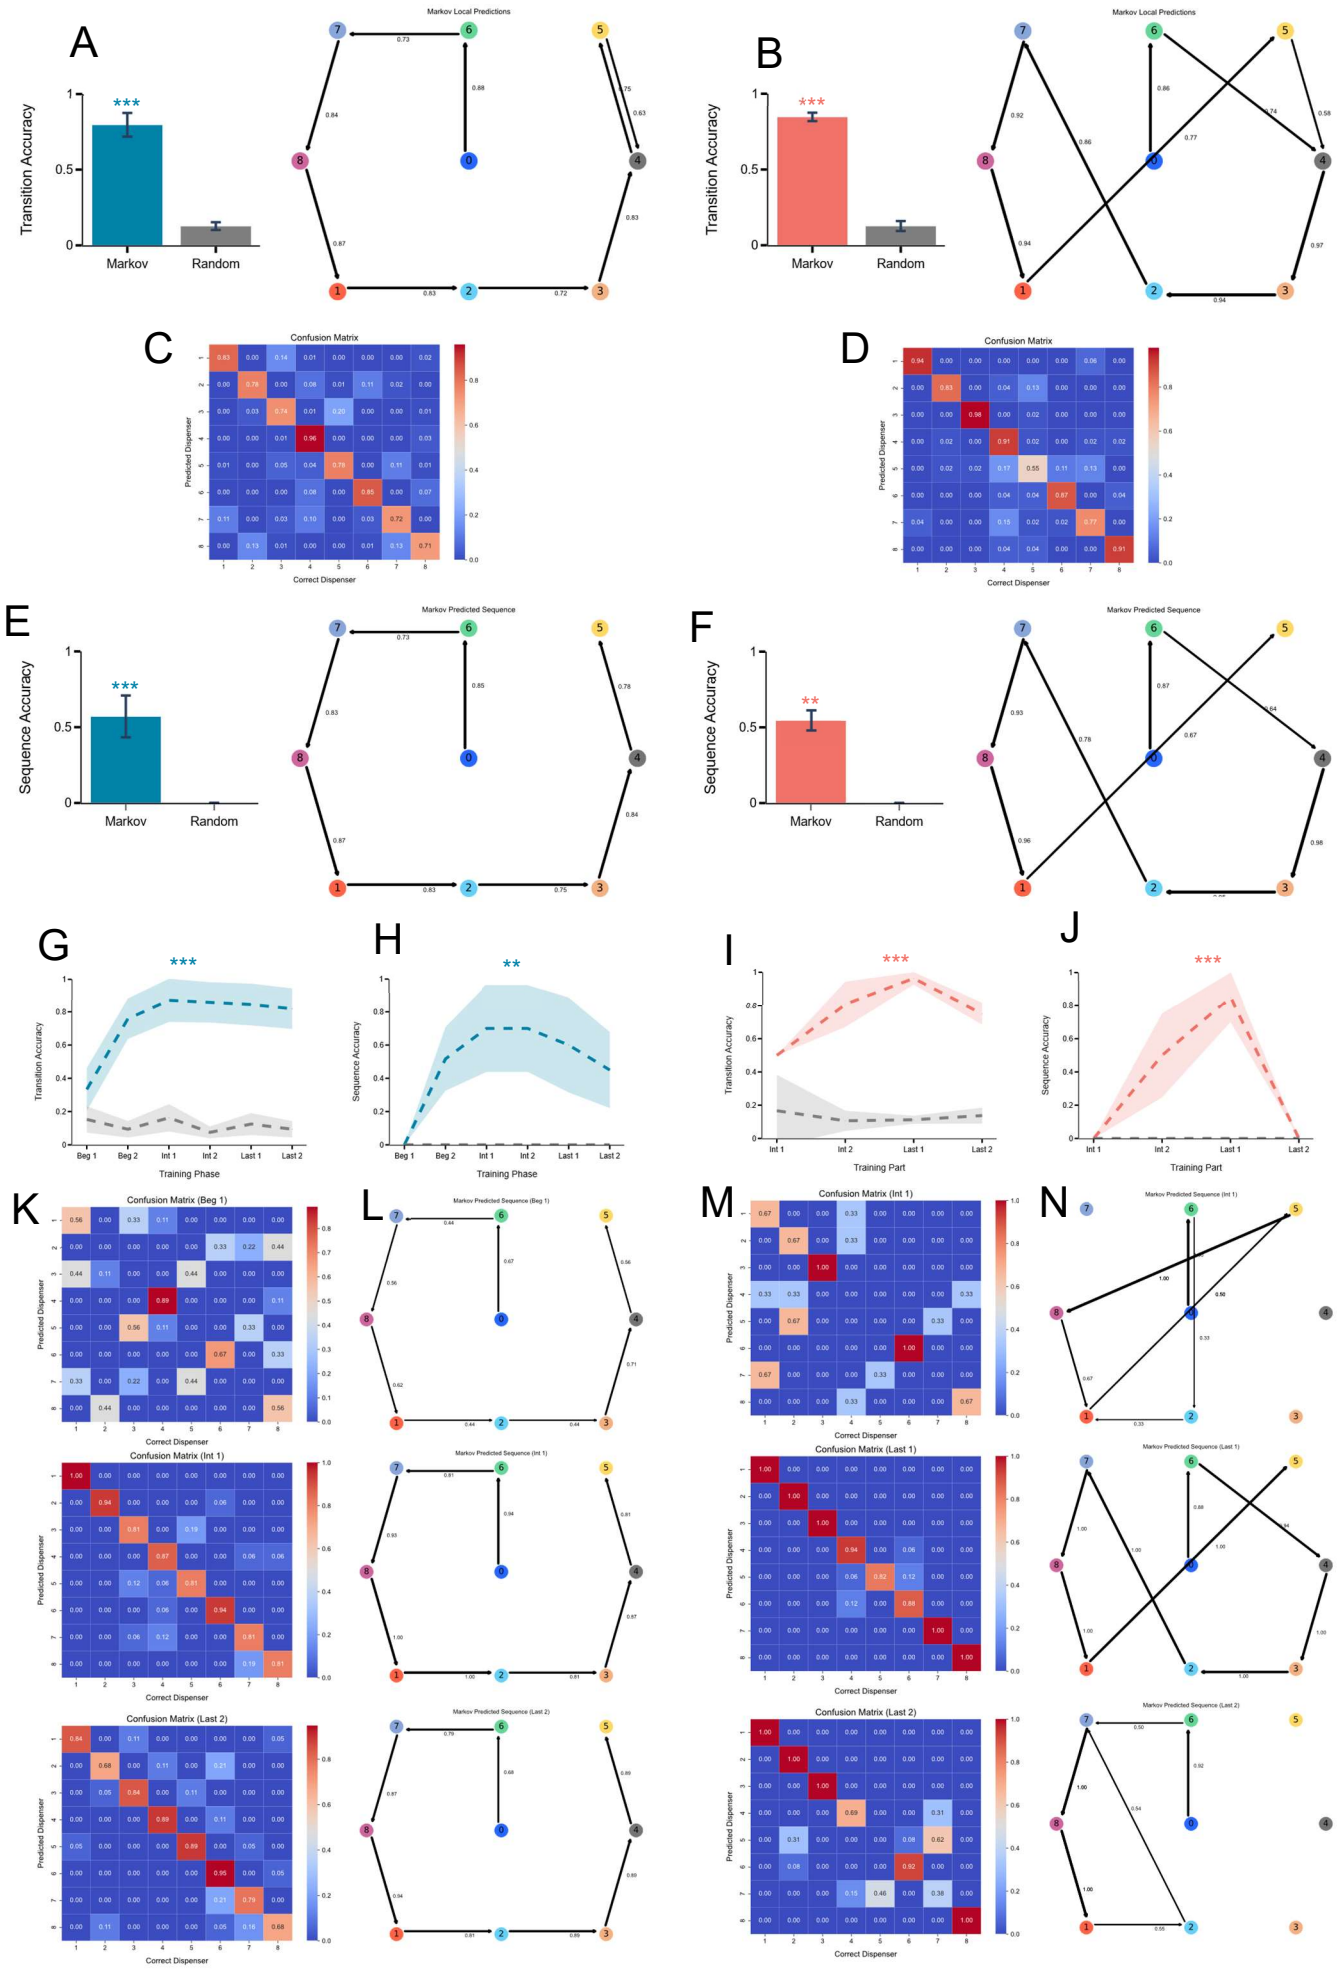
A
Transition Accuracy
Markov
Random
Markov Local Predictions
B
Transition Accuracy
Markov
Random
Markov Local Predictions
C
Confusion Matrix
Predicted Dispenser
Correct Dispenser
D
Confusion Matrix
Predicted Dispenser
Correct Dispenser
E
Sequence Accuracy
Markov
Random
Markov Predicted Sequence
F
Sequence Accuracy
Markov
Random
Markov Predicted Sequence
G
Transition Accuracy
Training Phase
H
Sequence Accuracy
Training Phase
I
Transition Accuracy
Training Part
J
Sequence Accuracy
Training Part
K
Confusion Matrix (Beg 1)
Confusion Matrix (Int 1)
Confusion Matrix (Last 2)
L
Markov Predicted Sequence (Beg 1)
Markov Predicted Sequence (Int 1)
Markov Predicted Sequence (Last 2)
M
Confusion Matrix (Int 1)
Confusion Matrix (Last 1)
Confusion Matrix (Last 2)
N
Markov Predicted Sequence (Int 1)
Markov Predicted Sequence (Last 1)
Markov Predicted Sequence (Last 2)

Figure 5. Markov model accuracy at local and global levels across training phases
A-B. Local transition accuracy and prediction map. *Left:* Bar plots showing the mean local transition accuracy (across cross-validation folds) for the Markov model and the random baseline. Error bars indicate the 95% confidence interval (CI). Statistical comparisons were made using Wilcoxon signed-rank tests with Bonferroni correction; significant p-values are indicated: ***p < 0.001, **p < 0.01, *p < 0.05. *Right:* Directed graphs depicting the most likely next transition from each dispenser based on transition probabilities estimated from training data. Nodes represent dispensers positioned according to their physical layout; arrow thickness reflects transition probability. Animal E (panel A), Animal F (panel B).
C-D. Normalized confusion matrices for local predictions. Confusion matrices showing predicted versus actual dispenser choices. Rows are normalized to reflect prediction distributions for each true dispenser. Diagonal values indicate correct predictions. Warmer colors represent higher frequencies. Animal E (panel C), Animal F (panel D).
E-F. Sequence prediction accuracy and transition maps. *Left:* Bar plots showing mean sequence accuracy across full trials for the Markov model and the random baseline. Error bars represent 95% CIs. Wilcoxon signed-rank tests with Bonferroni correction were used for significance testing. *Right:* Directed graphs illustrating the most probable sequences of up to eight transitions predicted by the Markov model, starting from state 0. Nodes correspond to dispensers, edges indicate predicted transitions. Animal E (panel E), Animal F (panel F).
G-I. Local transition accuracy across training phases. Line plots showing phase-wise Markov model accuracy (dashed lines) compared to the random baseline. Points correspond to mean accuracy per phase; shaded areas indicate 95% CIs. Kruskal–Wallis tests with Dunn's post-hoc comparisons assessed accuracy changes over phases; Wilcoxon signed-rank tests compared models within phases (all Bonferroni-corrected). Animal E (panel G), Animal F (panel I).
H-J. Sequence prediction accuracy across training phases. Same conventions as in panels G and I. Markov model sequence prediction accuracy is shown over training phases alongside the random baseline. Animal E (panel H), Animal F (panel J).
K-M. Confusion matrices across training phases. Confusion matrices for local predictions at three training stages: early (Beg 1: sessions 1–5), intermediate (Int 2: sessions 16–20), and late (Last 2: sessions 26–30). Color scale indicates prediction frequency per actual dispenser. Animal E (panel K), Animal F (panel M).
L-N. Sequence prediction maps across training phases. Directed graphs showing the most probable full sequences predicted by the Markov model across three training phases (Beg 1, Int 2, Last 2). Nodes represent dispensers; arrows indicate predicted transitions. Animal E (panel L), Animal F (panel N).

In early phases, the model did not outperform the random baseline (Animal E: 33.3% ± 12.9%, $p > 0.05$, Animal F: 50.0% ± 0.0%, $p > 0.05$), indicating initially unstructured transitions. Over time, accuracy improved significantly (Kruskal-Wallis: Animal E: $H = 22.3$, $p < 0.001$, Animal F: $H = 16.7$, $p < 0.001$), consistently exceeding the random baseline ($p < 0.05$ for all comparisons), and peaking at 86.8% ± 12.9% (Animal E) and 96.3% ± 3.7% (Animal F). Accuracy then declined slightly in later phases.

We then evaluated global sequence predictability (**Fig. 5H,L, J,N**). Initially, the model failed to predict any full sequences (Animal E: 0%, Animal F: 0%), reflecting highly variable behavior. Sequence accuracy increased significantly with training (Kruskal-Wallis: Animal E: $H = 17.8$, $p < 0.01$, Animal F: $H = 22.5$, $p < 0.001$), peaking at 70% ± 26.1% (Animal E, vs random baseline, $p < 0.01$) and 85% ± 14.96% (Animal F, vs random baseline, $p<0.01$), before sharply declining in the final phase (Animal E: 45% ± 22.8%, Animal F: 0%).

