

A multi-region recurrent circuit for evidence accumulation in rats

Diksha Gupta^{*1,2}, Charles D. Kopec¹, Adrian G. Bondy¹, Thomas Z. Luo¹, Verity A. Elliott¹, Carlos D. Brody^{*1,3}

¹*Princeton Neuroscience Institute, Princeton University, Princeton NJ, USA*

²*Present address: Sainsbury Wellcome Centre, University College London, London, UK*

³*Howard Hughes Medical Institute, Princeton University, Princeton NJ, USA*

* *Correspondence should be addressed to diksha.gupta@ucl.ac.uk, brody@princeton.edu*

Abstract

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 anterodorsal striatum (ADS) is categorized in prefrontal cortex (frontal orienting fields, 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 FOF in accumulation, silencing FOF to ADS projections impacted behavior throughout the accumulation period, even 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^{1,2,3,4} and empirical^{5,6,7,8} 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 advantages⁹; for instance it allows for easy dynamic modulation of decision thresholds based on behavioral demands while leaving the accumulation process unchanged^{10,11,12,3,13}.

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 process¹⁴, whereas FOF is known to represent the categorical choice and is causally necessary only at the very end of accumulation period when the graded evidence needs to be thresholded^{15,16,17,18}. 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, thalamus and/or superior colliculus 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 choice¹⁶, 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 FOF¹⁶, 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 hemisphere²¹. Since the network has enough time to recover early in the accumulation period but not towards 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^{22,23,24}. 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^{27,28,29,30,31}. Additionally, confirming that the feedback projection under this hypothesis i.e. the FOF \rightarrow 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 sub-computation.

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 evolves at comparable timescales. Further, simultaneous population recordings let us define the feedforward (ADS \rightarrow FOF) and feedback (FOF \rightarrow ADS) communication subspaces 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 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

In order to test whether the hypothesized transformation between ADS and FOF holds on single trials (Figure 1B-C), we recorded simultaneous population activity from ADS and FOF of rats trained to perform a previously developed decision-making task³² (Fig 1A). This task requires accumulation of auditory evidence over hundreds of milliseconds for good performance. On any given trial, rats are presented with two streams of randomly timed auditory clicks, one played from a speaker to their left and the other from a speaker to their right. At the end of the stimulus, rats are rewarded with a drop of water for correctly reporting the side which played the greater number of clicks.

While rats ($n = 5$) performed this task, we used Neuropixels probes^{33,34} to record neural activity from FOF and ADS simultaneously (Fig 1D-E). We recorded from 765 FOF and 1583 ADS neurons across 12 sessions. Cells with a mean firing rate of at least 1Hz during the stimulus period were considered active and included for further analysis (462 FOF, 559 ADS neurons). In both FOF and ADS, we found that neurons that significantly modulated their activity in a side-selective manner on average exhibited firing rates that ramped upwards for stimuli favoring their preferred side with a slope proportional to the stimulus strength (Fig 1F). These trial-averaged stimulus-dependent ramping responses have long been interpreted as a neural correlate of the evidence accumulation process⁵ and are consistent with previous reports in rats from FOF¹⁶ and ADS¹⁴, and analogous regions in primates (FEF³⁵ and caudate³⁶). However, both graded and categorical representations of accumulated evidence can give rise to such ramping responses^{37,16}, hence to distinguish between these possibilities we characterized the encoding and decoding profiles of task-relevant variables in the two regions with more directed analysis.

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 compared to FOF.

To test the validity of this hypothesis, first we sought to compare the differences in the extent of graded encoding between FOF and ADS. For this we use the method developed in Hanks 2015¹⁶. This method computes the accumulator ‘tuning curves’ i.e. maps of how firing rates of neurons vary as a function of the accumulator value. To do so, it leverages the behavioral model from Brunton 2013³² to obtain trial-by-trial, moment-by-moment estimates of the subject’s accumulator value and relates it to the simultaneously observed firing rates of neurons. Previous application of this method had found that tuning curves in FOF had high slopes, i.e. were more step-like or categorical, with accumulator values on one side of the decision boundary giving rise to one cluster of firing rates, and the values on the other side of the decision boundary giving rise to another cluster¹⁶. On the other hand, previous application of this method in ADS had found a prevalence of neurons whose firing rates had lower slopes i.e. varied smoothly with the graded

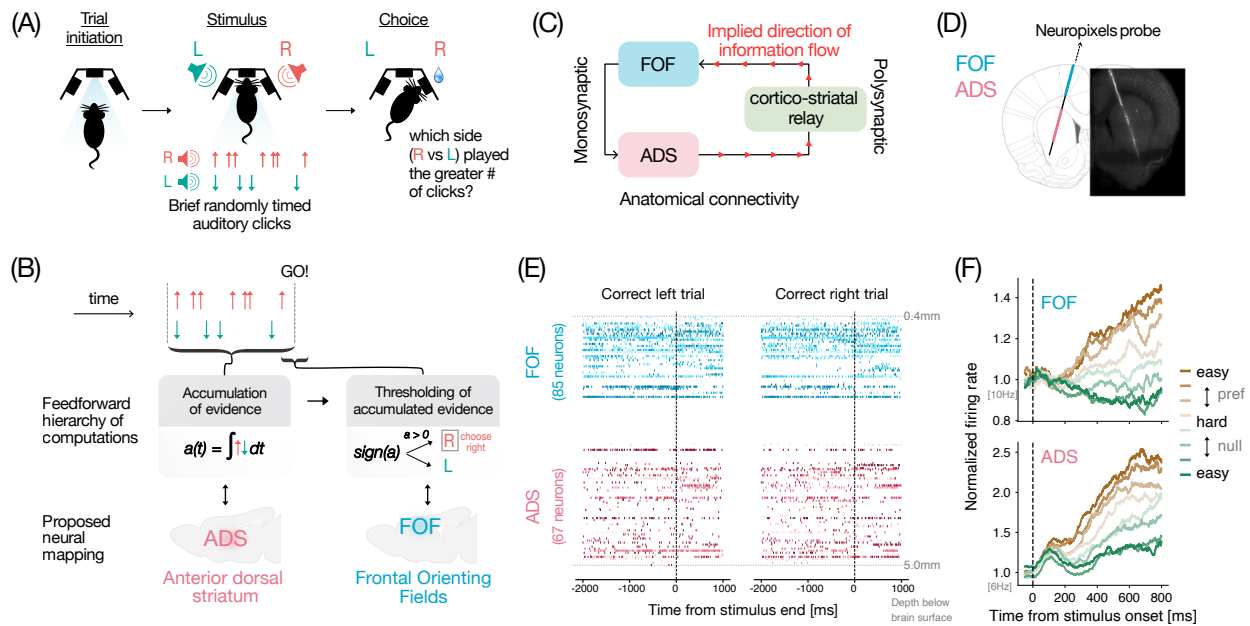


Figure 1 Experimental setup and simultaneous recordings from FOF and ADS (A) Schematic of the events in the rat evidence accumulation task (adapted from Brunton et al. 2013). From left to right: rats initiate a trial by poking their nose into the center port, two streams of randomly timed auditory clicks play from speakers on their left and right side, rats are rewarded for choosing the side with the greater number of clicks. (B) Schematic showing proposed implementation of the decision-making process^{16,14,38}. The anterior dorsal part of striatum (ADS, pink) is thought to be involved in gradual accumulation of sensory evidence and represents this accumulated value “*a*” in a graded fashion¹⁴. The accumulated evidence is thought to be thresholded, and represented in categorical form in the frontal orienting fields (FOF, blue), a part of the rat secondary motor cortex^{16,15,17}. (C) The pathway from ADS to FOF involves a multi-synaptic pathway (grey arrows through olive-green regions) through basal ganglia (BG) nuclei, the thalamus and the superior colliculus (SC) that could potentially support this transformation. (D) Schematic showing neuropixels targeting (AP 1.9, ML 1.15 at a 15 degrees angle in the coronal plane) to record simultaneously from FOF and ADS, overlaid on histology from an example rat. Brain atlas was adapted from Paxinos and Watson, 2006 (E) Spike rasters showing simultaneous neural activity recorded with a neuropixels probe from two example correct trials (left and right) from one session. Neurons are arranged in order of their estimated depth from the brain surface. Aqua blue shades represent FOF neurons and pink shades represent ADS neurons. (F) Average FOF (top) and ADS (bottom) population responses during stimulus period for different strengths of sensory evidence. Brown and green colors correspond to stimuli from preferred and non-preferred sides of the neurons, respectively. Darker hues correspond to easier trials and greyish hues correspond to harder trials. Both areas show stimulus strength dependent ramping activity. (FOF *n* = 219 neurons, ADS *n* = 235 active side-selective neurons from 5 rats).

value of the accumulator¹⁴. Here, simultaneous recordings allowed us to perform comparisons controlled for inter-subject variability.

When applying this analysis to side-selective neurons in our data, we found examples of both graded and categorically tuned neurons in FOF as well as ADS (Fig 2A). Surprisingly, we found no significant differences between FOF and ADS (Fig 2B, *P* = 0.13 Mann-Whitney U test). If anything, the median slope was slightly higher in ADS (0.32 ± 0.25 , median \pm std) compared to FOF (0.29 ± 0.22). Contrary to expectations from previous studies, these results demonstrate that when neural measurements from the two areas are made in the same subjects, the two regions do not appear to differ in their encoding of the accumulator variable.

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 (Fig 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 (Supp Fig 1A). 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 (Fig 2C; two way repeated measures ANOVA, $P = 0.93$ for region and $P = 0.36$ for region x time). These results did not change when choice information was controlled for, by separately decoding the stimulus on left and right choice trials (Supp Fig 1B). The results were similar also when instead of the veridical stimulus, the mean accumulated evidence inferred with the behavioral model was decoded from the two populations (not shown).

Additionally, we trained logistic decoders to predict other binary task-relevant variables such as choice and trial-history from the two regions. The two regions showed high decoding of choice (Supp Fig 1C) but once again with no significant differences in decoding performance over time (two way repeated measures ANOVA, $P = 0.90$ for region and $P = 0.54$ for region x time). The past trial's choice could be decoded well into the current trial with substantial accuracy from both regions (Supp Fig 1D), however there were no significant differences between the two regions (two way repeated measures ANOVA, $P = 0.13$ for region and $P = 0.37$ for region x time). 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 also supported by an independent analysis that decodes the decision variable from neural population activity. The distance of a population firing rate vector from the decision hyperplane of a logistic decoder that is trained to predict choice, reflects the model's prediction confidence. This neurally inferred variable, known as the decision variable (DV), estimates the animal's internal decision variable (log odds of making a choice) capturing the influence of sensory evidence and other choice determinants^{39,21,40}. This complements our previous analyses by relaxing assumptions of the behavioral model and not assuming a fixed stimulus encoding lag in the two regions. We sought to examine if these neurally inferred DVs from the two regions exhibit any systematic lead-lag relationships in their evolution, in order to infer the dominant direction of task-related information flow within the frontal cortico-striatal circuitry. We used the logistic decoder trained to predict the eventual choice (from Supp Fig 1C) and estimated the DV on single trials from the two regions (Supp Fig 2A). DVs from FOF and ADS showed good correspondence with each other and with the eventual true choice of the animal on both correct and error trials (Supp Fig 2B). Averaged across trials of different difficulties, DV values (unsigned) increased over the trial, with the rate of increase dependent on the strength of sensory evidence – exhibiting the classic ramping signature expected from evidence accumulation (Supp Fig 2C). Consistent

with our previous analyses, we did not find any lead-lag relationship in the evolution of DV in the two regions (Fig 2D). While the DVs showed high correlations at short latencies on average, the cross-correlogram of these DVs averaged over rats and sessions had a peak at a latency not significantly different from zero (mean latency \pm sem = -0.01 ± 0.02 s; $P = 0.47$ one-sided t-test, $n = 12$). This is consistent with no dominant feedforward or feedback relationship⁴¹. We conducted extensive simulations to validate our ability to infer lead-lag relationships in evolution of DV across varying number of simultaneously recorded neurons, trials and interaction lags (Supp Fig 3).

Next, we probed the content of the interaction between FOF and ADS during decision-making more generally, rather than restricting ourselves to activity along the decision hyperplane. To measure such interactions we used a generalized linear model with coupling filters across the two populations (similar to Pillow 2008⁴²) in addition to temporally-extended kernels for left and right clicks and other major task events. To overcome the explosion in the number of parameters with increasing numbers of cells, instead of modeling pairwise coupling terms, we projected the activity of the regressor population onto a (learnt) lower dimensional subspace, over which the coupling filters were learnt (Methods; Supp Fig 4A). We estimated the number of such latent factors required by evaluating the cross-validated log-likelihood of models fit with varying dimensions (Supp Fig 4B). The resultant model captured a high percentage of variance in average firing rates for most of the neurons in both areas (Supp Fig 4C). We compared the information content of activity in FOF \rightarrow ADS and ADS \rightarrow FOF communication subspaces (CS) to the whole population activity space by measuring the decoding accuracy of stimulus (Fig 2E). Remarkably, we found that a sizeable fraction of stimulus information in the two regions could also be decoded from the CS and this fraction increased over the course of the trial, with FOF sharing a slightly higher fraction of its information in its CS (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 noise fluctuations and therefore are likely communicated between the two.

Silencing FOF inputs to ADS disrupts decision-making

From the analysis of simultaneous activity of FOF and ADS we observed that the two regions have remarkably similar encoding and decoding of decision-relevant variables, this information is shared in their communication subspaces, and that there is no discernible lead-lag relationship between the representations of the evolving decision variable in the two regions. 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 interactions⁴¹ between FOF and ADS. Are these interactions necessary for making decisions? Establishing whether or not these interactions are important would put strong constraints on the network mechanisms that underlie decision-making and help further assess the feedforward, modular hypothesis.

