Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Selective Memory Operations Shape History-Dependent Population Representations and Choice under Sensory Ambiguity: A Hippocampal-Inspired Dynamic Gating Model

Mingzhe Ma
Xi'an Jiaotong University
Yibo Zhao
Xi'an Jiaotong University
Chen Qiao
qiaochen@mail.xjtu.edu.cn

Xi'an Jiaotong University

Research Article

Keywords:

Posted Date: November 2nd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11531942/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Selective Memory Operations Shape History-Dependent Population Representations and Choice under Sensory Ambiguity: A Hippocampal-Inspired Dynamic Gating Model

Mingzhe Ma[1], Yibo Zhao[1], Chen Qiao[1,*]

[1]School of Mathematics and Statistics,

Xi'an Jiaotong University, Xi'an, China

[*]Corresponding author: Chen Qiao

qiaochen@mail.xjtu.edu.cn

Abstract

How does past experience continue to shape behavior? The hippocampal–entorhinal system supports memory-guided behavior, yet how selective memory operations shape population activity and history-dependent choice remains unclear. We introduce a hippocampal-inspired dynamic gating model (HDGM) that integrates predictive recurrence, online associative memory, and feedback-dependent regulation to examine how past experience shapes population representations and subsequent choice under sensory ambiguity.

HDGM exhibited distinct history-conditioned population states despite identical current sensory input. Selective disruptions of encoding and associative processing altered subsequent population activity and impaired history-dependent choice. Output-gain modulation increased linear decoding accuracy for task-relevant classes, whereas cross-trial gate reassignment reduced delay-state separation with little change in mean choice performance. Across complementary tasks, population organization, linear information accessibility, and behavioral choice showed distinct responses to selective memory operations. Together, these findings provide a computational framework for investigating how selective memory operations shape the expression of past experience and its relationship to subsequent choice.

Keywords: Hippocampal modeling; Memory gating; History-dependent choice; Representational geometry; Linear information accessibility

Introduction

Past experience can guide action even when the events that established it are no longer present in current sensory input, so the same observation may call for different choices depending on the preceding sequence. Hippocampal–entorhinal population activity can reflect past experience, but detecting history-related information alone does not establish how that information is organized, whether task-relevant information can be recovered from the activity by a linear decoder, or how it relates to behavioral choice. We therefore examine how selective memory operations shape evolving population states, the linear accessibility of task-relevant information, and subsequent choice.

Neural recordings demonstrate that hippocampal population activity reflects not only the current environment but also the history through which it was encountered. Hippocampal responses distinguish different trajectories passing through the same spatial location (Ferbinteanu and Shapiro 2003; Keinath et al. 2020; Wood et al. 2000). Learning can separate representations of latent contexts associated with similar observations (Sun et al. 2025), while working-memory-related information can be embedded within place-field activity (Varga et al. 2024). These findings establish that preceding experience can organize neural population states beyond what is specified by current sensory input. Hippocampal population recordings further reveal geometrically organized representations of spatial position and accumulated evidence (Nieh et al. 2021). Representational geometry thus provides a population-level description of how task-related states are organized and separated, while their linear accessibility and relationship to subsequent choice remain distinct questions.

Selective input and circuit regulation contribute to the formation and modification of hippocampal representations (Grienberger and Magee 2022; Udakis et al. 2025). Computational models provide complementary accounts of these processes. The Tolman–Eichenbaum machine (TEM) combines predictive and associative computations to explain structured representations (Whittington et al. 2020). GATE investigates feedback-dependent regulation of memory maintenance, readout, and adaptive reuse (Liu et al. 2026), while remapping-based sequence learning connects context-dependent representations with flexible behavior (Ito and Toyoizumi 2026). These approaches establish important computational principles for representing and regulating past experience. They also motivate a more focused question concerning memory gating and

memory readout: how do selective memory operations affect the subsequent organization of population activity, the linear accessibility of task-relevant information, and behavioral choice? In particular, these functional consequences need not change in parallel within the same recurrent memory system.

Here, we introduce a hippocampal-inspired dynamic gating model (HDGM) that couples action-dependent predictive recurrence with online associative memory and feedback-dependent regulation. The model implements three functionally distinct operations: EC-inspired selection of sensory input and memory writing, CA3-inspired iterative associative processing, and CA1-inspired regulation of associative output. These operations provide a population-level computational abstraction of selective processing in the hippocampal–entorhinal system.

To assess the contribution of adaptive regulation beyond predictive and associative computation alone, we compared HDGM with a predictive–associative model (PAM) that retains action-dependent recurrence and online associative memory but lacks HDGM's adaptive gate operations. We then used temporally targeted interventions to examine how individual operations influence subsequent population organization, linear information accessibility, and behavioral choice. This approach allowed us to distinguish the consequences of modifying the overall model configuration from those of perturbing specific memory operations.

Across complementary tasks, HDGM developed spatially organized population activity and distinct history-conditioned states under convergent sensory observations. Selective disruptions of encoding and associative processing altered subsequent population activity and impaired history-dependent choice. Output-focused analyses further revealed distinct functional effects: gain regulation increased the linear decodability of task-relevant information, whereas cross-trial gate reassignment reorganized population geometry with little change in mean choice performance. Together, these findings indicate that selective memory operations can differentially shape population organization, linear information accessibility, and behavioral choice.

Materials and Methods

Model architecture and learning objectives

I. HDGM architecture and state dynamics

The hippocampal-inspired dynamic gating model (HDGM) integrates predictive recurrence, online associative memory, and feedback-dependent regulation to examine how preceding experience shapes population activity and subsequent choice. Its architecture distinguishes three functional operations: EC-inspired sensory selection and memory writing, CA3-inspired associative processing, and CA1-inspired memory-mediated readout and output regulation. These operations interact through evolving predictive, memory, and feedback states, providing distinct computational targets for functional interventions.

At each time step, HDGM receives a sensory observation $x_t \in \mathbb{R}^{d_x}$ and an action a_t. Its internal dynamics maintain a predictive state $g_t \in \mathbb{R}^{d_h}$, a feedback state $v_t \in \mathbb{R}^{d_v}$, and an online associative-memory matrix $M_t \in \mathbb{R}^{d_h \times d_h}$. Motivated by the contribution of entorhinal input to hippocampal representations and plasticity (Basu et al. 2016; Grienberger and Magee 2022), an EC-inspired gate selectively regulates sensory encoding while action-dependent recurrence

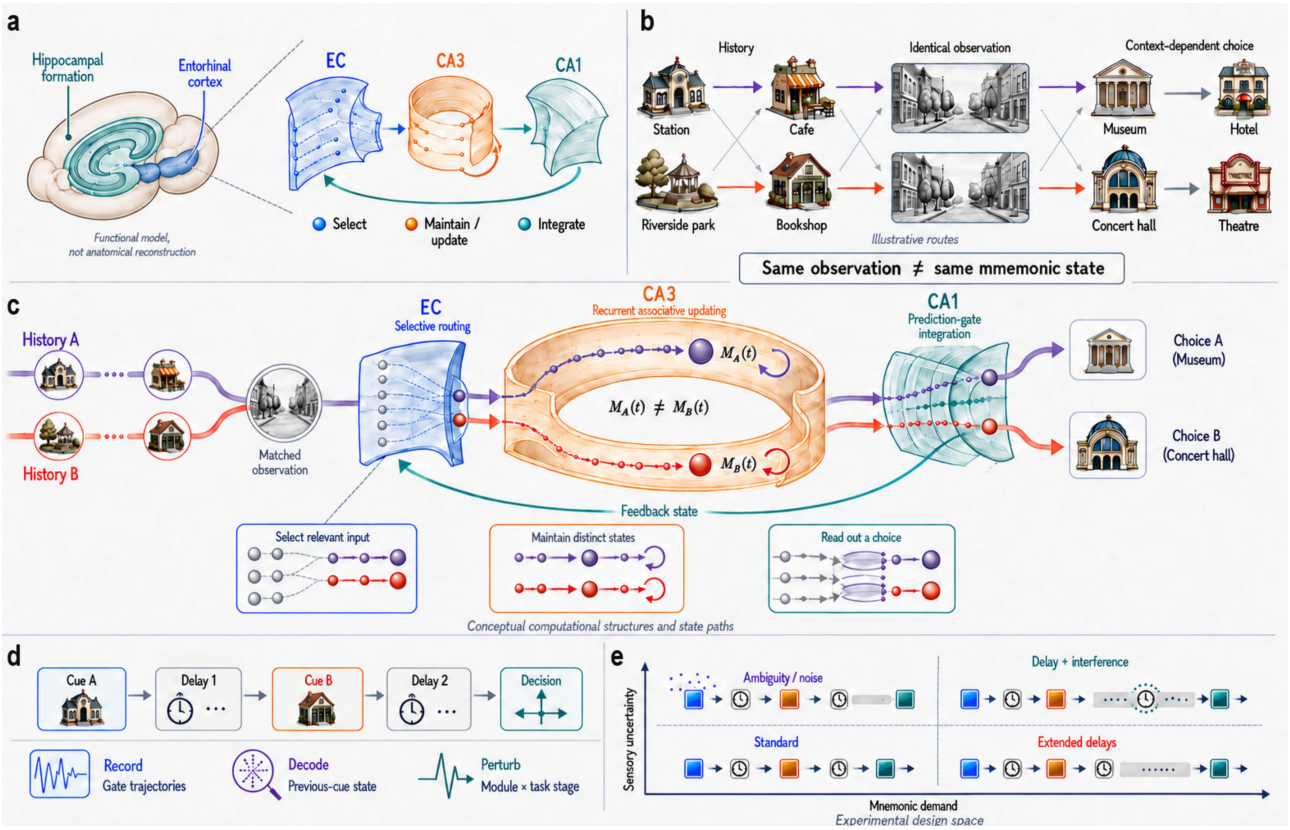


Fig. 1 Functional organization of history-dependent memory. (a) Biological inspiration and computational roles of EC, CA3, and CA1. (b) Different histories precede identical observations. (c) Selected input, associative memory, and feedback support distinct internal states. (d) Stage-specific recording, decoding, and perturbation design. (e) Schematic memory-demand and sensory-uncertainty conditions. Anatomical, route, and state panels are conceptual schematics. Author-designed schematic rendered using AI; see Methods for details

generates a predictive state:

$$\begin{cases} e_t = \text{sigmoid}(W_e[x_t, g_{t-1}, v_{t-1}] + b_e) \\ z_t = \tanh(W_x(e_t \odot x_t) + b_x) \\ g_t^- = \tanh(W_g g_{t-1} + W_a \,\text{Emb}(a_t) + b_g) \end{cases} \tag{1}$$

Here, $e_t \in (0,1)^{d_x}$ controls sensory selection, $z_t \in \mathbb{R}^{d_h}$ is the encoded sensory activity, and g_t^- is the predictive state before associative integration. All W and b terms denote learned projections and biases, respectively; Emb is a learned action embedding, brackets indicate concatenation, and $\odot$ denotes elementwise multiplication.

