

# 1 Data-derived agents reveal dynamical reservoirs in mouse cortex for adaptive behavior

2  
3 Siyan Zhou<sup>1</sup>, Ryan P. Badman<sup>1,2</sup>, Charlotte Arlt<sup>1</sup>, Kanaka Rajan<sup>1,2,\*</sup>, Christopher D. Harvey<sup>1,\*</sup>

4  
5 <sup>1</sup> Department of Neurobiology, Harvard Medical School, Boston, MA

6 <sup>2</sup> Kempner Institute for the Study of Natural and Artificial Intelligence at Harvard University, Boston, MA

7 \* Correspondence to: [harvey@hms.harvard.edu](mailto:harvey@hms.harvard.edu), [kanaka\\_rajan@hms.harvard.edu](mailto:kanaka_rajan@hms.harvard.edu)

## 8 9 **Abstract**

10 Animals generate behaviors that are robust to perturbations yet adaptable to changing conditions. How  
11 neural population dynamics support this balance between robustness and flexibility remains unclear. We  
12 address this question in goal-directed navigation by combining large-scale calcium imaging from mouse  
13 cortex with a data-derived modeling framework. We trained agents to navigate in a simulative  
14 environment while recapitulating mouse neural and behavioral data trial-by-trial. Data-derived agents  
15 discovered novel dynamics of chaotic attractors, characterized by intrinsically variable trajectories  
16 confined within overall goal-specific attracting landscapes. These dynamics support reliable goal  
17 achievement while maintaining a structured distribution of navigational trajectories. Circuit-level analyses  
18 and perturbations reveal mechanisms that stabilize chaos and enhance behavioral adaptability in the data-  
19 derived agents. Thus, through our new modeling approach that emphasizes closed-loop interactions  
20 between behavior and neural dynamics, we reveal chaos as a functional principle for flexible behavior.

## 21 22 **Introduction**

23 Brains generate behaviors that are robust to perturbations yet adaptable to changing conditions. Even for  
24 behavioral outcomes that are stable and robust, the activity in populations of neurons often exhibits  
25 substantial variability across repeated trials and conditions <sup>1-7</sup>. An interesting proposal is that the trial-to-  
26 trial variability in neural activity may play a functional role in flexible behavior, rather than reflecting  
27 noise that is simply averaged away.

28  
29 Existing theories of neural dynamics typically emphasize either robustness or flexibility, but not both  
30 simultaneously. Classical attractor-based frameworks achieve reliable outcomes by contracting neural  
31 trajectories toward fixed points or low-dimensional manifolds, a mechanism that necessarily suppresses  
32 within-condition variability <sup>8-15</sup>. Conversely, models that generate flexible or diverse trajectories often  
33 rely on transient or feedforward dynamics which are not designed to guarantee robustness to perturbations  
34 <sup>16-20</sup>. As a result, current theories lack a dynamical regime that can simultaneously support stable  
35 behavioral outcomes and rich neural and behavioral variability under the closed-loop demands that  
36 animals naturally operate under.

37  
38 Neural activity unfolds in continuous interaction with the environment, such that variability in neural  
39 dynamics shapes future sensory input, which in turn influences subsequent neural activity <sup>21-24</sup>. Yet many  
40 widely used models of neural dynamics are primarily evaluated in open-loop settings, decoupled from the

41 behavioral consequences of neural activity<sup>25,26</sup>. Severing this interaction obscures how neural dynamics  
42 support behavior over extended trajectories and might systematically mischaracterize the functional role  
43 of variability.

44

45 Goal-directed navigation naturally exposes this challenge. Animals reliably reach spatial goals while  
46 expressing diverse trajectories across trials, even under similar environmental conditions<sup>1,2,27,28</sup>. This  
47 dissociation between outcome reliability and path variability provides a clear setting in which to  
48 investigate how neural population dynamics can support stable behavioral goals without committing to  
49 fixed trajectories.

50

51 Here we show that chaotic neural dynamics provide a regime in which population activity supports reliable  
52 goal attainment while simultaneously generating diverse and adaptable trajectories. In closed-loop  
53 interaction, chaos enables structured exploration without compromising goal stability, rather than  
54 producing unstructured instability. This regime reconciles robustness and flexibility within a single set of  
55 neural dynamics.

56

57 Identifying this functional role of chaos requires models that are evaluated under closed-loop interaction  
58 with the environment. We achieved this by reconstructing the neural dynamics jointly with the closed-  
59 loop interactive behavior of navigation. In open-loop analyses that primarily evaluate trajectory  
60 predictability or stability, chaotic regimes can appear unstable or uninformative and are therefore often  
61 disfavored during model fitting or selection<sup>15,29,30</sup>. Closed-loop evaluation reveals that these same regimes  
62 can support reliable behavior while preserving structured variability, a distinction that cannot be recovered  
63 by fitting neural activity alone.

64

65 Together, our results establish chaos as a functional principle of neural population dynamics supporting  
66 flexible behavior. More broadly, they suggest that balancing stability and adaptability in closed-loop tasks  
67 may rely on dynamical regimes that are systematically overlooked by standard analytical approaches.

68

## 69 **Results**

### 70 **Closed-loop modeling of neural dynamics and behaviors**

71 In sequential behavior tasks like navigation, animals can achieve their goals through a variety of  
72 behavioral sequences. This is permitted by these tasks by definition, as they admit or even encourage  
73 variable behavioral sequences that fulfill a higher-level goal and comply with certain task constraints (Fig.  
74 1a). In navigation, the higher-level goal is to reach a destination location, and the task constraints manifest  
75 as the organization of the maze walls, as well as coupling laws of how locomotion actions move the  
76 animal's location in the environment and induce different future state distributions and requirements on  
77 subsequent actions. These constraints are embedded in the closed-loop interaction between the nervous  
78 system and the environment. Examining the neural activity alone without acknowledging the closed-loop  
79 interaction may obscure how neural dynamics support the temporally extended behavior.

80

81 To jointly model the underlying neural dynamics and the closed-loop interactive behavior, we developed  
82 a method that combines two commonly used modeling directions, neural dynamical systems  
83 reconstruction through data fitting<sup>19,31</sup>, and agent-based modeling<sup>32</sup>. Data-fitting models train dynamical  
84 systems like recurrent neural networks (RNN) to reconstruct experimentally recorded neural activity  
85 sequences. However, most existing approaches treat sensory inputs as independent processes, either pre-  
86 designed and fixed or co-optimized with the RNN weights as free parameters<sup>19,33,34</sup>. This assumption of  
87 independence neglects the closed-loop environment interaction in navigation. On the other hand, task-  
88 performing agents are trained (usually by reinforcement learning) to perform tasks by acting in an  
89 environment and receiving sensory inputs according to their own actions in a closed loop<sup>32,35-40</sup>. Yet these  
90 agents often aim to identify effective policies rather than reveal specific neural dynamics mechanisms. A  
91 common current practice is to compare internal representations of the agents with real animals in a *post*  
92 *hoc* manner, rather than using real neural dynamics to drive the agent's behavior<sup>41</sup>.

93  
94 We propose ARCTIC (Activity Reconstruction in Closed loop) that combines data fitting with closed-  
95 loop agents (Fig. 1b). The approach features an environment-interacting agent that is constrained by both  
96 the neural activity and behavioral outputs recorded from real task-performing animals. There are two key  
97 technical advancements. First, in the spirit of the closed-loop environment interaction, neural activity  
98 reconstruction and behavior output optimization are trained simultaneously in an online manner<sup>42</sup>. This  
99 helps to address error accumulation in the closed-loop setup. Second, the agent is trained on the neural  
100 and behavioral data of individual animal trials rather than averaged data. This enables the agent to  
101 recapitulate the distribution of possible sequences and the underlying neural mechanism for such  
102 variability (see Methods).

### 103 104 **Mice exhibit variable trajectories to the same choice in navigation**

105 We studied mouse navigation to understand how a distribution of trajectories arises from neural dynamics.  
106 Mice were trained to perform a navigation decision-making task in a virtual reality Y-maze<sup>2,43</sup> (Fig. 2a).  
107 On each trial, one of two visual cues was presented in the Y-stem, and mice learned that each cue was  
108 associated with either a left or right turn at the Y-intersection to receive a reward. As mice performed the  
109 task, cellular-resolution calcium imaging via a large field-of-view, random-access microscope<sup>44</sup> was used  
110 to measure the activity of thousands of individual neurons across four cortical areas: primary visual cortex,  
111 posterior parietal cortex, retrosplenial cortex, and secondary motor cortex (Fig. 2b, Extended Data Fig.  
112 1a). Mice performed the task with high accuracy across sessions (Fig. 2c). A portion of these experimental  
113 data was reported previously in ref.<sup>43</sup>.

114  
115 Neural activity unfolded as choice-selective sequences<sup>45</sup> (Fig. 2d). Individual neurons were transiently  
116 active at specific points in each trial of the task. Different sequences of neurons were active on trials  
117 associated with different behavioral choices. Although population activity followed distinct trajectories  
118 that were separable for different choices, it exhibited substantial variability across trials of the same  
119 choice. This variability was concentrated along task-relevant dimensions of neural activity. It was evident  
120 in both the leading principal components of population activity and the choice dimension, defined as the

121 axis that best separates neural activity for left versus right choices (Fig. 2e, f). Consistent with the neural  
122 variability, mice showed behavioral differences in left-right (lateral) running velocity across trials, even  
123 when navigating to the same goal location (Fig. 2g). This behavioral variability emerged because, by  
124 nature of the navigation task, the trial evolution is fully determined by the mouse. With the long running  
125 distance (2 m) and trial duration (5-10 s, Extended Data Fig. 1b), it is sensible that mice explored different  
126 trajectories and made variable movements in their running at different points across trials.

127

128 To isolate trial-to-trial neural variability, we subtracted the mean activity at each forward position in the  
129 maze from each individual trial, computed separately for left and right choice trials. The resulting residual  
130 activity reflects variability across repeated trials that is not explained by choice or forward position within  
131 the maze. If this trial-to-trial variance is unstructured, its projection onto the choice dimension would be  
132 small and near chance level. Instead, substantial residual variance, and its leading principal components,  
133 were strongly aligned with the choice dimension (Fig. 2h-i). This finding is consistent with previous work  
134 showing an alignment of signal and noise correlations in neural populations<sup>4,5,46</sup>. Together, these results  
135 indicate that the trial-to-trial variability is task-relevant and unlikely to arise from unstructured noise or  
136 encoding of irrelevant stimuli or behaviors.

137

### 138 **Data-derived agents recapitulate variable trajectories without added noise**

139 To investigate the mechanisms underlying these activity patterns, we applied ARCTIC to train data-  
140 derived navigation agents. Each agent has a densely connected RNN in which each unit corresponds to a  
141 single neuron recorded during a calcium imaging and behavioral session. Each unit was trained to  
142 reproduce the activity of a corresponding neuron<sup>19,29,31</sup>, while the network as a whole was optimized to  
143 match the mouse's forward and lateral running velocities (Fig. 2j). We trained one agent for each animal  
144 session. The agent interacted with a simulative environment analogously to how a mouse interacts with  
145 its environment. The agent received inputs from the environment corresponding to the visual cue and its  
146 current position in the maze. The agent's velocity outputs then updated its position, thereby determining  
147 the next set of inputs. For each trial, the activity of units in the RNN and the position in the environment  
148 were initialized to the first empirical datapoint of that trial. The trained agent then autonomously generated  
149 the full neural and behavioral sequence of the trial, meaning its outputs arose entirely from internal  
150 dynamics and closed-loop interaction with the environment. After training, the agents performed the task  
151 with high accuracy (Fig. 2k).

152

153 The agents produced single-neuron activity and behavioral outputs that recapitulated key features of the  
154 mouse data. Strikingly, this was the case even in test trials held out from training. First, the agents  
155 generated distinct neural sequences and running trajectories for each trial type, matching the  
156 experimentally observed separation by mouse choice (Fig. 2l-o). Second, the agents recapitulated trial-  
157 to-trial variability in both neural activity and behavior. This variation emerged even within a single trial  
158 type, despite the absence of added noise (Fig. 2m-o; Extended Data Fig. 1f). Like the experimental data,  
159 the residual variance, which was computed after subtracting trial-averaged activity, was well aligned with  
160 the choice dimension (Fig. 2p-q). Also, the agents captured a larger fraction of the residual variance along

161 the choice dimension compared to the overall residual variance, potentially because they ignored task-  
162 irrelevant components (Fig. 2r). Together, these results indicate that the data-derived agents developed  
163 intrinsic mechanisms for generating task-relevant trial-to-trial variability that was distinct from random  
164 noise, while maintaining high decision-making accuracy.

165

### 166 **Low-dimensional chaotic attractors**

167 To probe a data-derived agent's intrinsic dynamics, we perturbed its RNN activity along the choice  
168 dimension and tracked the evolution of the agent's average response over time. Small perturbations  
169 decayed rapidly and preserved behavioral accuracy, whereas larger perturbations caused the activity to  
170 jump to the opposite trial type, resulting in incorrect choices (Fig. 3a; Extended Data Fig. 2a). To quantify  
171 the average convergence of neural activity following a perturbation, we computed the norm deviation of  
172 trial-averaged neural activity (NDM), defined as the squared difference between perturbed and  
173 unperturbed means across trials of the same choice. NDM decreased rapidly after a perturbation in correct  
174 trials, indicating that mean neural activity returned toward its original trajectory (Fig. 3b). As the  
175 perturbation amplitude increased, choice accuracy decreased nonlinearly, suggesting a boundary  
176 separating neural activity for the two choices (Extended Data Fig. 2b). Perturbations had a greater effect  
177 during the delay period than during the cue period, consistent with the idea that visual cues constrain  
178 neural activity early in the trial (Extended Data Fig. 2a). Together, these findings support the existence of  
179 separate attractor states for left and right choices. This is consistent with previous theories that attractor  
180 dynamics support reliable decision making<sup>8-12</sup>.

181

182 We next examined a data-derived agent's response to perturbations on individual trials. Surprisingly, on  
183 individual trials, small perturbations caused deviations in neural activity patterns that diverged by  
184 increasing amounts over time and did not return to their original trajectories (Fig. 3c). We quantified these  
185 deviations using the mean norm deviation (MND) of individual trials. Because deviations occurred in  
186 different directions across trials, the mean network response appeared to return to the original average  
187 trajectory (Extended Data Fig. 2c). Notably, this divergence after perturbation occurred across all  
188 dimensions of neural activity, including components aligned or not aligned with the choice dimension  
189 (Extended Data Fig. 2c, d). These findings contrast with traditional attractor models, in which activity on  
190 individual trials reliably returns to the same trajectory after small perturbations (Fig. 3d). Instead, the  
191 progressive divergence of neural trajectories following a small perturbation is a hallmark of chaotic  
192 dynamics<sup>30,47</sup>. Thus, our results reveal the presence of chaotic dynamics within each choice-specific  
193 attractor, motivating the concept of chaotic attractors. Intuitively, these dynamics amplify small  
194 fluctuations in neural activity over time, intrinsically generating trial-to-trial variability. However, within  
195 the broader attracting landscape, these chaotic dynamics remain confined to the basin associated with a  
196 single choice.

197

198 To further characterize the chaotic attractors, we identified normal modes of the dynamics, which are  
199 characteristic directions along which evolution of neural activity is self-sustained, i.e., independent of  
200 other directions<sup>48,49</sup>. Along stable directions, infinitesimal perturbations decay exponentially, while along

201 unstable directions they grow exponentially (Fig. 3e). The exponential growth rate along each mode is  
202 quantified by its Lyapunov exponent, and the corresponding direction is referred to as the covariant  
203 Lyapunov vector. Lyapunov exponents and covariant Lyapunov vectors can be viewed as nonlinear  
204 analogs of eigenvalues and eigenvectors, with the key distinction that the vectors vary with system state  
205 due to nonlinearity (see Methods for details).

206  
207 We numerically computed the Lyapunov exponents of the data-derived agents and identified the span of  
208 the leading unstable covariant Lyapunov vectors<sup>48,49</sup> (Extended Data Fig. 2e, f). Only a small subset of  
209 directions exhibited instability, as indicated by their positive Lyapunov exponents (Fig. 3f). These  
210 unstable directions define a low-dimensional unstable manifold along which perturbations amplify over  
211 time. The directions of this manifold are state-dependent, reflecting the nonlinear nature of the underlying  
212 dynamics. Due to its nonlinear geometry, the low-dimensional unstable manifold is embedded within a  
213 much higher-dimensional neural space, consistent with amplified deviations appearing across many  
214 activity dimensions. The choice dimension projected strongly onto this unstable manifold, explaining its  
215 disproportionate share of trial-to-trial variance (Fig. 3g).

216  
217 Because chaotic dynamics are closely linked to trial-to-trial variability, one might wonder whether they  
218 arose because the agents were trained to reproduce individual-trial neural activity and behavior  
219 trajectories. Indeed, training on left- and right-trial-averaged neural and behavioral trajectories did not  
220 produce chaotic dynamics. Instead, the agents learned either a continuous attractor or a pair of point  
221 attractors, depending on the learning rate (Fig. 3h). These agents trained on trial-averaged data used the  
222 same architecture and optimization as the individual-trial agents, underscoring that it is essential to  
223 consider the trial-to-trial variation to reveal the underlying dynamics.

224  
225 One potential concern is whether partial, noisy observations of a stable system could be misinterpreted as  
226 chaos<sup>50</sup>. Trial-to-trial variation can arise either from intrinsic chaotic dynamics or from random noise.  
227 Additionally, our empirical data sample only a subset of neurons involved in the task's dynamics. To  
228 address this concern, we used an agent trained on trial-averaged data as the teacher. The teacher exhibited  
229 point attractor dynamics. We then generated individual trials from the teacher by adding i.i.d. Gaussian  
230 noise (Extended Data Fig. 3a). Thus, the resulting trial-to-trial variability was due entirely to injected  
231 noise, not intrinsic dynamics. We trained a student agent on trial-by-trial observations from only a fraction  
232 of the teacher's units. Despite partial and noisy observations, the student recapitulated the teacher's point  
233 attractor dynamics (Extended Data Fig. 3a, b). Therefore, the modeling framework can distinguish  
234 unstructured noise from chaos even with partial observations, reinforcing the evidence for chaotic  
235 dynamics in the agents trained on individual trials of mouse data.

236  
237 Thus, the modeled cortical network appears to solve the navigation task using chaotic attractors, which  
238 share features with, but also differ from, stable attractors proposed previously for decision-making. Like  
239 stable attractors, the chaotic attractors segregate choices and remain robust in the presence of ongoing  
240 variability, even from chaotic dynamics. Moreover, chaotic dynamics contribute to the variability along

241 the task-relevant dimensions and define the choice-specific dynamical reservoirs of possible navigation  
242 trajectories.

243

### 244 **Circuit connectivity motifs for trial-to-trial variability**

245 We examined the connectivity of the RNN in the data-derived agents, interpreting each connection weight  
246 as a functional interaction between the corresponding real neurons, rather than as a direct anatomical  
247 synapse<sup>19,31</sup>. We focused on connections between units whose activity patterns were likely important for  
248 task performance. First, we identified choice-selective units that could contribute to separating neural  
249 trajectories for the two choices (Fig. 4a, b; Extended Data Fig. 4a). Second, because trial-to-trial variability  
250 occurs in running trajectories within a trial type, we focused on choice-selective neurons that also encoded  
251 the mouse's lateral running velocity (Fig. 4a, b; Extended Data Fig. 4b). These neurons exhibited diverse  
252 lateral velocity preferences, forming a tuning spectrum that spanned the entire lateral velocity range (Fig.  
253 4c). Moreover, many of these neurons' lateral velocity tuning was modulated by spatial position in the  
254 maze (Fig. 4a, b; Extended Data Fig. 4c), indicating that they represented specific actions at specific  
255 locations for a given choice. These cells provide all the components needed to mediate a choice-specific  
256 navigational trajectory. Thus, we observed neurons that encoded not only discrete decisions but also the  
257 variability in how those decisions were executed as navigational trajectories. The data-derived agents  
258 reproduced the tuning of these neurons, even in test trials withheld from training (Fig. 4d, e; Extended  
259 Data Fig. 5a–d). The ability of the data-derived agents to recapitulate single-neuron tuning further supports  
260 their biological plausibility. This, in turn, enables direct examination of the inferred connectivity among  
261 neurons with specific task-related representations.

262

263 Analysis of the RNN's connection weights showed that choice-selective neurons formed two competing  
264 populations, with neurons of opposite choice preference inhibiting each other more strongly than neurons  
265 with the same choice preference (Fig. 4f, g). This motif is consistent with opponent inhibition, which is a  
266 hallmark of decision-making models<sup>8,9,51–53</sup>. Within each choice preference, neurons tended to form  
267 stronger excitatory connections with neurons active later in the sequence than with those active earlier,  
268 supporting forward propagation of the neural trajectory (Fig. 4g)<sup>19,20,54</sup>. Notably, neurons sharing the same  
269 choice and spatial selectivity could be further subdivided into competing subgroups based on their lateral  
270 velocity tuning. Within these subgroups, neurons with similar lateral velocity tuning excited each other  
271 more, whereas those with distinct tuning inhibited each other more (Fig. 4h, i). Thus, the data-derived  
272 agents' connectivity reveals a two-scale motif: (1) opponent inhibition between neurons with opposite  
273 choice preferences, potentially supporting decision-making and choice stability in attractors; and (2) fine-  
274 grained competition within each choice population, which could generate within-choice variability.