These findings reveal a dissociation in the dynamics of navigation structure: while local transitions became increasingly stereotyped over time, global sequence patterns fluctuated, reflecting the persistence of navigation flexibility despite local consolidation.

Sequences of dispenser visits are variable but share core transitions

To further understand the global-level variability observed in successful trials, we examined the structure and recurrence of complete sequences (**Fig. 6A**). Animal E used 22 unique sequences, and Animal F 11, but these were not used equally: a dominant sequence accounted for 65.1% of successful trials in Animal E and 56.5% in Animal F. To quantify this variability, we computed the normalized Shannon diversity index (range: 0-1), which captures both the number and distribution of unique sequences. The resulting values (Animal E: 0.59, Animal F: 0.66) confirmed moderate variability, indicating that while animals relied heavily on a dominant path, they also explored alternative solutions.

To probe structural similarity across sequences, we extracted all transitions (pairs of consecutive dispensers) and computed the Jaccard similarity index (range: 0-1), between sequences(**Fig. 6B-D**). For Animal E, the mean Jaccard similarity was 0.54 ± 0.006. A majority of sequence pairs (54.7%) shared at least 50% of their transitions, 40.4% fell between 60-70% similarity, and 33.8% were identical. Notably, 7 transitions appeared in over 70% of sequences, reflecting recurring motifs. For Animal F, the pattern was comparable: mean Jaccard similarity reached 0.58 ± 0.01, with 56.9% of sequence pairs sharing ≥50% of transitions, 38.4% between 60-70%, and 32.6% identical. Five transitions were present in over 75% of sequences.

These findings indicate that, while animals occasionally used alternative routes to complete the task, most sequences converged around a core set of transitions. This suggests that even when behavior appears variable at the sequence level, it is underpinned by consistent structural motifs.

Dynamics of sequence diversity and similarity

We next asked when this structural consistency emerged during training. Specifically, we examined how sequence diversity and similarity evolved across training phases to identify when animals began to converge on consistent behavioral motifs.

We first analyzed changes in sequence diversity using the normalized Shannon diversity index (**Fig. 6E**). Both animals showed a non-linear U-shaped trajectory across training. To quantify this pattern, we fitted a quadratic model to the diversity values (**Fig. S5A**). For Animal E, the model accounted for a large proportion of the variance ($R^2 = 0.96$, $p < 0.01$), confirming the significance of the U-shape. For Animal F, the fit was also strong ($R^2 = 0.86$) and outperformed a linear model ($R^2 = 0.4$), although it did not reach significance due to the limited number of phases with error-free trials ($p = 0.3$). Nonetheless, the improvement in fit suggests that a similar non-linear trend likely applies to both animals.

To further probe this trajectory, we examined how the use of the dominant sequence changed across training phases (**Fig. S5B**). In both animals, its frequency followed an inverted U-shaped pattern, initially low, rising to a peak mid-training, then declining in later phases. This trend paralleled changes in diversity: dominant sequence frequency was negatively correlated with the Shannon index (**Fig. S5C**, Animal E: $r = -0.95$, $p < 0.01$, Animal F: $r = -0.95$, $p = 0.05$), confirming that increased reliance on a single sequence coincided with reduced sequence diversity.

We then assessed how sequence similarity evolved by computing the Jaccard index between all sequence pairs within each training phase (**Fig. 6F-G**). In both animals, similarity increased significantly over training (Kruskal-Wallis, Animal E: $H = 50.38$, $p < 0.001$, Animal F: $H = 102.74$, $p < 0.001$), peaking in Int 1 (Animal E: 0.66 ± 0.03) and Last 1 (Animal F: 0.89 ± 0.013). Although similarity declined in the final phase (Dunn tests, Animal E: Int 1 vs Last 2, $p < 0.05$, Animal F: Last 1 vs Last 2, $p < 0.001$), it remained significantly higher than in early training (Animal E: Beg 1 vs Last 2, $p < 0.001$, Animal F: Int 1 vs Last 2, $p < 0.05$). These results indicate that, as training progressed, the sequences used in error-free trials within each phase increasingly shared common transitions, reflecting a growing consistency in the animals' transition patterns.

To assess whether this structural consistency persisted across phases, we computed the similarity between sequences from different training periods. In both animals, late-phase sequences shared a substantial proportion of transitions with earlier ones (**Fig. S5D),** indicating that core transitions were not confined to a specific phase but remained stable over time.

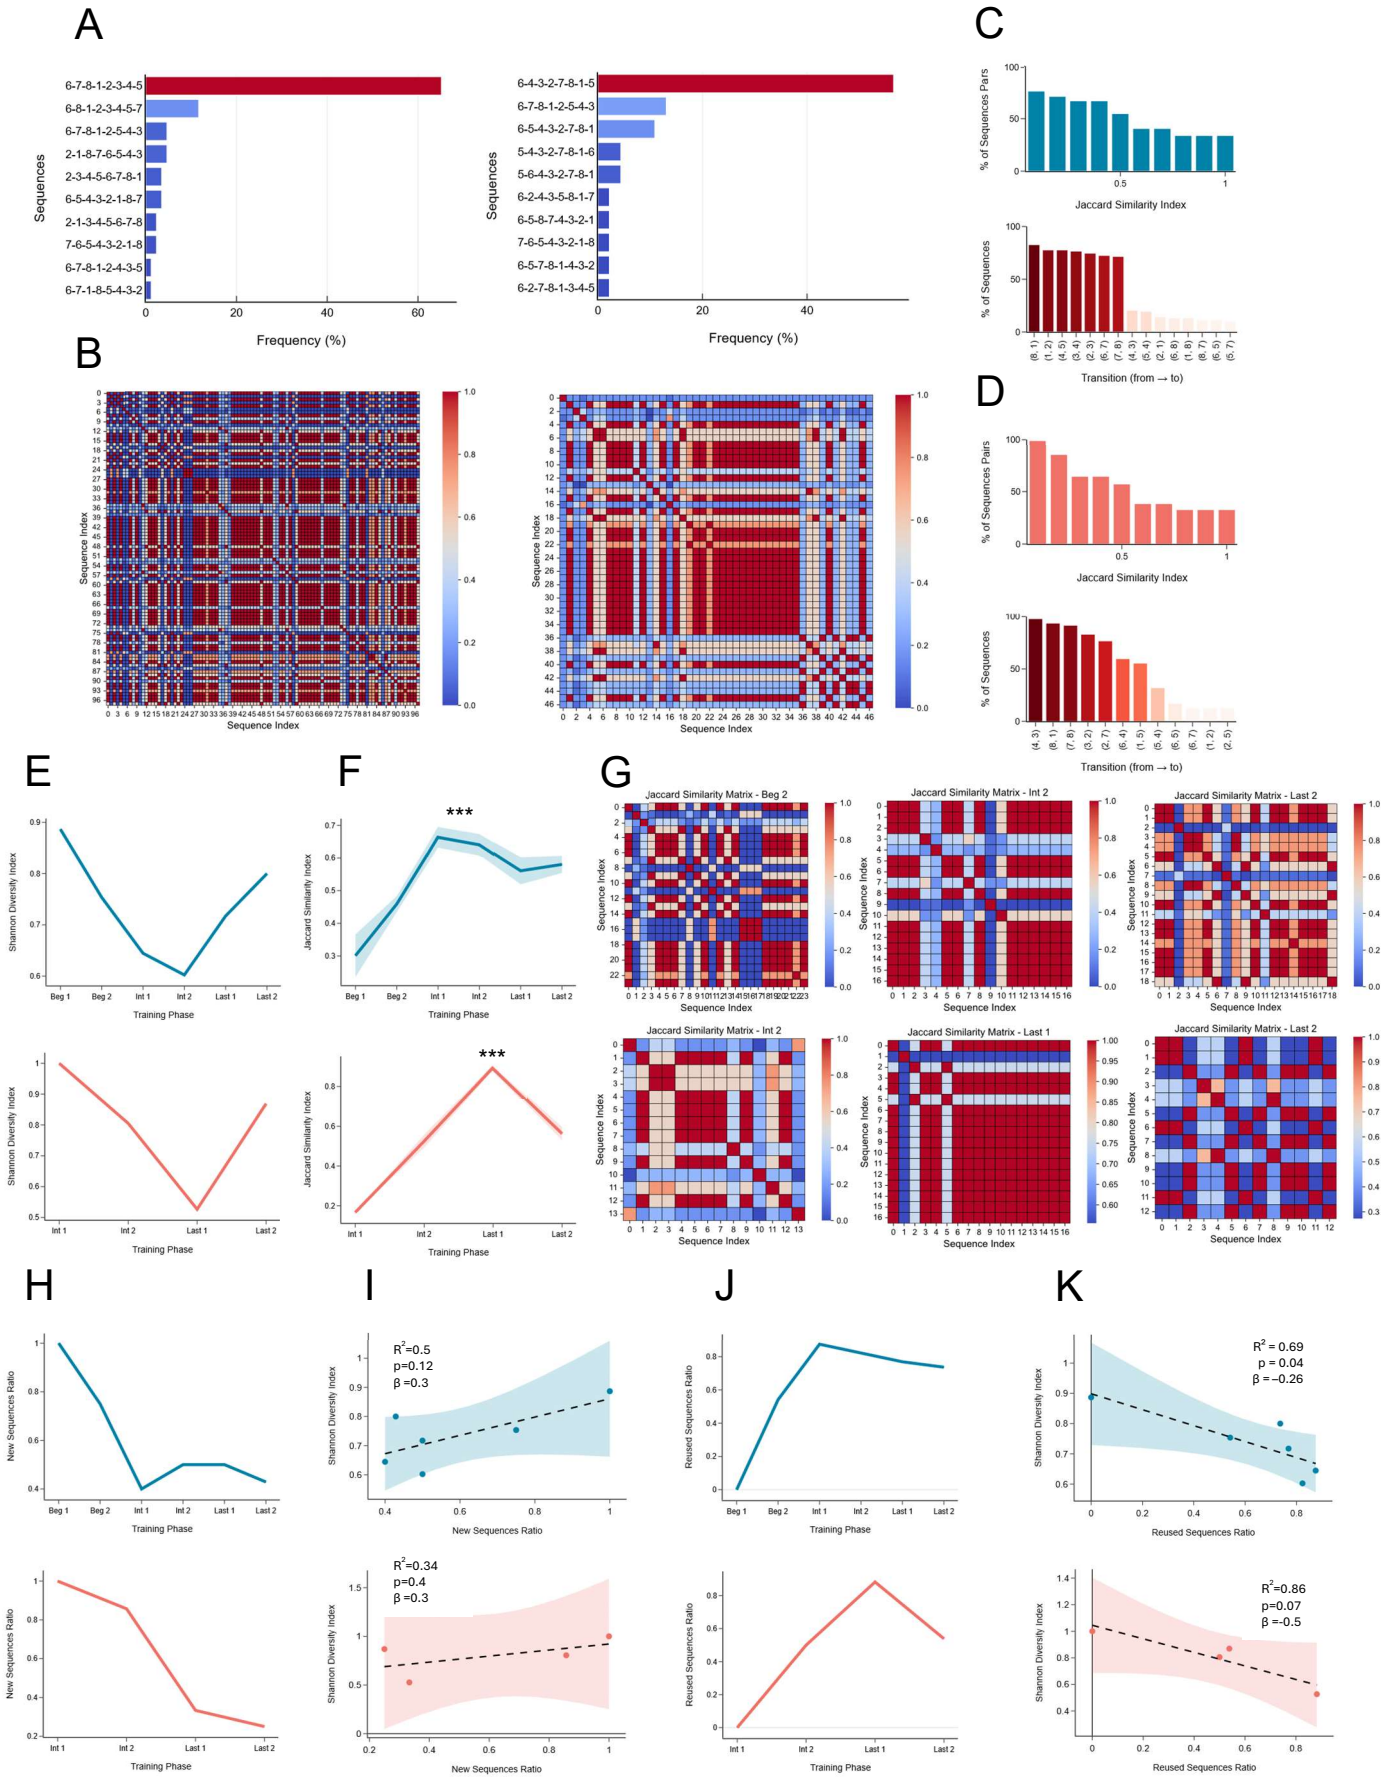
A
Sequences
Frequency (%)
B
Sequence Index
C
% of Sequences Pairs
Jaccard Similarity Index
% of Sequences
Transition (from → to)
D
E
Shannon Diversity Index
Training Phase
F
Jaccard Similarity Index
***
G
Jaccard Similarity Matrix - Beg 2
Jaccard Similarity Matrix - Int 2
Jaccard Similarity Matrix - Last 2
Jaccard Similarity Matrix - Last 1
H
New Sequences Ratio
I
R²=0.5
p=0.12
β =0.3
R²=0.34
p=0.4
β =0.3
J
Reused Sequences Ratio
K
R² = 0.69
p = 0.04
β = −0.26
R²=0.86
p=0.07
β =-0.5

