Neuron

Article

A multi-region recurrent circuit for evidence accumulation in rats

Highlights

- During decision-making, the prefrontal cortex was thought to categorize striatal evidence
- Instead, the cortex and striatum reciprocally communicate evidence-related signals
- Silencing cortex→striatum inputs impairs decisions, challenging the feedforward view
- Modeling shows that recurrence in striato-cortical circuit reconciles these findings

Authors

Diksha Gupta, Charles D. Kopec, Adrian G. Bondy, Thomas Z. Luo, Verity Elliott, Carlos D. Brody

Correspondence

diksha.gupta@ucl.ac.uk

In brief

Gupta et al. demonstrate that the rat prefrontal cortex and striatum jointly accumulate sensory evidence through recurrent, bidirectional communication rather than a feedforward hierarchy. Combining multi-region recordings, projection-specific perturbations, and biologically constrained recurrent neural network modeling, the study reveals distributed cortico-striatal mechanisms for decision-making that recover in the face of perturbations.

Gupta et al., 2026, Neuron *114*, 521–535
February 4, 2026
https://doi.org/10.1016/j.neuron.2025.12.029

CellPress

Neuron

CellPress
OPEN ACCESS

Article

A multi-region recurrent circuit for evidence accumulation in rats

Diksha Gupta,[1,3,6,*] Charles D. Kopec,[1] Adrian G. Bondy,[1] Thomas Z. Luo,[1,4] Verity Elliott,[1] and Carlos D. Brody[1,2,5]
[1]Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA
[2]Howard Hughes Medical Institute, Princeton University, Princeton, NJ, USA
[3]Present address: Sainsbury Wellcome Centre, University College London, London, UK
[4]Present address: School of Biological Sciences, University of Utah, Salt Lake City, UT, USA
[5]Senior author
[6]Lead contact
*Correspondence: diksha.gupta@ucl.ac.uk
https://doi.org/10.1016/j.neuron.2025.12.029

SUMMARY

Decision-making based on noisy evidence requires accumulating evidence and categorizing it to form a choice. Here, we evaluate a proposed feedforward and modular mapping of this process in rats: evidence accumulated in the anterior dorsal striatum (ADS) is categorized in the prefrontal cortex (frontal orienting fields or FOF). Contrary to this, we show that both regions appear to be indistinguishable in their encoding/decoding of accumulator value and communicate this information bidirectionally. Consistent with a role for the FOF in accumulation, silencing FOF to ADS projections impacted behavior throughout the accumulation period, while nonselective FOF silencing did not. We synthesize these findings into a multi-region recurrent neural network trained with a novel approach. *In silico* experiments reveal that multiple scales of recurrence in the cortico-striatal circuit rescue computation upon nonselective FOF perturbations. These results suggest that ADS and FOF accumulate evidence in a recurrent and distributed manner, yielding redundant representations and robustness to certain perturbations.

INTRODUCTION

Decision-making based on noisy perceptual evidence is a core cognitive process that is often conceptualized as a sequence of two sub-computations: gradual accumulation of evidence, followed by thresholding to commit to a categorical choice. Both theoretical[1–4] and empirical[5–8] studies of the neural implementation of this process have proposed candidate circuits with feedforward information flow by mapping the two sub-computations onto distinct brain areas. Such a modular implementation offers several adaptive advantages;[9] for instance, it allows for easy dynamic modulation of decision thresholds based on behavioral demands while leaving the accumulation process unchanged.[3,10–13]

In rats, the anterior dorsal striatum (ADS) and the frontal orienting fields (FOF) have been mapped onto these two theoretically defined sub-computations. ADS is known to represent the graded accumulated evidence and is causally required throughout the accumulation process,[14] whereas FOF is known to represent the categorical choice and is causally necessary only at the very end of the accumulation period when the graded evidence needs to be thresholded.[15–18] These observations provide an intriguing, yet untested, neural implementation of the decision process, in which ADS and FOF form a feedforward functional hierarchy—with evidence accumulated in ADS being thresholded in FOF.

While appealing, this proposal is surprising in several ways. First, the fact that FOF projects monosynaptically to ADS, but ADS does not project directly back to FOF, implies that this proposed implementation must map onto a multisynaptic pathway in which signals from ADS must be relayed through other basal ganglia nuclei, the thalamus, and/or the superior colliculus (SC) before they reach FOF. Furthermore, FOF is one of the main cortical inputs to ADS, but this salient FOF → ADS projection plays no role in the proposed implementation. Anatomically connected brain areas can selectively communicate information to their downstream targets,[19,20] so it is certainly possible that the information flow during evidence accumulation renders the monosynaptic projection from FOF to ADS non-functional and only involves the multisynaptic pathway from ADS to FOF. However, the substantial heterogeneity in neural encoding within FOF, with neurons representing both graded evidence and categorical choice,[16] along with FOF's anatomical placement, strongly suggests that FOF might have a role in gradual accumulation as well.

Second, some recent studies challenge the interpretation of transient unilateral optogenetic perturbations of FOF,[16] which had previously provided compelling evidence for confining FOF's role to the thresholding epoch of decision-making. It has been observed that, unlike its response to bilateral perturbations, the cortex can be robust to transient unilateral perturbations—with the network recovering through information from

the other hemisphere.[21] Since the network has enough time to recover early in the accumulation period but not toward the end, this interhemispheric compensation might underlie the lack of impairment observed during FOF's inactivations early in the decision process, obscuring its role. Furthermore, whole-region non-specific perturbations have been shown to conceal effects that can be revealed by more specific perturbations that target a particular projection.[22–24] This suggests that the previous unilateral inactivation studies may have missed aspects of FOF's involvement in the decision process by not using projection-specific and/or bilateral manipulations.

Therefore, it is crucial to directly evaluate the hypothesized feedforward and modular mapping of the decision process onto ADS and FOF. Testing this would require recording activity from both areas simultaneously[25,26] to determine if the distinct representations as well as the transformations and latencies expected from feedforward communication, are present across the two areas. Such a comparison would also control for slight differences in cognitive strategies or training histories between subjects that can vary neural representations and dynamics.[27–31] Additionally, confirming that the feedback projection under this hypothesis, i.e., the FOF → ADS projection, plays no causal role in evidence accumulation would add substantial support in favor of this hypothesis. Similarly, demonstrating undisrupted evidence accumulation during bilateral manipulations of FOF activity would address the aforementioned confounds and further validate that FOF's role is confined to the thresholding subcomputation.

Here, we follow this logic and apply simultaneous multi-region population recordings, bilateral, and projection-specific perturbations in rats performing an evidence accumulation task. The resulting data provide evidence against the modular, feedforward hypothesis in which evidence accumulated in ADS is thresholded in FOF. We find that ADS does not have privileged information about evidence. Rather, both regions carry similar amounts of information that evolve at comparable timescales. Further, simultaneous population recordings let us define the feedforward (ADS → FOF) and feedback (FOF → ADS) communication subspaces (CSs) under this hypothesis. We find that evidence can be substantially decoded from both these subspaces, in contradiction to the hypothesized information flow. Moreover, with projection-specific inactivations, we show that this ''feedback'' communication is necessary during the accumulation process for unimpaired decision-making.

Therefore, we revise the feedforward hypothesis and advocate for an alternate distributed implementation of the evidence accumulation process in the cortico-striatal circuitry. We instantiate this proposal using a multi-region recurrent neural network (RNN) model with biological constraints on connectivity. The model successfully reproduces empirical representational redundancies and effects of held-out perturbations when trained with a novel approach of capturing patterns of animal behavior on optogenetic perturbation trials alongside control trials. Finally, leveraging the RNN's observability and controllability, we predict that recovery dynamics are at play during non-selective silencing of FOF neurons and perform a number of *in silico* experiments to understand the contributions of different projections to this phenomenon. We show that recovery is supported by multiple scales of recurrence—including self-recurrence, inter-hemispheric recurrence, and striato-cortical recurrence. Taken together, our approach of using projection-specific perturbations in conjunction with multi-region RNN modeling offers a useful method to decipher the contributions of distributed and recurrently connected networks of brain regions to cognitive processes.

RESULTS

To test the hypothesized ADS-FOF transformation on single trials (Figures 1B and 1C), we recorded simultaneous population activity from both regions in rats performing a decision-making task[32] (Figure 1A). In this auditory evidence accumulation task, rats hear two streams of randomly timed clicks from left and right speakers over hundreds of milliseconds. The rat must accumulate evidence by tracking which side has more clicks and report their decision by moving to that side's port after the stimulus ends (Figure S1A). Correct choices (choosing the side with more clicks) are rewarded with water. Task difficulty varies with the difference in click rates between sides, requiring integration of evidence over time for good performance.

Using Neuropixels probes,[35,36] we recorded from 764 FOF and 1,528 ADS neurons across 12 sessions (n = 5 rats; Figures 1D and 1E). Both FOF and ADS populations exhibited diverse temporal response profiles during task performance, with neurons showing increasing, decreasing, or stable firing rate patterns across the trial (Figure S1C). Spectral clustering (Figures S2A and S2B) revealed six distinct clusters of activity patterns that were similarly represented across FOF and ADS (Figures S2C–S2E), with each cluster capturing different temporal dynamics and evidence sensitivity profiles (Figures S2F–S2H). Side-selective neurons in both regions exhibited ramping activity proportional to stimulus strength for their preferred side (Figure 1F). These stimulus-dependent ramping responses have long been interpreted as neural correlates of evidence accumulation.[5] However, both graded and categorical representations can produce such ramping,[16,37] so we characterized encoding and decoding profiles with more directed analysis to distinguish between these possibilities.

FOF and ADS carry redundant bidirectionally communicated task-related information

The feedforward functional hypothesis posits that ADS has graded encoding of evidence that is relayed to FOF, where this evidence is thresholded, and the categorical choice is represented. Therefore, the hypothesis predicts that evidence can be decoded from ADS earlier and with higher accuracy than from FOF.

To test this hypothesis, we compared graded encoding between FOF and ADS using accumulator tuning curves.[16] Previous studies found FOF neurons had categorical tuning,[16] while ADS neurons showed graded tuning.[14] Our simultaneous recordings allowed controlled comparisons without inter-subject variability. We found both graded and categorical neurons in both regions (Figure 2A). Surprisingly, slopes did not differ significantly between FOF (0.29 ± 0.22) and ADS (0.32 ± 0.25; p = 0.13,

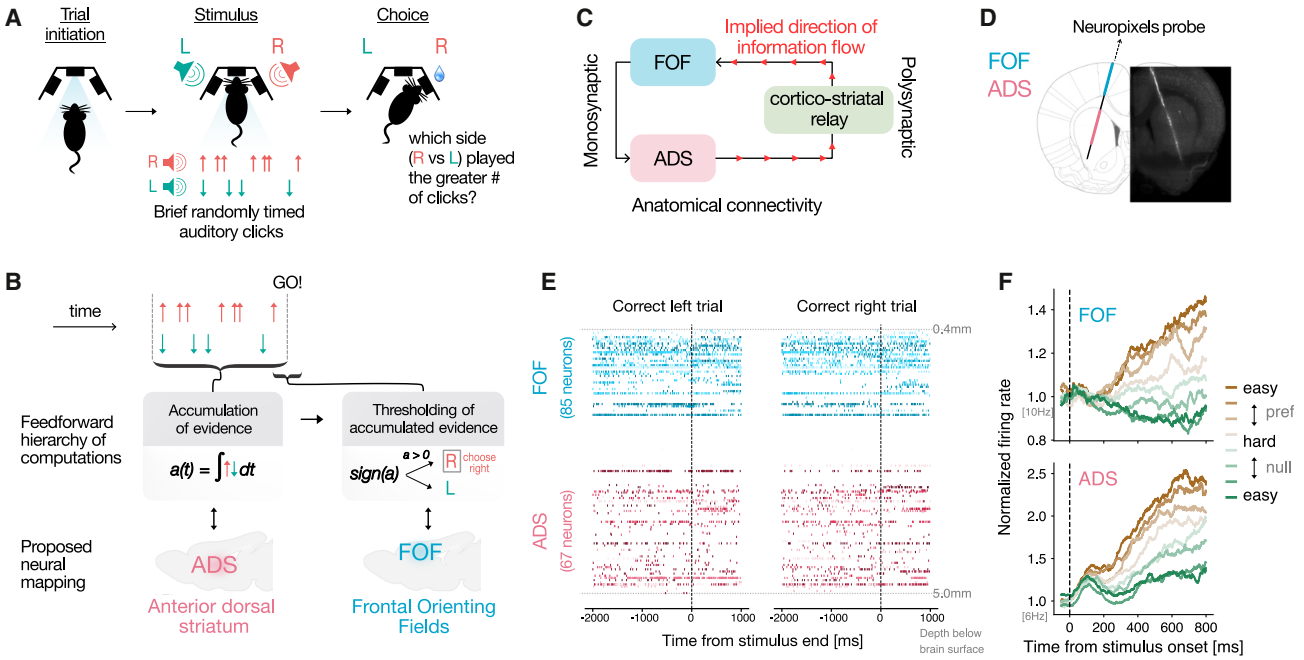

Figure 1. Experimental setup and simultaneous recordings from FOF and ADS

(A) Schematic of the events in the rat evidence accumulation task (adapted from Brunton et al.[32]).

(B) Schematic of proposed decision-making implementation.[14,16,33] The ADS (pink) is thought to be involved in gradual accumulation of sensory evidence and represents accumulated value[14] (*a*). The accumulated evidence is thought to be thresholded and represented in the FOFR[15–17] (blue).

(C) The pathway from ADS to FOF involves a multi-synaptic pathway (gray arrows through olive-green regions) through basal ganglia (BGs) nuclei, the thalamus, and the SC that could potentially support this transformation.

(D) Schematic showing neuropixels targeting overlaid on histology from an example rat. Brain atlas was adapted from Paxinos and Watson.[34]

(E) Spike rasters showing simultaneous neural activity from two example correct trials (left and right). Aqua blue shades represent FOF neurons, and pink shades represent ADS neurons.

(F) FOF (top) and ADS (bottom) population responses during stimulus presentation for different evidence strengths. Brown/green indicates preferred/non-preferred stimuli, and darker hues show easier trials, and gray shows harder trials. Both areas exhibit stimulus-dependent ramping. FOF $n = 219$ neurons, ADS $n = 235$ active side-selective neurons from 5 rats.

Mann-Whitney U test; Figure 2B). Contrary to previous studies, simultaneous measurements revealed no differences in accumulator encoding between regions.

Next, we compared the decodability of information about the stimulus and other task variables in the population activity of FOF and ADS. Despite the similarity in trial-averaged tuning to the accumulator variable (Figure 2B), it is still possible that this information is decodable earlier in ADS and/or with higher accuracy. For each region, we trained separate linear decoders to predict the cumulative difference in the number of right and left clicks at each time point during the trial (Figure S3A). Interestingly, both regions showed comparable decoding performance over time, as measured by the correlation between the true and predicted cumulative stimulus on held-out time points, with no significant differences between the regions (Figure 2C; two-way repeated measures ANOVA, $p = 0.93$ for region and $p = 0.36$ for region $\times$ time). Results were unchanged when controlling for choice information (Figure S3B).

We also decoded choice and trial history from both regions. Choice decoding was high in both regions with no significant differences over time (two-way repeated measures ANOVA, $p = 0.90$ for region, $p = 0.54$ for region $\times$ time; Figure S3C). Past trial choice was decoded well into the current trial from both regions with no regional differences (two-way repeated measures ANOVA, $p = 0.13$ for region, $p = 0.37$ for region $\times$ time; Figure S3D). Finally, all spectral clusters contributed to accumulator tuning and decoding, with evidence-selective clusters receiving slightly higher decoder weights (Figure S4). Altogether, these results show that, at least under the analyses carried out here, the two regions have indistinguishable representations of task-relevant variables, leading us to believe that their involvement in decision-making may be strongly related.

Our conclusions about the simultaneous evolution of evidence-related representations in FOF and ADS are further supported by an independent, model-free decoding of the decision variable (DV) from population activity. The DV—quantified as the distance to the choice decoder hyperplane—reflects the model's confidence and estimates the animal's internal DV, integrating sensory evidence and other choice determinants.[21,38,39] This complements previous model-based analyses by relaxing the behavioral model's assumptions and not fixing stimulus encoding lags between regions. DVs from both regions corresponded well with each other and eventual choices, exhibited classic evidence accumulation ramping (Figure S5), and showed

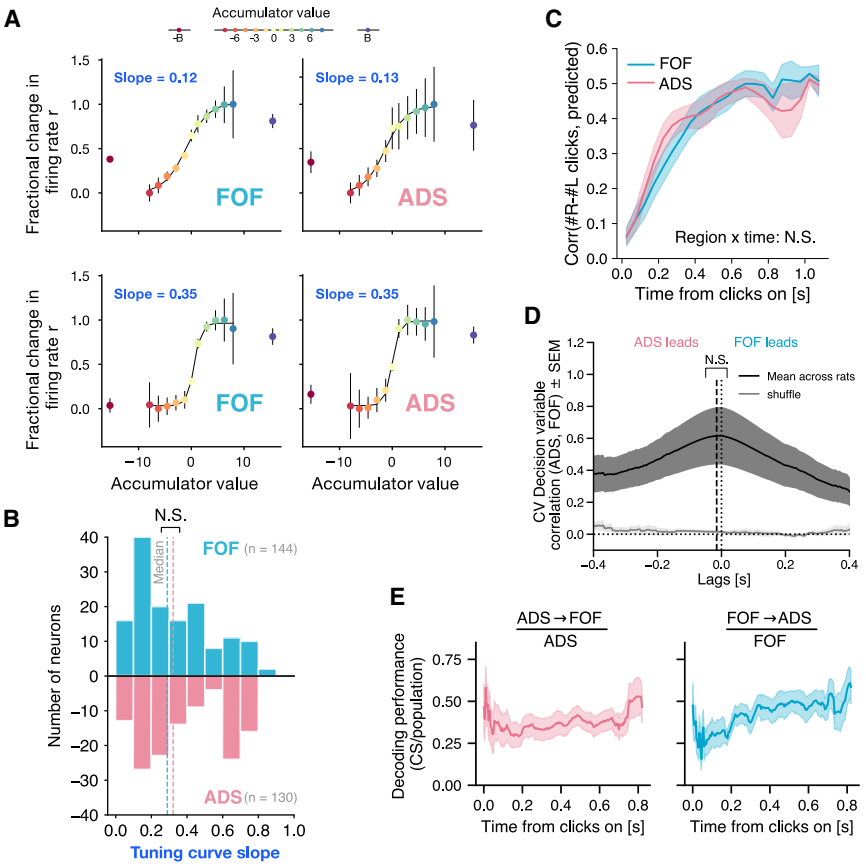

Figure 2. FOF and ADS carry redundant bidirectionally communicated task-related information

(A) Example neurons from FOF (left) and ADS (right) showing graded (top, shallow slope) vs. categorical (bottom, steep slope) tuning to the accumulator variable.

(B) Distribution of tuning slopes in FOF (aqua, n = 144) and ADS (n = 130). Medians: FOF = 0.29 ± 0.22; ADS = 0.32 ± 0.25 (not significantly different, p = 0.13, Mann-Whitney U test).

(C) Linear decoder performance for stimulus decoding over time (n = 12 sessions), measured as correlation between actual click difference trajectory and neural decoding predictions. No significant difference between FOF and ADS (p = 0.36, two-way RM ANOVA).

(D) Time-lagged cross-correlogram of DVs from FOF and ADS. y axis: correlation coefficient; x axis: time lags (positive = FOF leading ADS). Black line, mean across sessions/rats; dark gray, SEM; light gray, shuffled control. Symmetric peak at zero lag (−0.01 ± 0.02 s, p = 0.47 one-sided t test, n = 12), showing no lead-lag relationship.

(E) Stimulus decoding accuracy in CSs: ADS→FOF (pink) and FOF→ADS (blue), as a fraction of whole-population accuracy. Information present in both CSs increases over trial, with FOF sharing slightly more in its CS.

no lead-lag relationship (Figure 2D). The cross-correlogram peaked at zero lag (mean latency = − 0.01 ± 0.02s; p = 0.47 one-sided t test, n = 12), indicating no dominant feedforward or feedback relationship between regions.[40] Simulations validated our ability to detect lead-lag relationships across varying recording conditions (Figure S6).

We next probed FOF-ADS interactions during decision-making beyond activity along the decision hyperplane. We used a generalized linear model with coupling filters across populations (similar to Pillow et al.[41]) plus temporally extended kernels for clicks and task events. To manage parameter explosion with increasing cell numbers, we projected regressor population activity onto learned lower-dimensional subspaces rather than modeling pairwise coupling terms (STAR Methods; Figure S7A). Latent factor dimensions were estimated by cross-validated log-likelihood (Figure S7B). The model captured high variance in average firing rates for most neurons in both areas (Figure S7C). We compared stimulus decoding from FOF → ADS and ADS → FOF communication subspaces (CSs) vs. whole population activity (Figure 2E). Remarkably, sizeable fractions of stimulus information were decodable from CS, increasing over the trial, with FOF sharing slightly more (mean = 0.36 $\frac{(ADS \rightarrow FOF)}{ADS}$; mean = 0.45 $\frac{(FOF \rightarrow ADS)}{FOF}$). Altogether, these results indicate that the representations of evidence in the two regions exist in the space that captures the shared trial-by-trial fluctuations and therefore are likely communicated between the two.

To ensure robustness of our findings, we repeated all statistical analyses, treating rats rather than neurons or sessions as independent units, confirming that our conclusions remain unchanged (Table S1).

Silencing FOF inputs to ADS disrupts decision-making

FOF and ADS showed remarkably similar encoding and decoding of DVs and shared information in CSs, with no discernible lead-lag relationships. This redundancy could indicate that these signals are inherited by both FOF and ADS from a third (unknown) brain region. However, given the causal involvement of these regions during decision-making, a more interesting and likely possibility is that this redundancy reflects ongoing recurrent interactions[40] between FOF and ADS. To test whether these interactions are necessary for decision-making, we probed the monosynaptic FOF→ADS projection using unilateral optogenetics (Figure 3A). We injected the inhibitory opsin, halorhodopsin (AAV5-CaMKIIα-eNpHR3.0-eYFP) bilaterally in FOF and delivered light (25–33 mW, 594 nm) via a sharp fiber optic implanted in ADS to silence the activity of the FOF-ADS projection neurons' axon terminals.