We sought to probe this question in the experimentally tractable monosynaptic projection

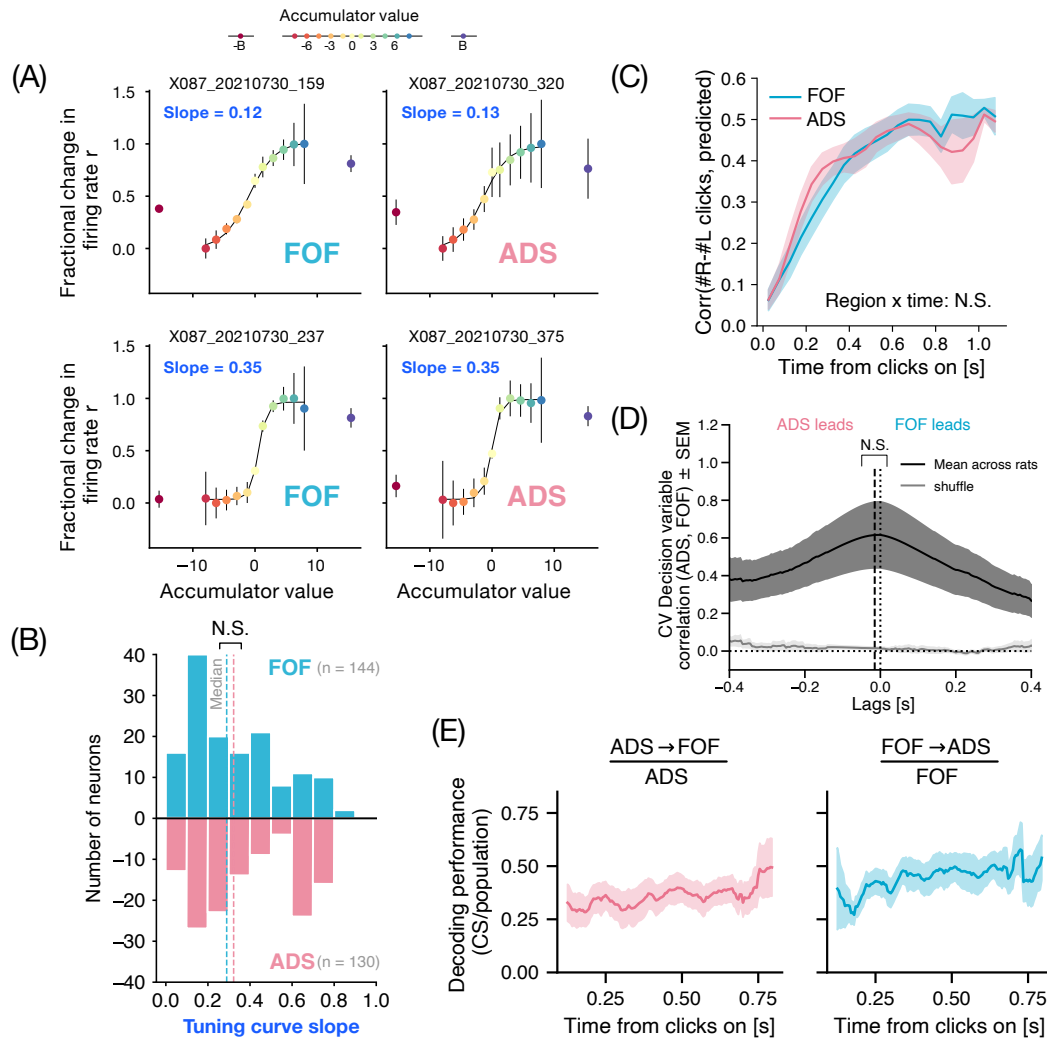


Figure 2 FOF and ADS carry redundant bidirectionally communicated task-related information (A) Tuning to the behaviorally inferred accumulator variable of two example neurons from FOF (left) and ADS (right) each. (Top panel) neuron exhibits a graded response to the accumulator variable, signified by the shallow slope of its tuning curve. (Bottom panel) exhibits a sharp dependence on the sign of the accumulator variable, reflected in the steeper slope of its tuning curve. **(B)** Histogram of tuning curve slopes in FOF (aqua blue; n = 144) and ADS (pink; n = 130). Dashed lines indicate the median slopes for the two populations (FOF = 0.29 ± 0.22 ; ADS = 0.32 ± 0.25 , median \pm std). The two populations have similar distribution of tuning to the accumulator variable, with neurons representing the accumulator value in both graded (smaller slopes) and categorical (larger slopes) fashion. The medians of the two populations were not significantly different ($P = 0.13$, Mann Whitney U test). **(C)** Stimulus decoding performance (mean \pm sem) of the linear decoder on held-out time points as a function of time from stimulus onset (n = 12 sessions). Linear decoders were fit independently to the two populations (FOF: aqua blue; ADS: pink) while controlling for the number of neurons from the two regions. Performance is measured as the correlation between the actual trajectory of difference in number of clicks and the prediction from linear decoding of neural responses. Decoding performance and its timecourse did not significantly differ between FOF and ADS ($P = 0.36$, two-way RM ANOVA) **(D)** Time-lagged cross-correlogram between decision-variable trajectories decoded from FOF and ADS, showing the correlation coefficient (y-axis) at different time lags (x-axis: positive (negative) values represent FOF trajectories leading (lagging) ADS). Solid black line indicates mean across sessions and rats, with dark-gray shaded region representing S.E.M. Light gray region indicates a shuffled control, computed by cross-correlating FOF activity with ADS activity on trials with shuffled (rather than matched) identities. (Caption continued on next page)

Figure 2 (Continued from previous page) The cross-correlogram is symmetric with a wide peak at a latency not significantly different from zero (mean latency \pm sem = -0.01 ± 0.02 s; $P = 0.47$ one-sided t-test, $n = 12$), hence showing no strong evidence of any lead-lag relationships. **(E)** Time-course of stimulus decoding accuracy present in the ADS \rightarrow FOF (pink) and FOF \rightarrow ADS (blue) communication subspaces (CSs), plotted as a fraction of stimulus decoding accuracy from the entire region's population activity. A fraction of 1 indicates that the decoding from CS is just as good as the entire population. Stimulus information is present in both regions' CS and increases over the course of the trial, with FOF sharing a slightly higher fraction of its information in the CS.

from FOF to ADS using unilateral optogenetic inactivations. We injected the inhibitory opsin, halorhodopsin (AAV5-CaMKII α -eNpHR3.0-eYFP) bilaterally in FOF and delivered light (25-33mW, 594nm) via a sharp fiber optic implanted in ADS to silence the activity of the FOF-ADS projection neurons' axon terminals (Fig 3A).

First, to test whether this projection has any role in the decision-making process, we silenced the activity of FOF \rightarrow ADS terminals in either hemisphere during the whole trial (2s inactivation starting 500ms before the stimulus period and ending 500ms after stimulus offset when the animal is free to make their response; Fig 3B,C) on a random subset of trials (25%). This inactivation significantly increased the proportion of choices ipsilateral to the laser side in rats injected with eNpHR3.0 expressing virus (Fig 3C; $P < 0.001$, nonparametric bootstrap test), with no significant effects on the controls injected with eYFP expressing virus (Fig 3C; $P = 0.78$, nonparametric bootstrap test), suggesting that the projection is involved in the decision process.

Next, we assessed the temporal extent of involvement of this projection, by restricting the inactivation to either early (first half of a 1s long stimulus) or late (second half of a 1s long stimulus) epochs of the evidence accumulation or stimulus period (Fig 3B-C; similar to^{16,14}). With this design, it is predicted that if the FOF \rightarrow ADS projection is involved in 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. Much to our surprise, when we restricted inactivations of the FOF \rightarrow ADS projection to the early and late epochs, we found that the inactivations led to an ipsilateral bias in both epochs (Fig 3C; early, $P = 0.002$; late, $P = 0.007$, nonparametric bootstrap test) with no significant differences in the degree of impairment ($P = 0.49$) between the two. This is in stark contrast to the effect of whole region FOF inactivations¹⁶ which caused an ipsilateral bias only during the late epoch. Instead, these results resemble ADS inactivations¹⁴ which caused an ipsilateral bias in both early and late epochs. Altogether, these results implicate the FOF \rightarrow ADS projection (i.e. feedback flow under the hypothesis) in the gradual accumulation process, revealing a previously obscured role for FOF neurons. Consistent with this, the involvement of cortico-striatal neurons throughout the stimulus period was also recently reported in mice performing a similar task using somatic inhibition⁴³.

Finally, we also examined the effect of inactivations on movements (Supp Fig 5). Terminal silencing slightly but significantly reduced mean movement times for choices both ipsiversive (P

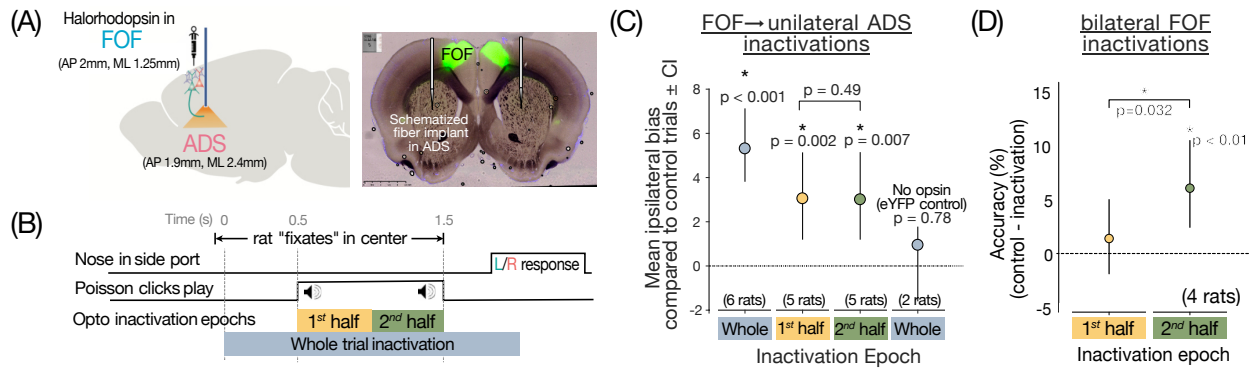


Figure 3 Silencing FOF inputs to ADS disrupts decision-making (A) Silencing FOF axon terminals in ADS: experimental setup (Left) Schematic showing viral delivery of Halorhodopsin (eNpHR3.0) to FOF and sharp optical fiber implant in ADS. Virus was injected bilaterally and 594nm wavelength laser (25-33mW) was delivered unilaterally through the fiber optic cable on 25% of trials to silence FOF axon terminals in ADS. (Right) Example histology from a rat expressing eNpHR3.0 bilaterally in FOF. (B) Timeline of task events showing perturbation epochs: FOF axon terminals in ADS were perturbed during one of three epochs: *Whole-trial* (grey bar) started as soon as the rat entered the center port and terminated 500ms after the stimulus ended. *Early* (yellow) spanned the first half of a 1s long stimulus and *late* (green) spanned the second half of a 1s long stimulus. It is predicted that if the FOF→ADS projection is involved in 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. (C) The ipsilateral bias, or the difference in percentage of trials the rat went to the side ipsilateral to the laser on inactivation v.s. control trials plotted for different inactivation epochs. Circles represent the mean ipsilateral bias and error bars represent 95% bootstrap confidence intervals. Whole-trial inactivations (6rats, 10 hemispheres) in rats injected with NpHR produced a significant impairment in behavior ($P < 0.001$, non-parametric bootstrap test) and no significant effects in the eYFP controls ($P = 0.78$) (2rats, 3 hemispheres). Early and late inactivations (5 rats, 8 hemispheres) also produced significant impairments ($P < 0.005$ non-parametric bootstrap test) with no significant differences between the effects of the two epochs ($P = 0.49$) suggesting that FOF-ADS projections are involved throughout the accumulation process. (D) Difference in accuracy between control and inactivation trials, plotted for temporally specific bilateral inactivations of FOF in early (yellow) or late (green) stimulus epochs. Circles represent mean accuracy difference across 4 rats, error bars represent 95% bootstrap confidence intervals. Only inactivations in the late epoch produce significant reductions in accuracy ($p < 0.01$, nonparametric bootstrap test), differing significantly from inactivations in early epochs ($p < 0.032$, nonparametric bootstrap test) in which there is no significant impairment.

= 0.05 Mann-Whitney U test) and contraversive ($P = 0.001$ Mann-Whitney U test) to the laser (Supp Fig 5A). However, silencing did not increase rates of fixation violations ($P = 0.58$, nonparametric bootstrap test; Supp Fig 5B). 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

Optogenetic inhibition of FOF inputs to ADS impairs decisions both in early and late stimulus epochs (Fig 3B-C), indicating a role for FOF in the gradual accumulation process, much like ADS. This is consistent with the highly similar neural representations observed in the two regions (Figs 1, 2) and suggests that rather than being incidental, evidence representations in FOF might be important for decision-making behavior. Interestingly, this result also raises the question: why don't

the unilateral whole-region FOF inactivations disrupt decisions during the accumulation period? In mouse M2 interhemispheric recovery was shown to be the reason for lack of impairments observed during transient unilateral inactivations in early task periods^{21,44}. So one hypothesis is that the FOFs in the two hemispheres are coupled and can support each other's recovery when aberrant inputs from the other hemisphere are detected. The localized nature of terminal inhibition (targeting FOF inputs just to ADS, and not to the other FOF) might bypass such a recovery.

