Neuron

NeuroResource

Inferring brain-wide interactions using data-constrained recurrent neural network models

Graphical abstract

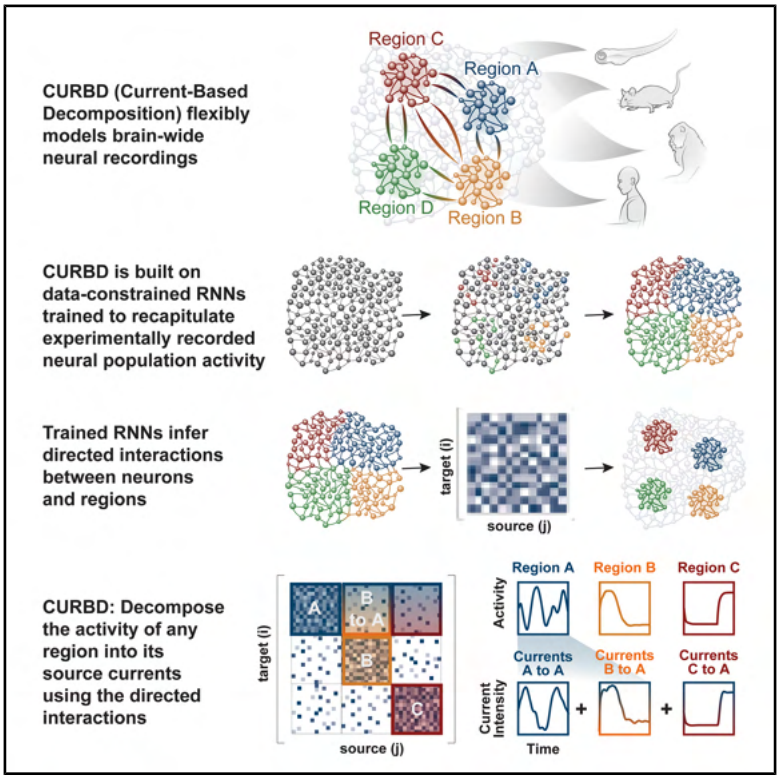

Authors

Matthew G. Perich, Charlotte Arlt, Sofia Soares, ..., Christopher D. Harvey, Karl Deisseroth, Kanaka Rajan

Correspondence

kanaka_rajan@hms.harvard.edu

In brief

Perich et al. introduce CURBD, a method for extracting multi-regional interaction structure from large-scale neural recordings. CURBD uses data-constrained recurrent neural network models to extract interactions at single-cell resolution. Perich et al. show the utility of CURBD with calcium and electrophysiology datasets from larval zebrafish, mice, monkeys, and humans.

Highlights

- The neural dynamics that drive behavior are distributed across many brain regions
- CURBD models brain-wide activity with data-constrained recurrent neural networks (RNNs)
- Reverse engineering CURBD RNNs uncovers the structure of brain-wide interactions
- We apply CURBD to recordings from larval zebrafish, mice, monkeys, and humans

Perich et al., 2026, Neuron *114*, 1–17
November 11, 2026
https://doi.org/10.1016/j.neuron.2026.07.016

CellPress

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

Neuron

CellPress
OPEN ACCESS

NeuroResource

Inferring brain-wide interactions using data-constrained recurrent neural network models

Matthew G. Perich,[1,2] Charlotte Arlt,[3] Sofia Soares,[3] Siyan Zhou,[3] Manuel Beiran,[4] Aaron S. Andalman,[5] Tyler Benster,[6] Megan E. Young,[7] Clayton P. Mosher,[7,8] Juri Minxha,[8,9] Eugene Carter,[7] Ueli Rutishauser,[8,9] Peter H. Rudebeck,[7] Christopher D. Harvey,[3] Karl Deisseroth,[5,10,11,13] and Kanaka Rajan[3,12,13,14,*]

[1]Département des neurosciences, Université de Montréal, Montréal, QC, Canada
[2]Quebec Artificial Intelligence Institute (Mila), Montréal, QC, Canada
[3]Department of Neurobiology, Harvard Medical School, Boston, MA 02115, USA
[4]Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY, USA
[5]Department of Bioengineering, Stanford University, Stanford, CA, USA
[6]Neurosciences Graduate Program, Stanford University, Stanford, CA, USA
[7]Nash Family Department of Neuroscience and Friedman Brain Institute, Icahn School of Medicine at Mount Sinai, New York, NY, USA
[8]Cedars-Sinai Medical Center, Los Angeles, CA, USA
[9]California Institute of Technology, Pasadena, CA, USA
[10]Department of Psychiatry and Behavioral Sciences, Stanford University, Stanford, CA, USA
[11]Howard Hughes Medical Institute, Stanford University, Stanford, CA, USA
[12]Kempner Institute for Natural and Artificial Intelligence, Harvard University, Boston, MA 02134, USA
[13]These authors contributed equally
[14]Lead contact
*Correspondence: kanaka_rajan@hms.harvard.edu
https://doi.org/10.1016/j.neuron.2026.07.016

SUMMARY

Behavior arises from the coordinated activity across anatomically and functionally distinct brain regions. Modern experimental tools allow unprecedented access to large neural populations spanning many interacting regions brain-wide. Yet, understanding such large-scale datasets necessitates robust, scalable computational models to extract meaningful features of inter-region communication and principled theories to interpret those features. Here, we introduce current-based decomposition (CURBD), an approach for inferring brain-wide interactions using data-constrained recurrent neural network models that autonomously produce dynamics consistent with experimentally obtained neural data. CURBD leverages the functional interactions inferred from such models to reveal directional currents between multiple brain regions simultaneously. We first show that CURBD accurately isolates inter-region currents in simulated, ground-truth networks with known connectivity and dynamics. We then apply CURBD to multi-region neural recordings obtained from many species—larval zebrafish, mice, macaques, and humans—to demonstrate the widespread applicability of CURBD in untangling brain-wide interactions and inter-area communication principles underlying behavior.

INTRODUCTION

During development, nervous systems organize into remarkably complex structures. Across the animal kingdom, brains exhibit modularity, both structural (e.g., regions and cell types) with phylogenetically determined specialization[1] and functional (e.g., interdigitated submodules within regions) with striking specialization and unique characteristics.[2] However, even at the level of individual brain regions, their activity cannot exist in isolation.[3] Regions are recurrently connected via direct projections, multi-synapse loops, and more widespread, indirect effects such as neuromodulator release.[4] Consequently, much of the brain is active during even simple behaviors that could, in theory, be mediated by only a smaller subset of regions.[5–7] Explaining the neural basis of behavior requires consideration of the distributed nature of brain-wide activity. However, we lack a comprehensive, unifying approach to infer brain-wide interactions. Here, we introduce current-based decomposition (CURBD), a computational framework that leverages recurrent neural network (RNN) models of multi-region neural recordings to infer the magnitude and directionality of the interactions between regions across the brain. While most neural data analysis and dimensionality reduction techniques[8] describe neuronal outputs (e.g., spiking activity), CURBD reconceptualizes the activity of a neural population in terms of neuronal inputs.[9,10] We apply CURBD to several experimentally obtained datasets, including calcium imaging data collected from larval zebrafish[11,12] and mice[13,14] and electrophysiology from rhesus macaques[15] and humans.[16]

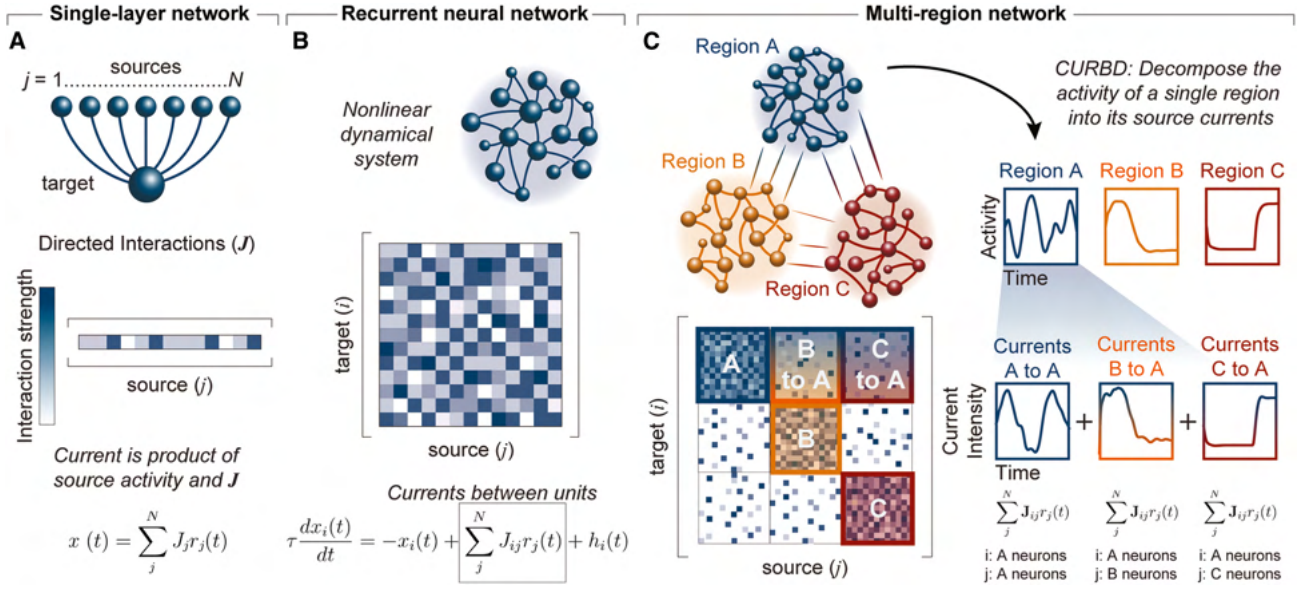

Figure 1. CURBD of multi-region interactions

(A) In a single-layer network, each of *N* source units connects to a target unit via a directed interaction weight given by vector *J*. Target unit activity, $x(t)$, derives from input currents, a weighted combination of source activity, $r(t)$.

(B) In RNN models, each unit receives inputs from and sends outputs to all other units. Directed interactions are summarized by matrix **J**, with each column containing a source unit's weights and each row those of a target.

(C) We model brains as "networks of networks" with interconnected but distinct regions. Directed interactions governing multiple regions are summarized by a single matrix **J**: diagonal submatrices correspond to within-region interactions and off-diagonal to between-region. Currents driving each target unit are the weighted sum of source activity from each submatrix. By multiplying weights in each submatrix by the source activity of each region individually, we decompose the total activity of any region (e.g., region A) into constituent source currents.

CURBD using RNNs

CURBD is built on the premise that currents between active units in a neural network can be computed from the activations and weights connecting the network. In a single-layer feedforward network, each unit is expressed as a weighted sum of the activity of the "source" units (Figure 1A). These weights correspond to interaction strengths, summarized by a vector with each source unit represented as a single entry. However, neural circuits in biological brains are typically recurrently connected,[3] yielding RNNs.[17,18] RNNs have numerous advantages, robustly generating time-varying activity patterns[10,19] with high trainability.[18,20] RNNs trained to produce desired behaviors[21–23] and tasks[24–30] or match neural data[3,11,28,31] (or both[9,32]) can be reverse-engineered to explore theories about neural computations.[33,34] As in the single-layer network, the activity of any unit in an RNN can be computed as a weighted sum of the activity of all other units in the network that provide its input (Figure 1B). The "directed interaction" matrix captures the magnitude and type (excitatory or inhibitory) of the weights between all unit pairs.

To implement CURBD, we model brain-wide circuitry as a single RNN comprising a "network of networks" across all recorded brain regions.[3] The activity of units in each simulated region of this RNN is shaped by excitatory and inhibitory source currents from all regions that provide inputs. If the connectivity relating these networks is known[9,35]—or in our case, inferred from the network model fit to time-series data—then the source currents into a target region from any other region can be estimated using only the corresponding submatrix of the directed interaction matrix and the activity of the source region (Figure 1C). When summed, these constituent source currents reconstruct the total activity observed in each neuron. However, CURBD allows the total activity to be decomposed into a set of source currents from all input regions (or cell types, submodules, etc.). Estimating population-wide source currents at this scale moves beyond simple correlations between pairs of interacting regions,[36–39] giving an unprecedented view into multi-region interactions.