Inspired by the recurrent associative organization of hippocampal CA3 (Nakazawa et al. 2002; Rolls 2013), HDGM iteratively combines encoded sensory activity with information retrieved from online memory:

$$\begin{cases} h_t^{(0)} = z_t \\ c_t^{(\ell)} = \text{sigmoid}\Big(W_c[h_t^{(\ell)}, g_{t-1}] + b_c\Big) \\ r_t^{(\ell)} = \tanh\Big(\kappa_{\text{CA3}} h_t^{(\ell)} + M_{t-1} h_t^{(\ell)}\Big) \\ h_t^{(\ell+1)} = (1 - c_t^{(\ell)}) \odot h_t^{(\ell)} + c_t^{(\ell)} \odot r_t^{(\ell)} \end{cases} \tag{2}$$

The iteration index $\ell = 0, \cdots, L-1$ denotes successive associative updates, with $p_t = h_t^{(L)}$ representing the final associative state. The gate $c_t^{(\ell)}$ controls the contribution of the candidate update $r_t^{(\ell)}$ relative to the preceding associative state. The coefficient κ_{CA3} scales the local-state contribution, whereas $M_{t-1} h_t^{(\ell)}$ represents memory-mediated processing. The predictive and memory states available at the beginning of the time step remain fixed during these iterations. The resulting associative activity is integrated with the predictive state before a separate memory-mediated readout. Motivated by the convergence of entorhinal and CA3 inputs onto hippocampal CA1 (Colgin et al. 2009), these computations are expressed as

$$\begin{cases} g_t = g_t^- + \eta p_t \\ q_t = \tanh(\kappa_{\text{CA1}} g_t + M_{t-1} g_t) \end{cases} \tag{3}$$

Here, η determines the associative contribution to the integrated predictive state, while q_t denotes the associative readout before output modulation. The direct predictive contribution

$\kappa_{\mathrm{CA1}} g_t$ and memory-mediated contribution $M_{t-1} g_t$ provide separable computational components for subsequent interventions.

To examine how output modulation shapes the expression of associative information, HDGM applies a CA1-inspired output gate. This abstraction is motivated by evidence that local CA1 circuits regulate place-cell plasticity and remapping (Udakis et al. 2025), and is defined as

$$\begin{cases} k_t = \mathrm{sigmoid}(W_k[g_t, q_t, v_{t-1}] + b_k) \\ o_t = \mathrm{ReLU}(q_t \odot (c_1 + c_2 k_t) - \beta) \end{cases} \tag{4}$$

The gate $k_t \in (0,1)^{d_h}$ regulates output gain, and $o_t \in \mathbb{R}^{d_h}$ is the resulting population activity. The coefficients c_1 and c_2 determine baseline and gate-dependent gain, respectively, while β is the output threshold. Separating the associative readout q_t from the regulated output o_t allows the influence of output modulation on population organization and linear information accessibility to be examined independently.

Output activity subsequently updates the feedback state, while EC-dependent selection regulates the writing of new associations:

$$\begin{cases} \bar{e}_t = \dfrac{1}{d_x} \displaystyle\sum_{j=1}^{d_x} e_{t,j} \\ v_t = \rho v_{t-1} + W_f o_t + b_f \\ M_t = \lambda M_{t-1} + \bar{e}_t p_t p_t^{\top} \end{cases} \tag{5}$$

Here, $\bar{e}_t$ is the mean EC gate activity and determines the strength of associative-memory writing. The coefficients ρ and λ control feedback and memory retention, respectively. Retrieval uses M_{t-1}, while newly written associations become available at the following time step. Predictive, feedback, and associative-memory states are initialized to zero at the beginning of each sequence.

The output population state also generates next-observation predictions through a learned linear mapping, $\hat{x}_{t+1} = W_d o_t + b_d$. Together, Eqs. (1)–(5) define a recurrent loop linking sensory selection, associative retrieval, output regulation, and memory writing. Their functional contributions are examined through the targeted interventions described in the experimental protocols.

II. Predictive–associative comparison model

The predictive and associative computations in HDGM provide a basis for examining the contribution of adaptive regulation. We therefore constructed a predictive–associative model (PAM) that retains action-dependent predictive recurrence, iterative associative processing, and online memory updating, while replacing the adaptive regulatory operations with ungated counterparts. The resulting differences in sensory encoding, associative updating, output activity, and memory writing are defined by

$$\begin{cases} z_t = \tanh(W_x x_t + b_x) \\ h_t^{(\ell+1)} = \tanh\left(\kappa_{\mathrm{CA3}} h_t^{(\ell)} + M_{t-1} h_t^{(\ell)}\right) \\ o_t = q_t \\ M_t = \lambda M_{t-1} + p_t p_t^{\mathsf{T}} \end{cases} \tag{6}$$

The associative iterations retain $h_t^{(0)} = z_t$ and $p_t = h_t^{(L)}$. The predictive-state, associative-readout, and observation-prediction updates follow the corresponding HDGM equations. PAM does not include the output-derived feedback state v_t.

III. Learning objectives and parameter settings

Trainable model parameters were optimized using task-specific learning objectives, whereas the associative-memory matrix was updated online through the state dynamics. Next-observation prediction was trained using mean-squared prediction error, with the temporal reduction convention determined by the corresponding task protocol. In the four-context assay, the event-aligned behavioral readout was trained using four-class cross-entropy. HDGM additionally received two auxiliary objectives that encouraged stronger EC-controlled writing during informative cue periods relative to selected non-cue intervals.

The four-context learning objectives were

$$\begin{cases} \mathcal{L}_{\mathrm{PAM}} = \mathcal{L}_{\mathrm{pred}} + \omega_{\mathrm{cls}} \mathcal{L}_{\mathrm{class}} \\ \mathcal{L}_{\mathrm{HDGM}} = \mathcal{L}_{\mathrm{pred}} + \omega_{\mathrm{cls}} \mathcal{L}_{\mathrm{class}} + \omega_B \mathcal{L}_B + \omega_{\mathrm{rank}} \mathcal{L}_{\mathrm{rank}} \end{cases} \tag{7}$$

Here, $\mathcal{L}_{\mathrm{pred}}$ is the next-observation prediction loss and $\mathcal{L}_{\mathrm{class}}$ is the four-context classification loss. The auxiliary terms $\mathcal{L}_B$ and $\mathcal{L}_{\mathrm{rank}}$ impose cue-related writing-strength margins, with

ω_{cls}, ω_B, and ω_{rank} determining their relative contributions. The corresponding loss weights, temporal reductions, and model parameters are reported in Table 1 (Appendix III). Training schedules, evaluation checkpoints, and decoder configurations are provided in Tables 2 and 3 (Appendix III).

The spatial assay used observation prediction without position supervision, while the shared-corridor assay combined prediction and choice cross-entropy. The stateless behavioral comparator used choice classification alone. Model parameters and task-specific loss conventions are summarized in Table 1, while training and decoder protocols are reported in Tables 2 and 3 (Appendix III).

Experimental design

Four complementary experimental designs examined how selective memory operations relate preceding experience to population organization, information accessibility, and behavioral choice. The experiments addressed history-dependent decisions under sensory ambiguity, the functional consequences of selective memory interventions, the organization of spatial and history-conditioned population activity, and the accessibility of historical information under sensory perturbations. Models were trained separately for each assay using the corresponding learning objectives defined above.

Population-level measurements and common evaluation principles

We distinguished three outcomes: population organization, linear information accessibility, and behavioral choice. Population organization describes the structure of model activity across task conditions, with population geometry quantified by the separation between condition-specific states. Linear information accessibility was measured by held-out accuracy of a specified linear decoder. Behavioral choice was generated by the task-specific readout and evaluated using choice accuracy and other reported performance metrics. These measures capture distinct aspects of memory-guided computation and were analyzed separately.

The model output state served as the primary population representation, with additional internal states examined when required by the experimental question. Linear probes were evaluated on independently generated training, validation, and test datasets, with normalization fitted on training data and decoder selection based on validation performance. Model parameters re-

mained fixed during population analyses, while recurrent and associative-memory states evolved according to the evaluation conditions.

Linear information accessibility was measured as the recoverability of task-relevant class labels by the specified linear decoder.

I. Four-context task and history-dependent choice

The four-context assay examined whether preceding sensory events could support different decisions when distinct experiences converged on the same current observation. Each trial contained two temporally separated binary cues, Cue A and Cue B, whose combined identities determined one of four possible choices. The target was defined as $y = 2A + B$, where $A, B \in \{0, 1\}$, corresponding to a chance classification level of 25%. Trials followed a shared sensory trajectory with variable delays separating cue presentation and decision. Cue identities did not alter the underlying path, action sequence, or non-cue observations, so successful classification required information about preceding events rather than the terminal sensory input alone.

History-dependent choice was generated using an event-aligned behavioral readout trained jointly with the recurrent model. Separate attention-pooling modules summarized output population activity during the Cue A and Cue B windows, while an additional representation was obtained immediately before the decision. These representations were transformed and combined to predict the four-class target. Event boundaries were supplied to the behavioral readout but not to the recurrent state updates. Training combined next-observation prediction with four-class cross-entropy, with HDGM additionally receiving the EC-controlled memory-writing objectives defined in the model section. HDGM and PAM were evaluated under matched training budgets and data-generation protocols, using choice accuracy, macro-F1, class-specific performance, and acquisition trajectories.

To determine whether classification depended on historical information, frozen-model evaluation compared intact history with history-reset and restricted-information controls. Decision-only, observation-only, and timing-only conditions assessed whether the task could be solved using restricted population activity, current sensory information, or temporal cues rather than preceding sensory history. The effects of extended memory demands were examined by introducing additional matched-observation repeats within the delay periods while keeping learned model parameters fixed. These controls assessed the contribution of preceding experience to behavioral

choice and its persistence across longer delays. Exact task-generation rules and readout specifications are provided in Appendix II (*Four-context task generation and behavioral readout*), with training settings summarized in Table 2 (Appendix III).