We tested this hypothesis by carrying out bilateral optogenetic inactivations in a separate cohort of rats (Fig 3D). We silenced FOF activity by delivering laser bilaterally either during the first half of the stimulus (early inactivation epoch), or the second (late inactivation epoch; same as Fig 3B). We observed a significant decrement in performance only in the late epoch ($P < 0.01$, nonparametric bootstrap test), while inactivation during the early accumulation epoch did not significantly hamper the performance ($P > 0.05$, nonparametric bootstrap test). The decrement in performance due to early epoch inactivations was significantly smaller than the decrement in response to late epoch inactivation ($P = 0.032$, nonparametric bootstrap test). This implies that interhemispheric compensation does not underlie the lack of impairment observed on inactivating FOF early during evidence accumulation. So the question of why projection specific inactivations impair decisions early during accumulation but nonspecific whole region FOF inactivations don't, remains outstanding.

A model of distributed implementation of evidence accumulation in corticostriatal circuitry

Our experimental findings are difficult to reconcile with a feedforward organisation of these brain regions, wherein regions have localized functions (as previously hypothesized). They might however be concordant with a circuit that accumulates evidence in a distributed fashion through recurrence and feedback, which allows for adaptive rescue of computations upon perturbation^{45,21,46,47,48}. Therefore, we translated our findings into a mechanistic multi-region recurrent neural network (RNN) model that instantiates a recurrent circuit between FOF and ADS, building on the success of RNN models in capturing animal behavior and neural representations during decision-making^{49,50,51,52,53}. We sought to analyze the behavior of this RNN in unperturbed and perturbed states, so as to identify mechanism(s) through which it can be sensitive to projection specific inactivations but resilient to nonselective whole region FOF inactivations during the accumulation period.

We imposed biological constraints on the RNN, by initializing its weights with possible anatomical connections while following Dale's law and maintained these constraints on the architecture throughout training (following methods developed in Song 2016⁵⁴). The model's architecture is schematized in Fig. 4A (also see Supp Fig 6A-B). Briefly, the network included two modules for the two brain hemispheres and each module in turn consisted of three submodules with 50 all-to-all connected neural units each - representing FOF, ADS and the multi-synaptic relay between ADS and FOF (i.e. other basal ganglia nuclei, thalamus, SC, etc.). In the FOF submodules 20% of the neurons were inhibitory whereas in ADS submodules 100% of the neurons were inhibitory, approximating the known distribution of E-I neurons in these brain regions. We modeled long-

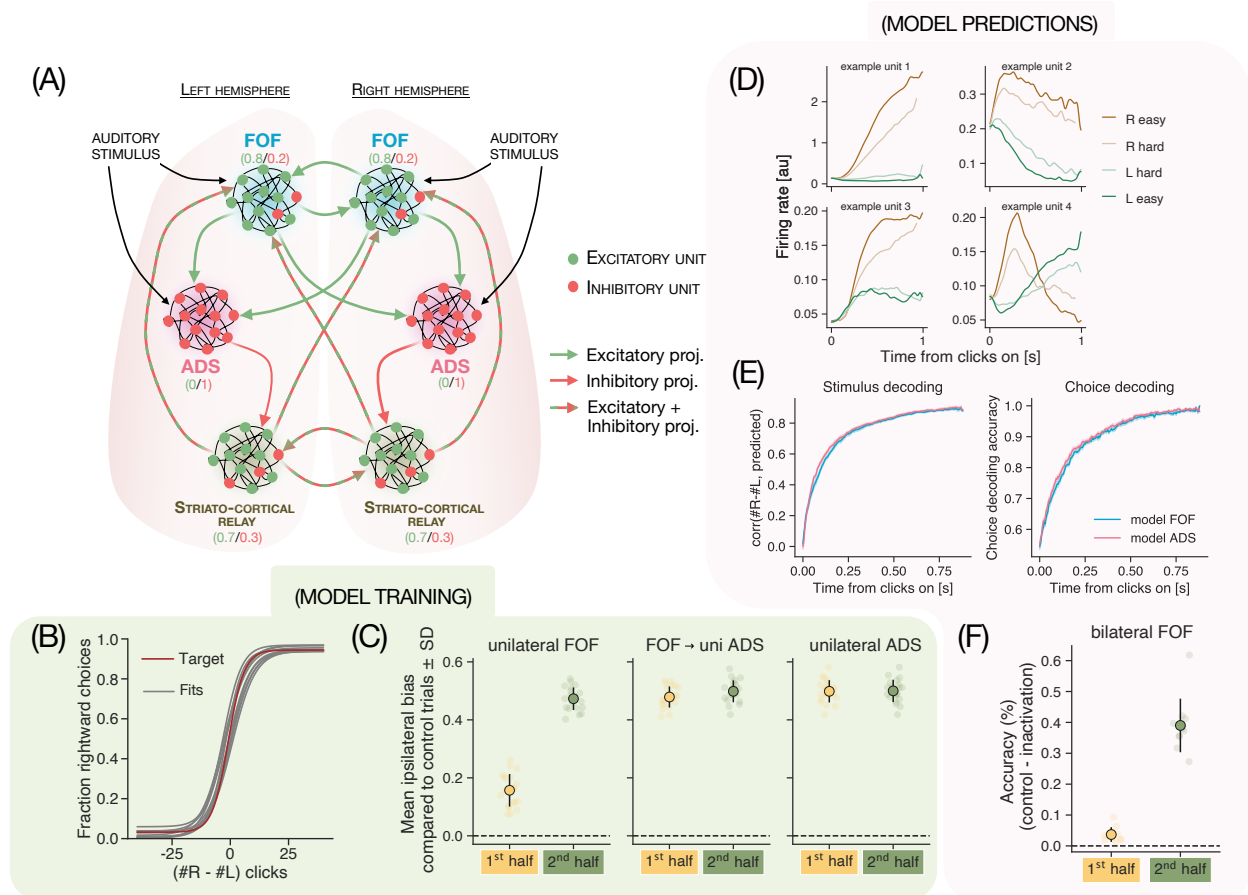


Figure 4 Multiregion recurrent neural network model successfully recapitulates empirical representational and robustness properties (A) Schematized model architecture. The model has two modules that represent the two brain hemispheres. Within each module are three submodules that correspond to FOF, ADS and the striato-cortical relay. Inhibitory projections and units are depicted in red, and excitatory projections and units are depicted in green. All units in the model followed Dale's law. The inputs were sent to all excitatory units in FOF and ADS and the outputs were read out from all units in the two hemispheres. (B) Psychometric curve for RNN model-produced choices (grey; $n = 10$) overlaid on target psychometric produced by the target choices from the Brunton model (red). (C) Average bias towards the ipsilateral choices exhibited upon unilateral (uni) inactivations of FOF, FOF→ADS projections and ADS (left to right) during either the first or second half of the stimulus. Bars represent the mean \pm CI across 10 trained networks, each dot represents results from inactivations in one of the hemispheres of the trained networks. The network recapitulates empirical data and shows high ipsilateral bias upon inactivation of FOF in the 2nd half compared to the 1st half. FOF→ADS and ADS inactivations produce ipsilateral bias both in the 1st and 2nd halves, resembling rat data. (D) PSTHs (aligned to stimulus onset, conditioned on stimulus difficulty) for example units from the model. Model units show variable timecourses and diversity in when the strength of evidence is encoded. (E) Model FOF and ADS have similar timecourse of stimulus and choice decoding (Left) Correlation between the true cumulative difference in number of right and left clicks at any time point during the trial and that predicted by a linear decoder trained on either model FOF activity (aqua blue) or model ADS activity (pink), as a function of time from stimulus onset. (Right) Accuracy of a logistic decoder trained to predict the eventual choice from responses of either model FOF (aqua blue) or ADS (pink) during the stimulus. (F) Model successfully predicts response to novel perturbations: difference in accuracy between model's choices on control trials and trials in which model FOF was inactivated bilaterally either during the first (orange) or the second (green) half of the stimulus. Bars represent the mean \pm CI across 10 trained networks, dots represent the difference for individual networks.

range projection patterns after empirical observations. The FOF submodule's projection patterns followed that of cortical long-range pyramidal neurons (IT and PT neurons), ADS projected unilaterally to the relay module and the relay module projected bilaterally to the two FOFs, completing the recurrent loop. The sensory inputs were sent to both FOF and ADS submodules, since it is known that the auditory cortex projects diffusely to both FOF and ADS in the rat brain. Outputs were read out from all neural units in the FOF, ADS and striato-cortical relay submodules.

In addition to this multi-region modular structure, our use of the RNN model differs from previous studies in two ways. First, we trained the RNN weights (using backpropagation through time) to perform the evidence accumulation task in a way that matches animal behavior, rather than like an ideal observer (Supp Fig 6B-D). To achieve this, we generated training data by sampling trials from the behavioral model of Brunton 2013³². Second, on a quarter of trials during training, we turned down the gain on weights either within an FOF submodule, ADS submodule or FOF inputs to an ADS submodule during either the first or the second half of the stimulus, essentially mimicking our unilateral inactivation experiments (Supp Fig 6E). On these inactivation trials we trained the network to produce an ipsilateral choice, if upon inactivation such a bias was observed empirically. Our use of perturbation data to introduce structure and specialization within a multi-region RNN model is novel, and complements previous approaches that have relied either on differences in neural representations or just the architecture to lend RNNs area-specific identities⁵⁵.

We trained the model to successfully reproduce choice behavior on control (Fig 4B) and inactivation trials (Fig 4C). In the model - much like the experimental data - unilateral FOF inactivations introduced an ipsilateral bias only in the second half of the stimulus, but inactivations of FOF to ADS projections and of ADS caused a bias throughout the stimulus. We probed the dynamical structure of these trained models in unperturbed states using the reverse-engineering approach from⁵⁶ and consistent with previous work modeling accumulation of evidence in RNNs^{50,57,31}, found approximate line attractors (not shown).

The RNN model's responses exhibited several features that resembled neural representations, despite not being explicitly trained to do so. The average firing rate of model units, just like neurons from FOF and ADS had heterogeneous temporally patterned responses (Fig 4D). Some neurons encoded the evidence strength transiently whereas other encoded it persistently, capturing the diversity observed in recorded neurons from FOF and ADS. Moreover, the timecourse of stimulus and choice decoding were similar across the model FOF and ADS, as we had observed in our population recordings (Fig. 4E, compare to Fig. 2C and Supp Fig. 1C).

In addition to these similarities in responses, the model was also successful at predicting effects of held-out perturbation. During training, the network was exposed to unilateral inactivations only, so bilateral inactivations of model FOF (Supp Fig 6F) formed a good test of the model's robustness properties. Prior modeling work comparing the robustness of different architectures²¹ has shown that modular and symmetric networks tend to be robust to unilateral inactivations but not bilateral ones. Therefore, we expected the model to deviate from our experimental findings and show disrupted performance during bilateral inactivations of FOF in the first half of stimulus - unless our novel training protocol with perturbation data was successful at instilling some robustness

properties of the real FOF-ADS circuitry into the network. To our surprise, much like the data, the model only showed performance disruptions in the second half of bilateral FOF inactivations, and not in the first half (Fig. 4F, for rat data see Fig. 3D).

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 (Fig 5A, B). 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 (Fig 5A), facilitating normal decision-making 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 towards the inactivated side (Fig 5B). A similar pattern is evident in network outputs, with inactivation trials initially diverging from control trials and subsequently recovering when FOF was inactivated in the first half (Fig 5C), but not when FOF→ADS projections were inactivated (Fig 5D).

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 different output projections from model FOF during the first half individually and in combination, to ascertain 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 (Fig 6A, Supp Fig 8A). We found that perturbing projections to the ipsilateral ADS had an inactivation effect in the first half, and invited no recovery in the second half (Supp Fig 8B). In contrast, perturbing projections to contralateral ADS had no inactivation effects at all (Supp Fig 8C). This suggests that the inactivation effects we had observed when perturbing FOF →ADS projections (which were projections to a single hemisphere's ADS from both FOFs; Supp Fig 6E) may have been largely driven by ipsilateral projections only. Next, we reasoned that inactivating FOF cell bodies should have an effect on the rest of the circuit equivalent to inactivating all of their outputs. We therefore asked, when inactivating the ipsilateral FOF→ADS projection, which FOF output projection inactivations led to the recovery observed when the FOF cells were inactivated. We found that additionally inactivating either of the self-recurrence or contralateral FOF projections led to only partial recovery, while additionally inactivating both of these invited complete recovery (Supp Fig 8D-E). We quantified these effects as the magnitude of the ipsilateral bias caused by these inactivations, relative to the bias caused by unilateral FOF inactivations. By this measure, each additional output inactivation invited significantly more recovery (Fig 6B).

Second, we selectively silenced different inputs to an FOF submodule during the second half, to ascertain which inputs contributed to the recovery in the second half. These included audi-

tory stimulus inputs, inputs from the contralateral FOF and inputs from the striato-cortical relay (Fig 6C). We found that removing any of these inputs individually had minor effects on the ipsilateral bias, only slightly reducing recovery. However, inactivating pairs of these inputs led to significantly more loss of recovery (Fig 6D).

Overall, these experiments suggest that recovery following unilateral FOF perturbations is supported by multiple scales of recurrence in the cortico-striatal circuit - self recurrence within a single hemisphere's FOF, inter-hemispheric recurrence and striato-cortical recurrence through the relay. Since our in-silico experiments target anatomically valid projections, they make empirically testable predictions about the *in-vivo* cortico-striatal circuitry.