Whole-trial FOF→ADS terminal silencing (2 s inactivation starting 500 ms before the stimulus period and ending 500 ms after stimulus offset when the animal is free to make their response; 25% of trials) in either hemisphere significantly increased ipsilateral choices in eNpHR3.0 rats (p < 0.001, nonparametric bootstrap test) but not eYFP controls (p = 0.78, nonparametric bootstrap test; Figure 3C), confirming involvement in decision-making.

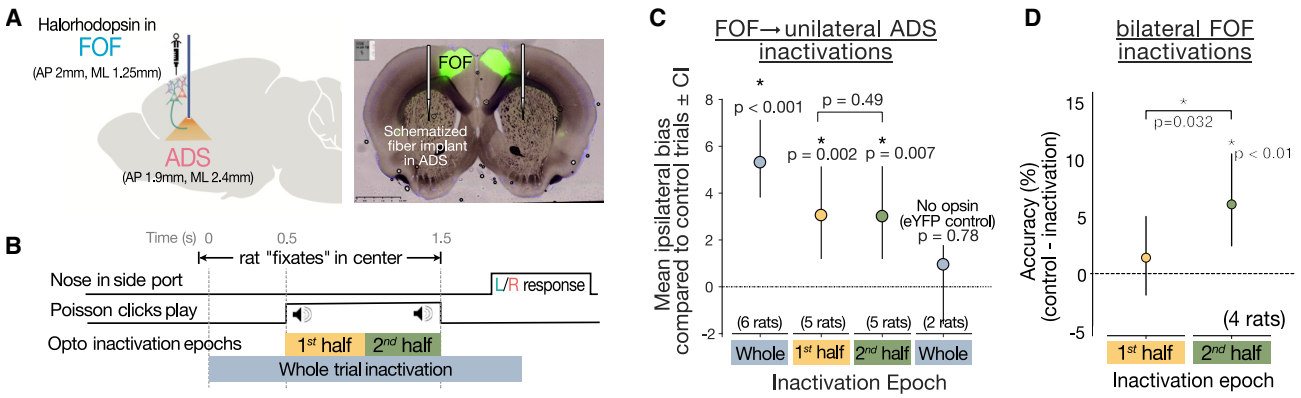

Figure 3. Silencing FOF inputs to ADS disrupts decision-making

(A) (Left) Silencing FOF axon terminals in ADS: experimental design. (Right) Example histology from a rat expressing eNpHR3.0 bilaterally in FOF.

(B) Perturbation epochs: whole-trial (gray), early (yellow), and late (green) during 1 s stimulus.

(C) Ipsilateral bias (mean ± 95% confidence interval [CI]) for FOF→ADS inactivation. Whole-trial: significant impairment (n = 6 rats, 10 hemispheres, $p <$ 0.001, non-parametric bootstrap test), no effect in eYFP controls (n = 2 rats, 3 hemispheres, p = 0.78). Early/late epochs both impaired (n = 5 rats, 8 hemispheres, $p <$ 0.005, non-parametric bootstrap test) with no difference between epochs (p = 0.49), suggesting FOF-ADS involvement throughout accumulation.

(D) Accuracy difference for bilateral FOF inactivation during early vs. late epochs (n = 4 rats). Only late epoch reduces accuracy significantly ($p <$ 0.01, nonparametric bootstrap test), differing from early epoch ($p <$ 0.032, nonparametric bootstrap test).

We then tested temporal involvement by restricting inactivations to early (first half of a 1 s long stimulus) or late (second half of a 1 s long stimulus) stimulus epochs. With this design, it is predicted that if the FOF→ADS projection is involved in the conversion of the graded evidence to a categorical choice, then the inactivation should have an effect on behavior only at the end of the stimulus presentation, i.e., during the late epoch. Whereas, if the projection is involved in evidence accumulation, a process that occurs throughout the stimulus presentation, inactivations during both early and late epochs should have an effect. Surprisingly, both early and late inactivations caused ipsilateral bias (early: p = 0.002; late: p = 0.007 nonparametric bootstrap test) with no difference between epochs (p = 0.49; Figures 3C and S8C). This contrasts with whole-region FOF inactivations that only affect late epochs[16] but resembles ADS inactivations affecting both epochs.[14] These results implicate FOF→ADS projections in gradual accumulation, revealing a previously hidden role for FOF neurons.

Finally, we also examined the effect of inactivations on movements (Figure S8). Terminal silencing slightly but significantly reduced mean movement times for choices both ipsiversive (p = 0.05 Mann-Whitney U test) and contraversive (p = 0.001 Mann-Whitney U test) to the laser (Figure S8A). However, silencing did not increase rates of fixation violations (p = 0.58, nonparametric bootstrap test; Figure S8B). This indicates that the laser did not induce any gross motor impairments affecting animals' ability to successfully complete trials.

Bilateral perturbations of FOF have effects similar to unilateral FOF

FOF→ADS inhibition impaired decisions in both early and late stimulus epochs (Figures 3B and 3C), indicating FOF's role in gradual accumulation like ADS. This supports the similar neural representations observed (Figures 1 and 2) and suggests FOF evidence representations are behaviorally important rather than incidental. However, this raises a puzzle: why don't unilateral whole-region FOF inactivations disrupt decisions during the accumulation period? In mouse M2, interhemispheric recovery has been shown to compensate for unilateral inactivations in early task periods.[21,42] One hypothesis is that the two hemispheres' FOFs are coupled and support recovery when aberrant inputs from the other hemisphere are detected. The localized nature of terminal inhibition (targeting FOF inputs just to ADS, not to the other FOF) might bypass such recovery.

We tested this with bilateral FOF inactivations during early or late stimulus epochs (Figure 3D). Only late-epoch inactivation significantly impaired performance ($p <$ 0.01, nonparametric bootstrap test), while early-epoch inactivation showed no significant impairment ($p >$ 0.05, nonparametric bootstrap test). Early-epoch impairment was significantly smaller than late-epoch (p = 0.032, nonparametric bootstrap test). This indicates that interhemispheric compensation does not underlie the lack of early-epoch impairment in whole-region FOF inactivations. The question of why projection-specific inactivations impair early accumulation decisions but nonspecific whole-region FOF inactivations don't remains outstanding.

A model of distributed implementation of evidence accumulation in cortico-striatal circuitry

Our experimental findings are difficult to reconcile with a feedforward organization wherein regions have localized functions (as previously hypothesized). They are, however, consistent with a circuit that accumulates evidence in a distributed fashion through recurrence and feedback, allowing adaptive rescue of computations upon perturbation.[21,43–46] We therefore developed a mechanistic multi-region RNN model instantiating recurrent FOF-ADS circuitry, building on the success of RNNs in capturing animal behavior and neural representations during decision-making.[47–51] We analyzed this RNN in unperturbed and perturbed states to identify mechanisms enabling sensitivity to

projection-specific inactivations but resilience to nonselective whole-region FOF inactivations during accumulation.

We imposed biological constraints by initializing network weights with anatomical connections following Dale's law and maintaining these constraints throughout training.[52] The architecture (Figures 4A and S9) comprised two hemispheres, each containing three 50-unit submodules representing FOF (20% inhibitory), ADS (100% inhibitory), and a multi-synaptic relay (basal ganglia nuclei, thalamus, SC). Long-range projections followed empirical patterns: FOF projected according to cortical pyramidal cell types (intratelencephalic [IT]and pyramidal tract [PT]), ADS projected unilaterally to the relay, and the relay projected bilaterally to both FOFs, thereby completing a recurrent cortico-striato-cortical loop. Sensory inputs were sent to both FOF and ADS submodules, consistent with diffuse auditory cortical projections to these regions in the rat brain, and network outputs were read from all submodules.

In addition to this biologically grounded architecture, we trained networks on both control and inactivation trials (25%), in which the gain of FOF, ADS, or FOF→ADS projections was transiently reduced during specific stimulus halves. This perturbation-based training encouraged the network to reproduce empirically observed ipsilateral biases during inactivation. Including perturbation data during training introduced functional specialization across regions, guiding the model toward circuit-level solutions consistent with experimental observations—complementing previous modeling efforts that relied solely on representational or architectural differences.[53]

The trained model accurately reproduced choice behavior on both control (Figure 4B) and inactivation trials (Figure 4C). Mirroring the experimental data, unilateral FOF inactivations caused ipsilateral bias only in the second stimulus half, while FOF→ADS and ADS inactivations caused bias throughout the stimulus. Dynamical analysis using reverse-engineering[54] revealed approximate line attractors (not shown), consistent with previous RNN models of evidence accumulation.[31,48,55]

By contrast, networks trained without perturbation constraints exhibited highly variable and often unrealistic responses to inactivation, including contralateral biases or no effect at all, with substantial variability across network instances (Figure S9G). This instability reflected an under-constrained optimization: when trained only on control behavior, networks could solve the task using arbitrary subsets of modules, often relying disproportionately on one region. Analyses of weight matrices post training revealed that networks trained with perturbation constraints developed markedly stronger and more structured connections to and from the striato-cortical relay (SCR), across both excitatory and inhibitory pathways (Figure S9H). These findings suggest that including perturbation data during training encourages the emergence of recurrent loops through the striato-cortical pathway, leading to more biologically realistic and robust solutions.

The RNN model's responses exhibited several features that resembled neural representations, despite not being explicitly trained to do so (Figure 4D). First, we examined stimulus and choice decoding time courses of model FOF and ADS (Figure 4E). Both model regions showed highly similar decoding performance and time courses: stimulus decoding (correlation between predicted and actual cumulative clicks) rose steadily over trials, while choice accuracy increased progressively throughout the stimulus, with nearly identical profiles between regions, mirroring neural data (Figure 2C).

We tested whether the RNN recapitulated bidirectional information sharing by analyzing CSs in trained networks. Having direct weight access, we computed inputs from FOF to ADS and from SCR to FOF.[56] Since SCR only receives ADS inputs, activity through SCR reflects ADS→FOF communication, creating functional bidirectionality. Dimensionality reduction captured 95% of variance (single dimension for 9/10 networks, 2 dimensions for one network). Remarkably, the model reproduced the key empirical finding (Figure 2E): both regions shared stimulus information bidirectionally in CSs, with information increasing over trials and FOF sharing slightly more than ADS. This correspondence emerged spontaneously despite training only on behavioral patterns and not neural population representations.

In addition to these similarities in responses, the model successfully predicted the outcome of a held-out perturbation: bilateral FOF inactivation (Figures 4G vs. 3D). Because training included only unilateral inactivations (Figure S9F), this provided a good test of the model's robustness. Prior work[21] showed that modular, symmetric networks are typically robust to unilateral but not bilateral inactivations. Thus, the model was expected to fail unless perturbation-based training had instilled genuine circuit-level robustness. Strikingly, the model reproduced the empirical finding, showing performance disruption only during the second stimulus half. This non-trivial generalization suggests that perturbation-constrained training led the network to develop recurrent compensatory mechanisms resembling those of the biological circuit.

Finally, we characterized RNN unit diversity (Figure S10) using the same spectral clustering as neural data (Figure S2). Cluster proportions were broadly similar across model FOF and ADS, including evidence-selective units with leftward/rightward preferences, mirroring neural data. However, RNN units differed in two respects: no non-selective clusters were identified (unlike neural clusters 1 and 3), and slow ramping dynamics prominent in neural data were absent. These differences likely stem from distinct training constraints. These differences likely stem from distinct training constraints. Unlike the brain, which coordinates movement, fixation, and contextual processes, our RNN was optimized only for task performance and perturbation robustness. This narrow objective may explain the absence of non-selective clusters and slow ramping dynamics, which in biological circuits depend on modulatory and preparatory inputs. This highlights additional complexities in biological implementations of evidence accumulation.

In silico experiments with the multiregion model reveal network recovery determinants

We examined the responses of neural units during different inactivations by projecting population activity from all submodules onto the first three principal components (PCs; Figures 5A and 5B). Examination of low-dimensional activity trajectories on trials with unilateral inactivation of FOF in the first half revealed that after an initial perturbation, the network activity recovers during the second half (Figure 5A), facilitating normal decision-making

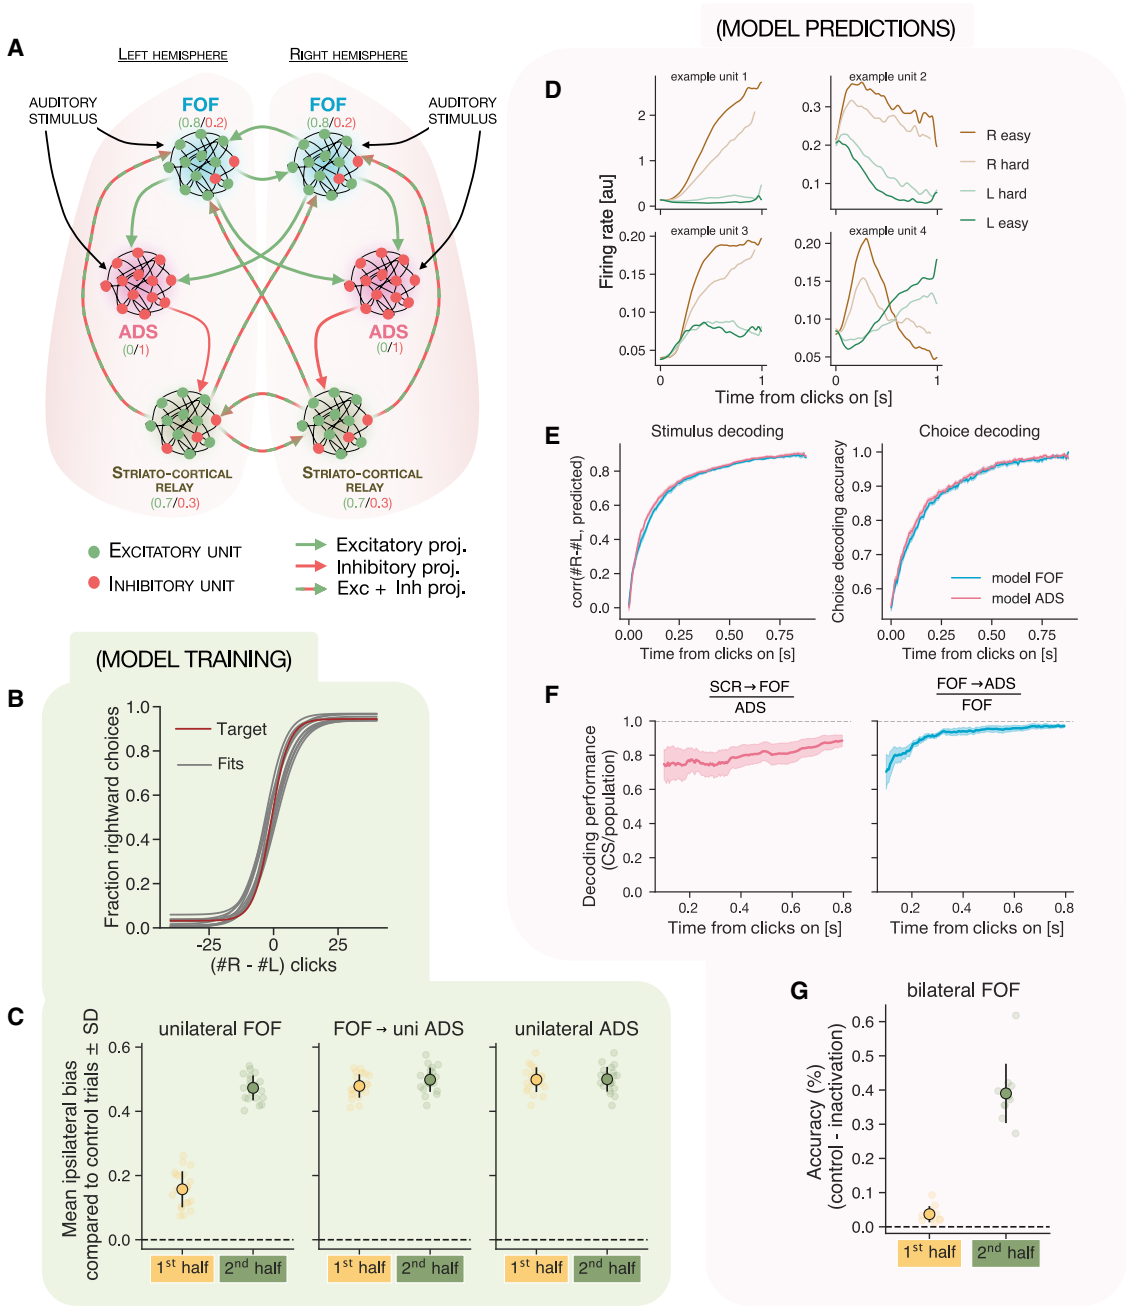

Figure 4. Multiregion recurrent neural network model successfully recapitulates empirical representational and robustness properties

(A) Model architecture with two hemispheres, each containing FOF, ADS, and SCR submodules. Red, inhibitory; green, excitatory. Inputs to FOF/ADS, outputs from all units.

(B) Model psychometric curve (gray; $n = 10$) vs. target curve (red).

(C) Ipsilateral bias (mean ± CI, $n = 10$ networks) for unilateral inactivations of FOF, FOF→ADS, and ADS during the 1st or 2nd stimulus half. Model recapitulates empirical data.

(D) Example model unit PSTHs (aligned to stimulus onset, conditioned on stimulus difficulty).

(E) Stimulus (left) and choice (right) decoding time course for model FOF (blue) and ADS (pink), showing similar performance.

(F) Fraction of stimulus information decoded from CSs over time. Left, ADS→FOF via SCR (pink); right, FOF→ADS (blue). Mean across 10 networks ± SEM. Model reproduces empirical bidirectional information sharing with a slight increase over time, with FOF sharing slightly more than ADS.

(G) Unseen perturbation prediction: accuracy difference for bilateral FOF inactivation during the 1st (orange) or 2nd (green) stimulus half ($n = 10$ networks).

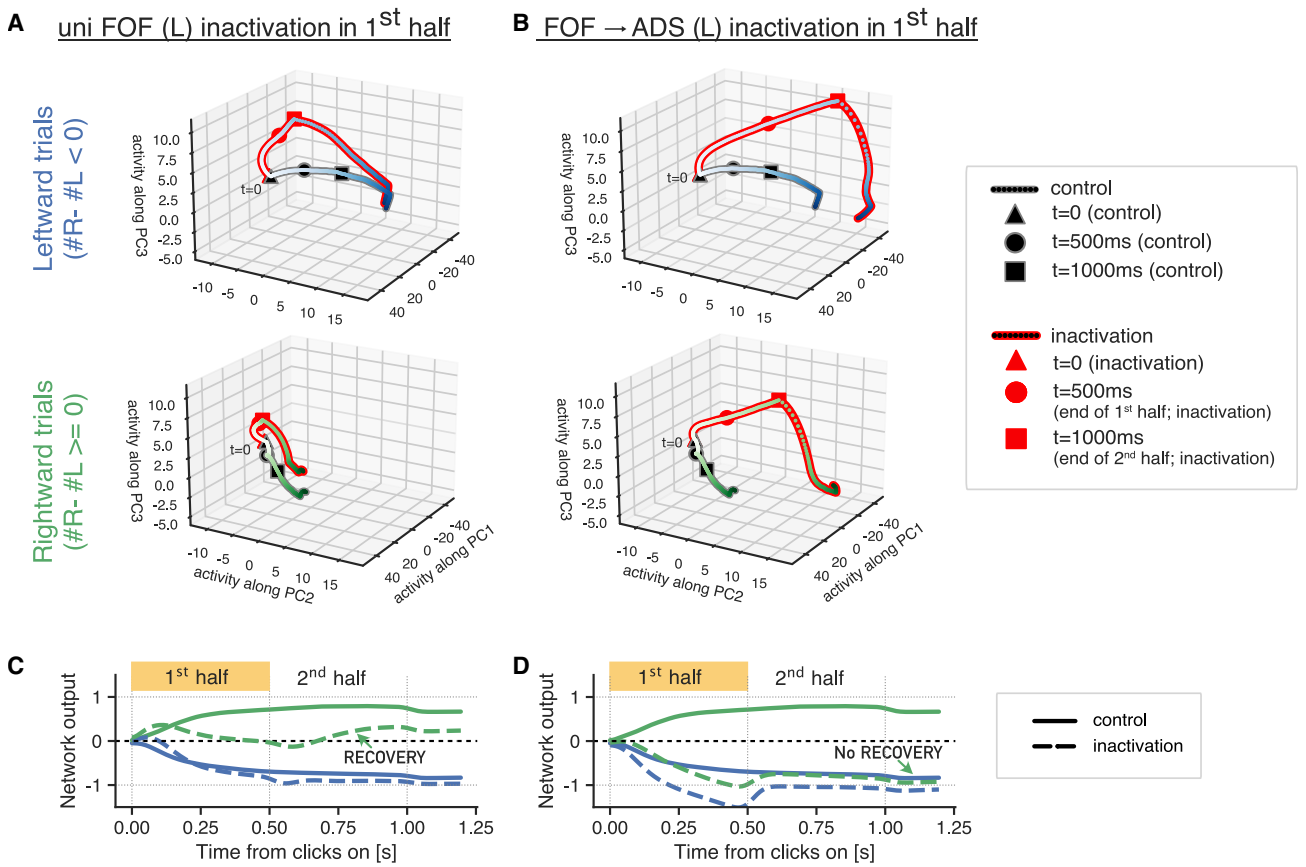

Figure 5. Network activity and outputs during nonselective FOF inactivations reveal recovery dynamics

Low-dimensional projections during left hemisphere inactivations of FOF (A and C) and FOF→ADS (B and D) in the 1st stimulus half. (A and B) Neural trajectories on leftward (top) and rightward (bottom) trials. Inactivation trajectories (red edges) diverge from control (gray) in 1st half (to circles). In 2nd half (circles to squares), FOF inactivations recover toward control (A), but FOF→ADS inactivations remain diverged, driving ipsilateral bias (B). The effect is stronger on rightward trials. (C and D) Network outputs show a similar pattern on rightward trials: FOF inactivations recover and drive correct choices (C), and FOF→ADS inactivations remain diverged (D).

behavior. Such recovery was absent during inactivations of FOF projections to an ADS submodule. Instead, these inactivations sent the population responses along the path that normally drives ipsilateral choices and subsequently gave rise to choices biased toward the inactivated side (Figure 5B). While recovery dynamics in the PC projections (Figures 5A and 5B) are subtle due to the high-dimensional nature of the neural activity, the phenomenon becomes clearly evident in the network outputs (Figures 5C and 5D), which capture the functionally relevant projection onto the decision-relevant subspace.