II. Memory-operation ablations and gate perturbations

To examine how distinct memory operations influence the later expression and behavioral use of past experience, we selectively disrupted EC-inspired sensory selection and memory writing, CA3-inspired associative processing, and CA1-inspired memory-mediated readout. Interventions were performed in trained four-context models with learned parameters held fixed and matched inputs supplied to control and intervention conditions. Functional ablations targeted cue encoding, delay maintenance, or terminal readout, allowing the consequences of disrupting the same computation at different task stages to be compared. Encoding interventions covered the Cue A and Cue B windows, maintenance interventions covered the intervening delays, and readout interventions targeted the terminal decision window.

Each functional ablation removed a defined contribution to the corresponding memory computation. The EC intervention suppressed selected sensory input and its associated online memory writing. The CA3 intervention removed the associative-memory contribution $M_{t-1}h_t^{(\ell)}$ while preserving the local-state contribution, learned iteration gate, and recurrent associative updates. The CA1 memory-mediated intervention removed $M_{t-1}g_t$ from associative readout while preserving the direct predictive contribution $\kappa_{\mathrm{CA1}}g_t$; a complementary intervention removed the predictive contribution while retaining memory-mediated retrieval.

The effects of memory-operation interventions were assessed using model-choice accuracy, population geometry, and linear decoding. Choice effects were calculated as control accuracy minus intervention accuracy, with positive values indicating an impairment under intervention. In the four-context assay, population geometry was quantified by normalized distances between Cue A population centroids during the first delay, following position-specific mean subtraction and standardization of trial-level responses. Linear probes assessed the recoverability of cue identity and four-context labels from activity at designated time windows.

Complementary gate perturbations examined the contributions of adaptive regulation. Within designated task windows, EC, CA3, and CA1 gates were replaced by zero, one, or independently calibrated per-unit mean values. These substitutions tested the consequences of suppressing

regulated computations, fixing gates open, or replacing dynamic gate activity with constant regulatory configurations. Fixed-gate retraining was conducted separately to examine learning under altered configurations rather than acute disruption of an established model. During retraining, the selected gate remained open while the remaining parameters adapted; the auxiliary writing penalties were removed in the fixed-EC condition. Cross-trial gate reassignment additionally disrupted the alignment between trial-specific gate activity and evolving population states by cyclically exchanging gate values across trials within the intervention window. Its effects were assessed using the same behavioral and population measurements.

An independent CA1 output-regulation analysis examined how gate-dependent output transformation influenced the linear accessibility of task-relevant information. Linear decoding was compared between pre-output associative activity q_t and output-regulated population activity o_t. The learned CA1 gate was further compared with fixed-open, calibrated-mean, causal running-mean, and one-step-lagged substitutions. Calibration, decoder fitting, and behavioral evaluation used separate data streams, with dynamic substitutions derived from intact reference gate sequences under matched inputs. These analyses distinguished the contribution of output regulation to the accessibility and organization of population information from its influence on behavioral choice. Detailed intervention and gate-substitution procedures are provided in Appendix II (*Memory-operation interventions and CA1 validation*), with training and decoder configurations summarized in Tables 2 and 3 (Appendix III).

III. Spatial and history-conditioned population representations

Spatial population organization was examined using two complementary tasks that distinguished representations of environmental position from representations conditioned on preceding experience. The first task used a partially observable ring environment with repeated sensory observations and intermittent landmarks. HDGM and PAM learned next-observation prediction without explicit position supervision, allowing spatial structure to be evaluated in population activity after learning. The second task used a branching environment in which distinct preceding trajectories converged onto shared corridor locations, providing a setting in which identical current sensory observations could correspond to different histories.

In the ring environment, spatial organization was assessed through population-map stability and position-dependent separation under matched sensory–action conditions. Split-half sta-

bility measured the reproducibility of position-related population activity across disjoint trial subsets, with a structured positional permutation serving as the reference distribution. Population separation was evaluated by comparing activity associated with different positions under matched sensory and action conditions. To distinguish position-related structure from components explained by immediate sensory input and action, population activity was additionally residualized using a fitted linear sensory–action model before the stability and separation analyses. Position decoding was evaluated using frozen-model representations, with position labels excluded from recurrent model training.

The shared-corridor task examined whether previous branch experience influenced population activity and subsequent choices at common spatial locations. Trials included alternative branch histories followed by traversal of the same corridor, with branch identity excluded from the recurrent model inputs. A task-specific choice readout was trained to predict history-dependent actions at designated decision states, using a combination of observation prediction and choice classification. Population activity was compared across intact-history, history-reset, and observation-only conditions. History reset cleared the relevant predictive, associative-memory, and feedback states, while linear probes assessed preceding branch identity at matched corridor positions and across selected internal model states.

History-conditioned population geometry was quantified by comparing the two branch-history representations at the same corridor position. Between-history separation was measured as the Euclidean distance between history-specific population centroids, normalized by the square root of the population dimension. Within-history dispersion was calculated from the average distance between individual population vectors and their corresponding history centroids. The ratio of between-history separation to within-history dispersion was then averaged across eligible corridor positions. The separation-to-dispersion ratio quantified geometric separation relative to within-history variability. Branch-identity decoding provided a complementary measure of linear decodability, while the task-specific choice readout quantified decision performance.

Exploratory structural controls examined how history-dependent representations and behavior changed when adaptive gating or online associative memory was modified during training. The variants retained predictive recurrence and were evaluated using pooled and condition-specific history probes. The exploratory structural-control analyses are presented in Appendix I. Pop-

ulation geometry and structured-control calculations are detailed in Appendix II (*Population measurements and spatial controls*).

IV. Sensory perturbations and history accessibility

Sensory perturbations examined how the structure and reliability of current observations influenced access to preceding experience and subsequent choice. Two independent perturbation families were considered: structural aliasing progressively increased similarity between distinct sensory cues, whereas observation noise introduced stochastic variation around the existing sensory representations. Both manipulations were evaluated from a shared reference condition using trained model checkpoints without additional parameter updates. Recurrent, associative-memory, and feedback states continued to evolve under each sensory condition, allowing the consequences of altered sensory information to be examined within the learned system.

History accessibility was evaluated using a fixed linear transfer decoder trained under the reference sensory condition. The decoder and its feature normalization were held constant across perturbations, allowing performance differences to be compared under a common readout. For each perturbation condition, the history-accessibility score was defined as the difference between full-history and history-reset decoding accuracy:

$$A(\xi) = \mathrm{Acc}_{\mathrm{full}}(\xi) - \mathrm{Acc}_{\mathrm{reset}}(\xi) \tag{8}$$

Here, $\mathrm{Acc}_{\mathrm{full}}(\xi)$ and $\mathrm{Acc}_{\mathrm{reset}}(\xi)$ denote preceding-history classification accuracy under intact and reset recurrent histories, respectively, for perturbation condition ξ. The history-accessibility score $A(\xi)$ therefore quantifies the difference in decoding accuracy between these two conditions. For each perturbation family, $A(\xi)$ was summarized by the normalized area under the curve (AUC) across the tested sensory conditions.

Behavioral choice accuracy was evaluated under the corresponding sensory conditions and compared with population-level accessibility measurements. To examine their relationship across perturbations, accessibility and choice accuracy were centered within each independent training run before estimating their descriptive association. Comparisons across HDGM, PAM, and the stateless behavioral model followed the corresponding matched evaluation protocols.

Statistical analysis and reproducibility

Independent training runs served as the statistical units for model comparisons and interventions. The primary experiments used ten independent training runs per protocol, with matched evaluations within each run; exploratory structural controls used three development runs. Paired differences were assessed using two-sided exact sign-flip tests, with Holm correction applied within predefined comparison families. Intervention effects were expressed as control minus intervention measurements, and standardized paired effects were calculated from the mean and sample standard deviation of run-level differences.

Uncertainty was quantified using the confidence-interval procedures specified for each analysis, including normal-approximation, Student-t, and percentile-bootstrap intervals. Bootstrap resampling used independent training runs as the sampling units and preserved paired observations and within-run dependence. The relationship between history accessibility and behavioral choice was evaluated descriptively using within-run centered measurements and seed-cluster bootstrap intervals. Comparison families, effect-size conventions, and uncertainty procedures are detailed in Table 4 (Appendix III).

All analyses were performed using stored simulation outputs, with population decoders fitted independently of core-model optimization. Figures were generated using Python, NumPy, and Matplotlib. Exploratory structural controls are presented in Appendix I, detailed experimental and analytical procedures in Appendix II, and model notation and analysis abbreviations in Table 5 (Appendix III).

AI-assisted procedures and manuscript preparation

The authors designed the study, developed the model and analysis procedures, conducted the simulations, and interpreted the results. OpenAI Codex, using GPT-5.6 and GPT-6, provided assistance with code implementation and debugging, inspection of archived outputs and organization of reproducibility materials. AI assistance followed computational procedures specified by the authors.

Fig. 1 was designed by the authors and rendered using OpenAI's image-generation tool accessed through GPT-6. The underlying image-generation model was not separately recorded. The design specifications and generation records were retained. The authors reviewed the illustration

for consistency with the model and task descriptions.

AI tools also assisted with language editing and translation. Numerical results were produced by simulation and analysis scripts. The authors reviewed the AI-assisted materials and retain responsibility for the accuracy, integrity, and interpretation of the work.

Results

Across complementary tasks, HDGM developed spatially organized and history-conditioned population activity. In the four-context task, encoding interventions linked sensory access and associative processing to subsequent population activity and choice. Cross-trial CA1 gate reassignment altered delay-state separation with little change in mean choice accuracy. A separate output validation showed that gain modulation and thresholding increased linear information accessibility. Sensory perturbations further examined how access to preceding experience depended on cue structure. Together, these analyses connect selective memory operations to the organization, accessibility, and behavioral use of past experience.

History-dependent choices required preceding cue information

HDGM learned to distinguish cue histories that required different choices despite identical current observations (Fig. 2a). Early acquisition exceeded PAM by 9.79 percentage points (pp) in paired AUC (95% bootstrap interval [3.02, 16.54] pp). At update 150, mean accuracy reached 47.42%, compared with 32.93% for PAM and a 25% chance level (Fig. 2b). Performance varied across contexts, with the lowest accuracy for A0B0 and the highest for A1B0 (Fig. 2c).