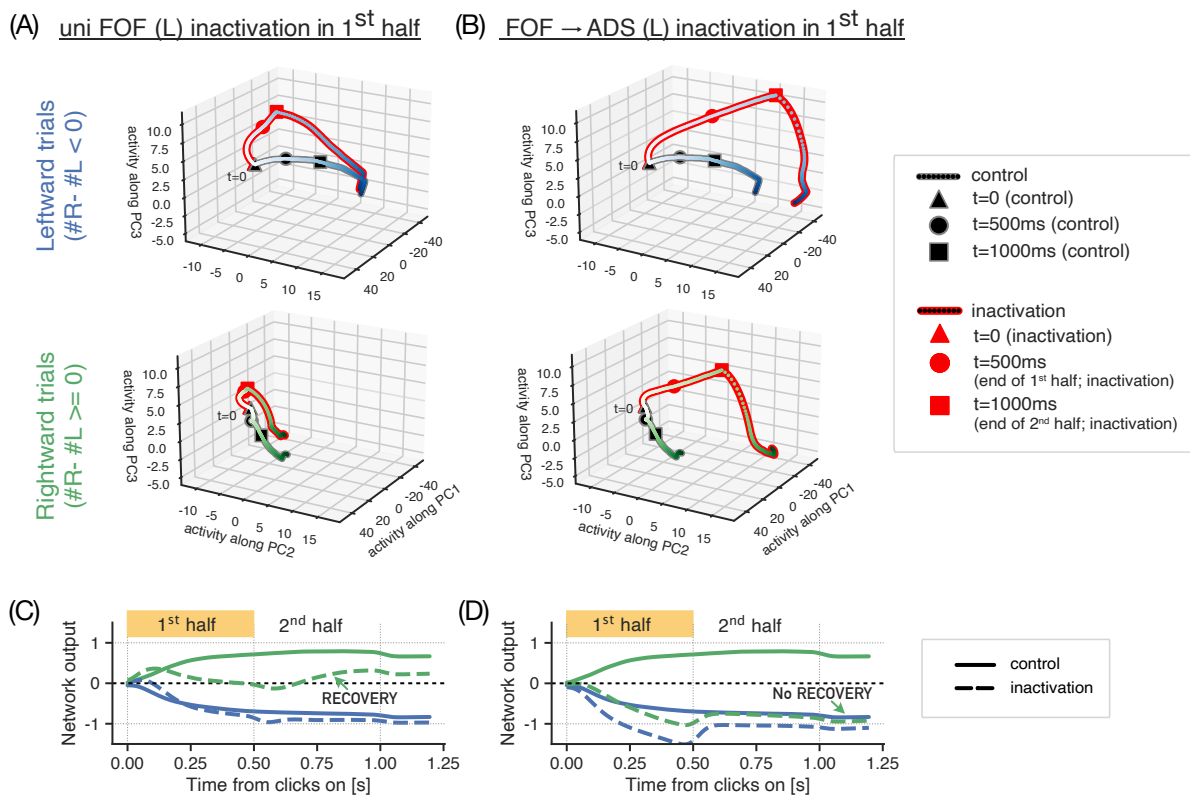


Figure 5 Network activity and outputs during nonselective FOF inactivations reveal recovery dynamics. Low-dimensional projections of network activity trajectories (A,B) and network outputs (C,D) during unilateral (uni; here left hemisphere) inactivations of FOF (A,C) and inactivations of its projections to left ADS (B,D) in the 1st half of the stimulus period. (A-B) Neural activity on leftward (top) and rightward (bottom) trials show that inactivation trajectories (lines with red edges, red boxes) diverge from control ones (lines/boxes with grey edges) for both experiments in the first half (until circle markers). However in the second half (from circle to square markers), these recover towards control trajectories during unilateral FOF inactivations (A), but not during FOF→ADS inactivations (B) driving leftward (or ipsilaterally biased) choices. The effect is more prominent in rightward trials (lower panels) (C-D) A similar pattern is seen in the network outputs on rightward trials (green lines in C,D), where for both experiments, outputs upon inactivation (dashed) diverge from control (solid) in the first half, recovering and driving correct choices during FOF inactivations (C) but not during FOF→ADS inactivations (D).

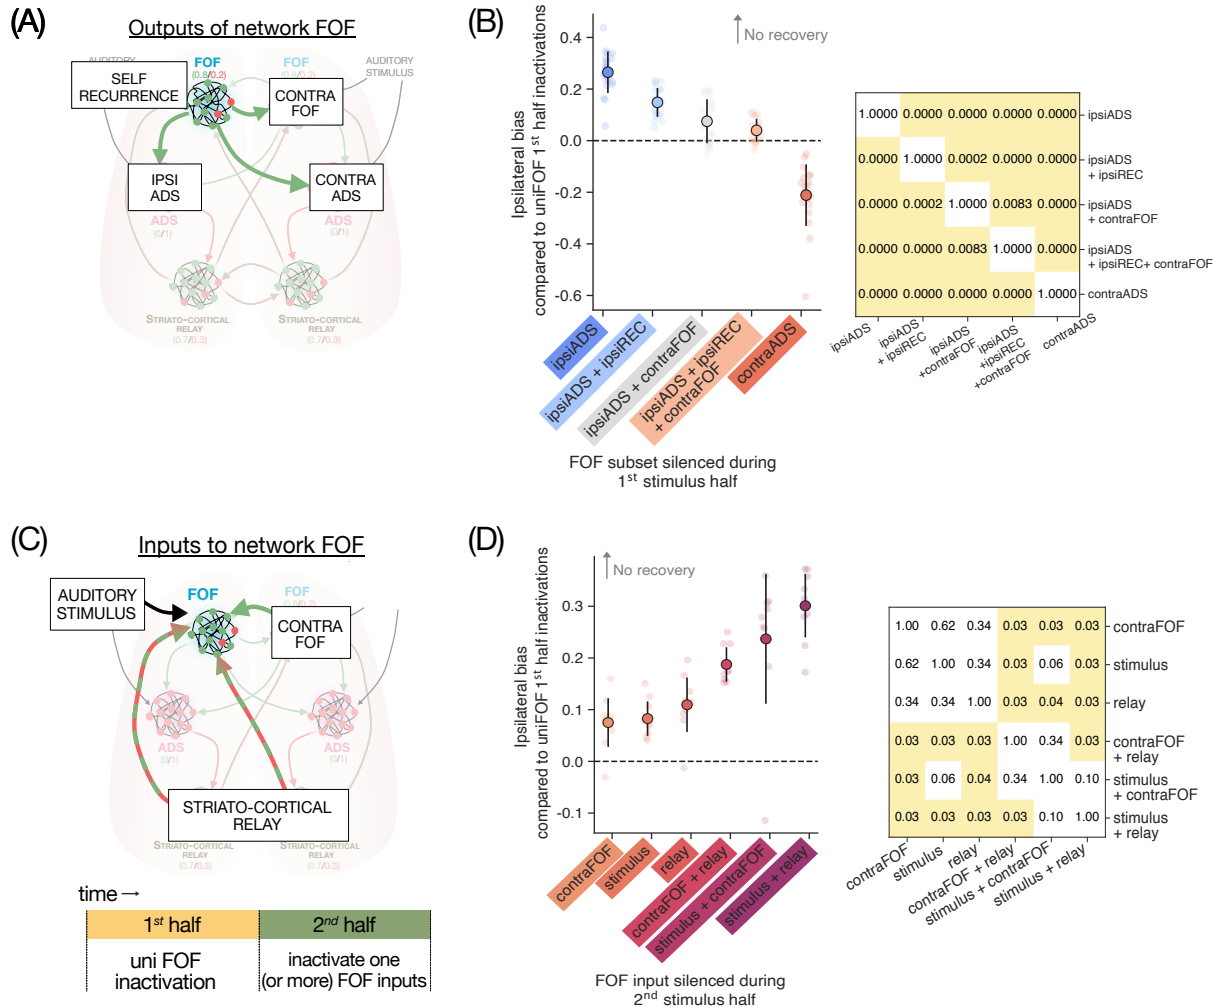


Figure 6 In-silico experiments with RNN to investigate contributions to inactivation dynamics (A-B): Selectively silencing of different FOF outputs in 1st half invites recovery differentially. **(A)** Schematic of model FOF output projections from a single hemisphere, targeted for in-silico perturbations in the 1st half. Different combinations of output projections were silenced to understand their influence on the inactivation dynamics. **(B)** Summary of in-silico perturbations, showing the ipsilateral bias observed when perturbing a subset of FOF output projections in the 1st half, relative to a baseline of perturbing all its outputs - higher values indicate larger effects and less recovery. Silencing ipsilateral ADS projections has the largest effect and invites the least recovery, additionally silencing recurrent connections and contralateral FOF projections invites partial recovery individually and full recovery in combination. Contralateral ADS projections have no effect to begin with, hence producing less bias than baseline (i.e. negative values). Inset: matrix of p-values for pairwise comparisons between effects, showing that all experiments have significantly different effects. **(C-D):** Selectively silencing FOF inputs in 2nd half shows multiple contributions to recovery **(C)** Schematic of model FOF inputs to a single hemisphere, targeted for in-silico perturbations in the 2nd half. **(D)** Summary of in-silico perturbations, showing the impact of silencing different input projections in the 2nd half on ipsilateral bias, relative to whole FOF silencing in the 1st half. Higher values of bias when a given projection is silenced mean less recovery, implicating that input projection in the recovery process. Inputs from contralateral FOF, auditory stimulus, and striato-cortical relay all have partial influence on recovery, with significantly bigger losses of recovery when they are silenced together. Inset: matrix of p-values for pairwise comparisons between effects, showing that many of these inputs have comparable influence on recovery. P-values were computed with posthoc Wilcoxon test with holm correction for multiple comparisons.

Discussion

Two separate brain regions in rats, the anterior dorsal stream (ADS) and the frontal orienting fields (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 synthesised 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 time scales⁵⁸ and that striatum reflects summed activity of its cortical inputs⁵⁹. Moreover, our results about encoding of sensory evidence in FOF are in agreement with a recent study of dorsal cortex’s role during decision-making in mice which reported that secondary motor cortex prominently encodes choice independent evidence-related signals⁶⁰. In the random dots motion task⁶¹, monkey FEF and caudate (the analogues 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 (FEF⁶², caudate³⁶, FEF and caudate^{35,63,64}). This similarity in evidence representation and functional role of 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 2006² are unclear, which predicts transient rather than sustained responses in striatum and

models striatum's role as the setting of evidence thresholds.

Despite these similarities in evidence accumulation signals between the two regions, some studies have indeed found differences between FEF and caudate during decision making, particularly in tasks with reward manipulations that require the adaptive tuning of decision thresholds^{63,64,65}. Such results fit with wider studies implicating the striatum in action selection⁶⁶ and reward-based habitual learning, and secondary motor cortex in planning goal-directed movements^{67,68}. 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. Indeed work from Luo and Kim⁶⁹ find differences in the geometry of dynamics in the two regions. Whether such distinctions influence the functional roles of these regions is an open question, and one that we leave for future work to investigate.

To understand if the evidence representations in FOF and ADS are related, we used reduced rank regression to identify dimensions in the neural population space that captured co-fluctuation in the trial-to-trial activity of FOF and ADS over and above those introduced by external task events and defined a functional communication subspace^{70,26,71,72,73,41}. We found that during the stimulus period, the communication subspace carries stimulus-related information such that evidence can be decoded fairly well from just the activity present in the communication subspace. This did not have to be the case, as using similar approaches recent work has identified patterns of activity or information that are shared between brain areas as well as those that are kept private within an area. For instance, it has been reported that majority of the variance in V1 is kept private and that V1 only selectively shares information related to a few of its activity patterns with V2²⁰. In rat M1 and M2 it has been shown that shared dynamics across areas identified using a related method (canonical correlation analysis; CCA) become more related to task performance with learning⁷⁴. Similarly in mice, CCA was recently used to show that the content of communication between A1 and mPFC changes during different contexts based on a control signal from mPFC, this helps select relevant sensory information during contextual decision making⁷⁵. Such communication subspaces defined based on noise correlations do not distinguish between correlations due to direct communication from those induced by communication arriving from common inputs to the two regions. While our projection-specific inactivations support a causal role for the interaction from FOF to ADS, our conclusions about communication of evidence between the two regions can be further strengthened if decrements in evidence-related information are observed upon perturbation of one of the areas, in paired perturbation and recording experiments.

Through specifically targeting the cortico-striatal pathway and silencing FOF axon terminals in ADS, we revealed a role for FOF throughout the gradual accumulation process. This finding might seem at odds with the past work that had 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 period¹⁶ – 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^{22,76,77,78,79,80,43,81,82,83,84} (further discussion^{66,85}) producing effects that could be otherwise obscured by whole region perturbations. A notable example of this comes from an auditory frequency discrimination task in rats⁷⁶, where perturbations of auditory cortex neurons with high (low) frequency tuning did not bias animals’ choices towards 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 mice⁴³, which has also proposed a role for corticostriatal projections from 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^{22,76,24}. Given recent controversies about the causal roles of different regions in evidence accumulation^{86,15,25,87,88}, 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.

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^{89,90,91,92,93}. 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^{94,95,96,47,97,48} (reviewed in^{98,55}) 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 (also see⁹⁹). 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.

References

- [1] Kong-Fatt Wong and Xiao-Jing Wang. “A Recurrent Network Mechanism of Time Integration in Perceptual Decisions”. en. In: *Journal of Neuroscience* 26.4 (2006). Publisher: Society for Neuroscience Section: Articles, pp. 1314–1328. DOI: 10.1523/JNEUROSCI.3733-05.2006.
- [2] Chung-Chuan Lo and Xiao-Jing Wang. “Cortico–basal ganglia circuit mechanism for a decision threshold in reaction time tasks”. en. In: *Nature Neuroscience* 9.7 (2006). Number: 7 Publisher: Nature Publishing Group, pp. 956–963. DOI: 10.1038/nn1722.