275

### 276 **A two-scale competition motif leads to chaotic dynamics in a toy model**

277 To examine how chaotic dynamics could arise in the context of the two-scale motif, we generated a “toy  
278 model” as a noise-free, 400-unit vanilla RNN with random weights (Fig. 5a). The network comprised 200  
279 “left-choice” and 200 “right-choice” units, each receiving a constant contextual input during left or right  
280 trials, respectively (Fig. 5a). This RNN's connectivity has a 2-by-2 block structure, where the on-diagonal

281 blocks are connections between units of same choice and the off-diagonal blocks are between units of  
282 opposite choices. The opponent inhibition between choice populations is modeled as the off-diagonal  
283 blocks being more inhibitory than the on-diagonal blocks. On top, we added fine-grained competition  
284 within each choice population. We sorted units in each choice population according to an additional tuning  
285 parameter (which can be generic, but we will call it lateral velocity tuning for connection to the empirical  
286 results). We added stronger inhibition between units with distinct lateral velocity tuning than between  
287 units with similar tuning. In this way, the units with both the same choice preference and similar lateral  
288 velocity tuning have the most positive (or least negative) weights, and we refer to them as the within-pool  
289 weights. All the other weights are the across-pool weights. This was reflected in the connectivity matrix  
290 as diagonal ridges in the on-diagonal blocks that were less inhibitory than anywhere else (Fig. 5a). The  
291 within-pool weights and across-pool weights were sampled from two Gaussian distributions: across-pool  
292 weights were drawn from  $\mathcal{N}(\sigma\beta, \sigma^2)$ , and within-pool weights were drawn from  $\mathcal{N}(\sigma(\mu + \beta), \sigma^2)$ .  $\beta$   
293 controlled the mean of the across-pool weights, and  $\mu$  determined the difference between the mean within-  
294 pool weights and across-pool weights (i.e., specificity of the connection). We used a positive value of  $\mu$ ,  
295 which made the within-pool weights more positive (or less negative) relative to the across-pool weights,  
296 generating opponent inhibition both between units of opposite choices and between units of the same  
297 choice but with distinct lateral velocity tuning (the two-scale competition motif).  $\sigma$  set the magnitude of  
298 the gaussian distributions, determining their shapes.

299  
300 With a positive value of  $\mu$  and certain choices of  $\sigma$  and  $\beta$  ( $\mu=0.3$ ,  $\sigma=1$ ,  $\beta=-0.15$ ), the simple two-scale  
301 connectivity motif strikingly led to chaotic dynamics in this toy model. Specifically, with a step input on  
302 left trials, left-choice units had higher activity than right-choice units but did not settle to a steady  
303 amplitude. Instead, they displayed complex temporal fluctuations (Fig. 5b). We summarized responses to  
304 contextual inputs as the mean activity difference between left- and right-choice units (Fig. 5c). Simulations  
305 with different random initial states produced distinct trajectories, consistent with chaotic dynamics. The  
306 resulting chaotic dynamics even had properties that matched those observed in the data-derived agents.  
307 Lyapunov exponent analysis revealed that, as in the data-derived agent, the toy model formed a low-  
308 dimensional unstable manifold, with only the first four exponents having a positive value (Fig. 5d). The  
309 leading principal components of trial-to-trial variation in the toy model were oriented toward the choice  
310 dimension, consistent with our observations in both the mouse data and the data-derived agents (Fig. 5e).

311  
312 Closer examination of the toy model parameters revealed a phase transition between fixed-point, chaotic,  
313 and runaway regimes (Fig. 5f; Extended Data Fig. 6). With  $\mu$  kept constant, which sets the specificity of  
314 within- versus across-pool weights, the toy model exhibited fixed-point dynamics when the weight  
315 magnitude ( $\sigma$ ) was small. As the weight magnitude increased, dynamics transitioned from fixed points to  
316 runaway activity, characterized by saturated activity in all units in the toy model. Runaway activity could  
317 be stabilized by an inhibitory background in the connection weights (negative  $\beta$ ). Chaotic dynamics  
318 emerged when a high weight magnitude was paired with sufficient inhibition, possibly through a nonlinear  
319 interplay between expanding and contracting dimensions. Together, these results highlight that chaotic  
320 dynamics can emerge from a combination of the two-scale competition motif and inhibition stabilization.

321

## 322 **Inhibition stabilization of chaotic dynamics in the data-derived agents**

323 Consistent with previous work on inhibition-stabilized networks<sup>53,55–57</sup>, the toy model suggests that  
324 inhibition plays a critical role in constraining neural instability. Particularly, there is an empirical phase  
325 transition between dynamical regimes modulated by the level of inhibition. Intermediate levels of  
326 inhibition permit chaotic dynamics, too little causes runaway activity, and high levels lead to fixed points.  
327 We tested these ideas from the toy model in our data-derived agents. We adjusted inhibition amplitude in  
328 the data-derived agents by adding a bias to all recurrent connection weights. Negative bias increased  
329 inhibition, and positive bias reduced it (Fig. 5g). To isolate transitions between dynamical regimes from  
330 mere transformations of the dynamics, we retrained the linear output weights to minimize changes in  
331 velocity outputs, reducing behavioral effects of potential rescaling, translations, or rotations in the RNN's  
332 neural activity.

333

334 Overall, observations supported the theory of a phase transition modulated by inhibition. For the  
335 behavioral outputs of the data-derived agents, increasing overall inhibition reduced variation and  
336 eventually produced converging running trajectories, consistent with fixed point attractors. Conversely,  
337 reducing overall inhibition blurred left–right trajectory separation, eventually eliminating it, consistent  
338 with runaway activity (Fig. 5h). Reducing inhibition impaired binary choice performance far more than  
339 increasing it (Fig. 5i). For the RNN unit activity of the data-derived agents, reducing inhibition increased  
340 activity but reduced choice separability and trial-to-trial variability within a choice, consistent with  
341 runaway activity (Fig. 5j–l). Thus, the level of inhibition is a critical parameter for neural dynamics, with  
342 chaotic regimes emerging only at intermediate inhibition.

343

## 344 **Chaotic dynamics facilitate behavioral adaptation**

345 Because chaotic dynamics generate variability across trials for the same decision and maintain a reservoir  
346 of possible trajectories, they may enhance behavioral adaptation, in particular to allow animals to achieve  
347 the same task goal in different ways. In navigation, the same discrete decision (turn left or right) can be  
348 achieved by many running trajectories, with changes in the environment altering which trajectories are  
349 preferable. To test if chaotic dynamics can serve this type of adaptation, we asked the data-derived agents  
350 to navigate around novel obstacles in the Y-maze to reach the goal location. Aided by the closed-loop  
351 design, the data-derived agents can be tested in environments beyond those originally used to collect the  
352 data for their training, which allows testing behavioral capacities of the data-derived neural mechanisms.

353

354 In the Y-maze, we placed two obstacles in the left arm, requiring new steering movements to reach the  
355 reward zone (Fig. 6a). The obstacles were positioned so they could not be bypassed by simply adding a  
356 bias to the velocity output. On first exposure, nearly all left-choice trials were blocked by the obstacles  
357 (Fig. 6a, middle). Then, with the RNN weights fixed, we used reinforcement learning to train only the  
358 linear output weights that translate the agent's internal dynamics into velocity outputs<sup>58</sup> (Fig. 6c,  
359 Methods). This preserved the data-derived dynamics obtained from fitting to empirical neural activity and  
360 tested their functional capacity for learning the new task. After re-training, the agent avoided the obstacles

361 and reached the reward location with high accuracy (Fig. 6a, bottom). To bypass the first obstacle, the  
362 agent steered more left early in the arm epoch, then steered more right later to avoid the second obstacle  
363 (Fig. 6b, top). At the same time, it steered more in trials expecting collisions and less so when far away  
364 from obstacles, avoiding overcompensation (Fig. 6b, middle and bottom).

365

366 To test whether this spatiotemporal adaptability is a property afforded by chaotic dynamics, we also  
367 evaluated agents pre-trained on trial-averaged data that exhibited continuous or point attractor dynamics  
368 (Fig. 3). For point-attractor agents, we added noise to match the trial-to-trial variability of the other agents  
369 for a fair comparison. After re-training with the same protocol, continuous- and point-attractor agents  
370 improved less than the agents with chaotic dynamics (Fig. 6c). Additionally, to assess the intrinsic  
371 expressivity of the neural dynamics independent of the complex learning dynamics of reinforcement  
372 learning, we took neural activity generated by the closed-loop agents in the original environment and held  
373 it fixed, and only retrained the mapping from neural activity to behavior by supervised learning.  
374 Consistently, the agents with chaotic dynamics achieved the best performance among the three types (Fig.  
375 6d). Thus, chaotic dynamics confer greater behavioral adaptability than traditional attractors.

376

377 The benefit of chaotic dynamics to adaptability is broadly consistent with the theory of reservoir  
378 computing, in which a weakly chaotic system works as a rich reservoir of nonlinear functions that can be  
379 linearly combined to generate diverse target outputs<sup>59-61</sup>. We examined this idea in the data-derived agents  
380 by applying dynamic mode decomposition (DMD)<sup>62</sup> on the RNN activity. DMD is a data-driven approach  
381 that uncovers the dominant oscillatory components of dynamical systems. The agents with chaotic  
382 dynamics had more slowly decaying oscillatory components compared to the ones with continuous or  
383 point attractors (Fig. 6e, f, Extended Data Fig. 7). In theory, by tuning the linear readout layer, these  
384 oscillatory components can be combined in various ways, with specific loading and phase offset of each  
385 component, therefore generating a large variety of outputs (see Methods for details).

386

## 387 **Discussion**

388 In many sequential behaviors, animals must achieve reliable outcomes while expressing variability in the  
389 specific trajectories used to reach those outcomes. Our results identify a dynamical regime in which these  
390 two demands can coexist. We show that neural population activity can operate in a chaotic regime that  
391 remains confined within choice-specific attractor landscapes, such that trial-to-trial variability is generated  
392 intrinsically while higher-level goals remain robust. In this regime, small perturbations lead to divergent  
393 neural trajectories on individual trials, yet population activity remains constrained within a basin  
394 associated with a given choice, preserving behavioral accuracy. These findings suggest that chaos can  
395 serve a viable functional role in neural population dynamics, enabling structured variability rather than  
396 reflecting unstructured instability or noise.

397

398 This functional regime - chaotic at the level of individual trajectories yet constrained at the level of  
399 behavioral goals - can be difficult to identify using standard approaches. When neural dynamics are  
400 evaluated in open-loop settings by fitting neural activity alone or prioritizing trajectory predictability, such

401 regimes often appear unstable, uninterpretable, or poorly predictive, and are therefore disfavored during  
402 model fitting or selection<sup>15,29</sup>. In contrast, when the same dynamics are evaluated under closed-loop  
403 interaction with the environment, their importance becomes apparent: chaotic dynamics can support  
404 reliable behavior while also generating rich trial-to-trial variability. This distinction helps explain why  
405 chaotic dynamical regimes, despite being theoretically well characterized<sup>30,47,63</sup>, have rarely been  
406 implicated as viable mechanisms for neural computation in behaviorally grounded models.

407  
408 The contrast between open-loop and closed-loop conditions reveals closed-loop interactions as a  
409 fundamental theoretical constraint on neural dynamics. In natural behavior, neural activity, behavioral  
410 outputs, and sensory inputs are mutually coupled, such that variability in neural state necessarily reshapes  
411 future inputs and behavioral demands<sup>21–24</sup>. When this coupling is ignored, the causal link between neural  
412 dynamics and behavior is broken, which can lead to systematic misclassification of which dynamical  
413 regimes are causally related to behavior. Our results show that identifying neural dynamics capable of  
414 supporting robust yet flexible behaviors over extended trajectories requires analyzing and modeling them  
415 in closed loop. Under this key constraint, dynamics that can appear unstable or uninformative in open-  
416 loop analyses instead emerge as well suited to behaviors that demand robustness at the level of  
417 goals/outcomes and flexibility at the level of trajectories.

418  
419 Our findings also clarify how chaotic dynamics relate to classic dynamical frameworks that emphasize  
420 either stability or sequential structure. A long line of work has proposed that attractor dynamics can  
421 segregate neural activity into discrete choices and support robust choice formation, including in decision-  
422 making models that rely on recurrent competition and slow reverberation<sup>8–12</sup>. At the same time, many  
423 models of sequential activity emphasize mechanisms that generate stable trajectories through recurrent  
424 structure and propagation<sup>13–15</sup>. In the regime we identify here, these perspectives can be integrated: the  
425 global attractor landscape provides robustness at the level of the choice or goal, while within each basin,  
426 low-dimensional chaotic dynamics generate rich within-choice variability in neural trajectories. In this  
427 sense, chaotic attractors allow a single network to maintain stable high-level outcomes while supporting  
428 flexible, variable realizations of those outcomes over extended trajectories.

429  
430 We also motivate a reframing of trial-to-trial variability in neural population activity. Prior studies across  
431 sensory representation, decision-making, and movement have emphasized different interpretations of  
432 variability. These include the possibility that variability aligned with coding dimensions can act as  
433 information-limiting noise correlations<sup>64,65</sup>, that variability may largely lie outside task-relevant  
434 dimensions and therefore be ignorable for downstream computations<sup>66–70</sup>, or that variability can reflect  
435 meaningful internal state or history dependence that supports computation over time<sup>1,3–5,71–73</sup>. Here we  
436 propose a complementary perspective grounded in task demands. In sequential behaviors that admit many  
437 valid trajectories toward the same goal, variability in how the goal is realized can be an essential  
438 component of the solution rather than a nuisance. In our setting, focusing only on the trial-averaged  
439 solution risks imposing an overly rigid view of the underlying dynamics. Consistent with this, restricting  
440 training to trial-averaged neural and behavioral trajectories has classically produced point-attractor or

441 continuous-attractor dynamics rather than chaotic dynamics, underscoring that trial-resolved variability is  
442 informative about the underlying regime and can be necessary to reveal it.

443

444 Our findings are broadly consistent with ideas from reservoir computing, which emphasize the  
445 computational capacity of rich recurrent dynamics and have highlighted “edge of chaos” regimes as  
446 potentially favorable for balancing memory and sensitivity to inputs<sup>29,59–61,74</sup>. In that literature, the term  
447 “reservoir” typically refers to a large recurrent network whose high-dimensional dynamics provide a rich  
448 set of nonlinear transformations, such that only a task-specific readout needs to be trained to realize  
449 particular computations. Here we adapt the term to be more behaviorally grounded and use “dynamical  
450 reservoirs” to describe how the chaotic attractors we identify can maintain a distribution of possible  
451 navigation trajectories that can be flexibly selected, and in our simulations adapted, to meet task  
452 constraints. Importantly, the variability generated by the dynamical reservoirs here differs from  
453 unstructured noise in two ways that matter for behavior. First, chaotic dynamics are deterministic, and the  
454 nonlinear functions embedded in their dynamics can be combined to generate structured output patterns  
455 in ways that injected noise typically cannot. Second, because the chaotic dynamics remain bounded by  
456 the overall attractor landscape, they can support variability without undermining the robustness of goal-  
457 level performance that unbounded instability or noise might compromise.

458

459 The circuit structure that emerged in our data-derived agents extends commonly studied motifs in a way  
460 that naturally supports constrained variability. Consistent with canonical decision-making models and  
461 recent experimental work, our modeling revealed an opponent inhibition motif between pools of neurons  
462 with opposite choice selectivity<sup>8,9,51–53</sup>. We also observed asymmetric connections that support sequence  
463 propagation, as proposed in many sequence-generation models<sup>13,20</sup>. Strikingly, within a group of neurons  
464 sharing the same choice selectivity and similar sequential position, we found additional competition motifs  
465 organized by tuning to specific running trajectories, suggesting a fine-grained structure that can support  
466 within-choice variability. This type of fine-grained competition is reminiscent of that found within  
467 populations for sensory perception<sup>75,76</sup>. In our toy model simulations, a two-scale competition motif  
468 coupled with inhibitory stabilization was sufficient to produce low-dimensional chaotic dynamics with  
469 bounded variability, and perturbations of inhibition in the data-derived agents produced corresponding  
470 transitions between fixed-point, chaotic, and runaway regimes. Together, these results suggest that  
471 constrained instability, and specifically chaos stabilized by inhibition, can provide a mechanistically  
472 plausible route to generating structured trial-to-trial variability while maintaining robust choice separation,  
473 consistent with broader ideas about inhibition-stabilized networks<sup>53,55,56,77</sup>.

474

475 Since closed-loop coupling is a defining constraint on viable neural dynamics, we developed a novel  
476 modeling setup of data-derived agents. This approach combines data-constrained dynamical systems  
477 reconstruction<sup>19,31</sup> with environment-interacting task-performing agents<sup>32,41</sup>. This setup allows modeling  
478 of behaviors that involve extended, self-paced action sequences while keeping the internal dynamics  
479 constrained by recorded neural activity and behavior on individual trials. The closed-loop design is  
480 essential because small action errors can otherwise compound into divergent environmental state

481 distributions, and the trial-resolved fitting is essential because it preserves the structured variability that is  
482 informative about the underlying dynamics. Within this data-constrained, closed-loop modeling objective,  
483 the data-derived agents developed here can be analyzed directly as dynamical systems. They learn a  
484 generative parameterization that captures a distribution of neural trajectories rather than reproducing exact  
485 single-trial trajectories, which is expected in the presence of chaos. This framework enables controlled  
486 perturbations and dynamical analyses, such as Lyapunov spectra and unstable manifold geometry, that  
487 characterize the stability structure of the learned regime and link it to circuit-level interaction motifs. It  
488 thereby provides a tractable setting in which candidate dynamical regimes and mechanisms can be  
489 identified, probed, and used to generate experimentally testable predictions.

490  
491 On the other hand, deep reinforcement learning agents, which operate in closed loop, are starting to be  
492 used in neuroscience to model behavior and representations in complex tasks<sup>35–40</sup>. However, because  
493 those agents are optimized for task performance rather than constrained by trial-resolved neural population  
494 dynamics from real recordings, they address a different modeling objective: learning effective policies,  
495 rather than identifying which internal dynamical regimes are compatible with observed neural population  
496 activity during behavior.

497  
498 Despite long-standing theoretical interest, the presence of chaos in the brain remains controversial, and  
499 experimentally establishing chaotic dynamics *in vivo* can be practically challenging<sup>78–82</sup>. While our data-  
500 derived agents provide evidence that chaotic regimes are consistent with the neural and behavioral  
501 constraints in this task, directly demonstrating chaos in biological circuits will likely require perturbations  
502 and single-trial analyses that can distinguish unstable dynamics from noise. In this respect, the circuit  
503 motifs suggested by the agents and toy models, such as structured competition within choice-selective  
504 populations and inhibition-dependent regime transitions, may offer more accessible near-term targets for  
505 experimental testing than direct dynamical reconstruction alone. More broadly, our recordings span  
506 multiple cortical areas, and we modeled all recorded neurons together without area distinctions. An  
507 interesting topic is whether different areas of the cortex contribute in distinct ways to the regime we  
508 identified<sup>31,83,84</sup>. In this work, we were limited in our ability to address this question because our data had  
509 imbalances in sampling across areas and limited sampling in some areas on individual sessions, likely due  
510 to inhomogeneities of imaging quality across the large cranial window (Extended Data Fig. 1a). As a  
511 result, it was challenging to address the contributions of some individual areas in some sessions and to  
512 ensure that any differences were genuine instead of due to the number of neurons from an area included  
513 in the agent. Future work with improved multi-area sampling or model structures that pool neurons across  
514 sessions could clarify whether and how different areas contribute distinctly to the dynamics.

## 515 516 **Acknowledgements**

517 We thank members of the Harvey and Rajan labs for discussions and feedback on the manuscript. This  
518 work was supported by NIH grants R37NS089521 (C.D.H), R01NS143140 (C.D.H), DP1MH125776  
519 (C.D.H.), RF1DA056403 (K.R.), U01NS136507 (K.R.), and funding from James S. McDonnell  
520 Foundation (220020466 to K.R.), Simons Foundation (Pilot Extension-00003332-02 to K.R.), McKnight

521 Endowment Fund (K.R.), CIFAR Azrieli Global Scholar Program (K.R.), NSF (2046583 to K.R.),  
522 Harvard Medical School Neurobiology Lefler Small Grant Award (K.R.), and Harvard Medical School  
523 Dean's Innovation Award (K.R.). This work has been made possible in part by a gift from the Chan  
524 Zuckerberg Initiative Foundation to establish the Kempner Institute for the Study of Natural and Artificial  
525 Intelligence at Harvard University.