In order to investigate which parts of the cortico-striatal circuit could potentially be involved in these recovery dynamics, we took advantage of the precise control we had over RNN activity and used it as a testbed for performing several *in silico* perturbations that would be time- and cost-intensive to run *in vivo*.

First, we selectively perturbed individual output projections from the model FOF during the first stimulus half, alone and in combination, to determine what subset of these projections needed to be inactivated to invite the recovery seen in whole-region FOF inactivations. This included self-recurrence, projections to contralateral FOF, and projections to the ipsilateral and contralateral ADS (Figures 6A and S12A). Perturbing projections to the ipsilateral ADS produced strong early inactivation effects with no late recovery (Figure S12B), whereas perturbing projections to the contralateral ADS had no measurable effect (Figure S12C). This suggests that the FOF→ADS inactivation effects observed previously (Figure S9E) were primarily driven by ipsilateral projections.

We next reasoned that inactivating FOF cell bodies should mimic the combined loss of all their outputs. To identify which outputs were necessary for the recovery seen after FOF inactivation, we examined combinations of output perturbations. When the ipsilateral FOF→ADS projection was silenced, additionally inactivating either self-recurrence or the contralateral FOF projection produced partial recovery, while inactivating both restored full recovery (Figures S12D and S12E). Quantifying the resulting ipsilateral bias relative to unilateral FOF inactivation revealed that each additional output inactivation significantly increased recovery (Figure 6B). Thus, the model makes a counterintuitive

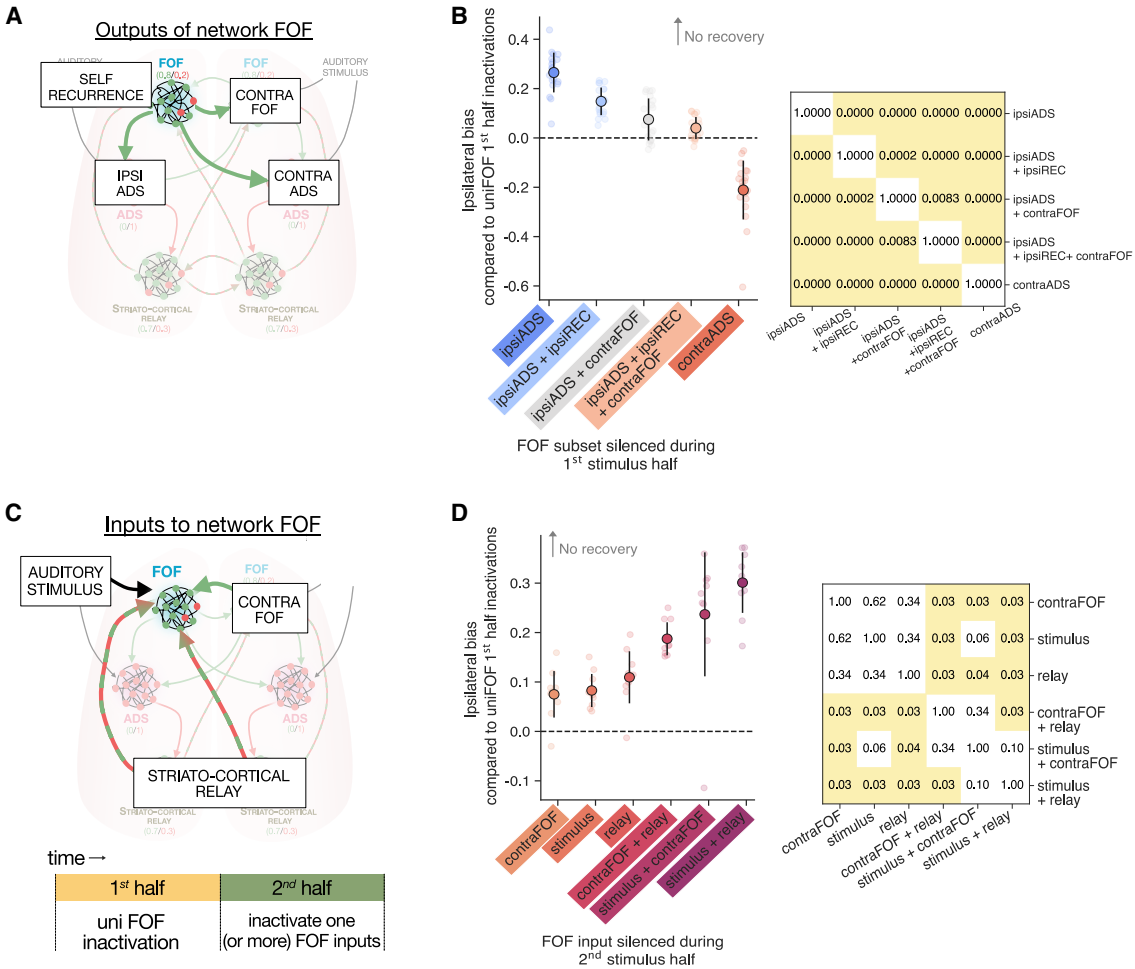

Figure 6. *In silico* experiments with RNN to investigate contributions to inactivation dynamics

(A and B) Selective silencing of different FOF outputs in the 1st half invites differential recovery. (A) FOF output projections: ipsilateral ADS, contralateral ADS, contralateral FOF, and recurrent connections.

(B) Ipsilateral bias for different output combinations relative to full FOF silencing. Silencing ipsilateral ADS projections alone produces the largest effect with the least recovery, and adding recurrent connections or contralateral FOF projections enables partial recovery individually and full recovery when combined. Contralateral ADS projections alone produce negative bias. Inset: *p* values for pairwise comparisons.

(C and D) Selectively silencing FOF inputs in 2nd half shows multiple contributions to recovery. (C) FOF inputs: contralateral FOF, auditory stimulus, and SCR. (D) Impact on ipsilateral bias relative to the 1st half of whole FOF silencing. Silencing contralateral FOF, auditory stimulus, or SCR inputs each reduces recovery partially, and combined silencing produces greater recovery loss. Inset: *p* values (Wilcoxon test with Holm correction).

prediction: simultaneous inactivation of additional FOF outputs (contralateral FOF + ipsilateral ADS) should paradoxically eliminate the behavioral deficits caused by ipsilateral ADS inactivation alone.

Second, we tested which inputs to an FOF submodule contributed to late-phase recovery by selectively silencing them during the second stimulus half. These included auditory stimulus inputs, inputs from the contralateral FOF, and inputs from the SCR (Figure 6C). Removing any single input had only minor effects on recovery, but inactivating input pairs produced substantially greater impairments (Figure 6D). These results predict that ongoing sensory input is critical for second-half recovery and that both relay and contralateral FOF inputs contribute cooperatively to this process.

Together, these analyses suggest that recovery following unilateral FOF perturbations depends on multiple scales of recurrence within the cortico-striatal network: self-recurrence within each FOF, interhemispheric FOF-FOF coupling, and striato-cortical recurrence through the relay. They further yield several falsifiable predictions: (1) contralateral ADS projections (modeling IT-type cortical outputs) are not necessary for recovery; (2) unilateral FOF recovery arises because contralateral FOF outputs are perturbed, whereas within-FOF perturbations have a smaller impact; (3) ongoing stimulus input during the second half is essential for robust recovery; and (4) inputs from both the SCR and contralateral FOF jointly support second-half recovery. Together, these findings define the key circuit interactions underlying recovery and establish concrete hypotheses for future experimental tests.

DISCUSSION

Two separate brain regions in rats, the ADS and the FOF, have been, respectively, linked to the two distinct computations necessary for decision-making based on noisy evidence: the gradual accumulation of evidence and the categorization of that accumulated evidence to reach a final decision.[14,16] However, this implied feedforward functional hierarchy going from ADS to FOF has never been directly evaluated. We sought to fill this important gap by simultaneously recording population activity from these two regions during decision-making (Figure 1). Contrary to the feedforward hypothesis, we found that the two regions do not differ in their accumulator encoding, and evidence can be decoded from both regions with comparable accuracy and at comparable timescales. Moreover, we show that this information is shared between the two regions and shows no discernible lead-lag relationship, consistent with a recurrent interaction (Figure 2). Further, we optogenetically silenced FOF axon terminals in ADS (i.e., the feedback projection under the hypothesis) during time periods that differentially overlap with the two sub-computations and found that such a manipulation impairs decision-making during all accumulation periods, unlike nonselective FOF perturbations (Figure 3). Together, these results indicate that the cortico-striatal circuitry participates in evidence accumulation in a distributed manner with ongoing recurrent interactions. We synthesized these findings into a multi-region recurrent neural network model. We successfully trained this RNN with the novel objective of reproducing behavior on control and inactivation trials and found that it recapitulated neural responses and perturbation effects that it wasn't trained on (Figure 4). Finally, we investigated the circuit's response to perturbations by examining the RNN's activity (Figure 5) and performing *in silico* experiments on the RNN, targeting different sets of projections (Figure 6). We found that FOF's robustness to nonselective perturbations was supported by multiple scales of recurrence in the cortico-striatal circuit.

The similarities we observed in the encoding of task-relevant features in FOF and its downstream target ADS during decision formation are largely consistent with recent reports in mice in simple decision-making tasks, which found that choice signals in secondary motor cortex (that FOF is a part of) and striatum emerge at indistinguishable timescales[57] and that striatum reflects the summed activity of its cortical inputs.[58] Moreover, our results about the encoding of sensory evidence in FOF are in agreement with a recent study of the dorsal cortex's role during decision-making in mice, which reported that the secondary motor cortex prominently encodes choice-independent evidence-related signals.[59] In the random dots motion task,[60] monkey frontal eye fields (FEF) and caudate (the analogs to rat FOF and ADS) have respectively been reported to carry similar accumulation-related signals early during the stimulus presentation with comparable prevalence of trial-difficulty and choice-modulated signals during the stimulus epoch (FEF,[61] caudate,[62] FEF, and caudate[63–65]). Further, this similarity in evidence representation and functional role of the cortex (FOF) and striatum has been predicted and is therefore consistent with previous theoretical accounts that have sought to embed the decision-making computation into the cortico-basal ganglia network.[3,4] However, similarities with Lo and Wang[2] are unclear, which predicts transient rather than sustained responses in the striatum and models the striatum's role as the setting of evidence thresholds. While these parallels point to a shared cortico-striatal implementation of evidence accumulation, other studies suggest that the relative contributions of cortical and striatal regions may diverge under certain behavioral contexts.

Some studies have found differences between FEF and caudate during decision-making, particularly in tasks with reward manipulations that require the adaptive tuning of decision thresholds.[64–66] Such results fit with wider studies implicating the striatum in action selection[67] and reward-based habitual learning and the secondary motor cortex in planning goal-directed movements.[68,69] Such distinctions are important avenues for future study and could reveal wider non-overlapping roles of the two regions but are out of scope for our study that involves a task with fixed, symmetric rewards and is not designed to distinguish between these other functions. Moreover, we have focused our analysis on the encoding and decoding properties of specific task variables in the two regions, such as accumulator tuning, choice, and history. We acknowledge the possibility that other views of the neural activity have the potential to reveal subtle differences where we have found none. Whether such distinctions influence the functional roles of these regions is an open question, and one that we leave for future work to investigate. Although our analyses were not designed to disentangle these broader functional distinctions, examining the specific pathways linking cortex and striatum can help clarify how their interactions shape decision computations.

To directly probe these circuit interactions, we targeted the cortico-striatal pathway by silencing FOF axon terminals in ADS and revealed a role for FOF throughout the gradual accumulation process. This finding might seem at odds with past work that has concluded, using unilateral silencing of FOF somas, that FOF's causal role is restricted to the decision commitment period and does not extend to the accumulation period[16]—however, it is in line with recent exciting work highlighting how different projection pathways from the same cortical region can carry different information and differentially impair behavior upon perturbation[22,70–79] (further discussion[67,80]), producing effects that could be otherwise obscured by whole-region perturbations. A notable example of this comes from an auditory frequency discrimination task in rats,[70] where perturbations of auditory cortex neurons with high (low) frequency tuning did not bias animals' choices toward high (low) responses, but selective activation of striatum projecting auditory cortex neurons with high (low) frequency tuning biased the animals' choices as predicted. Also of note is a recent study in mice,[75] which has also proposed a role for cortico-striatal projections from the frontal cortex throughout gradual accumulation of evidence. Our results add an important piece of evidence in favor of the emerging principle that projection-specific perturbations can unmask effects not visible in whole-region perturbations.[22,24,70] Given recent controversies about the causal roles of different regions in evidence accumulation,[15,25,81–83] our approach might offer a way to further clarify the roles of these brain regions and help with identifying the key communication channels operant during decision-making.

Together with previous temporally precise perturbation studies (FOF[16] and ADS[14]), our findings highlight the value of leveraging the temporal structure of decision subcomputations to infer causal roles. This approach has proven powerful not only for dissecting the neural basis of decision-making but also for understanding other temporally separable computations such as working memory.[43,84] Yet, gradual accumulation itself comprises multiple subcomputations—computing the evidence, incorporating prior biases, setting the relevant timescale, etc.[5,6,32] It is possible that despite being involved in accumulation, different brain areas differentially affect one or more of these computational components.[7,14,25,85–87] Disentangling which of these components are altered by inactivation is statistically demanding and has been feasible primarily with large, cross-session datasets. In rats, muscimol studies illustrate this challenge well: unilateral FOF inactivation was best explained by a post-categorization bias rather than changes to accumulation per se, whereas bilateral FOF effects could be captured by changes to integration time constants and related parameters.[88] By contrast, ADS/striatal inactivation increased accumulation-related noise consistent with a direct role in the accumulation computation.[14] By comparison, temporally specific optogenetics—ours included—yields far fewer trials because of opsin expression windows, multiple conditions, and interleaved controls; therefore, full parameter recovery is underpowered, and robust summary metrics (e.g., ipsilateral bias) are preferable. In retrospect, it seems that the results of temporally specific optogenetics led us to over-rely on the interpretation from unilateral effects—of FOF contributing largely post-categorization. In line with our findings of FOF being involved in accumulation, a recent study[89] showed graded encoding of utility in FOF, and recent large-scale recordings suggest coordinated, brain-wide timing of decision commitment across cortico-striatal loops rather than a strict serial hierarchy.[90,91] Further work will be needed using scaled-up trial counts to fit the full behavioral model to different perturbations and understand if the different parameters of accumulation are controlled modularly by different brain areas.

Beyond its role in decision formation, the observed reduction in movement times following terminal silencing raises important questions about the broader role of cortico-striatal projections in action control. This effect occurred for choices for both ipsiversive and contraversive to the laser, suggesting it reflects a change in movement vigor rather than the decision process itself, given the concurrent process of an ipsilateral choice bias. Importantly, violation rates were unaffected, indicating no gross motor impairments. This interpretation aligns with prior work demonstrating that cortico-striatal circuits modulate movement kinematics independently of decision content.[92–94] Our results suggest that even unilaterally suppressing FOF to ADS projections may remove a layer of top-down modulation over movement gain, implicating this circuit in both decision formation and its execution.

Finally, our modeling results complement these empirical findings by demonstrating how biologically constrained RNNs can capture both neural and behavioral effects of perturbation. While the encoding/decoding properties of units in our network resemble those of the empirical data, a complementary approach would be to train the RNN such that individual units in the network additionally match observed neural firing rates.[56,95–98] It remains an exciting empirical question of how these different training objectives, algorithms, and constraints alter the dynamics of the network and which objectives make better predictions. Past studies seeking to model the complex web of multi-region interactions that give rise to cognition[45,46,99–102] (reviewed in Yang and Molano-Mazón[53] and Perich and Rajan[103]) by training RNNs have relied either on differences in neural representations or just the architecture to lend RNNs area-specific identities. We believe our use of projection-specific perturbation data to introduce structure is novel and is likely to tighten the correspondence between the models and data (see also Cowley et al.[104]). We are hopeful that this approach of using multi-region RNNs in tandem with projection-specific perturbations and simultaneous recording will offer a powerful framework for generating hypotheses and intuitions about how recurrently connected brain regions perform decision-making in a distributed fashion.

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to Diksha Gupta (diksha.gupta@ucl.ac.uk).

Materials availability

This study did not generate new materials.

Data and code availability

- The processed data have been deposited at figshare and are publicly available as of the date of publication at Figshare: https://doi.org/10.6084/m9.figshare.30369064.
- All original code has been deposited at GitHub and is publicly available at Zenodo: https://doi.org/10.5281/zenodo.10161052.
- Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

ACKNOWLEDGMENTS

We thank members of the Brody lab for experimental support and helpful feedback throughout the project. We also thank Jovanna Teran and Brody lab technicians for assistance with rat training. We are grateful to Jonathan Pillow, Ilana Witten, Joshua Gold, David Zoltowski, and Aditi Jha for their suggestions and discussions. This work was supported by NIH grant R01MH108358 awarded to C.D.B. as well as a grant from the Simons Foundation (grant number: NC-GB-CULM-00003118-03) awarded to C.D.B.

AUTHOR CONTRIBUTIONS

D.G., C.D.K., and A.G.B. collected data. T.Z.L., A.G.B., and V.E. helped with data collection and pre-processing. D.G. analyzed and modeled the data. D.G. wrote the initial draft of the manuscript. All authors provided feedback on the manuscript. C.D.B. oversaw all aspects of the project.

DECLARATION OF INTERESTS

The authors declare no competing interests.

DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During manuscript preparation, the authors used Anthropic's Claude to help reduce word count. The authors reviewed and edited all output and take full responsibility for the content.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
  - Neuronal recordings and analysis
  - Optogenetic perturbations and analysis
  - Multi-region recurrent neural network model
- QUANTIFICATION AND STATISTICAL ANALYSIS

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.neuron.2025.12.029.

Received: December 16, 2024
Revised: September 24, 2025
Accepted: December 17, 2025
Published: January 21, 2026

REFERENCES

1. Wong, K.F., and Wang, X.J. (2006). A Recurrent Network Mechanism of Time Integration in Perceptual Decisions. J. Neurosci. *26*, 1314–1328. https://doi.org/10.1523/JNEUROSCI.3733-05.2006.
2. Lo, C.C., and Wang, X.J. (2006). Cortico–basal ganglia circuit mechanism for a decision threshold in reaction time tasks. Nat. Neurosci. *9*, 956–963. https://doi.org/10.1038/nn1722.
3. Bogacz, R., and Gurney, K. (2007). The Basal Ganglia and Cortex Implement Optimal Decision Making Between Alternative Actions. Neural Comput. *19*, 442–477. https://doi.org/10.1162/neco.2007.19.2.442.
4. Rao, R.P.N. (2010). Decision Making Under Uncertainty: A Neural Model Based on Partially Observable Markov Decision Processes. Front. Comput. Neurosci. *4*, 146. https://doi.org/10.3389/fncom.2010.00146.
5. Gold, J.I., and Shadlen, M.N. (2007). The Neural Basis of Decision Making. Annu. Rev. Neurosci. *30*, 535–574. https://doi.org/10.1146/annurev.neuro.29.051605.113038.
6. Hanks, T.D., and Summerfield, C. (2017). Perceptual Decision Making in Rodents, Monkeys, and Humans. Neuron *93*, 15–31. https://doi.org/10.1016/j.neuron.2016.12.003.
7. Jun, E.J., Bautista, A.R., Nunez, M.D., Allen, D.C., Tak, J.H., Alvarez, E., and Basso, M.A. (2021). Causal role for the primate superior colliculus in the computation of evidence for perceptual decisions. Nat. Neurosci. *24*, 1121–1131. https://doi.org/10.1038/s41593-021-00878-6.
8. Stine, G.M., Trautmann, E.M., Jeurissen, D., and Shadlen, M.N. (2023). A neural mechanism for terminating decisions. Neuron *111*, 2601–2613.e5. https://doi.org/10.1016/j.neuron.2023.05.028.
9. Okazawa, G., and Kiani, R. (2023). Neural Mechanisms that Make Perceptual Decisions Flexible. Annu. Rev. Physiol. *85*, 191–215. https://doi.org/10.1146/annurev-physiol-031722-024731.
10. Gold, J.I., and Shadlen, M.N. (2002). Banburismus and the Brain: Decoding the Relationship between Sensory Stimuli, Decisions, and Reward. Neuron *36*, 299–308. https://doi.org/10.1016/S0896-6273(02)00971-6.
11. Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J.D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. Psychol. Rev. *113*, 700–765. https://doi.org/10.1037/0033-295X.113.4.700.
12. Simen, P., Cohen, J.D., and Holmes, P. (2006). Rapid decision threshold modulation by reward rate in a neural network. Neural Netw. *19*, 1013–1026. https://doi.org/10.1016/j.neunet.2006.05.038.
13. Drugowitsch, J., Moreno-Bote, R., Churchland, A.K., Shadlen, M.N., and Pouget, A. (2012). The Cost of Accumulating Evidence in Perceptual Decision Making. J. Neurosci. *32*, 3612–3628. https://doi.org/10.1523/JNEUROSCI.4010-11.2012.
14. Yartsev, M.M., Hanks, T.D., Yoon, A.M., and Brody, C.D. (2018). Causal contribution and dynamical encoding in the striatum during evidence accumulation. eLife *7*, e34929. https://doi.org/10.7554/eLife.34929.
15. Erlich, J.C., Brunton, B.W., Duan, C.A., Hanks, T.D., and Brody, C.D. (2015). Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. eLife *4*, e05457. https://doi.org/10.7554/eLife.05457.
16. Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature *520*, 220–223. https://doi.org/10.1038/nature14066.
17. Piet, A.T., Erlich, J.C., Kopec, C.D., and Brody, C.D. (2017). Rat Prefrontal Cortex Inactivations during Decision Making Are Explained by Bistable Attractor Dynamics. Neural Comput. *29*, 2861–2886. https://doi.org/10.1162/neco_a_01005.
18. Boyd-Meredith, J.T., Piet, A.T., Dennis, E.J., El Hady, A., and Brody, C.D. (2022). Stable choice coding in rat frontal orienting fields across model-predicted changes of mind. Nat. Commun. *13*, 3235. https://doi.org/10.1038/s41467-022-30736-3.
19. Kaufman, M.T., Churchland, M.M., Ryu, S.I., and Shenoy, K.V. (2014). Cortical activity in the null space: permitting preparation without movement. Nat. Neurosci. *17*, 440–448. https://doi.org/10.1038/nn.3643.
20. Semedo, J.D., Zandvakili, A., Machens, C.K., Yu, B.M., and Kohn, A. (2019). Cortical Areas Interact through a Communication Subspace. Neuron *102*, 249–259.e4. https://doi.org/10.1016/j.neuron.2019.01.026.
21. Li, N., Daie, K., Svoboda, K., and Druckmann, S. (2016). Robust neuronal dynamics in premotor cortex during motor planning. Nature *532*, 459–464. https://doi.org/10.1038/nature17643.
22. Warden, M.R., Selimbeyoglu, A., Mirzabekov, J.J., Lo, M., Thompson, K.R., Kim, S.Y., Adhikari, A., Tye, K.M., Frank, L.M., and Deisseroth, K. (2012). A prefrontal cortex-brainstem neuronal projection that controls response to behavioural challenge. Nature *492*, 428–432. https://doi.org/10.1038/nature11617.
23. Otis, J.M., Namboodiri, V.M.K., Matan, A.M., Voets, E.S., Mohorn, E.P., Kosyk, O., McHenry, J.A., Robinson, J.E., Resendez, S.L., Rossi, M.A., et al. (2017). Prefrontal cortex output circuits guide reward seeking through divergent cue encoding. Nature *543*, 103–107. https://doi.org/10.1038/nature21376.
24. Kim, C.K., Ye, L., Jennings, J.H., Pichamoorthy, N., Tang, D.D., Yoo, A.W., Ramakrishnan, C., and Deisseroth, K. (2017). Molecular and Circuit-Dynamical Identification of Top-Down Neural Mechanisms for Restraint of Reward Seeking. Cell *170*, 1013–1027.e14. https://doi.org/10.1016/j.cell.2017.07.020.
25. Katz, L.N., Yates, J.L., Pillow, J.W., and Huk, A.C. (2016). Dissociated functional significance of decision-related activity in the primate dorsal stream. Nature *535*, 285–288. https://doi.org/10.1038/nature18617.
26. Semedo, J.D., Gokcen, E., Machens, C.K., Kohn, A., and Yu, B.M. (2020). Statistical methods for dissecting interactions between brain areas. Curr. Opin. Neurobiol. *65*, 59–69. https://doi.org/10.1016/j.conb.2020.09.009.
27. Gilad, A., Gallero-Salas, Y., Groos, D., and Helmchen, F. (2018). Behavioral strategy determines frontal or posterior location of short-term memory in neocortex. Neuron *99*, 814–828.e7. https://doi.org/10.1016/j.neuron.2018.07.029.
28. Latimer, K.W., and Freedman, D.J. (2021). Low-dimensional encoding of decisions in parietal cortex reflects long-term training history. Preprint at bioRxiv. https://doi.org/10.1101/2021.10.07.463576.
29. Arlt, C., Barroso-Luque, R., Kira, S., Bruno, C.A., Xia, N., Chettih, S.N., Soares, S., Pettit, N.L., and Harvey, C.D. (2022). Cognitive experience