- [3] Rafal Bogacz and Kevin Gurney. “The Basal Ganglia and Cortex Implement Optimal Decision Making Between Alternative Actions”. In: *Neural Computation* 19.2 (2007), pp. 442–477. DOI: 10.1162/neco.2007.19.2.442.
- [4] Rajesh P. N. Rao. “Decision Making Under Uncertainty: A Neural Model Based on Partially Observable Markov Decision Processes”. In: *Frontiers in Computational Neuroscience* 4 (2010), p. 146. DOI: 10.3389/fncom.2010.00146.
- [5] Joshua I. Gold and Michael N. Shadlen. “The Neural Basis of Decision Making”. In: *Annual Review of Neuroscience* 30.1 (2007), pp. 535–574. DOI: 10.1146/annurev.neuro.29.051605.113038.
- [6] Timothy D. Hanks and Christopher Summerfield. “Perceptual Decision Making in Rodents, Monkeys, and Humans”. en. In: *Neuron* 93.1 (2017), pp. 15–31. DOI: 10.1016/j.neuron.2016.12.003.
- [7] Elizabeth J. Jun et al. “Causal role for the primate superior colliculus in the computation of evidence for perceptual decisions”. en. In: *Nature Neuroscience* 24.8 (2021). Publisher: Nature Publishing Group, pp. 1121–1131. DOI: 10.1038/s41593-021-00878-6.
- [8] Gabriel M. Stine et al. “A neural mechanism for terminating decisions”. eng. In: *Neuron* 111.16 (2023), 2601–2613.e5. DOI: 10.1016/j.neuron.2023.05.028.
- [9] Gouki Okazawa and Roozbeh Kiani. “Neural Mechanisms that Make Perceptual Decisions Flexible”. In: *Annual Review of Physiology* 85.1 (2023), null. DOI: 10.1146/annurev-physiol-031722-024731.
- [10] Joshua I Gold and Michael N Shadlen. “Banburismus and the Brain: Decoding the Relationship between Sensory Stimuli, Decisions, and Reward”. In: *Neuron* 36.2 (2002), pp. 299–308.
- [11] Rafal Bogacz et al. “The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks”. In: *Psychological Review* 113.4 (2006). Place: US Publisher: American Psychological Association, pp. 700–765. DOI: 10.1037/0033-295X.113.4.700.
- [12] Patrick Simen, Jonathan D. Cohen, and Philip Holmes. “Rapid decision threshold modulation by reward rate in a neural network”. In: *Neural Networks. Neurobiology of Decision Making* 19.8 (2006), pp. 1013–1026. DOI: 10.1016/j.neunet.2006.05.038.
- [13] J. Drugowitsch et al. “The Cost of Accumulating Evidence in Perceptual Decision Making”. en. In: *Journal of Neuroscience* 32.11 (2012), pp. 3612–3628. DOI: 10.1523/JNEUROSCI.4010-11.2012.
- [14] Michael M Yartsev et al. “Causal contribution and dynamical encoding in the striatum during evidence accumulation”. In: *eLife* 7 (2018). Ed. by Joshua I Gold et al. Publisher: eLife Sciences Publications, Ltd, e34929. DOI: 10.7554/eLife.34929.

- [15] Jeffrey C Erlich et al. “Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat”. In: *eLife* 4 (2015), p. 8166.
- [16] Timothy D Hanks et al. “Distinct relationships of parietal and prefrontal cortices to evidence accumulation”. In: *Nature* 520.7546 (2015), pp. 220–223.
- [17] Alex T. Piet et al. “Rat Prefrontal Cortex Inactivations during Decision Making Are Explained by Bistable Attractor Dynamics”. en. In: *Neural Computation* 29.11 (2017), pp. 2861–2886. DOI: 10.1162/neco_a_01005.
- [18] J. Tyler Boyd-Meredith et al. “Stable choice coding in rat frontal orienting fields across model-predicted changes of mind”. en. In: *Nature Communications* 13.1 (2022). Number: 1 Publisher: Nature Publishing Group, p. 3235. DOI: 10.1038/s41467-022-30736-3.
- [19] Matthew T. Kaufman et al. “Cortical activity in the null space: permitting preparation without movement”. en. In: *Nature Neuroscience* 17.3 (2014). Number: 3 Publisher: Nature Publishing Group, pp. 440–448. DOI: 10.1038/nn.3643.
- [20] João D. Semedo et al. “Cortical Areas Interact through a Communication Subspace”. en. In: *Neuron* 102.1 (2019), 249–259.e4. DOI: 10.1016/j.neuron.2019.01.026.
- [21] Nuo Li et al. “Robust neuronal dynamics in premotor cortex during motor planning”. In: *Nature* 532.7600 (2016), pp. 459–464.
- [22] Melissa R Warden et al. “A prefrontal cortex-brainstem neuronal projection that controls response to behavioural challenge”. In: *Nature* 492.7429 (2012), pp. 428–432.
- [23] James M Otis et al. “Prefrontal cortex output circuits guide reward seeking through divergent cue encoding”. In: *Nature* 543.7643 (2017), pp. 103–107.
- [24] Christina K Kim et al. “Molecular and Circuit-Dynamical Identification of Top-Down Neural Mechanisms for Restraint of Reward Seeking”. In: *Cell* (2017).
- [25] Leor N. Katz et al. “Dissociated functional significance of decision-related activity in the primate dorsal stream”. en. In: *Nature* 535.7611 (2016). Number: 7611 Publisher: Nature Publishing Group, pp. 285–288. DOI: 10.1038/nature18617.
- [26] João D Semedo et al. “Statistical methods for dissecting interactions between brain areas”. en. In: *Current Opinion in Neurobiology*. Whole-brain interactions between neural circuits 65 (2020), pp. 59–69. DOI: 10.1016/j.conb.2020.09.009.
- [27] Ariel Gilad et al. “Behavioral Strategy Determines Frontal or Posterior Location of Short-Term Memory in Neocortex”. eng. In: *Neuron* 99.4 (2018), 814–828.e7. DOI: 10.1016/j.neuron.2018.07.029.
- [28] Kenneth W. Latimer and David J. Freedman. *Low-dimensional encoding of decisions in parietal cortex reflects long-term training history*. en. Pages: 2021.10.07.463576 Section: New Results. 2021. DOI: 10.1101/2021.10.07.463576.

- [29] Charlotte Arlt et al. “Cognitive experience alters cortical involvement in goal-directed navigation”. In: *eLife* 11 (2022). Ed. by Mathieu Wolff, Joshua I Gold, and Mathieu Wolff. Publisher: eLife Sciences Publications, Ltd, e76051. DOI: 10.7554/eLife.76051.
- [30] Scott S. Bolkan et al. “Opponent control of behavior by dorsomedial striatal pathways depends on task demands and internal state”. en. In: *Nature Neuroscience* 25.3 (2022). Number: 3 Publisher: Nature Publishing Group, pp. 345–357. DOI: 10.1038/s41593-022-01021-9.
- [31] Marino Pagan et al. *A new theoretical framework jointly explains behavioral and neural variability across subjects performing flexible decision-making*. en. Pages: 2022.11.28.518207 Section: New Results. 2022. DOI: 10.1101/2022.11.28.518207.
- [32] Bingni W. Brunton, Matthew M. Botvinick, and Carlos D. Brody. “Rats and Humans Can Optimally Accumulate Evidence for Decision-Making”. en. In: *Science* 340.6128 (2013). Publisher: American Association for the Advancement of Science Section: Report, pp. 95–98. DOI: 10.1126/science.1233912.
- [33] James J. Jun et al. “Fully integrated silicon probes for high-density recording of neural activity”. en. In: *Nature* 551.7679 (2017). Number: 7679 Publisher: Nature Publishing Group, pp. 232–236. DOI: 10.1038/nature24636.
- [34] Thomas Zhihao Luo et al. *An approach for long-term, multi-probe Neuropixels recordings in unrestrained rats*. en. Publisher: eLife Sciences Publications Limited. 2020. DOI: 10.7554/eLife.59716.
- [35] Long Ding and Joshua I. Gold. “Neural Correlates of Perceptual Decision Making before, during, and after Decision Commitment in Monkey Frontal Eye Field”. In: *Cerebral Cortex* 22.5 (2012), pp. 1052–1067. DOI: 10.1093/cercor/bhr178.
- [36] Long Ding and Joshua I. Gold. “Caudate Encodes Multiple Computations for Perceptual Decisions”. en. In: *Journal of Neuroscience* 30.47 (2010). Publisher: Society for Neuroscience Section: Articles, pp. 15747–15759. DOI: 10.1523/JNEUROSCI.2894-10.2010.
- [37] Kenneth W. Latimer et al. “Single-trial Spike Trains in Parietal Cortex Reveal Discrete Steps During Decision-making”. In: *Science (New York, N.Y.)* 349.6244 (2015), pp. 184–187. DOI: 10.1126/science.aaa4056.
- [38] Carlos D Brody and Timothy D Hanks. “Neural underpinnings of the evidence accumulator”. en. In: *Current Opinion in Neurobiology*. Neurobiology of cognitive behavior 37 (2016), pp. 149–157. DOI: 10.1016/j.conb.2016.01.003.
- [39] Roozbeh Kiani et al. “Dynamics of Neural Population Responses in Prefrontal Cortex Indicate Changes of Mind on Single Trials”. en. In: *Current Biology* 24.13 (2014), pp. 1542–1547. DOI: 10.1016/j.cub.2014.05.049.

- [40] Diogo Peixoto et al. “Decoding and perturbing decision states in real time”. en. In: *Nature* 591.7851 (2021). Number: 7851 Publisher: Nature Publishing Group, pp. 604–609. DOI: 10.1038/s41586-020-03181-9.
- [41] João D. Semedo et al. “Feedforward and feedback interactions between visual cortical areas use different population activity patterns”. en. In: *Nature Communications* 13.1 (2022). Number: 1 Publisher: Nature Publishing Group, p. 1099. DOI: 10.1038/s41467-022-28552-w.
- [42] Jonathan W. Pillow et al. “Spatio-temporal correlations and visual signalling in a complete neuronal population”. en. In: *Nature* 454.7207 (2008). Publisher: Nature Publishing Group, pp. 995–999. DOI: 10.1038/nature07140.
- [43] Simon Musall et al. *Pyramidal cell types drive functionally distinct cortical activity patterns during decision-making*. en. Tech. rep. Company: Cold Spring Harbor Laboratory Distributor: Cold Spring Harbor Laboratory Label: Cold Spring Harbor Laboratory Section: New Results Type: article. 2021, p. 2021.09.27.461599. DOI: 10.1101/2021.09.27.461599.
- [44] Guang Chen et al. “Modularity and robustness of frontal cortical networks”. en. In: *Cell* 184.14 (2021), 3717–3730.e24. DOI: 10.1016/j.cell.2021.05.026.
- [45] Charles D. Kopec et al. “Cortical and Subcortical Contributions to Short-Term Memory for Orienting Movements”. en. In: *Neuron* 88.2 (2015), pp. 367–377. DOI: 10.1016/j.neuron.2015.08.033.
- [46] Chunyu A. Duan et al. “Collicular circuits for flexible sensorimotor routing”. en. In: *Nature Neuroscience* 24.8 (2021). Number: 8 Publisher: Nature Publishing Group, pp. 1110–1120. DOI: 10.1038/s41593-021-00865-x.
- [47] Colin J. Bredenberg, Cristina Savin, and Roozbeh Kiani. *Recurrent neural circuits overcome partial inactivation by compensation and re-learning*. en. Tech. rep. Company: Cold Spring Harbor Laboratory Distributor: Cold Spring Harbor Laboratory Label: Cold Spring Harbor Laboratory Section: New Results Type: article. 2021, p. 2021.11.12.468273. DOI: 10.1101/2021.11.12.468273.
- [48] Jorge F Mejías and Xiao-Jing Wang. “Mechanisms of distributed working memory in a large-scale network of macaque neocortex”. In: *eLife* 11 (2022). Ed. by Tatiana Pasternak and Tirin Moore. Publisher: eLife Sciences Publications, Ltd, e72136. DOI: 10.7554/eLife.72136.
- [49] Mattia Rigotti et al. “Internal representation of task rules by recurrent dynamics: the importance of the diversity of neural responses”. eng. In: *Frontiers in Computational Neuroscience* 4 (2010), p. 24. DOI: 10.3389/fncom.2010.00024.
- [50] Valerio Mante et al. “Context-dependent computation by recurrent dynamics in prefrontal cortex”. en. In: *Nature* 503.7474 (2013), pp. 78–84. DOI: 10.1038/nature12742.