526

### 527 **Author Contributions**

528 S.Z., K.R., C.D.H conceived of the project. S.Z. developed the modeling approach, implemented all  
529 models, and performed the data analysis with guidance from K.R. and C.D.H. R.P.B. provided input on  
530 the model development and implementation. C.A. collected the experimental data with guidance from  
531 C.D.H. S.Z., K.R., and C.D.H. wrote the manuscript with feedback from all authors.

532

### 533 **Data Availability**

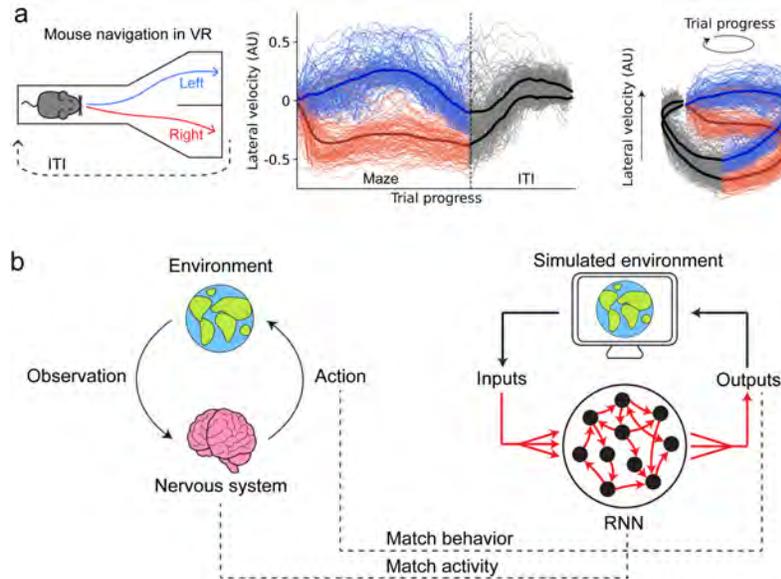
534 Data are available upon request to the corresponding authors.

535

### 536 **Code Availability**

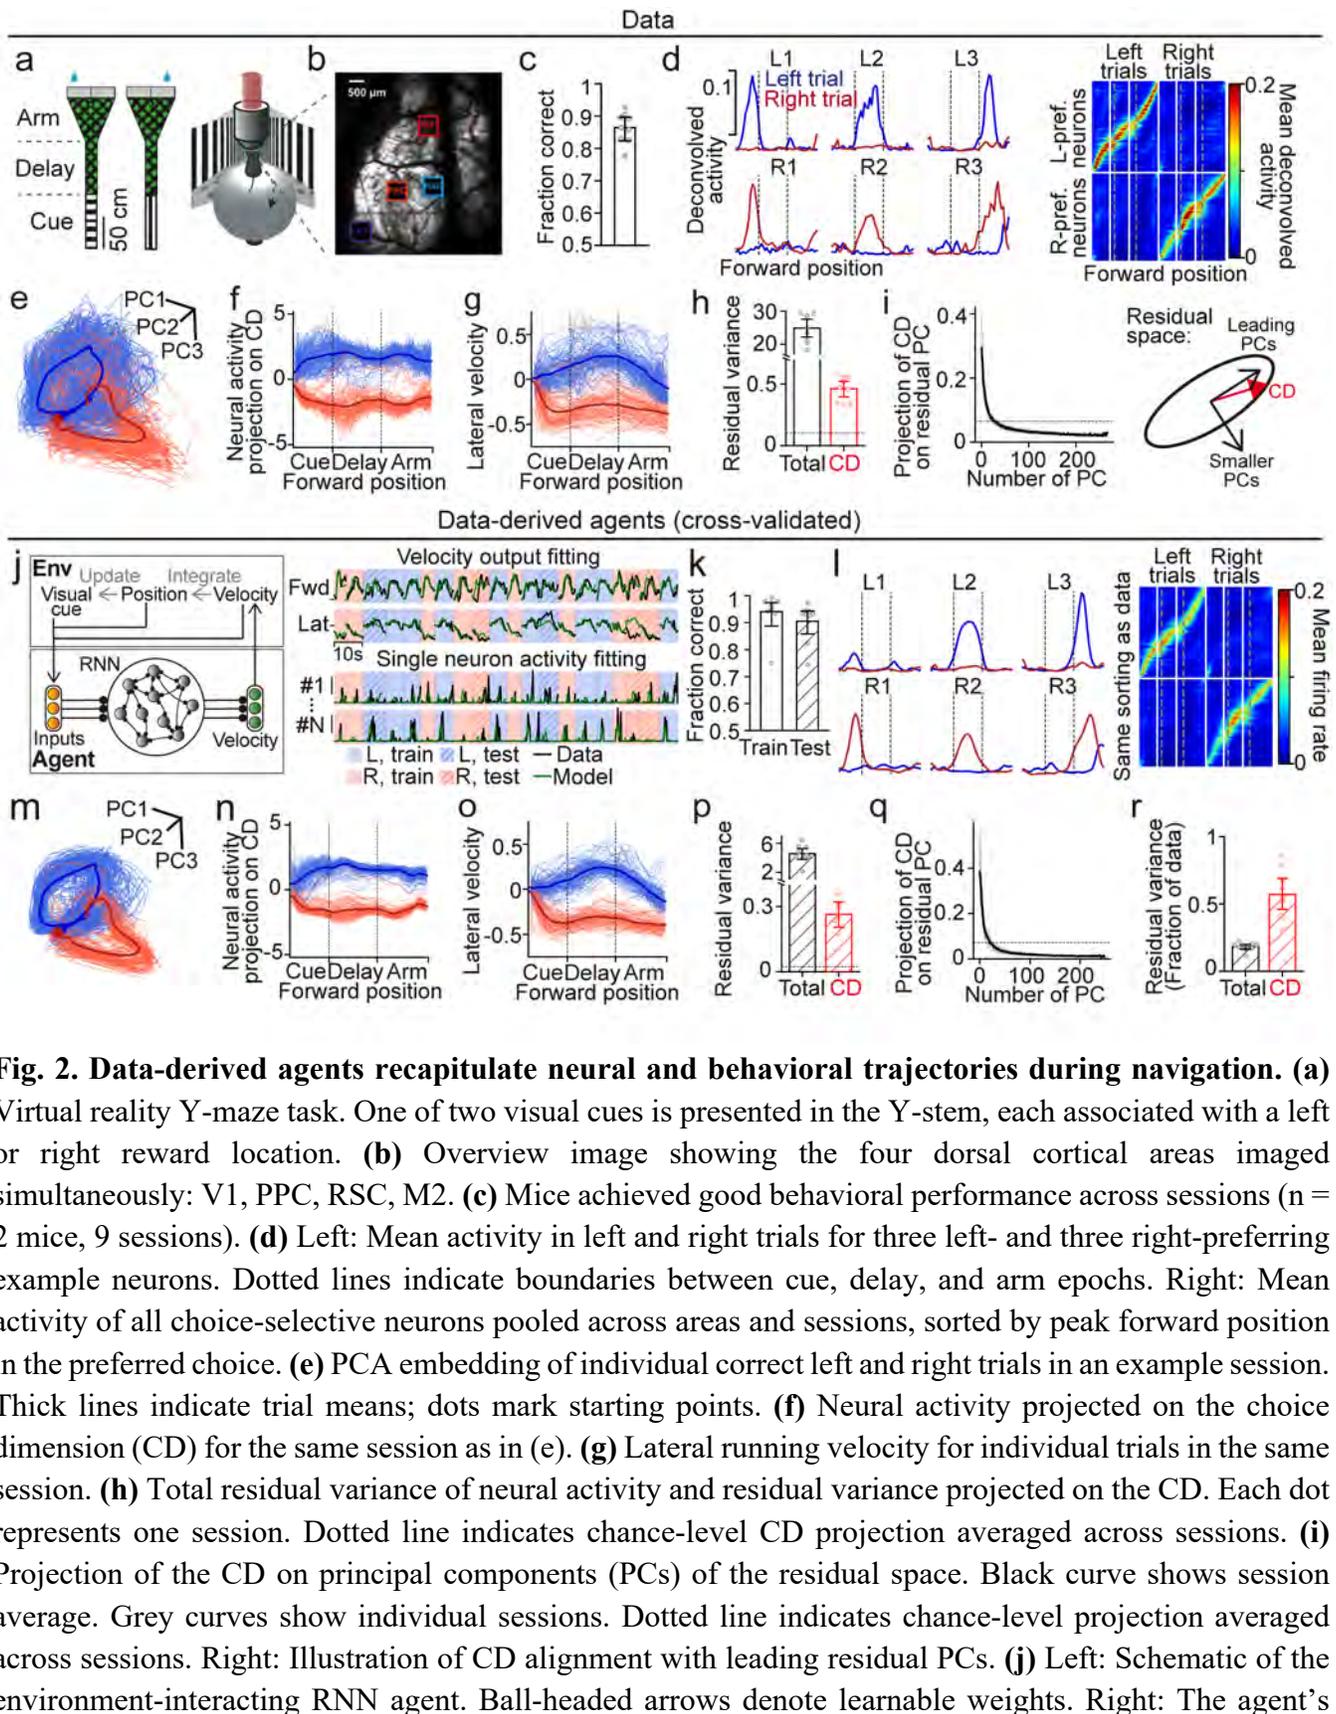
537 All models and analyses were implemented in Python (<https://www.python.org>). The code to train data-  
538 derived agents with ARCTIC will be made publicly available upon publication.

539



540

541 **Fig. 1. Data-derived agents jointly model neural dynamics and closed-loop behavior.** (a) Left and  
542 middle: Mouse navigation in a virtual reality (VR) Y-maze shows diverse locomotor trajectories toward  
543 the same discrete choice. Thin lines indicate individual left or right trials in one mouse session; thick lines  
544 indicate trial-averaged trajectories. ITI, inter-trial interval. Right: Representing trial progress in polar  
545 coordinates reveals cycle-like geometry, but trajectories do not converge to fixed limit cycles, instead  
546 forming choice-specific distributions. (b) Schematic of ARCTIC (Activity Reconstruction in Closed  
547 loop). RNN-based agents are trained to perform cognitive tasks such as navigation in closed-loop  
548 simulated environments, while constrained by neural and behavioral data from real task-performing  
549 animals. Red arrows denote learnable connection weights.



550

551 **Fig. 2. Data-derived agents recapitulate neural and behavioral trajectories during navigation.**

552 Virtual reality Y-maze task. One of two visual cues is presented in the Y-stem, each associated with a left

553 or right reward location. **(b)** Overview image showing the four dorsal cortical areas imaged

554 simultaneously: V1, PPC, RSC, M2. **(c)** Mice achieved good behavioral performance across sessions (n =

555 2 mice, 9 sessions). **(d)** Left: Mean activity in left and right trials for three left- and three right-preferring

556 example neurons. Dotted lines indicate boundaries between cue, delay, and arm epochs. Right: Mean

557 activity of all choice-selective neurons pooled across areas and sessions, sorted by peak forward position

558 in the preferred choice. **(e)** PCA embedding of individual correct left and right trials in an example session.

559 Thick lines indicate trial means; dots mark starting points. **(f)** Neural activity projected on the choice

560 dimension (CD) for the same session as in (e). **(g)** Lateral running velocity for individual trials in the same

561 session. **(h)** Total residual variance of neural activity and residual variance projected on the CD. Each dot

562 represents one session. Dotted line indicates chance-level CD projection averaged across sessions. **(i)**

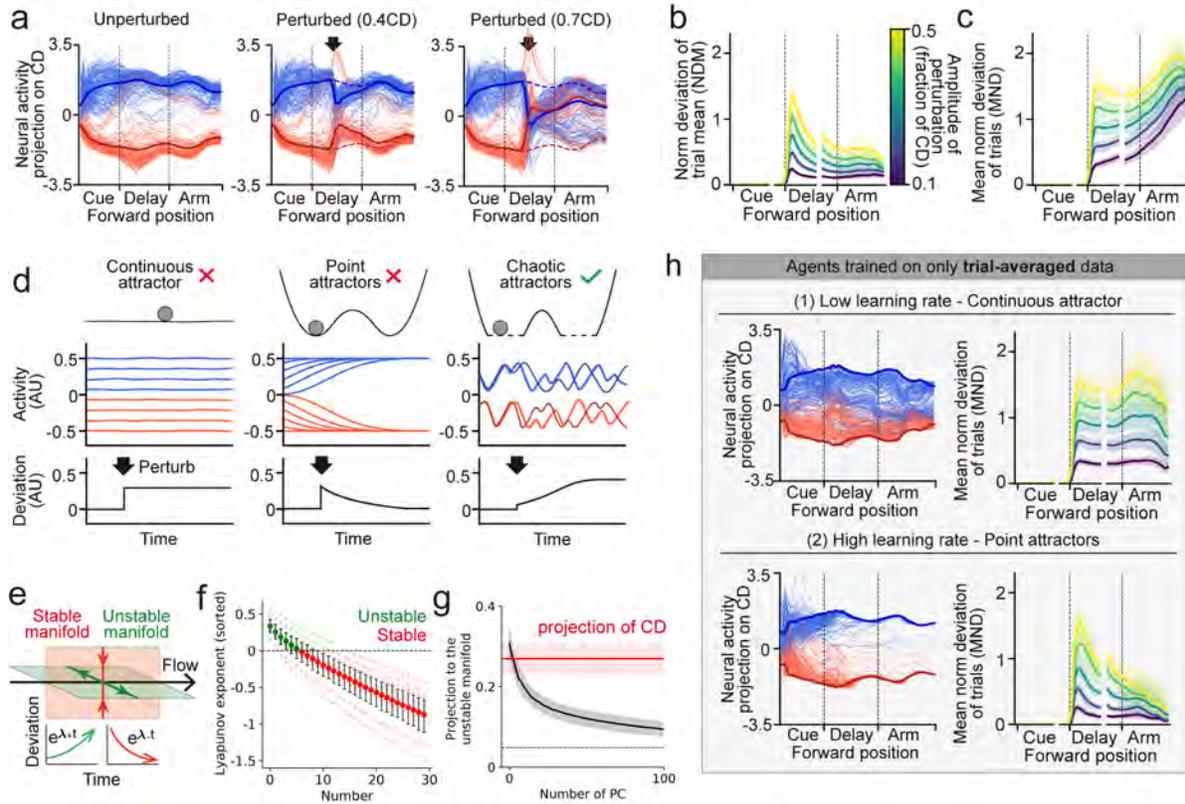
563 Projection of the CD on principal components (PCs) of the residual space. Black curve shows session

564 average. Grey curves show individual sessions. Dotted line indicates chance-level projection averaged

565 across sessions. Right: Illustration of CD alignment with leading residual PCs. **(j)** Left: Schematic of the

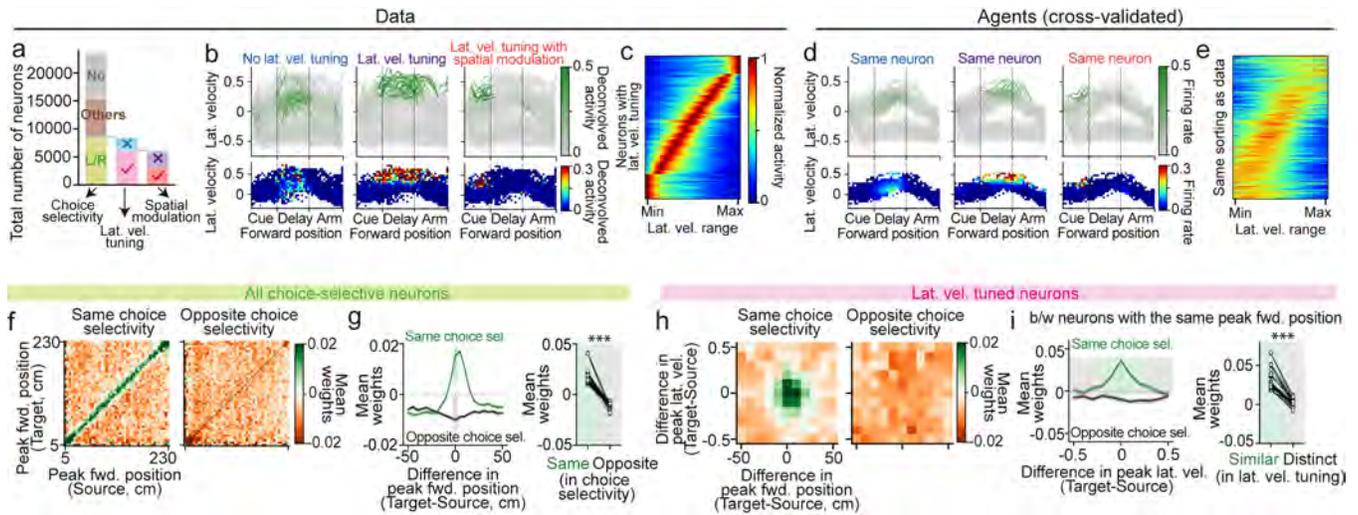
566 environment-interacting RNN agent. Ball-headed arrows denote learnable weights. Right: The agent's

567 velocity output and single-neuron activity are fit to single-trial data. One agent is trained per cross-  
568 validation fold for each session. **(k)** Agent performance on training and test trials (fraction correct). Each  
569 dot represents one session, averaged across five-fold cross-validation. In the following panels (l-r), the  
570 agents are evaluated for their neural activity and behavior in cross-validated test trials. **(l)** Left: Agent-  
571 predicted trial-mean activity for the six example neurons in (d). Right: Trial-mean activity of units pooled  
572 across agents, indexed as in (d). **(m-n)** Single-trial agent activity from the same session as in (e) projected  
573 on the PC space of the data (m) or the CD of the data (n). **(o)** Agent-generated lateral velocity in the same  
574 session. **(p)** Total residual variance generated by agents and its projection on the CD. Dotted line indicates  
575 chance level. **(q)** Projection of the CD on residual PCs of the agents. **(r)** Fraction of total residual variance  
576 in the data captured by the agents, compared with the fraction of residual variance on CD captured by the  
577 agents. Error bars in this and the following figures indicate 95% confidence interval obtained by  
578 bootstrapping.  
579



580

581 **Fig. 3. Perturbative analyses reveal low-dimensional chaotic attractors in data-derived agents.** (a)  
 582 Neural activity generated by a data-derived agent in an example session (only including correct trials when  
 583 unperturbed), projected on the choice dimension (CD) of the agent. Left: Unperturbed. Middle and right:  
 584 Perturbed along the CD toward the opposite trial type (amplitude 0.4 or 0.7× the difference between mean  
 585 left and right CD projections). Black arrow indicates perturbation time. (b) Evolution of the norm  
 586 deviation of the trial-mean (NDM) activity following perturbations of varying amplitudes. Deviations for  
 587 left and right trial means are computed separately and then averaged. (c) Evolution of the mean norm  
 588 deviation (MND) of individual-trial activity. In (b) and (c), only trials with correct post-perturbation  
 589 choices are included, so the deviation reveals dynamics within each choice. (d) Schematics of different  
 590 dynamical regimes. Top: Potential landscapes. In chaotic attractors dotted lines are used to indicate non-  
 591 canonical landscapes. Middle: Temporal evolution of activity from different initial states. Bottom:  
 592 Responses to perturbation at the black arrow. Deviation remains constant for continuous attractors, decays  
 593 for point attractors, and shows bounded amplification for chaotic attractors (e) Decomposition of  
 594 dynamics onto stable and unstable manifolds. Inset: Evolution of perturbations on the stable and unstable  
 595 manifolds. (f) First 30 Lyapunov exponents of data-derived agents. Light dots represent agents trained on  
 596 individual sessions; solid dots indicate session average. (g) Projection of residual-space principal  
 597 components (PCs) and the CD on the unstable manifold, computed per session and averaged. Dotted line  
 598 indicates chance-level projection. (h) Learned dynamics of agents trained on trial-averaged data (low or  
 599 high learning rate, top and bottom). Left: The thick blue (red) line shows the neural trajectory simulated  
 600 from the averaged initial states of left (right) trials; thin lines show neural trajectories simulated from the  
 601 initial states of individual trials. Right: Responses to perturbation.