Implementation of CURBD

CURBD is based on the directed interaction matrix, **J**, which estimates the effective strength and type—excitatory ($J_{ij} > 0$) or inhibitory ($J_{ij} < 0$)—of interactions between all pairs (i to j) of neurons. **J** governs the dynamical system over time but cannot be easily recorded experimentally. We thus estimate **J** through data-constrained model RNNs (Figure 2A). In brief, we first initialize a model RNN with random connectivity (see STAR Methods). Model RNNs typically contain a number of units equal to the number of neurons in the dataset to be modeled with a 1:1 correspondence, but larger or subsampled variants can be employed.[28] The goal is to learn a directed interaction matrix such that, after training, the model RNN autonomously reproduces the time-series activity of the recorded neurons given the initial state. At each time step, the output activity at the current time point *t* in the model RNN, $r[t]$, can be computed

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

CellPress
OPEN ACCESS

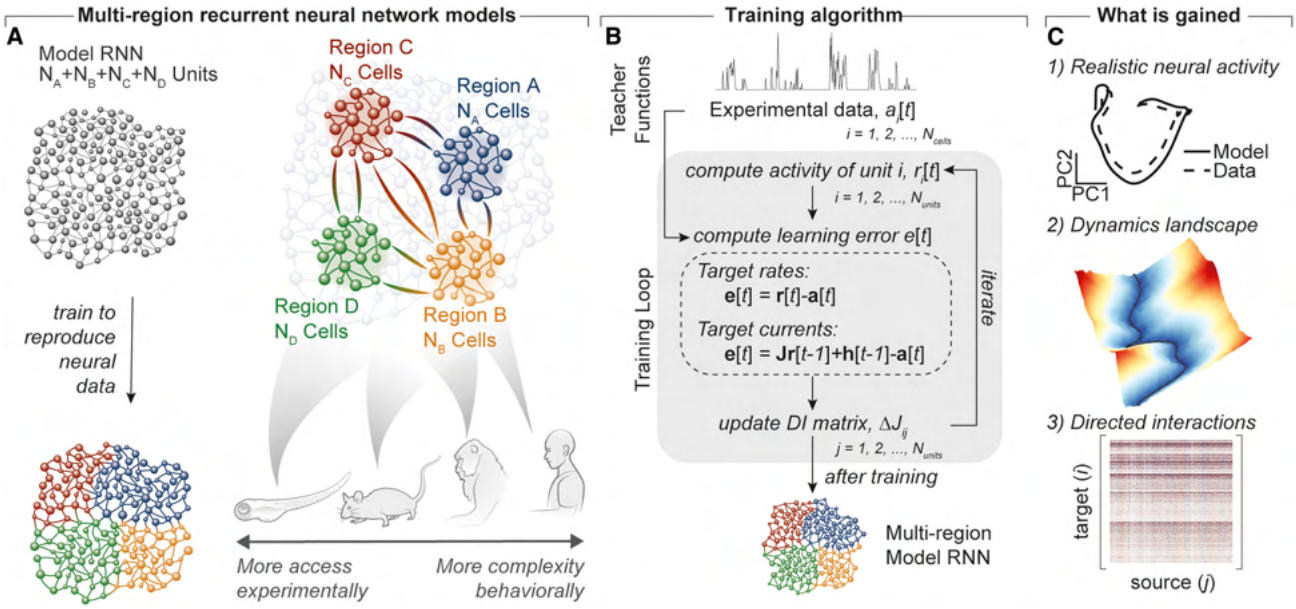

Figure 2. Data-constrained multi-region RNN overview

(A) CURBD is a flexible framework implemented through a model RNN constrained by experimental data. We construct model RNNs such that each unit is trained to match a single recorded neuron from behaving animals.

(B) Training occurs iteratively, where **J** is modified until RNN activity matches the data. Training can be in the space of rates or currents.

(C) Advantageous outcomes: (1) the trained model exhibits realistic dynamics consistent with experimental data; (2) the dynamical landscape, beyond covariance structure, can be analyzed after training; and (3) **J** gives unique insight into the functional connectivity responsible for the observed dynamics, including strength and type (excitatory/inhibitory, unidirectional/feedback) of interactions, within and across regions.

based on the sum of the previous state of population activity—$r[t\text{-}1]$, a nonlinear function ϕ of $x[t\text{-}1]$, the currents driving the RNN units—weighted by **J**. In our notation, square brackets denote discrete time, and parentheses denote continuous time.

Training proceeds iteratively[11,21,28] (Figure 2B) to minimize the instantaneous error between the activity of each recorded neuron ($a_i[t]$) and its corresponding model RNN unit. The error can be computed in the space of currents driving the RNN units ($\mathbf{x}[t]$) or the rates resulting from these currents ($\mathbf{r}[t]$). At each time step of the learning process, the directed interaction matrix **J** is updated by $\Delta\mathbf{J}$, a function of this error (Equations 1 and 2). Thus,

$$\boldsymbol{J}[t] = \boldsymbol{J}[t\ -\ 1]\ -\ \Delta\boldsymbol{J}[t] \qquad \text{(Equation 1)}$$

where the update is given by

$$\Delta\boldsymbol{J}[t] = f(\boldsymbol{r}[t]\ -\ a[t]\,)\ or\ f(x[t]\ -\ a[t]\,). \qquad \text{(Equation 2)}$$

Note that model RNNs train equally well with trial-averaged[28] or continuous neural data and that during training we make no assumptions about the identity of the modeled neurons (e.g., their corresponding brain region). After training, we obtain an *in silico* model of the experimentally recorded activity with advantageous properties (Figure 2C): (1) the model RNN generates realistic patterns of neural activity[9,11,28]; (2) training quenches the spontaneous chaotic dynamics of the randomly initialized network,[21] improving reproducibility; and (3) the trained network model yields the estimated directed interaction matrix, capturing effective functional interactions between each recorded neuron. These are notably distinct from true anatomical connectivity since they can include long-range or polysynaptic pathways and indirect effects such as neuromodulator release.

The directed interaction matrix is the heart of CURBD. The current $I_i(t)$ received by any one target unit $i = 1, 2, \ldots, N$ is the weighted sum of the N source units connected to each target unit.[9,40]

$$I_i(t) = \sum_j^N \boldsymbol{J}_{ij}\boldsymbol{r}_j(t) = \boldsymbol{J}_{i1}\boldsymbol{r}_1(t) + \boldsymbol{J}_{i2}\boldsymbol{r}_2(t) + \ldots + \boldsymbol{J}_{iN}\boldsymbol{r}_N(t) \qquad \text{(Equation 3)}$$

By definition, summing all source currents fully reconstructs the activity of the target unit. However, by restricting the summation to source units, we can isolate the currents into the target from specific sources (Figure 1C). In this paper, we focus on the case where the matrix **J** can be broken into M^2 submatrices, where M is the number of regions, corresponding to all pairs of source/target-region interactions in the region (Figure 1C), though the approach works equally well in principle when discretizing with finer granularity (e.g., by cortical layer or cell type). This separation of currents thus represents a decomposition of the activity of the target-region neurons based on the relative contributions of each source region. Direct analysis of the characteristics (e.g., strength, direction, or timing) of the disparate current inputs can help identify new functional subcircuits in multi-region data and unveil inter-region communication mechanisms that may be inaccessible through experimental measurements alone.

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

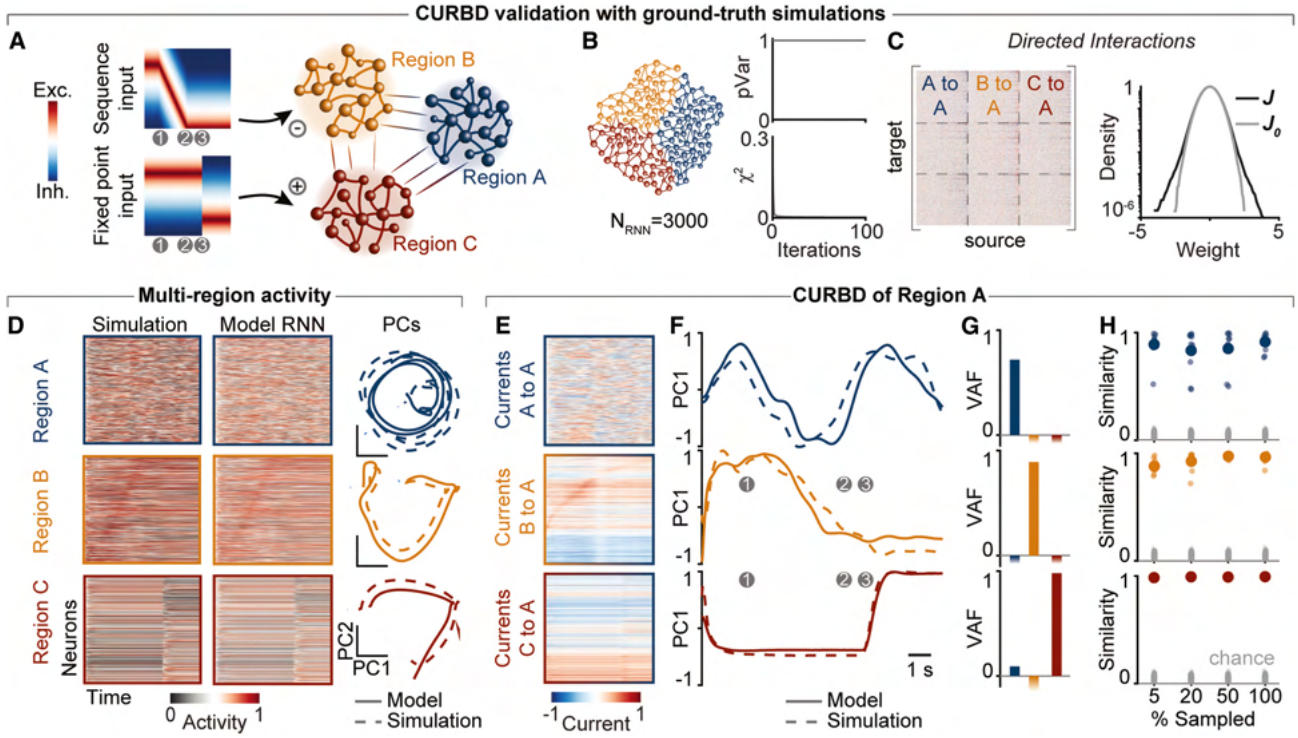

Figure 3. Validation of CURBD on ground-truth datasets

(A) We generated three RNNs representing three regions. Region B was externally driven by a sequentially active population, region C by a population generating fixed points, and region A only through interactions with B and C. Time points 1–3 (gray circles): key events in external inputs.

(B) Model RNN fit to time-series data of all three regions reached high variance explained (pVar = 0.99) and low error ($\chi 2 = 7.4 \times 10^{-6}$).

(C) (Left) Trained J. (Right) Normalized distribution (log scale) of weights in **J** (black) vs. randomly initialized $\mathbf{J_0}$ (gray).

(D) The trained model RNN reproduced single-unit activity of all three regions (left) and population trajectories in leading principal components (PCs) (right; solid: model; dashed: simulation).

(E) Individual current sources into region A showed similar patterns to source regions not apparent in its population activity.

(F) Dominant PC of each source current (solid) and ground truth from known weights (dashed). Gray circles: same events as (A).

(G) VAF for the first PC of each source current vs. ground truth ($VAF_{AtoA} = 0.72$; $VAF_{BtoA} = 0.89$; $VAF_{CtoA} = 0.98$).

(H) We randomly sampled subpopulations of the 1000-unit networks from 5% to 100% of units. Small dots: similarity from 10 repetitions; large dots: average; gray: lower bound from 100 random shuffles.

RESULTS

Validation of CURBD on ground-truth datasets

We validated CURBD in synthetic datasets with known ground-truth inter-region currents. We created a generator model comprising 3 chaotic RNNs[41] representing 3 regions and connected through sparse inter-region connection weights (Figure 3A). Region B was externally driven by a sequential external input, region C was externally driven by fixed-point inputs, and region A was driven only by inputs from regions B and C. We tuned the model such that the simulated activity in region A was visibly chaotic without any clear representation of either the sequential or the fixed-point external input driving the regions connected to it (Figure 3D).

We trained a model RNN (Figure 3B) to reproduce the activity of all 3 regions (Figures 3C and 3D) with the current-based learning rule. Using the respective submatrices of the directed interaction matrix inferred from the trained model RNN (Figure 3C), we decomposed the activity of region A into the currents from each source region. These currents reflected qualitatively similar activation patterns to those in the source regions (Figure 3E)—sequential and fixed point, respectively—even though these patterns were not readily apparent in the activity of region A. We summarized the population-wide currents inferred by CURBD using principal-component analysis (PCA) and compared them to the ground-truth currents in region A using variance accounted for (VAF). CURBD accurately reconstructed each current source driving region A (Figures 3F and 3G).

We performed several control analyses to confirm CURBD's performance. First, we confirmed CURBD uncovers ground-truth interactions in the presence of Poisson noise (Figure S1C) and in the presence of uncorrelated distractor inputs (Figure S1D). Second, we compared the performance of CURBD to a correlation-based method (canonical correlation analysis [CCA]), which has previously been used to identify individual subspaces that capture the shared dynamics.[42,43] CCA did not accurately capture the ground-truth currents (Figure S1E), as it was easily confused by the recurrence in the generator network.[44] Third, we showed that CURBD is robust to undersampling of the true population.

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

In real datasets, despite explosive advances in recording technologies,[6,7,45–48] experimenters typically only have access to a small percentage of neurons in a given region.[49] CURBD accurately estimated the current dynamics even when the network was highly undersampled, as low as 5% of the population, based on a similarity metric between the inferred dynamics and the ground-truth dynamics.[23,42,50]

Lastly, we explored possible failure modes of CURBD with a second simulation. We created 2 reciprocally connected RNNs, each receiving sinusoidal inputs of different frequencies (Figures S1F–S1L). Since the sinusoidal inputs can mix with the ongoing chaotic dynamics in recurrent networks,[41,51] they provide a more challenging paradigm to assess CURBD. We found that CURBD was most effective when the intrinsic dynamics of the two RNNs were distinct—resulting from differences in external drive, internal structure, or connectivity, or both—with sparse inter-region connectivity (Figure S1L). Altogether, our ground-truth simulations illustrate that CURBD can infer unobserved source currents between multiple brain regions under a range of experimentally realistic conditions.

Comparing different learning rules

The ground-truth datasets above used current-based learning to train model RNNs to match currents flowing through the simulated RNNs.[11,28] However, many real datasets capture the outputs (e.g., spikes) of neurons. Furthermore, current-based fitting makes assumptions about the nonlinearity in the biological neuron (a specific mathematical function[52]), which needs to be inverted to compute the error (see Materials and Methods of Rajan et al.[28]). Consequently, rate-based fitting makes fewer assumptions and might infer a more unbiased estimate of the population interactions. We thus compared current- and rate-based learning on an experimental dataset (Figures 4A, 4B, and S2). In brief, we optically recorded fluorescence from genetically encoded calcium sensors to simultaneously track the activity of thousands of neurons in five larval zebrafish expressing GCaMP6.[11,53] We fit model RNNs to the neural recordings from the telencephalon and thalamus of each fish (Figures 4C and 4D), chosen based on the high activity variance and large numbers of neurons. We compared the model RNN fits obtained using both rate-based and current-based learning (see STAR Methods) for five training runs starting from $\mathbf{J_0}$ (Figure 4E). All initialization and input parameters were identical for both learning rules. We found that the two learning rules gave comparable performance overall. Rate-based learning, however, had two advantages: (1) higher explained variance (Figure 4F) and (2) greater consistency across runs (Figure 4G). Given the qualitatively similar current estimates from both learning rules (Figures 4H–4J), we chose to employ rate-based learning for the following applications of CURBD.