The population readout revealed a corresponding dependence on preceding cue information. Observation-only and timing-only probes remained near chance, whereas cue- and decision-aligned activity supported higher classification accuracy (Fig. 2d). Thus, the event-aligned population responses contained history-related information beyond that provided by the current observation or elapsed time alone.

HDGM also maintained relatively stable choice accuracy across the tested delay extensions, while PAM showed a greater decline (Fig. 2e). The observed difference in delay-related performance changes was 6.1 pp ($p_{\mathrm{Holm}} = .070$).

These results motivated a closer examination of how selective memory operations contribute to subsequent population activity and choice.

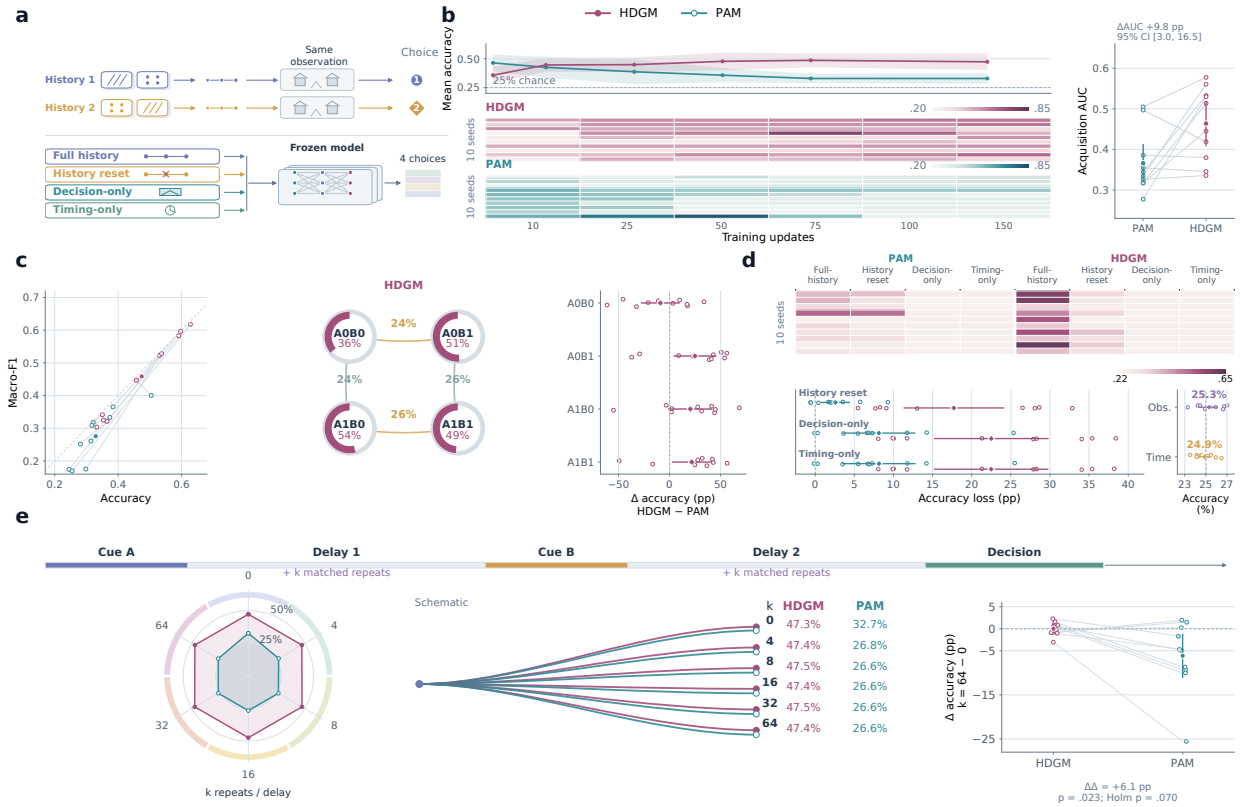


Fig. 2 History-dependent classification and information controls. (a) Four-context task and controls. (b) Acquisition curves, seed heatmaps, and acquisition area under the curve (AUC). (c) Accuracy, macro-F1, confusion structure, and class-wise contrasts. (d) Frozen history controls and separately trained observation/timing probes. (e) Frozen evaluation with additional matched-observation repeats in each delay. Ten paired seeds; chance accuracy is 25%. Purple: HDGM; teal: PAM; pp: percentage points. Accuracy intervals are mean $\pm 1.96 s/\sqrt{10}$. AUC and other panel-specific statistics are defined in Materials and Methods

Selective memory operations shaped population activity, accessibility, and choice

Selective memory operations during cue encoding made substantial contributions to subsequent history-dependent choice. Disrupting EC sensory access and memory writing, CA3 associative iteration, or the CA1 memory-mediated readout reduced normal-delay accuracy by 22.44, 20.38, and 13.30 percentage points (pp), respectively (Fig. 3a). All three effects survived Holm correction ($p_{\text{Holm}} = .018$). Corresponding interventions during maintenance and terminal readout produced smaller mean losses, with all six paired stage contrasts surviving correction ($p_{\text{Holm}} = .018$, exploratory family). These stage-dependent effects identify cue encoding as a critical period for memory operations contributing to later choice.

Encoding-stage disruptions also changed the population states expressed after the cues. The combined EC intervention reduced Cue A decoding during the first delay, delay-state separation, and pre-decision class decoding ($p_{\text{Holm}} = .018$ for each). Disrupting the CA1 memory-mediated readout likewise reduced delay-state separation. Within this readout, removing the associative contribution $M_{t-1}g_t$ produced a 13.32 pp greater behavioral loss than removing the direct predictive contribution $0.5g_t$ ($p_{\text{Holm}} = .018$, exploratory family). These results link selective sensory encoding and associative processing to subsequent changes in population geometry and the linear decodability of history-related information.

Gate manipulations revealed how the operating conditions of memory processing influenced learned behavior. Encoding-window gate zeroing produced substantial accuracy losses, whereas per-unit mean replacement preserved much of the learned performance (Fig. 3b). Acute maximal opening also impaired choice, reducing accuracy by 23.72 pp for EC and 25.23 pp for CA3 ($p_{\text{Holm}} = .018$ and .012, respectively). Retraining with individual gates fixed open supported acquisition under the altered configurations, with fixed-EC training additionally removing the writing penalties. These findings distinguish the gate settings supporting an established solution from the configurations under which a new solution can be learned.

In the cross-trial reassignment experiment, exchanging CA1 gain signals across trials during maintenance reduced delay-state separation ($d_z = 0.98$, $p_{\text{Holm}} = .018$), while mean choice accuracy changed by less than 0.01 percentage points (Fig. 3c). Trial-specific gain alignment thus shaped population separation alongside largely preserved mean choice performance.

In a separate CA1 output-validation analysis, gain modulation and thresholding increased linear class-decoding accuracy by 6.45 percentage points relative to pre-output associative activity ($p_{\text{Holm}} = .008$; Fig. 3d). Decoding with the learned gate also exceeded the fixed-open and independently calibrated mean-gate conditions by approximately 6–7 percentage points ($p_{\text{Holm}} = .031$ for both comparisons; Fig. 3e). Mean decoding was higher with the running-mean substitute than with the learned gate ($p_{\text{Holm}} = .070$); the one-step-lagged comparison yielded $p_{\text{Holm}} = .50$. These results show that output-gain configuration influenced the linear decodability of task-relevant classes in the model.

Cross-trial reassignment altered population separation with little change in mean choice accuracy, while independent output validation identified changes in linear decoding.

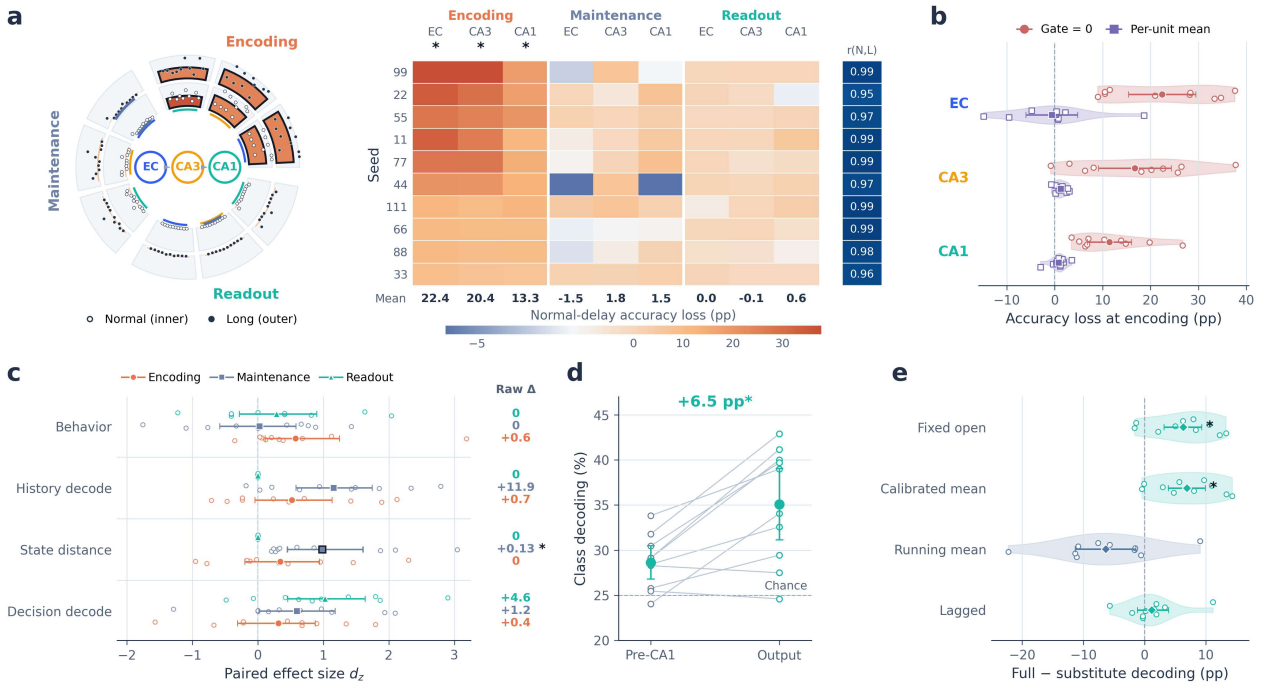


Fig. 3 Stage-specific operation and gate interventions. (a) Operation cuts: normal/long-delay rings and uncapped normal-delay losses. (b) Encoding-window gate zeroing and per-unit mean clamping. (c) Cross-trial CA1 gate reassignment effects on behavior and population measures. (d) Pre-CA1 versus output class decoding. (e) Full-output decoding minus substitute-gate decoding. Ten paired seeds per experiment; points show seeds and lines join paired observations. Positive losses indicate impairment. Chance decoding in (d) is 25%; stars and black outlines denote panel-specific Holm-adjusted $p < .05$. Intervals and display transformations are specified in Materials and Methods. Panels (d)–(e) use a separate CA1 validation experiment

Population activity encoded spatial structure and preceding experience

HDGM developed reproducible spatial population organization through observation-prediction learning. In the 20-position ring, position-related activity emerged without explicit position labels (Fig. 4a,b). After linear sensory–action residualization, split-half population-map stability exceeded the structured null ($d_z = 2.77$, $p_{\text{Holm}} = .004$). Spatial separation under matched sensory–action conditions also exceeded its structured null ($d_z = 2.52$, $p_{\text{Holm}} = .004$; Fig. 4f). These effects reveal spatial population structure beyond the components captured by a linear sensory–action mapping.