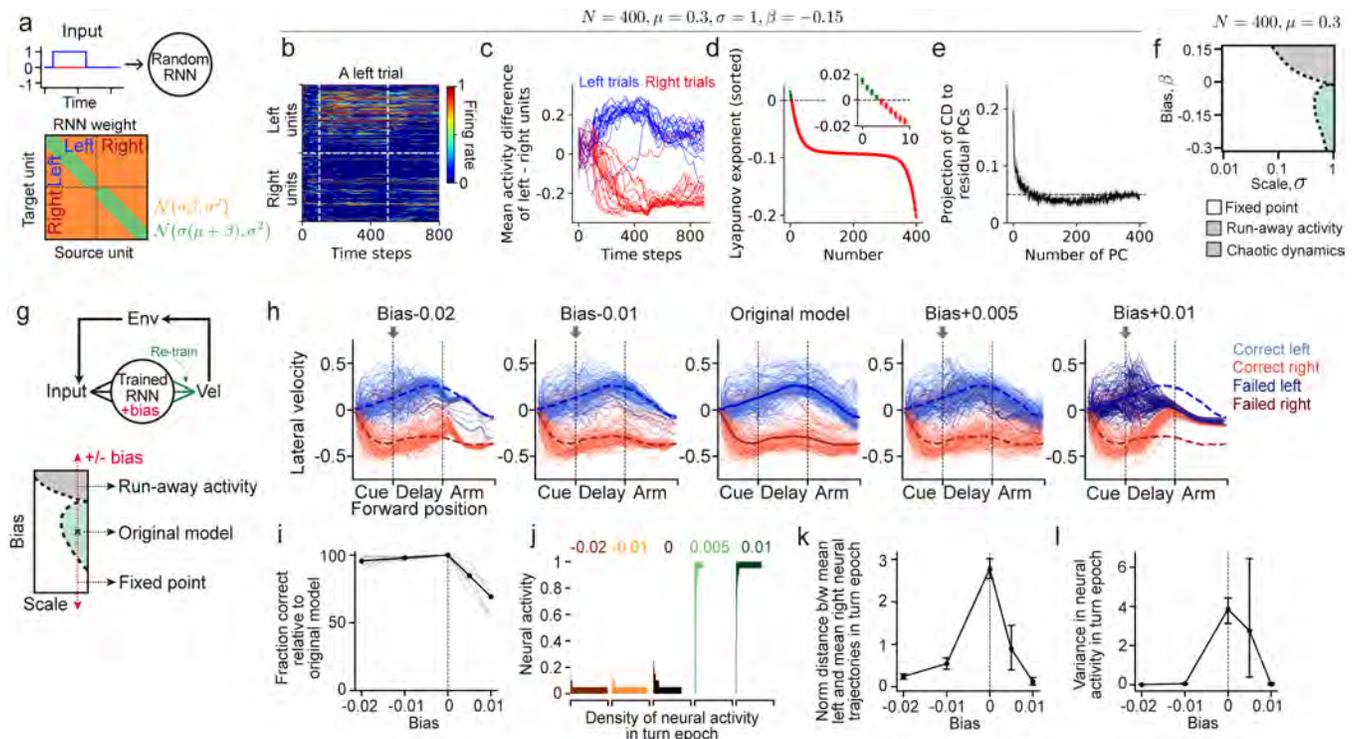


602

603

**Fig. 4. Weight analyses reveal competitive connectivity motifs on two scales. (a)** Classification of neuron representations, pooled across areas and sessions. Neurons are first classified by choice selectivity (no selectivity, left/right selectivity, or others with multiple fields). Left or right choice-selective neurons are further classified by lateral velocity tuning, and velocity-tuned neurons are further subdivided by whether their tuning to lateral velocity is modulated by forward position in the maze. **(b)** Activity of three example left choice-selective neurons. Top: Running trajectories color-coded by instantaneous neural activity. Bottom: Bin-averaged neural activity in conjunctive bins of forward position and lateral velocity. **(c)** Neurons have diverse lateral velocity tuning that spanned the entire lateral velocity range. Activity within each neuron's position field is plotted over its available lateral velocity range, defined by the minimum (most rightward) and maximum (most leftward) lateral velocity within the position field. Neurons are sorted by their peaks. Activity is normalized. **(d)** Corresponding RNN unit activity for the three example neurons in (b), from cross-validated test trials pooled across folds. **(e)** Lateral velocity tuning of RNN units pooled across agents, in cross-validated test trials. Units are indexed as in (c) and exhibit similar tuning as the data. **(f)** Mean connection weight between neurons with the same or opposite choice selectivity, organized by peak forward positions of source and target neurons. **(g)** Left: Mean weight as a function of the difference in peak forward position between target and source neurons. Right: Mean weight between same- or opposite-choice neurons with matched peak forward position. Dots represent individual cross-validation folds of all sessions. Shaded areas in the left panel indicate the subsets analyzed in the right panel. **(h)** Mean weight as a function of the differences in peak forward position and peak lateral velocity between target and source neurons. **(i)** Left: Mean weight between neurons with matched peak forward position, plotted against the difference in peak lateral velocity. Right: Between neurons with the same choice selectivity and matched peak forward position, mean weight is further modulated by the difference in peak lateral velocity. Shaded areas in the left panel indicate subsets analyzed in the right panel.  $***P < 0.001$ , one-sided Wilcoxon signed-rank test.

626



627

628 **Fig. 5. Competitive connectivity motifs and inhibitory stabilization give rise to chaotic dynamics. (a)**

629 Top: Schematic of the toy model. Blue and red lines represent an external input to left and right units

630 during left trials. Bottom: Schematic of the connectivity structure of the toy model. Recurrent weights are

631 drawn from two Gaussian distributions (green and orange partitions). Units within each choice population

632 are ordered according to a hypothetical tuning parameter. A positive  $\mu$  imposes competition both between

633 choice populations (on- vs. off-diagonal blocks) and within each choice population (diagonal ridges within

634 on-diagonal blocks). (b-e) Analyses of the toy model with  $\mu = 0.3, \sigma = 1, \beta = -0.15$ . (b) Unit activity

635 in response to the external input during a left trial. Vertical white dotted lines indicate input onset and

636 offset. (c) Difference between mean activity of left and right units, in 20 left and 20 right trials simulated

637 using identical connectivity but different activity initializations. (d) Lyapunov exponent spectrum of the

638 toy model, averaged across 50 random weight realizations. (e) Projection of the choice dimension (CD)

639 on principal components (PCs) of the residual space. (f) Phase diagram of how the dynamical regimes

640 transition with different values of  $\sigma$  and  $\beta$ . Regimes are determined by simulations of 10 random weight

641 realizations (Extended Data Fig. 6). (g) Top: Schematic of the weight perturbation procedure in data-

642 derived agents. A bias is added to the recurrent weights at delay onset, followed by re-training only the

643 output weights to minimize changes in velocity outputs. Bottom: Schematic of the hypothesized

644 modulatory effect of inhibition. (h) Running trajectories for individual left and right trials in an example

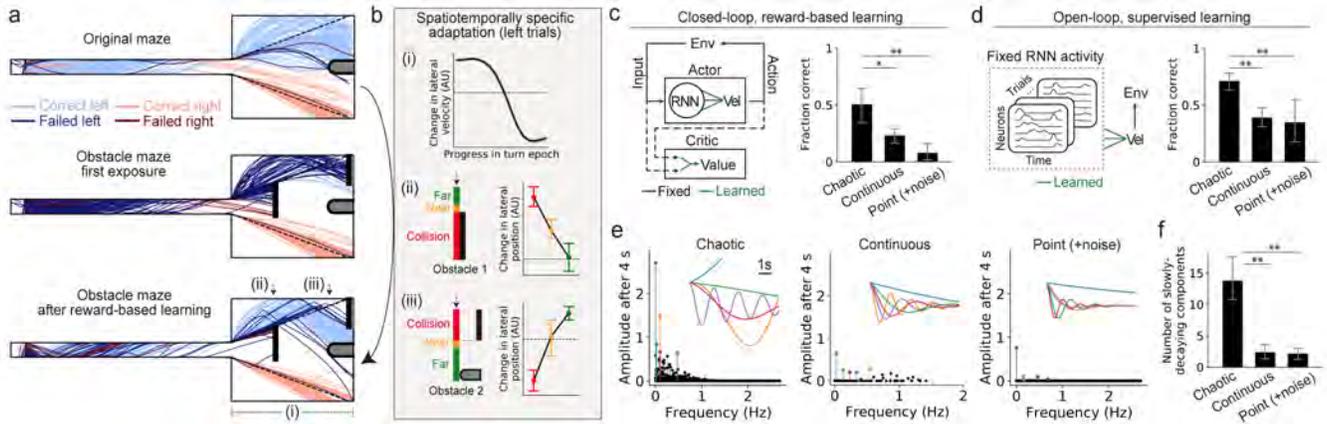
645 session under different weight biases. Arrow marks where the bias begins to take effect. Blue and red

646 dotted lines indicate trial-averaged trajectories in the unmodified agent. (i) Choice correctness relative to

647 the unmodified agents as a function of added bias. Thin lines represent individual sessions; the thick line

648 shows session average. (j-l) Effects of inhibition level on neural activity amplitude (j), separation between

649 left and right trial activity ( $k$ ), and trial-to-trial variance of neural activity ( $l$ ) in the agents during the arm  
650 epoch. Results illustrate transitions between fixed-point, chaotic and runaway regimes.

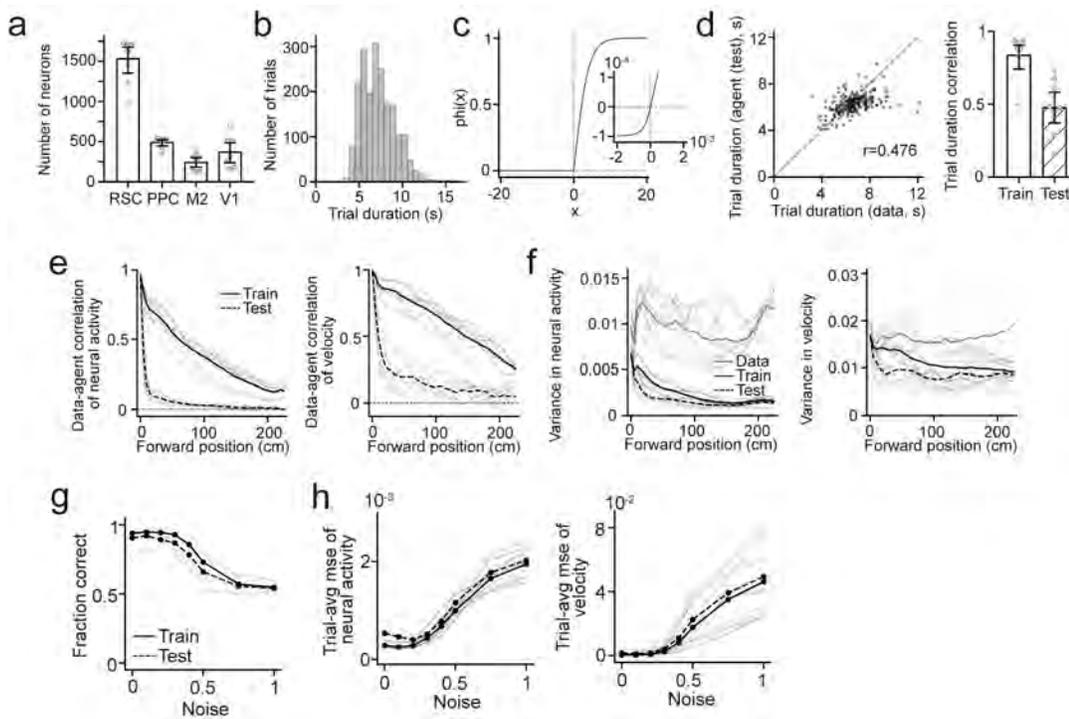


651  
652

653 **Fig. 6. Chaotic dynamics facilitate flexible adaptation in a novel obstacle environment. (a)** Running  
654 trajectories from an example session. Top: A data-derived agent's running trajectories in the original maze  
655 where the neural and behavioral data used to train the agent was collected. Middle: On first exposure to  
656 the obstacle maze, most left-choice trials are blocked by the obstacles (failed trials in dark blue, obstacles  
657 shown as black vertical lines). Bottom: After reward-based re-training of only output weights, with  
658 recurrent weights held fixed, the agent avoids obstacles while maintaining high choice accuracy. **(b)**  
659 Spatiotemporally specific adaptation in the example session. Top: Change in lateral velocity (before vs.  
660 after reward-based learning) as a function of forward progress during the arm epoch, averaged across all  
661 correct left trials. Middle and bottom: Change in lateral position in front of the first or second  
662 obstacle depends on whether trajectories would have collided or were far from the obstacles. **(c)** Left: Schematic  
663 of the actor-critic algorithm used for reward-based learning. RNN weights are fixed while output weights  
664 are further trained to produce actions with higher values. Right: Performance in the obstacle maze (left  
665 trials only) compared across agents with chaotic dynamics (derived from individual-trial data) and agents  
666 with continuous or point attractor dynamics (derived from trial-averaged data). Performance is computed  
667 on cross-validated test trials and averaged across sessions. **(d)** Left: Supervised learning setup. RNN  
668 activity sequence generated by data-derived agents in the original maze is recorded and fixed while output  
669 weights are retrained to produce desired velocity outputs. Right: Performance comparison of the three  
670 agent types in the obstacle maze (left trials only) under supervised learning. Performance is computed  
671 on cross-validated test trials and averaged across sessions. **(e)** DMD eigenvalue spectra of RNN activity in  
672 agents with different dynamics (example session). The amplitudes and frequencies of components are  
673 determined by the real and imaginary part of the eigenvalues of a fitted linear operator. Insets show  
674 temporal evolution of the top five components. **(f)** Number of slowly decaying components in agents with  
675 different dynamics. Components with growth rates between  $-0.2$  and  $0.2$   $s^{-1}$  are considered as slowly  
676 decaying. Bars show session average.  $*P < 0.05$ ;  $**P < 0.01$ , one-sided Wilcoxon signed-rank test.

677

678



679

680

681

**Extended Data Fig. 1. Additional data information and evaluations of the data-derived agents.**

682

(a) Number of neurons recorded in each of the four brain areas. Dots are individual sessions. (b)

683

Distribution of the trial duration (time to traverse the maze), pooled across sessions. (c)

684

Custom activation function for agent's RNN units. Inserted panel shows the function around origin. See Methods for details. (d)

685

The time taken by the data-derived agents to traverse the maze is also variable and correlates with the data. Left: Duration of cross-validated test trials from the data-derived agent plotted against the

686

corresponding trial duration of the mouse trials in one example session. Right: Correlation of trial duration

687

of all sessions in training and testing trials. (e) Residual correlation between data-derived agents and data,

688

which reflects how the trial-to-trial variation of the agents matched that of the data. Calculated from

689

position-binned variables. Thin lines are individual sessions, averaged across each session's five cross-

690

validation folds. Thick black lines represent the average across all sessions. As discussed in the main text,

691

it is challenging for the data-derived agents to reproduce the exact trial-by-trial trajectories, instead it

692

learned a generalizable distribution. (f) Residual variance in neural activity and velocity. As in each trial

693

the agents are initialized with the first data point of that trial, the residual variance of the agents starts at a

694

similar level to the data and later shows some decrease. It converges to a non-zero value later in the maze.

695

This potentially indicates that a fraction of (but not all) the variance in the data is recapturable by internal

696

dynamics. (g-h) Robustness of the data-derived agents against random noise. Note that the noise level

697

during training is 0.1. As the data-derived agents are evaluated with increasing noise, they initially stay

698

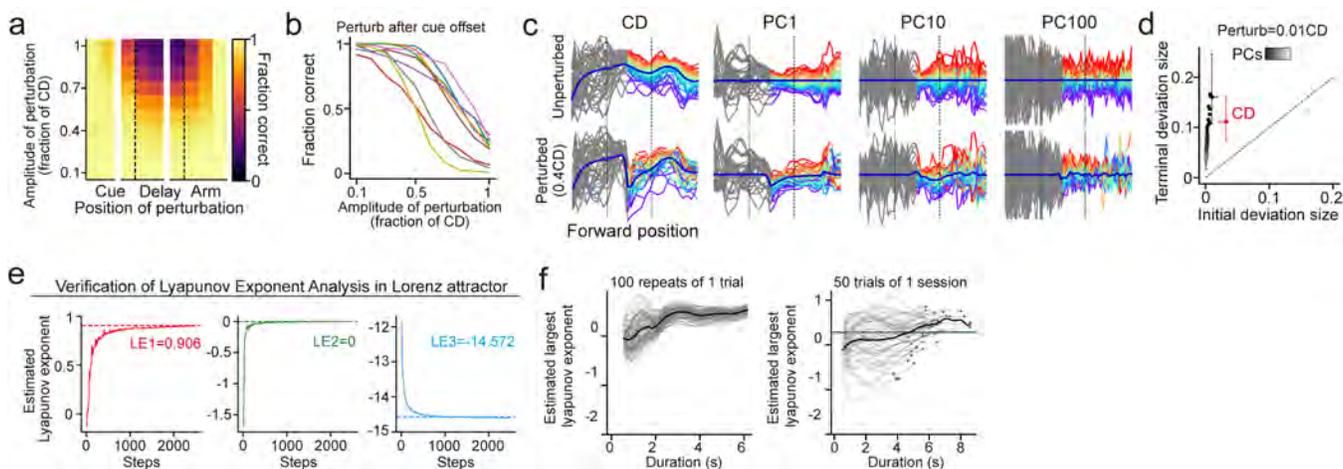
robust in behavioral performance, trial-averaged neural activity and trial-averaged velocity for up to three

699

times the training noise (0.3). Beyond that the agents' performance starts to drop to chance level.

700

701



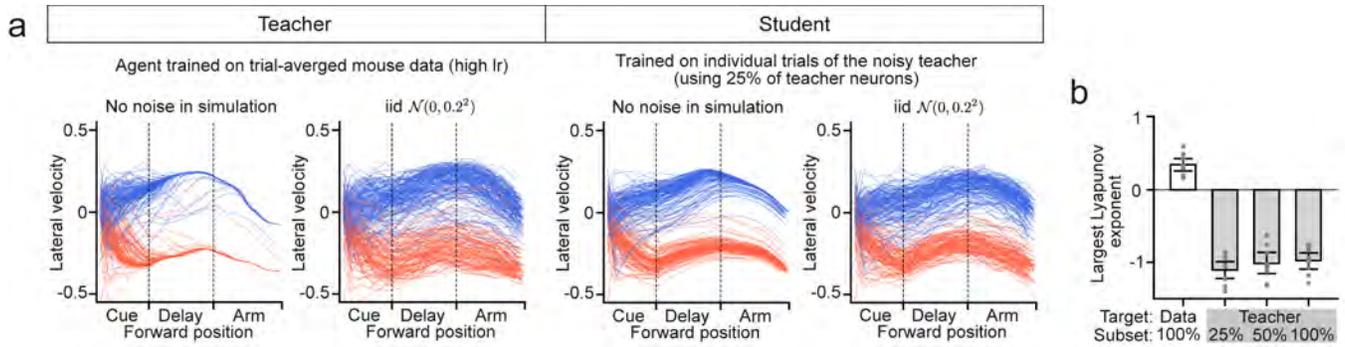
702

703

704 **Extended Data Fig. 2. Additional results of perturbations in data-derived agents and Lyapunov**  
 705 **exponent analysis**

706 **(a)** Effect of perturbations of different amplitudes applied at different positions in the maze on the  
 707 behavioral performance of the data-derived agents. Perturbation is more effective in the delay epoch than  
 708 in the cue epoch, suggesting visual cues constrain neural dynamics into a single attractor, while during  
 709 delay there are two choice-specific attractors across which a strong perturbation can drive trajectories to  
 710 switch. **(b)** Perturbation effect on behavioral performance in individual sessions. **(c)** While the trial  
 711 average recovers after perturbation, individual trials do not. Individual left trials in a session are color-  
 712 coded with their position-wise rank in the dimension evaluated (CD, PC1, PC10, PC100). Following  
 713 perturbation, trajectories of individual trials experience shuffling (as indicated by the mixed colormap)  
 714 while the trial average (thick blue curve) recovers. **(d)** Divergence occurs on all dimensions. Initial  
 715 perturbation is along the CD, and evolution of the deviation on all the PC dimensions are evaluated. **(e)**  
 716 Verification of the numerical method calculating Lyapunov exponents in the known system of Lorenz  
 717 attractor. The dotted line marks the theoretical value of the 3 Lyapunov exponents. The curves show how  
 718 the estimated values converge to the theoretical values during simulation. **(f)** For the data-derived agents,  
 719 the Lyapunov exponents are estimated by running 100 repeats of 50 randomly chosen trials of each  
 720 session. The estimated Lyapunov exponents show trends of convergence over the trial duration. For some  
 721 trials, the estimated Lyapunov exponents might have not fully converged, making our estimation a  
 722 potential lower bound. Left: Thin lines show individual repeats; the thick line shows average across all  
 723 repeats for the given trial. Right: Thin lines show individual trials, thick line shows average across the 50  
 724 trials, green dots show the final estimates for each trial, and green line shows the final estimated largest  
 725 Lyapunov exponent for this session, which takes the average of the estimates for individual trials.