alters cortical involvement in goal-directed navigation. eLife *11*, e76051. https://doi.org/10.7554/eLife.76051.

30. Bolkan, S.S., Stone, I.R., Pinto, L., Ashwood, Z.C., Iravedra Garcia, J.M., Herman, A.L., Singh, P., Bandi, A., Cox, J., Zimmerman, C.A., et al. (2022). Opponent control of behavior by dorsomedial striatal pathways depends on task demands and internal state. Nat. Neurosci. *25*, 345–357. https://doi.org/10.1038/s41593-022-01021-9.
31. Pagan, M., Tang, V.D., Aoi, M.C., Pillow, J.W., Mante, V., Sussillo, D., and Brody, C.D. (2025). Individual variability of neural computations underlying flexible decisions. Nature *639*, 421–429. https://doi.org/10.1038/s41586-024-08433-6.
32. Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and humans can optimally accumulate evidence for decision-making. Science *340*, 95–98. https://doi.org/10.1126/science.1233912.
33. Brody, C.D., and Hanks, T.D. (2016). Neural underpinnings of the evidence accumulator. Curr. Opin. Neurobiol. *37*, 149–157. https://doi.org/10.1016/j.conb.2016.01.003.
34. Paxinos, G., and Watson, C. (2006). The Rat Brain In Stereotaxic Coordinates, Sixth edition (Amsterdam: Elsevier, Academic Press).
35. Jun, J.J., Steinmetz, N.A., Siegle, J.H., Denman, D.J., Bauza, M., Barbarits, B., Lee, A.K., Anastassiou, C.A., Andrei, A., Aydın, Ç., et al. (2017). Fully integrated silicon probes for high-density recording of neural activity. Nature *551*, 232–236. https://doi.org/10.1038/nature24636.
36. Luo, T.Z., Bondy, A.G., Gupta, D., Elliott, V.A., Kopec, C.D., and Brody, C.D. (2020). An approach for long-term, multi-probe Neuropixels recordings in unrestrained rats. Elife *9*, e59716. https://doi.org/10.7554/eLife.59716.
37. Latimer, K.W., Yates, J.L., Meister, M.L.R., Huk, A.C., and Pillow, J.W. (2015). Neuronal Modeling. Single-trial Spike Trains in Parietal Cortex Reveal Discrete Steps During Decision-making. Science *349*, 184–187. https://doi.org/10.1126/science.aaa4056.
38. Kiani, R., Cueva, C.J., Reppas, J.B., and Newsome, W.T. (2014). Dynamics of Neural Population Responses in Prefrontal Cortex Indicate Changes of Mind on Single Trials. Curr. Biol. *24*, 1542–1547. https://doi.org/10.1016/j.cub.2014.05.049.
39. Peixoto, D., Verhein, J.R., Kiani, R., Kao, J.C., Nuyujukian, P., Chandrasekaran, C., Brown, J., Fong, S., Ryu, S.I., Shenoy, K.V., et al. (2021). Decoding and perturbing decision states in real time. Nature *591*, 604–609. https://doi.org/10.1038/s41586-020-03181-9.
40. Semedo, J.D., Jasper, A.I., Zandvakili, A., Krishna, A., Aschner, A., Machens, C.K., Kohn, A., and Yu, B.M. (2022). Feedforward and feedback interactions between visual cortical areas use different population activity patterns. Nat. Commun. *13*, 1099. https://doi.org/10.1038/s41467-022-28552-w.
41. Pillow, J.W., Shlens, J., Paninski, L., Sher, A., Litke, A.M., Chichilnisky, E.J., and Simoncelli, E.P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. Nature *454*, 995–999. https://doi.org/10.1038/nature07140.
42. Chen, G., Kang, B., Lindsey, J., Druckmann, S., and Li, N. (2021). Modularity and robustness of frontal cortical networks. Cell *184*, 3717–3730.e24. https://doi.org/10.1016/j.cell.2021.05.026.
43. Kopec, C.D., Erlich, J.C., Brunton, B.W., Deisseroth, K., and Brody, C.D. (2015). Cortical and Subcortical Contributions to Short-Term Memory for Orienting Movements. Neuron *88*, 367–377. https://doi.org/10.1016/j.neuron.2015.08.033.
44. Duan, C.A., Pagan, M., Piet, A.T., Kopec, C.D., Akrami, A., Riordan, A.J., Erlich, J.C., and Brody, C.D. (2021). Collicular circuits for flexible sensorimotor routing. Nat. Neurosci. *24*, 1110–1120. https://doi.org/10.1038/s41593-021-00865-x.
45. Bredenberg, C., Savin, C., and Kiani, R. (2024). Recurrent neural circuits overcome partial inactivation by compensation and re-learning. J. Neurosci. *44*, e1635232024. https://doi.org/10.1523/JNEUROSCI.1635-23.2024.
46. Mejías, J.F., and Wang, X.J. (2022). Mechanisms of distributed working memory in a large-scale network of macaque neocortex. eLife *11*, e72136. https://doi.org/10.7554/eLife.72136.
47. Rigotti, M., Ben Dayan Rubin, D., Wang, X.J., and Fusi, S. (2010). Internal representation of task rules by recurrent dynamics: the importance of the diversity of neural responses. Front. Comput. Neurosci. *4*, 24. https://doi.org/10.3389/fncom.2010.00024.
48. Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature *503*, 78–84. https://doi.org/10.1038/nature12742.
49. Carnevale, F., de Lafuente, V., Romo, R., Barak, O., and Parga, N. (2015). Dynamic Control of Response Criterion in Premotor Cortex during Perceptual Detection under Temporal Uncertainty. Neuron *86*, 1067–1077. https://doi.org/10.1016/j.neuron.2015.04.014.
50. Remington, E.D., Egger, S.W., Narain, D., Wang, J., and Jazayeri, M. (2018). A Dynamical Systems Perspective on Flexible Motor Timing. Trends Cogn. Sci. *22*, 938–952. https://doi.org/10.1016/j.tics.2018.07.010.
51. Kleinman, M., Chandrasekaran, C., and Kao, J. (2021). A mechanistic multi-area recurrent network model of decision-making. Adv. Neural Inf. Process. Syst. *34*, 23152–23165.
52. Song, H.F., Yang, G.R., and Wang, X.J. (2016). Training Excitatory-Inhibitory Recurrent Neural Networks for Cognitive Tasks: A Simple and Flexible Framework. PLOS Comput. Biol. *12*, e1004792. https://doi.org/10.1371/journal.pcbi.1004792.
53. Yang, G.R., and Molano-Mazón, M. (2021). Towards the next generation of recurrent network models for cognitive neuroscience. Curr. Opin. Neurobiol. *70*, 182–192. https://doi.org/10.1016/j.conb.2021.10.015.
54. Sussillo, D., and Barak, O. (2013). Opening the Black Box: Low-Dimensional Dynamics in High-Dimensional Recurrent Neural Networks. Neural Comput. *25*, 626–649. https://doi.org/10.1162/NECO_a_00409.
55. Maheswaranathan, N., Williams, A.H., Golub, M.D., Ganguli, S., and Sussillo, D. (2019). Universality and individuality in neural dynamics across large populations of recurrent networks. Adv. Neural Inf. Process. Syst. *2019*, 15629–15641.
56. Perich, M.G., Arlt, C., Soares, S., Young, M.E., Mosher, C.P., Minxha, J., Carter, E., Rutishauser, U., Rudebeck, P.H., Harvey, C.D., et al. (2020). Inferring brain-wide interactions using data-constrained recurrent neural network models. Preprint at bioRxiv. https://doi.org/10.1101/2020.12.18.423348.
57. Steinmetz, N.A., Zatka-Haas, P., Carandini, M., and Harris, K.D. (2019). Distributed coding of choice, action and engagement across the mouse brain. Nature *576*, 266–273. https://doi.org/10.1038/s41586-019-1787-x.
58. Peters, A.J., Fabre, J.M.J., Steinmetz, N.A., Harris, K.D., and Carandini, M. (2021). Striatal activity topographically reflects cortical activity. Nature *591*, 420–425. https://doi.org/10.1038/s41586-020-03166-8.
59. Orsolic, I., Rio, M., Mrsic-Flogel, T.D., and Znamenskiy, P. (2021). Mesoscale cortical dynamics reflect the interaction of sensory evidence and temporal expectation during perceptual decision-making. Neuron *109*, 1861–1875.e10. https://doi.org/10.1016/j.neuron.2021.03.031.
60. Roitman, J.D., and Shadlen, M.N. (2002). Response of Neurons in the Lateral Intraparietal Area during a Combined Visual Discrimination Reaction Time Task. J. Neurosci. *22*, 9475–9489. https://doi.org/10.1523/JNEUROSCI.22-21-09475.2002.
61. Kim, J.N., and Shadlen, M.N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat. Neurosci. *2*, 176–185. https://doi.org/10.1038/5739.
62. Ding, L., and Gold, J.I. (2010). Caudate Encodes Multiple Computations for Perceptual Decisions. J. Neurosci. *30*, 15747–15759. https://doi.org/10.1523/JNEUROSCI.2894-10.2010.
63. Ding, L., and Gold, J.I. (2012). Neural Correlates of Perceptual Decision Making before, during, and after Decision Commitment in Monkey

Frontal Eye Field. Cereb. Cortex *22*, 1052–1067. https://doi.org/10.1093/cercor/bhr178.

64. Ding, L. (2015). Distinct dynamics of ramping activity in the frontal cortex and caudate nucleus in monkeys. J. Neurophysiol. *114*, 1850–1861. https://doi.org/10.1152/jn.00395.2015.
65. Fan, Y., Gold, J.I., and Ding, L. (2020). Frontal eye field and caudate neurons make different contributions to reward-biased perceptual decisions. eLife *9*, e60535. https://doi.org/10.7554/eLife.60535.
66. Fan, Y., Doi, T., Gold, J.I., and Ding, L. (2024). Neural Representations of Post-Decision Accuracy and Reward Expectation in the Caudate Nucleus and Frontal Eye Field. J. Neurosci. *44*, e0902232023. https://doi.org/10.1523/JNEUROSCI.0902-23.2023.
67. Cox, J., and Witten, I.B. (2019). Striatal circuits for reward learning and decision-making. Nat. Rev. Neurosci. *20*, 482–494. https://doi.org/10.1038/s41583-019-0189-2.
68. Ebbesen, C.L., Insanally, M.N., Kopec, C.D., Murakami, M., Saiki, A., and Erlich, J.C. (2018). More than Just a "Motor": Recent Surprises from the Frontal Cortex. J. Neurosci. *38*, 9402–9413. https://doi.org/10.1523/JNEUROSCI.1671-18.2018.
69. Yang, J.H., and Kwan, A.C. (2020). Secondary motor cortex: Broadcasting and biasing animal's decisions through long-range circuits. Int. Rev. Neurobiol. *158*, 443–470. https://doi.org/10.1016/bs.irn.2020.11.008.
70. Znamenskiy, P., and Zador, A.M. (2013). Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. Nature *497*, 482–485. https://doi.org/10.1038/nature12077.
71. Li, N., Chen, T.W., Guo, Z.V., Gerfen, C.R., and Svoboda, K. (2015). A motor cortex circuit for motor planning and movement. Nature *519*, 51–56.
72. Murugan, M., Jang, H.J., Park, M., Miller, E.M., Cox, J., Taliaferro, J.P., Parker, N.F., Bhave, V., Hur, H., Liang, Y., et al. (2017). Combined Social and Spatial Coding in a Descending Projection from the Prefrontal Cortex. Cell *171*, 1663–1677.e16. https://doi.org/10.1016/j.cell.2017.11.002.
73. Hwang, E.J., Link, T.D., Hu, Y.Y., Lu, S., Wang, E.H.J., Lilascharoen, V., Aronson, S., O'Neil, K., Lim, B.K., and Komiyama, T. (2019). Corticostriatal Flow of Action Selection Bias. Neuron *104*, 1126–1140.e6. https://doi.org/10.1016/j.neuron.2019.09.028.
74. Hirokawa, J., Vaughan, A., Masset, P., Ott, T., and Kepecs, A. (2019). Frontal cortex neuron types categorically encode single decision variables. Nature *576*, 446–451. https://doi.org/10.1038/s41586-019-1816-9.
75. Musall, S., Sun, X.R., Mohan, H., An, X., Gluf, S., Li, S.J., Drewes, R., Cravo, E., Lenzi, I., Yin, C., et al. (2023). Pyramidal cell types drive functionally distinct cortical activity patterns during decision-making. Nat. Neurosci. *26*, 495–505. https://doi.org/10.1038/s41593-022-01245-9.
76. Tervo, D.G.R., Kuleshova, E., Manakov, M., Proskurin, M., Karlsson, M., Lustig, A., Behnam, R., and Karpova, A.Y. (2021). The anterior cingulate cortex directs exploration of alternative strategies. Neuron *109*, 1876–1887.e6. https://doi.org/10.1016/j.neuron.2021.03.028.
77. Duan, C.A., Pan, Y., Ma, G., Zhou, T., Zhang, S., and Xu, N.L. (2021). A cortico-collicular pathway for motor planning in a memory-dependent perceptual decision task. Nat. Commun. *12*, 2727. https://doi.org/10.1038/s41467-021-22547-9.
78. Wolff, S.B.E., Ko, R., and Ölveczky, B.P. (2022). Distinct roles for motor cortical and thalamic inputs to striatum during motor skill learning and execution. Sci. Adv. *8*, eabk0231. https://doi.org/10.1126/sciadv.abk0231.
79. Park, J., Phillips, J.W., Guo, J.Z., Martin, K.A., Hantman, A.W., and Dudman, J.T. (2022). Motor cortical output for skilled forelimb movement is selectively distributed across projection neuron classes. Sci. Adv. *8*, eabj5167. https://doi.org/10.1126/sciadv.abj5167.
80. Jazayeri, M., and Ostojic, S. (2021). Interpreting neural computations by examining intrinsic and embedding dimensionality of neural activity. Preprint at arXiv. https://doi.org/10.48550/arXiv.2107.04084.
81. Raposo, D., Kaufman, M.T., and Churchland, A.K. (2014). A category-free neural population supports evolving demands during decision-making. Nat. Neurosci. *17*, 1784–1792. https://doi.org/10.1038/nn.3865.
82. Licata, A.M., Kaufman, M.T., Raposo, D., Ryan, M.B., Sheppard, J.P., and Churchland, A.K. (2017). Posterior Parietal Cortex Guides Visual Decisions in Rats. J. Neurosci. *37*, 4954–4966. https://doi.org/10.1523/JNEUROSCI.0105-17.2017.
83. Jeurissen, D., Shushruth, S., El-Shamayleh, Y., Horwitz, G.D., and Shadlen, M.N. (2022). Deficits in decision-making induced by parietal cortex inactivation are compensated at two timescales. Neuron *110*, 1924–1931.e5. https://doi.org/10.1016/j.neuron.2022.03.022.
84. Guo, Z.V., Inagaki, H.K., Daie, K., Druckmann, S., Gerfen, C.R., and Svoboda, K. (2017). Maintenance of persistent activity in a frontal thalamocortical loop. Nature *545*, 181–186. https://doi.org/10.1038/nature22324.
85. Hanks, T.D., Ditterich, J., and Shadlen, M.N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. Nat. Neurosci. *9*, 682–689. https://doi.org/10.1038/nn1683.
86. Doi, T., Fan, Y., Gold, J.I., and Ding, L. (2020). The caudate nucleus contributes causally to decisions that balance reward and uncertain visual information. eLife *9*, e56694. https://doi.org/10.7554/eLife.56694.
87. Stine, G.M., Zylberberg, A., Ditterich, J., and Shadlen, M.N. (2020). Differentiating between integration and non-integration strategies in perceptual decision making. eLife *9*, e55365. https://doi.org/10.7554/eLife.55365.
88. Erlich, J.C., Bialek, M., and Brody, C.D. (2011). A Cortical Substrate for Memory-Guided Orienting in the Rat. Neuron *72*, 330–343. https://doi.org/10.1016/j.neuron.2011.07.010.
89. Bao, C., Zhu, X., Mōller-Mara, J., Li, J., Dubroqua, S., and Erlich, J.C. (2023). The rat frontal orienting field dynamically encodes value for economic decisions under risk. Nat. Neurosci. *26*, 1942–1952. https://doi.org/10.1038/s41593-023-01461-x.
90. Khilkevich, A., Lohse, M., Low, R., Orsolic, I., Bozic, T., Windmill, P., and Mrsic-Flogel, T.D. (2024). Brain-wide dynamics linking sensation to action during decision-making. Nature *634*, 890–900. https://doi.org/10.1038/s41586-024-07908-w.
91. Bondy, A.G., Charlton, J.A., Luo, T.Z., Kopec, C.D., Stagnaro, W.M., Venditto, S.J.C., Lynch, L., Janarthanan, S., Oline, S.N., Harris, T.D., et al. (2025). Brain-wide coordination of decision formation and commitment. Preprint at bioRxiv. https://doi.org/10.1101/2024.08.21.609044.
92. Pandarinath, C., O'Shea, D.J., Collins, J., Jozefowicz, R., Stavisky, S.D., Kao, J.C., Trautmann, E.M., Kaufman, M.T., Ryu, S.I., Hochberg, L.R., et al. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. Nat. Methods *15*, 805–815. https://doi.org/10.1038/s41592-018-0109-9.
93. Yttri, E.A., and Dudman, J.T. (2016). Opponent and bidirectional control of movement velocity in the basal ganglia. Nature *533*, 402–406. https://doi.org/10.1038/nature17639.
94. Dudman, J.T., and Krakauer, J.W. (2016). The basal ganglia: from motor commands to the control of vigor. Curr. Opin. Neurobiol. *37*, 158–166. https://doi.org/10.1016/j.conb.2016.02.005.
95. Fisher, D., Olasagasti, I., and Tank, D.W. (2013). A modeling framework for deriving the structural and functional architecture of a short-term memory microcircuit. Neuron *79*, 987–1000. https://doi.org/10.1016/j.neuron.2013.06.041.
96. Rajan, K., Harvey, C.D., and Tank, D.W. (2016). Recurrent Network Models of Sequence Generation and Memory. Neuron *90*, 128–142. https://doi.org/10.1016/j.neuron.2016.02.009.
97. Andalman, A.S., Burns, V.M., Lovett-Barron, M., Broxton, M., Poole, B., Yang, S.J., Grosenick, L., Lerner, T.N., Chen, R., Benster, T., et al. (2019). Neuronal Dynamics Regulating Brain and Behavioral State Transitions. Cell *177*, 970–985.e20. https://doi.org/10.1016/j.cell.2019.02.037.