Figure 6: Analysis of sequence diversity, similarity, and strategy reuse across training phases
A. Top 10 unique sequences in error-free trials. Horizontal bar plots showing the most frequently used transition sequences during error-free trials. Sequences are sorted by decreasing frequency, with percentages indicating their relative occurrence. *Left*: Animal E; *Right*: Animal F.
B. Sequence similarity (Jaccard index). Heatmaps displaying pairwise Jaccard similarity between all error-free sequences. Warmer colors indicate greater structural similarity (1 = identical, 0 = no overlap). *Left*: Animal E; *Right*: Animal F.
C. *Top*: **Pairwise sequence similarity distribution.** Bar plot showing the percentage of sequence pairs with Jaccard similarity equal to or greater than a range of thresholds. This illustrates the degree of structural redundancy or diversity across error-free sequences for Animal E. *Bottom*: **Most frequent transitions.** Bar plot showing the frequency of individual transitions (dispenser-to-dispenser moves) across all error-free sequences. Color intensity reflects recurrence, with darker red bars indicating higher frequency.
D. Same analyses as in panel C, shown for Animal F.
E. Evolution of sequence diversity. Line plots showing normalized Shannon diversity index across training phases, with higher values indicating greater sequence variability. *Top*: Animal E; *Bottom*: Animal F.
F. Sequence similarity across training phases. Line plots showing the evolution of mean Jaccard similarity (± SEM) across training phases. Asterisks indicate significant differences (Kruskal–Wallis test with Bonferroni-corrected Dunn's post hoc), *** $p < 0.001$. *Top*: Animal E; *Bottom*: Animal F.
G. Phase-specific similarity matrices. Heatmaps showing pairwise Jaccard similarity between all error-free sequences within individual training phases. These plots highlight structural convergence in behavior over time. *Top*: Animal E; *Bottom*: Animal F.
H. Exploration of new strategies across training phases. Line plots showing the evolution of the New Sequence Ratio (NSR), defined as the proportion of sequences in each phase that were not observed in any preceding phase. Higher values indicate greater exploration. *Top*: Animal E; *Bottom*: Animal F.
I. Association between sequence novelty and diversity. Scatter plots showing linear regressions between NSR and the normalized Shannon diversity index across training phases. Dashed lines represent regression fits; shaded areas indicate 95% confidence intervals. R^2, p-values, and regression slope (β) quantify the strength and direction of the association. *Top*: Animal E; *Bottom*: Animal F.
J. Reuse of past strategies across training phases. Line plots showing the evolution of the Reused Sequence Ratio (RSR), defined as the proportion of sequences in each phase repeated from earlier phases. Higher values indicate greater reliance on previously learned sequences. *Top*: Animal E; *Bottom*: Animal F.
K. Association between sequence reuse and diversity. Scatter plots showing linear regressions between RSR and the normalized Shannon diversity index across training phases. Dashed lines represent regression fits; shaded areas indicate 95% confidence intervals. R^2, p-values, and regression slope (β) quantify the strength and direction of the association. *Top*: Animal E; *Bottom*: Animal F.

Together, these results show that structural consistency in behavioral sequences emerged progressively over the course of training. Early phases were characterized by high variability and low similarity between sequences, suggesting that animals initially deployed a broad set of behavioral solutions. In later phases, they increasingly converged on dominant, structurally similar trajectories, reflecting a gradual stabilization of their navigation patterns. Although sequence diversity rose again in the final phase, core transition motifs remained stable both within and across phases. This temporal persistence suggests that animals increasingly relied on a stable set of transitions, rather than transient, phase-specific strategies.

Marmosets' shift from novel to reused sequences across training phases

The observed changes in sequence diversity and similarity indicate that structural motifs became increasingly consistent within and across training phases. However, these patterns do not reveal how such stability emerged. Was this convergence driven by the continuous generation of new but structurally similar sequences, or by the repeated use of previously established ones? To address this, we sought to characterize how animals balanced exploration and exploitation over training, that is, to what extent variability arose from the search for novel strategies versus the selective reuse of learned ones.

We quantified the relative contributions of novel and repeated sequences over time using two complementary metrics. The New Sequence Ratio (NSR) captured the proportion of unique sequences in each phase that had not been used previously, while the Reused Sequence Ratio (RSR) quantified how often sequences from earlier phases reappeared. High NSR values reflect active exploration of novel solutions, whereas high RSR values indicate growing reliance on familiar trajectories retrieved from past experiences.

Both animals displayed a similar progression (**Fig. 6H,J**). Early training phases were marked by high NSR and low RSR values, consistent with broad exploration. Over time, NSR progressively declined while RSR increased, indicating a gradual shift toward the reuse of previously successful sequences and reflecting a growing tendency to repeat previous sequences within each training phase.

To determine which behavioral metric best explained within-phase sequence diversity, we fitted separate linear models using each metric as a predictor of Shannon diversity index (**Fig. 6I,K**). The RSR emerged as the strongest predictor of variability, particularly in Animal E ($R^2 = 0.69$, $p = 0.04$, $\beta = -0.26$). In Animal F, a similar negative trend was observed ($R^2 = 0.86$, $p = 0.07$, $\beta = -0.51$), though it did not reach statistical significance, likely due to the small sample size ($n = 4$). In contrast, the NSR showed a weaker association with diversity (Animal E: $R^2 = 0.5$, $p = 0.12$, $\beta = 0.3$, Animal F: $R^2 = 0.34$, $p = 0.4$, $\beta = 0.3$), suggesting that exploration of novel solutions contributed less to variability.

Together, these findings reveal that, over the course of training, animals progressively shifted from generating new solutions to flexibly reusing learned ones, reflecting a dynamic balance between exploration and exploitation. Moreover, behavioral diversity appeared more strongly shaped by reuse of previous sequences than by the continued generation of novel ones.

Discussion

Understanding how animals organize their spatial behavior in complex environments remains a central challenge in cognitive neuroscience. In this study, we examined how freely moving marmosets progressively structured their navigation through repeated experience in a semi-naturalistic task. Without instruction or shaping, animals transitioned from exploratory to organized behavior, progressing through discrete states that reflected engagement and execution quality. Learning occurred selectively within engaged states, marked by reduced errors and the emergence of efficient navigation patterns. Over time, animals consolidated a stable repertoire of locally stereotyped transitions, while preserving global flexibility through the recombination of familiar elements. Rather than relying on fixed routes, they reused successful sequences in a dynamic and adaptive manner. These findings offer a rare behavioral framework for understanding how structured yet flexible navigation strategies emerge through self-guided experience in non-human primates.

A naturalistic paradigm to study spatial memory in marmosets

Unlike classical marmoset tasks that investigate memory using touchscreens, constrained locomotion, or simplified mazes (Izumi et al., 2013; MacDonald et al., 1994; Nakamura et al., 2018; Spinelli et al., 2004), our paradigm focuses on spatial navigation in a complex, three-dimensional environment, allowing animals to engage in unconstrained, ethologically relevant behaviors. This setting is well suited to investigate how spatial memory systems operate under conditions that approximate real-world challenges. The value of such naturalistic designs has been increasingly recognized across cognitive and systems neuroscience, as core functions like memory, decision-making, and learning emerge through continuous interactions between perception, action, and environmental context (Dennis et al., 2021; Krakauer et al., 2017; Lanzarini et al., 2025; Miller et al., 2022).

Our framework contributes to this shift by combining control over task contingencies with the freedom required to observe the spontaneous emergence of adaptive structure. In a semi-naturalistic setting without instructions or shaping, we show that marmosets progressively self-organize their navigation through repeated experience. This unsupervised process leads to the gradual development of efficient, memory-guided strategies that remain flexible, revealing how

cognitive structure can emerge from free interaction with a complex environment. By following behavior over extended training, our paradigm offers rare longitudinal insight into the co-evolution of navigational structure and memory performance in freely moving non-human primates.

Memory-guided learning emerges from distinct states of engagement

A major challenge in minimally constrained settings is accounting for fluctuations in motivation and task engagement. In our dataset, animals displayed substantial trial-by-trial variability in performance. To dissect this variability, we used differences across error types as markers of distinct motivational and cognitive states. Repetition and rule errors were positively correlated with each other and with visit counts, suggesting they reflect active but imprecise task execution. In contrast, omission errors were uncorrelated with other error types and likely signaled disengagement or failure to initiate the task.

Based on this dissociation, we applied unsupervised clustering using omission errors and visit counts. This revealed four consistent behavioral states: Clusters 0 and 1 both reflected active engagement, but differed in performance, with Cluster 0 corresponding to efficient execution and Cluster 1 to more error-prone behavior. In contrast, Clusters 2 and 3 captured partial or full disengagement. Although the clustering relied only on two features, the resulting states showed distinct error profiles, supporting their interpretation as reflecting both engagement and performance levels. Importantly, only the engaged clusters evolved with training: Cluster 0 (efficient execution) increased, while Cluster 1 (error-prone execution) decreased suggesting a shift toward more efficient task execution. Disengaged states remained stable, consistent with spontaneous fluctuations unrelated to learning.

This segmentation offers a more precise view of learning-related dynamics than traditional trial-averaging approaches, which risk obscuring meaningful variability. It shows that learning was not uniformly distributed across trials but emerged specifically within engaged states. By isolating learning within periods of clear engagement, we were able to disentangle memory-driven improvements from spontaneous motivational fluctuations, a distinction critical for interpreting behavior in unconstrained, naturalistic settings. This approach aligns with recent advances in computational neuroethology, where unsupervised clustering and behavioral state-space modeling are increasingly used to capture the richness of naturalistic behavior (Datta et al., 2019; Pereira et al., 2020).

Parallel refinement of spatial working and reference memory

During active engagement, we observed the parallel refinement of different complementary memory systems. Importantly, this refinement occurred in the absence of explicit instruction or corrective feedback, underscoring the capacity of marmosets to learn spatial structure through self-guided experience alone. The decline in repetition and rule errors over time reflects improved spatial working memory and a more accurate application of task structure, consistent with previous evidence that marmosets possess robust spatial working memory capacities (MacDonald et al., 1994). In parallel, animals increasingly reused previously successful sequences across sessions, as captured by the Reused Sequence Ratio (RSR), suggesting the gradual consolidation of spatial reference memory. They also progressively stabilized specific transition motifs between dispenser locations, as revealed by sequence similarity analysis and the high performance of the first-order Markov model at the local level. Rather than relying exclusively on working memory, animals appeared to progressively retrieve and apply long-term spatial knowledge accumulated through prior experience, at both global and local levels of organization. These patterns indicate that animals combined short-term tracking and long-term storage to optimize behavior, demonstrating that distinct memory systems can co-develop in naturalistic settings without external guidance, a dynamic rarely documented in freely moving non-human primates in captivity. Rather than emerging from a single learning mechanism,

efficient navigation appears to result from the gradual coordination of mnemonic processes operating at different timescales. This distinction underscores the importance of differentiating within-trial working memory from across-session regularities when analyzing spatial behavior in unconstrained, complex environments.

Emergence of structured navigation reflects consolidation of movement strategies

A key question in understanding memory-guided behavior is how performance improvements relate to the structuring of spatial exploration. In our task, animals were free to choose their paths, allowing us to ask whether learning involved not only reducing errors, but also organizing movement patterns more efficiently over time.

Across training, marmoset navigation exhibited a marked shift from dispersed exploration to more structured and selective transitions between dispenser locations. This was captured by a significant decrease in both global and local entropies, which strongly correlated with performance gains, suggesting that behavioral structure was not just a byproduct of learning, but an active contributor to task efficiency and memory performance.

To further characterize this process, we analyzed transition matrices across sessions. Early transitions were variable and inconsistent, reflecting exploration. With experience, these matrices stabilized, showing increased similarity across sessions and reduced Euclidean distances, indicating a consolidation of preferred transitions into structured navigational routines. Importantly, this structuring unfolded within a shared movement repertoire across engaged states. Clusters 0 and 1 differed in execution quality but relied on similar transition profiles. Cluster 1, however, maintained higher entropy and greater transition diversity, consistent with a transitional state where structure was present but inconsistently deployed. The shift from Cluster 1 to Cluster 0 across training suggests that learning involves not only discovering efficient routes, but also stabilizing their use over time.

Our paradigm provides the first longitudinal evidence for the emergence of structured navigation in freely moving non-human primates. While similar dynamics have been observed in rodents (Dupret et al., 2010; Wathen and Roberts, 1994; Yoerg and Kamil, 1982), our approach uniquely captures the temporal unfolding of structure through detailed trajectory analysis. These findings highlight a central feature of spatial cognition: efficient navigation arises not only from reducing errors, but from consolidating reproducible movement routines. This trajectory-level perspective offers new insight into how navigational structure emerges through sustained interaction with complex environments.

Structured yet flexible navigation suggests route-based strategies anchored in topological representations

The emergence of structured navigation routines raises a deeper question: what kind of spatial strategies support this organization? Classically, two broad types of strategies have been proposed: route-based navigation, which relies on memorized sequences of movements or landmark associations, and map-like representations, which support flexible path planning based on an internal model of spatial relationships. Distinguishing between these strategies provides insight not only into the behavioral algorithms used by the animals, but also into the underlying neural computations, as they are thought to recruit distinct cognitive processes and brain systems (O'Keefe and Nadel, 1978; Parra-Barrero et al., 2023; Poucet, 1993; Redish, 1999).