Decomposition of three regions from larval zebrafish during behavioral challenge

Here, we use CURBD to infer mechanisms underlying responses to a behavioral challenge paradigm in larval zebrafish. In brief, larval zebrafish (6 days post fertilization) expressing GCaMP6s[11,53] were partially embedded in agar and presented with inescapable electric shocks (Figure 5A). Activity from 3 regions (habenula, raphe nucleus, and telencephalon[54]) was simultaneously imaged (Figures 5B and S3A–S3D). Previous work demonstrated that the inescapable shock paradigm induced depression-like behaviors in the fish.[11] As stress accumulates, the fish gradually lapse to passivity, a maladaptive state that minimizes energy expenditure, akin to learned helplessness in depression.[55,56] This behavioral state transition is mediated by interactions between the habenula and the raphe nucleus,[11] though it remains unclear how habenula is recruited during this process.

We applied CURBD to infer the multi-regional inputs that drive habenula to cause the behavioral state transition. We fit model RNNs to the 3 regions and performed CURBD to isolate the current sources driving habenular activity (Figures 5C–5G). We then studied the relationship between the fish's behavioral response and the temporal dynamics of each current source to the habenula during the behavioral challenge paradigm (Figure 5H). We found that the shocked fish gradually decreased their tail movements compared with a control group as they adopted a passive coping strategy (Figure 5I). Strikingly, we found two timescales of habenular inputs corresponding to distinct behavioral states: initial habenular activity driven primarily by the raphe nucleus (Figures 5J, S3E, and S3F) followed by ramp-like recruitment of habenular neurons by the telencephalon and intra-habenular inputs as the fish transitioned into passive coping.

The model RNNs give us access not only to the estimated current dynamics between regions but also to the strengths of interactions via the weights of the interaction matrix. We analyzed the distribution of inferred interaction weights between the raphe nucleus and habenula before and after passivity (Figure 5K). Before passivity, a phase marked by large inputs from the raphe nucleus to the habenula, had no change in weights compared with the pre-shock baseline. After the transition to passivity, we saw a clear change in interaction weights. The joint insights from the current dynamics and interaction weights show that the habenular responses to the behavioral challenge are driven by two distinct mechanisms operating on two different timescales.

CURBD untangles brain-wide currents during spontaneous movement in mice

Here, we demonstrate that CURBD untangles source currents from multi-region calcium imaging in mice.[13,14] Mice expressing GCaMP6s[57] spontaneously ran on top of an air-supported ball in complete darkness (Figure 6A). Using a large field-of-view two-photon microscope,[58] we imaged neural activity simultaneously from four regions (Figures 6A and 6B): primary visual cortex (V1), secondary motor cortex (M2), posterior parietal cortex (PPC), and retrosplenial cortex (RSC). Together, these regions contribute to complex functions such as navigation, decision-making, and movement.[59–61] Mice exhibited spontaneous bouts of running behavior (Figure 6C), with complex patterns of neural activity observed across all four brain regions during these bouts. Consistent with recent studies (e.g., Musall et al.,[6] Stringer et al.,[7] and Steinmetz et al.[62]), we observed a high degree of activity even in V1 despite the fact that the mice received little visual input.

We tested whether CURBD could isolate different sources of behaviorally relevant information in regions such as V1. We trained model RNNs to reproduce the neural data from the 4

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

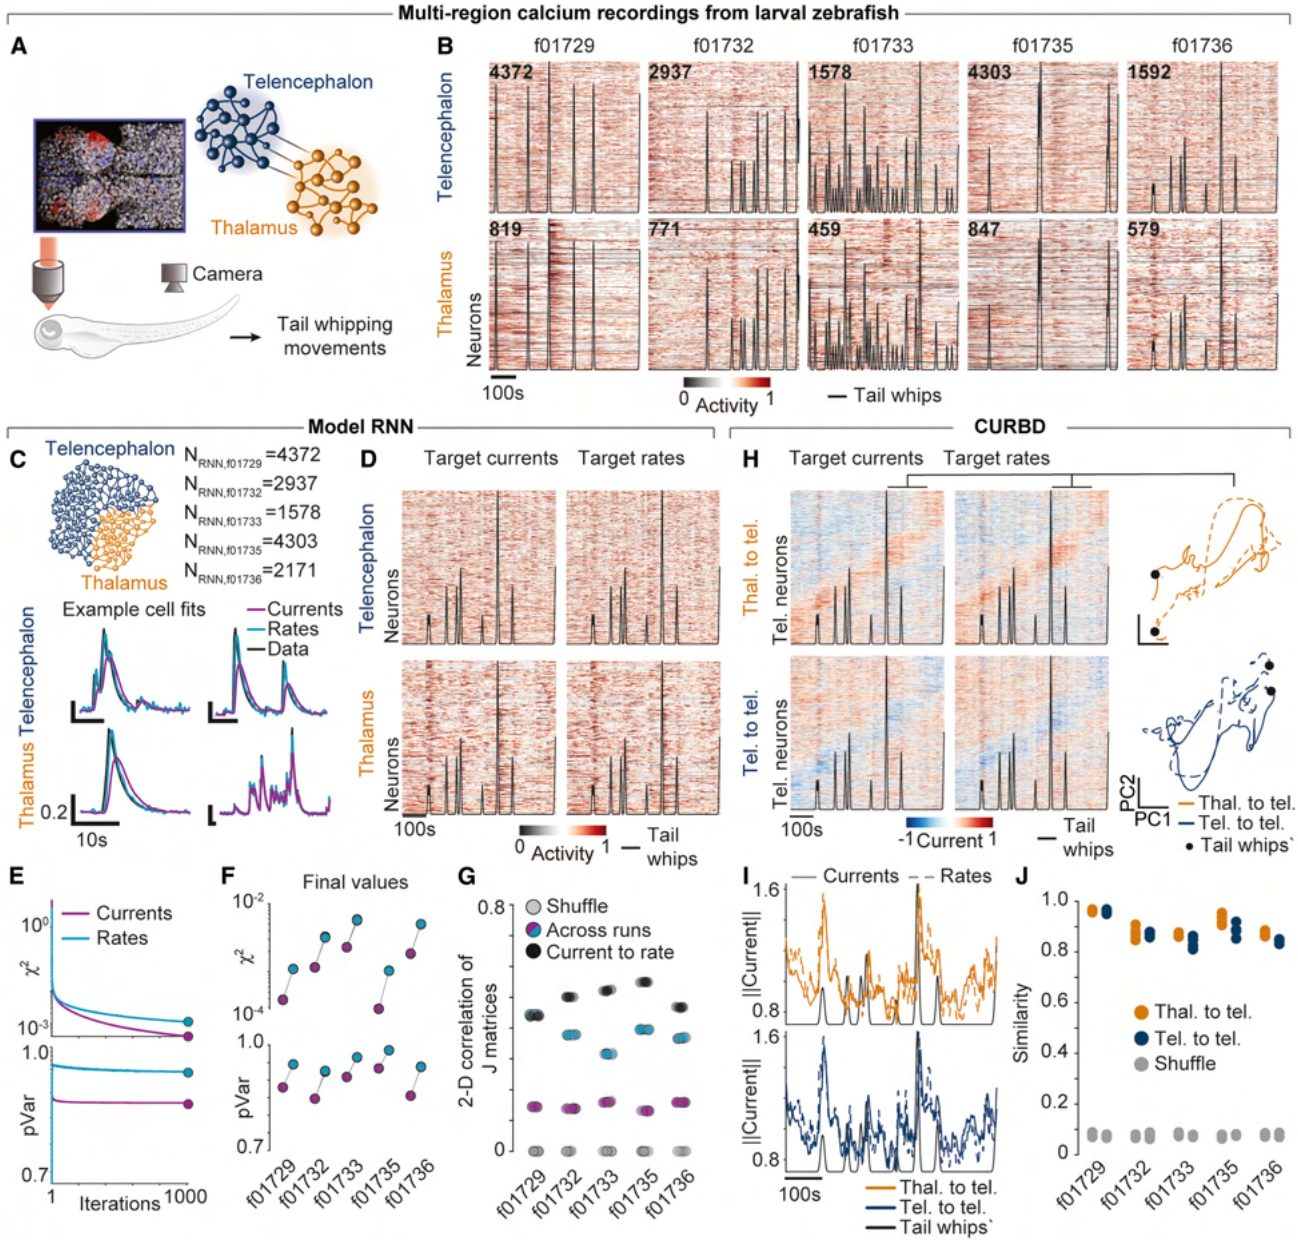

Figure 4. Comparing learning rules using calcium recordings from larval zebrafish

(A) Five larval zebrafish expressing GCaMP6s were head-fixed with tails free for multi-region imaging. Tail-whip timing was recorded with an overhead camera.

(B) Example activity (color bar) of telencephalon and thalamus populations for the five fish. Black trace: tail whips. Inset: neuron count.

(C) (Top) Model RNNs fit to telencephalon and thalamus. (bottom) Example single neuron fits from the telencephalon (top) and thalamus (bottom) using current-based (magenta) and rate-based (blue) learning.

(D) Model RNN activity (f01736) for fits to (B) using current- (left) and rate-based (right) learning.

(E) Example training curves from f01736 for the two rules.

(F) Final performance for both rules across five fish.

(G) 2D correlation of J matrices comparing rules on the same run (black), across runs with current- (magenta) or rate-based (cyan) training, against shuffled (gray). Similarity between the two rules ranged from 0.45 to 0.6.

(H) Inter-region current dynamics inferred by CURBD using both rules. (left) Heatmaps: inhibitory (blue) and excitatory (red) currents driving each telencephalon neuron, sorted by peak excitatory thalamic input time. (right) Leading two PCs comparing current (solid) and rate (dashed) training for a brief window (black bars). Black dots: state at each tail whip.

(I) Magnitude of current inputs from current (solid) and rate (dashed) training.

(J) CCA similarity (STAR Methods) between CURBD output using current and rate training. Currents were highly similar (mean correlations > 0.8) across all five fish.

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

CellPress
OPEN ACCESS

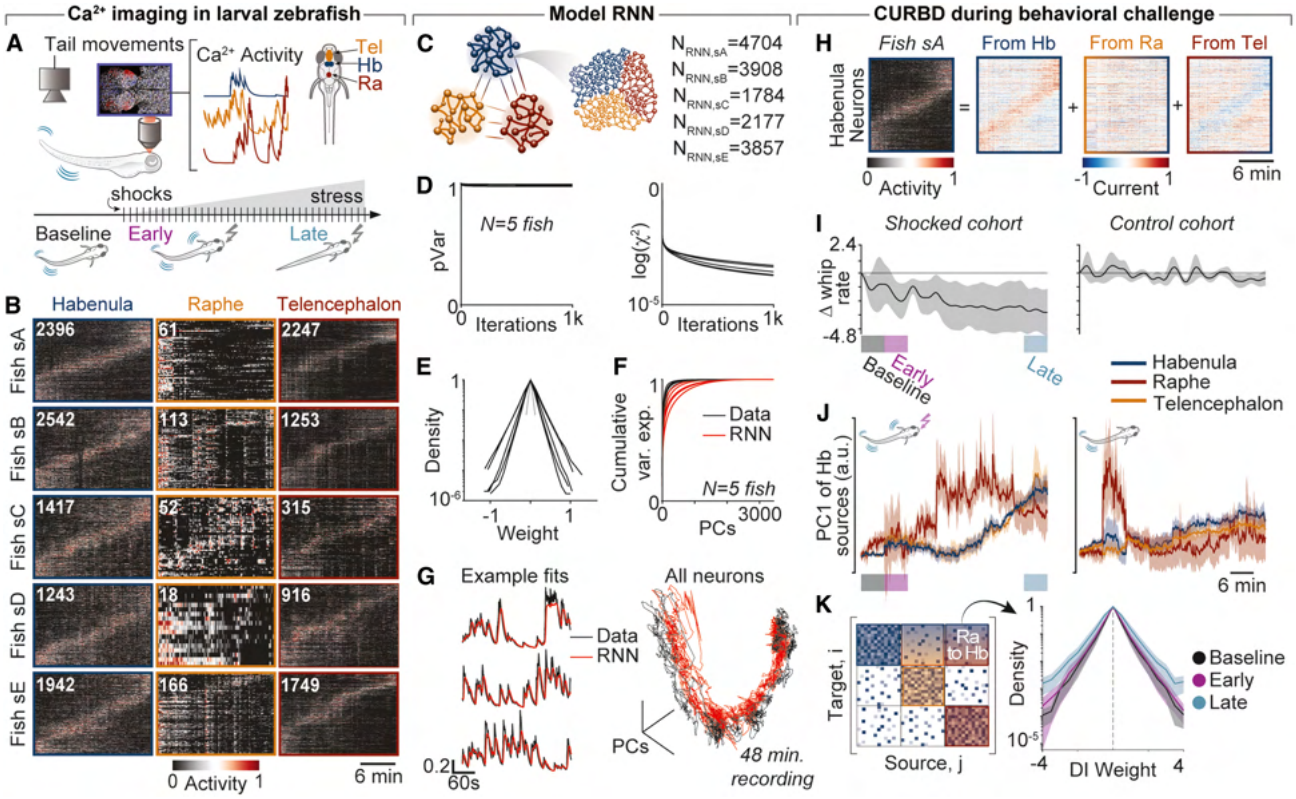

Figure 5. Isolating source currents into the habenula from multi-region calcium recordings in larval zebrafish

(A) (Top) Multi-regional Ca2+ imaging in larval zebrafish expressing GCaMP6s with simultaneous tail recording. (Bottom) Fish experienced gentle but inescapable shocks, which, over time, induced a behavioral state transition toward passivity.