98. Duncker, L., O'Shea, D.J., Shenoy, K.V., and Sahani, M. (2020). A dynamical model with E/I balance explains robustness to optogenetic stimulation in motor cortex. In Computational and Systems Neuroscience (COSYNE) (Denver, Co.) http://web.stanford.edu/~shenoy/GroupPublications/DunckerEtAlCOSYNE2020.pdf.

99. Murray, J.D., Jaramillo, J., and Wang, X.J. (2017). Working Memory and Decision-Making in a Frontoparietal Circuit Model. J. Neurosci. *37*, 12167–12186. https://doi.org/10.1523/JNEUROSCI.0343-17.2017.

100. Kleinman, M., Chandrasekaran, C., and Kao, J.C. (2019). Recurrent neural network models of multi-area computation underlying decision-making. Preprint at bioRxiv. https://doi.org/10.1101/798553.

101. Pinto, L., Rajan, K., DePasquale, B., Thiberge, S.Y., Tank, D.W., and Brody, C.D. (2019). Task-Dependent Changes in the Large-Scale Dynamics and Necessity of Cortical Regions. Neuron *104*, 810–824.e9. https://doi.org/10.1016/j.neuron.2019.08.025.

102. Márton, C.D., Schultz, S.R., and Averbeck, B.B. (2020). Learning to select actions shapes recurrent dynamics in the corticostriatal system. Neural Netw. *132*, 375–393. https://doi.org/10.1016/j.neunet.2020.09.008.

103. Perich, M.G., and Rajan, K. (2020). Rethinking brain-wide interactions through multi-region 'network of networks' models. Curr. Opin. Neurobiol. *65*, 146–151. https://doi.org/10.1016/j.conb.2020.11.003.

104. Cowley, B.R., Calhoun, A.J., Rangarajan, N., Ireland, E., Turner, M.H., Pillow, J.W., and Murthy, M. (2024). Mapping model units to visual neurons reveals population code for social behaviour. Nature *629*, 1100–1108. https://doi.org/10.1038/s41586-024-07451-8.

105. Putzeys, J., Raducanu, B.C., Carton, A., De Ceulaer, J., Karsh, B., Siegle, J.H., Van Helleputte, N., Harris, T.D., Dutta, B., Musa, S., et al. (2019). Neuropixels Data-Acquisition System: A Scalable Platform for Parallel Recording of 10 000+ Electrophysiological Signals. IEEE Trans. Biomed. Circuits Syst. *13*, 1635–1644. https://doi.org/10.1109/TBCAS.2019.2943077.

106. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine Learning in Python. J. Mach. Learn. Res. *12*, 2825–2830. http://jmlr.org/papers/v12/pedregosa11a.html.

107. Vallat, R. (2018). Pingouin: statistics in Python. J. Open Source Softw. *3*, 1026. https://doi.org/10.21105/joss.01026.

108. Park, I.M., Meister, M.L.R., Huk, A.C., and Pillow, J.W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. Nat. Neurosci. *17*, 1395–1403. https://doi.org/10.1038/nn.3800.

109. Bondy, A.G., Luo, T.Z., Gupta, D., and Brody, C.D. (2020). Anterior and posterior striatum play distinct roles in evidence accumulation. In Computational and Systems Neuroscience Meeting (Cosyne Abstracts).

110. Kingma, D.P., and Ba, J. (2017). Adam: A Method for Stochastic Optimization. Preprint at arXiv. https://doi.org/10.48550/arXiv.1412.6980.

111. Ansel, J., Yang, E., He, H., Gimelshein, N., Jain, A., Voznesensky, M., Bao, B., Bell, P., Berard, D., Burovski, E., et al. (2024). PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation. In ASPLOS '24: Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems.

112. Dudman, J.T., and Gerfen, C.R. (2015). Chapter 17. The Basal Ganglia. In The Rat Nervous System, Fourth Edition, G. Paxinos, ed. (Academic Press), pp. 391–440. https://doi.org/10.1016/B978-0-12-374245-2.00017-6.

113. Glorot, X., and Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. In Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, JMLR Workshop and Conference Proceedings-7228, pp. 249–256.

114. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. (2016). (TensorFlow): a system for (Large-Scale) machine learning. In 12th USENIX symposium on operating systems design and implementation (OSDI 16), pp. 265–283.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental models: Organisms/strains		
Rat: Long Evans	Hilltop	http://www.hilltoplabs.com/public/rats.html
Software and algorithms		
MATLAB	Mathworks	https://www.mathworks.com/products/matlab.html
Bcontrol		https://brodylabwiki.princeton.edu/bcontrol/index.php/Main_Page
SpikeGLX	HHMI/Janelia Research Campus	https://billkarsh.github.io/SpikeGLX/
JRClust	HHMI/Janelia Research Campus	https://github.com/JaneliaSciComp/JRCLUST
Kilosort 2	Marius Pachitariu	https://github.com/jamesjun/Kilosort2

EXPERIMENTAL MODEL AND SUBJECT DETAILS

A total of 19 male Long-evans rats (*Rattus norvegicus* from Hilltop, PA) were used for this study. Of these, 5 animals were used for neural recordings and 14 for optogenetic experiments. Investigators were not blinded to experimental groups during data collection or analyses. Animal use procedures were approved by the Princeton University Institutional Animal Care and Use Committee (IACUC #1853) and were carried out in accordance with National Institute of Health standards.

Rats were housed in Techniplast cages and all training and testing were conducted during the dark cycle. Animals that trained during the day were housed in a 12 hr reverse light cycle room. Animals were pair housed whenever possible, but were always single housed after optic fiber or Neuropixels implantation to prevent damage to the implant. Rats had free access to food but were placed on a controlled water schedule in order to motivate them to participate in behavioral task for water rewards. They obtained water rewards during behavioral training sessions (1-5hrs). Following behavioral training rats received either an unlimited water supplement in an ad lib access period of up to 1hr or a controlled supplement so that their total access to water within a period of 24hrs equalled 3% of their body weight.

METHOD DETAILS

Neuronal recordings and analysis

Acquisition and pre-processing

We followed the general surgery and Neuropixels probe implantation methods for recording chronically in freely moving animals from[36] to record unihemispherically from FOF and ADS of 5 rats. The Neuropixels probes were assembled in a compact 3D printed casing, printed in-house using Formlabs SLA printer. This assembly allowed the probe to be stereotaxically manipulated, provided electromagnetic shielding, prevented it from contacting biological fluids or other adhesive materials applied during surgery and imparted mechanical protection to the implant for robust chronic tethered recordings. We targeted FOF and ADS with one penetration, by implanting the probe at AP 1.9mm and ML ± 1.15mm from bregma, at an angle of 15° in the coronal plane. Before implantation, the probe shank was dipped for several seconds in a dye (DiI, 1-2 mg/ml in isopropyl alcohol) to allow for histological reconstruction of probe location. We inserted at least 6mm of the probe shank into the brain; the actual depth of implantation varied between animals.

Electrophysiological recordings were made using either commercially available Neuropixels 1.0 acquisition hardware[105] or the earliest test-phase IMEC acquisition hardware. We used SpikeGLX software to acquire the data. On every session we recorded from 384 sites spanning banks 0 and 1 of the probe. The data were automatically spike sorted offline using Kilosort2 with default parameters and then manually created using JRClust's GUI.[35] During manual curation, each unit detected by Kilosort2 was inspected, if the events comprising the unit had near zero amplitudes or non-physiological waveforms then the entire unit was discarded. If the unit was judged to be comprised of multiple distinct waveform then it was split into two or more units and finally spatially neighboring units were compared using cross-correlograms, drift patterns and waveform similarity to determine if they should be merged. After following these steps, all units with < 5% refractory violations were considered as reflecting spikes from a single neuron.

Analysis of neural recordings

Neuron selection criteria. All population-level analyses (e.g., clustering, decoding, decision variable cross-correlation, and RR-GLM) included all recorded neurons with mean firing rate > 1 Hz, regardless of side-selectivity or tuning profile. For analyses requiring individual neuron characterization (tuning curves, population PSTH), we applied analysis-relevant selection criteria as specified below.

Average PSTH of side-selective neurons. To compute the average population response to different stimulus strengths, we selected neurons that had large differences in their firing rates during the stimulus period for trials that subsequently resulted in rightward versus leftward choices ($|auc-0.5|>0.05, P<0.05$, Receiver operating characteristic). The choice that produced the larger response was defined as the preferred side. For each neuron, firing rates for individual trials were calculated by smoothing the spike trains with a causal half-Gaussian filter with 0.05s standard deviation, and then normalizing by neuron's mean response at stimulus onset.

Neural clustering. We systematically characterized the diversity of neural response profiles by performing unsupervised spectral clustering.[74] For each neuron, we aligned spike trains to stimulus onset and computed average firing rates in 50ms bins over a 700ms window, separately for leftward and rightward trials. The resulting two time series were concatenated into a single feature vector and normalized to the [0,1] range to emphasize temporal structure while minimizing effects of firing rate scale. We applied spectral clustering using a correlation-based k-nearrest neighbors (k-NN) affinity matrix, symmetrized to ensure a valid similarity graph. Feature vectors from all neurons with firing rate >1Hz from FOF and ADS were included. Clustering was performed on the resulting graph using the normalized Laplacian and k-means for label assignment. To select hyperparameters, we performed a grid search over the number of clusters (3–15) and neighbors (5–50). For each parameter combination, we assessed clustering stability using the adjusted Rand Index (ARI) across 100 bootstrap iterations, each using 90% subsampling without replacement. Based on ARI stability and interpretability, we selected 6 clusters as optimal. We quantified the proportion of FOF and ADS neurons assigned to each cluster and tested for regional differences using chi-square tests. For each identified cluster, we characterized: (1) average temporal dynamics by plotting mean normalized firing rates over time for leftward/rightward trials, (2) choice selectivity by computing the distribution of AUC values within each cluster, (3) contribution to evidence sensitivity by examining tuning curve slopes for each cluster, and (4) contribution to population-level decoding by examining the distribution of decoder weights assigned to neurons in each cluster. The clustering of RNN responses was performed in an identical fashion, using the same temporal binning, normalization, and spectral clustering parameters to ensure direct comparability with the neural data.

Neural tuning curves. Following Hanks et al.,[16] we fit the trial-by-trial behavioral model developed in Brunton et al.[32] separately to each rats' choices. In order to best describe the behavior on recording sessions and account for nonstationarities, we first fit the model to choices from sessions (min 30 sessions, upto 90) prior to recording. We then used the best-fit parameters and their estimated confidence intervals to impose Gaussian priors on the parameters and refit the model to behavior from just the recording days. The mean of the Gaussian prior was set to the maximum likelihood estimate of the parameters, and the standard deviation was set to 2 times the estimated confidence interval. Once fit, we obtained the trial-by-trial, moment-by-moment estimate of the likely accumulator variable trajectories while taking into account the choice made by the rat from the model. This was computed by selectively propagating accumulator values that are consistent with rat's choice - at the end of the stimulus - backwards in time.

We collated across trials and used these estimates of the accumulator value to compute the joint probability distribution between accumulator value, time and observed firing rates for each neuron. Following past work, the neural time was assumed to lag behind the time in the model by 100ms. Our results are invariant to small deviations (50 − 100ms) from this assumed lag. Firing rates were computed by smoothing the spike trains with a Gaussian filter with 0.75s standard deviation. We marginalized the joint probability distribution over time to extract the response of the neuron given the value of the accumulator. We analysed the time period between 0.15 and 0.5s from stimulus onset and responses from side-selective neurons ($|auc-0.5|>0.075, P<0.05$). This slightly higher AUC threshold yielded firing rate maps that were largely stable in choice-selectivity over time, as this stability is necessary to avoid artificially flattening tuning curve slopes when averaging across time points. We accounted for differences in the dynamic ranges of the neural response, by scaling the responses to span the range from 0 to 1. We then fit the estimated relationship between scaled neural responses r and the accumulator value a with a four-parameter sigmoid:

$$r = k_1 + \frac{k_2}{1 + e^{-k_3(a-k_4)}}$$

At zero-crossing, the slope of this equation is given by $k_2 k_3/4$. We compared this slope parameter for individual neurons across the two regions and used the Mann-Whitney U test to determine whether the medians of each region's population distribution were significantly different.

Decoding analyses. We trained linear decoders to predict the cumulative difference in number of right and left clicks (Δclicks) from neural population activity. To be able to compare decoding performance from the two regions, we controlled for differences in the number of recorded neurons by sub-sampling neurons from the region with higher number, while maintaining the overall distribution of side selectivity. We considered time points aligned to stimulus onset, and assumed a neural encoding lag of 100ms. Δclicks during the trial and neural responses were both binned into 50ms bins, neural responses were further smoothed with a Gaussian kernel with 75ms standard deviation. Neural responses were then assembled into a design matrix X with dimensions $T \times N$, where T is equal to the total number of time points (time points from stimulus onset to termination × number of trials), and N is the number of neurons. Similarly, Δclicks were concatenated across trials to yield Y with dimensions $T \times 1$. We fit one set of regression weights β_{st} for all time points such that $Y = X\beta_{st}$ and imposed L1 penalty. We performed 10-fold cross-validation to assess the decoder's performance. Within each fold the strength of regularization was estimated independently using 10-fold cross-validation. Fitting was performed using the sklearn package in Python.[106] To determine if the decoding performance or its timecourse was different between the two regions, we performed a two-way repeated measures ANOVA with time (time bins from stimulus onset) and area (FOF, ADS) as factors using Pingouin package in Python.[107]

To decode binary variables, choice and past trial's choice we used logistic decoders. We assembled the design matrix similarly, however without any neural lags since the target variables are constant for a trial. We duplicated the eventual or past choice for all time bins from a trial to construct the *Y* variable. We fit the decoders with L2 penalty, once again with 10-fold cross-validation. Again, we used sklearn package in Python for fitting.[106] While estimating the decision variables or DVs using the decoders trained to predict choice, we binned the neural activity into 1ms bins and convolved them with a Gaussian kernel with 25ms standard deviation before projecting onto the decision hyperplane. These parameters were picked based on recovery analysis on simulated synthetic datasets. However, our results are not sensitive to small variations in these parameters.

Reduced rank Genralized linear model (RR-GLM). We used RR-GLM to define the FOF → ADS and ADS→FOF communication subspaces. For this we fit a GLM with a softplus link function to the entire population activity. The model was optimized to estimate the firing rate of neurons at each time point *t*, assuming Poisson emissions and given a set of time-varying task parameters. Neural spikes during a period of 2s from initiation of a trial were fit. The following task variables were included: trial start, stimulus start, left and right clicks and current trial's choice (separate variables for the left and right choice were included). The task variables were linearly convolved with a kernel to allow for time-varying affect on neural spiking. Kernels for choice were allowed to be acausal, whereas all other kernels were causal in their timing. Following past work,[108,109] to enforce smoothness onto these temporal kernels, we parameterized each kernel using raised cosine bases. Hyper-parameters, such as number and duration of these kernels, were optimized through cross-validation on a small grid. Additionally, temporal kernels with raised cosine bases were fit for the simultaneously recorded population from the other region, incorporating a one-time step delay between the regressor and regressand populations. The regressor activity was first projected onto a lower rank subspace to which the convolutional kernels were applied. We binned the neural spiking activity into 5ms bins. Both the linear weights for the kernels and the dimensions spanning the lower-rank subspace were fit simultaneously by minimizing negative Poisson log-likelihood of the spikes with an L2 penalty, using the Adam optimizer[110] in Pytorch.[111] The regularization strength was chosen using cross-validation. For decoding from the communication subspace, procedures identical to full population space (section above) were followed.

Optogenetic perturbations and analysis

Optical fiber chemical sharpening

We followed methods similar to Hanks et al. (2015)[16] and used standard off the shelf 50-125μm FC-FC duplex fiber optic cables (FiberCables.com). The metal casing of the connector and the outer protective layers of the cable were removed, yielding a 1.5cm of fiber optic cable with inner plastic coating (typically clear) intact. To etch the fiber tip, 2-2.5mm of the fiber tip was submerged in a 5ml Eppendorf tube containing 48% hydrofluoric acid, topped with mineral oil. Over the course of 17mins, the fiber tip was slowly pulled out of the acid using a motor (Narishige) producing a taper in the inner plastic coating. Then the speed of the motor was increased and the motor was run until the tip was entirely removed from the acid (typically 13mins), producing a sharp fiber tip, usually with uniform and broad light scatter. We measured the laser beam transmission efficiency of the etched fibers using an integrating sphere photodiode power sensor (Thor Labs). Fibers that did not produce sufficiently broad or uniform scatter or had efficiency $< 85\%$ were discarded.

Optogenetic virus injection and fiber implantation

We followed previously described procedures for general surgery and virus injections[16,88]; here we describe procedures specific to this study. We bilaterally injected adeno-associated virus, either AAV5-CaMKIIa-eYFP-eNpHR3.0 or AAV2/5-CamMKIIa-EYFP-WPRE-hGHpA using a Nanoject (Drummond Scientific) in FOF. Five closely spaced injection tracts were made in each FOF's craniotomy. The tracts were typically arranged 500μm apart at the 4 vertices and the center of a cross spanning AP, ML axes, centered at AP 2mm ML $\pm$1.3mm. In each tract, 16 injections of 14.1nL were made every 100μm in depth, starting at 200μm below brain surface i.e. from 0.2-1.8mm. Virus was expelled at 20nL/s and the injections were made once every 10s. At the final injection in a tract the pipette was left in place for at least 2min before removal. In each craniotomy a total of 1.128μl of virus was injected over a 30min period consisting of 80 separate injections.

For bilateral FOF silencing, chemically sharpened optical fibers were then lowered down the central injection tracts to a depth of 1mm. For inhibition of FOF terminals in ADS we made additional craniotomies centered over ADS (AP 1.9mm, ML $\pm$2.4mm) and lowered etched fibers to a depth of 3.5-3.9mm. The craniotomies were filled with kwik-sil (World Precision Instruments) and the fibers were secured to the skull with viscous dental composite (Absolute Dentin, Parkell). Optical fibers were then enclosed in a custom 3D printed casing and the implant was filled with dental acrylic, allowing only the FC connectors to protrude. Before behavioral experiments, viral constructs were allowed to develop for 6-8 weeks for FOF silencing, and for 10-12 weeks for terminal silencing. Accurate viral injection targeting and expression was verified histologically.

Optogenetic perturbation

For optogenetically perturbing neural activity during behavior, a fiber rotary joint (Princetel) was mounted on the ceiling of the behavioral chamber. In the beginning of a session, the FC connectors in the rat's implant were connected to the patch cables attached to the fiber rotary joint. The rotary joint was connected to the laser beam source. For FOF inactivations, the laser beam from a 200mW, 532nm laser (OEM Laser Systems) was used. To deliver illumination bilaterally, the laser beam was split into two beams of roughly equal power ($\sim$ 25mW) using a beam splitter (Doric). An OBIS 594nm (Coherent) laser beam of power 25-33mW was used for silencing of FOF axon terminals in either left or right ADS (unilateral silencing). Illumination was delivered continuously on a subset of trials (25%) by opening and closing a shutter with a 5V TTL pulse.

In the inactivated trials, light was delivered during one of three different epochs. The first epoch was 2000ms long, starting on trial initiation and ending 500ms after the termination of the stimulus. This "whole-trial" inactivation was delivered to a cohort of 6 rats (10 hemispheres). The second and third epochs were 500ms long and spanned the first and second halves respectively of a 1s long auditory stimulus. These early and late inactivations were done in sessions separate from the whole trial inactivations in 5 rats (8 hemispheres; 3 of these rats also belonged to whole-trial inactivation cohort). In sessions with early and late inactivations, 50% of the trials had 1s long stimulus. Duration of the stimulus on rest of the trials was uniformly distributed between 0.2-1s. On half of these 1s long stimulus trials, light was delivered either during early or late epoch, and the other half served as control trials.

We measured behavioral bias resulting from optogenetic inactivation following methods used in.[14,16] First, we binned the trials on the basis of stimulus strength and for each of the 10 binned stimulus strengths, we computed the percentage of trials during which the rats chose the side ipsilateral to the optically stimulated side on control and inactivation trials. The mean difference across stimulus strengths in this 'go ipsi' rate between inactivation and control trials measured the bias caused by optogenetic inactivation. A positive bias represents an increase in ipsilateral choices upon laser stimulation. Confidence intervals and statistical comparisons for this bias parameter were calculated using nonparametric bootstrap procedures.

Multi-region recurrent neural network model

We trained an RNN with biological constraints on connectivity to perform the evidence accumulation task from Brunton et al. (2013).[32] The RNN is composed of standard firing-rate units ($N = 300$) with a rectifying nonlinearity and $\tau = 30ms$ as the time constant of decay of network units. RNN dynamics are governed by the equations:

$$\tau\dot{\boldsymbol{x}} = -\boldsymbol{x} + W^{rec}\boldsymbol{r} + W^{in}\boldsymbol{u} + \sqrt{2\tau\sigma^2_{rec}}\xi \quad \text{where} \quad \mathbf{r} = [\boldsymbol{x}]_+$$

where the variable $\boldsymbol{x}(t)$ is an N dimensional vector containing the activation of neural units in the network. The activations are mapped to the corresponding firing rates $\boldsymbol{r}(t)$ by passing $\boldsymbol{x}$ through the threshold linear nonlinearity which maps individual activations to positive firing rates: x_i if $x_i > 0$ and 0 otherwise for $i = 1,\ldots,N$. Recurrent connection weights in the network are given by the $N \times N$ matrix W^{rec} and the connection weights from the inputs to network units are given by the $N \times N^{in}$ matrix W^{in}. The network receives 2-dimensional ($N^{in} = 2$) time-varying inputs, $\boldsymbol{u}(t) = [s_L(t) \quad s_R(t)]^\top$, which represent the left clicks and right clicks respectively "(Figures S9B–S9D)". Noise intrinsic to the network is represented by $N \times 1$ dimensional $\xi(t)$, independent white noise with zero mean and unit variance. The network output $z(t)$ is read out linearly, as a weighted sum of the firing rates of network units, with weights W^{out} and a bias b:

$$z = W^{out}\boldsymbol{r} + b$$

Network inputs, perturbations and targets

The network received 2 inputs $\boldsymbol{u}(t) = [s_L(t) \quad s_R(t)]^\top$ at every time step. The two inputs correspond to the randomly timed discrete auditory evidence that is presented to the rats during the task, from speakers to their left and right sides respectively. These inputs are zero except when a left/right click occurrs. The click times were sampled from Poisson processes, with Poisson rates for left/right inputs sampled independently on each trial, from the same distribution as the one used for rats. At the times when left clicks occurred the average input was equal to -1, however we added independent Gaussian noise η with zero mean and variance $\sigma^2_S = 2$ to each click, hence the input for any given left click was $-(1+\eta)$. σ^2_S represents the variance of sensory noise accompanying each sensory input, and was set to match the average sensory noise variance measured in rats performing this task.[32] Similarly, when a right click occurred, the magnitude of input s_R was set to $1+\eta$. The overall timing of these stimulus inputs mirrored that of the real task. On each trial, the stimulus duration i.e. the time during which clicks played (in s) was sampled uniformly from the interval $[0.2, 1.0]$ and was preceded by a variable delay such that the stimulus always ended at 1.5s from the trial start.