To probe the nature of the underlying strategy, we applied a first-order Markov model to error-free trials, quantifying how well local transitions between dispensers could be predicted based on the immediately preceding location. This modeling approach is well suited to detect route-based strategies, as it captures local sequential dependencies. The model achieved a high prediction accuracy (~80%), indicating that animals relied on a restricted set of stereotyped transitions. Notably, this local predictability increased over the course of training, reflecting the progressive consolidation of these patterns into stable navigational routines. These dominant transitions thus

appear to be reinforced through repeated experience, supporting the view that marmosets progressively assembled a repertoire of habitual local routes. Importantly, these core transitions were not confined to specific training phases but remained stable across training, indicating that navigational routines were consolidated into a temporally persistent structure.

However, when the same model was applied to predict full movement sequences, its performance dropped markedly. This drop suggests that while transitions were highly constrained at the local level, animals did not rigidly repeat the same full path on each trial. Instead, they used a limited set of frequently reused transitions that could be flexibly recombined. Sequence similarity analyses revealed that successful paths were not identical but were built from overlapping transitions, suggesting that animals flexibly recombined familiar elements rather than executing a fixed route. Further evidence for flexible recombination comes from spatial entropy analyses. Although global entropy decreased with training, indicating growing regularity in transitions, local entropy remained heterogeneous across dispensers. Some nodes retained higher entropy throughout, suggesting that they served as flexible decision points within otherwise stable routes. This combination of structure and variability is characteristic of topological representations, in which navigation is supported by a graph-like map of spatial connections rather than by fixed trajectories (Parra-Barrero et al., 2023; Poucet, 1993).

Altogether, these findings point toward a hybrid strategy: marmosets gradually assembled a structured set of route segments through repeated experience but retained the flexibility to navigate using different combinations of these segments. This organization is reminiscent of navigation in wild marmosets, which reuse habitual travel paths while adapting to environmental contingencies and maintaining flexible points within a network of interconnected paths (Abreu et al., 2021). It also aligns with theoretical accounts of topological navigation in rodents and humans, in which local regularities support efficient action while preserving global flexibility (Epstein et al., 2017; Poucet, 1993; Redish, 1999). Thus, the structured yet variable behavior observed in our task suggests that spatial learning in marmosets involves the progressive reinforcement of reliable local transitions nested within a flexible, topologically organized framework.

Dynamic shifts in sequence use reflect an exploration-exploitation trade-off

A striking feature of our results is the evolving balance between behavioral stability and flexibility over training. Early on, animals explored a broad set of options, generating diverse and structurally distinct sequences. With experience, they converged on a core set of strategies, marked by increased reuse of familiar trajectories and the emergence of a dominant sequence. Yet this consolidation did not yield rigidity: in later phases, animals flexibly recombined familiar transitions into new configurations, maintaining variability within a stabilized topological scaffold.

This shift was mirrored in the predictability of full movement sequences derived from the Markov model, which increased during consolidation and declined again in late training. This U-shaped trajectory in sequence diversity and model performance reflects the classical exploration-exploitation trade-off (Cohen et al., 2007; Mehlhorn et al., 2015; Sutton and Barto, 1998): initial variability promotes solution discovery, mid-training reflects the exploitation of reinforced paths, and late-stage flexibility enables adaptation without losing structural consistency.

Importantly, this renewed variability did not stem from the exploration of novel paths. The new sequence ratio (NSR) steadily declined over training, while the reused sequence ratio (RSR) increased and emerged as the strongest predictor of sequence diversity. This indicates that late-stage variability was not driven by novelty-seeking, but by the flexible recombination of previously learned sequences. Supporting this interpretation, the dominant sequence followed an inverted U-shaped trajectory, and sequence similarity remained high during the final training phases. Together, these patterns suggest that animals retained a consolidated repertoire of transition motifs, which they flexibly redeployed in novel configurations, rather than adopting entirely new strategies. Together, these findings extend classical reinforcement learning frameworks beyond

isolated choices to the structural level of full behavioral sequences. While exploration-exploitation dynamics are well documented in discrete decision-making tasks, especially in primates (Costa and Averbeck, 2020; Ferecatu and De Bruyn, 2022; Khamassi et al., 2010), our results demonstrate that similar principles govern naturalistic spatial navigation. Specifically, they show how structured action patterns can emerge through self-guided learning and later be flexibly reorganized. These findings resonate with recent theoretical accounts that emphasize the role of internal planning and information search in naturalistic behavior (Hunt et al., 2021) and suggest that late-stage variability may reflect proactive sampling within a consolidated internal model, balancing efficiency and flexibility through the recombination of learned elements. Altogether, our findings provide a dynamic view of adaptive behavior as an emergent property of internally regulated variability, bridging ecological navigation and computational models of learning beyond classical discrete-choice paradigms.

Neural correlates of memory-guided navigation in primates

The behavioral transformations observed in our study, ranging from performance gains and memory refinement to the consolidation of structured yet flexible trajectories, likely reflect the interplay of multiple cognitive and neural systems involved in memory, decision-making, and strategy selection.

Traditionally, spatial learning has been conceptualized within a dual-system framework, in which the hippocampus supports flexible, relational representations, while the striatum underlies reinforcement-based route acquisition and automatization (Goodman, 2021, O'Keefe and Nadel, 1978). However, accumulating evidence in both rodents and primates supports a more integrated and parallel organization, in which these systems interact dynamically throughout learning rather than operating in strict competition (Brown et al., 2012; Chersi and Burgess, 2015; Gahnstrom and Spiers, 2020; Goodroe et al., 2018; Sadeh et al., 2011). This integrative view aligns with our findings: rather than undergoing a linear shift from cognitive to habitual strategies, marmosets exhibited a dynamic balance between behavioral stability and flexibility, consistent with ongoing coordination between hippocampal and striatal systems.

This coordination is likely further modulated by the prefrontal cortex, which plays a critical role in integrating contextual goals, maintaining task-relevant information, and regulating the exploration-exploitation trade-off (Constantinidis and Procyk, 2004; Daw et al., 2006; Domenech et al., 2020; Jahn et al., 2023). In our task, the ability to flexibly adapt familiar routes, follow task-relevant rules, and alternate between variability and consolidation suggests the involvement of prefrontal mechanisms in orchestrating behavioral flexibility across timescales, a pattern consistent with integrative models of prefrontal-hippocampo-striatal interaction during spatial navigation (Khamassi and Humphries, 2012).

While the hippocampal-striatal-prefrontal architecture underlying spatial learning may be conserved across mammals, key sensorimotor differences shape how these computations are implemented in primates. Unlike rodents, which rely heavily on olfaction and path integration, primates, including marmosets, engage in visually guided navigation, combining landmark use with active gaze exploration (MacDonald et al., 2004; Martinez-Trujillo, 2025). These behavioral differences are reflected in distinct neural coding schemes: in rodents, hippocampal place cells encode allocentric position within 2D environments (O'Keefe and Dostrovsky, 1971), whereas primate hippocampal neurons show strong modulation by gaze direction and visual viewpoint, supporting more relational, view-dependent representations of space (Buffalo, 2025; Mao, 2022; Rolls et al., 1997; Rolls and Wirth, 2018). Notably, such coding has now been demonstrated in freely moving marmosets: hippocampal cells encode 3D position as well as head- and gaze-centered views (Piza et al., 2024), directly linking visual exploration to hippocampal dynamics under naturalistic conditions.

These findings position the marmoset as a key model not only for uncovering primate-specific coding principles, but also for bridging rodent and human studies of spatial cognition. Recent comparative work shows that marmosets display visual object recognition behavior strikingly

similar to that of humans and macaques (Kell et al., 2023), underscoring shared perceptual strategies across simian primates. This behavioral continuity suggests that the spatial navigation mechanisms at play in marmosets may reflect core components of primate cognition that are less accessible in traditional rodent models.

Limitations and future directions

While our paradigm enabled a rare longitudinal view of spatial learning in freely moving primates, some limitations should be acknowledged. First, the study was conducted on a small number of animals. However, this constraint is partly mitigated by the consistency of behavioral patterns across individuals and the richness of within-subject data collected over extended periods. This approach is common in primate neuroscience, where methodological and ethical constraints often limit sample size (Buffalo et al., 2019; Laurens, 2022) Importantly, detailed longitudinal tracking in small-N designs has yielded valuable insights into cognitive processes in non-human primates, particularly when intra-individual dynamics are the primary focus. Second, our semi-naturalistic setting maximized ecological validity in a laboratory environment, but limited experimental control over the sensory cues guiding navigation. In this open setup, animals could rely on any combination of visual, proprioceptive, or olfactory inputs, making it difficult to precisely dissociate their respective contributions. Future studies could address this by selectively manipulating sensory information (e.g., masking olfactory cues, altering visual landmarks) or introducing spatial perturbations to assess cue weighting and integration. Nonetheless, this openness allowed animals to develop ethologically relevant strategies that may be overlooked in more constrained paradigms, highlighting a trade-off between ecological richness and experimental control. Finally, our study did not include neural recordings, and circuit-level interpretations therefore remain hypothetical. While our behavioral findings align with theoretical models of hippocampal, striatal, and prefrontal contributions to navigation, direct evidence is needed to validate these mechanisms. Integrating chronic electrophysiological recordings within our behavioral framework will be essential to test how these circuits interact during self-guided navigation, and to uncover how conserved and species-specific neural dynamics support flexible spatial behavior in primates.

Conclusion

Our study reveals how structured, memory-guided behavior can emerge spontaneously in primates through repeated, self-guided interaction with a complex environment. By combining ethological richness with longitudinal precision, our paradigm captures how distinct memory systems co-develop to support flexible yet stable navigation strategies, learning processes rarely accessible in primates. Rather than acquiring fixed routes or relying solely on short-term memory, marmosets progressively assembled and flexibly reused a repertoire of spatial motifs embedded within a topological structure. These findings offer a powerful framework for linking ecological behavior to neural computations, and open new directions for exploring how sensorimotor constraints shape the neural bases of spatial cognition across species. This approach paves the way for bridging naturalistic ethology and systems neuroscience in primates, opening experimental avenues for integrating neural dynamics with flexible, structured behavior.

Methods

1. Ethical Statement

All animal procedures in this study were approved by the relevant regional government authority (F149/1012) and adhered to all applicable German and European regulations on animal husbandry and welfare. A small-sample approach was adopted, in line with previous research on spatial behavior in freely moving non-human primates (Buffalo et al., 2019; Ludvig et al., 2003; MacDonald et al., 2004, 1994). This approach allowed for a substantial amount of data to be collected per animal, minimizing the need for a larger sample while maintaining methodological rigor (Laurens, 2022).

2. Animals and Housing

Two adult male common marmosets (*Callithrix jacchus*, mean age: 3 years) participated in this study. The animals were pair-housed in a cage measuring 180 cm (H) × 90 cm (W) × 60 cm (D) under a 12-hour light-dark cycle (07:00-19:00). Room temperature was maintained at 24-25°C with 60% relative humidity. The marmosets were obtained from the German Primate Center (Göttingen, Germany). Experimental sessions took place in the morning. In addition to the gum arabic collected during the task, animals were fed a daily meal consisting of pellets, a dietary supplement primarily composed of quark, yogurt, banana and vitamins, and fresh vegetables. Water was provided ad libitum.

3. Experimental Apparatus and Data Collection

The experimental setup consisted of a wire mesh enclosure measuring 1.8 m × 1.4 m × 1.6 m (total volume: 4 cubic meters). The wire mesh walls facilitated climbing and allowed the marmosets to navigate freely within the space. The entrance was located in a designated corner, where a tunnel system enabled transition from the transport cage to the setup. To encourage locomotion and exploratory behavior, the enclosure was enriched with wooden sticks, semi-rigid ropes, and suspended enrichment items (**Fig.1A**). Additional environmental cues were introduced around the enclosure, including four white walls with landscape posters, four colored curtains, and large plastic plants.

During each experimental session, rewards were dispensed using custom-made, 3D-printed, remotely controlled dispensers (**Fig. 1A, E**). Each dispenser was calibrated to deliver 3 ml of liquid gum arabic per session, in increments of approximately 0.15 ml per activation. To ensure accurate dispensing, the system was tested and calibrated prior to each session.

Two experimenters conducted the sessions individually, alternating roles between sessions. During each session, the experimenter was positioned at a fixed location at a predetermined distance from the setup, manually scoring dispenser visits and task latencies in real time using BORIS, a well-established behavioral event logging software (Friard and Gamba, 2016). A visit to a dispenser was defined as the marmoset pausing in front of the dispenser and attempting to retrieve the reward. This criterion ensured that only purposeful interactions with the dispenser were recorded as visits. Additionally, a wide-angle lens camera was mounted above the setup, continuously recording the sessions for further analysis.