(B) Neural population activity from the three regions during 48 min of recording in five shocked fish. Inset: neuron count.

(C) Model RNNs fit sequentially in 6-min epochs to all neurons from the three regions.

(D) pVar (left) and $\chi 2$ (right) for the baseline epoch of all five shocked fish.

(E) Normalized distribution of weights (log scale) for the five model RNNs before (gray) and after (black) training.

(F) Cumulative variance explained by all PCs for data (black) and model RNNs (red), and each line is one fish.

(G) (Left) Example single neuron fits from Fish sA: data (black), model RNN (red). (Right) Leading three PCs over 48 min for the example fish.

(H) CURBD decomposition for Fish sA. (Left) Heatmap of RNN activity for the habenula. (Right) Heatmaps for each of three source currents into habenula.

(I) Change relative to baseline of tail-whip movements over time in shocked (left) and control (right) cohorts. Mean ± SEM across 5 fish per cohort. Blocks: baseline (black), early shock (magenta), late shock (cyan).

(J) Temporal dynamics of the first PC capturing inputs to habenula from habenula (blue), raphe (red), and telencephalon (yellow). Mean ± SEM across 5 shocked (left) and control (right) fish.

(K) Normalized distribution of weights in the raphe-to-habenula submatrix for baseline (black), early (magenta), and late (cyan) epochs. Mean ± SEM across 5 shocked fish.

regions (Figures 6D–6F). Applying CURBD, we identified strikingly different patterns of excitation and inhibition during running bouts for the sixteen source currents (Figures 4G, S4A, and S4C). We focused on the currents into V1 to isolate sources of movement-related signals, computing the relative variance explained by each source current of the full V1 population activity (Figure 6H). We originally predicted that currents from M2 and PPC, which are closely involved in planning and producing behavior,[59,63] would increase during bouts of running. However, we saw no clear relationship between the variance captured by each source current and the running speed. Instead, source currents into V1 from each of the four brain regions occurred in similar proportions, indicating that currents flow between the four regions even while the mice are at rest.

To analyze the population-wide dynamics of these currents, we computed separate low-dimensional neural manifolds[8,64–66] spanning each source current using PCA. The trajectories within these manifolds (Figures 6I, S4B, and S4D) capture the dominant dynamics of each source current into the target region. Studying the dynamics of the source currents into V1 (Figure 6I), we observed that M2 and PPC, but not RSC, currents showed large deviations in their trajectories during running bouts. Using linear decoders that predicted running speed based on each source current to V1, we found that the currents from M2 contained the most information about running speed (Figures 6J–6L and S5). These results illustrate the potential of CURBD to untangle the complex, multi-regional interactions underlying behavior using RNNs based on multi-region calcium imaging data.

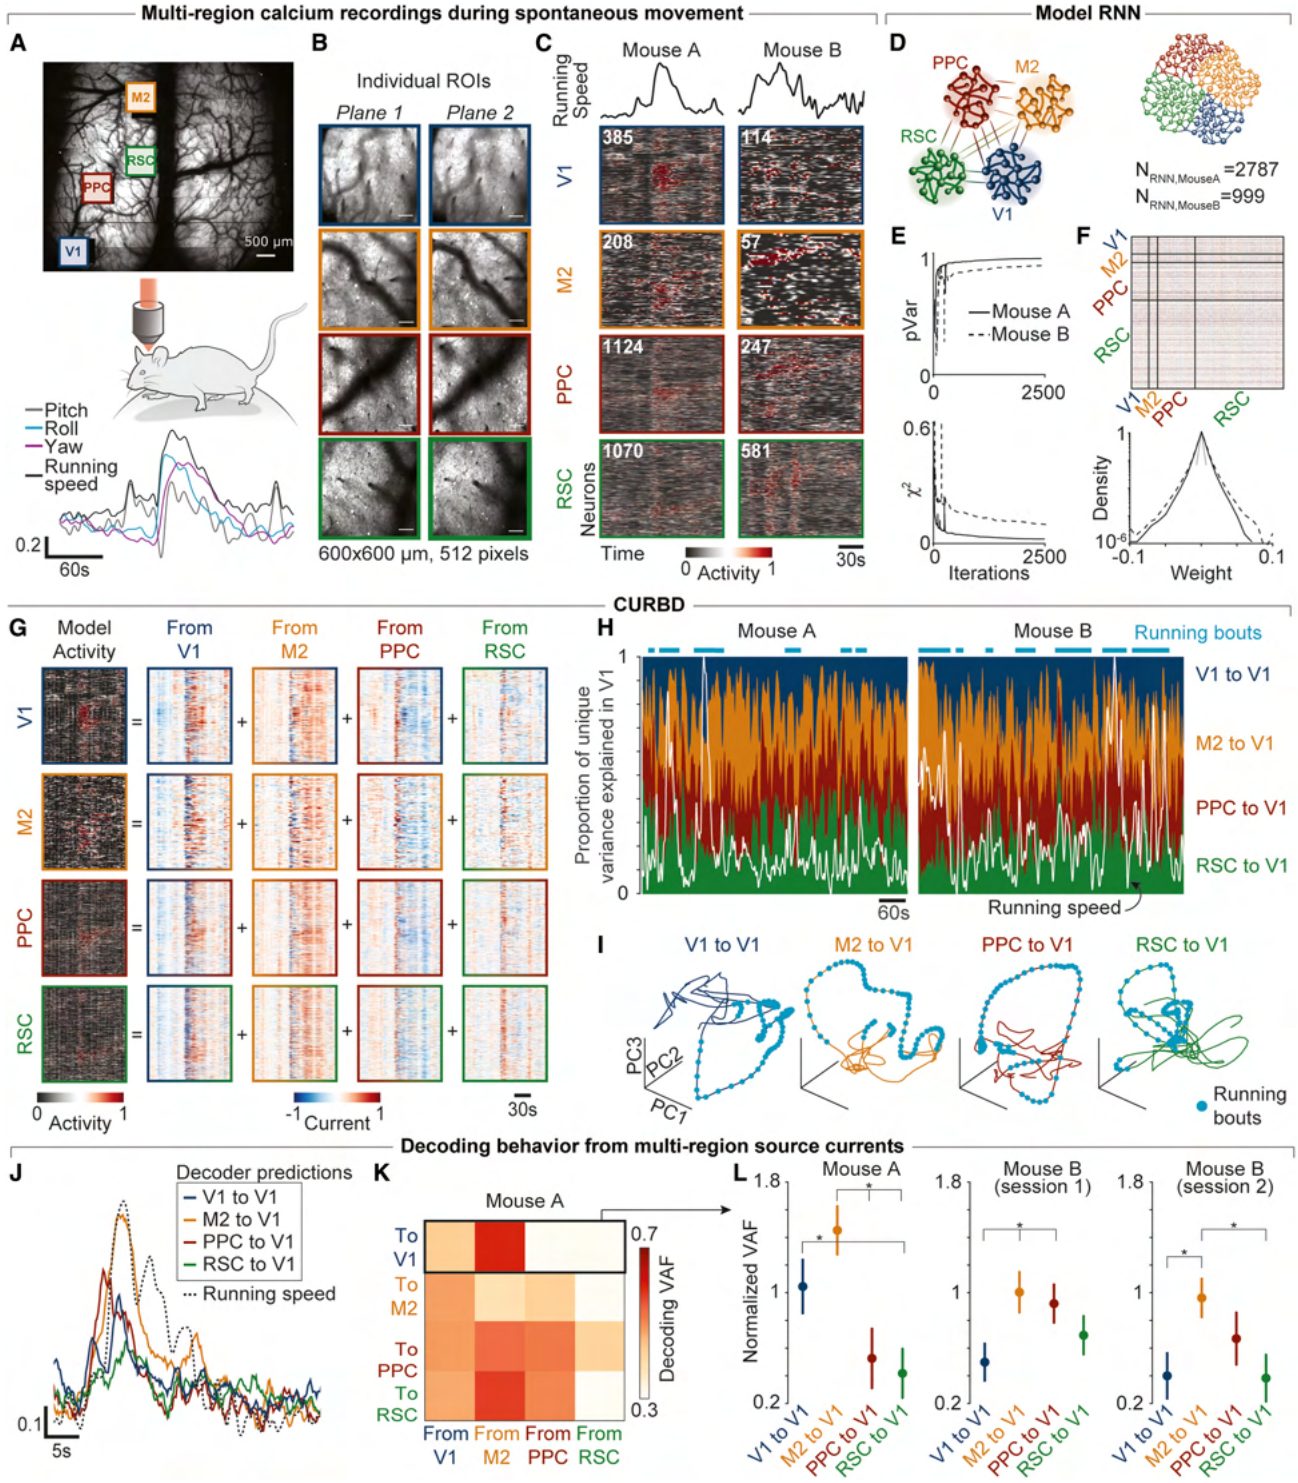

Figure 6. Isolating source currents from multi-region calcium recordings in mice

(A) We recorded neural activity from four regions in two head-fixed mice expressing GCaMP6s on an air-supported ball in complete darkness. Running was tracked via sensors recording pitch (gray), roll (cyan), and yaw (magenta). The magnitude gives running speed (black).

(B) Two example imaging planes per region (regions of interest, ROIs).

(C) Behavior and neural activity during spontaneous running for two mice. Inset: neuron count.

(D) We fit a model RNN to 10 min of activity per mouse.

(legend continued on next page)

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

CellPress
OPEN ACCESS

CURBD applied to spiking data collected during Pavlovian conditioning in monkeys

We next applied CURBD to spiking activity acquired by electrophysiological recordings in macaques. We obtained neural data from two monkeys (*Macaca mulatta*) performing a Pavlovian conditioning task (Figure 7A).[15] The monkeys learned to associate three conditioning stimuli with three different reward levels of increasing desirability: no reward, water, and juice. Using intracranial electrodes, we recorded from neurons in the amygdala, subcallosal anterior cingulate cortex (ACC), and rostromedial striatum, regions important for reward processing and affective behaviors.[67] Since we could record only small numbers of neurons on a given session, we constructed "pseudopopulations" by averaging neural activity across all trials for each condition on each session (343 neurons for monkey D and 199 neurons for monkey H). All three regions also displayed condition-specific, time-ordered sequence-like activity (Figures 7B and S7).

We trained model RNNs to reproduce the neural data from all three regions for each monkey (Figures 7C and S6). The models fit well even though the neurons were not simultaneously collected, highlighting the potential of CURBD for pseudopopulation data. The nine source currents inferred by CURBD showed distinct dynamics for each region in the circuit (Figure 7D). Prior work shows that neurons in the ACC directly project to the rostromedial striatum,[68] and we focused our analysis on these two regions. Intriguingly, we found that the strength of interactions between the striatum and ACC was also asymmetric (Figures 7E and S6). CURBD applied to this pathway revealed a high magnitude of currents from the striatum to the ACC following the water stimulus (Figures 7F and 7G), but no corresponding current from the ACC to the striatum. However, on the juice trials, we saw strong bidirectional currents, indicating context-specific computations. Crucially, the currents inferred through CURBD were consistent across RNNs based on data from the two monkeys, as well as across five different random initializations of the model RNN.

Since the pseudopopulations are constructed post hoc, their size and the specific neurons that are chosen for inclusion in the population can sometimes be prone to empirical biases or the specific limitations of the experiments. We tested whether CURBD infers the same population-wide current dynamics with pseudopopulations constructed by sampling different subsets of neurons from the full population by subsampling neurons from each region to create pseudopopulations of different sizes (between 60% and 90% of the total numbers recorded). Using the similarity metric employed in the ground-truth and zebrafish datasets above, we found a high degree of similarity in the identified currents even when using just 60% of the available neurons (Figure S7). Thus, CURBD can be readily applied to pseudopopulations comprising non-simultaneous recordings, yielding robust estimates of the interactions between regions, even under different partial sampling conditions.

CURBD applied to single-cell spiking data from humans during memory retrieval

We next applied CURBD to multi-region spiking data from five human participants performing a set of two memory tasks in eight interleaved blocks (Figures 8A–8C).[16] In the first task, participants categorized images based on high-level sensory features. In the second, participants were presented with an image and reported whether or not they had seen the image before. As the participants performed these tasks, we recorded the activity of neurons in two frontal cortical regions—the pre-supplementary motor area (preSMA) and the dorsal ACC (dACC)—and two subcortical regions—the hippocampus and amygdala (H/A)—using hybrid depth electrodes.[69] We constructed pseudopopulations from neurons recorded from between two and five sessions in each participant. Since some datasets had few recorded neurons in either the hippocampus or the amygdala, we combined data from the two subcortical regions.[16] We focused our analyses on the memory task. Since memory retrieval may be mediated by interactions between frontal cortices and the H/A,[16] we hypothesized that CURBD can separate currents related to memory retrieval and memory formation within these regions.