Choice targets for each trial were obtained by passing stimulus, history and lapsing inputs to the behavioral model from Brunton et al. (2013).[32] If the target choice was towards right (left), then the target was set to $+1$ (-1) in the 200ms following termination of stimulus, i.e. from 1.5 to 1.7s. On a quarter of trials, we turned down the gain of certain recurrent weights in W^{rec} by a factor of 0.1 so as to simulate unilateral inactivations of FOF, ADS or FOF inputs to ADS during either the first or second half of a 1s long stimulus. On these trials, if an ipsilateral bias was observed experimentally then the choice targets were not obtained from the model but instead always set to the ipsilateral side.

Network connectivity and training

We engineered the recurrent weight matrix W^{rec} so as to capture the major anatomical features of the FOF-ADS circuitry. The network included two modules for the two brain hemispheres (150 units each) and each module in turn consisted of three submodules with 50 all-to-all connected neural units - representing FOF, ADS and the multi-synaptic relay between ADS and FOF (i.e. other basal ganglia nuclei, thalamus, SC, etc.). All units in the network followed Dale's law. In other words, all recurrent projections from a unit were either positive or negative but never both. We refer to the units which send out positive recurrent projections as excitatory units and the ones with negative projections as inhibitory units. Within each submodule, we set the ratio of excitatory to inhibitory units to match the known distribution of E-I neurons in these brain regions. Therefore, FOF submodules had 20% inhibitory units and ADS had 100%. For the relay submodule we set the ratio to 30%.

In FOF submodules, inhibitory units projected locally within the submodule, and excitatory units projected to other submodules in a manner aiming to mimic the long-range projections of cortical pyramidal neurons.[112] About 30% of the excitatory neurons from an FOF submodule followed the pattern of projections observed in cortical IT neurons - they sent projections to the other FOF submodule and the two ADS submodules. Another 30% followed projection patters of cortical PT neurons - they sent projects to the ipsilateral ADS submodule. The remaining FOF excitatory units only had local recurrent connections. The ADS submodules projected ipsilaterally to the relay module, since we are not aware of any cross-hemipshere projections between the two hemispheres. We do not differentiate units in ADS as belonging to D1/D2 subtypes etc. Given the abstract nature of the relay submodule, we put minimal constraints on its recurrent connections and it was allowed to project bilaterally to the two FOFs and to the other relay submodule with both excitatory and inhibitory projections.

To initialize W^{rec}, we first sampled a matrix of values from a normal distribution with zero mean and variance $1/N$ while enforcing the connectivity patterns as detailed above. Then we balanced the average weights of excitatory and inhibitory inputs into each unit and set the spectral radius to 1.3.

All units in the FOF and ADS submodules received the inputs i.e. W^{in} projected only to FOF and ADS units, and outputs were read out from all units in the network. Input and outputs weight were initialized using samples from Glorot normal distributions.[113]

While training, to enforce the Dale constraints on connectivity we parameterized the recurrent weight matrix W^{rec} as a product of a non-negative matrix $W^{rec,+}$ and a diagonal matrix D of 1s and $-$ 1s identifying excitatory and inhibitory units i.e. $W^{rec} = W^{rec,+}D$. We maintained the sparse connectivity across submodules by further performing element-wise multiplication with a connectivity mask M^{rec}, such that $W^{rec} = M^{rec} \odot (W^{rec,+}D)$. We simulated the network's dynamics using Euler updates with $\Delta t = 5$ms and generated $240,000$ trials for this, with minibatches of 64 trials each. We optimized $W^{in}, W^{rec}, W^{out}, b$ and $x(0)$ to minimize the mean-squared error (MSE) between the network output and targets, using Adam[110] in TensorFlow.[114] The methods used for analysis of activity and perturbation response of the multi-region models were identical to that used for rat data.

QUANTIFICATION AND STATISTICAL ANALYSIS

The details of all statistical tests performed are described in the method details section above. The number of cells, animals, and sessions used for each analysis are reported in the text and figure captions. Error bars and shaded regions on plots reflect the standard error of the mean unless otherwise noted in figure captions.

Neuron, Volume *114*

Supplemental information

A multi-region recurrent circuit for evidence accumulation in rats

Diksha Gupta, Charles D. Kopec, Adrian G. Bondy, Thomas Z. Luo, Verity Elliott, and Carlos D. Brody

Supplementary Information

A multi-region recurrent circuit for evidence accumulation in rats

Diksha Gupta, Charles D. Kopec, Adrian G. Bondy, Thomas Z. Luo, Verity Elliott, Carlos D. Brody

List of Figures

1 Neural responses in FOF and ADS during the evidence accumulation task 2
2 Neural clustering reveals similar temporal activity profiles across FOF and ADS . 3
3 Similar stimulus, choice and history decoding in FOF and ADS 4
4 Contribution of neural clusters to accumulator tuning and decoding analyses . . . 5
5 Decoding decision variable from FOF and ADS population activity 6
6 Simulations used to validate the decision variable cross-correlation analysis. . . . 7
7 Reduced rank GLM captures the low-dimensional communication subspace . . . 8
8 Silencing FOF axon terminals in ADS: effect on movement and psychometric parameters 10
9 Training details of multi-region recurrent neural network model 11
10 RNN model recapitulates functional clustering 13
11 Network dynamics during 1st and 2nd half inactivations 14
12 Network outputs in response to in-silico perturbations 15

List of Tables

1 Summary of Neural Statistical Analysis 16

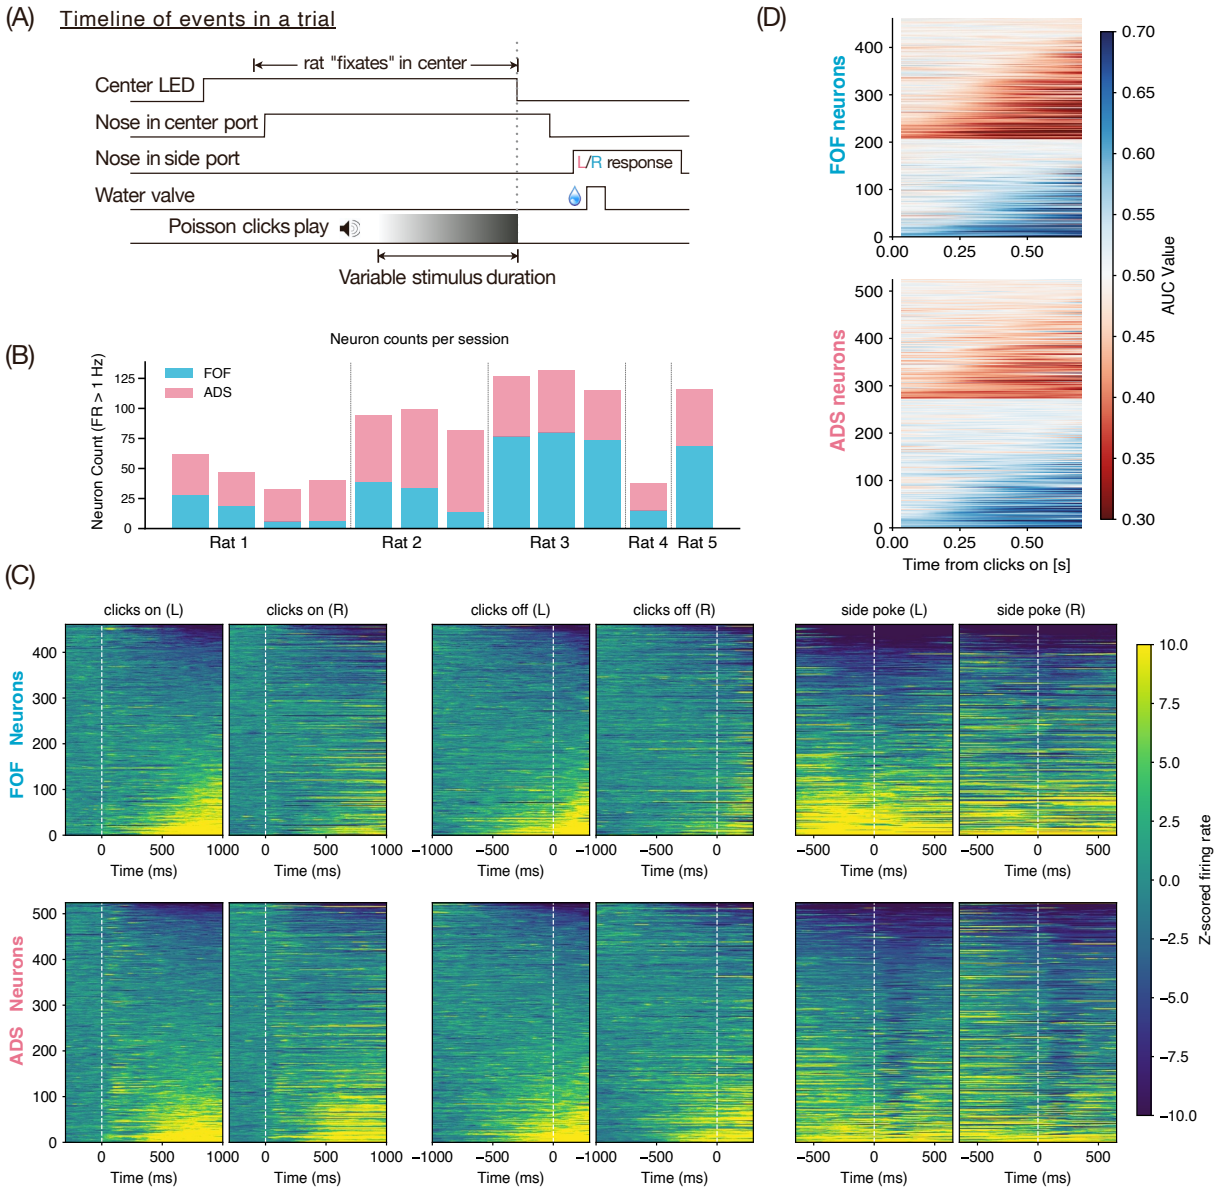

Supplementary Figure 1: **Neural responses in FOF and ADS during the evidence accumulation task (related to Figure 1) (A)** Timeline of events in a trial. Rats initiate trials by nose-poking into the center port. After a variable delay, Poisson clicks were delivered from left and right speakers; stimulus duration varied across trials. After stimulus offset, rats indicated their decision via a side poke to receive a reward. **(B)** Number of recorded neurons (with mean firing rate >1 Hz) in each of the 12 recording sessions from 5 rats from FOF (blue) and ADS (pink). **(C)** Trial-averaged, z-scored firing rates of all neurons with firing rate >1 Hz in FOF (top) and ADS (bottom), aligned to key task events (stimulus onset, stimulus offset, and side poke), shown separately for leftward (L) and rightward (R) trials. Neurons are sorted by their projection onto the first principal component of leftward-trial responses. Both FOF and ADS include neurons that ramp up, down or maintain their firing rates over the trial. **(D)** Heat maps showing area under the ROC curve (AUC) for discriminating between left and right choice trials during the stimulus period. Neurons are sorted by AUC value, with values >0.5 indicating preference for rightward choices (red) and <0.5 indicating preference for leftward choices (blue). Both regions show comparable distributions of side-selective neurons.

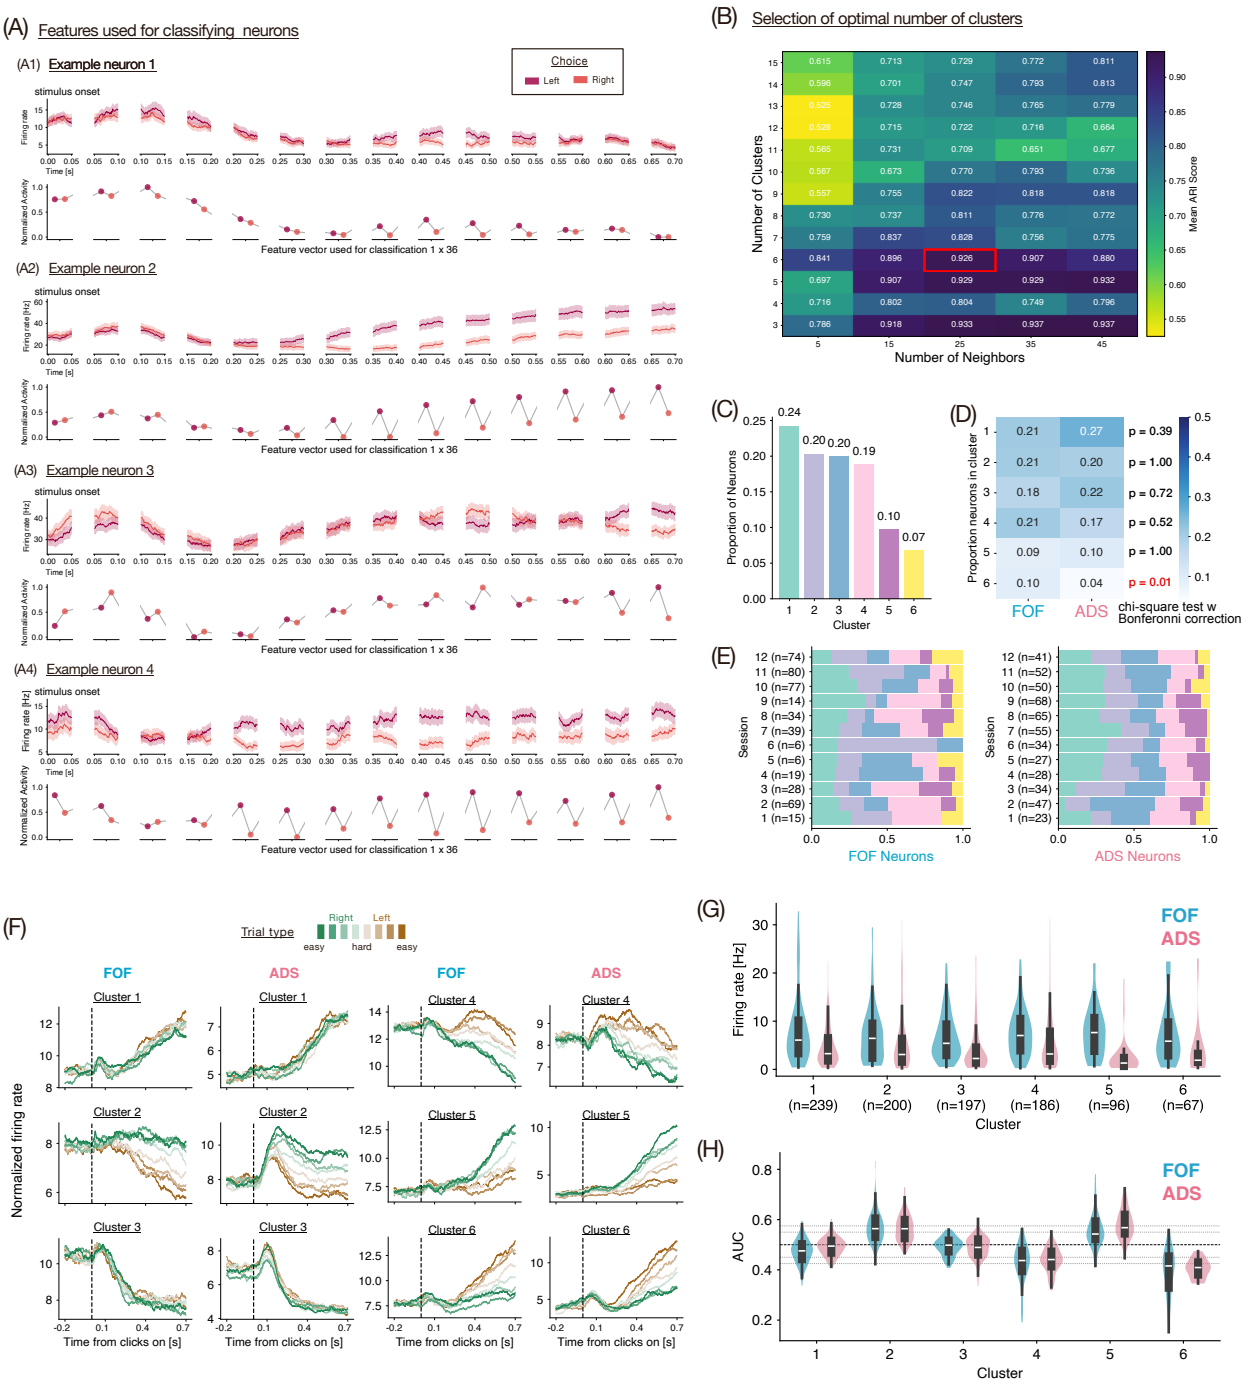

Supplementary Figure 2: **Neural clustering reveals similar temporal activity profiles across FOF and ADS (related to Figure 2) (A1-A4)** Example construction of the input feature vector from four single neurons. Top panels: PSTHs (50ms bins) aligned to stimulus onset, separately for left (maroon) and right-choice (red) trials. Bottom panels: corresponding 36-dimensional feature vectors obtained by concatenating across left/right choices, creating a feature vector per neuron. **(B)** Heatmap of mean Adjusted Rand Index (ARI) across 100 bootstrap iterations as a function of k-nearest-neighbors (x-axis) and number of clusters (y-axis). We selected 6 clusters to capture maximum diversity in the neural population while maintaining robust cluster separation (marked in red). **(C)** Number of neurons identified in each cluster across FOF and ADS ordered by prevalence **(D)** Relative abundance of the six clusters is similar in FOF and ADS. Numbers show the proportion; χ^2 tests with Bonferroni correction reveal no significant region difference except for cluster 6 (P = 0.01; all others P>0.39) **(E)** Fraction of neurons assigned to each cluster across sessions. Each row is a session. Colors correspond to cluster IDs from (C). The near-identical pattern across regions and sessions demonstrate that the cluster response profiles are not driven by a single animal or day. (Caption continued on next page)

Supplementary Figure 2: (Continued from previous page) **(F)** Cluster-averaged PSTHs (mean ± SEM) aligned to stimulus onset for each stimulus strength. Each cluster captures a distinct pattern of slow dynamics and evidence sensitivity. Clusters 1, 5 and 6 ramp up, whereas clusters 2–4 show decreasing profiles; Clusters 2, 4, 5, and 6 show robust tuning to evidence strength, while Clusters 1 and 3 do not. **(G)** Distributions of average firing rates within cluster. ADS neurons have lower firing rates across clusters. **(H)** Distributions of choice selectivity (AUC) within cluster. The dotted lines mark the thresholds used for population PSTH (0.05 + |auc - 0.5|; Fig 1F) and neural tuning curve analysis (0.075+ |auc - 0.5|; Fig 2A-B)

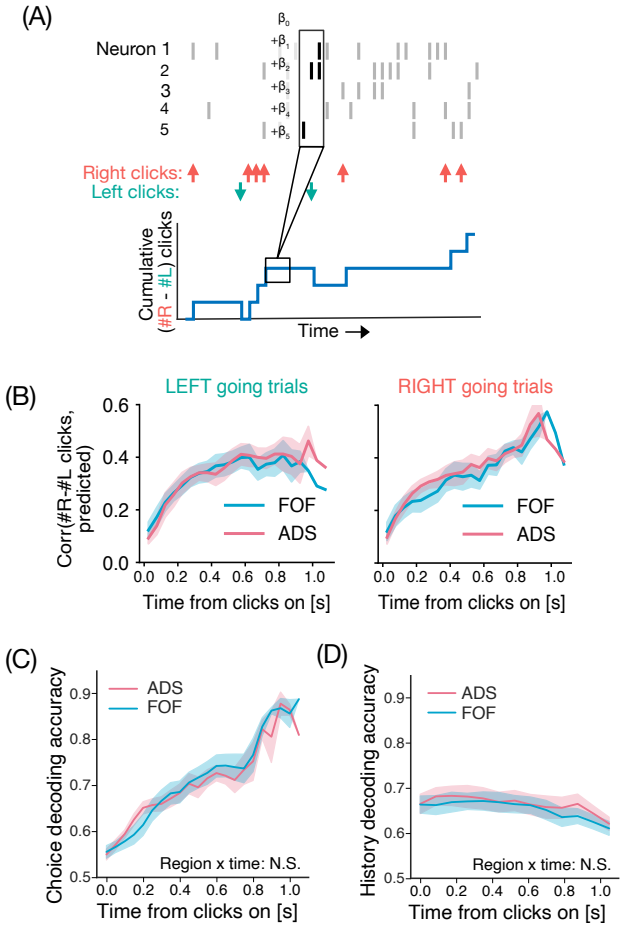

Supplementary Figure 3: **Similar strength and timecourse of stimulus, choice and history decoding in FOF and ADS (related to Figure 2) (A)** Schematic showing the linear decoder designed to predict the cumulative click difference in the number of right and left clicks at any given time during a trial using a linear combination of neural activity from FOF or ADS. We assumed a neural encoding lag of 100ms and regularized the linear weights on neurons with L1 penalty using cross-validation. **(B)** Stimulus decoding performance (mean ± sem) of the linear decoder on held-out time points as a function of time from stimulus onset (n = 12 sessions) performed separately for left (left panel) and right (right panel) going trials to control for choice. The two areas show high decoding performance which evolves with a similar timecourse ($P > 0.05$, two-way RM ANOVA). **(C)** Cross-validated choice decoding accuracy as a function of time from stimulus onset for FOF (aqua blue) and ADS (pink) neural activity (n = 12 sessions). Time courses and strength of choice decoding from the two populations did not differ significantly ($P > 0.05$ two-way RM ANOVA). Decoding was performed with an L2 regularized logistic decoder, while controlling for number of neurons from the two areas. **(D)** Same as C but decoding for choice from past trial (n = 12 sessions). No significant differences were observed between the two populations ($P > 0.05$, two-way RM ANOVA)