726



727

728

729

**Extended Data Fig. 3. The data-derived agents distinguish unstructured noise from chaos despite partial observations.**

730

731

**(a)** An agent trained on trial-averaged activity and exhibited point attractor dynamics is used as the teacher.

732

The teacher is simulated with injected unstructured noise to generate trial-to-trial variability. A student

733

agent trained on the individual trials of the noisy teacher is able to recover that the teacher has stable

734

dynamics, even when the student is only given a subset of teacher neurons. Figures show velocity outputs

735

of the teacher and student agents, without or with injected random noise. **(b)** Largest Lyapunov exponent

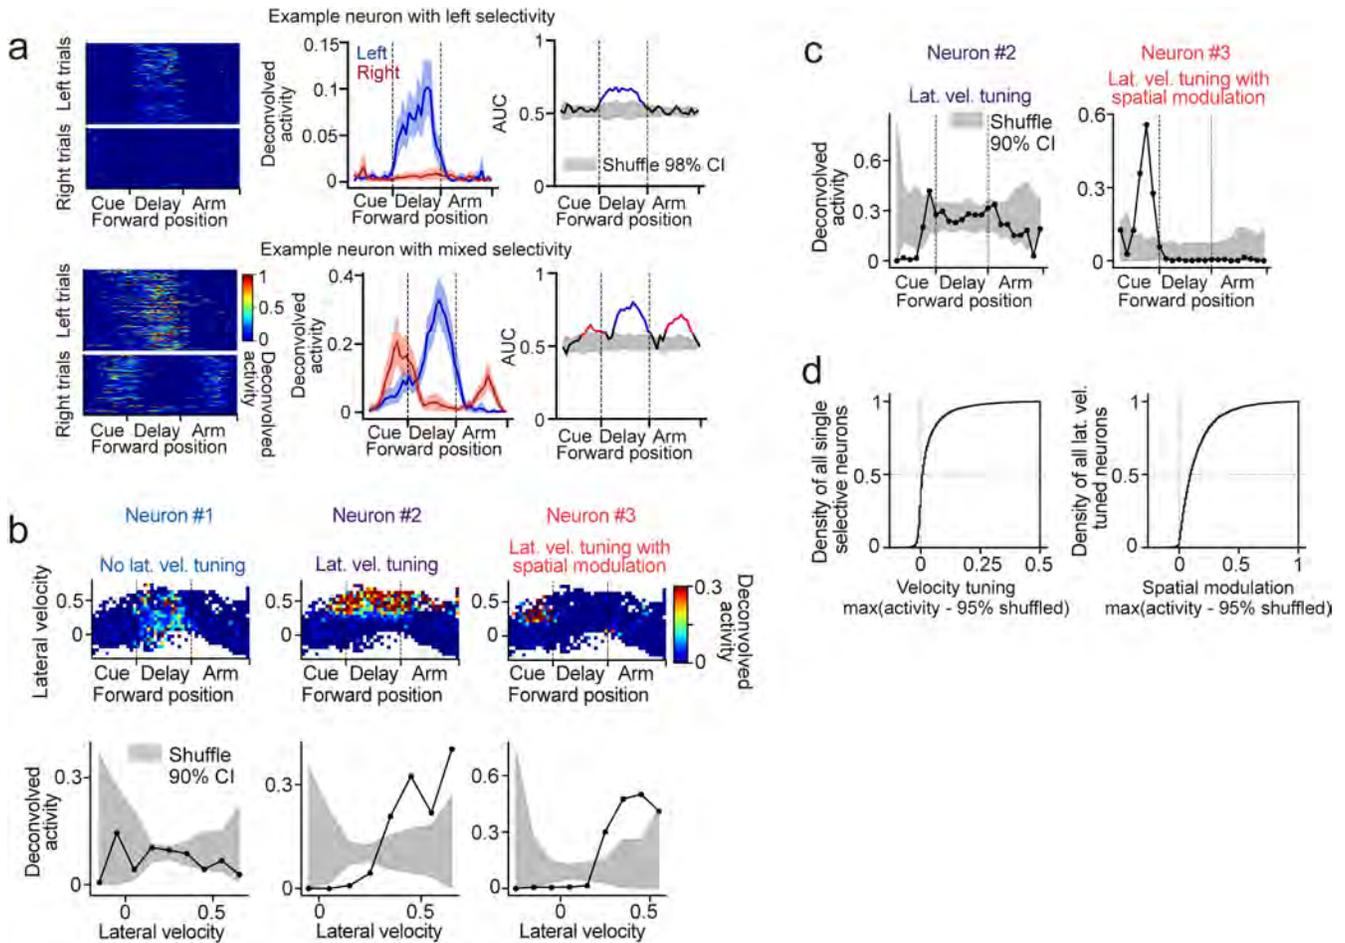
736

of agents fit to the individual trials of different sources (data, or a subset of neurons from the teacher). Dot

737

are individual sessions. Negative value indicates stable dynamics.

738



739

740

741

**Extended Data Fig. 4. Evaluation of single neuron representations.**

742

**(a)** Two example neurons with choice selectivity. Left: Neural activity of individual left and right trials,

743

plotted against the forward position in the maze. Middle: Activity averaged across all left or right trials.

744

Right: Selectivity is determined by area under curve (AUC) of a support vector classifier. Shaded area

745

shows null distribution of the AUC. Colored segments mark the selective fields of the neurons. **(b)** Lateral

746

velocity tuning of neurons. Top: Bin-averaged neural activity in conjunctive bins of forward position and

747

lateral velocity (Same as in Fig. 4b). Bottom: Activity around the peak forward position is plotted as a

748

function of lateral velocity. The shuffles are generated by shuffling the velocity labels of time steps. A

749

neuron is considered to have significant lateral velocity tuning if there is at least one lateral velocity bin

750

where the activity exceeds 95% of the shuffled distribution. **(c)** Spatial modulation of lateral velocity

751

tuning. Activity at time steps when the lateral velocity is within the neuron's selective field of lateral

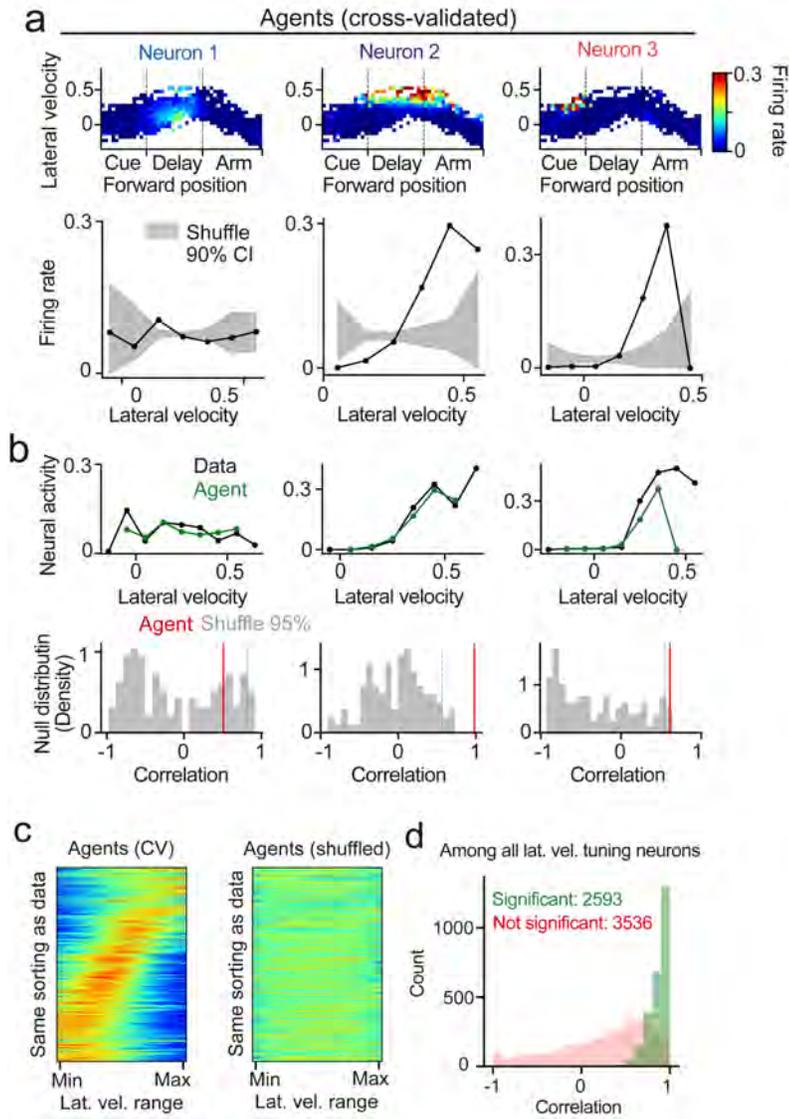
752

velocity is plotted against the forward position. Neuron #3 but not #2 shows large spatial modulation. **(d)**

753

Distribution of lateral velocity tuning and its spatial modulation. See Methods for details.

754



755

756

757

**Extended Data Fig. 5. Data-derived agents recapitulate lateral velocity tuning of single neurons.**

758

(a) Similar to Extended Data Fig. 4b, the same analysis is repeated for the corresponding units in a data-derived agent to quantify their lateral velocity tuning. In the bottom row, the black curve is the tuning curve of the RNN unit in the data-derived agent, and the shade is the 90% CI of the shuffled tuning curves.

759

(b) To evaluate how well the data-derived agents recapitulate lateral velocity tuning, correlation between the tuning curves of the data and the agents is computed. The distribution of shuffled correlation is calculated based on shuffled tuning curves. The fitting is considered significant if the correlation is greater than 95% of the null distribution.

760

(c) Lateral velocity tuning generated by the data-derived agents forms a spectrum over the lateral velocity range. This feature disappears in the shuffled tuning curve, suggesting it is not due to trivial reasons (e.g., correlation between lateral velocity and forward position). Neurons are sorted with the same indexing as in the data (Fig. 3c)

761

(d) Distribution of the data-agent correlation of lateral velocity tuning curves.

762

763

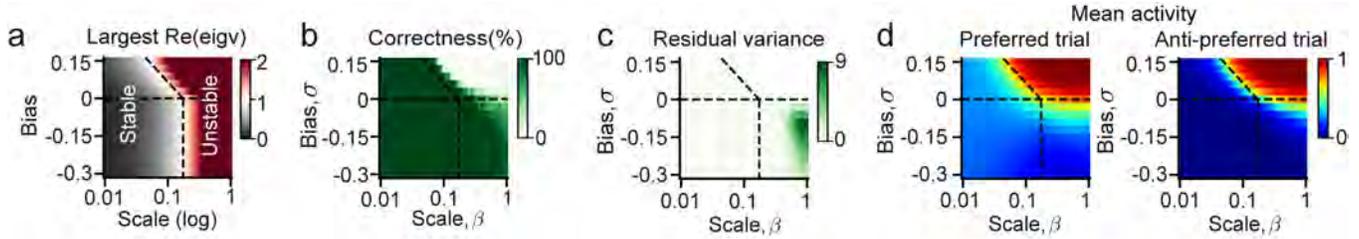
764

765

766

767

768



769

770

771

**Extended Data Fig. 6. Evidence for the empirical phase transition in Fig 4f.**

772

**(a)** Largest eigenvalue of the connection matrix. Values greater than one suggest instability. **(b)**

773

Left/right classification correctness, evaluated by whether the left or right neurons have greater activity.

774

**(c)** Residual variance in neural activity. **(d)** Mean activity of all units in their preferred or anti-preferred

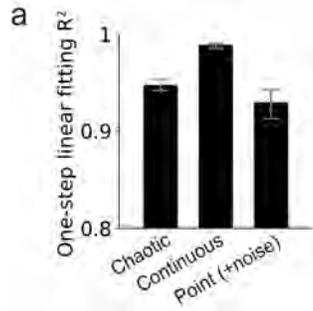
775

trials. Note that (a-d) show results averaged across 10 random weights realization. (b-d) show average

776

across 100 left and 100 right trials with random activity initializations.

777



778

779

780 **Extended Data Fig. 7. Additional information on dynamic mode decomposition (DMD)**

781 **(a)** One-step fitting performance of the linear operator in DMD, in  $R^2$ , for the RNN activity sequence in  
782 the data-derived agents with different dynamics. Averaged across all sessions.

783

784

785

## 786 **Methods**

### 787 **Multi-region calcium imaging in mice**

788 All procedures were approved by the Harvard Medical School Institutional Animal Care and Use Com-  
789 mittee and were performed in compliance with the Guide for the Care and Use of Laboratory Animals.  
790 Detailed experimental procedures was described in reference <sup>43</sup>, where the data was previously reported.  
791 In brief, two female C57BL/6J-Tg (Thy1-GCaMP6s) GP4.3Dkim/J mice (The Jackson Laboratory, stock  
792 024275) were trained to navigate in a virtual reality maze operated in VirRMEn (Virtual Reality Mouse  
793 Engine) <sup>85, 86</sup>. Mice were head-fixed atop an 8-inch Styrofoam ball serving as a spherical treadmill. Move-  
794 ments of the treadmill were measured and converted into pitch, roll, and yaw velocities, which controlled  
795 the mouse's forward and lateral translocation in the virtual environment. The virtual environment was  
796 projected by a micro projector onto a 15-inch diameter half-cylindrical screen. Mice were trained in a  
797 Y-maze. The total length of the Y-maze was 250 cm, with a 150-cm-long, 20-cm-wide Y-stem corridor,  
798 followed by a 100-cm-long, 80-cm-wide Y-arm funnel. The virtual location of mice could not get within  
799 5 cm of the walls, and the virtual view angle of mice was fixed. In all trials, one of two visual cues was  
800 shown on the wall, with a consistent association between each cue and its corresponding rewarded arm in  
801 the Y-maze. Horizontal gratings signaled a left reward, whereas vertical gratings signaled a right reward.  
802 The visual cue was presented only during the first part of the maze, followed by a delay period where a  
803 neutral visual pattern not predictive of the choice was shown on the wall. For sessions used in this paper,  
804 the delay onset ranges from 70 to 100 cm in the Y-stem.

805  
806 Mice were fitted with a chronic cranial window implant, exposing the dorsal surface of both cortical  
807 hemispheres <sup>87</sup> or only the left hemisphere. Data were acquired with a large field of view two-photon  
808 microscope assembled as reported previously <sup>44</sup>, which allows random access imaging across a 5-mm-  
809 diameter field with 1 mm depth. Four regions (primary visual cortex, posterior parietal cortex, retrosplenial  
810 cortex, and secondary motor cortex) in the left cortical hemisphere were imaged and identified by retino-  
811 topic mapping. For each region, imaging was performed in layer 2/3 across two  $600 \mu\text{m} \times 600 \mu\text{m}$  planes  
812 (with a resolution of  $512 \times 512$  pixels) separated by  $50 \mu\text{m}$  in depth, at 5.36 Hz per plane. We performed  
813 motion correction on the original images <sup>75</sup> and used Suite2P to extract the regions of interest <sup>88</sup>. A custom  
814 convolutional neural network implemented in MATLAB was used to classify somatic sources <sup>75</sup>. The  
815 mean fluorescence for each region of interest was calculated and transformed into a normalized change  
816 in fluorescence ( $\Delta F/F$ ) and then deconvolved by the constrained OASIS AR1 method <sup>88</sup>. We used the  
817 deconvolved activity for all subsequent analyses.

### 818 819 **Data-derived agent**

#### 820 ***Model equation:***

821 The agent involves a recurrent neural network (RNN) interacting with a task environment in closed loop.  
822 The RNN consists of  $N$  unit, where  $N$  is chosen to match the number of neurons recorded from the mouse  
823 cortex in a given session. The state of each unit is represented by a variable  $x_i$ , and the dynamics of the

824 RNN is defined as

$$\tau \frac{\mathbf{x}(t)}{dt} = -\mathbf{x}(t) + \mathbf{J}\phi(\mathbf{x}(t)) + \mathbf{W}\mathbf{u}(t) + \boldsymbol{\nu}(t) \quad (1)$$

825 where  $\mathbf{x} \in \mathbb{R}^N$  denotes the states of all RNN units,  $\mathbf{u} \in \mathbb{R}^E$  denotes the environment observations,  
 826  $\boldsymbol{\nu} \in \mathbb{R}^N$  is a per-unit random noise drawn from an iid Gaussian distribution  $\mathcal{N}(0, \sigma^2)$ ,  $\mathbf{J} \in \mathbb{R}^{N \times N}$  is the  
 827 recurrent connection between RNN units, and  $\mathbf{W} \in \mathbb{R}^{N \times E}$  is the weight of a linear input layer from the  
 828 environment to the RNN.  $\phi(\cdot)$  is a nonlinear activation function and  $r = \phi(x)$  gives the firing rate of  
 829 units. For biological plausibility<sup>90</sup>, we designed  $\phi(\cdot)$  to be a modified tanh function controlled by two  
 830 hyper-parameters,  $r_0 = 0.0001$ ,  $r_1 = 4$  (Supplementary Figure 1c). In this way, the firing rate is modulated  
 831 to be between  $-r_0$  and 1, while the resting firing rate (when a unit receives no input from the other units  
 832 nor the environment) is 0. The negative regime was further modulated by a function  $g(x) = x/(1 - 500x)$   
 833 for better numerical stability.  $g(x)$  extends the regime of negative  $x$  that does not suffer from numerical  
 834 saturation. In practice, this helped to make  $\phi(\cdot)$  invertible on its full domain that is used in data fitting.

$$\phi(x) = \begin{cases} r_0 \cdot \tanh(g(x)/r_0/r_1) & \text{for } x \leq 0 \\ \tanh(x/r_1) & \text{for } x > 0 \end{cases} \quad (2)$$

### 835 **Environment:**

836 The environment observations are composed of three components: locomotion velocities  $\mathbf{v}$ , positions  
 837 in the maze  $\mathbf{p}$ , and an abstracted visual cue  $\mathbf{c}$ , i.e.,  $\mathbf{u}(t) = [\mathbf{v}(t) \quad \mathbf{p}(t) \quad \mathbf{c}(t)]^\top$ . **Velocity:**  $\mathbf{v}(t) =$   
 838  $[v_f(t) \quad v_l(t) \quad v_y(t)]$  are the forward(pitch), lateral(roll) and yaw locomotion velocity at time  $t$ , gener-  
 839 ated from the RNN via a linear output layer:  $\mathbf{v}(t) = \mathbf{W}^{\text{out}}\phi(\mathbf{x}(t))$ . **Position:**  $\mathbf{p}(t) = [\mathbf{r}(p_f(t)) \quad p_l(t)]$   
 840 are the basis-expanded forward position and lateral position in the maze. The position is updated in  
 841 the environment by integrating the locomotion velocity generated by the RNN,  $p_f(t) = \int_0^t v_f(\tau)d\tau$   
 842 and  $p_l(t) = \int_0^t v_l(\tau)d\tau$ . In this way, the RNN agent interacts with the environment in a closed loop.  
 843 For biological plausibility, we used a set of basis functions  $p_f(t) \rightarrow \mathbf{r}(p_f(t)) : \mathbb{R} \rightarrow \mathbb{R}^M$  to mimic  
 844 how  $M$  idealized place cells would represent the forward position. The basis functions are designed  
 845 to be cosine bumps with uniform spacing and a width-to-spacing ratio of 2. The cosine bumps are  
 846  $r_i(p_f(t)) = 0.5 \times \cos(2\pi(p_f - \hat{p}_i)/w) + 0.5$  if  $|p_f - \hat{p}_i| \leq w/2$  else 0, where  $\hat{p}_i$  is the preferred position  
 847 of the  $i$ -th "place cell" and  $w$  is the width of the field. We chose  $M = 5$ . **Visual cue:** For simplicity  
 848 we abstracted the visual cue observed from the environment as a binarized variable  $\mathbf{c} \in \{0, 1\}^2$ , where  
 849  $\mathbf{c} = [1 \ 0]$  if in the cue epoch and the cue is vertical bar,  $\mathbf{c} = [0 \ 1]$  if in the cue epoch and the cue is  
 850 horizontal bar, and  $\mathbf{c} = [0 \ 0]$  if not in the cue epoch. In practice, all environment observations are rescaled  
 851 so their amplitudes are approximately 1.

852

### 853 **ARCTIC (Activity ReConstruction in Closed loop) via online imitation learning**

#### 854 **Training procedure and loss:**

855 We trained one agent for each cross validation (CV) of each mouse session. The agent features an

856 RNN with the same number of neurons as in the calcium imaging recording, and the RNN units are  
857 one-to-one matched with the real neurons. The input ( $\mathbf{W}$ ), recurrent ( $\mathbf{J}$ ) and output weights ( $\mathbf{W}^{\text{out}}$ ) of  
858 the agent are all initialized to  $\mathbf{0}$  and are then trained to reproduce single-unit activity and velocity out-  
859 puts on a trial-by-trial basis. More precisely, for any trial  $k$  in this session, suppose it takes  $\hat{T}^k$  time  
860 steps to finish (reaching the end the the maze), we use  $\{\hat{r}_i^k[t = 0, 1, \dots, \hat{T}^k]\}$  to represent the ground  
861 truth de-convolved activity of neuron  $i$  ( $i = 1, \dots, N$ ) and  $\{\hat{v}_a^k[t = 0, 1, \dots, \hat{T}^k]\}$  ( $a = \text{forward, lateral}$   
862 and yaw) for the ground truth velocity in this trial. Then the purpose of training is if we initialize the  
863 agent's unit activity  $r_i[t = 0]$  to be the initial ground truth neural activity,  $r_i[t = 0] \leftarrow \hat{r}_i[t = 0]$ , and  
864 initialize the environment such that  $\mathbf{v}[t = 0]$ ,  $\mathbf{p}[t = 0]$  and  $\mathbf{c}[t = 0]$  all become the initial environment  
865 seen by the mouse, then starting from  $t = 0$  the agent would autonomously generate an entire sequence  
866 of  $\{r_i^k[t = 0, 1, \dots, T^k]\}$  and  $\{v_a^k[t = 0, 1, \dots, T^k]\}$  such that both the difference of the neural trajec-  
867 tories  $\mathbb{D}(\{r_i^k[t = 0, 1, \dots, T^k]\}, \{\hat{r}_i^k[t = 0, 1, \dots, \hat{T}^k]\})$  and the difference of the velocity trajectories  
868  $\mathbb{D}(\{v_a^k[t = 0, 1, \dots, T^k]\}, \{\hat{v}_a^k[t = 0, 1, \dots, \hat{T}^k]\})$  are minimized for all  $i$  and  $a$ . Because the trial duration  
869  $\hat{T}^k$  varies a lot among different trials (Supplementary Figure 1b), it is challenging for the agent to reproduce  
870 the trial duration precisely for all the trials, so that  $T^k$  and  $\hat{T}^k$  are usually not exactly the same (but can  
871 be correlated, see Supplementary Figure 1d). Therefore, it would make less sense if the loss is defined  
872 per time step as  $\sum_{i=1}^N \sum_{t=1}^{\min(T^k, \hat{T}^k)} (r_i^k[t] - \hat{r}_i^k[t])^2$ . Instead, we took the idea from dynamic time warping  
873 (DTW) and aligned the trajectories to the forward progress in the maze, and we compared the aligned  
874 trajectories for the fitting error.