We fit model RNNs to the pseudopopulation datasets from each participant (Figures 8D–8F and S8) to estimate the directed interaction matrices. We then performed CURBD to infer the currents driving H/A following the presentation of familiar or novel images. The state space trajectories in the first two PCs of the full H/A activity and each source current showed distinct current dynamics between the categorization and memory tasks after image presentation (Figure S9). Within the memory task, the currents within the circuit also changed between novel vs. familiar image conditions, with familiar images causing a small response in the frontal cortex to H/A currents and novel images causing a large response in all currents (Figure 8G). We quantified these effects using the Mahalanobis distance from the cluster of resting state activity (Figure 8H). Across all five participants, CURBD identified change in currents (relative to baseline) from preSMA to H/A when viewing familiar images and smaller changes in the other source currents. Viewing novel images, on the other hand, caused large sustained currents throughout the whole

(E) pVar and $\chi 2$ for RNNs trained on mouse A (solid) and two sessions from mouse B (dashed).
(F) (Top) Example **J** for mouse A. (Bottom) Normalized distribution of weights (log scale) for the 3 **J** matrices before (gray) and after (black) training.
(G) CURBD decomposition for mouse A. (Left) Heatmaps of RNN activity for the four regions. (Right) Heatmap of current decomposition for each of the sixteen sources.
(H) Proportion of unique variance explained by each source current of total V1 activity. White: running speed; cyan: running bouts.
(I) V1 source current trajectories in three leading PCs for mouse A. Cyan dots: running bouts.
(J) Linear decoders predicting running speed from each source current. Example predictions (colored) vs. measured speed (dashed) for the four source currents into V1, mouse A.
(K) VAF for all sixteen source current decoders for mouse A.
(L) Decoder performance (mean and s.d. across 1,000 cross-validated test sets; see STAR Methods) for source currents into V1 for mouse A and two sessions from mouse B. *: significance at 95% confidence interval (CI) on VAF.

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

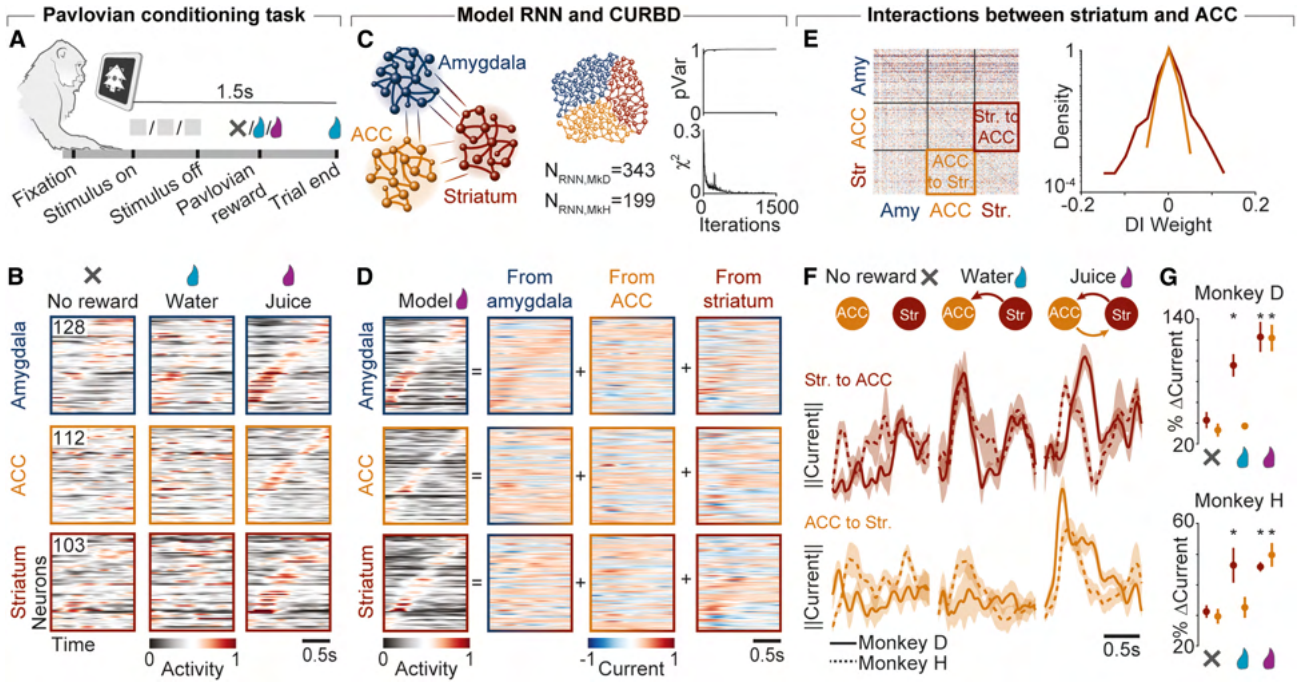

Figure 7. CURBD of three-region pseudopopulation recordings in monkeys

(A) Monkeys performed a Pavlovian conditioning task with three stimuli (unconditioned, water, juice).

(B) Trial-averaged firing rates in the pseudopopulation for monkey D for amygdala, subcallosal ACC, and striatum during unconditioned (left, inset: neuron count), water (middle), and juice (right) stimuli. Neurons aligned on stimulus presentation, sorted by peak activity in the juice condition.

(C) Model RNNs fit to the three-region dataset. (Left) Schematic. (Right) pVar (top) and $\chi 2$ (bottom) vs. training iterations.

(D) CURBD of activity in each region for juice trials. Left heatmaps: full model RNN activity. Right heatmaps: decomposition for each of sixteen source currents.

(E) (Left) **J** for an example model RNN from monkey D. (Right) Distribution of weights (log scale) in striatum-to-ACC (red) and ACC-to-striatum (yellow) submatrices.

(F) Magnitude of bidirectional currents from striatum to ACC (red, top) and ACC to striatum (yellow, bottom) during the three stimuli. Solid: monkey D; dashed: monkey H. Error bars: s.d. across five initializations. Schematics (top): dominant currents (magnitude and direction) inferred by CURBD.

(G) Percent change in total current in the first 1 s of each condition vs. mean in the unrewarded condition. Points: mean; lines: SEM. $^*p < 0.05$, *t* test.

network. These results suggest a specificity in the inter-region interactions inferred by CURBD: frontal cortex provides input to H/A during memory retrieval, while all pathways are recruited following a novel image to encode new information into memory.

DISCUSSION

Advantages of CURBD

Typical population activity analysis approaches study the experimentally measured outputs from a neural circuit (e.g., action potentials or calcium fluorescence). Dimensionality reduction[8] estimates low-dimensional neural manifolds[64,66,70] capturing patterns of coactivation between neurons. However, the covariance driving this phenomenon is shaped by the inputs driving that population,[71] including recurrent interactions within and across brain regions.[45,46,58,72–75] CURBD offers a unique view by decomposing population activity into putative inputs or source currents. Rather than identifying a single manifold capturing the measured outputs of active units within a given region, we can use such source currents to compute separate manifolds for each source current inferred. CURBD allows us to reconceptualize population activity as numerous manifolds embedded in the space of neural activity, each capturing the dynamics of isolated sources of input.

CURBD builds on and extends existing approaches using data-constrained models to understand the functional architecture of neural circuits or the statistics of multi-region interactions.[3,9,11,28,33,35,76] Reverse engineering our data-constrained multi-region RNNs reveals quantities that are inaccessible from measurements alone: (1) the directionality, type, and magnitudes of the interactions responsible for the observed neural dynamics[11] and (2) the time-varying behavior of different inter-area currents. CURBD addresses gaps in existing approaches to analyze modern, large-scale experimental datasets.[45,58,72–75] Common methods to study interactions between brain regions such as linear regression,[36,39] CCA,[43] constrained dimensionality reduction,[37,38] linear discriminant analysis,[77] generalized linear models (GLMs),[36] or Granger causality[78] rely on correlative analysis of neural data with several challenges. First, correlation-based inference of effective connectivity cannot distinguish common inputs from directed inter-region interactions (but can be partially ameliorated with additional covariates[36,79]). Second, correlation does not indicate directionality, though careful assessment of spike latencies unveils possible directional effects.[80] Third, correlative analyses are difficult to extend to

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

CellPress
OPEN ACCESS

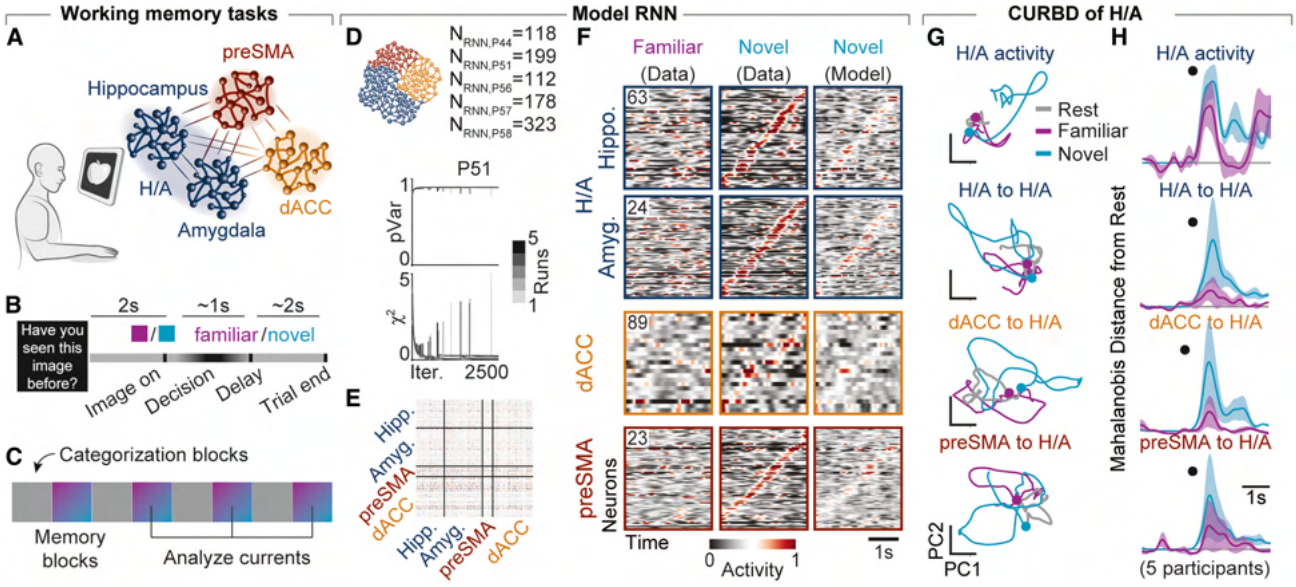

Figure 8. CURBD of four-region recordings of humans during long-term memory

(A) Human participants were implanted with depth electrodes to record single-unit activity from the hippocampus and amygdala (combined as H/A), preSMA, and dACC.

(B) Trial structure per memory block. Familiar or novel images were presented as participants reported familiarity.

(C) Each session comprised eight blocks. Odd blocks: image categorization; even blocks (as in B): novel/familiar judgment. We used blocks 4, 6, and 8, where performance was highest.

(D) Training performance (pVar and $\chi 2$) for model RNNs in P51; gray shades: five random initializations.

(E) **J** of a model RNN trained on P51.

(F) Pseudopopulation activity during the memory task (Block 4) from P51 following familiar (left, inset: neuron count), novel (middle), and corresponding model RNN activity on novel trials (right) for the four regions. Neurons sorted by peak activity in data on novel trials.

(G) Population trajectories in the leading two PCs for H/A neurons during pre-stimulus baseline (Rest, gray) and in response to familiar (magenta) and novel (cyan) stimuli, and source current trajectories within H/A. Dot: state at stimulus onset.

(H) Mahalanobis distance from pre-stimulus rest over time for each source current into H/A. Dot: stimulus onset. Lines show mean ± SEM.

data from multiple interacting regions, though recent work on switching dynamical systems shows promise.[81,82] CURBD addresses these limitations by building and analyzing RNNs that are trained to match experimental time series[3,9] in a largely unbiased manner.

CURBD explicitly models, in an unbiased way, recurrence between all recorded neurons, capturing all possible multi-region interactions in the dataset. Importantly, the inferred directed interaction matrices are not necessarily symmetric, allowing directional estimates of the inferred functional interactions (e.g., Figure 7E). In contrast to dynamic causal modeling,[83,84] CURBD does not necessarily require known perturbations or inputs and can, in principle, model any dataset. Indeed, CURBD scales natively to arbitrarily large datasets with any number of regions, even to whole-brain recordings now available from *C. elegans*,[85] fruit fly,[86] and larval zebrafish.[11,72,73]

Limitations of CURBD

CURBD currently ignores known structure in the data, giving relatively unconstrained solutions beyond the optimization target to reproduce the neural dynamics. Since the dynamics in neuroscience datasets are typically relatively low-dimensional relative to the full population size,[8,49,64] this gives a degenerate solution space where small variations in initialization can give large changes in individual weights. As we demonstrate, CURBD can robustly estimate low-dimensional dynamics and is well-suited to explore scientific questions formulated at this level. However, the models are unable to return specific neuron-to-neuron connectivity robustly, though admittedly such a task may not be possible.[44] Indeed, the directed interactions inferred by CURBD need not capture causal relationships or relate to actual synaptic connectivity. While a direct connection between two neurons should increase the directed interaction weight, strong interactions could arise indirectly.[87] For example, polysynaptic pathways or neuromodulator release could enable one neuron to exert an influence on other neurons that would manifest as an inferred directed interaction weight.[11,88] This also allows individual neurons to functionally have both excitatory and inhibitory effects on different downstream cells, since an excitatory cell could have an inhibitory effect on a neighbor by recruiting an intermediate inhibitory cell.

CURBD currently struggles with sparse activity given the objective to explain population variance. Additionally, CURBD currently fits a single dynamical system for a given sample of data and is unable to natively adapt to drifts, dropout, or other potentially aberrant features of the neural data. One consequence of this direct fitting of a single system to data is the difficulty in crafting train-test-validation sets common in machine

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

learning. While implementing training/validation sets will help avoid overfitting, this cannot necessarily be achieved in all scenarios where CURBD is useful. Indeed, we designed CURBD to help infer structure in neural data in conditions of spontaneous behavior or other continuous processes where an explicit trial structure—which allows repetition for model building and statistics—is not achievable. In some scenarios (e.g., repeated reaching to target locations while recording from the motor cortex[42]), the training/validation/test splits could be employed while training the RNNs. However, it may be challenging for existing CURBD models trained in this manner to capture the intricacies of single-trial dynamics. When proper training and validation sets cannot be readily defined, we do think that carefully designed regularizers or complex training approaches could bring us closer to that goal, for example, conditioning connectivity on connectomes as priors to modify the gradient. As described below, future extensions of CURBD could at least partially resolve these limitations.