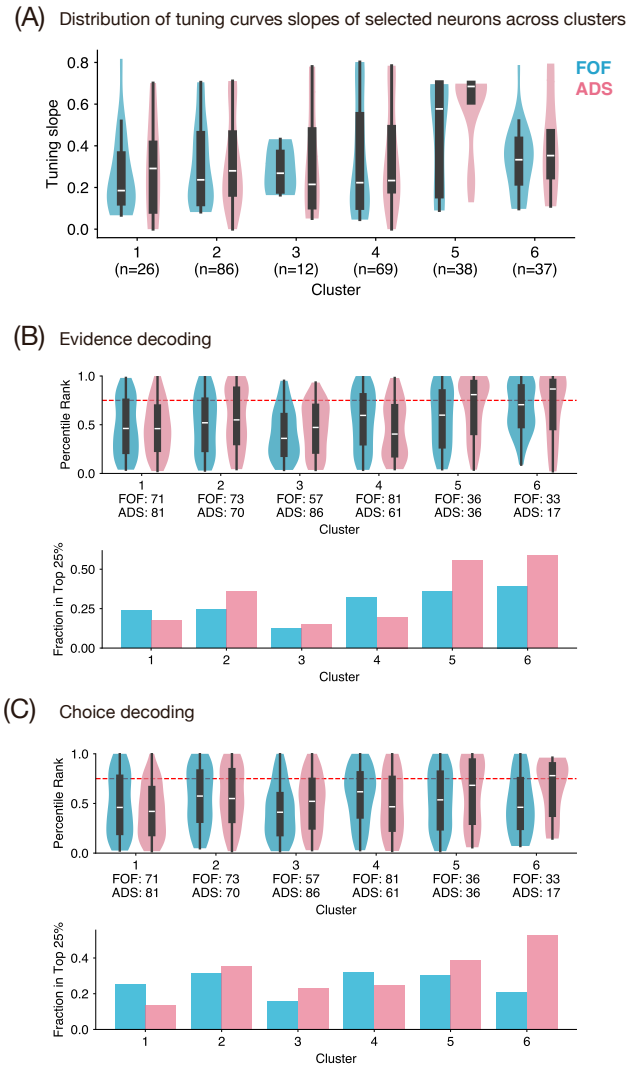

Supplementary Figure 4: **Contribution of neural clusters to accumulator tuning and decoding analyses (related to Figure 2) (A)** Violin plots showing the distribution of accumulator tuning curve slopes for neurons from each of the 6 identified clusters, separated by brain region (FOF in blue, ADS in pink). Numbers below indicate total neurons per cluster included in the analysis. All clusters contribute to the tuning curve analysis, with evidence-selective clusters (2,4,5,6) showing greater representation. **(B)** (Top) Violin plots showing the percentile rank of decoding weights assigned to neurons from each cluster during evidence decoding, with red dashed line indicating the 75th percentile. (Bottom) Bar plots showing the fraction of neurons from each cluster that fall within the top 25% of decoding weights. Evidence-selective clusters (2, 4, 5, 6) receive slightly higher decoding weights, but neurons from all clusters contribute to evidence decoding. **(C)** Same as (B) but for choice decoding. All clusters contribute broadly to choice decoding . The broad participation across clusters suggests that decision variable cross-correlations reflect population-level recurrence rather than contributions from specific functional cell types.

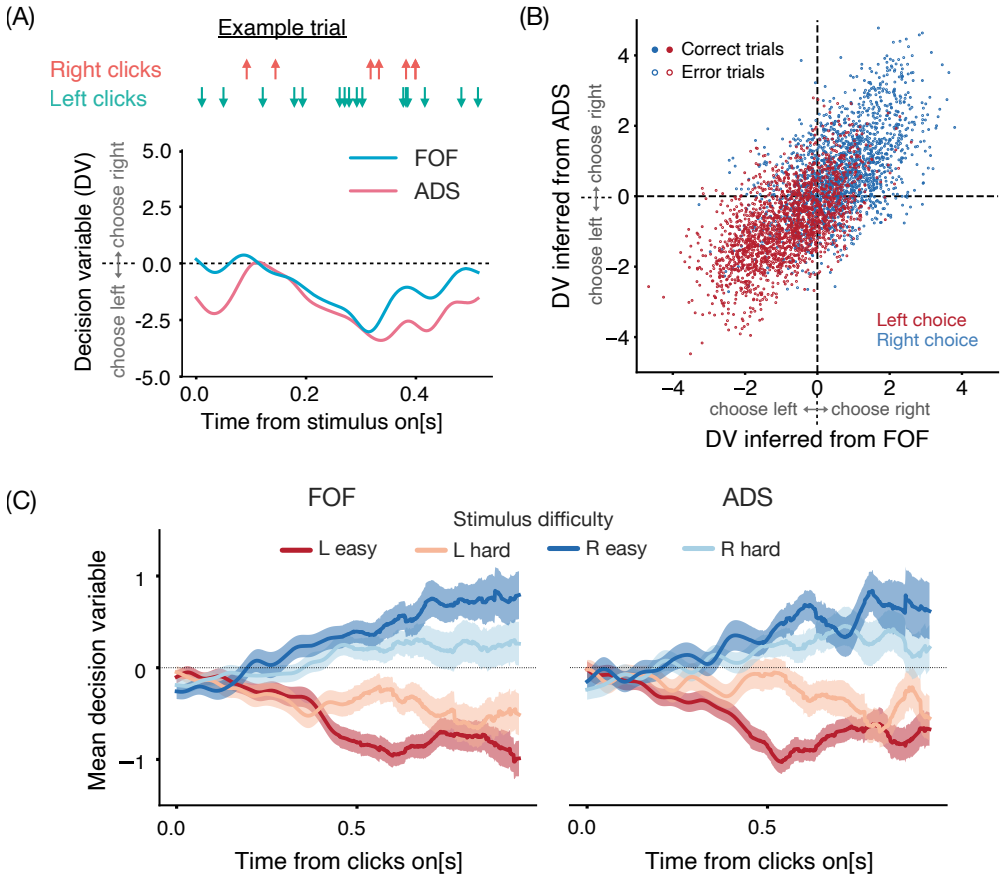

Supplementary Figure 5: **Decoding decision variable from FOF and ADS population activity (related to Figure 2) (A)** Example trial showing the trajectory of the decision variable (DV) decoded from the simultaneous population activity of FOF and ADS (blue, pink line respectively), using the logistic decoder (see Methods). The DV represents the log odds ratio of making one choice over another, with higher positive (negative) values indicating a stronger likelihood of rightward (leftward) choices. Both regions have comparable DV time-courses, and seem to be influenced by the sequence of right and left clicks (coral, green arrows respectively). **(B)** Scatter plot comparing DV values inferred from FOF (x-axis) and ADS (y-axis) activity at the same time-points from an example session. Dots represent time-points across trials, with filled (unfilled) circles indicating correct (incorrect) trials. Points are color-coded according to the actual choice of the animal, with blue (red) indicating rightward (leftward) choices. The DV inferred from both areas shows good correspondence with each other and the true choice. **(C)** Mean DV trajectories from FOF (left) and ADS (right) over time from an example session and separated by stimulus difficulty (red - easy leftward stimulus, light red - hard leftward, blue - easy rightward, light blue - hard rightward). Shaded regions indicate S.E.M. Both regions show similar DV time-courses and similar dependence on stimulus difficulty.

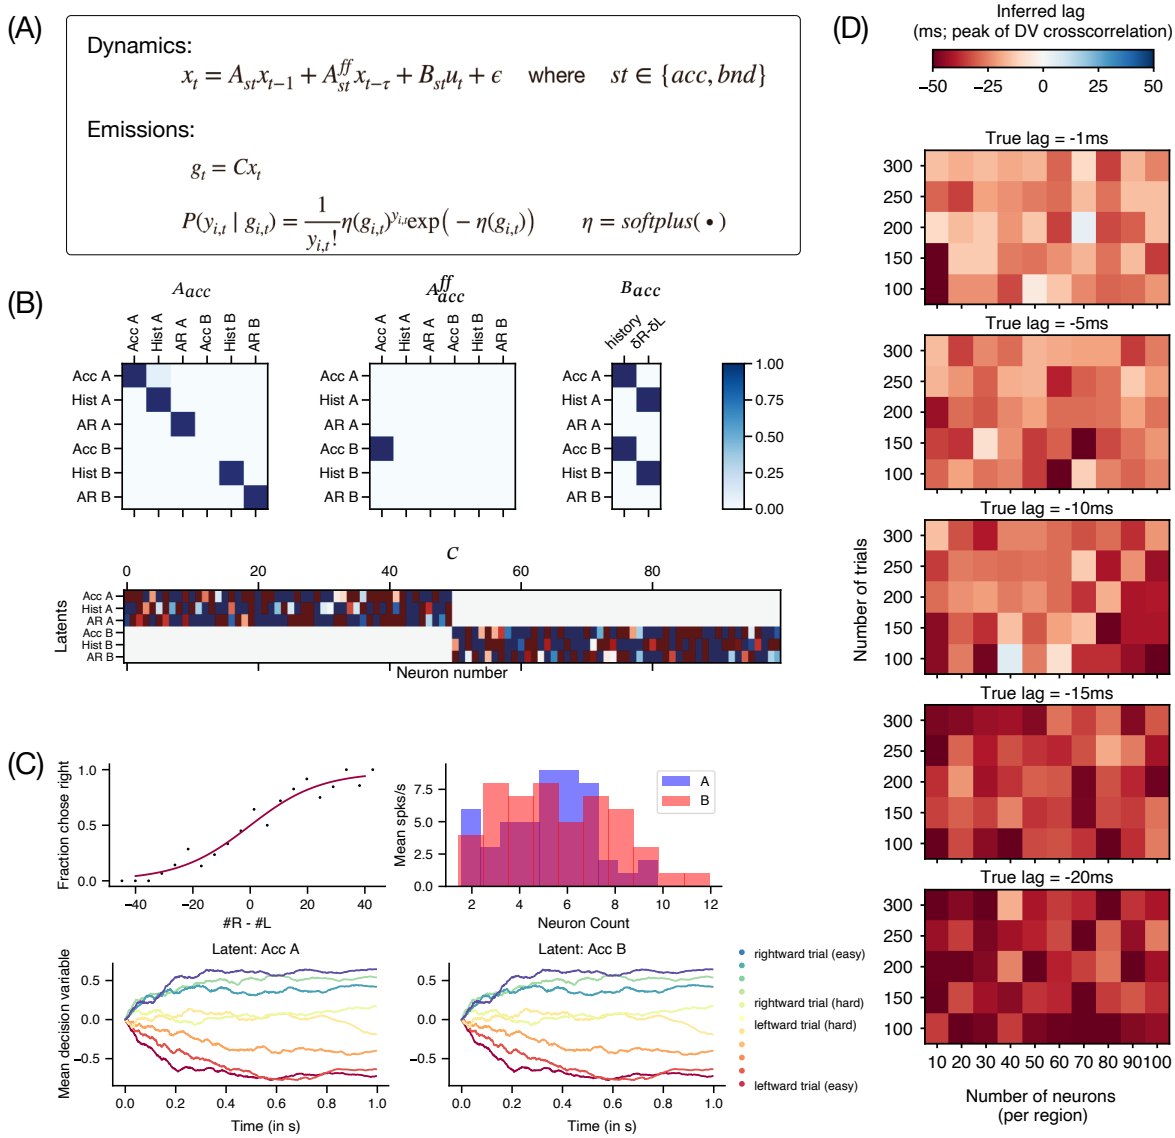

Supplementary Figure 6: **Simulations used to validate the decision variable cross-correlation analysis (related to Figure 2) (A)** Equations used to simulate two regions with feedforward flow of accumulator dynamics. (Top) The low-dimensional latents in the two regions x_t represent accumulator variables (Acc), history (Hist) and autoregressive (AR) terms that evolve according to linear dynamics with dynamics A_{st} that are different during accumulation ($st = acc$) and bound-hitting ($st = bnd$). During accumulation, one region's accumulator dynamics (region B) are influenced by feed-forward inputs from the other region's accumulator (region A) according to A^{ff}_{acc} with delay τ. Both regions receive external inputs u_t (including stimulus clicks and history) according to B_{st} and Gaussian noise ϵ. (Bottom) The mean firing rates g_t of neurons in the two regions reflect a high-dimensional projection C of the aforementioned latents x_t. These are then transformed through a softplus function (to enforce positive firing rates) into Poisson distributed spikes y_t. **(B)** (Top) Schematized A_{acc}, A^{ff}_{acc} and B_acc (Bottom) Linear weights C that transform latents (y-axis) into mean firing rates of individual neurons (x-axis). **(C)** (Top left) Fraction chose right based on thresholded values of the accumulator, showing a monotonic dependence on the stimulus. (Top right) Firing rate distributions in the two regions showing substantial heterogeneity and overlap. (Bottom) Mean decision variables across trials of the same stimulus difficulty for the accumulator latents from the two regions, showing a separation by difficulty. **(D)** (Top to bottom) Inferred lag, measured as the peak of the cross-correlation between decision variables from the two areas, shown as a function of the number of neurons (x-axis) and trials (y-axis), for increasing true lags between the regions. Even though the exact value of lag is not recovered for low trial and/or neuron counts, the direction of flow can be accurately inferred. Increasing color intensities show that the measured peak lag is sensitive to the true lag values.

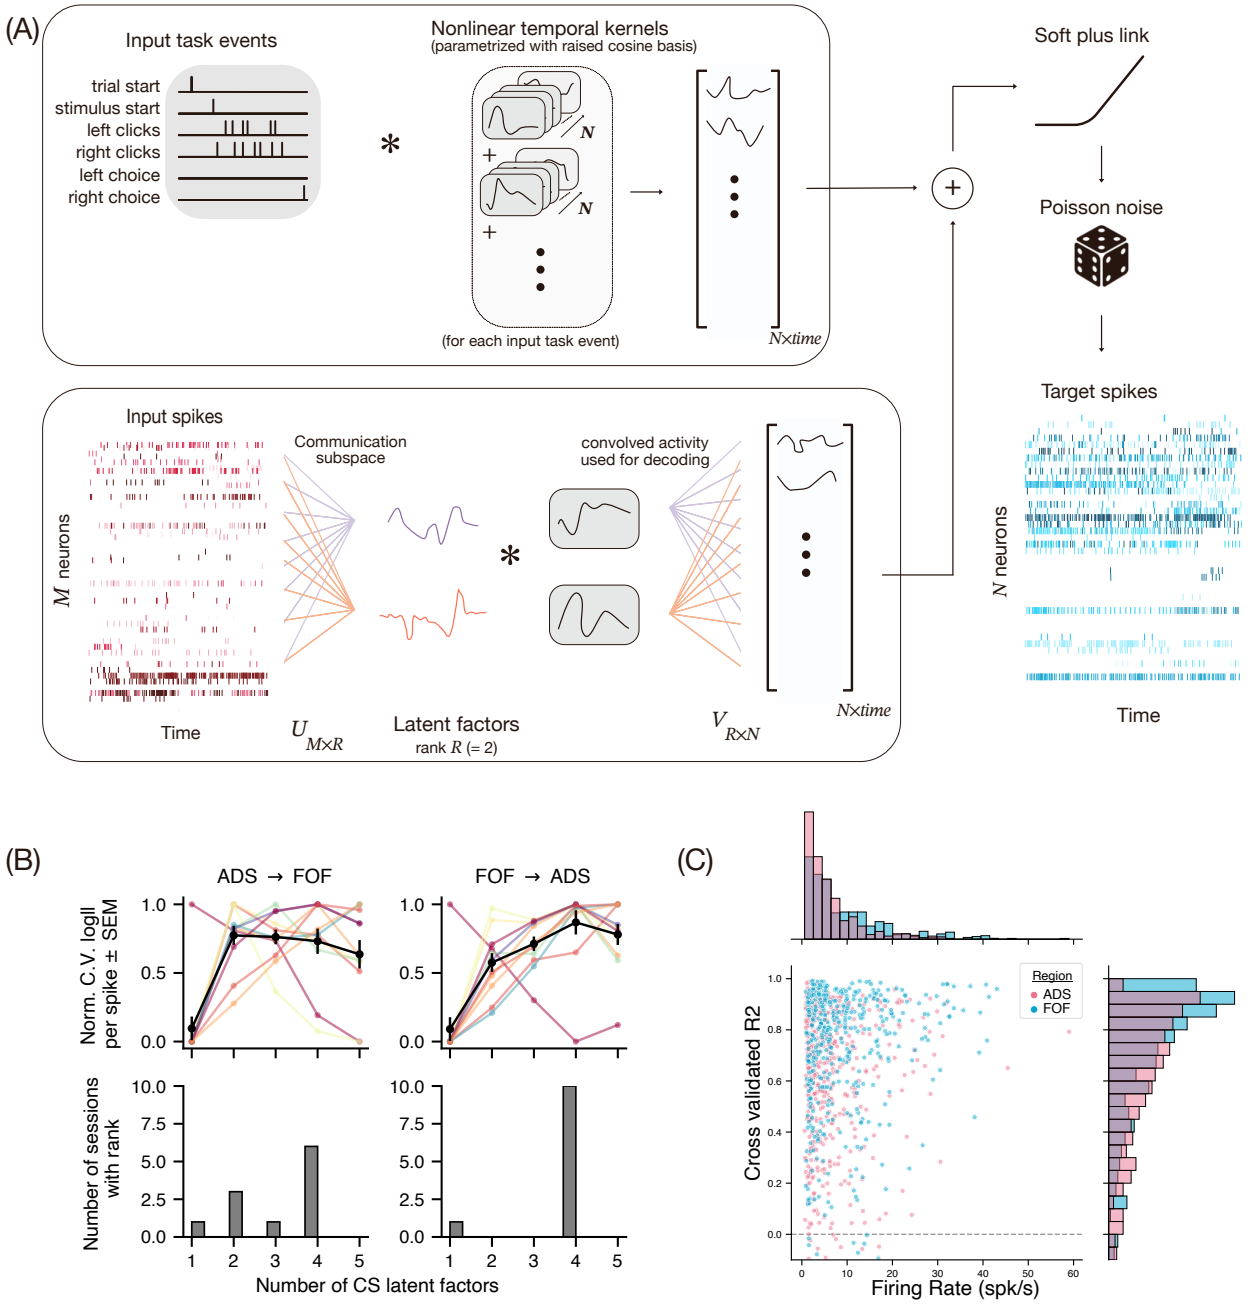

Supplementary Figure 7: **Reduced rank GLM captures the low-dimensional communication subspace between FOF and ADS (related to Figure 2) (A)** Schematic of the reduced rank GLM approach. The target variables are sequences of spikes for each of N neurons in the output region (blue rasters), which are modelled as a generalized linear model with a soft plus link function and poisson noise distribution. The predictors consist of task events (top box, left panel) and sequences of spikes from M neurons in the input region (bottom box, red rasters). Task events are convolved with nonlinear temporal kernels parameterized with raised cosine bases - one per task event per output neuron - allowing them to have temporally extended effects (top box, grey boxes). Spikes from the M neurons in input region are first projected into a reduced R-dimensional communication subspace using learnt weights $U_{M\times R}$, yielding a small number of latent factors (bottom box, colored curves). These latent factors are then convolved with temporal kernels (grey boxes) and influence the N neurons in the output region through learnt output weights $V_{R\times N}$. **(B)** Required dimensionality of the communication subspace. (Top) Normalised cross-validated log likelihood per spike, as a function of the dimensionality of the communication subspace. Colored curves individual sessions, black represents mean across sessions. (Bottom) Distribution of the optimal dimensionality of the communication subspace across sessions. (Caption continued on next page).

Supplementary Figure 7: (Continued from previous page) Based on both these measures, the smallest dimension required (≤ 4) was used for each session, with most sessions warranting a 4-dimensional communication subspace. **(C)** Cross validated r-squared values across sessions of the reduced rank GLM (ADS = 0.64 $\pm$ 0.01, FOF = 0.73 $\pm$ 0.01 mean $\pm$ sem), plotted as a function of the session average firing rates - colors represent different directions of communication, with ADS (pink) or FOF (blue) as output regions. In ADS (FOF), 0.12 (0.09) fraction of cells had R2 < 0. These cells were excluded while computing the mean R2.