- [51] Federico Carnevale et al. “Dynamic Control of Response Criterion in Premotor Cortex during Perceptual Detection under Temporal Uncertainty”. In: *Neuron* 86.4 (2015), pp. 1067–1077.
- [52] Evan D. Remington et al. “A Dynamical Systems Perspective on Flexible Motor Timing”. en. In: *Trends in Cognitive Sciences*. Special Issue: Time in the Brain 22.10 (2018), pp. 938–952. DOI: 10.1016/j.tics.2018.07.010.
- [53] Michael Kleinman, Chandramouli Chandrasekaran, and Jonathan Kao. “A mechanistic multi-area recurrent network model of decision-making”. In: *Advances in Neural Information Processing Systems*. Vol. 34. Curran Associates, Inc., 2021, pp. 23152–23165.
- [54] H. Francis Song, Guangyu R. Yang, and Xiao-Jing Wang. “Training Excitatory-Inhibitory Recurrent Neural Networks for Cognitive Tasks: A Simple and Flexible Framework”. en. In: *PLOS Computational Biology* 12.2 (2016). Publisher: Public Library of Science, e1004792. DOI: 10.1371/journal.pcbi.1004792.
- [55] Guangyu Robert Yang and Manuel Molano-Mazón. “Towards the next generation of recurrent network models for cognitive neuroscience”. en. In: *Current Opinion in Neurobiology*. Computational Neuroscience 70 (2021), pp. 182–192. DOI: 10.1016/j.conb.2021.10.015.
- [56] David Sussillo and Omri Barak. “Opening the Black Box: Low-Dimensional Dynamics in High-Dimensional Recurrent Neural Networks”. In: *Neural Computation* 25.3 (2013), pp. 626–649. DOI: 10.1162/NECO_a_00409.
- [57] Niru Maheswaranathan et al. “Universality and individuality in neural dynamics across large populations of recurrent networks”. In: *Advances in Neural Information Processing Systems*. Vol. 32. Curran Associates, Inc., 2019.
- [58] Nicholas A. Steinmetz et al. “Distributed coding of choice, action and engagement across the mouse brain”. en. In: *Nature* 576.7786 (2019). Number: 7786 Publisher: Nature Publishing Group, pp. 266–273. DOI: 10.1038/s41586-019-1787-x.
- [59] Andrew J. Peters et al. “Striatal activity topographically reflects cortical activity”. en. In: *Nature* 591.7850 (2021). Number: 7850 Publisher: Nature Publishing Group, pp. 420–425. DOI: 10.1038/s41586-020-03166-8.
- [60] Ivana Orsolich et al. “Mesoscale cortical dynamics reflect the interaction of sensory evidence and temporal expectation during perceptual decision-making”. English. In: *Neuron* 109.11 (2021). Publisher: Elsevier, 1861–1875.e10. DOI: 10.1016/j.neuron.2021.03.031.
- [61] Jamie D. Roitman and Michael N. Shadlen. “Response of Neurons in the Lateral Intraparietal Area during a Combined Visual Discrimination Reaction Time Task”. en. In: *Journal of Neuroscience* 22.21 (2002). Publisher: Society for Neuroscience Section: ARTICLE, pp. 9475–9489. DOI: 10.1523/JNEUROSCI.22-21-09475.2002.

- [62] Jong-Nam Kim and Michael N. Shadlen. “Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque”. en. In: *Nature Neuroscience* 2.2 (1999). Number: 2 Publisher: Nature Publishing Group, pp. 176–185. DOI: 10.1038/5739.
- [63] Long Ding. “Distinct dynamics of ramping activity in the frontal cortex and caudate nucleus in monkeys”. In: *Journal of Neurophysiology* 114.3 (2015). Publisher: American Physiological Society, pp. 1850–1861. DOI: 10.1152/jn.00395.2015.
- [64] Yunshu Fan, Joshua I Gold, and Long Ding. “Frontal eye field and caudate neurons make different contributions to reward-biased perceptual decisions”. In: *eLife* 9 (2020). Ed. by Daeyeol Lee et al. Publisher: eLife Sciences Publications, Ltd, e60535. DOI: 10.7554/eLife.60535.
- [65] Yunshu Fan et al. “Neural Representations of Post-Decision Accuracy and Reward Expectation in the Caudate Nucleus and Frontal Eye Field”. en. In: *The Journal of Neuroscience* 44.2 (2024), e0902232023. DOI: 10.1523/JNEUROSCI.0902-23.2023.
- [66] Julia Cox and Ilana B. Witten. “Striatal circuits for reward learning and decision-making”. en. In: *Nature Reviews Neuroscience* 20.8 (2019). Number: 8 Publisher: Nature Publishing Group, pp. 482–494. DOI: 10.1038/s41583-019-0189-2.
- [67] Christian L. Ebbesen et al. “More than Just a “Motor”: Recent Surprises from the Frontal Cortex”. en. In: *Journal of Neuroscience* 38.44 (2018). Publisher: Society for Neuroscience Section: Symposium and Mini-Symposium, pp. 9402–9413. DOI: 10.1523/JNEUROSCI.1671-18.2018.
- [68] Jen-Hau Yang and Alex C. Kwan. “Secondary motor cortex: Broadcasting and biasing animal’s decisions through long-range circuits”. en. In: *International Review of Neurobiology*. Elsevier, 2020, S0074774220301537. DOI: 10.1016/bs.irn.2020.11.008.
- [69] Thomas Zhihao Luo et al. *Non-canonical attractor dynamics underlie perceptual decision-making*. en. Pages: 2023.10.15.562427 Section: New Results. 2023. DOI: 10.1101/2023.10.15.562427.
- [70] Harold Hotelling. “Relations Between Two Sets of Variates”. In: *Biometrika* 28.3/4 (1936). Publisher: [Oxford University Press, Biometrika Trust], pp. 321–377. DOI: 10.2307/2333955.
- [71] Adam Kohn et al. “Principles of Corticocortical Communication: Proposed Schemes and Design Considerations”. en. In: *Trends in Neurosciences* (2020), S016622362030165X. DOI: 10.1016/j.tins.2020.07.001.
- [72] Byungwoo Kang and Shaul Druckmann. “Approaches to inferring multi-regional interactions from simultaneous population recordings”. en. In: *Current Opinion in Neurobiology*. Whole-brain interactions between neural circuits 65 (2020), pp. 108–119. DOI: 10.1016/j.conb.2020.10.004.

- [73] Stephen L Keeley et al. “Modeling statistical dependencies in multi-region spike train data”. en. In: *Current Opinion in Neurobiology*. Whole-brain interactions between neural circuits 65 (2020), pp. 194–202. DOI: 10.1016/j.conb.2020.11.005.
- [74] T. L. Veuthy et al. “Single-trial cross-area neural population dynamics during long-term skill learning”. en. In: *Nature Communications* 11.1 (2020). Number: 1 Publisher: Nature Publishing Group, p. 4057. DOI: 10.1038/s41467-020-17902-1.
- [75] Joao Barbosa et al. *Flexible selection of task-relevant features through across-area population gating*. en. Pages: 2022.07.21.500962 Section: New Results. 2022. DOI: 10.1101/2022.07.21.500962.
- [76] Petr Znamenskiy and Anthony M. Zador. “Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination”. en. In: *Nature* 497.7450 (2013). Number: 7450 Publisher: Nature Publishing Group, pp. 482–485. DOI: 10.1038/nature12077.
- [77] N Li et al. “Amotor cortex circuit for motor planning and movement”. In: *Nature* (2015).
- [78] Malavika Murugan et al. “Combined Social and Spatial Coding in a Descending Projection from the Prefrontal Cortex”. en. In: *Cell* 171.7 (2017), 1663–1677.e16. DOI: 10.1016/j.cell.2017.11.002.
- [79] Eun Jung Hwang et al. “Corticostriatal Flow of Action Selection Bias”. en. In: *Neuron* 104.6 (2019), 1126–1140.e6. DOI: 10.1016/j.neuron.2019.09.028.
- [80] Junya Hirokawa et al. “Frontal cortex neuron types categorically encode single decision variables”. en. In: *Nature* 576.7787 (2019). Number: 7787 Publisher: Nature Publishing Group, pp. 446–451. DOI: 10.1038/s41586-019-1816-9.
- [81] D. Gowanlock R. Tervo et al. “The anterior cingulate cortex directs exploration of alternative strategies”. en. In: *Neuron* 109.11 (2021), 1876–1887.e6. DOI: 10.1016/j.neuron.2021.03.028.
- [82] Chunyu A. Duan et al. “A cortico-collicular pathway for motor planning in a memory-dependent perceptual decision task”. en. In: *Nature Communications* 12.1 (2021). Number: 1 Publisher: Nature Publishing Group, p. 2727. DOI: 10.1038/s41467-021-22547-9.
- [83] Steffen B. E. Wolff, Raymond Ko, and Bence P. Ölveczky. “Distinct roles for motor cortical and thalamic inputs to striatum during motor skill learning and execution”. In: *Science Advances* 8.8 (2022). Publisher: American Association for the Advancement of Science, eabk0231. DOI: 10.1126/sciadv.abk0231.
- [84] Junchol Park et al. “Motor cortical output for skilled forelimb movement is selectively distributed across projection neuron classes”. In: *Science Advances* 8.10 (2022), eabj5167. DOI: 10.1126/sciadv.abj5167.

- [85] Mehrdad Jazayeri and Srdjan Ostojic. “Interpreting neural computations by examining intrinsic and embedding dimensionality of neural activity”. In: *arXiv:2107.04084 [q-bio]* (2021). arXiv: 2107.04084.
- [86] David Raposo, Matthew T. Kaufman, and Anne K. Churchland. “A category-free neural population supports evolving demands during decision-making”. en. In: *Nature Neuroscience* 17.12 (2014). Number: 12 Publisher: Nature Publishing Group, pp. 1784–1792. DOI: 10.1038/nn.3865.
- [87] Angela M. Licata et al. “Posterior Parietal Cortex Guides Visual Decisions in Rats”. en. In: *Journal of Neuroscience* 37.19 (2017). Publisher: Society for Neuroscience Section: Research Articles, pp. 4954–4966. DOI: 10.1523/JNEUROSCI.0105-17.2017.
- [88] Danique Jeurissen et al. “Deficits in decision-making induced by parietal cortex inactivation are compensated at two timescales”. English. In: *Neuron* 110.12 (2022). Publisher: Elsevier, 1924–1931.e5. DOI: 10.1016/j.neuron.2022.03.022.
- [89] Dimitry Fisher et al. “A Modeling Framework for Deriving the Structural and Functional Architecture of a Short-Term Memory Microcircuit”. English. In: *Neuron* 79.5 (2013). Publisher: Elsevier, pp. 987–1000. DOI: 10.1016/j.neuron.2013.06.041.
- [90] Kanaka Rajan, Christopher D Harvey, and David W Tank. “Recurrent Network Models of Sequence Generation and Memory”. In: *Neuron* 0.0 (2016), pp. 128–142.
- [91] Aaron S. Andalman et al. “Neuronal Dynamics Regulating Brain and Behavioral State Transitions”. en. In: *Cell* 177.4 (2019), 970–985.e20. DOI: 10.1016/j.cell.2019.02.037.
- [92] Lea Duncker et al. “A dynamical model with E/I balance explains robustness to optogenetic stimulation in motor cortex.” In: *Computational and Systems Neuroscience (COSYNE)*. Denver, CO, 2020.
- [93] Matthew G. Perich et al. *Inferring brain-wide interactions using data-constrained recurrent neural network models*. en. Tech. rep. Company: Cold Spring Harbor Laboratory Distributor: Cold Spring Harbor Laboratory Label: Cold Spring Harbor Laboratory Section: New Results Type: article. 2021, p. 2020.12.18.423348. DOI: 10.1101/2020.12.18.423348.
- [94] John D. Murray, Jorge Jaramillo, and Xiao-Jing Wang. “Working Memory and Decision-Making in a Frontoparietal Circuit Model”. en. In: *Journal of Neuroscience* 37.50 (2017). Publisher: Society for Neuroscience Section: Research Articles, pp. 12167–12186. DOI: 10.1523/JNEUROSCI.0343-17.2017.
- [95] Michael Kleinman, Chandramouli Chandrasekaran, and Jonathan C. Kao. *Recurrent neural network models of multi-area computation underlying decision-making*. en. Pages: 798553 Section: New Results. 2019. DOI: 10.1101/798553.

- [96] Lucas Pinto et al. “Task-Dependent Changes in the Large-Scale Dynamics and Necessity of Cortical Regions”. en. In: *Neuron* 104.4 (2019), 810–824.e9. DOI: 10.1016/j.neuron.2019.08.025.
- [97] Christian D. Márton, Simon R. Schultz, and Bruno B. Averbeck. “Learning to select actions shapes recurrent dynamics in the corticostriatal system”. en. In: *Neural Networks* 132 (2020), pp. 375–393. DOI: 10.1016/j.neunet.2020.09.008.
- [98] Matthew G Perich and Kanaka Rajan. “Rethinking brain-wide interactions through multi-region ‘network of networks’ models”. en. In: *Current Opinion in Neurobiology*. Whole-brain interactions between neural circuits 65 (2020), pp. 146–151. DOI: 10.1016/j.conb.2020.11.003.
- [99] Benjamin R. Cowley et al. *One-to-one mapping between deep network units and real neurons uncovers a visual population code for social behavior*. en. Pages: 2022.07.18.500505 Section: New Results. 2022. DOI: 10.1101/2022.07.18.500505.
- [100] Jan Putzeys et al. “Neuropixels Data-Acquisition System: A Scalable Platform for Parallel Recording of 10 000+ Electrophysiological Signals”. eng. In: *IEEE transactions on biomedical circuits and systems* 13.6 (2019), pp. 1635–1644. DOI: 10.1109/TBCAS.2019.2943077.
- [101] Fabian Pedregosa et al. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12.85 (2011), pp. 2825–2830.
- [102] Raphael Vallat. “Pingouin: statistics in Python”. en. In: *Journal of Open Source Software* 3.31 (2018), p. 1026. DOI: 10.21105/joss.01026.
- [103] Il Memming Park et al. “Encoding and decoding in parietal cortex during sensorimotor decision-making”. en. In: *Nature Neuroscience* 17.10 (2014). Number: 10 Publisher: Nature Publishing Group, pp. 1395–1403. DOI: 10.1038/nn.3800.
- [104] Adrian Gopnik Bondy et al. “Anterior and posterior striatum play distinct roles in evidence accumulation”. In: Denver, CO, 2020.
- [105] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. arXiv:1412.6980 [cs]. 2017. DOI: 10.48550/arXiv.1412.6980.
- [106] Jason Ansel et al. *PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation*. Publication Title: 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS ’24) original-date: 2016-08-13T05:26:41Z. 2024. DOI: 10.1145/3620665.3640366.
- [107] Jeffrey C Erlich, Max Bialek, and Carlos D Brody. “A Cortical Substrate for Memory-Guided Orienting in the Rat”. In: *Neuron* 72.2 (2011), pp. 330–343.