Interpreting the model RNNs

Our model RNNs are specific to the neural dynamics of the data they were trained to reproduce. This represents a difference between CURBD and other approaches that seek generative models of many samples of neural dynamics.[33] However, even though our models were fit to single data samples,[89] we identify consistent solutions across iterations, e.g., at the level of statistical distributions of groups of interaction weights (Figure S6) and at the level of currents inferred by CURBD (Figure 7F). CURBD uses the product of the directed interaction weight matrix and the source activity to compute currents, in practice averaging out estimation noise and lending robustness to the inferred currents. Ultimately, the model RNNs are best viewed as an *in silico* representation of the experimental data itself. This enables a deeper dive into the data than we can typically access experimentally.[20] Our current approach assumes that a single directed interaction matrix captures the dynamics for the whole duration of the data. Factors such as learning[13,90] or behavioral state changes—such as those induced by stress,[11] depression,[91] or even glia[92]—could change the dynamical rules governing the interactions among different neural populations *in vivo*. State changes can be easily accounted for by fitting different model RNNs on different samples of data (e.g., periods of time and task conditions), with total currents reconstructed by stitching the resulting currents together in time. The training process could also be modified to more elegantly identify state changes and adjust the directed interaction matrix over time in a partially unsupervised, adaptive manner consistent with biological plasticity mechanisms.

Data and computational requirements

CURBD flexibly models datasets of arbitrary size by leveraging the statistics across neurons to infer dynamics. This is a major advantage for continuous datasets that cannot be trial-averaged, such as during complex and naturalistic behaviors. A downside, however, is that CURBD requires sufficiently large populations of neurons, either simultaneously recorded or with pseudopopulations. The exact number will be dependent on many features of the data experiments (e.g., complexity and dimensionality of dynamics, number of conditions, and sparsity of neural activity). We find that CURBD is most reliable and effective with larger numbers of neurons. This may be counter-intuitive since the number of parameters scales quadratically with the number of neurons, but each neuron also provides a useful additional constraint on the space of solutions that captures the lower-dimensional dynamics robustly. The number of required data points (time samples) poses several tradeoffs: too few samples prevent the model from converging robustly, but longer time periods require the assumption that the underlying dynamics are unchanged throughout the range, which may be difficult to guarantee in realistic experimental scenarios. Further, modeling longer time-series requires significantly increased compute time given the serialized nature of the recursive least squares computations in the current formulation of CURBD. The ideal length of time to model will thus be strongly dependent on the peculiarities of the experiment and is an important hyperparameter.

Extensions of CURBD

While the model RNNs used here made no assumptions about the structure of the directed interaction matrix or inter-region connectivity, brains have numerous anatomical constraints that could be incorporated into the model in the future. For example, the effect of a given neuron on its numerous downstream targets can be either excitatory or inhibitory.[93] This could be incorporated into the learning rule by restricting columns of the directed interaction matrix to either positive or negative weights. Additionally, while our RNNs are weighted all-to-all, inter-regional connections in biological brains are highly structured. For example, long-range connections between regions are likely more sparse than local connections.[94] Sparsity could be induced in an unsupervised manner by applying an L1 norm on the weights of specific inter-region submatrices in the cost function or by extending our previously developed "PINning" approach,[28] restricting the fraction of recurrent weights to be trained in inter-region submatrices. Brain-wide connectomics data[87,95,96] could also be leveraged to constrain structure a priori in the directed interaction matrix. Future extensions of this work could focus on developing faster, online training and inference pipelines to see how well the directed interaction matrices inferred from data-constrained multi-region RNNs are either constrained by or predictive of the respective structural connectomes. Lastly, future extensions of CURBD could allow more temporally precise directed interaction estimates by incorporating models of spiking statistics into the model.[79,97]

In this work, the region identity of each experimental neuron was known using anatomical landmarks or electrode implantation sites,[11,13,15,16] allowing us to divide the directed interaction matrix into region-specific blocks. In future work CURBD could provide a basis for functional clustering that goes beyond anatomical designations by applying clustering or tensor decomposition methods[98] directly to the inferred currents in an unsupervised manner. This could, for example, help identify interdigitated submodules within single regions[99–101] or new brain-wide functional circuits.[102]

The possible use cases of CURBD need not be confined to cellular resolution data. For example, model RNNs can be fit to signals such as local field potentials, which can serve as

macroscopic estimates of similar population dynamics.[103] Relevant, non-neuronal data, such as behavioral data derived from pose detection methods,[104,105] can also be incorporated as additional constraints on the model RNNs during training to help account for unobserved common inputs.[106] Static labels representing experimental metadata (behavioral task, stimulus condition, etc.) could help compensate for brain-wide state changes. Importantly, all of the extensions described above do not change the fundamental principles underlying CURBD, which at its core relies on straightforward and interpretable matrix multiplication. They only serve to provide a more constrained estimate of the biological system's directed interactions.[107]

Here, we use data from four species to show how CURBD can flexibly reverse-engineer neural population interactions at scale. Many structural and functional aspects of brain organization are conserved across species even with divergent phylogenetic trees.[108,109] Understanding the commonalities (or differences) in interactions across different species[109] will be critical to uncover fundamental principles of neural computation.[1,12] This requires an analytical framework such as CURBD that robustly and flexibly scales across a range of different experimental factors—e.g., methodologies, levels of granularity, sampling densities, and spatiotemporal resolutions—such as those encountered when comparing different species ranging from smaller, highly sampled nervous systems (e.g., *Caenorhabditis elegans*, *Drosophila*, and larval zebrafish) to larger, less sampled brains (e.g., rodents, monkeys, and humans) (Figure 2A). Thus, CURBD provides a powerful new approach for comparative studies over time, across individuals, across scales of neural function, or even across species.

RESOURCE AVAILABILITY

Lead contact

Any additional information required to reanalyze the data reported in this work is available from the lead contact upon request (kanaka_rajan@hms.harvard.edu).

Materials availability

This work did not generate new, unique materials.

Data and code availability

The human datasets are publicly available at https://osf.io/u3kcp/. The remaining fish, mouse, and monkey datasets can be shared upon reasonable request. All modeling and analysis in this manuscript were done in MATLAB (The MathWorks, Inc.) or Python. Code to train multi-region model RNNs based on experimental recordings and perform CURBD is available at https://github.com/rajanlab/CURBD.

ACKNOWLEDGMENTS

This work was supported by the Fonds de Recherche du Québec – Santé (chercheurs-boursiers en intelligence artificielle) (M.G.P.); a Louis Perry Jones, an Alice and Joseph Brooks, and a Mahoney Postdoctoral Fellowship (C.A.); a European Molecular Biology Organization (EMBO) Postdoctoral Fellowship (S.S.); the Gatsby Charitable Foundation GAT3708 and NIH R01EB029858 (M.B.); the National Institutes of Health (NIH) (R01 MH110831 and U01 NS117839) and National Science Foundation (NSF) (BCS-1554105) (U.R.); a National Institute of Mental Health BRAINS award (R01MH118638 and R01MH132064) and a young investigator grant from the Brain and Behavior Foundation (NARSAD) (P.R.); an NIH Director's Pioneer Award (DP1 MH125776) and NINDS R37 NS089521 (C.D.H.); an NIH TR01 grant (K.D., T.B., and A.S.A.); and awards from the NIH (R01 EB028166 and RF1DA056403), the James S. McDonnell Foundation (220020466), the McKnight Endowment Fund, and the NSF (1926800 and 2427124) (K.R.). We thank Drs. Juan A. Gallego, Raeed H. Chowdhury, L.F. Abbott, Sheila Cherry, and Christian D. Marton for comments on earlier drafts of this manuscript. We are indebted to and inspired by Dr. C.R. Rajan and Prabha Rajan for sharing their undying love of learning.

AUTHOR CONTRIBUTIONS

M.G.P. and K.R. conceived of CURBD. M.G.P. analyzed datasets and generated figures. M.G.P., M.B., K.D., and K.R. wrote the manuscript. M.G.P., M.B., C.A., S.S., S.Z., M.B., C.P.M., J.M., U.R., P.H.R., K.D., C.D.H., K.D., and K.R. edited the manuscript. E.C. ran cloud simulations and wrote code. A.S.A., T.B., and K.D. provided the larval zebrafish dataset. C.A., S.S., and C.D.H. provided the mouse dataset. M.E.Y., C.P.M., and P.H.R. provided the macaque dataset. J.M. and U.R. provided the human dataset. K.D. and K.R. supervised this work.

DECLARATION OF INTERESTS

The authors declare no competing interests.

STAR★METHODS

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

- KEY RESOURCES TABLE
- METHOD DETAILS
  - Multi-region recurrent neural networks
  - Computing current sources to specific brain regions
  - Two-region model producing idealized, ground truth, simulated data to validate CURBD
  - Three-region ground truth simulation to validate CURBD
  - Multi-region calcium fluorescence recordings in larval zebrafish
  - Multi-region calcium fluorescence recordings in mice
  - Multi-region electrophysiology recordings in monkeys
  - Multi-region electrophysiology recordings in humans

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.neuron.2026.07.016.

Received: March 25, 2024
Revised: February 8, 2026
Accepted: July 10, 2026

REFERENCES

1. Krubitzer, L.A., and Seelke, A.M.H. (2012). Cortical evolution in mammals: the bane and beauty of phenotypic variability. Proc. Natl. Acad. Sci. USA *109*, 10647–10654. https://doi.org/10.1073/pnas.1201891109.
2. Béna, G., and Goodman, D.F.M. (2025). Dynamics of specialization in neural modules under resource constraints. Nat Commun *16*, 187. https://doi.org/10.1038/s41467-024-55188-9.
3. Perich, M.G., and Rajan, K. (2020). Rethinking brain-wide interactions through multi-region "network of networks" models. Curr. Opin. Neurobiol. *65*, 146–151. https://doi.org/10.1016/j.conb.2020.11.003.
4. Tian, J., Huang, R., Cohen, J.Y., Osakada, F., Kobak, D., Machens, C.K., Callaway, E.M., Uchida, N., and Watabe-Uchida, M. (2016). Distributed and Mixed Information in Monosynaptic Inputs to Dopamine Neurons. Neuron *91*, 1374–1389. https://doi.org/10.1016/j.neuron.2016.08.018.
5. Schröder, S., Steinmetz, N.A., Krumin, M., Pachitariu, M., Rizzi, M., Lagnado, L., Harris, K.D., and Carandini, M. (2020). Arousal Modulates

Retinal Output. Neuron *107*, 487–495.e9. https://doi.org/10.1016/j.neuron.2020.04.026.