4. Behavioral Training

4.1 Habituation

Before the experiment, the marmosets underwent a habituation phase to ensure they did not exhibit anxiety and displayed natural exploratory behavior in the setup. This process took place over multiple sessions and was conducted by both experimenters, who had prior experience training the animals. As a result, the marmosets were well-acquainted with both experimenters before the start of behavioral testing. Initially, the two marmosets were trained to be comfortably transported from the animal facility to the laboratory in a transport cage. They were then introduced together into the setup, where they could explore and discover hidden rewards (e.g., marshmallows) placed in various locations. In the final habituation phase, only one marmoset entered the setup at a time, while the second remained in the transport cage nearby, maintaining visual contact. Initially, the separation lasted only a few minutes but gradually increased until the marmoset could remain alone for at least 30 minutes without signs of discomfort. Throughout this process, the experimenter monitored welfare by assessing vocalizations and behaviors based on established ethological criteria for captive marmosets (Schiel and Souto, 2017).

4.2 Home Dispenser Training

Since the marmosets explored the setup freely, they needed to be trained to voluntarily move to the center of the apparatus when required, ensuring that each trial started from a standardized position during the spatial working memory task. To achieve this, a food dispenser (i.e the "Home Dispenser (HD)") was installed upside down at the center of the ceiling (**Fig.1E**). The HD contained liquid gum arabic and was equipped with a blue LED and a buzzer, providing visual and auditory cues upon activation. Using positive reinforcement, the monkeys were trained to associate these cues with the HD activation and to reach it promptly.

Each marmoset was tested individually. At the start of each trial, the HD was activated for ~3 seconds, signaling the monkey to move to the center. The subject had up to 60 seconds to reach the HD for the trial to be considered successful. If the monkey failed to do so within this timeframe, the trial was marked as unsuccessful. A session was terminated if the monkey failed to complete two consecutive trials. Performance was evaluated based on the number of successful trials per session and the latency to reach the HD. Training was considered complete once the monkey achieved over 80% successful trials in two consecutive sessions. On average, each session consisted of 20.4 ± 1.19 trials for Animal E and 16.05 ± 1.2 trials for Animal F. To reach the training completion criterion, Animal E required 10 sessions, while Animal F needed 18 sessions.

5. Spatial memory task (Behavioral testing)

The experiment assessed spatial working memory by requiring marmosets to retrieve rewards from eight dispensers while remembering which ones they had already visited.

During the testing phase, eight additional dispensers were introduced, placed in fixed locations along the perimeter of the setup to ensure uniform spatial coverage (**Fig.1A**). Specifically, three dispensers were placed equidistantly along each of the two long sides, while one dispenser was positioned at the midpoint of each short side. This configuration ensured consistent spatial cues and a uniform distribution of reward locations across trials.

At the start of each trial, the Home Dispenser (HD) was activated, prompting the monkey to move to the center of the apparatus within 60 seconds. Once the monkey reached the HD, all eight dispensers were remotely activated simultaneously (**Fig.1B**)., each releasing ~0.15 ml of gum arabic as a reward. The food was concealed within the spout of each dispenser and not visible to the monkey, requiring them to actively search for and retrieve it. The reward remained available until collected.

To successfully complete the trial, the monkey had to visit all eight dispensers at least once, retrieving the reward from each. The trial ended upon the first visit to the eighth unique dispenser, regardless of whether the monkey revisited some dispensers along the way. However, if the monkey failed to visit all eight dispensers within 10 minutes, the trial was automatically terminated. After trial completion or timeout, the HD was reactivated, signaling the start of the next trial.

Each monkey underwent 30 training sessions, with an average of 7.96±0.4 trials for Animal E and 8±0.4 trials for Animal F per session. A session was terminated if the monkey failed to complete two consecutive trials.

6. Data analysis

All data analyses were performed in Python using custom-made scripts. Given the single-subject design, non-parametric tests were preferred. When needed, parametric analyses were conducted with model assumptions systematically verified, and robust standard errors applied when necessary. Bonferroni corrections were used to control for multiple comparisons. Statistical significance was set at $\alpha = 0.05$, with additional thresholds reported for $p < 0.01$ and $p < 0.001$ where relevant.

6.1 Spatial distribution of visits and omissions

To assess whether the marmosets interacted equally with all dispensers, we analyzed the number of visits and omissions for each of the eight dispensers. We first used Chi-squared tests to compare the observed visit and omission distributions to an expected uniform distribution, testing whether interaction patterns were evenly distributed. If a significant deviation was detected, binomial tests were conducted for each dispenser to identify which ones had visit frequencies that significantly differed from chance. A Bonferroni correction was applied to control for multiple comparisons.

6.2. Error Analysis and Regression Across Sessions

To assess error patterns across training sessions, errors were categorized into three types: rule errors (returning to the home dispenser during a trial), repetition errors (revisiting a dispenser within the same trial), and omission errors (failing to visit a dispenser). The total number of errors per trial was computed and averaged per session to obtain the mean number of errors per session and the standard error of the mean (SEM).

A linear regression analysis was performed to evaluate changes in the mean number of errors over time, with session number as the independent variable. The slope of the regression line was examined to determine whether errors decreased with training. Model fit was assessed using R-squared values, and confidence intervals were computed to visualize prediction uncertainty. To ensure the validity of the regression model, key assumptions were tested: normality of residuals was assessed using the Omnibus test, independence of errors was verified with the Durbin-

Watson statistic, and homoscedasticity was examined using the Breusch-Pagan test. Heteroscedasticity, if detected, was addressed by computing robust standard errors.

To further investigate factors influencing error rates, a multiple linear regression was conducted with session number, number of trials, number of days since the last session, and experimenter identity as independent variables. Model assumptions were systematically tested to ensure the validity of the regression analysis. Normality of residuals was assessed using the Shapiro-Wilk test, homoscedasticity was evaluated with the Levene test, and multicollinearity was examined through the Variance Inflation Factor (VIF). Ordinary Least Squares (OLS) regression with heteroscedasticity-robust standard errors was applied to account for variance inconsistencies. To control for multiple comparisons, p-values were adjusted using Bonferroni corrections.

6.3 Error Type Dynamics and Correlation Analysis

To examine the evolution of error types across training, we calculated the mean number and SEM of rule errors, repetition errors, and omission errors per session. A Kruskal-Wallis test was used to assess whether these error distributions differed significantly across sessions, followed by Dunn's post hoc test with Bonferroni correction for multiple comparisons. To explore relationships between error types, two correlation matrices were computed: one at the session level, where values were averaged per session, and one at the trial level, using individual trial data. Both matrices included total errors, total number of visits, rule errors, repetition errors, and omission errors. Pearson's correlation coefficients were used to assess linear associations, with Bonferroni correction applied to control for multiple comparisons.

6.4. Clustering analysis

6.4.1 Clustering procedure

We conducted a k-means clustering analysis to identify distinct behavioral patterns in marmoset task execution based on omission errors (failing to visit a dispenser) and total visits to dispensers. These two variables were selected to capture complementary aspects of task engagement, with omission errors reflecting disengagement and total visits indicating active task participation. Prior to clustering, the data were standardized to ensure equal weighting of both variables. The optimal number of clusters was determined by balancing silhouette scores and inertia, which assess cluster separation and within-cluster cohesion, respectively. Each trial was then assigned to a cluster based on proximity to the cluster centroids, representing the mean omission errors and total visits per cluster.

6.4.2 Statistical Comparison Between Clusters

To assess differences in task performance among the clusters identified by K-Means clustering, we analyzed several performance variables, including trial duration, total number of visits, and error types (rule errors, repetition errors, and omission errors). A Kruskal-Wallis test was conducted to determine whether these variables differed significantly across clusters. If significant differences were detected ($p < 0.05$), Dunn's post-hoc test was performed for pairwise comparisons, with a Bonferroni correction applied to control for multiple comparisons.

6.4.3 Regression Analysis of Cluster Proportions Across Sessions

A linear regression analysis was performed to evaluate changes in the proportion of trials assigned to each cluster over training, with session number as the independent variable. Separate

ordinary least squares (OLS) regression models were fitted for each cluster. The slope of the regression line was examined to determine whether the proportion of trials in each cluster changed significantly over time. Model fit was assessed using R-squared values, and confidence intervals were computed to visualize prediction uncertainty. To ensure the validity of the regression models, normality of residuals was assessed, independence of errors was verified using the Durbin-Watson statistic, and homoscedasticity was examined using the Breusch-Pagan test. When necessary, heteroscedasticity was addressed by computing robust standard errors.

6.5 Analysis of performance over the first eight choices

To assess task performance, we analyzed the success rate over the first eight choices, defined as the proportion of correct responses within these first eight selections. A response was considered correct if the marmoset selected a dispenser that had not been visited previously in the trial. This analysis was performed in two contexts: (1) comparison between clusters 0 and 1, and (2) evolution of performance across training phases.

6.5.1 Cluster Performance Analysis

To determine whether performance within each cluster differed from what would be expected under random selection, we computed the expected random performance assuming choices were made randomly among nine available options. Since the number of remaining correct options decreases at each step, the probability of selecting a correct option follows this sequence:

Equation 1:

$$P_{\text{random}} = \frac{8}{9} \times \frac{7}{9} \times \frac{6}{9} \times \frac{5}{9} \times \frac{4}{9} \times \frac{3}{9} \times \frac{2}{9} \times \frac{1}{9}$$

The expected average success rate was then computed as:

Equation 2:

$$E\left(\text{performance}_{\text{random}}\right) = \frac{1}{8} \sum_{i=1}^{8} \frac{9-i}{9} = 0.5$$

Thus, under random selection, the expected success rate is 50%, meaning that, on average, a marmoset would make 4 correct choices out of 8 by chance. This value served as a baseline for evaluating observed performance in each cluster.

For each cluster, we calculated the mean success rate and SEM. To assess whether cluster performance significantly deviated from the random expectation, we performed a Wilcoxon signed-rank test against the 50% baseline, with Bonferroni correction applied for multiple comparisons. Pairwise comparisons between clusters were conducted using a Mann-Whitney U test.

6.5.2 Performance Evolution Across Training Phases

To evaluate changes in performance over training, we divided the training sessions into six phases, each consisting of five sessions: the first ten sessions were divided into Beg 1 (S1-S5) and

Beg 2 (S6-S10), the ten intermediate sessions into Int 1 (S11-S15) and Int 2 (S16-S20), and the last ten sessions into Last 1 (S21-S25) and Last 2 (S26-S30). For each phase, we computed the mean performance rate and SEM. Performance rates were compared across these six phases using a Kruskal-Wallis test to determine whether task performance significantly improved over time.

6.6. Transition Patterns Analysis

6.6.1 Cluster Analysis

6.6.1.1 Transition Probability Computation

Each trial was represented as a sequence of visits to dispensers. Transition probabilities between locations were calculated by determining the frequency of transitions from one dispenser to another. The probability of transitioning from dispenser i to dispenser j was defined as:

Equation 3:

$$P(i \rightarrow j) = \frac{\text{Number of transitions from } i \text{ to } j}{\sum_k \text{Number of transitions from } i \text{ to } k}$$

Where $P(i \rightarrow j)$ is the probability of transitioning from dispenser i to dispenser j, and the denominator is the total number of transitions starting from dispenser i to any possible dispenser (the sum over all possible dispensers k).

To ensure comparability across clusters, we considered all possible transitions-including those with zero occurrences-and then normalized the transition probability matrix on a row-by-row basis. In other words, for each dispenser (row), we scaled the counts so that the sum of the probabilities for all transitions starting from that dispenser equals one.

6.6.1.2 Global Entropy and Global Transition Diversity

For the calculation of global entropy and global transition diversity, the transition probability matrix was further normalized globally. That is, after computing the row-wise transition probabilities for each dispenser, we rescaled the entire matrix so that the sum of all transition probabilities within each cluster equals 1. This global normalization allows for a direct comparison of transition patterns between clusters.

We calculated global entropy based on Shannon entropy's formula:

Equation 4:

$$H = -\sum P(i \rightarrow j) \log_2 P\,(i \rightarrow j)$$

Where $P(i \rightarrow j)$ represents the probability of transitioning from dispenser i to dispenser j. In this context, global entropy quantifies the overall unpredictability of transitions across all dispensers: higher entropy values indicate that transitions are more varied and less predictable, while lower values suggest a more structured, stereotyped pattern dominated by a few transitions.

In addition, we calculated global transition diversity using Simpson's Diversity Index to assess the evenness of transition probabilities between dispensers. The index was computed as:

Equation 5:

$$D = 1 - \sum P(i \rightarrow j)^2$$

In the context of transition patterns, global transition diversity represents the overall distribution of transition probabilities across all dispensers. A higher diversity value indicates that transitions are more evenly spread across dispensers, meaning that no single transition overwhelmingly dominates the movement pattern. Conversely, lower diversity suggests that only a few transitions occur more frequently, leading to more repetitive and less varied behavior.

To estimate the variability and compute confidence intervals for global entropy and global transition diversity, we performed a bootstrap resampling procedure with 1,000 iterations. In each iteration, transition probabilities were resampled with replacement, and entropy and diversity values were recomputed. The 95% confidence intervals (CI) were derived from the 2.5th and 97.5th percentiles of the bootstrap distributions.