- [108] Joshua T. Dudman and Charles R. Gerfen. “Chapter 17 - The Basal Ganglia”. en. In: *The Rat Nervous System (Fourth Edition)*. Ed. by George Paxinos. San Diego: Academic Press, 2015, pp. 391–440. DOI: 10.1016/B978-0-12-374245-2.00017-6.
- [109] Xavier Glorot and Yoshua Bengio. “Understanding the difficulty of training deep feedforward neural networks”. en. In: *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. ISSN: 1938-7228. JMLR Workshop and Conference Proceedings, 2010, pp. 249–256.
- [110] Martin Abadi et al. “TensorFlow: A system for large-scale machine learning”. en. In: (2015), p. 21.

Methods

Subjects and housing

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.

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 from³⁴ 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 ± 1.15 mm 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 hardware¹⁰⁰ or the earliest test-phase IMEC acquisition hardware. We used SpikeGLX software to acquire the data. On every session we recorded from 384 sites arranged in a checkerboard pattern spanning banks 0 and 1 of the probe. The data were automatically spike sorted offline using Kilosort2 with default parameters and then manually curated using JR-Clust's GUI³³. 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 waveforms 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

For all analyses, we analyzed neurons that had an average firing rate greater than 1Hz during the stimulus period, neural responses were aligned to stimulus onset and included only up to the time of movement onset for each trial. 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 ($|\text{auc} - 0.5| > 0.05, P < 0.05$, 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 tuning curves: Following Hanks 2015¹⁶, we fit the trial-by-trial behavioral model developed in³² separately to each rat's 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, 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 of side-selective neurons ($|\text{auc} - 0.5| > 0.075, P < 0.05$). These settings yielded firing rate maps that were largely stable in time for different accumulator values. 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 \times 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 Python¹⁰¹. 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 Python¹⁰².

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 fitting¹⁰¹. 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 Generalized linear model (RR-GLM): We used RR-GLM to define the FOF \rightarrow ADS and ADS \rightarrow 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^{103,104}, 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 optimizer¹⁰⁵ in Pytorch¹⁰⁶. 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 2015¹⁶ and used standard off the shelf 50-125 μ m 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^{107,16}; here we describe procedures specific to this study. We bilaterally injected adeno-associated virus, either AAV5-CaMKIIa-eYFP-eNpHR3.0 or AAV2/5-CaMKIIa-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 $500\mu\text{m}$ apart at the 4 vertices and the center of a cross spanning AP, ML axes, centered at AP 2mm ML $\pm 1.3\text{mm}$. In each tract, 16 injections of 14.1nL were made every $100\mu\text{m}$ in depth, starting at $200\mu\text{m}$ 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 $1.128\mu\text{l}$ 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 $\pm 2.4\text{mm}$) 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 ($\sim 25\text{mW}$) 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. We call this "whole-trial" inactivation and this type of stimulation 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. We refer to them as "early" and "late" inactivation epochs respectively. 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 by following methods used in¹⁶ and¹⁴. 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 from³². 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:

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

where the variable $\mathbf{x}(t)$ is an N dimensional vector containing the activation of neural units in the network. The activations are mapped to the corresponding firing rates $\mathbf{r}(t)$ by passing \mathbf{x} through the threshold linear nonlinearity which maps individual activations to positive firing rates: x_i if $x_i > 0$ and 0 otherwise for $i = 1, \dots, N$. 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, $\mathbf{u}(t) = [s_L(t) \ s_R(t)]^T$, which represent the left clicks and right clicks respectively (Supp Fig 6 B-D). 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 :

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

Network inputs, perturbations and targets

The network received 2 inputs $\mathbf{u}(t) = [s_L(t) \ s_R(t)]^T$ 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 occurs. 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 $\sigma_S^2 = 2$ to each click, hence the input for any given left click was $-(1 + \eta)$. σ_S^2 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 task³². 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 2013³². 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 neurons¹⁰⁸. 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 patterns 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-hemisphere 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 distributions¹⁰⁹.

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 -1 s 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 Adam¹⁰⁵ in TensorFlow¹¹⁰. The methods used for analysis of activity and perturbation response of the multi-region models were identical to that used for rat data.

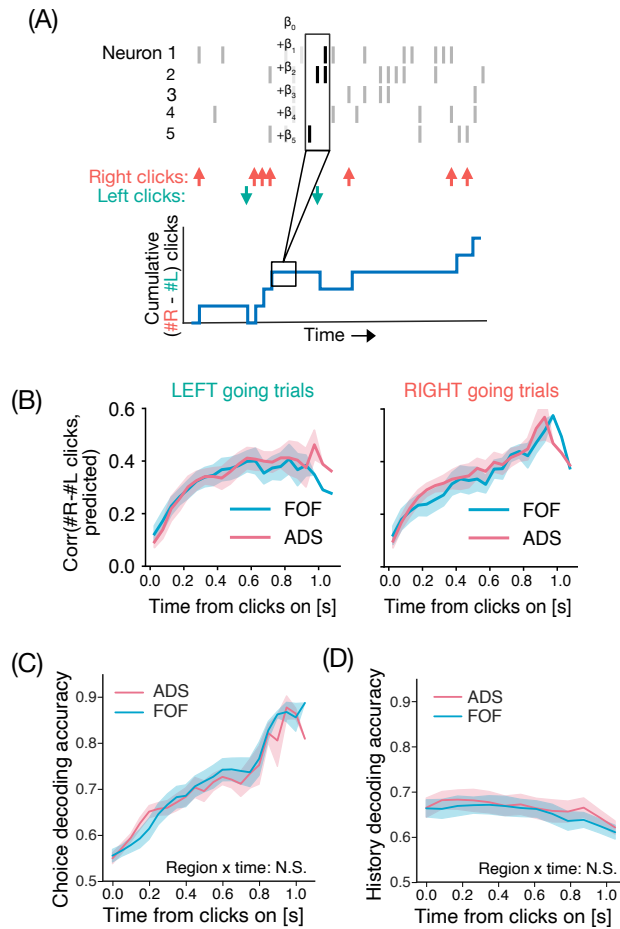
Acknowledgements

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.

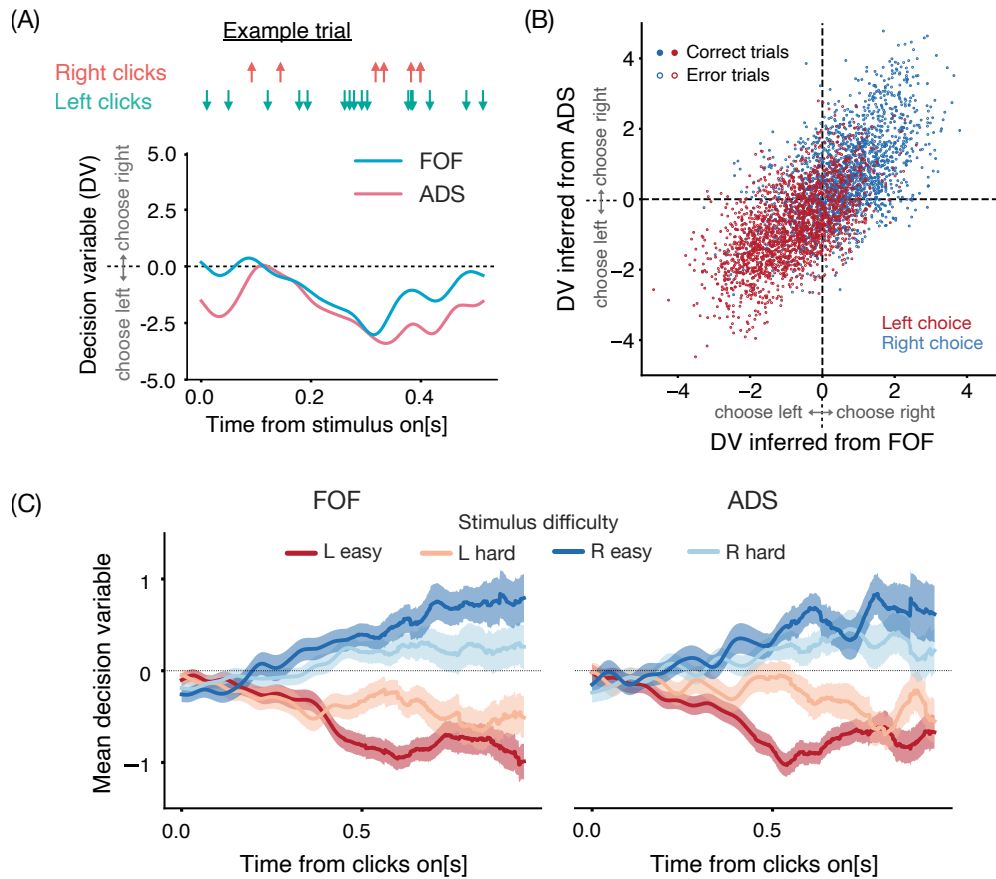
Author Contributions

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. did the analysis and modeling, and wrote the initial draft of the manuscript. All authors provided feedback on the manuscript. C.D.B. oversaw all aspects of the project.

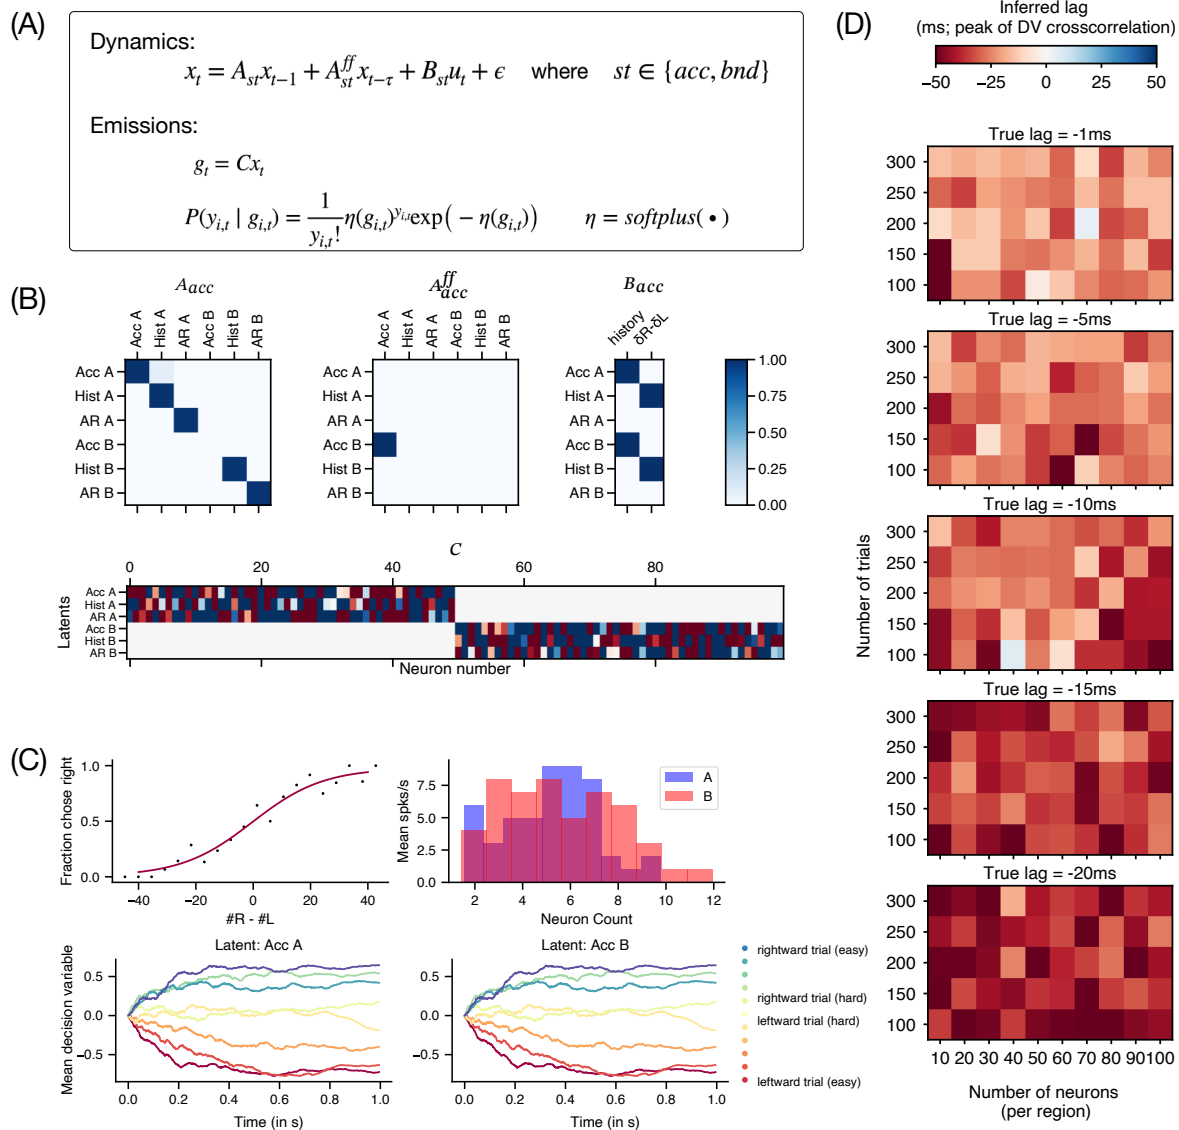
Supplementary materials



Supplementary Figure 1 **Similar strength and timecourse of stimulus, choice and history decoding in FOF and ADS** (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 \pm 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)