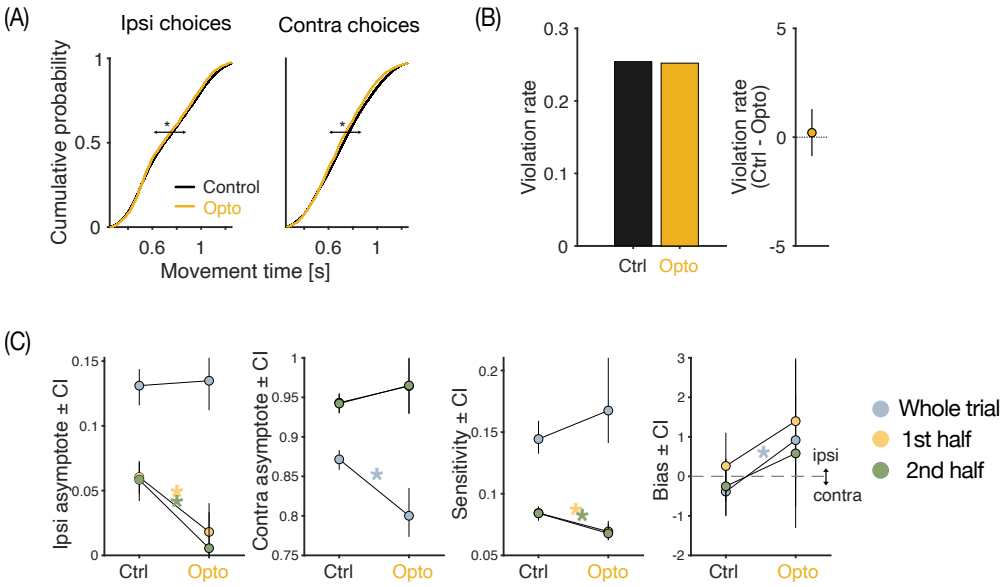

Supplementary Figure 8: **Silencing FOF axon terminals in ADS: effect on movement and psychometric parameters (related to Figure 3) (A)** Silencing reduces movement times. Cumulative distribution of movement times across (n = 6) rats on control (black) and whole-trial inactivation (orange) trials. Movement time is defined as the time taken by the rats to leave the center port and enter one of the two side ports to report their choice. Silencing of FOF axon terminals in ADS reduced movement times for contraversive choices (right, median: 0.725 ± 0.378 vs 0.698 ± 0.352, median ± IQR; P = 0.001, Mann-Whitney U test) as well as ipsiversive choices (left, median: 0.700 ± 0.441 vs 0.689 ± 0.422, median ± IQR; P = 0.05, Mann-Whitney U test). **(B)** Silencing doesn't affect rate of trial completion (left) Mean violation rates observed across rats (n=6) on control (black) and whole-trial inactivation (orange) trials. Rats are expected to "fixate" at the center port for 1.5s from trial initiation, failure to do so aborts the trial. Violation rates measure the rate of premature exits from the center port. An increased violation rate might reflect an inability to complete trials due to gross motor/cognitive impairment induced by laser. (right) Mean violation rates between control and opto trial were not different (P = 0.58, non-parametric bootstrap test). **(C)** Effects of optogenetic silencing during different trial epochs on parameters of the psychometric curve. Data shown for whole trial (gray), first half (yellow), and second half (green) inactivation conditions. (Left) Asymptote (lapse rate) parameter on the ipsilateral side (mean $\pm$ confidence interval). Whole trial inactivation produced no significant change (Control: 0.131, Opto: 0.135, $p > 0.05$), whereas first-half and second-half inactivations significantly reduced ipsilateral lapse rates (first half: -0.041, $p < 0.05$; second half: -0.051, $p < 0.05$). (Middle-left) Contralateral asymptote parameter $\pm$ confidence intervals. Whole trial inactivation produced a significant decrease (Control: 0.871, Opto: 0.800, difference: -0.071, $p < 0.05$), while epoch-specific inactivations showed no significant effects. (Middle-right) Sensitivity parameter $\pm$ confidence intervals. Whole trial inactivation showed no significant change (difference: $+0.023$, $p > 0.05$), while both first half and second half inactivations significantly reduced sensitivity (first half: -0.015, $p < 0.05$; second half: -0.017, $p < 0.05$). (Right) Bias parameter $\pm$ confidence intervals, showed significant positive bias toward ipsilateral choices during whole trial inactivation ($+1.354$, $p < 0.05$), with non-significant trends in the same direction for epoch-specific inactivations. Error bars represent 95% bootstrap confidence intervals. The observed differences in effects between whole trial and sub-trial inactivations (increased sensitivity with reduced contra asymptote in whole trial condition) are consistent with known parameter tradeoffs in psychometric curve fitting.

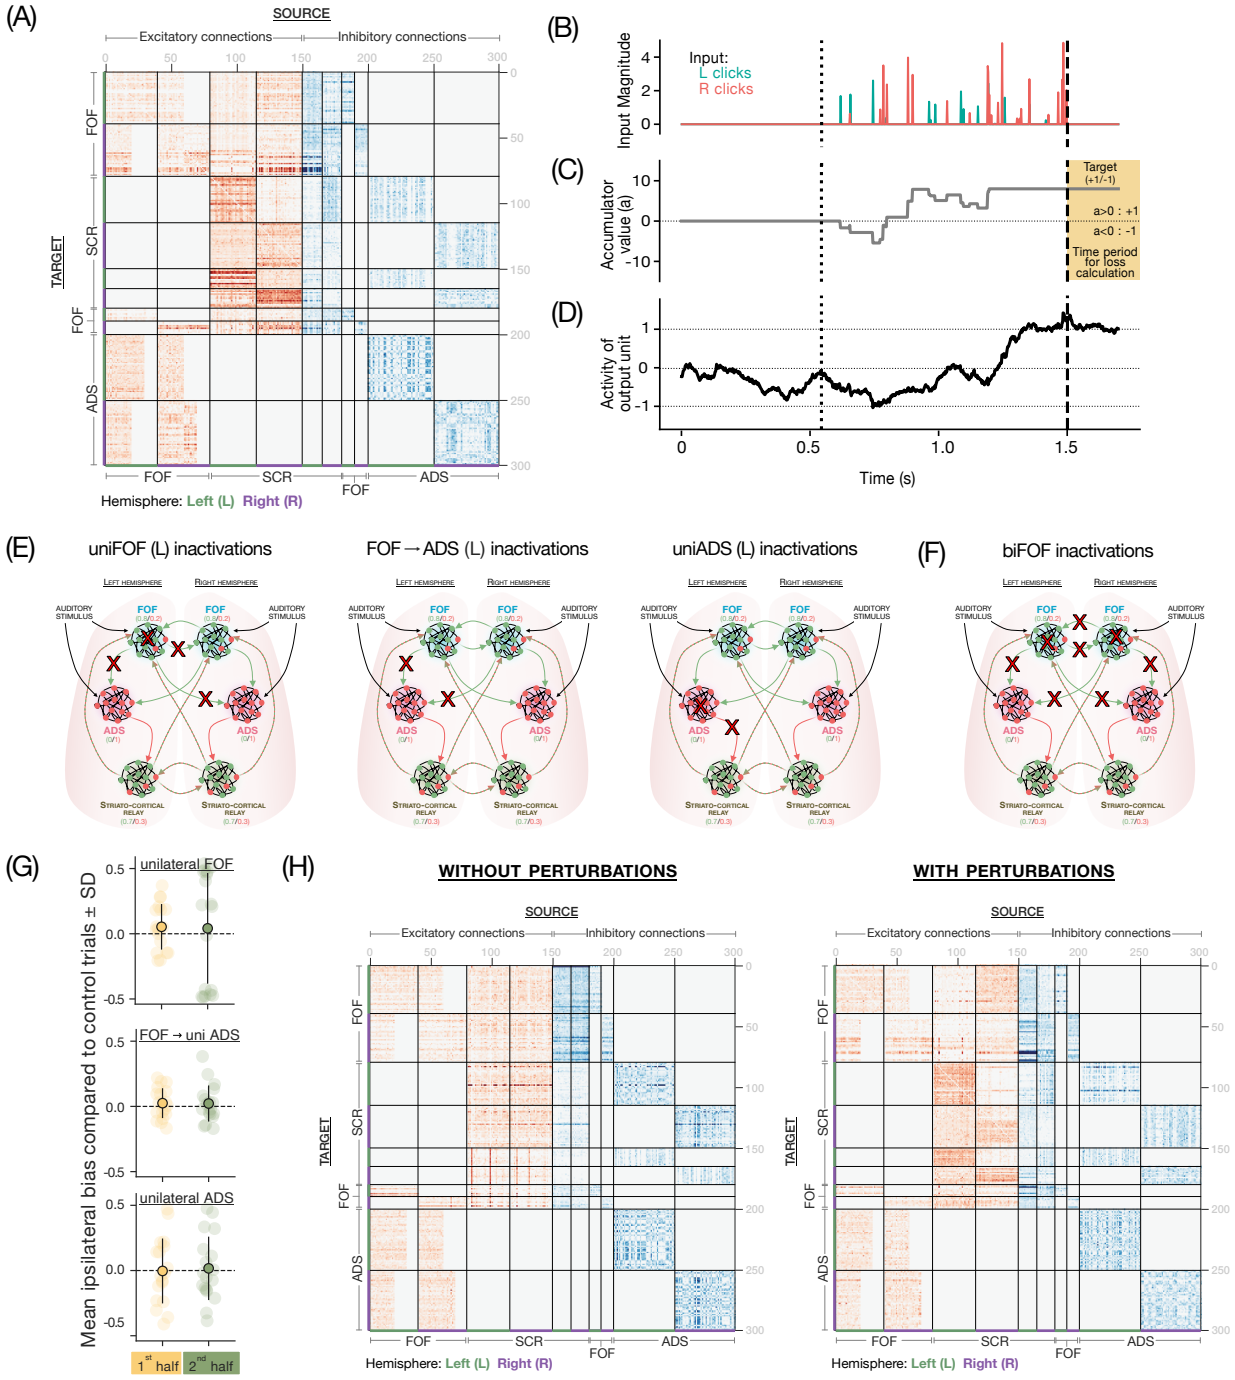

Supplementary Figure 9: **Training details of multi-region recurrent neural network model (related to Figure 4)** **(A)** Connectivity diagram of recurrent neural network, showing modular structure and following Dale's law. Excitatory connections are shown in red, and inhibitory in blue, with lines demarcating individual modules in a hemisphere. **(B)** Example trace of stimulus inputs to the network, showing leftward and rightward clicks with variable timing and magnitudes. **(C)** Accumulator value computed from stimulus trace above. The final thresholded accumulator value acts as the training target for network outputs on this trial. **(D)** Example activity of a network output unit that has successfully matched the target output on this trial. **(E)** Schematics showing all activity and projections that were silenced in unilateral inactivation experiments during training. (Left to Right): unilateral FOF inactivations, FOF → ADS inactivations, unilateral ADS inactivations. Red crosses indicate inactivated elements. **(F)** Same as E, but for bilateral FOF inactivations. **(G)** Networks trained without perturbation constraints produce variable and unrealistic inactivation effects, including contralateral biases and large variability across network instances. This reflects under-constrained optimization from training only on control behavior. (Caption continued on next page)

Supplementary Figure 9: (Continued from previous page) Adding perturbation data guides networks toward empirically consistent, lateralized organization. **(H)** Comparison of connectivity patterns from networks trained without (left) and with perturbation constraints (right). The two weight matrices show that the most prominent difference occurs in connections to and from the striato-cortical relay (SCR). Networks trained with perturbation constraints develop stronger and more structured connectivity involving the SCR, across both excitatory and inhibitory pathways, suggesting increased reliance on recurrent striato-cortical loops.

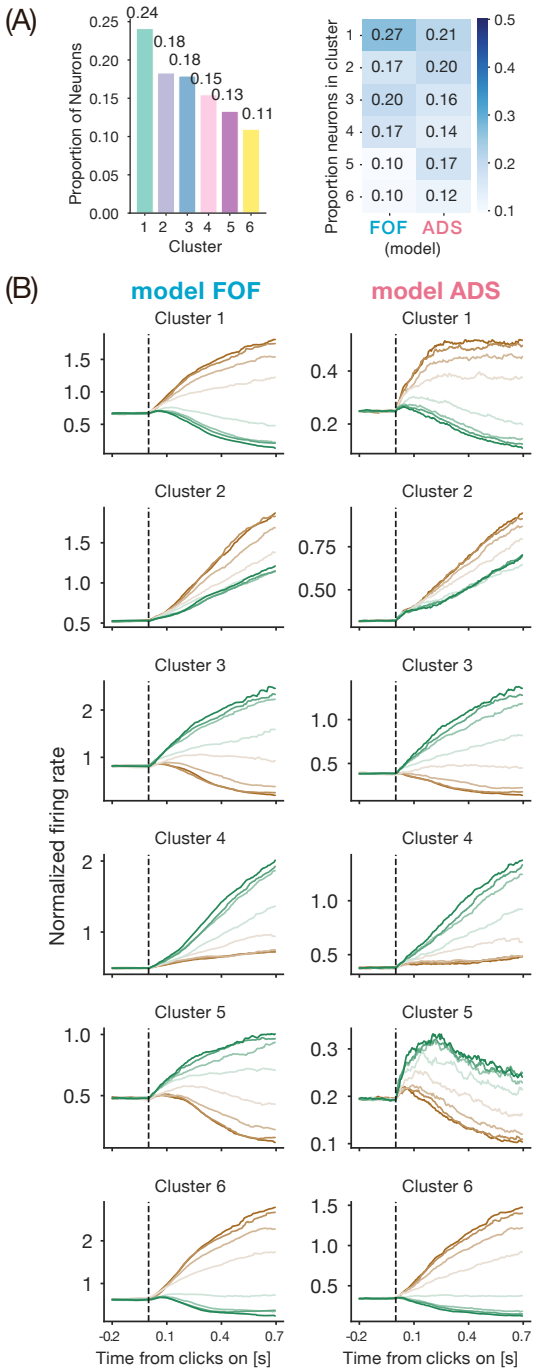

Supplementary Figure 10: **RNN model recapitulates functional clustering but lacks non-selective units and slow dynamics (related to Figure 4) (A)** (Left) Overall distribution of the 6 identified clusters across all model units. (Right) Heatmap showing the proportion of FOF and ADS model units in each cluster. Cluster proportions were broadly similar across model FOF and ADS. See Supp Fig 2 and *Neural clustering (Methods)* for additional details. **(B)** Stimulus onset aligned PSTHs for each of the 6 clusters, shown separately for model FOF (left column) and model ADS (right column). Green traces represent units preferring rightward choices, orange traces represent units preferring leftward choices, with lighter shades showing trials with low stimulus strength and darker shades trial with high evidence strength. Dashed vertical line indicates stimulus onset. Cluster responses largely mirror neural data (Supp Fig 2F), with two major differences: no non-selective clusters were identified (unlike neural clusters 1 and 3), and slow ramping dynamics prominent in neural data were absent. These differences likely reflect differing objectives of RNN and the brain. While the brain coordinates multiple functions—movement preparation, fixation, and processing broader contextual information—our RNN was optimized to model task performance and perturbation robustness.

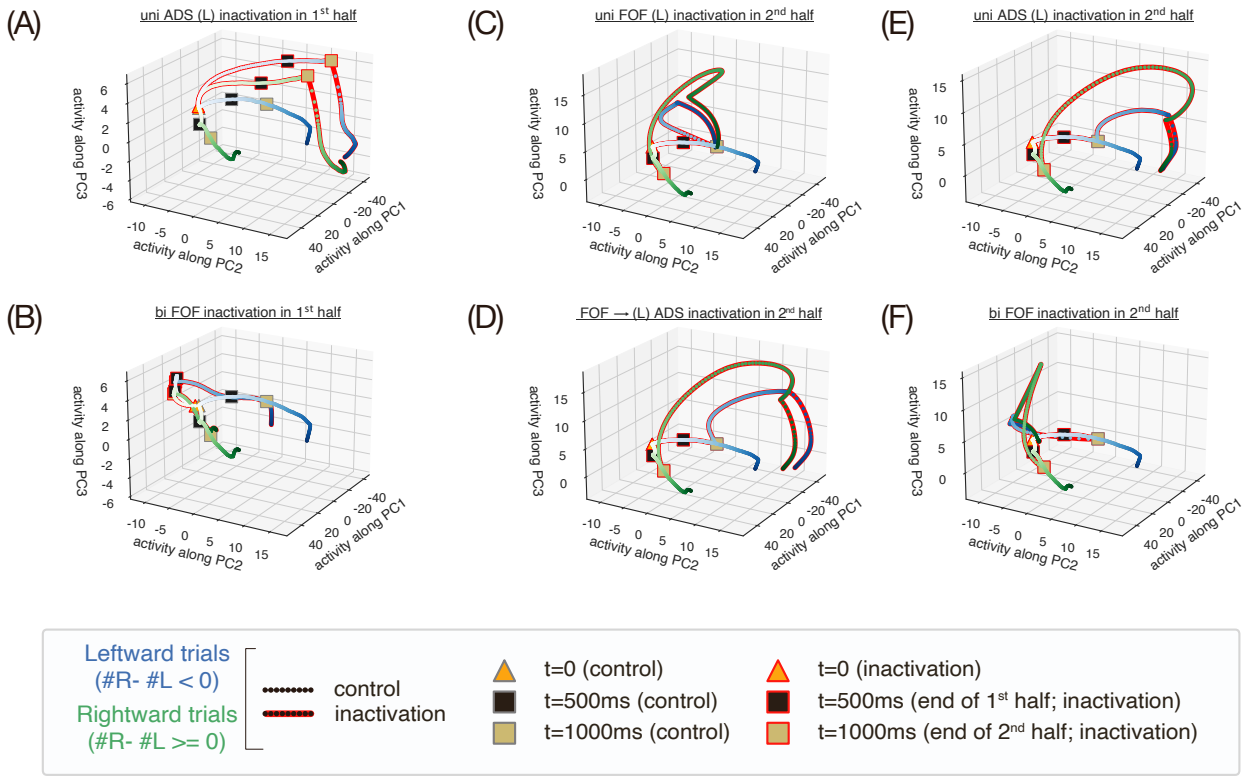

Supplementary Figure 11: **Network dynamics during 1st and 2nd half inactivations (related to Figure 5)** Projections of network activity onto first 3 principal components showing trajectories on rightward (green) and leftward (blue) trials. Red edges represent inactivation trials, and black/yellow boxes represent end of 1st/2nd half respectively. Unilateral inactivation of ADS in the 1st half (A) leads to rightwards trajectories diverging from control and incorrectly driving leftward choices, while bilateral inactivation of FOF in the 1st half (B) recovers towards correct choices subsequently. Unilateral inactivation of all regions in the second half (FOF (C), FOF→ADS projections (D), ADS (E) or bilateral inactivations of FOF (F)) drives incorrect choices and shows no recovery.

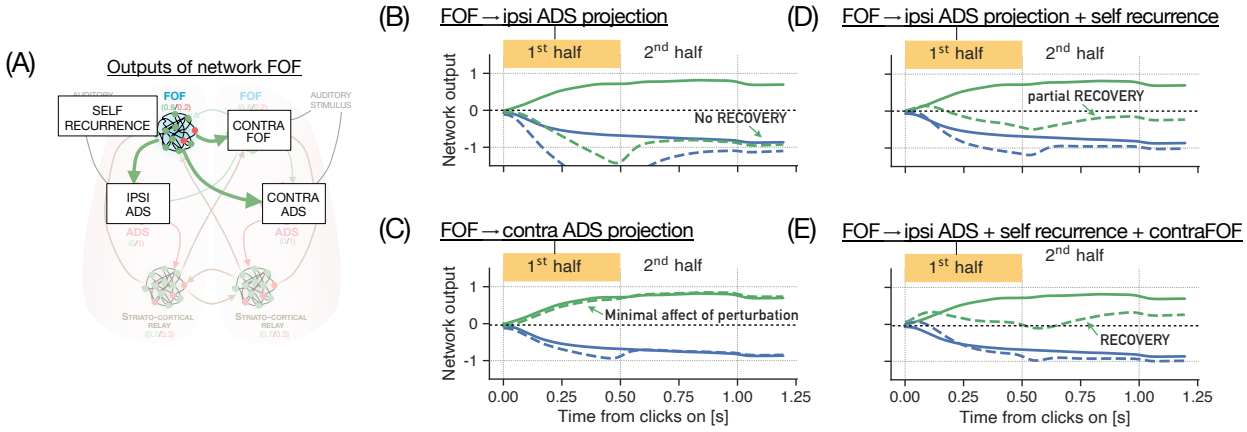

Supplementary Figure 12: **Network outputs in response to in-silico perturbations (related to Figure 6)** **(A)** Schematic of model FOF output projections from a single hemisphere, targeted for in-silico perturbations in the 1st half. (**B**) Perturbations of left FOF's projections to ipsilateral ADS led to network outputs diverging from control on rightward inactivation trials (dotted green) with no recovery, (**C**) while perturbations of the projections to contralateral ADS had no inactivation effect to begin with - suggesting that the ipsilateral ADS projections were largely responsible for inactivation effects. (**D**) Additionally perturbing the self recurrence projection or the contralateral FOF projection (not shown) invites partial recovery, (**E**) while perturbing both invites nearly complete recovery, with outputs crossing the decision threshold (y=0 line).

Supplementary Table

Table 1: Summary of Neural Statistical Analysis (related to Figure 2)

Analysis Type	**Figure**	**Neurons/Sessions as Independent Units**	**Rats as Independent Units**	**Significant?**
Accumulator Tuning Curves	Fig 2B	Mann-Whitney U test, $P = 0.13$ ($n = 144$ FOF, 130 ADS neurons)	Mann-Whitney U test, $P = 0.84$ ($n = 5$ rats)	No
Stimulus Decoding Performance	Fig 2C	Two-way RM ANOVA, Region: $P = 0.93$, Region×Time: $P = 0.36$ ($n = 12$ sessions)	Two-way RM ANOVA, Region: $P = 0.40$, Region×Time: $P = 0.58$ ($n = 5$ rats)	No
Choice Decoding Accuracy	Supp Fig 1C	Two-way RM ANOVA, Region: $P = 0.90$, Region×Time: $P = 0.54$ ($n = 12$ sessions)	Two-way RM ANOVA, Region: $P = 0.44$, Region×Time: $P = 0.49$ ($n = 5$ rats)	No
History Decoding Accuracy	Supp Fig 1D	Two-way RM ANOVA, Region: $P = 0.13$, Region×Time: $P = 0.37$ ($n = 12$ sessions)	Two-way RM ANOVA, Region: $P = 0.29$, Region×Time: $P = 0.64$ ($n = 5$ rats)	No
Decision Variable Cross-Correlation Lags	Fig 2D	One-sided t-test, $P = 0.47$ ($n = 12$ sessions)	One-sided t-test, $P = 0.74$ ($n = 5$ rats)	No
Communication Subspace Decoding	Fig 2E	Two-way RM ANOVA, Region: $P = 0.09$, Region×Time: $P = 0.65$ ($n = 12$ sessions)	Two-way RM ANOVA, Region: $P = 0.25$, Region×Time: $P = 0.86$ ($n = 5$ rats)	No

* RM ANOVA = Repeated Measures ANOVA

** All statistical tests were performed at $\alpha = 0.05$ significance level