6. Musall, S., Kaufman, M.T., Juavinett, A.L., Gluf, S., and Churchland, A.K. (2019). Single-trial neural dynamics are dominated by richly varied movements. Nat. Neurosci. *22*, 1677–1686. https://doi.org/10.1038/s41593-019-0502-4.
7. Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C.B., Carandini, M., and Harris, K.D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. Science *364*, 255. https://doi.org/10.1126/science.aav7893.
8. Cunningham, J.P., and Yu, B.M. (2014). Dimensionality reduction for large-scale neural recordings. Nat. Neurosci. *17*, 1500–1509. https://doi.org/10.1038/nn.3776.
9. Fisher, D., Olasagasti, I., Tank, D.W., Aksay, E.R.F., and Goldman, M.S. (2013). A modeling framework for deriving the structural and functional architecture of a short-term memory microcircuit. Neuron *79*, 987–1000. https://doi.org/10.1016/j.neuron.2013.06.041.
10. Vogels, T.P., Rajan, K., and Abbott, L.F. (2005). Neural network dynamics. Annu. Rev. Neurosci. *28*, 357–376. https://doi.org/10.1146/annurev.neuro.28.061604.135637.
11. Andalman, A.S., Burns, V.M., Lovett-Barron, M., Broxton, M., Poole, B., Yang, S.J., Grosenick, L., Lerner, T.N., Chen, R., Benster, T., et al. (2019). Neuronal Dynamics Regulating Brain and Behavioral State Transitions. Cell *177*, 970–985.e20. https://doi.org/10.1016/j.cell.2019.02.037.
12. Lovett-Barron, M., Andalman, A.S., Allen, W.E., Vesuna, S., Kauvar, I., Burns, V.M., and Deisseroth, K. (2017). Ancestral Circuits for the Coordinated Modulation of Brain State. Cell *171*, 1411–1423.e17. https://doi.org/10.1016/j.cell.2017.10.021.
13. Driscoll, L.N., Pettit, N.L., Minderer, M., Chettih, S.N., and Harvey, C.D. (2017). Dynamic Reorganization of Neuronal Activity Patterns in Parietal Cortex. Cell *170*, 986–999.e16. https://doi.org/10.1016/j.cell.2017.07.021.
14. Arlt, C., Barroso-Luque, R., Kira, S., Bruno, C.A., Xia, N., Chettih, S.N., Soares, S., Pettit, N.L., and Harvey, C.D. (2022). Cognitive experience alters cortical involvement in goal-directed navigation. eLife *11*, e76051. https://doi.org/10.7554/eLife.76051.
15. Young, M.E., Spencer-Salmon, C., Mosher, C., Tamang, S., Rajan, K., and Rudebeck, P.H. (2023). Temporally specific patterns of neural activity in interconnected corticolimbic structures during reward anticipation. Neuron *111*, 3668–3682.e5. https://doi.org/10.1016/j.neuron.2023.07.012.
16. Minxha, J., Adolphs, R., Fusi, S., Mamelak, A.N., and Rutishauser, U. (2020). Flexible recruitment of memory-based choice representations by the human medial frontal cortex. Science *368*, eaba3313. https://doi.org/10.1126/science.aba3313.
17. Barak, O. (2017). Recurrent neural networks as versatile tools of neuroscience research. Curr. Opin. Neurobiol. *46*, 1–6. https://doi.org/10.1016/j.conb.2017.06.003.
18. Yang, G.R., and Wang, X.-J. (2020). Artificial Neural Networks for Neuroscientists: A Primer. Neuron *107*, 1048–1070. https://doi.org/10.1016/j.neuron.2020.09.005.
19. Sompolinsky, H., Crisanti, A., and Sommers, H.J. (1988). Chaos in random neural networks. Phys. Rev. Lett. *61*, 259–262. https://doi.org/10.1103/PhysRevLett.61.259.
20. Sussillo, D., and Barak, O. (2013). Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks. Neural Comput. *25*, 626–649. https://doi.org/10.1162/NECO_a_00409.
21. Sussillo, D., and Abbott, L.F. (2009). Generating coherent patterns of activity from chaotic neural networks. Neuron *63*, 544–557. https://doi.org/10.1016/j.neuron.2009.07.018.
22. DePasquale, B., Cueva, C.J., Rajan, K., Escola, G.S., and Abbott, L.F. (2018). Full-force: a target-based method for training recurrent networks. PLoS One *13*, e0191527. https://doi.org/10.1371/journal.pone.0191527.
23. Sussillo, D., Churchland, M.M., Kaufman, M.T., and Shenoy, K.V. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. Nat. Neurosci. *18*, 1025–1033. https://doi.org/10.1038/nn.4042.
24. Remington, E.D., Narain, D., Hosseini, E.A., and Jazayeri, M. (2018). Flexible Sensorimotor Computations through Rapid Reconfiguration of Cortical Dynamics. Neuron *98*, 1005–1019.e5. https://doi.org/10.1016/j.neuron.2018.05.020.
25. Wang, J., Narain, D., Hosseini, E.A., and Jazayeri, M. (2018). Flexible timing by temporal scaling of cortical responses. Nat. Neurosci. *21*, 102–110. https://doi.org/10.1038/s41593-017-0028-6.
26. Sohn, H., Narain, D., Meirhaeghe, N., and Jazayeri, M. (2019). Bayesian Computation through Cortical Latent Dynamics. Neuron *103*, 934–947.e5. https://doi.org/10.1016/j.neuron.2019.06.012.
27. Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature *503*, 78–84. https://doi.org/10.1038/nature12742.
28. Rajan, K., Harvey, C.D., and Tank, D.W. (2016). Recurrent Network Models of Sequence Generation and Memory. Neuron *90*, 128–142. https://doi.org/10.1016/j.neuron.2016.02.009.
29. Michaels, J.A., Schaffelhofer, S., Agudelo-Toro, A., and Scherberger, H. (2020). A goal-driven modular neural network predicts parietofrontal neural dynamics during grasping. Proc. Natl. Acad. Sci. USA *117*, 32124–32135. https://doi.org/10.1073/pnas.2005087117.
30. Codol, O., Krishna, N.H., Lajoie, G., and Perich, M.G. (2024). Brain-like neural dynamics for behavioral control develop through reinforcement learning. Preprint at bioRxiv. https://doi.org/10.1101/2024.10.04.616712.
31. Jamkhandi, A., Korojy, A., Codol, O., Lajoie, G., and Perich, M.G. (2026). JEDI: Jointly embedded inference of neural dynamics. arXiv. https://doi.org/10.48550/arXiv.2603.10489.
32. Cohen, Z., DePasquale, B., Aoi, M.C., and Pillow, J.W. (2020). Recurrent dynamics of prefrontal cortex during context-dependent decision-making. Preprint at bioRxiv. https://doi.org/10.1101/2020.11.27.401539.
33. Pandarinath, C., O'Shea, D.J., Collins, J., Jozefowicz, R., Stavisky, S.D., Kao, J.C., Trautmann, E.M., Kaufman, M.T., Ryu, S.I., Hochberg, L.R., et al. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. Nat. Methods *15*, 805–815. https://doi.org/10.1038/s41592-018-0109-9.
34. Maheswaranathan, N., Williams, A.H., Golub, M.D., Ganguli, S., and Sussillo, D. (2019). Reverse engineering recurrent networks for sentiment classification reveals line attractor dynamics. Adv. Neural Inf. Process. Syst. *32*, 15696–15705.
35. Miri, A., Daie, K., Arrenberg, A.B., Baier, H., Aksay, E., and Tank, D.W. (2011). Spatial gradients and multidimensional dynamics in a neural integrator circuit. Nat. Neurosci. *14*, 1150–1159. https://doi.org/10.1038/nn.2888.
36. Perich, M.G., Gallego, J.A., and Miller, L.E. (2018). A Neural Population Mechanism for Rapid Learning. Neuron *100*, 964–976.e7. https://doi.org/10.1016/j.neuron.2018.09.030.
37. Semedo, J.D., Zandvakili, A., Machens, C.K., Yu, B.M., and Kohn, A. (2019). Cortical Areas Interact through a Communication Subspace. Neuron *102*, 249–259.e4. https://doi.org/10.1016/j.neuron.2019.01.026.
38. Perich, M.G., Conti, S., Badi, M., Bogaard, A., Barra, B., Rajan, K., Bloch, J., Courtine, G., Micera, S., Capogrosso, M., et al. (2020). Motor cortical dynamics are shaped by multiple distinct subspaces during naturalistic behavior. Preprint at bioRxiv. https://doi.org/10.1101/2020.07.30.228767.
39. Kaufman, M.T., Churchland, M.M., Ryu, S.I., and Shenoy, K.V. (2014). Cortical activity in the null space: permitting preparation without movement. Nat. Neurosci. *17*, 440–448. https://doi.org/10.1038/nn.3643.

40. Abbott, L.F., and Dayan, P. (2005). Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems (MIT Press).

41. Rajan, K., Abbott, L.F., and Sompolinsky, H. (2010). Stimulus-dependent suppression of chaos in recurrent neural networks. Phys. Rev. E Stat. Nonlin. Soft Matter Phys. *82*, 011903. https://doi.org/10.1103/PhysRevE.82.011903.

42. Gallego, J.A., Perich, M.G., Chowdhury, R.H., Solla, S.A., and Miller, L.E. (2020). Long-term stability of cortical population dynamics underlying consistent behavior. Nat. Neurosci. *23*, 260–270. https://doi.org/10.1038/s41593-019-0555-4.

43. Veuthey, T.L., Derosier, K., Kondapavulur, S., and Ganguly, K. (2020). Single-trial cross-area neural population dynamics during long-term skill learning. Nat. Commun. *11*, 4057. https://doi.org/10.1038/s41467-020-17902-1.

44. Das, A., and Fiete, I.R. (2020). Systematic errors in connectivity inferred from activity in strongly recurrent networks. Nat. Neurosci. *23*, 1286–1296. https://doi.org/10.1038/s41593-020-0699-2.

45. Jun, J.J., Steinmetz, N.A., Siegle, J.H., Denman, D.J., Bauza, M., Barbarits, B., Lee, A.K., Anastassiou, C.A., Andrei, A., Aydın, Ç., et al. (2017). Fully integrated silicon probes for high-density recording of neural activity. Nature *551*, 232–236. https://doi.org/10.1038/nature24636.

46. Kim, T.H., Zhang, Y., Lecoq, J., Jung, J.C., Li, J., Zeng, H., Niell, C.M., and Schnitzer, M.J. (2016). Long-Term Optical Access to an Estimated One Million Neurons in the Live Mouse Cortex. Cell Rep. *17*, 3385–3394. https://doi.org/10.1016/j.celrep.2016.12.004.

47. Dombeck, D.A., Khabbaz, A.N., Collman, F., Adelman, T.L., and Tank, D.W. (2007). Imaging large-scale neural activity with cellular resolution in awake, mobile mice. Neuron *56*, 43–57. https://doi.org/10.1016/j.neuron.2007.08.003.

48. Engelhard, B., Finkelstein, J., Cox, J., Fleming, W., Jang, H.J., Ornelas, S., Koay, S.A., Thiberge, S.Y., Daw, N.D., Tank, D.W., et al. (2019). Specialized coding of sensory, motor and cognitive variables in VTA dopamine neurons. Nature *570*, 509–513. https://doi.org/10.1038/s41586-019-1261-9.

49. Gao, P., Trautmann, E., Yu, B., Santhanam, G., Ryu, S., Shenoy, K., and Ganguli, S. (2017). A theory of multineuronal dimensionality, dynamics and measurement. Preprint at bioRxiv. https://doi.org/10.1101/214262.

50. Gallego, J.A., Perich, M.G., Naufel, S.N., Ethier, C., Solla, S.A., and Miller, L.E. (2018). Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. Nat. Commun. *9*, 4233. https://doi.org/10.1038/s41467-018-06560-z.

51. Abbott, L.F., Rajan, K., and Sompolinsky, H. (2011). Interactions between Intrinsic and Stimulus-Evoked Activity in Recurrent Neural Networks. In The Dynamic Brain (Oxford University Press), pp. 65–82. https://doi.org/10.1093/acprof:oso/9780195393798.003.0004.

52. Maheswaranathan, N., Williams, A.H., Golub, M.D., Ganguli, S., and Sussillo, D. (2019). Universality and individuality in neural dynamics across large populations of recurrent networks. Adv. Neural Inf. Process. Syst. *2019*, 15629–15641.

53. Vladimirov, N., Mu, Y., Kawashima, T., Bennett, D.V., Yang, C.-T., Looger, L.L., Keller, P.J., Freeman, J., and Ahrens, M.B. (2014). Light-sheet functional imaging in fictively behaving zebrafish. Nat. Methods *11*, 883–884. https://doi.org/10.1038/nmeth.3040.

54. Aoki, T., Kinoshita, M., Aoki, R., Agetsuma, M., Aizawa, H., Yamazaki, M., Takahoko, M., Amo, R., Arata, A., Higashijima, S.-I., et al. (2013). Imaging of neural ensemble for the retrieval of a learned behavioral program. Neuron *78*, 881–894. https://doi.org/10.1016/j.neuron.2013.04.009.

55. Beck, A.T., Steer, R.A., Kovacs, M., and Garrison, B. (1985). Hopelessness and eventual suicide: a 10-year prospective study of patients hospitalized with suicidal ideation. Am. J. Psychiatry *142*, 559–563. https://doi.org/10.1176/ajp.142.5.559.

56. Maier, S.F., and Seligman, M.E.P. (2016). Learned helplessness at fifty: Insights from neuroscience. Psychol. Rev. *123*, 349–367. https://doi.org/10.1037/rev0000033.

57. Chen, T.-W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V., et al. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. Nature *499*, 295–300. https://doi.org/10.1038/nature12354.

58. Sofroniew, N.J., Flickinger, D., King, J., and Svoboda, K. (2016). A large field of view two-photon mesoscope with subcellular resolution for in vivo imaging. eLife *5*, e14472. https://doi.org/10.7554/eLife.14472.

59. Harvey, C.D., Coen, P., and Tank, D.W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. Nature *484*, 62–68. https://doi.org/10.1038/nature10918.

60. Vann, S.D., Aggleton, J.P., and Maguire, E.A. (2009). What does the retrosplenial cortex do? Nat. Rev. Neurosci. *10*, 792–802. https://doi.org/10.1038/nrn2733.

61. Ebbesen, C.L., and Brecht, M. (2017). Motor cortex — to act or not to act? Nat. Rev. Neurosci. *18*, 694–705. https://doi.org/10.1038/nrn.2017.119.

62. Steinmetz, N.A., Zatka-Haas, P., Carandini, M., and Harris, K.D. (2019). Distributed coding of choice, action and engagement across the mouse brain. Nature *576*, 266–273. https://doi.org/10.1038/s41586-019-1787-x.

63. Gremel, C.M., and Costa, R.M. (2013). Premotor cortex is critical for goal-directed actions. Front. Comput. Neurosci. *7*, 110. https://doi.org/10.3389/fncom.2013.00110.

64. Gallego, J.A., Perich, M.G., Miller, L.E., and Solla, S.A. (2017). Neural Manifolds for the Control of Movement. Neuron *94*, 978–984. https://doi.org/10.1016/j.neuron.2017.05.025.

65. Fortunato, C., Bennasar-Vázquez, J., Park, J., Chang, J.C., Miller, L.E., Dudman, J.T., Perich, M.G., and Gallego, J.A. (2024). Nonlinear manifolds underlie neural population activity during behaviour. Preprint at bioRxiv. https://doi.org/10.1101/2023.07.18.549575.

66. Perich, M.G., Narain, D., and Gallego, J.A. (2025). A neural manifold view of the brain. Nat. Neurosci. *28*, 1582–1597. https://doi.org/10.1038/s41593-025-02031-z.

67. Lindquist, K.A., Wager, T.D., Kober, H., Bliss-Moreau, E., and Barrett, L.F. (2012). The brain basis of emotion: a meta-analytic review. Behav. Brain Sci. *35*, 121–143. https://doi.org/10.1017/S0140525X11000446.

68. Haber, S.N., Kim, K.-S., Mailly, P., and Calzavara, R. (2006). Reward-related cortical inputs define a large striatal region in primates that interface with associative cortical connections, providing a substrate for incentive-based learning. J. Neurosci. *26*, 8368–8376. https://doi.org/10.1523/JNEUROSCI.0271-06.2006.

69. Minxha, J., Mamelak, A.N., and Rutishauser, U. (2018). Surgical and Electrophysiological Techniques for Single-Neuron Recordings in Human Epilepsy Patients. Neuromethods, 267–293. https://doi.org/10.1007/978-1-4939-7549-5_14.

70. Sadtler, P.T., Quick, K.M., Golub, M.D., Chase, S.M., Ryu, S.I., Tyler-Kabara, E.C., Yu, B.M., and Batista, A.P. (2014). Neural constraints on learning. Nature *512*, 423–426. https://doi.org/10.1038/nature13665.