6.6.1.3 Simulation of Random Behavior for Baseline Comparison

To establish a baseline for global entropy and global transition diversity under random movement conditions, we simulated random sequences of visits and computed their expected entropy and diversity values. A total of 1,000 simulations were performed, each consisting of 1,000 randomly generated sequences of 8 visits (excluding the starting location). In each sequence, dispensers were selected randomly with replacement, ensuring no systematic preference for specific locations. Transition counts were collected across all simulated sequences to compute empirical random transition probabilities. These probabilities were globally normalized to ensure comparability with the observed data. From these simulations, we calculated the mean global entropy and global transition diversity, and derived confidence intervals (at the 95%, 99%, and 99.9% levels) based on the corresponding percentiles of the bootstrap distributions.

6.6.1.4 Permutation Test for Statistical Comparison

Permutation tests were applied to compare the observed global entropy and global transition diversity of each cluster to both the random baseline and between clusters. In 1,000 iterations, the observed data and simulated baseline were pooled, and data were randomly reassigned to two groups of equal size. The mean difference in global entropy and global transition diversity was calculated for each permutation, generating a null distribution. Bonferroni corrections were applied to adjust for multiple comparisons.

6.6.1.5 Transition Matrices

To investigate the behavioral transition patterns, transition matrices were calculated for each cluster, representing the transition probabilities between dispensers across trials. These matrices were computed for both global transitions (including all observed transitions) and significant transitions (those that were significantly different from random chance).

Transition matrices were constructed by counting the frequency of transitions between dispensers across all trials. Transition probabilities were then normalized so that the sum of probabilities for each row (representing outgoing transitions from a dispenser) equaled 1.

To identify statistically significant transitions, binomial tests were conducted for each transition in the matrix, comparing the observed transition probabilities to a uniform random baseline of 1/9 (assuming each dispenser is equally likely to be chosen). Bonferroni correction was applied to adjust for multiple comparisons. Only transitions with corrected p-values less than 0.05 were considered significant.

The similarity between the transition matrices of different clusters was assessed by computing the correlation between their flattened versions. A permutation test with 1,000 iterations was applied to randomly shuffle the values in each matrix pair and compare the correlation of the shuffled matrices to the observed correlation. This process allowed the calculation of a p-value for the correlation, which quantified the likelihood of obtaining the observed correlation by chance. Additionally, the Euclidean distance between the transition matrices was calculated to measure the dissimilarity between matrices, with smaller distances indicating greater similarity.

6.6.1.6 Transition Probability Mapping

To visualize the structure of movement patterns, transition maps were generated for each cluster based on the first 8 choices of each trial. These patterns were visualized using directed graphs where all transition probabilities were represented by a weighted graph, with edge thickness proportional to the transition probability. Significant transitions were represented as a subset of transitions meeting statistical significance.

6.6.2 Analysis by Training Phase

6.6.2.1 Transition Probability Computation

The computation of transition probabilities followed the methodology described in section 6.6.1.1. Briefly, each trial was represented as a sequence of dispenser visits, and transition probabilities were calculated from the frequency of transitions between dispensers. To ensure comparability across training phases, we considered all possible transitions-even those with zero occurrences-and normalized the resulting probabilities on a row-by-row basis (i.e., for each dispenser, the sum of the probabilities of transitions from that dispenser was scaled to equal 1).

6.6.2.2 Global Entropy

Global entropy analysis for training phases was conducted in the same way as in Section 6.6.1.2 section. Transition probabilities were globally normalized, so that the sum of all transition probabilities within each cluster equaled 1 and Shannon entropy was computed to measure the unpredictability of transition patterns within each training phase. Bootstrap resampling with 1,000 iterations was used to estimate the variability of entropy values, as described in 6.6.1.2. The 95% confidence intervals (CI) for entropy were derived from the 2.5th and 97.5th percentiles of the bootstrap distributions. Permutation tests were applied to compare the observed entropy between training phases. In 1,000 iterations, data were pooled and randomly reassigned to two groups. The mean difference in entropy was calculated for each permutation, and Bonferroni corrections were applied to adjust for multiple comparisons.

6.6.2.3 Overlap Index Calculation

To assess the degree of similarity between entropy distributions from different training phases, we calculated the Overlap Index. This index measures the shared area between two probability distributions and provides an indication of how similar the transition behaviors are between

phases. The index ranges from 0 to 1, with 0 indicating no similarity and 1 indicating complete similarity in the transition patterns. For each pair of training phases, the entropy distributions were calculated, and histograms were constructed with 100 bins, covering the range of observed entropy values. The Overlap Index was then determined by summing, across all bins, the minimum value between the corresponding bins of the two histograms. Higher Overlap Index values indicate greater similarity in the transition patterns between dispensers across training phases

6.6.2.4 Correlation between global entropy and performance rate

To evaluate the linear relationship between the mean performance rate (across the first 8 choices) and the mean global entropy, their correlation was assessed using Pearson's correlation coefficient across all training phases. Bonferroni correction was applied to account for multiple comparisons. Additionally, a linear regression model was fitted to the data to visually represent and further assess the relationship between these two variables.

6.6.2.5 Transition Matrices

Transition matrices for each training phase were calculated following the same approach as in the 6.6.1.5 section. Significant transitions were identified using binomial tests, with Bonferroni corrections for multiple comparisons.

6.6.2.6 Transition Probability Mapping

Transition maps were generated for each training phase in the same manner as in Section 6.6.1.6. Significant transitions were identified based on binomial tests, and Bonferroni corrections were applied for multiple comparisons.

6.6.2.7 Local Entropy

Local entropy quantifies the unpredictability of transitions at the individual dispenser level. For each dispenser, we first calculated the transition probabilities, that is, the likelihood of transitioning from that dispenser to each of the other dispensers. The local entropy for a dispenser was then computed using Shannon's entropy formula:

Equation 6:

$$H(X) = -\sum_{i} p(x_i) \log_2 p\,(x_i)$$

Where $p(x_i)$ represents the probability of transitioning to dispenser i. For dispensers with no outgoing transitions, the entropy was set to zero. In this context, higher local entropy values indicate greater variability (i.e., less predictable and more diverse transitions), whereas lower values suggest more structured, predictable transition patterns.

To assess whether local entropy values differed significantly across training phases, we applied the Kruskal-Wallis test. When significant differences were detected, pairwise comparisons were performed using Dunn's post-hoc test with Bonferroni correction.

In addition, to evaluate whether local entropy values differed from what would be expected under random behavior, we performed a bootstrap resampling analysis at the local level. For each dispenser, we resampled the observed transitions 1,000 times with replacement and computed a distribution of local entropy values. From these distributions, we derived 95% confidence

intervals (CI) based on the 2.5th and 97.5th percentiles. This allowed us to determine whether the observed local entropy for each dispenser significantly deviated from chance-level expectations.

6.6.2.8 Error Probability

For each dispenser, we calculated error probability as the proportion of errors that occurred following a visit to that dispenser. Specifically, for each training phase, we counted the number of errors that immediately followed visits to each dispenser and divided that number by the total number of trials in that phase. This provided an estimate of the likelihood of making an error after visiting a given dispenser, relative to the overall number of trials-in other words, the probability of an error in any given trial when that dispenser is involved. To assess whether error probabilities differed significantly across training phases, we applied the Kruskal-Wallis test. When significant differences were detected, Dunn's post-hoc test with Bonferroni correction was used to identify which specific phases differed.

6.6.2.9 Correlation between Local Entropy and Error Probability

To evaluate the linear relationship between mean local entropy and mean error probability, Pearson's correlation coefficient was computed using data from all training phases combined. A Bonferroni correction was applied to adjust for multiple comparisons. Additionally, a linear regression model was fitted to the data to visually represent and further assess the relationship between these two variables.

6.7. Markov Model Analysis

The Markov model was trained exclusively on error-free trials to focus on transition patterns observed during successful task performance. This section describes the methods used to construct the transition matrix, predict transitions, evaluate performance, and compare the Markov model to a random baseline.

6.7.1 Transition Prediction and Accuracy

Each trial was represented as a sequence of visits to dispensers. Transition probabilities between dispensers were computed by tracking the frequency of transitions from one dispenser to another across all trials. The transition probability from dispenser i to dispenser j was defined following *Equation 3* (see Section 6.6.1.1).

After calculating these transition frequencies, each row of the transition matrix was normalized to ensure that the sum of outgoing transition probabilities from each dispenser was equal to 1. This normalization ensures that the transition matrix represents valid probability distributions for the Markov model, where each row is a probability distribution over the next possible dispensers.

6.7.1.1 Transition prediction Procedure

To evaluate the model's predictive capabilities, we compared the Markov model against a random baseline model. For the Markov model, the next dispenser choice was predicted based on the transition matrix computed from the data. Specifically, the prediction of the next dispenser for a given state i (current dispenser) is made by selecting the dispenser j with the highest transition probability:

Equation 7:

$$\hat{j} = \arg\max_{j} P\,(i \rightarrow j)$$

Where $\hat{j}$ is the predicted next dispenser and $P(i \rightarrow j)$ is the transition probability from dispenser i to dispenser j. As a baseline, a random model was implemented, where the next dispenser was selected uniformly at random from all dispensers. Prediction accuracy was calculated as the proportion of correct predictions, i.e., the proportion of times the predicted next dispenser matched the actual next dispenser.

6.7.1.2 Cross-Validation

To evaluate the robustness of the Markov model, we performed 10-fold cross-validation. The dataset was randomly split into 10 subsets (folds), with 9 folds used for training and the remaining fold used for testing. This process was repeated 10 times, with each fold serving as the test set once. For each fold, the transition matrix was computed based on the training data, and the prediction accuracy for both the Markov model and a random baseline model was evaluated on the test set. Performance metrics, including prediction accuracy, were then averaged across all folds to provide a reliable measure of the model's generalization capability.

In addition, confusion matrices were generated for each fold to analyze the predictions made by the Markov model. The final confusion matrix was obtained by averaging the confusion matrices across all folds, and each row was normalized so that the sum of probabilities for transitions from each actual dispenser equals 1. This normalization accounts for variations in transition frequencies and provides a clear picture of the model's performance on a per-dispenser basis.

Similarly, an error matrix was constructed by aggregating the counts of mispredicted dispensers across all folds. The error matrix was then normalized globally-by dividing each entry by the total number of errors across all folds-to yield a measure of the relative frequency of errors associated with each dispenser. The normalized error matrix highlights which dispensers are most prone to errors, offering insights into specific weaknesses in the model's predictions.

6.7.1.3 Evaluation and Statistical Comparison Between Models

To determine whether the Markov model significantly outperformed the random baseline in predicting individual transitions, we conducted a Wilcoxon signed-rank test on the accuracy scores obtained across the cross-validation folds. For both models, mean transition prediction accuracy was computed across folds to provide an overall performance measure. Variability was quantified by calculating the standard deviation and 95% confidence intervals (CI) of these accuracy scores. The Wilcoxon signed-rank test was then applied to assess whether the Markov model's transition accuracy was significantly higher than that of the random model, ensuring a robust statistical comparison.

6.7.1.4 Transition Matrix Visualization

To visualize the transition patterns learned by the Markov model, we generated predicted transition maps. These maps were represented as directed graphs, where nodes corresponded to dispensers and edges depicted the predicted transitions. The thickness of each edge was proportional to the corresponding transition probability, providing an intuitive visual representation of the local behavioral patterns captured by the model.

6.7.2 Sequence Prediction and Accuracy

To further evaluate the performance of the Markov model, we assessed its ability to predict full sequences of dispenser visits. Specifically, we focused on predicting sequences of 8 consecutive choices made by the marmoset. Instead of predicting transitions independently, we treated each sequence as a complete set of transitions to be predicted. Importantly, the same transition matrix computed from the training data in Section 6.7.1 was used for sequence predictions, ensuring consistency across both levels of analysis.

6.7.2.1 Sequence Prediction Procedure

For each sequence, predictions were generated using the same transition matrix from the training phase. The sequence was predicted iteratively, beginning with the first dispenser and selecting the next dispenser at each step based on the highest transition probability. Sequence prediction accuracy was then calculated as the proportion of correctly predicted sequences, representing the percentage of cases where the predicted sequence matched the actual sequence across all eight transitions. This approach assessed the Markov model's ability to capture global dependencies between transitions. To establish a baseline for comparison, a random model was also evaluated, where each dispenser was predicted with equal probability across all possible dispensers. The accuracy of the random model was computed using the same methodology, allowing for a direct performance comparison.

6.7.2.2 Cross-Validation for Sequence Prediction

To evaluate the generalization ability of the Markov model for sequence prediction, we applied 10-fold cross-validation, ensuring consistency with the methodology used in Section 6.7.1. The dataset was randomly divided into 10 folds, with each fold serving as a test set once while the remaining folds were used for training. For each fold, the transition matrix was computed from the training sequences and used to predict sequences in the test set. Prediction accuracy was assessed for both the Markov and random models, allowing for a comparative evaluation of performance.

6.7.2.3 Evaluation and Statistical Comparison Between Models

To determine whether the Markov model significantly outperformed the random baseline in predicting sequences of dispenser visits, we conducted a Wilcoxon signed-rank test on sequence accuracy scores. Mean sequence accuracy across folds was computed for both models, providing an overall measure of prediction performance. Variability was quantified by calculating the standard deviation and 95% confidence intervals (CI) of these accuracy scores. The Wilcoxon signed-rank test was then applied to assess whether the Markov model's sequence accuracy was significantly higher than that of the random model, ensuring a robust statistical comparison.