875  
876 We did 5-fold CV. For each CV fold the training set contains 80% of all the correct trials in a ses-  
877 sion. We trained the agent for 30 epochs, where for each epoch the sequence of the training trials presented  
878 is shuffled. For each trial, the agent and environment are initialized to the first data-point of this trial.  
879 As the agent simulates forward, the agent weights are updated for each time step  $t$  (online learning<sup>42</sup>).  
880 At each agent simulation time  $t$ , an instantaneous fitting error is evaluated and the weights are update  
881 accordingly. To evaluate the fitting error at agent time  $t$ , we need to find a target data time  $\hat{t}$  with which  
882 the agent trajectory and the data trajectory are aligned according to the forward progress in the maze. We  
883 define the target data time  $\hat{t}$  as  $\hat{t} = 1 + \arg \min_{\tau} |p_f(t - 1) - \hat{p}_f(\tau)|$ . The fitting error is then computed  
884 as  $\epsilon_i^{\text{neu}}[t] = r_i[t] - \hat{r}_i[\hat{t}]$  (neural error) and  $\epsilon_a^{\text{beh}}[t] = v_a[t] - \hat{v}_a[\hat{t}]$  (behavioral error). Note that the weight  
885 update at agent simulation time  $t$  would be immediately in effect for the simulation from time  $t$  onward, so  
886 that it timely affected the agent's internal activity, behavioral outputs, and resulting environmental inputs  
887 at time  $t + 1$ .

888  
889 **Optimization method:**

890 To optimize both the instantaneous neural error ( $\epsilon_i^{\text{neu}}[t]$ ) and the behavioral error ( $\epsilon_a^{\text{beh}}[t]$ ) at the same time,  
891 we use the recursive least square (RLS) method as in<sup>19, 29</sup>, which is a gradient-free online learning algo-  
892 rithm. Consider all weights that go into the RNN units as  $\mathbf{W}^{\text{neu}} = \begin{bmatrix} \mathbf{J} & \mathbf{W} \end{bmatrix} \in \mathbb{R}^{N \times (N+E)}$ , where  $N$  is the

893 number of RNN units and  $E$  is the number of environment observations, then the update of  $\mathbf{W}^{neu}$  follows:

$$\forall i = [1, \dots, N], j = [1, \dots, N + E], \quad \Delta W_{ij}^{neu}[t] = -\eta^{neu} \cdot \epsilon_i^{neu} \sum_{k=1}^{N+E} P_{jk}[t] z_k[t] \quad (3)$$

894 where  $\eta^{neu}$  is the learning rate.  $\mathbf{z}^\top[t] = [\phi(\mathbf{x}[t])^\top \quad \mathbf{u}[t]^\top]$  is the concatenation of the firing rate of all RNN  
 895 units and the environment observations, which can be regarded as the "presynaptic" activity of the weights  
 896  $\mathbf{W}^{neu}$ .  $\mathbf{P}[t] \in \mathbb{R}^{(N+E) \times (N+E)}$  is the inverse of the online correlation matrix of  $\mathbf{z}[t]$ , that is,  $\mathbf{P}[t] = \mathbf{C}[t]^{-1}$   
 897 and  $\mathbf{C}[t] = \mathbf{C}[t-1] + \mathbf{z}[t]\mathbf{z}[t]^\top$ . In practice, instead of performing the computationally heavy inversion of  
 898 the correlation matrix  $\mathbf{C}$  for every single step, we use the Matrix Inverse Lemma:

$$\mathbf{P}[t] = \mathbf{P}[t-1] - \frac{\mathbf{P}[t-1]\mathbf{z}[t]\mathbf{z}[t]^\top\mathbf{P}[t-1]}{1 + \mathbf{z}[t]^\top\mathbf{P}[t-1]\mathbf{z}[t]} \quad (4)$$

899 This algorithm requires  $\mathbf{P}$  to be initialized properly, and we follow the convention to set  $\mathbf{P}[0]$  as a diagonal  
 900 matrix  $\mathbf{I}/\alpha$ . It turns out that  $\alpha$  imposes an L2 regularization on  $\mathbf{W}^{neu}$  with strength  $\alpha$ . In practice, we  
 901 would like to acquire a solution in which the contribution from RNN units and that from the environ-  
 902 ment are of a similar scale. We did this by applying extra regularization to the environment observations  
 903 through  $\mathbf{P}[0] = \mathbf{A}\mathbf{I}$ , where  $\mathbf{A} = \text{Diag}([\alpha_1, \dots, \alpha_{N+E}])$ ,  $\alpha_1 = \dots = \alpha_N = 1$ ,  $\alpha_{N+1} = \dots = \alpha_{N+E} = 500$ .  
 904

905 Another modification we did to the conventional method is that we intentionally removed self-projections  
 906 of the RNN units. Autapses (self-synapses) are relatively rare *in vivo*, yet data-derived RNN models  
 907 tend to learn an unproportionally strong self-inhibition. For better biological plausibility, we fixed all  
 908 self-projections to zero during training. We found that following the conventional update rule (3) and  
 909 then resetting self-projections to 0 after each weight update significantly compromised the performance.  
 910 Instead, we calculate a separate matrix  $\mathbf{P}_i$  for each unit  $i$  that only tracks the non-self inputs. An efficient  
 911 algorithm to achieve this is presented in the Supplementary Text.

912

913 Finally, the velocity output weights,  $\mathbf{W}^{out} \in \mathbb{R}^{V \times N}$  where  $V$  is the number of velocity output channels, are  
 914 updated using a similar rule

$$\forall i = [1, \dots, V], j = [1, \dots, N], \quad \Delta W_{ij}^{out}[t] = -\eta^{out} \cdot \epsilon_i^{beh} \sum_{l=1}^N \tilde{P}_{jk}[t] \phi(x_k)[t] \quad (5)$$

915 where  $\tilde{\mathbf{P}}$  tracks the inverse correlation matrix of  $\phi(\mathbf{x})[t]$ .

916

## 917 Statistics

918 All error bars in the figures represented 95% confidence interval (CI) of the mean unless otherwise men-  
 919 tioned. The CI was calculated by bootstrapping. For a set of measurements with size  $N$ , 1000 random  
 920 samples of size  $N$  were drawn from the set of measurements with replacement. The mean of each random

921 sample was calculated, forming a distribution of the mean. The lower and upper bounds of the 95% CI of  
922 the mean were then calculated as the 2.5% and 97.5% percentile of this distribution.

923

## 924 **Neural representations**

925 For population representation, choice dimension was computed per forward position bins of 5 cm. In each  
926 position bin, mean population activity of all correct left trials and all correct right trials were calculated  
927 and their difference (left-right) was defined as the choice dimension associated with that position.

928

929 The representations of individual neurons were considered in figure 4 and supplementary figure 4, which  
930 were later used to understand the connectivity motifs. The methods of computing single neuron represen-  
931 tations are detailed below:

932

933 **Choice selectivity:** Neural activity of each trial was binned into 5 cm bins according to the forward  
934 position in the maze. Choice-selectivity only considered correct trials. Significance of choice selectivity  
935 was calculated based on AUC (Area Under the Curve). For each session, as there were usually different  
936 numbers of left and right trials, the trial-type with more trials was down-sampled randomly such that the left  
937 and right trials used for classification were balanced. A Support Vector Classifier was trained per position  
938 bin to predict the choice based on single neuron's activity in that position bin. An ROC (Receiver-operating  
939 characterization) curve was constructed and the area under this curve was used to quantify the classifier's  
940 performance. To compute the null distribution, the choice labels were shuffled among trials. Both the  
941 original and shuffled AUC were calculated for 100 repeats, with random sub-sampling of trials. A neuron  
942 was considered to have significant prediction of choice in a given position bin if the averaged original AUC  
943 was greater than 99% of the shuffled distribution. A neuron was then considered to have choice selectivity  
944 in a position field if (1) the neuron had significant prediction of choice in a consecutive field of at least 5  
945 bins, and (2) the neuron had increased activity in that field, where period of increased activity was defined  
946 as the positions where the trial-averaged activity of a given choice is greater than 0.01. Neurons with  
947 left or right selectivity were defined as having only one position field selective for either the left or right  
948 choice. Neurons with multiple choice-selective position fields, whether those were for the same choice or  
949 difference choices, were classified as "others". We focused on neurons with left or right choice selectivity  
950 in the following analyses.

951

952 **Lateral velocity selectivity:** For simplicity, lateral velocity selectivity was only considered for neurons  
953 with left or right choice selectivity. For a left (right) selective neuron, we considered all time points in  
954 left (right) trials within a 25 cm window centered at its peak forward position. The neuron's activity  
955 was bin-averaged in conjunctive bins of forward position and lateral velocity, with the bin size of forward  
956 position being 5 cm and that of lateral velocity being 0.1. The neuron's lateral velocity tuning curve was  
957 obtained by averaging out the forward position axis. To compute the shuffle distribution, the lateral velocity  
958 was shuffled across all time points we considered, and the tuning curve was computed according to the  
959 shuffled lateral velocity. This was repeated 1000 times to get the distribution of the shuffled tuning curve.

960 A neuron was considered to have significant lateral velocity tuning if there was at least one lateral velocity  
961 bin where the original activity was greater than the 95% of the shuffled curves.

962

963 ***Spatial modulation of lateral velocity tuning:*** Spatial modulation was considered for neurons that have  
964 significant lateral velocity selectivity. We took a similar approach as above. For each neuron with lateral  
965 velocity selectivity, we considered all time points when the lateral velocity was in the neuron's selective  
966 window, defined as lateral velocity bins where the activity was greater than the 95% of the shuffle. Activity  
967 was then binned into conjunctive bins of forward position and lateral velocity, where we used 10 cm forward  
968 position bins here. We shuffled the forward position across time points to get the null distribution. The  
969 modulatory effect of forward position was calculated by averaging-out the lateral velocity axis. A lateral  
970 velocity tuned neuron was considered to be modulated by forward position if there was at least one forward  
971 position bin where the original activity was greater than the 95% of the shuffled curves by at least 0.1.

972

### 973 **Agent evaluation**

974 We claimed that the data-derived agents learned a distribution of neural trajectories from the training set and  
975 generalized such distribution to the test set. In figure 2m-q we evaluated such distribution on the population  
976 level through position-wise residual variance structures. In figure 2l, figure 4 and supplementary figure 5,  
977 we evaluated how well the agents recapitulated single neuron representations, which can be interpreted as  
978 a conditional distribution of neural activity given specific behavioral variables. In both cases, we treated  
979 each mouse session with 5-fold CV, and assembled the test trials from every fold to recover the original  
980 number of trials per session for evaluation.

981

982 For analyses of the residual variance structure, we started by binning neural activity in each trial ac-  
983 cording to the forward position in the maze into 5 cm bins. We then computed the mean population  
984 activity of all correct left or right trials. The mean activity at each position bin was subtracted from each  
985 individual trial, isolating the trial-to-trial neural activity residuals (we only considered correct trials). The  
986 total residual variance took the mean square of the residuals, summed over all neurons and averaged across  
987 position bins. We projected the residuals on choice dimension by taking the inner product, and the residual  
988 variance on choice dimension took the mean square of the projected residuals, averaged across position  
989 bins. Projected variance at chance level was calculated by dividing the total residual variance by the rank  
990 of the residuals, which was the number of correct trials as it was always smaller than the number of neurons  
991 in our data. Residual principal components were then calculated per position bins. The projection of  
992 choice dimension on each residual principal component took their inner product, and we then took the  
993 quadratic mean across position bins. The chance level projection took the square root of 1 over the rank  
994 of the residuals (number of correct trials). Evaluation of the data-derived agent took a similar procedure,  
995 while we only considered the trials where the agent made a correct choice.

996

997 For evaluation of single neuron representations, we focused on the neurons' tuning of choice and lateral  
998 velocity. Choice selective neurons had spatially confined activity bumps, which could be characterized by

999 a neuron’s trial-averaged activity. We sorted the neurons preferring left or right choice according to their  
 1000 peak forward positions, resulting in the characteristic choice-selective sequences. The same sequences  
 1001 were observed in the agents (Fig. 2l). The tuning to lateral velocity closely related to the trial-to-trial  
 1002 variation. We first visualized the lateral velocity tuning spectrum of the population. Each neuron’s lateral  
 1003 velocity tuning curve had the range of the possible lateral velocity visited by the mouse or the agent in  
 1004 the 25 cm window centered at this neuron’s peak forward position. For visualization purpose, the tuning  
 1005 curves were interpolated to all have 20 bins, and were normalized for each neuron. The neurons were sorted  
 1006 according to which one of the 20 bins had the peak activity in their tuning curves. We also compared the  
 1007 lateral velocity tuning curve of single neurons in the data and in the agent, by taking their Pearson corre-  
 1008 lation coefficient. A null distribution was generated by correlating the data tuning curve with the shuffled  
 1009 tuning curve from the agent. A neuron’s lateral velocity tuning was considered to be significantly reca-  
 1010 pitulated by the agent if the true Pearson correlation coefficient was greater than 95% of the null distribution.  
 1011

### 1012 **In-model perturbation**

1013 To perturb the neural activity in the data-derived agents, we added an instantaneous vector deviation to the  
 1014 RNN activity in individual trials at a certain forward position, chosen from [5, 30, 55, 80, 105, 130, 155, 180, 205]  
 1015 (in cm). The perturbation vector was chosen to be the position-specific choice dimension. The direction  
 1016 of perturbation was towards the opposite trial type, such that on a left (right) trial, the perturbation moved  
 1017 the RNN activity closer to that of the right (left) trials. The amplitude of perturbation was defined as  
 1018 fractions of the norm of the difference between mean left and mean right trial activity, chosen from  
 1019 [0.01, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0], such that a perturbation of size 0.5 means that we moved  
 1020 the RNN activity to roughly the separation plane of the mean left and mean right trial activity. The effect  
 1021 of perturbation was evaluated using two metrics quantifying the deviation of perturbed neural trajectory,  
 1022 aligned to forward position. The first metric is norm deviation of trial mean,

$$\text{NDM} = \left( \left\| \frac{\sum_{k \in \text{left trials}} (\mathbf{r}_{p_f}^k - \hat{\mathbf{r}}_{p_f}^k)}{\# \text{ left trials}} \right\|_2^2 + \left\| \frac{\sum_{k \in \text{right trials}} (\mathbf{r}_{p_f}^k - \hat{\mathbf{r}}_{p_f}^k)}{\# \text{ right trials}} \right\|_2^2 \right) / 2 \quad (6)$$

1023 with  $\mathbf{r}_{p_f}^k$  denotes the original agent activity in trial  $k$  at forward position  $p_f$  and  $\hat{\mathbf{r}}_{p_f}^k$  denotes the perturbed  
 1024 agent activity. The second metric is mean norm deviation of individual trials,

$$\text{MND} = \sum_{k=1}^K \left\| \mathbf{r}_{p_f}^k - \hat{\mathbf{r}}_{p_f}^k \right\|_2^2 / K \quad (7)$$

### 1025 **Lyapunov exponents (LEs) and covariant Lyapunov vectors (CLVs)**

1026 In a dynamical system, the exponential growth rates of infinitesimal perturbations in different directions  
 1027 in the state space are described by the Lyapunov exponent (LE) spectrum,  $\lambda_1, \dots, \lambda_n$  (sorted from highest  
 1028 to lowest) <sup>48, 49</sup>. Positive and negative LEs are associated with the unstable and stable manifolds, whose  
 1029 directions are indicated by the covariant Lyapunov vectors (CLVs),  $\gamma_1, \dots, \gamma_n$ .

1030

1031 It is worth pointing out that LEs and CLVs can be viewed as a nonlinear equivalent of the log-eigenvalues  
 1032 ( $\sigma_i$ ) and eigenvectors ( $\mathbf{v}_i$ ) of a linear operator, with the difference that CLVs are state-varying due to the  
 1033 nonlinearity. Let  $\mathbf{x}(t)$  be a temporally evolving variable from some initial state  $\mathbf{x}(0)$ , then:

$$\text{For a nonlinear operator } \mathbf{f}(\cdot): \quad \mathbf{f}_t(\mathbf{x}(0)) = \mathbf{x}(t) \quad \Rightarrow \quad \mathbf{f}_t(\mathbf{x}(0) + \gamma_i(\mathbf{x}(0))) = \mathbf{x}(t) + e^{\lambda_i t} \gamma_i(\mathbf{x}(t)) \quad (8)$$

$$\text{For a linear operator } \mathbf{A}: \quad \mathbf{A}^t \cdot \mathbf{x}(0) = \mathbf{x}(t) \quad \Rightarrow \quad \mathbf{A}^t \cdot (\mathbf{x}(0) + \mathbf{v}_i) = \mathbf{x}(t) + \sigma_i^t \mathbf{v}_i \quad (9)$$

1034 The algorithm to calculate the LEs and CLVs is detailed in <sup>48,49</sup>. We described it briefly here. Notice that  
 1035 any random perturbation ( $\delta\mathbf{x}(0)$ ) would eventually converge to the direction of the fastest growth ( $\gamma_1$ ) with  
 1036 growth rate of the leading LE ( $\lambda_1$ ). Therefore, the leading LE and CLV can be calculated from numerically  
 1037 simulating the evolution of random perturbations:

$$\lambda_1 = \lim_{T \rightarrow \infty} \frac{1}{T} \log \frac{\|\delta\mathbf{x}(T)\|}{\|\delta\mathbf{x}(0)\|} \quad (10)$$

1038 To calculate the rest of the LE spectrum ( $\lambda_2, \dots, \lambda_N$ ), the challenge is that unless the initial perturbation  
 1039 happens to be exactly perpendicular to the direction of the leading CLV, then the evolution of this per-  
 1040 turbation would eventually be dominated by the leading direction. In other words, to reveal the growth  
 1041 rate of the second leading direction, we need to remove the effect of the leading direction by perturbing  
 1042 in a subspace orthogonal to the leading direction. In practice, to numerically calculate the first  $K$  LEs,  
 1043 we simulated the evolution of  $K$  small random perturbations on a chosen trajectory in parallel. During  
 1044 simulation, we periodically re-orthogonalized the  $K$  deviation vectors to prevent their collapse onto the  
 1045 fastest-growing direction. For orthogonalization, we did QR factorization in a way that the  $i$ th deviation  
 1046 vector is projected to the subspace orthogonal to all the  $1, \dots, i-1$  vectors. The growth rate of such  
 1047 projection gave the  $i$ th LE. For each session, we randomly picked  $A = 50$  correctly performed trials and  
 1048 used the agent simulations as the unperturbed trajectories. For each trajectory, we repeated the above  
 1049 process with random initial perturbations for  $B = 100$  times. In figure 3 we averaged across all sessions,  
 1050 trajectories and repeats:

$$\lambda_1 = \frac{1}{A \cdot B \cdot T \Delta t} \sum_{a=1}^A \sum_{b=1}^B \sum_{t=1}^T \log \frac{\|\mathbf{q}_i[t]\|}{\|\mathbf{q}_i[t-1]\|} \quad (11)$$

1051 where  $\mathbf{q}_i$  is the  $i$ th deviation vector after orthogonalization. However, note that  $\mathbf{q}_i$ 's are no longer covariant  
 1052 with the dynamics due to the orthogonalizations. To reveal the true CLVs, we need to recombine  $\{\mathbf{q}_i[t]\}$   
 1053 by:

$$\mathbf{\Gamma}[t] = \mathbf{Q}[t] \mathbf{C}[t] \quad (12)$$

1054 where  $\mathbf{\Gamma}[t]$  contains column vectors of the CLVs ( $\lambda_i[t]$ ) and  $\mathbf{Q}[t]$  contains column vectors of  $\mathbf{q}_i[t]$ . The  
 1055 coefficient  $\mathbf{C}[t]$  is an upper triangular matrix that recombines the orthogonal  $\{\mathbf{q}_1[t], \dots, \mathbf{q}_i[t]\}$  to recover  
 1056 the  $i$ th CLV  $\gamma_i[t]$ . This is to counteract the orthogonalizing process during forward simulation.  $\mathbf{C}[t]$  is

1057 calculated in reverse iterations using  $\mathbf{C}[t] = \mathbf{R}^{-1}[t]\mathbf{C}[t+1]\mathbf{D}[t]$ , where  $\mathbf{R}[t]$  is the upper triangular matrix  
1058 in QR factorization that orthogonalizes  $\mathbf{q}_i[t]$ 's during simulation. To ensure that the CLVs have unit norms,  
1059  $\mathbf{C}[t]$  is normalized column-wise by  $\mathbf{D}[t]$ , a diagonal matrix containing column norms of  $\mathbf{R}^{-1}[t]\mathbf{C}[t+1]$ .