Preceding experience further distinguished population states at shared sensory locations. In the separately trained shared-corridor task, full-history decoding reached 76.83%, whereas history-reset and observation-only conditions remained near the 50% chance level (Fig. 4c,f). Both paired contrasts survived Holm correction ($p_{\text{Holm}} = .004$). Resetting the predictive, associative-

memory, and feedback states reduced history decoding to near chance under unchanged current observations.

HDGM developed splitter-like, history-conditioned population responses that distinguished preceding branches at the same location. Branch identity was linearly decodable from both output activity and the output-derived feedback state at corridor and choice locations (Fig. 4d). Position-matched population geometry further quantified the separation between history-conditioned states relative to within-branch dispersion (Fig. 4e). This pattern captures a population-level signature of history-dependent coding: preceding trajectory information remains expressed in model activity when the current location is shared.

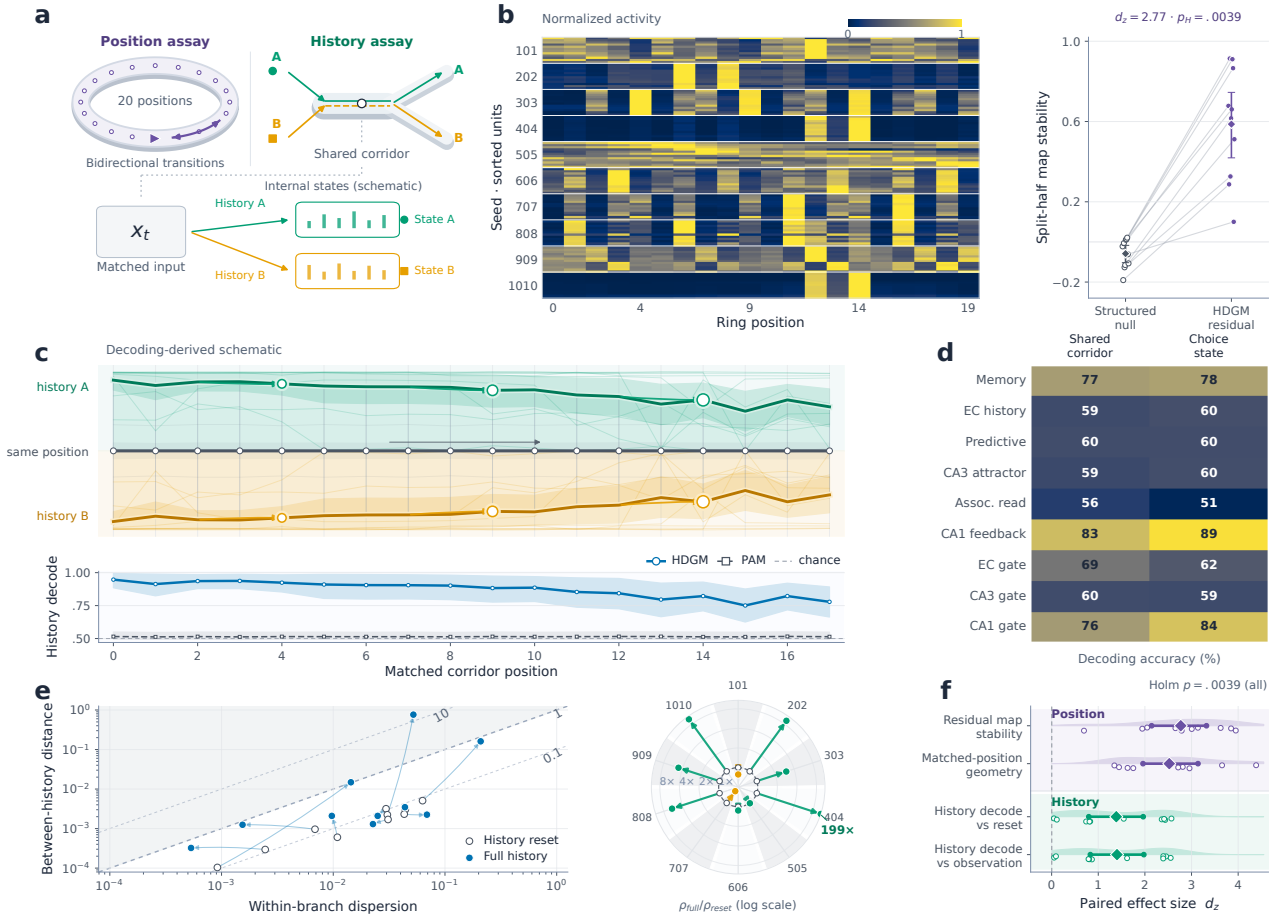


Fig. 4 Spatial and history-conditioned population organization. (a) Separately trained ring-position and shared-corridor assays. (b) Position maps and residual-map stability versus a structured null. (c) Decoding-derived history schematic and measured position-wise decoding. (d) Mean history decoding across model states and gates. "Memory," "CA3 attractor," and "CA1 feedback" denote o_t, $p_t = h_t^{(5)}$, and v_t, respectively; values are rounded percentages. (e) Between-history separation versus within-branch dispersion and full/reset separation-ratio gains. (f) Paired standardized effects against the specified controls; "Matched-position geometry" denotes spatial separation under matched sensory–action conditions. Ten seeds per assay; chance history decoding is 50%. Open/filled symbols in e denote reset/full history. Intervals and Holm families are specified in Materials and Methods. The upper branches in (c) schematically visualize decoding-derived history information

History-conditioned activity was accompanied by history-dependent choice. At the later behavioral checkpoint, HDGM reached approximately 64% mean choice accuracy, compared with 52% for PAM ($p_{\text{Holm}} = .063$), while the stateless comparator remained near the 50% chance level. The task thus connected the distinction between preceding branches to a subsequent behavioral response.

Exploratory structural controls revealed additional differences in how learned solutions expressed and used history. Across three seeds, high pooled history decoding accompanied variable choice performance; the exploratory results are presented in Appendix I. In one seed, predictive recurrence retained history-dependent choice after adaptive gates and online associative memory were removed. Position- and trip-specific probes also recovered history information under conditions with weak pooled decoding or low choice accuracy. These observations highlight the contribution of recurrent dynamics and readout configuration to the measured accessibility and behavioral relevance of historical information.

Sensory perturbations differentially shaped history accessibility and choice

The history-accessibility score showed distinct responses to cue aliasing and observation noise. Increasing cue similarity reduced history decoding and shifted choice accuracy toward chance, whereas observation noise preserved greater access to preceding experience across the tested conditions (Fig. 5 a–c,f). These results show that the two sensory perturbations differentially affected history-related information recovered by the fixed transfer decoder.

HDGM yielded a larger history-accessibility score than PAM across both perturbation scans. The normalized AUC difference was 0.0411 under cue aliasing and 0.2152 under observation noise (Fig. 5d). Both contrasts survived Holm correction ($p_{\text{Holm}} = .008$). Under the tested conditions, HDGM therefore showed a larger decoding advantage for intact history relative to the history-reset condition.

History accessibility also showed a positive association with behavioral performance. Across seven unique conditions per seed, within-seed centered accessibility and choice accuracy yielded a correlation of $r = .69$ (95% seed-cluster bootstrap interval $[-.05, .94]$; Fig. 5e). Conditions with greater accessible history information tended to exhibit higher choice accuracy, linking

population readout and behavioral performance across the tested sensory conditions.

Together, these findings reveal that access to past experience depends on sensory structure and the computations expressing historical information. Across intervention, population, and perturbation analyses, HDGM showed dissociable effects on population organization, information accessibility, and choice, highlighting how memory processing shapes the behavioral relevance of representations. Exploratory structural-control results are presented in Appendix I (Fig. 6).



Fig. 5 History accessibility under aliasing and observation noise. (a) Independent perturbations of one frozen system. (b) History-decoding versus behavioral coordinates. (c) Seed-wise and mean choice accuracy. (d) History access: full-history minus history-reset transfer decoding. (e) Seed-centered access–behavior association. (f) Standardized full-minus-reset access and separation effects, and choice accuracy relative to 50% chance. Ten seeds; seven unique conditions per seed; alias $\alpha = 0$ and noise $\sigma = .05$ share the baseline. Chance choice accuracy is 50%. Intervals, Holm families, and graphical transformations are defined in Materials and Methods. Shaded regions and connecting branches are illustrative guides to the displayed conditions

Discussion

HDGM describes how stored experience participates in ongoing computation: selected input updates associative memory, stored associations shape subsequent population activity, and output regulation changes the information available to readout. TEM formalizes how predictive and associative computations support structured representations (Whittington et al. 2020), while GATE examines feedback-dependent regulation of memory processing (Liu et al. 2026). Encoding and associative interventions altered subsequent population states and choice, whereas CA1-inspired output regulation produced distinct effects on population geometry and linear information accessibility.

The shared-corridor result provides a computational counterpart to hippocampal history-dependent coding, in which neural activity can distinguish trajectories or episodes occurring at the same location (Ferbinteanu and Shapiro 2003; Keinath et al. 2020; Wood et al. 2000). In HDGM, branch-specific population states emerged through the interaction of predictive recurrence, associative memory, and feedback, offering a candidate account of how preceding trajectory information remains expressed under shared sensory input. The distinct effects of output regulation further suggest that representational organization and linear information accessibility are separable properties, consistent with population-coding accounts in which representational changes can coexist with stable functional outputs (Kaufman et al. 2014; Martín-Sánchez et al. 2025).