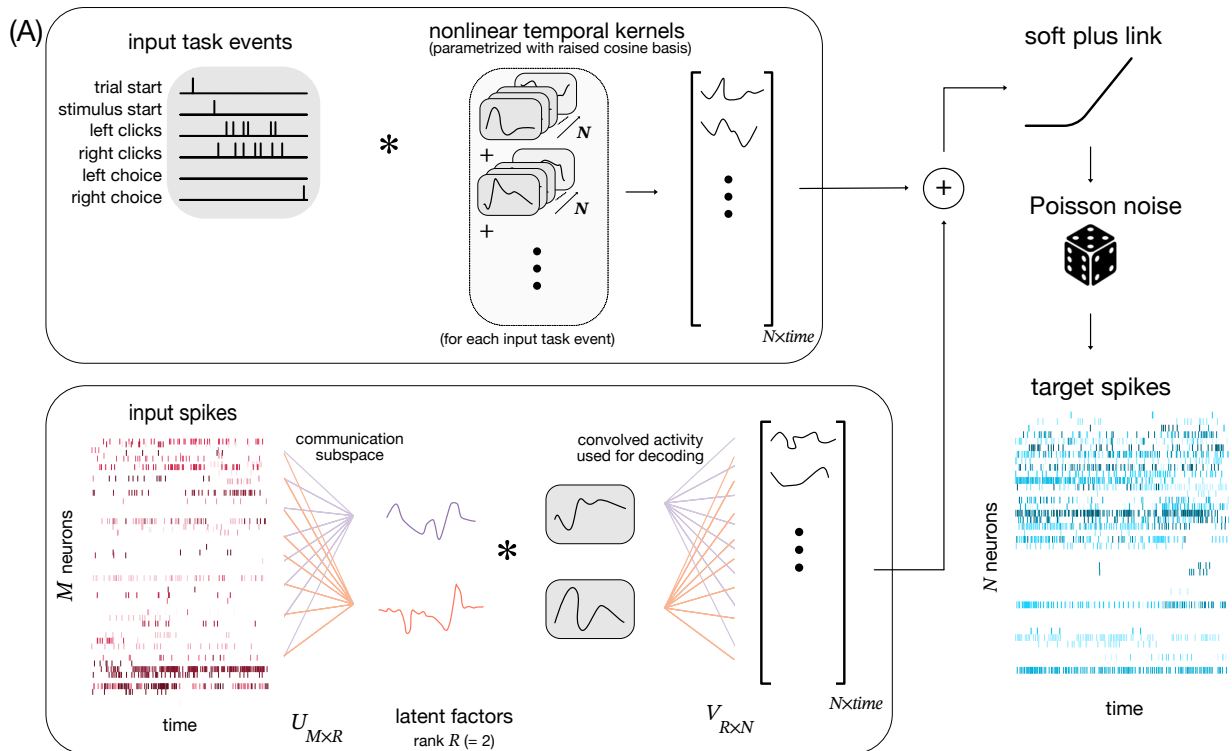


Supplementary Figure 2 **Decoding decision variable from FOF and ADS population activity** (A) Example trial showing the trajectory of the decision variable (DV) decoded from the simultaneous population activity of FOF and ADS (blue, pink line respectively), using the logistic decoder (see Methods). 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 3 Simulations used to validate the decision variable cross-correlation analysis.

(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_{acc}^{ff} 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_{acc}^{ff} 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 two accumulator latents, showing a separation by difficulty. **(D)** (Top to bottom) Inferred lag measured by the peak of the decision-variable cross-correlation, for increasing true lags between the regions as a function of number of neurons (x-axis) and trials (y-axis). 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.

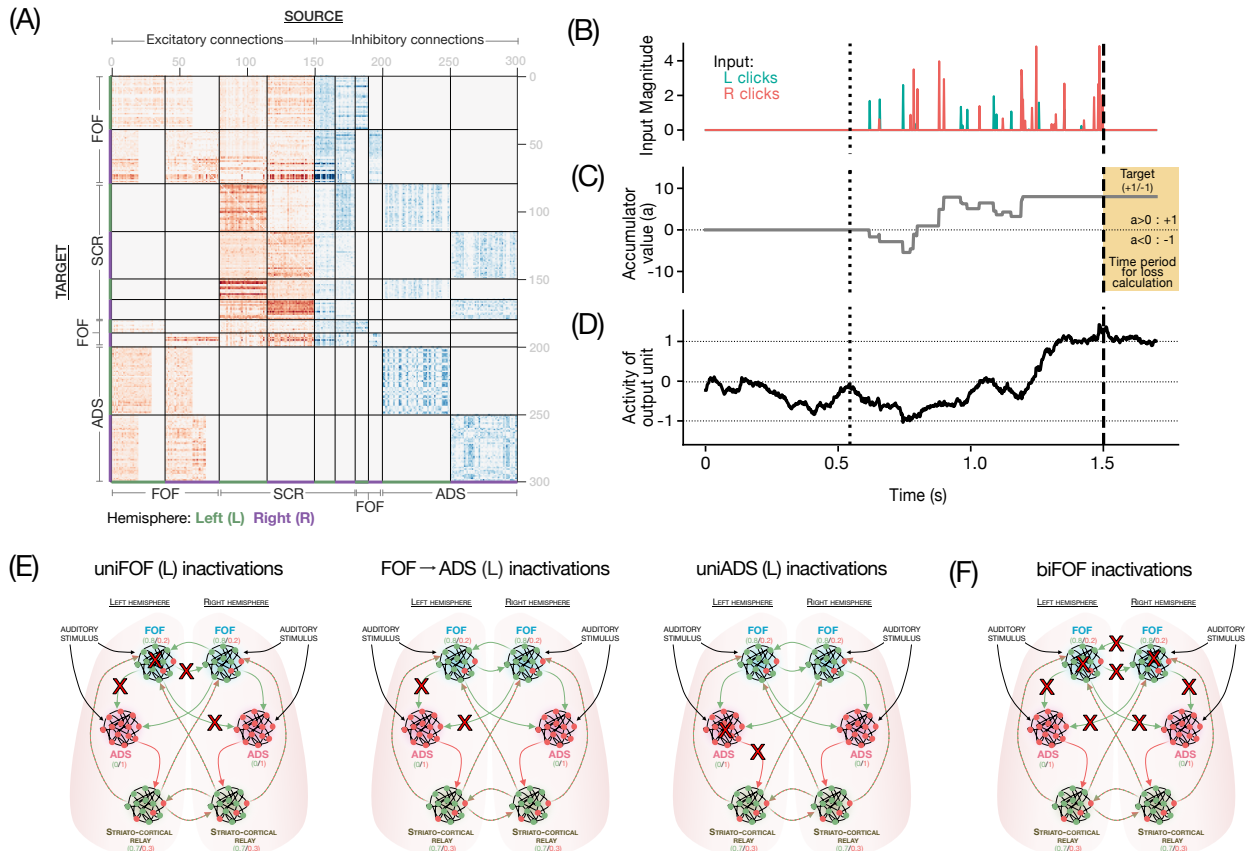


Supplementary Figure 4 **Reduced rank GLM captures the low-dimensional communication subspace between FOF and ADS** (A) Schematic of the reduced rank GLM approach. The target variables are sequences of spikes for each of N neurons in the output region (blue rasters), which are modelled as a generalized 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).

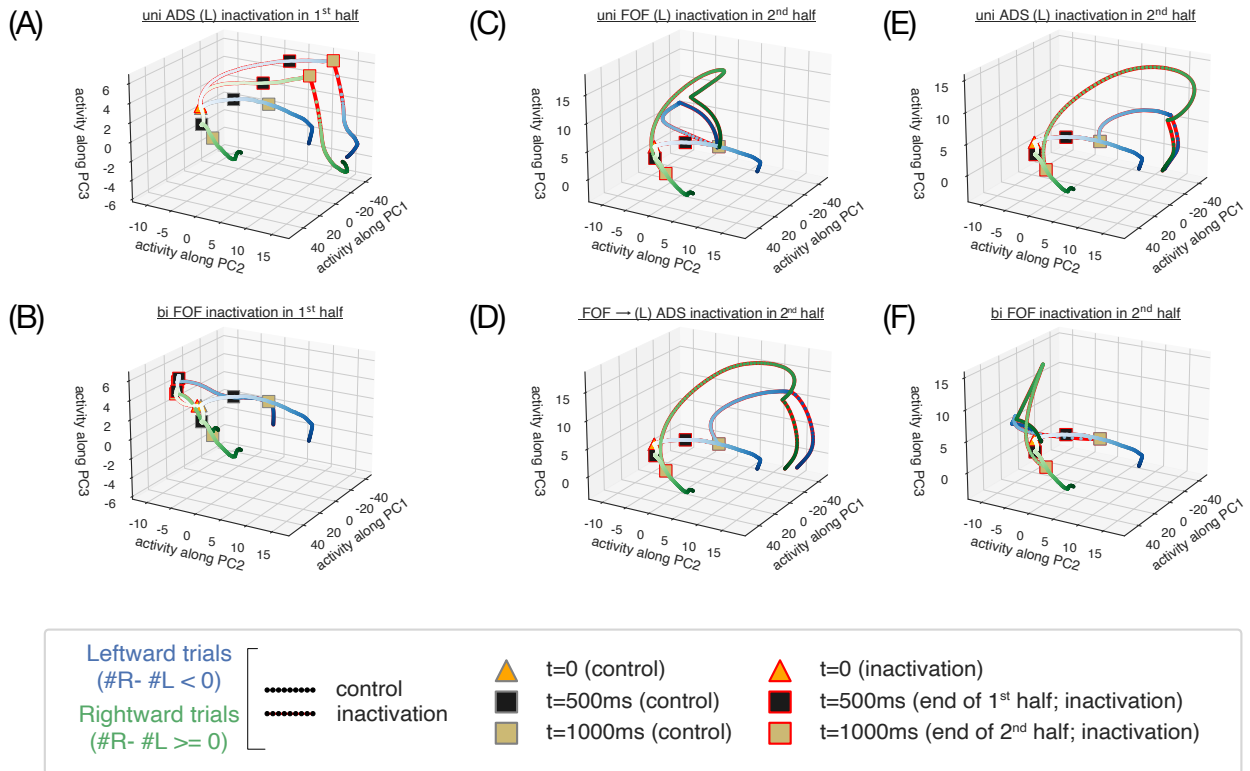
Supplementary Figure 4 (Continued from previous page) Based on both these measures, the smallest dimension required (≤ 4) was used for each session, with most sessions warranting a 4-dimensional communication subspace. (C) Cross validated r-squared values across sessions of the reduced rank GLM (ADS = 0.64 ± 0.01 , FOF = 0.73 ± 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 $R^2 < 0$. These cells were excluded while computing the mean R^2 .



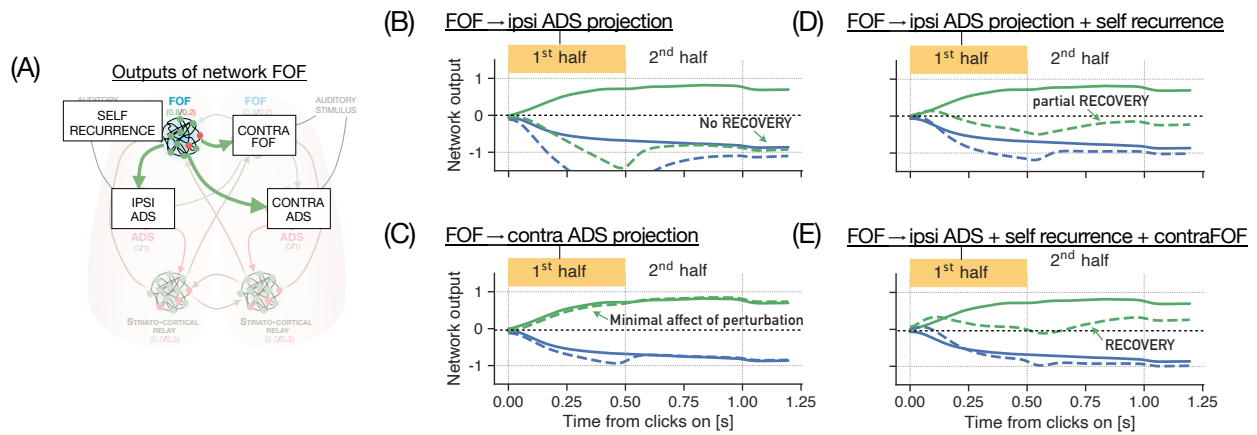
Supplementary Figure 5 **Silencing FOF axon terminals in ADS: effect on movement** (A) Silencing reduces movement times. Cumulative distribution of movement times across ($n = 6$) rats on control (black) and whole-trial inactivation (orange) trials. Movement time is defined as the time taken by the rats to leave the center port and enter one of the two side ports to report their choice. Silencing of FOF axon terminals in ADS significantly reduced the mean movement times for both choices ipsiversive (left, $P = 0.05$ Mann-Whitney U test) and contraversive (right, $P = 0.001$ Mann-Whitney U test) to the laser. (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).



Supplementary Figure 6 Training details of multi-region recurrent neural network model (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 above³². 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 \rightarrow ADS inactivations, unilateral ADS inactivations. Red crosses indicate inactivated elements. **F:** Same as E, but for bilateral FOF inactivations.



Supplementary Figure 7 Network dynamics during 1st and 2nd half inactivations 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 8 Network outputs in response to in-silico perturbations (A) Schematic of model FOF output projections from a single hemisphere, targeted for in-silico perturbations in the 1st half. Perturbations of left FOF's projections to ipsilateral ADS (B) led to network outputs diverging from control on rightward inactivation trials (dotted green) with no recovery, while perturbations of the projections to contralateral ADS (C) had no inactivation effect to begin with - suggesting that the ipsilateral ADS projections were largely responsible for inactivation effects. Additionally perturbing the self recurrence projection (D) or the contralateral FOF projection (not shown) invites partial recovery, while perturbing both (E) invites nearly complete recovery, with outputs crossing the decision threshold ($y=0$ line).