1060

### 1061 **Weight analyses**

1062 For RNN weight  $J_{ij}$ , we call neuron  $i$  its target neuron and neuron  $j$  its source neuron. RNN weights were  
1063 grouped according to the source and target neurons' peak forward positions and peak lateral velocities.  
1064 For all panels in Figure 4 f-i except the right plots in panel 4g and 4i, we pooled weights across all agents  
1065 for individual sessions and all the CV folds. For figure 4f-g, we took a forward position bin of 5 cm, and  
1066 determined the peak position of neurons based on their trial-averaged activity. For figure 4 h-i, we took  
1067 a wider forward position bin of 10 cm and used the lateral velocity bin of 0.1, and determined the peak  
1068 forward position and peak lateral velocity of a neuron as the 2-D indexes of its peak conjunctive activity in  
1069 bins of forward position and lateral velocity.

1070

### 1071 **Adding bias to the weights in the data-derived agents**

1072 The influence of inhibition strength on neural dynamics was examined for the trial period after cue offset,  
1073 in order to isolate the constraining effect of visual cues. Therefore, the neural and behavioral trajectories  
1074 in the cue period were held fixed, and re-training and evaluations were performed following cue offset.  
1075 We started with the input, recurrent, and output weights of the data-derived agent. A bias (chosen from  
1076  $[-0.02, -0.01, 0.005, 0.01]$ ) was added to the recurrent weights from the time step of cue offset, instantly  
1077 changing the neural activity and behavioral outputs. Then, the same procedure of online imitation  
1078 learning was performed. During this re-training phase, the input and biased recurrent weights were held  
1079 fixed and only the velocity output weights were re-trained to minimize the difference between the new  
1080 velocity outputs (forward and lateral) and those of the original agent. Evaluations of neural effects (Fig.  
1081 5j-l) were performed for the arm epoch.

1082

### 1083 **Toy model with fixed random weights**

1084 For the toy model in Figure 5a-f, the within-pool weights were drawn from  $\mathcal{N}(\sigma(\beta + \mu), \sigma^2)$  and the  
1085 across-pool weights were drawn from  $\mathcal{N}(\sigma\beta, \sigma^2)$ . Here, the mean was intentionally scaled by  $\sigma$ , as the  
1086 eigen-spectrum properties of Gaussian matrices usually depend on not the absolute mean, but its propor-  
1087 tion to the standard deviation. A fixed positive  $\mu = 0.3$  was considered to impose the competition motifs.  
1088 To evaluate the dynamical properties of a specific combination of  $\mu, \sigma$  and  $\beta$ , a specific weight matrix  
1089 realization was drawn from the Gaussian distributions, and multiple trial trajectories were simulated based  
1090 on the same weight realization from different random activity initializations. The same activation function,  
1091 time constant and integration time step as in the data-derived agents were used for the toy model. Phase  
1092 transitions (Fig. 5f, Supplementary Fig. 6) combined the results from 10 random weight realizations for  
1093 each parameter combination of  $\sigma$  and  $\beta$ , each simulated for 100 left and 100 right trials, with a constant cue  
1094 signal presented for 400 time steps. For Lyapunov exponent analyses (Fig 5. d,e), numerical simulation  
1095 with a given weight realization and constant cue inputs was performed from a single activity initialization

1096 for 50000 time steps to provide thorough sampling of the state space and complete convergence of the  
 1097 estimated Lyapunov exponents. Residual PCs were computed based on the last 20000 time steps of the  
 1098 simulation. Reported Lyapunov exponent analyses results combined simulations from 50 random weights  
 1099 realizations.

1100

### 1101 **Reward-based learning with the Actor-Critic algorithm**

1102 For reward-based learning of obstacle avoidance behavior, two obstacles were placed at forward position  
 1103 180, lateral span  $[-7, 16]$ , and forward position 230, lateral span $[15, 35]$  in the arm epoch. Note that the  
 1104 arm epoch started at forward position 150 and ended at 235, with allowed range for lateral movement  
 1105 spanning  $[-35, 35]$ , all units in cm. Because the two obstacles overlapped in their lateral spans and the  
 1106 upper one connected with the upper boundary of the arena, such arrangement of obstacles made sure that  
 1107 they could not be avoided by simply adding a bias to the lateral velocity output.

1108

1109 The input weights, recurrent weights, and output weights to forward velocity of the data-derived agent were  
 1110 held fixed while the output weights to lateral velocity were re-trained using the actor-critic method as in  
 1111 reference<sup>58</sup>. Here we described the setup briefly. The actor, simply being the lateral velocity output layer,  
 1112 generated lateral velocity  $a_t$  from RNN activity  $\mathbf{r}_t$ . The actor weights were initialized as that of the original  
 1113 data-derived agent. The critic was a multi-layer perceptron (MLP) with 2 hidden layers (layer width 64  
 1114 and 16), and was trained to predict a value  $Q_t$  based on the state  $\mathbf{s}_t$ , action  $\mathbf{a}_t$  and desired choice  $\mathbf{c}$ . Here  
 1115 the state was defined as the forward and lateral positions in the maze, both expanded by a set of 5 cosine  
 1116 basis functions, and the action was the lateral velocity, expanded by a set of 10 cosine basis functions. The  
 1117 choice was represented as a binary variable,  $[1\ 0]$  or  $[0\ 1]$ . These together constituted the input ( $\in \mathbb{R}^{22}$ ) for  
 1118 the critic network. For training, we considered trials simulated from the initial conditions sampled from  
 1119 the first data-points of mouse trials. Trained together, the actor was optimized to maximize the value of  
 1120 its actions, judged by the critic, while the critic was optimized to minimize the reward prediction error.  
 1121 The actor and critic networks were updated once per trial, as they were trained on a replay buffer which  
 1122 contained the entire input and output sequences for both networks at each time point of the trial, as well as  
 1123 the reward history if any.

$$\delta_t = r_t + \gamma Q(\mathbf{s}_{t+1}, \mathbf{a}_{t+1}, \mathbf{c}) - Q(\mathbf{s}_t, \mathbf{a}_t, \mathbf{c}) \quad (13)$$

$$\Delta \boldsymbol{\theta} = \alpha_{\boldsymbol{\theta}} \cdot \sum_{t=1}^{T-1} \nabla_{\boldsymbol{\theta}} \mathbf{a}_t(\mathbf{r}_t, \boldsymbol{\theta}) \cdot \nabla_{\mathbf{a}} Q(\mathbf{s}_t, \mathbf{a}_t, \mathbf{c}) \quad (14)$$

$$\Delta \mathbf{w} = \alpha_{\mathbf{w}} \cdot \sum_{t=1}^{T-1} \delta_t \nabla_{\mathbf{w}} Q(\mathbf{s}_t, \mathbf{a}_t, \mathbf{c}) \quad (15)$$

1124 where  $\boldsymbol{\theta}$  stands for the parameterization of the actor, and  $\mathbf{w}$  stands for the parameterization of the critic.  
 1125 The reward,  $r_t$ , might occur at 3 positions along the maze: (1) Reward of getting over the first obstacle,  
 1126 evaluated at 181 cm. Successful avoidance led to a transient reward of  $r^1$ . Hitting the obstacle led to

1127 game-over and resulted in a negative reward of  $r^{\text{neg}}$ . (2) Reward of entering the correct arm evaluated  
1128 at 220 cm. Correctness led to a transient reward of  $r^2$ , while incorrectness led to a transient negative  
1129 reward of  $r^{\text{neg}}$ . (3) Reward of getting over the second obstacle and getting to the correct goal location,  
1130 evaluated at the end of maze. If success, agent received a reward of  $r^3$ . If hitting the second obstacle or  
1131 choosing the wrong goal, agent received  $r^{\text{neg}}$ . We used the Adam optimization method implemented in  
1132 PyTorch. 5-fold CV was performed with the same ways of train-test splits as in the data-derived agent.  
1133 Hyper-parameters, including the amplitude of the rewards  $[r^1, r^2, r^3, r^{\text{neg}}]$  and the learning rates of the actor  
1134 and critic networks  $[\alpha_\theta, \alpha_w]$ , were searched based on the training performance in the first session of each  
1135 mouse and then generalized to the other sessions. To facilitate learning, we implemented a curriculum  
1136 learning procedure, that the length of the obstacles were gradually increased to the full length. Agents  
1137 were trained in each curriculum stage for a maximum of 30 epochs and would be promoted early to the  
1138 next stage if training performance reached 85%. We also implemented curriculum rollback, that the agents  
1139 would be reverted to the previous checkpoint if performance stayed low for more than 4 epochs. The best  
1140 checkpoint for each session was selected based on training performance across left and right trials, and we  
1141 reported the cross-validated testing performance for left trials.

1142

#### 1143 **Supervised learning in the obstacle maze**

1144 RNN activity in the original maze was recorded and fixed. Output weights to the lateral velocity were  
1145 retrained to minimize the loss in the obstacle maze. The loss was computed in the following way: Assuming  
1146 the pace of forward progressing was unchanged, use the new lateral velocity to simulate lateral locomotion  
1147 in the maze. The loss of a trial evaluated whether the trajectory arrived at the right Y-arm and whether the  
1148 trajectory went through an obstacle, where the later term took the form of  $\text{softplus}(\min(y^{\text{upper}} - y, y - y^{\text{lower}}))$ .  
1149 Here  $y^{\text{upper}}$  and  $y^{\text{lower}}$  are the upper and lower lateral position boundary of an obstacle, and  $y$  is the lateral  
1150 position when a trajectory arrived at the obstacle. The term is positive when  $y$  is between  $y^{\text{upper}}$  and  $y^{\text{lower}}$ ,  
1151 which means the trajectory collided with the obstacle. The output weights were trained on both left and  
1152 right trials for 10 epochs with Adam optimization in Pytorch. We reported the cross-validated testing  
1153 performance of left trials, averaged across 10 repeats with random seeds.

1154

#### 1155 **Dynamic Mode Decomposition (DMD)**

1156 We performed exact DMD<sup>62</sup> on the RNN activity of the data-derived agents. In brief, DMD computes  
1157 the best-fit one-step linear map of the data. Note that DMD does not assume linear dynamics, but rather  
1158 finds the linear transformation that best captures the dominant, repeatable patterns in nonlinear dynamics.  
1159 We started by arranging the RNN activity  $\{\mathbf{x}_t^m\}$  (time step  $t$  in trial  $m$ ) into data matrices  $\mathbf{X}$  and  $\mathbf{Y}$ :  
1160  $\mathbf{X} = [\mathbf{x}_0^0 \dots \mathbf{x}_{T_0-1}^0 \quad \mathbf{x}_0^1 \dots \mathbf{x}_{T_1-1}^1 \dots \mathbf{x}_0^M \dots \mathbf{x}_{T_M-1}^M]$ ,  $\mathbf{Y} = [\mathbf{x}_1^0 \dots \mathbf{x}_{T_0}^0 \quad \mathbf{x}_1^1 \dots \mathbf{x}_{T_1}^1 \dots \mathbf{x}_1^M \dots \mathbf{x}_{T_M}^M]$ , so that  $\mathbf{X}$   
1161 and  $\mathbf{Y}$  concatenate all trials with a 1-step offset. We define operator  $\mathbf{A}$  to be the least-squares solution to the  
1162 problem  $\mathbf{A}\mathbf{X} = \mathbf{Y}$ . Then the DMD modes and eigenvalues of the system are given by the eigenvectors and  
1163 eigenvalues of  $\mathbf{A}$ . In practice, exact DMD was calculated by the following algorithm from reference<sup>62</sup>, de-  
1164 scribed briefly here: (1) Compute the reduced SVD of  $\mathbf{X}$ :  $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^*$  that preserves 99.9% of the energy.  
1165 (2) Define matrix  $\tilde{\mathbf{A}} = \mathbf{U}^*\mathbf{Y}\mathbf{V}\Sigma^{-1}$ . Compute eigen-decomposition of  $\tilde{\mathbf{A}}$ :  $\tilde{\mathbf{A}}\mathbf{w} = \lambda\mathbf{w}$ . Each nonzero  $\lambda$  is

1166 a DMD eigenvalue. (3) The DMD mode corresponding to  $\lambda$  is computed as  $\phi = \mathbf{Y}\mathbf{V}\Sigma^{-1}\mathbf{w}$ . The DMD  
 1167 eigenvalue,  $\lambda$ , can be written in polar coordinate  $\lambda = \rho e^{it\theta}$ , where the real part  $\log \rho$  gives the exponential  
 1168 growth rate and the imaginary part  $\theta$  gives the oscillation frequency. One-step DMD prediction error was  
 1169 calculated by  $R^2 = 1 - \|\mathbf{Y} - \Phi\Lambda\Phi^\dagger\mathbf{X}\|_F^2 / \|\mathbf{Y}\|_F^2$ , where  $\Phi = [\phi_1 \dots \phi_r]$  and  $\Lambda = \text{Diag}(\lambda_1 \dots \lambda_r)$ .

1170

1171 In theory, by tuning the linear readout weight from the RNN, one could combine the oscillatory com-  
 1172 ponents with specific loading and phase offset for each component, creating a wide range of output  
 1173 functions. Here we sketch out the main idea. Any initial state can be expressed with the basis of DMD  
 1174 modes,  $\mathbf{x}_0 = \sum_{r=1}^R b_r \phi_r$ . The linear evolution of the activity is then given by  $\mathbf{x}_t = \sum_{r=1}^R \lambda_r^t b_r \phi_r$ , and the  
 1175 linear readout with weight  $\mathbf{c}$  is  $z_t = \mathbf{c}^\top \mathbf{x}_t = \sum_{r=1}^R \lambda_r^t b_r \mathbf{c}^\top \phi_r$ . For a pair of complex eigenvalues,  $\lambda = \rho e^{it\theta}$   
 1176 with mode  $\phi$  and amplitude  $b$ , and its conjugate  $\bar{\lambda}$ ,  $\bar{\phi}$ ,  $\bar{b}$ , their contribution to the output is given by

$$z_t = 2 \cdot \text{Re}(\lambda^t b \mathbf{c}^\top \phi) = 2\rho^t |b \mathbf{c}^\top \phi| \cos(t\theta + \beta), \quad \beta = \arg(b \mathbf{c}^\top \phi) \quad (16)$$

1177 Thus, specific choice of  $\mathbf{c}$  determines  $b \mathbf{c}^\top \phi$ , whose magnitude and argument sets the loading and phase  
 1178 offset of the component associated with  $\phi$ . Therefore, having more slowly decaying components (that we  
 1179 define as  $-0.2 \leq \rho \leq 0.2$ ) is beneficial to the system's capacity to generate arbitrary targets.

## 1180 Supplementary Information

### 1181 Hyperparameters

	Data-derived agents	Toy model
Weights trainable	Yes	No
External inputs	Dependent on RNN outputs	Fixed pattern
Number of RNN units	2628±506 (mean±s.d.)	400
Number of input channels	11	1
Number of output channels	3	0
1182 Integration time step ( $dt$ )	0.0093s	0.0093s
Target time step	0.186s	N/A
Time constant ( $\tau$ )	0.1s	0.1s
Learning rate (recurrent weights)	5	N/A
Learning rate (output weights)	1	N/A
Noise during training	0.1	N/A
Additional L2 reg. on input weights	500	N/A

### 1183 Removal of RNN self-projections

1184 For this section, we only consider the recurrent weights. In an RNN with  $N$  units, for each unit, there are  
 1185 only  $N - 1$  pre-synaptic units since we remove its self-projection. So the corresponding inverse correlation  
 1186 matrix for post-synaptic unit  $i$  (without considering itself) should be:

$$\mathbf{P}^i = (\mathbf{C}^i + \mathbf{I}/p_0)^{-1} \in \mathbb{R}^{N \times N}, \quad \text{where } C_{mn}^i = \begin{cases} \langle r_m r_n \rangle, & m, n \neq i \\ 0 & (m \text{ or } n = i) \end{cases} \quad (17)$$

1187

$$J_{ij}[t] = J_{ij}[t-1] - (1 - \delta_{ij})e_i[t] \sum_{k=1}^N P_{jk}^i[t]r_k[t-1] \quad (18)$$

1188 Note that although  $\mathbf{C}^i$  can be constructed from the sub-matrix of  $\mathbf{C}$ ,  $\mathbf{P}^i$  can NOT be constructed from the  
 1189 sub-matrix of  $\mathbf{P}$ .

1190

1191 Method 1: Directly update individual  $\mathbf{P}^i$  's. For notation simplicity we define

$$\mathbf{r}^i = [r_1, \dots, r_{i-1}, 0, r_{i+1}, \dots, r_N] \in \mathbb{R}^N \quad (19)$$

1192 Then the update rule should be

$$\mathbf{P}^i[t] = \mathbf{P}^i[t-1] - \frac{\mathbf{P}^i[t-1]\mathbf{r}^i[t-1]\mathbf{r}^{i,\top}[t-1]\mathbf{P}^i[t-1]}{1 + \mathbf{r}^{i,\top}[t-1]\mathbf{P}^i[t-1]\mathbf{r}^i[t-1]} \quad (20)$$

1193 Limitation: This method requires maintaining all the  $\{\mathbf{P}^i\}_{i=1}^N$  in the memory. When  $N$  is large ( $\sim 3000$  in  
1194 our case), this would require  $\sim 200$  GB memory.

1195

1196 Method 2: Derive  $\{\mathbf{P}^i\}$  from the original  $\mathbf{P}$ . Another idea is to do a bit more computation in exchange  
1197 for less required memory.

1198

1199 For notation simplicity, define  $\mathbf{C} + \mathbf{I}/p_0 := \tilde{\mathbf{C}}$ . Then

$$\mathbf{P} = \tilde{\mathbf{C}}^{-1} \in \mathbb{R}^{N \times N}, \quad \text{where } \tilde{\mathbf{C}}_{mn} = \langle r_m r_n \rangle + \delta_{mn}/p_0 \quad (21)$$

1200

$$\mathbf{P}^i = (\tilde{\mathbf{C}}^i)^{-1} \in \mathbb{R}^{N \times N}, \quad \text{where } \tilde{\mathbf{C}}_{mn}^i = \begin{cases} \langle r_m r_n \rangle + \delta_{mn}/p_0, & m, n \neq i \\ \delta_{mn}/p_0 & (m \text{ or } n = i) \end{cases} \quad (22)$$

1201 Note that  $\tilde{\mathbf{C}}^i$  can be constructed from the sub-matrix of  $\tilde{\mathbf{C}}$ :

$$\tilde{\mathbf{C}}^i = \begin{bmatrix} & & & 0 & & & \\ & \tilde{\mathbf{C}}_{11}^i & & 0 & & \tilde{\mathbf{C}}_{12}^i & \\ & & & \vdots & & & \\ 0 & 0 & \dots & 1/p_0 & \dots & 0 & \\ & \tilde{\mathbf{C}}_{21}^i & & \vdots & & \tilde{\mathbf{C}}_{22}^i & \\ & & & 0 & & & \end{bmatrix}, \quad \tilde{\mathbf{C}} = \begin{bmatrix} & & & \tilde{\mathbf{C}}_{1,i} & & & \\ & \tilde{\mathbf{C}}_{11}^i & & \tilde{\mathbf{C}}_{2,i} & & \tilde{\mathbf{C}}_{12}^i & \\ & & & \vdots & & & \\ \tilde{\mathbf{C}}_{i,1} & \tilde{\mathbf{C}}_{i,2} & \dots & \tilde{\mathbf{C}}_{i,i} & \dots & \tilde{\mathbf{C}}_{i,N} & \\ & \tilde{\mathbf{C}}_{21}^i & & \vdots & & \tilde{\mathbf{C}}_{22}^i & \\ & & & \tilde{\mathbf{C}}_{N,i} & & & \end{bmatrix} \quad (23)$$

1202 Note that  $\tilde{\mathbf{C}}^i$  differs from  $\tilde{\mathbf{C}}$  by a rank-two operation:

$$\tilde{\mathbf{C}}^i = \tilde{\mathbf{C}} - \mathbf{U}^i \mathbf{V}^i, \quad (24)$$

1203 where

$$\mathbf{U}^i = \begin{bmatrix} 0 & \tilde{\mathbf{C}}_{1,i} \\ \vdots & \vdots \\ 0 & \tilde{\mathbf{C}}_{i-1,i} \\ 1 & \tilde{\mathbf{C}}_{i,i} - 1/p_0 \\ 0 & \tilde{\mathbf{C}}_{i+1,i} \\ \vdots & \vdots \\ 0 & \tilde{\mathbf{C}}_{N,i} \end{bmatrix} \in \mathbb{R}^{N \times 2}, \quad \mathbf{V}^i = \begin{bmatrix} \tilde{\mathbf{C}}_{i,1} & \dots & \tilde{\mathbf{C}}_{i,i-1} & 0 & \tilde{\mathbf{C}}_{i,i+1} & \dots & \tilde{\mathbf{C}}_{i,N} \\ 0 & \dots & 0 & 1 & 0 & \dots & 0 \end{bmatrix} \in \mathbb{R}^{2 \times N} \quad (25)$$

1204 Using the Woodbury formula,  $(\tilde{\mathbf{C}}^i)^{-1}$  can be derived from  $\tilde{\mathbf{C}}^{-1}$ :

$$\begin{aligned} \mathbf{P}^i &= (\tilde{\mathbf{C}}^i)^{-1} = (\tilde{\mathbf{C}} - \mathbf{U}^i \mathbf{V}^i)^{-1} = \tilde{\mathbf{C}}^{-1} + \tilde{\mathbf{C}}^{-1} \mathbf{U}^i (\mathbf{I} - \mathbf{V}^i \tilde{\mathbf{C}}^{-1} \mathbf{U}^i)^{-1} \mathbf{V}^i \tilde{\mathbf{C}}^{-1} \\ &= \mathbf{P} + \mathbf{P} \mathbf{U}^i \underbrace{(\mathbf{I} - \mathbf{V}^i \mathbf{P} \mathbf{U}^i)^{-1}}_{\mathbb{R}^{2 \times 2}} \mathbf{V}^i \mathbf{P} \end{aligned} \quad (26)$$

1205 The advantage of this method is that all the  $\{\mathbf{P}^i\}_{i=1}^N$  are temporary variables and do not need to constantly  
1206 occupy the memory.