During history-dependent choice, perturbing cue-related entorhinal input is predicted to alter subsequent history-conditioned activity and behavior, whereas manipulating CA1-related output regulation may reorganize population geometry without necessarily changing mean choice performance. Behavioral changes may depend more closely on disruptions of task-relevant readout than on overall population separation. Simultaneous hippocampal–prefrontal recordings and temporally targeted circuit perturbations could test these predictions, extending existing evidence for entorhinal contributions to hippocampal plasticity and hippocampal–prefrontal communication (Grienberger and Magee 2022; Young et al. 2025). HDGM links the expression of past experience in population activity to the operations that shape its accessibility and contribution to choice, providing an operation-level account of memory-guided behavior under sensory ambiguity.

Statements and Declarations

Funding Information

This work was supported by the National Natural Science Foundation of China (Grant Nos. 12671633, 12226007, and 12271429) and the National Key R&D Program of China (Grant No. 2024YFA1012000).

Competing Interests

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

Data and Code Availability

The source code and publicly released reproducibility materials, including simulation data and analysis outputs supporting the reported findings, are available at https://github.com/M0425/hdgm-hippocampal-memory.

Author Contributions

Conceptualization, M.M. and C.Q.; methodology and model design, M.M. and C.Q.; software, M.M.; validation, M.M. and Y.Z.; formal analysis, M.M. and Y.Z.; investigation, M.M.; data curation, M.M. and Y.Z.; writing—original draft preparation, M.M.; writing—review and editing, M.M., Y.Z., and C.Q.; visualization and figure preparation, M.M. and Y.Z.; supervision, C.Q.; funding acquisition, C.Q. All authors read and approved the final manuscript.

Acknowledgments

The authors acknowledge the computational resources provided by their research group, which supported model training and analysis. We thank the researchers who made the TEM and GATE publications and associated implementations available as resources for this study. We also acknowledge the developers and maintainers of the open-source scientific computing and visualization tools used in this work.

Appendix I. Supplementary structural analysis

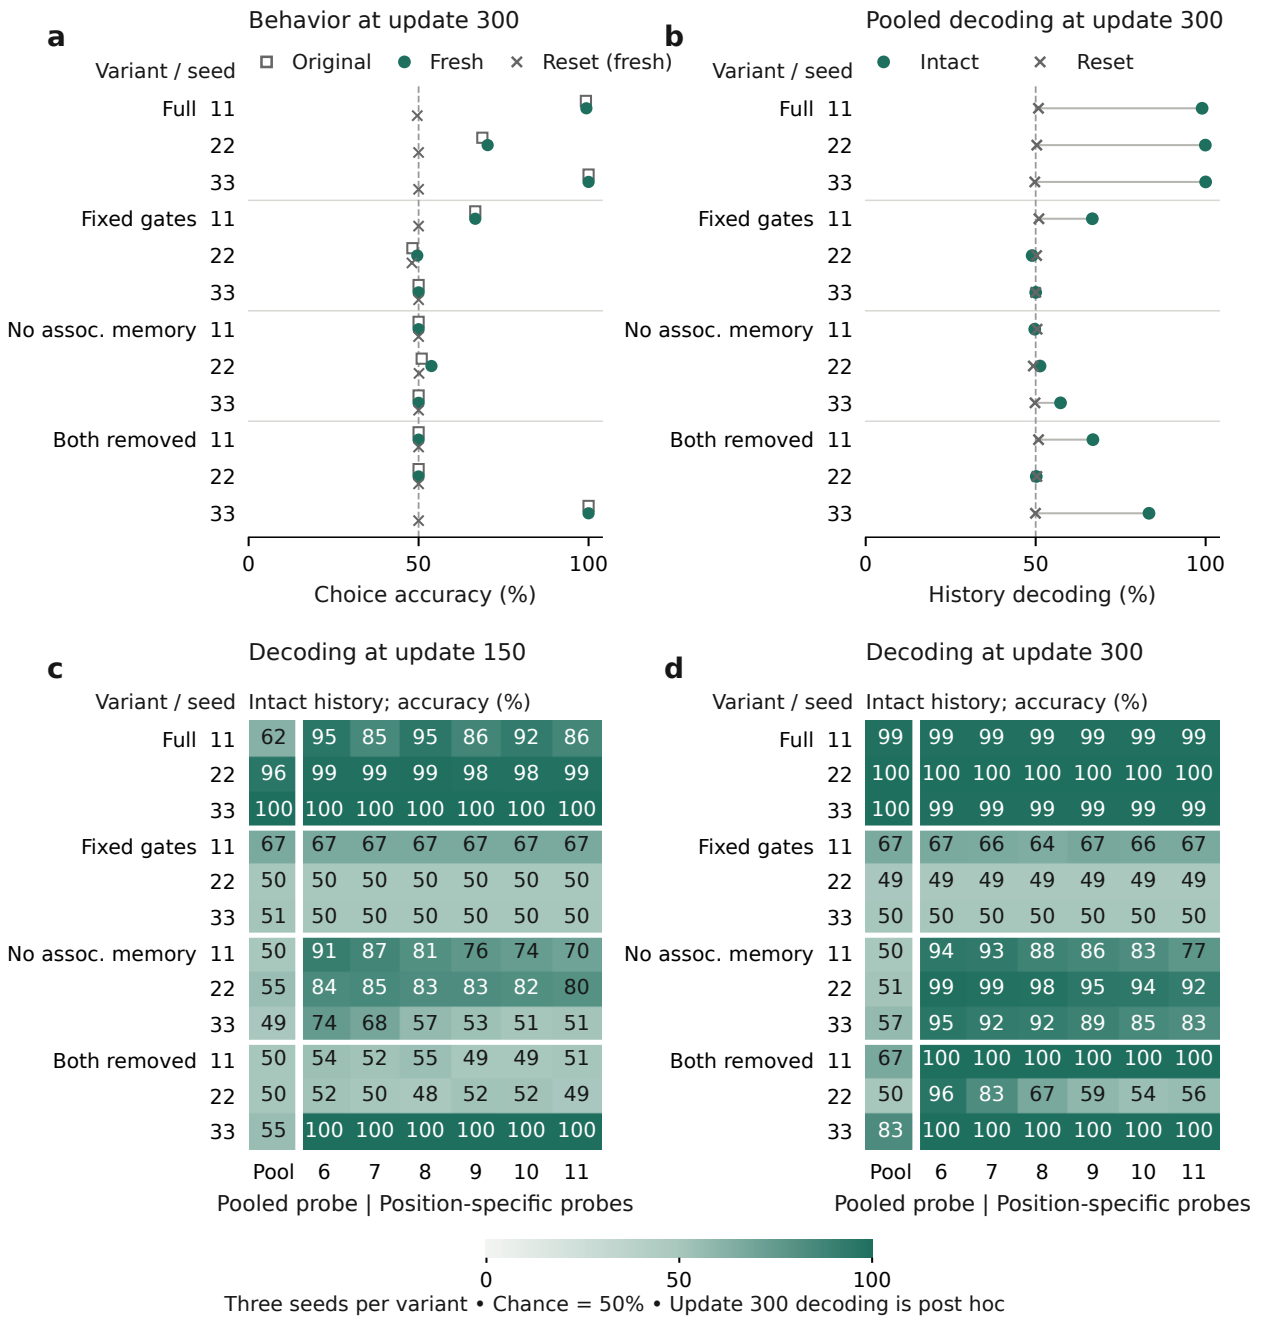


Fig. 6 Exploratory structural controls and history readout. (a) Choice accuracy at update 300 on original, fresh, and history-reset evaluation trials. (b) Pooled history decoding with intact and reset histories. (c–d) Pooled and position-specific decoding at updates 150 and 300. Rows represent individual seeds across four model variants. All variants retain predictive recurrence; "No assoc. memory" disables online associative memory, and "Both removed" additionally fixes all adaptive gates open. Three development seeds per variant; chance accuracy is 50%. Update-300 and condition-specific decoding are exploratory analyses

Appendix II. Experimental and analytical protocols

Structural-control sampling and readouts

Structural controls used three development seeds. Intact-history and history-reset probes were fitted separately. Pooled probes used class-balanced observations, whereas position- and trip-conditioned probes used samples from their corresponding conditions. The update-150 representation endpoint was retained, with update-300 measurements included as exploratory comparisons. Behavioral and decoder sample sizes are reported in Table 3 (Appendix III).

Four-context task generation and behavioral readout

The four-context task generated 80-observation sequences with independently sampled binary cues $A, B \in \{0, 1\}$ and target $y = 2A + B$. Using zero-based indices, Cue A began uniformly at positions 3–14, and each cue lasted four samples. The two delay lengths were sampled uniformly from 8–20 and 10–24 samples, respectively. Cue B followed the first delay, and the decision followed the second; sequences were retained when the decision index fell within the observation sequence.

Each position received a fixed 24-dimensional standard-normal sensory vector, with probability 0.10 of copying another position's vector during environment construction. Fixed cue-identity vectors were added during their corresponding windows, with independent Gaussian noise of standard deviation 0.08. All samples used the same forward action. Task phases comprised Cue A, delay 1, Cue B, delay 2, and decision; baseline included samples before Cue A and after the decision.

The event-aligned classifier combined separately attention-pooled Cue A and Cue B output activity with the output state immediately preceding the decision. Three feedforward branches transformed these representations before four-class classification. Event indices were provided to the classifier, while the recurrent core received observations and actions. The readout was jointly trained with the model and retained previously recorded cue-window activity during subsequent interventions.