6.7.2.4 Transition Matrix Visualization

To illustrate the predicted sequences, we constructed directed graphs in which nodes represent dispensers and edges represent the corresponding transition probabilities. In these graphs, only the sequence predicted by the Markov model was displayed. The thickness of each edge was scaled according to its transition probability, thereby highlighting the most likely transitions. This visualization provides an intuitive depiction of the predicted sequence and of how effectively the model captures global behavioral patterns.

6.7.3 Transition Prediction Accuracy Across Training Phases

To assess the evolution of the Markov model's performance over time, we evaluated its predictive accuracy separately for each training phase. This analysis allowed us to examine how the model adapted across different learning phases and whether prediction accuracy improved as training progressed.

6.7.3.1 Local Transition Prediction and Accuracy by Training Phase

The same Markov modeling approach described in Section 6.7.1 was applied independently to each training phase, with separate transition matrices computed for each phase. This ensured that any variations in transition dynamics across training could be captured. For each training phase, we calculated transition probabilities, predicted the next dispenser based on the highest transition probability, and compared the Markov model's accuracy to a random baseline.

6.7.3.2 Cross-Validation by Training Phase

For cross-validation, the number of folds was dynamically adjusted based on the available data in each training phase. When sufficient sequences were available, 10-fold cross-validation was performed, otherwise, the number of folds was reduced accordingly to ensure reliable evaluation. Accuracy scores were averaged across folds to provide a robust measure of generalization within each training phase. Additionally, confusion matrices were computed separately for each training phase to analyze errors and evaluate model performance over time.

6.7.3.3 Evaluation and Statistical Comparison Between Models by Training Phase

To determine whether the Markov model's performance improved throughout training, we compared accuracy across training phases as well as against the random baseline for each phase. Variability was quantified by calculating the standard deviation and 95% confidence intervals (CI) of these accuracy scores. For each training phase, a Wilcoxon signed-rank test with Bonferroni correction was conducted to assess whether the Markov model significantly outperformed the random model. In addition, a Kruskal-Wallis test was performed to evaluate whether differences in accuracy across training phases were statistically significant. When the Kruskal-Wallis test indicated significant differences, Dunn's post-hoc test with Bonferroni correction was applied to identify which specific phases differed.

6.7.3.4 Transition Matrix Visualization by Training Phase

To visualize how the Markov model's predictions evolved over training, we generated predicted transition maps for each training phase. These maps were represented as directed graphs, where nodes corresponded to dispensers and edges depicted the predicted transitions for that phase. The thickness of each edge was proportional to the corresponding transition probability, providing an intuitive visual representation of the transition patterns during each phase. This approach allowed us to track the evolution of local behavioral patterns as training progressed.

6.7.4 Sequence Prediction Accuracy Across Training Phases

Beyond individual transitions, we also evaluated the Markov model's ability to predict full sequences of dispenser visits separately for each training phase. This analysis allowed us to assess whether the model became more effective at capturing global behavioral patterns over time. Importantly, the same transition matrices computed from the training data for each training

phase in Section 6.7.3 were used for sequence predictions, ensuring consistency across both levels of analysis.

6.7.4.1 Sequence Prediction and Cross-Validation by Training Phase

For each training phase, sequences were extracted and analyzed separately, following the same Markov modeling approach as in Section 6.7.2. We used the transition matrices computed in Section 6.7.3 for each training phase. This ensured consistency across both transition- and sequence-level analyses.

Cross-validation was adapted dynamically to match the dataset size in each training phase. 10-fold cross-validation was applied when enough sequences were available, while fewer folds were used when necessary to maintain a robust evaluation. In each fold, the precomputed transition matrix from the training phase was used to predict full sequences of up to 8 transitions in the test set. The Markov model's sequence accuracy was compared to that of a random baseline model, which predicted each transition with equal probability.

6.7.4.2 Evaluation and Statistical Comparison of Sequence Prediction

To assess whether the Markov model's sequence prediction accuracy significantly improved over training, a Wilcoxon signed-rank test with Bonferroni correction was conducted for each training phase to determine if the Markov model significantly outperformed the random baseline. Additionally, a Kruskal-Wallis test was applied to analyze whether sequence prediction accuracy evolved significantly across training phases. When the Kruskal-Wallis test indicated significant differences, Dunn's post-hoc test with Bonferroni correction was used to identify which specific training phases differed.

6.7.4.3 Sequence Transition Visualization

To illustrate the predicted sequences for each training phase, we constructed directed graphs in which nodes represent dispensers and edges represent the corresponding transition probabilities. In these graphs, only the sequence predicted by the Markov model was displayed. The thickness of each edge was scaled according to its transition probability, thereby highlighting the most likely transitions for each training phase.

6.8. Analysis of Sequence Variability and Similarity

To further explore the structure of dispensers transition sequences, we analyzed the variability and similarity of observed sequences across error-free trials.

6.8.1 Identification of Unique Sequences and Frequency Analysis

Transition sequences were extracted from error-free trials, where each sequence was defined as a structured set of consecutive transitions between dispensers. The number of unique sequences was identified, and their relative usage frequencies were computed to evaluate the diversity of sequence patterns. To determine whether certain sequences were dominant behavioral patterns, we identified the most frequently used sequence for each animal and calculated its proportion relative to all observed sequences.

6.8.2 Analysis of Sequence Diversity

To quantify the diversity of navigation sequences used in error-free trials, we computed the normalized Shannon diversity index, which is derived from Shannon entropy and normalized by its theoretical maximum. The normalized Shannon index was defined as:

Equation 8:

$$H_{\text{norm}} = \frac{-\sum_{i=1}^{N} p_i \log_2 p_i}{\log_2 N}$$

where N is the total number of unique sequences, and p_i represents the relative frequency of sequence i, computed as the number of occurrences of that sequence divided by the total number of sequences.

This normalization ensures that H_{norm} ranges between 0 (when only one sequence is used, indicating no diversity) and 1 (when all sequences occur with equal probability, indicating maximum diversity).

6.8.3. Analysis of Sequence Similarity

To assess the structural similarity between sequences, we examined common transitions, defined as pairs of consecutive dispensers appearing within sequences. Sequence similarity was quantified using the Jaccard similarity index, calculated as:

Equation 9:

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|}$$

Where A and B represent the sets of transitions in two sequences, and |A ∩ B| denotes the number of shared transitions between them. This index varies between 0 and 1, with 0 indicating no shared transitions and 1 indicating complete overlap. This measure provided insight into the degree of transition overlap, indicating whether sequences followed similar transition patterns or exhibited high variability.

We computed the mean Jaccard similarity across all sequence pairs and analyzed the distribution of similarity levels. Specifically, we determined the proportion of sequence pairs exceeding similarity thresholds ranging from 0.1 to 1.0, providing a detailed view of how frequently sequences shared a given percentage of their transitions. This approach allowed us to evaluate sequence overlap and assess the structural consistency of navigation patterns over trials.

To further examine sequence similarity patterns, the Jaccard similarity matrix was represented as a heatmap, where each cell indicated the similarity score between two sequences. This visualization provided a structured overview of sequence similarity distributions and facilitated the identification of clusters of highly similar sequences.

Additionally, to quantify the occurrence of individual transitions, we counted the frequency of each transition (i.e., movement from one dispenser to another) across all sequences. These frequencies were normalized by computing the percentage of sequences containing each transition relative to the total number of sequences analyzed.

6.8.4 Sequence Diversity Across Training Phase

To examine the evolution of sequence diversity over training, we computed the normalized Shannon diversity index separately for each training phase (as described in Section 6.8.2). To model the trajectory of diversity across training, we applied weighted least squares (WLS) quadratic regression, with weights determined by the number of trials in each phase to account for sample size differences. We compared the performance of a simple linear regression to that of the quadratic model to determine whether a non-linear (e.g., U-shaped) trend provided a better fit. Model fit was evaluated using R^2 values, and the statistical significance of the quadratic term was examined. Prior to interpreting the results, we verified that the model assumptions-normality of residuals, homoscedasticity, independence of errors, and absence of multicollinearity-were met, confirming the reliability of our approach.

6.8.5 Dominant Sequence Frequency, New Sequence Ratio, and Reused Sequence Ratio Across Training Phases

To examine how sequence selection evolved throughout training, we quantified three key metrics: dominant sequence frequency, new sequence ratio, and reused sequence ratio. These measures were computed separately for each training phase to assess changes in the patterns of sequence selection over time.

6.8.5.1 Dominant Sequence Frequency

The dominant sequence was identified as the most frequently occurring sequence across all error-free trials. For each training phase, we computed its frequency of use, defined as the proportion of trials in which the dominant sequence was executed.

6.8.5.2 New Sequence Ratio

The New Sequence Ratio (NSR) was computed as the number of unique sequences in each training phase that had not been observed in any previous phase. This metric was calculated iteratively by maintaining a cumulative set of previously encountered unique sequences and identifying, for each phase, the fraction of new unique sequences relative to all unique sequences observed during that phase

6.8.5.3 Reused Sequence Ratio

The Reused Sequence Ratio (RSR) was computed as the proportion of sequences in each training phase that had already been observed in previous phases. For each training phase, reused sequences were identified by comparing the sequences in the current phase to a cumulative set of sequences encountered in earlier phases. The reused sequence ratio was then calculated as the total number of reused sequences divided by the total number of sequences in that phase.

6.8.5.4 Linear Relationship Between Sequence Diversity and Sequence Selection Strategies

To assess the relationship between sequence diversity and sequence selection patterns, we conducted linear regression analyses using the normalized Shannon Diversity Index as the independent variable and either the New Sequence Ratio (NSR) or the Reused Sequence Ratio (RSR) as the dependent variable. Separate ordinary least squares (OLS) regression models were fitted for each metric and each animal to evaluate whether higher diversity was associated with a greater tendency to generate novel sequences (NSR) or reuse existing ones (RSR). Model fit was

assessed using R-squared values, and 95% confidence intervals were computed to visualize prediction uncertainty. To validate the models, residual normality was assessed using the Shapiro-Wilk test, independence of errors was verified using the Durbin-Watson statistic, and homoscedasticity was tested using the Breusch-Pagan test.

6.8.6 Sequence Similarity Within and Across Training Phases

To assess how sequence similarity evolved throughout training, we applied the same Jaccard similarity analysis described in Section 6.8.3, computed separately for each training phase. For each phase, sequences were extracted and the pairwise Jaccard similarity index was calculated to quantify the overlap in transitions between sequences. Mean similarity scores and SEM were computed to evaluate whether sequences became more homogeneous or remained diverse over time. Additionally, we assessed sequence similarity between training phases by computing pairwise Jaccard similarity scores between sequences belonging to different phases. This analysis allowed us to examine whether similar sequences tended to reoccur in later phases, indicating consolidation or reuse of specific patterns. For visualization, Jaccard similarity matrices were generated for each phase and across-phase comparisons, and represented as heatmaps to illustrate the structural similarity of transition patterns over time. To determine whether sequence similarity differed significantly across phases, Kruskal-Wallis tests were performed on similarity scores, when significant effects were found, Dunn's post-hoc tests with Bonferroni correction were applied to identify specific pairwise differences.

Data availability

All analysis code and data will be made publicly available on a GitHub repository upon acceptance of the manuscript. We are happy to provide the code upon request during peer review.

Acknowledgments

We thank the animal care staff at the Ernst Strüngmann Institute for their assistance in ensuring the welfare of the marmosets. A specific thank to Colleen Illing for supporting the marmosets' training. This work was supported by Bundesministerium für Bildung und Forschung (ERANET NEURON 01EW2110).

Author contributions

Conceptualization: N.E.M., F.L., J.L., Methodology: N.E.M., F.L., Investigation: N.E.M., F.L., Formal Analysis: N.E.M., Technical Support: F.Z.N., D.S., Writing - Original Draft: N.E.M., F.L., Writing - Review & Editing: N.E.M., F.L., F.Z.N., D.S., J.L., Funding Acquisition: J.L., Supervision: J.L. All authors read and approved the final manuscript.

Declaration of interests

The authors declare no competing financial or non-financial interests related to this work.

Supplemental information

Document S1. Figures S1-S5, Tables S1 and S2.