1207

1208 We can use a compact matrix formula to speed up the computation, using  $\odot$  to denote element-wise  
1209 product (broadcast may apply) and  $\times 1$  in the matrix shape to denote the operation of adding a dimension.

1210 Say there are  $M$  post-synaptic units in total.

$$\hat{\mathbf{P}} \in \mathbb{R}^{M \times N \times N}, \quad \text{where } \hat{\mathbf{P}}_{i,:,:} = \mathbf{P}^i \quad (27)$$

1211

$$\mathbf{U} \in \mathbb{R}^{M \times N \times 2}, \quad \text{where } \mathbf{U}_{i,:,:} = \mathbf{U}^i \quad (28)$$

1212

$$\mathbf{V} \in \mathbb{R}^{M \times 2 \times N}, \quad \text{where } \mathbf{V}_{i,:,:} = \mathbf{V}^i \quad (29)$$

1213

$$\hat{\mathbf{P}}[t] = \underbrace{\mathbf{P}[t]}_{\mathbb{R}^{M \times N \times N}} + \underbrace{\mathbf{P}[t]}_{\mathbb{R}^{N \times N}} \underbrace{\mathbf{U}[t]}_{\mathbb{R}^{N \times N} \mathbb{R}^{M \times N \times 2}} (\underbrace{\mathbf{I} - \mathbf{V}[t] \mathbf{P}[t] \mathbf{U}[t]}_{\mathbb{R}^{M \times 2 \times 2}})^{-1} \underbrace{\mathbf{V}[t]}_{\mathbb{R}^{M \times 2 \times N}} \underbrace{\mathbf{P}[t]}_{\mathbb{R}^{N \times N}} \quad (30)$$

1214

$$\begin{aligned} \underbrace{\mathbf{J}[t]}_{\mathbb{R}^{M \times N}} &= \mathbf{J}[t-1] - \underbrace{(\mathbf{1} - \mathbf{I}_N)_{:M,:}}_{\mathbb{R}^{M \times N}} \odot \underbrace{\mathbf{e}[\mathbf{t}]}_{\mathbb{R}^{M \times 1}} \odot \underbrace{\hat{\mathbf{P}}[t]}_{\mathbb{R}^{M \times N \times N}} \underbrace{\mathbf{r}[t-1]}_{\mathbb{R}^N} \\ &= \mathbf{J}[t-1] - \underbrace{(\mathbf{1} - \mathbf{I}_N)_{:M,:}}_{\mathbb{R}^{M \times N}} \odot \underbrace{\mathbf{e}[\mathbf{t}]}_{\mathbb{R}^{M \times 1}} \odot \left( \underbrace{\mathbf{P}[t] \mathbf{r}[t-1]}_{\mathbb{R}^N} + \underbrace{\mathbf{P}[t] \mathbf{U}[t] (\mathbf{I} - \mathbf{V}[t] \mathbf{P}[t] \mathbf{U}[t])^{-1} \mathbf{V}[t] \mathbf{P}[t]}_{\mathbb{R}^{M \times N \times 2}} \underbrace{\mathbf{r}[t-1]}_{\mathbb{R}^{M \times 2 \times 1}} \right) \end{aligned} \quad (31)$$

1215 As indicated by the under-braces, if one performs the matrix multiplications in a proper order, we can  
1216 avoid directly computing the big  $\hat{\mathbf{P}}$  matrix, and there is no need to save it in the memory even temporarily!

1217 Therefore the memory/space required is not much different from the normal training algorithm.

1218

1219 Furthermore,  $\mathbf{P} \mathbf{U}$  and  $\mathbf{V} \mathbf{P} \mathbf{U}$  can be hand-calculated to speed up the program (as they are the only

1220 two  $\mathcal{O}(MN^2)$  heavy calculations involved):

$$\begin{aligned} (\mathbf{PU}^i)_{m,1} &= P_{m,i} \\ (\mathbf{PU}^i)_{m,2} &= \sum_{k=1}^N P_{m,k} (\tilde{C}_{k,i} - \delta_{k,i}/p_0) = \delta_{m,i} - P_{m,i}/p_0 \end{aligned} \quad (32)$$

1221

$$\begin{aligned} (\mathbf{V}^i \mathbf{PU}^i)_{1,1} &= \sum_{m \neq i}^N \tilde{C}_{i,m} P_{m,i} = 1 - \tilde{C}_{i,i} P_{i,i} \\ (\mathbf{V}^i \mathbf{PU}^i)_{1,2} &= \sum_{m \neq i}^N \tilde{C}_{i,m} (\delta_{m,i} - P_{m,i}/p_0) = \tilde{C}_{i,i} P_{i,i}/p_0 - 1/p_0 \\ (\mathbf{V}^i \mathbf{PU}^i)_{2,1} &= P_{i,i} \\ (\mathbf{V}^i \mathbf{PU}^i)_{2,2} &= 1 - P_{i,i}/p_0 \end{aligned} \quad (33)$$

## References

1. Morcos, A. S. & Harvey, C. D. History-dependent variability in population dynamics during evidence accumulation in cortex. *Nat. Neurosci.* **19**, 1672–1681 (2016).
2. Tseng, S.-Y., Chettih, S. N., Arlt, C., Barroso-Luque, R. & Harvey, C. D. Shared and specialized coding across posterior cortical areas for dynamic navigation decisions. *Neuron* **110**, 2484–2502.e16 (2022).
3. Marcos, E. *et al.* Neural variability in premotor cortex is modulated by trial history and predicts behavioral performance. *Neuron* **78**, 249–255 (2013).
4. Valente, M. *et al.* Correlations enhance the behavioral readout of neural population activity in association cortex. *Nat. Neurosci.* **24**, 975–986 (2021).
5. Panzeri, S., Moroni, M., Safaai, H. & Harvey, C. D. The structures and functions of correlations in neural population codes. *Nat. Rev. Neurosci.* **23**, 551–567 (2022).
6. Wu, H. G., Miyamoto, Y. R., Castro, L. N. G., Ölveczky, B. P. & Smith, M. A. Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nat. Neurosci.* **17**, 312–321 (2014).
7. Shamash, P., Lee, S., Saxe, A. M. & Branco, T. Mice identify subgoal locations through an action-driven mapping process. *Neuron* **111**, 1966–1978.e8 (2023).
8. Wang, X.-J. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* **36**, 955–968 (2002).
9. Wong, K.-F. & Wang, X.-J. A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* **26**, 1314–1328 (2006).
10. Hopfield, J. J. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U. S. A.* **79**, 2554–2558 (1982).
11. Inagaki, H. K., Fontolan, L., Romani, S. & Svoboda, K. Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature* **566**, 212–217 (2019).
12. Khona, M. & Fiete, I. R. Attractor and integrator networks in the brain. *Nat. Rev. Neurosci.* **23**, 744–766 (2022).
13. Kleinfeld, D. & Sompolinsky, H. Associative neural network model for the generation of temporal patterns. Theory and application to central pattern generators. *Biophys. J.* **54**, 1039–1051 (1988).
14. Chandra, S., Sharma, S., Chaudhuri, R. & Fiete, I. Episodic and associative memory from spatial scaffolds in the hippocampus. *Nature* **638**, 739–751 (2025).
15. Laje, R. & Buonomano, D. V. Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat. Neurosci.* **16**, 925–933 (2013).
16. Vyas, S., Golub, M. D., Sussillo, D. & Shenoy, K. V. Computation through neural population

- dynamics. *Annu. Rev. Neurosci.* **43**, 249–275 (2020).
17. Remington, E. D., Narain, D., Hosseini, E. A. & Jazayeri, M. Flexible sensorimotor computations through rapid reconfiguration of cortical dynamics. *Neuron* **98**, 1005–1019.e5 (2018).
  18. Ikegaya, Y. *et al.* Synfire chains and cortical songs: temporal modules of cortical activity. *Science* **304**, 559–564 (2004).
  19. Rajan, K., Harvey, C. D. & Tank, D. W. Recurrent network models of sequence generation and memory. *Neuron* **90**, 128–142 (2016).
  20. Fiete, I. R., Senn, W., Wang, C. Z. H. & Hahnloser, R. H. R. Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity. *Neuron* **65**, 563–576 (2010).
  21. Fuster, J. M. Upper processing stages of the perception-action cycle. *Trends Cogn. Sci.* **8**, 143–145 (2004).
  22. Cisek, P. & Kalaska, J. F. Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* **33**, 269–298 (2010).
  23. Schroeder, C. E., Wilson, D. A., Radman, T., Scharfman, H. & Lakatos, P. Dynamics of Active Sensing and perceptual selection. *Curr. Opin. Neurobiol.* **20**, 172–176 (2010).
  24. Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A. & Poeppel, D. Neuroscience needs behavior: Correcting a reductionist bias. *Neuron* **93**, 480–490 (2017).
  25. Pandarinath, C. *et al.* Inferring single-trial neural population dynamics using sequential auto-encoders. *Nat. Methods* **15**, 805–815 (2018).
  26. Linderman, S., Nichols, A., Blei, D., Zimmer, M. & Paninski, L. Hierarchical recurrent state space models reveal discrete and continuous dynamics of neural activity in *C. elegans*. *bioRxiv* 621540 (2019) doi:10.1101/621540.
  27. Tolman, E. C. Cognitive maps in rats and men. *Psychol. Rev.* **55**, 189–208 (1948).
  28. Pfeiffer, B. E. & Foster, D. J. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* **497**, 74–79 (2013).
  29. Sussillo, D. & Abbott, L. F. Generating coherent patterns of activity from chaotic neural networks. *Neuron* **63**, 544–557 (2009).
  30. Sompolinsky, H., Crisanti, A. & Sommers, H. J. Chaos in random neural networks. *Phys. Rev. Lett.* **61**, 259–262 (1988).
  31. Perich, M. G. *et al.* Inferring brain-wide interactions using data-constrained recurrent neural network models. *bioRxiv* (2020) doi:10.1101/2020.12.18.423348.
  32. Yang, G. R., Joglekar, M. R., Song, H. F., Newsome, W. T. & Wang, X.-J. Task representations in neural networks trained to perform many cognitive tasks. *Nat. Neurosci.* **22**, 297–306 (2019).

33. Kim, T. D. *et al.* Flow-field inference from neural data using deep recurrent networks. in *Proceedings of the 42nd International Conference on Machine Learning* (2025).
34. Schimel, M., Kao, T.-C., Jensen, K. T. & Hennequin, G. iLQR-VAE : control-based learning of input-driven dynamics with applications to neural data. in *International Conference on Learning Representations* (2022).
35. Simmons-Edler, R. *et al.* Deep RL needs deep behavior analysis: Exploring implicit planning by model-free agents in open-ended environments. *arXiv [cs.AI]* (2025) doi:10.48550/arXiv.2506.06981.
36. Botvinick, M., Wang, J. X., Dabney, W., Miller, K. J. & Kurth-Nelson, Z. Deep reinforcement learning and its neuroscientific implications. *Neuron* **107**, 603–616 (2020).
37. Jensen, K. T. An introduction to reinforcement learning for neuroscience. *Neurons, Behavior, Data analysis, and Theory* (2024) doi:10.51628/001c.127771.
38. Merel, J. *et al.* Deep neuroethology of a virtual rodent. *arXiv [q-bio.NC]* (2019).
39. Aldarondo, D. *et al.* A virtual rodent predicts the structure of neural activity across behaviours. *Nature* **632**, 594–602 (2024).
40. Cross, L., Cockburn, J., Yue, Y. & O’Doherty, J. P. Using deep reinforcement learning to reveal how the brain encodes abstract state-space representations in high-dimensional environments. *Neuron* **109**, 724–738.e7 (2021).
41. Hennig, J. A. *et al.* Emergence of belief-like representations through reinforcement learning. *PLoS Comput. Biol.* **19**, e1011067 (2023).
42. Ross, S., Gordon, G. J. & Bagnell, J. A. A reduction of imitation learning and structured prediction to no-regret online learning. *arXiv [cs.LG]* 627–635 (2010).
43. Arlt, C. *et al.* Cognitive experience alters cortical involvement in goal-directed navigation. *Elife* **11**, (2022).
44. Sofroniew, N. J., Flickinger, D., King, J. & Svoboda, K. A large field of view two-photon mesoscope with subcellular resolution for in vivo imaging. *Elife* **5**, (2016).
45. Harvey, C. D., Coen, P. & Tank, D. W. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484**, 62–68 (2012).
46. Averbeck, B. B., Latham, P. E. & Pouget, A. Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* **7**, 358–366 (2006).
47. van Vreeswijk, C. & Sompolinsky, H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **274**, 1724–1726 (1996).
48. Ginelli, F., Chaté, H., Livi, R. & Politi, A. Covariant Lyapunov vectors. *J. Phys. A Math. Theor.* **46**, 254005 (2013).

49. Kuptsov, P. V. & Parlitz, U. Theory and computation of covariant lyapunov vectors. *J. Nonlinear Sci.* **22**, 727–762 (2012).
50. Qian, W., Zavatone-Veth, J. A., Ruben, B. S. & Pehlevan, C. Partial observation can induce mechanistic mismatches in data-constrained models of neural dynamics. in *Advances in Neural Information Processing Systems* (2024). doi:10.1101/2024.05.24.595741.
51. Machens, C. K., Romo, R. & Brody, C. D. Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science* **307**, 1121–1124 (2005).
52. Kuan, A. T. *et al.* Synaptic wiring motifs in posterior parietal cortex support decision-making. *Nature* **627**, 367–373 (2024).
53. Roach, J. P., Churchland, A. K. & Engel, T. A. Choice selective inhibition drives stability and competition in decision circuits. *Nat Commun* **14**, 147 (2023).
54. Goldman, M. S. Memory without feedback in a neural network. *Neuron* **93**, 715 (2009).
55. Ozeki, H., Finn, I. M., Schaffer, E. S., Miller, K. D. & Ferster, D. Inhibitory stabilization of the cortical network underlies visual surround suppression. *Neuron* **62**, 578–592 (2009).
56. Sadeh, S. & Clopath, C. Inhibitory stabilization and cortical computation. *Nat. Rev. Neurosci.* **22**, 21–37 (2021).
57. Rubin, D. B., Van Hooser, S. D. & Miller, K. D. The stabilized supralinear network: a unifying circuit motif underlying multi-input integration in sensory cortex. *Neuron* **85**, 402–417 (2015).
58. Silver, D. *et al.* Deterministic policy gradient algorithms. in *Proceedings of the 31st International Conference on International Conference on Machine Learning* (2014).
59. Maass, W., Natschläger, T. & Markram, H. Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.* **14**, 2531–2560 (2002).
60. Lukoševičius, M. & Jaeger, H. Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.* **3**, 127–149 (2009).
61. Legenstein, R. & Maass, W. Edge of chaos and prediction of computational performance for neural circuit models. *Neural Netw.* **20**, 323–334 (2007).
62. H. Tu, J. *et al.* On dynamic mode decomposition: Theory and applications. *J. Comput. Dyn.* **1**, 391–421 (2014).
63. Lorenz, E. N. Deterministic nonperiodic flow. *J. Atmos. Sci.* **20**, 130–141 (1963).
64. Moreno-Bote, R. *et al.* Information-limiting correlations. *Nat. Neurosci.* **17**, 1410–1417 (2014).
65. Quiñ Quiroga, R. & Panzeri, S. Extracting information from neuronal populations: information theory and decoding approaches. *Nat. Rev. Neurosci.* **10**, 173–185 (2009).
66. Romyantsev, O. I. *et al.* Fundamental bounds on the fidelity of sensory cortical coding. *Nature* **580**, 100–105 (2020).

67. Stringer, C. *et al.* Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, 255 (2019).
68. Niell, C. M. & Stryker, M. P. Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* **65**, 472–479 (2010).
69. Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* **17**, 440–448 (2014).
70. Stavisky, S. D., Kao, J. C., Ryu, S. I. & Shenoy, K. V. Motor cortical visuomotor feedback activity is initially isolated from downstream targets in output-null neural state space dimensions. *Neuron* **95**, 195–208.e9 (2017).
71. Seo, H., Barraclough, D. J. & Lee, D. Dynamic signals related to choices and outcomes in the dorsolateral prefrontal cortex. *Cereb. Cortex* **17 Suppl 1**, i110–7 (2007).
72. Hattori, R., Danskin, B., Babic, Z., Mlynaryk, N. & Komiyama, T. Area-specificity and plasticity of history-dependent value coding during learning. *Cell* **177**, 1858–1872.e15 (2019).
73. Akrami, A., Kopec, C. D., Diamond, M. E. & Brody, C. D. Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature* **554**, 368–372 (2018).
74. Jaeger, H. & Haas, H. Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science* **304**, 78–80 (2004).
75. Chettih, S. N. & Harvey, C. D. Single-neuron perturbations reveal feature-specific competition in V1. *Nature* **567**, 334–340 (2019).
76. Ogando, M. B. *et al.* Feature-specific inhibitory connectivity augments the accuracy of cortical representations. *bioRxiv* (2025) doi:10.1101/2025.08.02.668307.
77. Sanzeni, A. *et al.* Inhibition stabilization is a widespread property of cortical networks. *Elife* **9**, (2020).
78. Faure, P. & Korn, H. Is there chaos in the brain? I. Concepts of nonlinear dynamics and methods of investigation. *C. R. Acad. Sci. III* **324**, 773–793 (2001).
79. Korn, H. & Faure, P. Is there chaos in the brain? II. Experimental evidence and related models. *C. R. Biol.* **326**, 787–840 (2003).
80. O’Byrne, J. & Jerbi, K. How critical is brain criticality? *Trends Neurosci.* **45**, 820–837 (2022).
81. Vignesh, D., He, S. & Banerjee, S. A review on the complexities of brain activity: insights from nonlinear dynamics in neuroscience. *Nonlinear Dyn.* (2024) doi:10.1007/s11071-024-10558-2.
82. Eckmann, J.-P. & Ruelle, D. Fundamental limitations for estimating dimensions and Lyapunov exponents in dynamical systems. *Physica D* **56**, 185–187 (1992).
83. Koay, S. A., Charles, A. S., Thiberge, S. Y., Brody, C. D. & Tank, D. W. Sequential and efficient neural-population coding of complex task information. *Neuron* **110**, 328–349.e11 (2022).

84. Pinto, L. *et al.* Task-dependent changes in the large-scale dynamics and necessity of cortical regions. *Neuron* **104**, 810–824.e9 (2019).
85. Harvey, C. D., Collman, F., Dombek, D. A. & Tank, D. W. Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature* **461**, 941–946 (2009).
86. Aronov, D. & Tank, D. W. Engagement of neural circuits underlying 2D spatial navigation in a rodent virtual reality system. *Neuron* **84**, 442–456 (2014).
87. Kim, T. H. *et al.* Long-term optical access to an estimated one million neurons in the live mouse cortex. *Cell Rep.* **17**, 3385–3394 (2016).
88. Pachitariu, M. *et al.* Suite2p: beyond 10,000 neurons with standard two-photon microscopy. *bioRxiv* (2016) doi:10.1101/061507.
89. Friedrich, J., Zhou, P. & Paninski, L. Fast online deconvolution of calcium imaging data. *PLoS Comput. Biol.* **13**, e1005423 (2017).
90. Rajan, K., Abbott, L. F. & Sompolinsky, H. Stimulus-dependent suppression of chaos in recurrent neural networks. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **82**, 011903 (2010).