Auxiliary memory-writing objectives

Four-context HDGM training included two EC-controlled writing penalties. Writing strength was $w_{i,t} = \bar{e}_{i,t}$ for the first 79 observation updates. For sequence i and phase s, the phase mean and general margin objective were

$$m_{i,s} = \frac{1}{|\mathcal{T}_{i,s}|} \sum_{t \in \mathcal{T}_{i,s}} w_{i,t}$$

$$\mathcal{L}_j = \frac{1}{|\mathcal{V}_j|} \sum_{(i,s,r) \in \mathcal{V}_j} [\delta_j - (m_{i,s} - m_{i,r})]_+ \qquad j \in \{B, \text{rank}\}$$

Here, $\mathcal{T}_{i,s}$ contains recorded updates from phase s, $[u]_+ = \max(0, u)$, and $\mathcal{V}_j$ contains sequence–phase pairs from the designated comparison set with both phases nonempty. The comparison sets were

$$\mathcal{P}_B = \{(B, \text{base}), (B, D_2)\}$$

$$\mathcal{P}_{\text{rank}} = \{(A, \text{base}), (B, \text{base}), (A, D_1), (B, D_2),$$
$$(A, \text{dec}), (B, \text{dec})\}$$

The margins were $\delta_B = 0.05$ and $\delta_{\text{rank}} = 0.02$. Losses were averaged over valid sequence–phase comparisons and set to zero for empty valid-comparison sets. Their weights are reported in Table 1 (Appendix III), with the combined objective defined in Eq. (7).

Memory-operation interventions and CA1 validation

Frozen-model interventions targeted cue encoding, delay maintenance, or the final four readout samples. The EC cut suppressed selected sensory input and memory writing; the CA3 cut removed $M_{t-1} h_t^{(\ell)}$ while retaining the local term, learned gate, and associative iterations; and the CA1 memory cut removed $M_{t-1} g_t$ while retaining $\kappa_{\text{CA1}} g_t$. A complementary CA1 cut removed the direct predictive contribution. Model parameters remained fixed during these interventions.

Gate substitutions used zero, one, or calibrated per-unit means. Mean gates were estimated from independent intact-model activity averaged over sequences and time, with CA3 values first averaged across associative iterations. Fixed-gate retraining held one gate open during optimization, with writing penalties omitted for fixed-EC retraining. Cross-trial reassignment cyclically shifted online-computed gates by one trial at each time step and, for CA3, at each

associative iteration within the selected intervention window.

Independent CA1 validation compared decoding from q_t and o_t and evaluated fixed-open, calibrated-mean, causal running-mean, and one-step-lagged output gates. Dynamic substitutes used intact reference sequences from matched inputs: the running mean included values through time t, while the lagged gate used time $t-1$, initialized with the independently calibrated mean. Reference schedules were replayed during evaluation. Calibration, decoding, and behavioral datasets were independent, with configurations reported in Tables 2 and 3 (Appendix III).

Population measurements and spatial controls

Four-context delay-state separation was computed after subtracting position-specific mean activity, averaging residual activity within trials, and standardizing features across trials. Standard deviations below 10^{-6} were replaced with one. Separation was the Euclidean distance between Cue A history centroids normalized by the square root of population dimension.

Ring-position map stability was the Pearson correlation between flattened position-by-unit maps from alternating trial subsets after unit-wise standardization. Its structured null permuted position correspondence between maps. Matched sensory–action spatial separation normalized between-position centroid distances by within-position dispersion, with a corresponding within-condition sample-permutation null. Both measurements were repeated after residualization using a linear sensory–action model fitted on training data.

For shared-corridor histories $b \in \{0,1\}$ at position p, between-history distance, within-history dispersion, and the separation ratio were defined as

$$
\begin{aligned}
D_p &= \frac{\|\mu_{1,p} - \mu_{0,p}\|_2}{\sqrt{d}} \\
S_p &= \frac{1}{2\sqrt{d}} \left[\frac{1}{n_{0,p}} \sum_{i \in \mathcal{I}_{0,p}} \|u_i - \mu_{0,p}\|_2 + \frac{1}{n_{1,p}} \sum_{i \in \mathcal{I}_{1,p}} \|u_i - \mu_{1,p}\|_2 \right] \\
\varrho &= \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} \frac{D_p}{S_p + 10^{-8}}
\end{aligned}
$$

Here, u_i is a population vector, $\mu_{b,p}$ its history-conditioned centroid, and $\mathcal{I}_{b,p}$ and $n_{b,p}$ denote the corresponding samples and counts. The set $\mathcal{P}$ contains positions with at least two samples from each history.

Appendix III. Reproducibility tables

Table 1: Model and task parameters

Parameter	Value	Role and scope
Observation dimension d_x	24	Sensory vector dimension in the reported implementations.
Population dimension d_h	96	Predictive, associative, output, and hidden-gate vectors.
Action embedding dimension	32	Learned action representation.
Feedback dimension d_v	96	HDGM output-derived feedback state.
Memory matrix size	96×96	Online associative state maintained per sequence.
Associative iterations	5	Iteration indices $\ell = 0, \ldots, 4$.
Associative mixing η	0.32	Contribution of retrieved state to predictive state.
CA3/CA1 local-state coefficients $\kappa_{\mathrm{CA3}}, \kappa_{\mathrm{CA1}}$	0.5, 0.5	Local-state contribution to CA3 associative updating and direct predictive contribution to CA1 associative readout.
Memory retention λ	0.92	Retention of the preceding memory matrix.
Feedback retention ρ	0.84	HDGM feedback-state retention.
Output gain c_1, c_2	0.2, 1	HDGM baseline and gate-dependent gain.
Output threshold β	0.1	HDGM threshold before rectification.
Initial states	$g_0 = M_0 = v_0 = 0$	Sequence initialization; history reset jointly clears these states at trip boundaries.
Cue-B margin	0.05	Cue B versus baseline and delay 2; four-context HDGM only.
Phase-ranking margin	0.02	Cue A versus baseline, delay 1, decision; cue B versus baseline, delay 2, decision. Nonempty windows only.
Classification-loss weight ω_{cls}	0.7	Four-context cross-entropy weight shared by HDGM and PAM.
Writing-loss weights $\omega_B, \omega_{\mathrm{rank}}$	0.08, 0.12	Weights of $\mathcal{L}_B$ and $\mathcal{L}_{\mathrm{rank}}$ in four-context HDGM training only.
Ring positions	20	Position labels used for analysis, not core training.
Movement probabilities	0.45 / 0.45 / 0.10	Forward / backward / stay.
Landmark positions	0, 7, 13	Zero-based ring positions.
Ring sequence length	100	Observations per training sequence.
Corridor positions	18	Locations used in history analyses.
Trips / arm length / endpoint hold	4 / 7 / 2	Trips per trial / positions per arm / endpoint samples.
Shared alias interval	6–11	Zero-based corridor indices.
Aliasing levels α	$0, 0.25, 0.50, 0.75$	Frozen perturbation scan.
Noise levels σ	$0.02, 0.05, 0.08, 0.12$	Frozen perturbation scan.
Shared baseline	$\alpha = 0$, $\sigma = 0.05$	Seven unique conditions per run across both scans.

Prediction losses were summed over time in the four-context assay and averaged over time in the spatial and shared-corridor assays. Auxiliary writing penalties were used only in four-context HDGM training.

Table 2: Training and frozen-evaluation protocols

Assay	Core training and checkpoint	Readout and evaluation
Four-context acquisition (Fig. 2)	Adam, 6×10^{-4}; batch 32; 150 updates. Snapshots: 10, 25, 50, 75, 100, 150. Objectives in Eq. (7).	Linear-probe banks: 20/5/20 batches of 64; Adam .01, weight decay 10^{-4}, 200 epochs. Best mean validation accuracy across three targets.
Stage interventions (Fig. 3a–c)	Frozen four-context core at update 150. Stage-specific operation cuts, gate substitutions, and trial shuffling.	Paired control/intervention evaluations within each seed. Mean-clamp calibration uses an independent analysis bank.
CA1 validation (Fig. 3d–e)	Frozen CA1 output-gain validation analysis core at update 100; source training run maximum 150.	Independent label-free calibration: 10×64; behavior: 30×64; probe banks: $20/5/20 \times 64$, as above.
Position (Fig. 4)	Adam, 7×10^{-4}; batch 24; prediction only; sequence length 100. Representation update 150.	Banks: $8/4/8 \times 32$; per-position sample caps 240/120/160. Linear probe: at most 180 epochs, patience 25.
Shared corridor (Fig. 4)	Adam, core 6×10^{-4}, choice head .002; batch 16. Loss: $.20\mathcal{L}_{\text{pred}} + 1.50\mathcal{L}_{\text{choice}}$. Representation/behavior: updates 150/300.	Banks: $8/4/8 \times 32$; balanced binary probes, at most 160 epochs, patience 25. Stateless comparator learning rate .002.
Perturbation scans (Fig. 5)	Frozen checkpoints for perturbation analysis; representation/behavior: updates 150/300. Four alias levels and four noise levels share one baseline.	Banks per condition: $8/4/8 \times 32$. Reference probe: at most 140 epochs, patience 22; weights and normalization fixed across scans.

Bank counts indicate training/validation/test. Four-context and CA1 analyses used seeds 11–99 in steps of 11 plus 111; position, corridor, and perturbation analyses used seeds 101–1010 in steps of 101. Core gradient-norm clipping was 1, with zero weight decay for position and corridor cores. Position, corridor, and transfer probes used Adam at .01 with weight decay 10^{-4}. Core checkpoints were fixed by update, and probe checkpoints were selected by validation performance. Fixed-gate retraining followed the four-context schedule; acute gate-opening evaluation used update 100.

Table 3: Linear-probe fitting and independent evaluation datasets

Analysis	Data and normalization	Fitting and selection
Four-context	Train/validation/test: 20/5/20 batches of 64 sequences. Training mean and population SD; SD floor 10^{-4}.	Adam 0.01; weight decay 10^{-4}; 200 epochs. Sum of cue-A, cue-B and class cross-entropies. Best mean validation accuracy across three targets.
Position	8/4/8 batches of 32 sequences; no-replacement caps 240/120/160 per position. Training mean and SD plus 10^{-6}.	Adam 0.01; weight decay 10^{-4}; maximum 180 epochs; validation patience 25.
Corridor	8/4/8 batches of 32 trials; class-balanced samples; training-derived standardization.	Adam 0.01; weight decay 10^{-4}; maximum 160 epochs; validation patience 25.
CA1 validation	Calibration: 10×64, label-free. Behavior: 30×64. Probe banks: $20/5/20 \times 64$ sequences.	Independent calibration, probe and behavior streams. Standardization and fitting follow four-context protocol.
Perturbation transfer	8/4/8 batches of 32 trials per condition. Fit and select using reference-condition train/validation data only.	Adam 0.01; weight decay 10^{-4}; maximum 140 epochs; patience 22. Weights and training normalization fixed across conditions.
Structural controls	Behavior: 256 trials, 768 scored choices per bank. Pooled test: 4,608 observations; per-position test: 768 observations.	Intact/reset probes fitted separately. Pooled probes class-balanced; conditional probes use all available samples. Update-300 and conditional analyses exploratory.