References

Abreu, F., De la Fuente, M.F.C., Schiel, N., Souto, A., 2016. Feeding ecology and behavioral adjustments: flexibility of a small neotropical primate (Callithrix jacchus) to survive in a semiarid environment. Mamm Res 61, 221–229. https://doi.org/10.1007/s13364-016-0262-4

Abreu, F., Garber, P.A., Souto, A., Presotto, A., Schiel, N., 2021. Navigating in a challenging semiarid environment: the use of a route-based mental map by a small-bodied neotropical primate. Anim Cogn 24, 629–643. https://doi.org/10.1007/s10071-020-01465-1

Abreu, F., Souto, A., Schiel, N., 2020. Wild common marmosets (Callithrix jacchus) employ spatial cognitive abilities to improve their food search and consumption: an experimental approach in small-scale space. Primates 61, 807–816. https://doi.org/10.1007/s10329-020-00826-1

Arleo, A., Rondi-Reig, L., 2007. Multimodal sensory integration and concurrent navigation strategies for spatial cognition in real and artificial organisms. J. Integr. Neurosci. 06, 327–366. https://doi.org/10.1142/S0219635207001593

Brown, T.I., Ross, R.S., Tobyne, S.M., Stern, C.E., 2012. Cooperative interactions between hippocampal and striatal systems support flexible navigation. NeuroImage 60, 1316–1330. https://doi.org/10.1016/j.neuroimage.2012.01.046

Buffalo, E.A., 2025. Visual Exploration and the Primate Hippocampal Formation. Hippocampus 35, e23673. https://doi.org/10.1002/hipo.23673

Buffalo, E.A., Movshon, J.A., Wurtz, R.H., 2019. From basic brain research to treating human brain disorders. Proceedings of the National Academy of Sciences 116, 26167–26172. https://doi.org/10.1073/pnas.1919895116

Chersi, F., Burgess, N., 2015. The Cognitive Architecture of Spatial Navigation: Hippocampal and Striatal Contributions. Neuron 88, 64–77. https://doi.org/10.1016/j.neuron.2015.09.021

Cohen, J.D., McClure, S.M., Yu, A.J., 2007. Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. Philosophical Transactions of the Royal Society B: Biological Sciences 362, 933–942. https://doi.org/10.1098/rstb.2007.2098

Constantinidis, C., Procyk, E., 2004. The primate working memory networks. Cognitive, Affective, & Behavioral Neuroscience 4, 444–465. https://doi.org/10.3758/CABN.4.4.444

Costa, V.D., Averbeck, B.B., 2020. Primate Orbitofrontal Cortex Codes Information Relevant for Managing Explore–Exploit Tradeoffs. J. Neurosci. 40, 2553–2561. https://doi.org/10.1523/JNEUROSCI.2355-19.2020

Datta, S.R., Anderson, D.J., Branson, K., Perona, P., Leifer, A., 2019. Computational Neuroethology: A Call to Action. Neuron 104, 11–24. https://doi.org/10.1016/j.neuron.2019.09.038

Daw, N.D., O'Doherty, J.P., Dayan, P., Seymour, B., Dolan, R.J., 2006. Cortical substrates for exploratory decisions in humans. Nature 441, 876–879. https://doi.org/10.1038/nature04766

Dennis, E.J., Hady, A.E., Michaiel, A., Clemens, A., Tervo, D.R.G., Voigts, J., Datta, S.R., 2021. Systems Neuroscience of Natural Behaviors in Rodents. J. Neurosci. 41, 911–919. https://doi.org/10.1523/JNEUROSCI.1877-20.2020

Domenech, P., Rheims, S., Koechlin, E., 2020. Neural mechanisms resolving exploitation-exploration dilemmas in the medial prefrontal cortex. Science 369, eabb0184. https://doi.org/10.1126/science.abb0184

Dupret, D., O'Neill, J., Pleydell-Bouverie, B., Csicsvari, J., 2010. The reorganization and reactivation of hippocampal maps predict spatial memory performance. Nat Neurosci 13, 995–1002. https://doi.org/10.1038/nn.2599

Epstein, R.A., Patai, E.Z., Julian, J.B., Spiers, H.J., 2017. The cognitive map in humans: spatial navigation and beyond. Nat Neurosci 20, 1504–1513. https://doi.org/10.1038/nn.4656

Ferecatu, A., De Bruyn, A., 2022. Understanding Managers' Trade-Offs Between Exploration and Exploitation. Marketing Science 41, 139–165. https://doi.org/10.1287/mksc.2021.1304

Friard, O., Gamba, M., 2016. BORIS: a free, versatile open-source event-logging software for video/audio coding and live observations. Methods in Ecology and Evolution 7, 1325–1330. https://doi.org/10.1111/2041-210X.12584

Gahnstrom, C.J., Spiers, H.J., 2020. Striatal and hippocampal contributions to flexible navigation in rats and humans. Brain and Neuroscience Advances 4, 2398212820979772. https://doi.org/10.1177/2398212820979772

Garber, P.A., 1992. Vertical clinging, small body size, and the evolution of feeding adaptations in the Callitrichinae. American Journal of Physical Anthropology 88, 469–482. https://doi.org/10.1002/ajpa.1330880404

Goodman, J., 2021. Place vs. Response Learning: History, Controversy, and Neurobiology. Front. Behav. Neurosci. 14. https://doi.org/10.3389/fnbeh.2020.598570

Goodroe, S.C., Starnes, J., Brown, T.I., 2018. The Complex Nature of Hippocampal-Striatal Interactions in Spatial Navigation. Front. Hum. Neurosci. 12. https://doi.org/10.3389/fnhum.2018.00250

Hunt, L.T., Daw, N.D., Kaanders, P., MacIver, M.A., Mugan, U., Procyk, E., Redish, A.D., Russo, E., Scholl, J., Stachenfeld, K., Wilson, C.R.E., Kolling, N., 2021. Formalizing planning and information search in naturalistic decision-making. Nat Neurosci 24, 1051–1064. https://doi.org/10.1038/s41593-021-00866-w

Izumi, A., Tsuchida, J., Yamaguchi, C., 2013. Spontaneous alternation behavior in common marmosets (Callithrix jacchus). Journal of Comparative Psychology 127, 76–81. https://doi.org/10.1037/a0026797

Jahn, C.I., Grohn, J., Cuell, S., Emberton, A., Bouret, S., Walton, M.E., Kolling, N., Sallet, J., 2023. Neural responses in macaque prefrontal cortex are linked to strategic exploration. PLOS Biology 21, e3001985. https://doi.org/10.1371/journal.pbio.3001985

Kell, A.J.E., Bokor, S.L., Jeon, Y.-N., Toosi, T., Issa, E.B., 2023. Marmoset core visual object recognition behavior is comparable to that of macaques and humans. iScience 26, 105788. https://doi.org/10.1016/j.isci.2022.105788

Khamassi, M., Humphries, M.D., 2012. Integrating cortico-limbic-basal ganglia architectures for learning model-based and model-free navigation strategies. Front. Behav. Neurosci. 6. https://doi.org/10.3389/fnbeh.2012.00079

Khamassi, M., Quilodran, R., Enel, P., Dominey, P., Procyk, E., 2010. Role of the frontal cortex in solving the exploration-exploitation trade-off, in: Cinquième Conférence Plénière Française de Neurosciences Computationnelles, "Neurocomp'10". Lyon, France.

Krakauer, J.W., Ghazanfar, A.A., Gomez-Marin, A., MacIver, M.A., Poeppel, D., 2017. Neuroscience Needs Behavior: Correcting a Reductionist Bias. Neuron 93, 480–490. https://doi.org/10.1016/j.neuron.2016.12.041

Lanzarini, F., Maranesi, M., Rondoni, E.H., Albertini, D., Ferretti, E., Lanzilotto, M., Micera, S., Mazzoni, A., Bonini, L., 2025. Neuroethology of natural actions in freely moving monkeys. Science 387, 214–220. https://doi.org/10.1126/science.adq6510

Laurens, J., 2022. The statistical power of three monkeys. https://doi.org/10.1101/2022.05.10.491373

Ludvig, N., Tang, H.M., Eichenbaum, H., Gohil, B.C., 2003. Spatial memory performance of freely-moving squirrel monkeys. Behavioural Brain Research 140, 175–183. https://doi.org/10.1016/S0166-4328(02)00325-X

MacDonald, S.E., Pang, J.C., Gibeault, S., 1994. Marmoset (Callithrix jacchus jacchus) spatial memory in a foraging task: Win-stay versus win-shift strategies. Journal of Comparative Psychology 108, 328–334. https://doi.org/10.1037/0735-7036.108.4.328

MacDonald, S.E., Spetch, M.L., Kelly, D.M., Cheng, K., 2004. Strategies in landmark use by children, adults, and marmoset monkeys. Learning and Motivation 35, 322–347. https://doi.org/10.1016/j.lmot.2004.03.002

Mao, D., 2022. Neural Correlates of Spatial Navigation in Primate Hippocampus. Neurosci Bull 39, 315–327. https://doi.org/10.1007/s12264-022-00968-w

Martinez-Trujillo, J., 2025. Why do primates have view cells instead of place cells? Trends in Cognitive Sciences 29, 226–229. https://doi.org/10.1016/j.tics.2024.12.007

Mehlhorn, K., Newell, B.R., Todd, P.M., Lee, M.D., Morgan, K., Braithwaite, V.A., Hausmann, D., Fiedler, K., Gonzalez, C., 2015. Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. Decision 2, 191–215. https://doi.org/10.1037/dec0000033

Miller, C.T., Gire, D., Hoke, K., Huk, A.C., Kelley, D., Leopold, D.A., Smear, M.C., Theunissen, F., Yartsev, M., Niell, C.M., 2022. Natural behavior is the language of the brain. Current Biology 32, R482–R493. https://doi.org/10.1016/j.cub.2022.03.031

Nakamura, K., Koba, R., Miwa, M., Yamaguchi, C., Suzuki, H., Takemoto, A., 2018. A Method to Train Marmosets in Visual Working Memory Task and Their Performance. Front. Behav. Neurosci. 12. https://doi.org/10.3389/fnbeh.2018.00046

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

O'Keefe, J., Nadel, L., 1978. The hippocampus as a cognitive map. Clarendon Press ; Oxford University Press, Oxford : New York.

Olton, D.S., 1977. Spatial Memory. Scientific American 236, 82–99.

Olton, D.S., Becker, J.T., Handelmann, G.E., 1979. Hippocampus, space, and memory. Behav Brain Sci 2, 313–322. https://doi.org/10.1017/S0140525X00062713

Parra-Barrero, E., Vijayabaskaran, S., Seabrook, E., Wiskott, L., Cheng, S., 2023. A map of spatial navigation for neuroscience. Neuroscience & Biobehavioral Reviews 152, 105200. https://doi.org/10.1016/j.neubiorev.2023.105200

Pereira, T.D., Shaevitz, J.W., Murthy, M., 2020. Quantifying behavior to understand the brain. Nat Neurosci 23, 1537–1549. https://doi.org/10.1038/s41593-020-00734-z

Piza, D.B., Corrigan, B.W., Gulli, R.A., Do Carmo, S., Cuello, A.C., Muller, L., Martinez-Trujillo, J., 2024. Primacy of vision shapes behavioral strategies and neural substrates of spatial navigation in marmoset hippocampus. Nat Commun 15, 4053. https://doi.org/10.1038/s41467-024-48374-2

Poucet, B., 1993. Spatial cognitive maps in animals: New Hypotheses on their structure and neural mechanisms. Psychological review 100, 163–82. https://doi.org/10.1037/0033-295X.100.2.163

Redish, A.D., 1999. Beyond the Cognitive Map: From Place Cells to Episodic Memory. A Bradford Book, Cambridge, MA, USA.

Rolls, E.T., Robertson, R.G., Georges-François, P., 1997. Spatial View Cells in the Primate Hippocampus. European Journal of Neuroscience 9, 1789–1794. https://doi.org/10.1111/j.1460-9568.1997.tb01538.x

Rolls, E.T., Wirth, S., 2018. Spatial representations in the primate hippocampus, and their functions in memory and navigation. Progress in Neurobiology 171, 90–113. https://doi.org/10.1016/j.pneurobio.2018.09.004

Sadeh, T., Shohamy, D., Levy, D.R., Reggev, N., Maril, A., 2011. Cooperation between the Hippocampus and the Striatum during Episodic Encoding. Journal of Cognitive Neuroscience 23, 1597–1608. https://doi.org/10.1162/jocn.2010.21549

Schiel, N., Souto, A., 2017. The common marmoset: An overview of its natural history, ecology and behavior. Developmental Neurobiology 77, 244–262. https://doi.org/10.1002/dneu.22458

Spinelli, S., Pennanen, L., Dettling, A.C., Feldon, J., Higgins, G.A., Pryce, C.R., 2004. Performance of the marmoset monkey on computerized tasks of attention and working memory. Cognitive Brain Research 19, 123–137. https://doi.org/10.1016/j.cogbrainres.2003.11.007

Sutton, R.S., Barto, A.G., 1998. Reinforcement Learning: An Introduction. The MIT Press, Cambridge.

Trapanese, C., Meunier, H., Masi, S., 2019. What, where and when: spatial foraging decisions in primates. Biological Reviews 94, 483–502. https://doi.org/10.1111/brv.12462

Wathen, C.N., Roberts, W.A., 1994. Multiple-pattern learning by rats on an eight-arm radial maze. Animal Learning & Behavior 22, 155–164. https://doi.org/10.3758/BF03199915

Xavier, D.P., Abreu, F., Souto, A., Schiel, N., 2024. Choosing the best way: how wild common marmosets travel to efficiently exploit resources. Anim Cogn 27, 20. https://doi.org/10.1007/s10071-024-01864-8

Yoerg, S.I., Kamil, A.C., 1982. Response strategies in the radial arm maze: Running around in circles. Animal Learning & Behavior 10, 530–534. https://doi.org/10.3758/BF03212295

Zhang, R., Quan, H., Wang, Y., Luo, F., 2023. Neurogenesis in primates versus rodents and the value of non-human primate models. Natl Sci Rev 10, nwad248. https://doi.org/10.1093/nsr/nwad248