Dataset counts indicate training/validation/test. Decoder checkpoints were selected using validation data.

Table 4: Statistical units, comparison families, and uncertainty definitions

Analysis	Estimand and test family	Displayed uncertainty
Acquisition and delay (Fig. 2)	Seed-paired model differences. Delay: exploratory family of three change contrasts, reconstructed during audit.	Accuracy: normal 95% interval. AUC difference: 20,000 seed-bootstrap resamples, percentile 95%; seed 24680.
Operation cuts (Fig. 3a)	Control minus intervention. Nine operation-by-stage comparisons per outcome and delay regime. Additional nine direct stage/component contrasts are exploratory.	Run-level paired losses; the independent units are training runs.
Gate substitutions and shuffling (Fig. 3b–c)	b: descriptive encoding losses. c: paired effects, nine-condition Holm family for each outcome.	b: normal 95%. c: 5,000 percentile resamples of standardized differences; original paired SD held fixed.
CA1 output (Fig. 3d–e)	d: two source-to-output contrasts. e: four substitute-gate contrasts per outcome.	d: normal 95%. e: 50,000 paired-seed percentile resamples, seed 62010.
Population assays (Fig. 4)	Separate two-test Holm families for position and history; behavior has a separate two-test family.	b–c: 12,000 seed resamples. f: 20,000 resamples of standardized differences, original SD fixed. Percentile 95% intervals.
Perturbation behavior and access (Fig. 5c–d)	Access is full minus reset decoding. Post hoc family: two access-AUC and two stability-AUC contrasts.	c: Student-t 95%, 9 degrees of freedom. d: 16,000 seed resamples, percentile 95%.
Access–behavior coupling (Fig. 5e)	Descriptive correlation after centering within seed; 70 observations clustered within ten independent training runs.	12,000 seed-cluster resamples for correlation and regression band; percentile 95%.
Perturbation effects (Fig. 5f)	24-cell Holm family. Access/separation: full minus reset; choice: HDGM minus .50. Baseline appears on both scan axes.	$d_z = \bar{d}/s_d$ and ten seed marks per cell; no cell-wise confidence interval.

Seeds are the independent units. For the original ten-seed analyses, paired tests enumerate all 2^{10} sign assignments and are two-sided. The three-seed structural pilot is descriptive and is not included in these test families. Holm correction applies within the specified family, not across the entire manuscript. Intervals are unadjusted and can exclude zero when a Holm-adjusted test is nonsignificant. Normal intervals are $\bar{x} \pm 1.96s/\sqrt{10}$. Bootstrap intervals resample seeds, preserving paired observations; standardized-mean intervals with fixed SD are not bootstrap intervals that re-estimate d_z in every draw. Percentage-point differences are distinguished from relative percentage changes.

Table 5: Notation for model updates and analyses

Symbol	Dimension / type	Definition
t	Step index	Observation/update step within a sequence.
ℓ	Iteration index	Associative iteration, $0, \ldots, 4$; distinct from output gate k_t.
x_t, $\hat{x}_{t+1}$	$\mathbb{R}^{d_x}$	Current observation and predicted next observation.
a_t, $\mathrm{Emb}(a_t)$	Action / $\mathbb{R}^{32}$	Supplied action and learned embedding.
g_t^-, g_t	$\mathbb{R}^{d_h}$	Predictive state before/after associative integration.
M_t	$\mathbb{R}^{d_h \times d_h}$	Online associative-memory matrix.
v_t	$\mathbb{R}^{d_v}$	Output-derived feedback state.
e_t, $\overline{e}_t$	$\mathbb{R}^{d_x}$ / scalar	Input gate and its per-sequence component mean.
z_t	$\mathbb{R}^{d_h}$	Encoded selected observation; initial associative state.
$h_t^{(\ell)}$, $r_t^{(\ell)}$	$\mathbb{R}^{d_h}$	Associative iterate and candidate update.
$c_t^{(\ell)}$	$\mathbb{R}^{d_h}$	Associative-iteration gate.
$p_t = h_t^{(5)}$	$\mathbb{R}^{d_h}$	Retrieved associative state after five iterations.
q_t	$\mathbb{R}^{d_h}$	Associative readout before output modulation.
k_t	$\mathbb{R}^{d_h}$	CA1-inspired output gate.
o_t	$\mathbb{R}^{d_h}$	Output population activity after gain and thresholding.
W, b	Learned parameters	Linear maps and biases; subscripts identify operations.
$[\cdot]$, $\odot$	Operators	Concatenation and elementwise multiplication.
$\eta, \kappa_{\mathrm{CA3}}, \kappa_{\mathrm{CA1}}, \lambda, \rho$	Constants	Associative mixing, CA3/CA1 local-state coefficients, and memory/feedback retention.
c_1, c_2, β	Constants	Output-gain coefficients and threshold.
α, σ	Perturbation parameters	Cue aliasing strength and observation-noise level.
$\mathcal{L}_{\mathrm{pred}}$, $\mathcal{L}_{\mathrm{class}}$, $\mathcal{L}_{\mathrm{choice}}$	Losses	Observation prediction, four-class cross-entropy, and choice cross-entropy.
$\mathcal{L}_B$, $\mathcal{L}_{\mathrm{rank}}$	Writing penalties	Cue-B margin and phase-ranking penalties.
$\omega_{\mathrm{cls}}, \omega_B, \omega_{\mathrm{rank}}$	Loss weights	Four-context classification and auxiliary memory-writing weights.
ϱ	Separation ratio	Mean across positions of between-history centroid distance divided by within-history dispersion.
$d_z = \overline{d}/s_d$	Paired effect size	Mean run-level paired difference divided by its sample SD.
s, s_d	Sample SDs	Across-run values and paired differences, respectively.
p_{Holm}	Adjusted probability	Holm-adjusted test p value within the stated family.
pp, AUC, CI	Abbreviations	Percentage points, area under the curve, confidence interval.

References

Basu J, Zaremba JD, Cheung SK, Hitti FL, Zemelman BV, Losonczy A, Siegelbaum SA (2016) Gating of hippocampal activity, plasticity, and memory by entorhinal cortex long-range inhibition. Science 351:aaa5694. https://doi.org/10.1126/science.aaa5694

Colgin LL, Denninger T, Fyhn M, Hafting T, Bonnevie T, Jensen O, Moser MB, Moser EI (2009) Frequency of gamma oscillations routes flow of information in the hippocampus. Nature 462:353–357. https://doi.org/10.1038/nature08573

Ferbinteanu J, Shapiro ML (2003) Prospective and retrospective memory coding in the hippocampus. Neuron 40:1227–1239. https://doi.org/10.1016/S0896-6273(03)00752-9

Grienberger C, Magee JC (2022) Entorhinal cortex directs learning-related changes in CA1 representations. Nature 611:554–562. https://doi.org/10.1038/s41586-022-05378-6

Ito Y, Toyoizumi T (2026) Modeling flexible behavior with remapping-based hippocampal sequence learning. eLife 14:RP106506. https://doi.org/10.7554/eLife.106506.4

Kaufman MT, Churchland MM, Ryu SI, Shenoy KV (2014) Cortical activity in the null space: permitting preparation without movement. Nat Neurosci 17:440–448. https://doi.org/10.1038/nn.3643

Keinath AT, Nieto-Posadas A, Robinson JC, Brandon MP (2020) DG–CA3 circuitry mediates hippocampal representations of latent information. Nat Commun 11:3026. https://doi.org/10.1038/s41467-020-16825-1

Liu Y, Wang Z, Qiao C, Xu Z (2026) GATE: adaptive learning with working memory by information gating in multi-lamellar hippocampal formation. PLoS Comput Biol 22:e1014438. https://doi.org/10.1371/journal.pcbi.1014438

Martín-Sánchez G, Machens CK, Podlaski WF (2025) Three types of remapping with linear decoders: a population-geometric perspective. PLoS Comput Biol 21:e1013545. https://doi.org/10.1371/journal.pcbi.1013545

Nakazawa K, Quirk MC, Chitwood RA, Watanabe M, Yeckel MF, Sun LD, Kato A, Carr CA, Johnston D, Wilson MA, Tonegawa S (2002) Requirement for hippocampal CA3 NMDA receptors in associative memory recall. Science 297:211–218. https://doi.org/10.1126/science.1071795

Nieh EH, Schottdorf M, Freeman NW, Low RJ, Lewallen S, Koay SA, Pinto L, Gauthier JL, Brody CD, Tank DW (2021) Geometry of abstract learned knowledge in the hippocampus. Nature 595:80–84. https://doi.org/10.1038/s41586-021-03652-7

Rolls ET (2013) The mechanisms for pattern completion and pattern separation in the hippocampus. Front Syst Neurosci 7:74. https://doi.org/10.3389/fnsys.2013.00074

Sun W, Winnubst J, Natrajan M, Lai C, Kajikawa K, Bast A, Michaelos M, Gattoni R, Stringer C, Flickinger D, Fitzgerald JE, Spruston N (2025) Learning produces an orthogonalized state machine in the hippocampus. Nature 640:165–175. https://doi.org/10.1038/s41586-024-08548-w

Udakis M, Claydon MDB, Zhu HW, Oakes EC, Mellor JR (2025) Hippocampal OLM interneurons regulate CA1 place cell plasticity and remapping. Nat Commun 16:9912. https://doi.org/10.1038/s41467-025-64859-0

Varga V, Petersen P, Zutshi I, Huszar R, Zhang Y, Buzsáki G (2024) Working memory features are embedded in hippocampal place fields. Cell Rep 43:113807. https://doi.org/10.1016/j.celrep.2024.113807

Whittington JCR, Muller TH, Mark S, Chen G, Barry C, Burgess N, Behrens TEJ (2020) The Tolman–Eichenbaum machine: unifying space and relational memory through generalization in the hippocampal formation. Cell 183:1249–1263.e23. https://doi.org/10.1016/j.cell.2020.10.024

Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H (2000) Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron 27:623–633. https://doi.org/10.1016/S0896-6273(00)00071-4

Young RA, Shin JD, Guo Z, Jadhav SP (2025) Hippocampal–prefrontal communication subspaces align with behavioral and network patterns in a spatial memory task. eNeuro 12:ENEURO.0336-24.2025. https://doi.org/10.1523/ENEURO.0336-24.2025