71. Stevenson, I.H., London, B.M., Oby, E.R., Sachs, N.A., Reimer, J., Englitz, B., David, S.V., Shamma, S.A., Blanche, T.J., Mizuseki, K., et al. (2012). Functional connectivity and tuning curves in populations of simultaneously recorded neurons. PLoS Comput. Biol. *8*, e1002775. https://doi.org/10.1371/journal.pcbi.1002775.

72. Ahrens, M.B., Orger, M.B., Robson, D.N., Li, J.M., and Keller, P.J. (2013). Whole-brain functional imaging at cellular resolution using light-sheet microscopy. Nat. Methods *10*, 413–420. https://doi.org/10.1038/nmeth.2434.

73. Ahrens, M.B., Li, J.M., Orger, M.B., Robson, D.N., Schier, A.F., Engert, F., and Portugues, R. (2012). Brain-wide neuronal dynamics during motor

adaptation in zebrafish. Nature *485*, 471–477. https://doi.org/10.1038/nature11057.

74. Scott, B.B., Thiberge, S.Y., Guo, C., Tervo, D.G.R., Brody, C.D., Karpova, A.Y., and Tank, D.W. (2018). Imaging Cortical Dynamics in GCaMP Transgenic Rats with a Head-Mounted Widefield Macroscope. Neuron *100*, 1045–1058.e5. https://doi.org/10.1016/j.neuron.2018.09.050.
75. Prevedel, R., Yoon, Y.-G., Hoffmann, M., Pak, N., Wetzstein, G., Kato, S., Schrödel, T., Raskar, R., Zimmer, M., Boyden, E.S., et al. (2014). Simultaneous whole-animal 3D imaging of neuronal activity using light-field microscopy. Nat. Methods *11*, 727–730. https://doi.org/10.1038/nmeth.2964.
76. Li, N., Daie, K., Svoboda, K., and Druckmann, S. (2016). Robust neuronal dynamics in premotor cortex during motor planning. Nature *532*, 459–464. https://doi.org/10.1038/nature17643.
77. Javadzadeh, M., and Hofer, S.B. (2022). Dynamic causal communication channels between neocortical areas. Neuron *110*, 2470–2483.e7. https://doi.org/10.1016/j.neuron.2022.05.011.
78. Seth, A.K., Barrett, A.B., and Barnett, L. (2015). Granger Causality Analysis in Neuroscience and Neuroimaging. J. Neurosci. *35*, 3293–3297. https://doi.org/10.1523/JNEUROSCI.4399-14.2015.
79. Lawlor, P.N., Perich, M.G., Miller, L.E., and Kording, K.P. (2018). Linear-nonlinear-time-warp-poisson models of neural activity. J. Comput. Neurosci. *45*, 173–191. https://doi.org/10.1007/s10827-018-0696-6.
80. Sirota, A., Csicsvari, J., Buhl, D., and Buzsáki, G. (2003). Communication between neocortex and hippocampus during sleep in rodents. Proc. Natl. Acad. Sci. USA *100*, 2065–2069. https://doi.org/10.1073/pnas.0437938100.
81. Glaser, J.I., Whiteway, M., Cunningham, J.P., Paninski, L., and Linderman, S.W. (2020). Recurrent Switching Dynamical Systems Models for Multiple Interacting Neural Populations. Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS).
82. Linderman, S., Nichols, A., Blei, D., Zimmer, M., and Paninski, L. (2019). Hierarchical recurrent state space models reveal discrete and continuous dynamics of neural activity in C. elegans. Preprint at bioRxiv. https://doi.org/10.1101/621540.
83. Friston, K.J., Harrison, L., and Penny, W. (2003). Dynamic causal modelling. Neuroimage *19*, 1273–1302. https://doi.org/10.1016/S1053-8119(03)00202-7.
84. Cao, X., Sandstede, B., and Luo, X. (2019). A Functional Data Method for Causal Dynamic Network Modeling of Task-Related fMRI. Front. Neurosci. *13*, 127. https://doi.org/10.3389/fnins.2019.00127.
85. Nguyen, J.P., Shipley, F.B., Linder, A.N., Plummer, G.S., Liu, M., Setru, S.U., Shaevitz, J.W., and Leifer, A.M. (2016). Whole-brain calcium imaging with cellular resolution in freely behaving Caenorhabditis elegans. Proc. Natl. Acad. Sci. USA *113*, E1074–E1081. https://doi.org/10.1073/pnas.1507110112.
86. Schaffer, E.S., Mishra, N., Whiteway, M.R., Li, W., Vancura, M.B., Freedman, J., Patel, K.B., Voleti, V., Paninski, L., Hillman, E.M.C., et al. (2023). The spatial and temporal structure of neural activity across the fly brain. Nat Commun *14*, 5572. https://doi.org/10.1038/s41467-023-41261-2.
87. Turner, M.H., Mann, K., and Clandinin, T.R. (2021). The connectome predicts resting-state functional connectivity across the Drosophila brain. Curr. Biol. *31*, 2386–2394.e3. https://doi.org/10.1016/j.cub.2021.03.004.
88. Falkner, A.L., Wei, D., Song, A., Watsek, L.W., Chen, I., Chen, P., Feng, J.E., and Lin, D. (2020). Hierarchical Representations of Aggression in a Hypothalamic-Midbrain Circuit. Neuron *106*, 637–648.e6. https://doi.org/10.1016/j.neuron.2020.02.014.
89. Hosseini, M., Powell, M., Collins, J., Callahan-Flintoft, C., Jones, W., Bowman, H., and Wyble, B. (2020). I tried a bunch of things: The dangers of unexpected overfitting in classification of brain data. Neurosci. Biobehav. Rev. *119*, 456–467. https://doi.org/10.1016/j.neubiorev.2020.09.036.
90. Peters, A.J., Chen, S.X., and Komiyama, T. (2014). Emergence of reproducible spatiotemporal activity during motor learning. Nature *510*, 263–267. https://doi.org/10.1038/nature13235.
91. Hultman, R., Ulrich, K., Sachs, B.D., Blount, C., Carlson, D.E., Ndubuizu, N., Bagot, R.C., Parise, E.M., Vu, M.T., Gallagher, N.M., et al. (2018). Brain-wide Electrical Spatiotemporal Dynamics Encode Depression Vulnerability. Cell *173*, 166–180.e14. https://doi.org/10.1016/j.cell.2018.02.012.
92. Mu, Y., Bennett, D.V., Rubinov, M., Narayan, S., Yang, C.-T., Tanimoto, M., Mensh, B.D., Looger, L.L., and Ahrens, M.B. (2019). Glia Accumulate Evidence that Actions Are Futile and Suppress Unsuccessful Behavior. Cell *178*, 27–43.e19. https://doi.org/10.1016/j.cell.2019.05.050.
93. Wilmes, K.A., and Clopath, C. (2019). Inhibitory microcircuits for top-down plasticity of sensory representations. Nat. Commun. *10*, 5055. https://doi.org/10.1038/s41467-019-12972-2.
94. Bassett, D.S., and Bullmore, E. (2006). Small-World Brain Networks. Neuroscientist *12*, 512–523. https://doi.org/10.1177/1073858406293182.
95. Abbott, L.F., Bock, D.D., Callaway, E.M., Denk, W., Dulac, C., Fairhall, A.L., Fiete, I., Harris, K.M., Helmstaedter, M., Jain, V., et al. (2020). The Mind of a Mouse. Cell *182*, 1372–1376. https://doi.org/10.1016/j.cell.2020.08.010.
96. Scheffer, L.K., Xu, C.S., Januszewski, M., Lu, Z., Takemura, S.-Y., Hayworth, K.J., Huang, G.B., Shinomiya, K., Maitlin-Shepard, J., Berg, S., et al. (2020). A connectome and analysis of the adult Drosophila central brain. eLife *9*, e57443. https://doi.org/10.7554/eLife.57443.
97. Pillow, J.W., Shlens, J., Paninski, L., Sher, A., Litke, A.M., Chichilnisky, E.J., and Simoncelli, E.P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. Nature *454*, 995–999. https://doi.org/10.1038/nature07140.
98. Williams, A.H., Kim, T.H., Wang, F., Vyas, S., Ryu, S.I., Shenoy, K.V., Schnitzer, M., Kolda, T.G., and Ganguli, S. (2018). Unsupervised Discovery of Demixed, Low-Dimensional Neural Dynamics across Multiple Timescales through Tensor Component Analysis. Neuron *98*, 1099–1115.e8. https://doi.org/10.1016/j.neuron.2018.05.015.
99. Inagaki, H.K., Fontolan, L., Romani, S., and Svoboda, K. (2019). Discrete attractor dynamics underlies persistent activity in the frontal cortex. Nature *566*, 212–217. https://doi.org/10.1038/s41586-019-0919-7.
100. Beyeler, A., Namburi, P., Glober, G.F., Simonnet, C., Calhoon, G.G., Conyers, G.F., Luck, R., Wildes, C.P., and Tye, K.M. (2016). Divergent Routing of Positive and Negative Information from the Amygdala during Memory Retrieval. Neuron *90*, 348–361. https://doi.org/10.1016/j.neuron.2016.03.004.
101. Jetti, S.K., Vendrell-Llopis, N., and Yaksi, E. (2014). Spontaneous activity governs olfactory representations in spatially organized habenular microcircuits. Curr. Biol. *24*, 434–439. https://doi.org/10.1016/j.cub.2014.01.015.
102. Bruno, A.M., Frost, W.N., and Humphries, M.D. (2015). Modular deconstruction reveals the dynamical and physical building blocks of a locomotion motor program. Neuron *86*, 304–318. https://doi.org/10.1016/j.neuron.2015.03.005.
103. Gallego-Carracedo, C., Perich, M.G., Chowdhury, R.H., Miller, L.E., and Gallego, J.A. (2022). Local field potentials reflect cortical population dynamics in a region-specific and frequency-dependent manner. eLife *11*, e73155. https://doi.org/10.7554/eLife.73155.
104. Mathis, A., Mamidanna, P., Cury, K.M., Abe, T., Murthy, V.N., Mathis, M.W., and Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. Nat. Neurosci. *21*, 1281–1289. https://doi.org/10.1038/s41593-018-0209-y.
105. Pereira, T.D., Aldarondo, D.E., Willmore, L., Kislin, M., Wang, S.S.-H., Murthy, M., and Shaevitz, J.W. (2019). Fast animal pose estimation using

deep neural networks. Nat. Methods *16*, 117–125. https://doi.org/10.1038/s41592-018-0234-5.

106. Khajehnejad, M., Habibollahi, F., Nock, R., Arabzadeh, E., Dayan, P., and Dezfouli, A. (2022). Proceedings of the 39th International Conference on Machine Learning, PMLR *162*, 10974–10996.
107. Cao, R., and Yamins, D. (2021). Explanatory models in neuroscience: Part 2 – constraint-based intelligibility. arXiv, 2104.01489. https://doi.org/10.48550/arXiv.2104.01490.
108. Cisek, P. (2019). Resynthesizing behavior through phylogenetic refinement. Atten. Percept. Psychophys. *81*, 2265–2287. https://doi.org/10.3758/s13414-019-01760-1.
109. Codol, O., Asclipe, M., Sobinov, A.R., Chen, Z., Park, J., Hatsopoulos, N.G., Dudman, J.T., Gallego, J.A., Lajoie, G., and Perich, M.G. (2026). Evolutionarily conserved neural dynamics across mice, monkeys, and humans. Preprint at bioRxiv. https://doi.org/10.64898/2026.03.06.709637.
110. Pachitariu, M., Stringer, C., Dipoppa, M., Schröder, S., Rossi, L.F., Dalgleish, H., Carandini, M., and Harris, K.D. (2017). Suite2p: beyond 10,000 neurons with standard two-photon microscopy. Preprint at bioRxiv. https://doi.org/10.1101/061507.
111. Chettih, S.N., and Harvey, C.D. (2019). Single-neuron perturbations reveal feature-specific competition in V1. Nature *567*, 334–340. https://doi.org/10.1038/s41586-019-0997-6.
112. Friedrich, J., Zhou, P., and Paninski, L. (2017). Fast online deconvolution of calcium imaging data. PLoS Comput Biol *13*, e1005423. https://doi.org/10.1371/journal.pcbi.1005423.
113. Haykin, S.S. (2002). Adaptive Filter Theory (Prentice Hall).
114. Safaie, M., Chang, J.C., Park, J., Miller, L.E., Dudman, J.T., Perich, M.G., and Gallego, J.A. (2023). Preserved neural dynamics across animals performing similar behaviour. Nature *623*, 765–771. https://doi.org/10.1038/s41586-023-06714-0.
115. Peng, Y., Ganesh, A., Wright, J., Xu, W., and Ma, Y. (2012). RASL: robust alignment by sparse and low-rank decomposition for linearly correlated images. IEEE Trans. Pattern Anal. Mach. Intell. *34*, 2233–2246. https://doi.org/10.1109/TPAMI.2011.282.
116. Marshel, J.H., Garrett, M.E., Nauhaus, I., and Callaway, E.M. (2011). Functional specialization of seven mouse visual cortical areas. Neuron *72*, 1040–1054. https://doi.org/10.1016/j.neuron.2011.12.004.
117. Kalatsky, V.A., and Stryker, M.P. (2003). New paradigm for optical imaging: temporally encoded maps of intrinsic signal. Neuron *38*, 529–545. https://doi.org/10.1016/s0896-6273(03)00286-1.
118. Glaser, J.I., Benjamin, A.S., Chowdhury, R.H., Perich, M.G., Miller, L.E., and Kording, K.P. (2020). Machine Learning for Neural Decoding. eNeuro *7*. ENEURO.0506-19.2020. https://doi.org/10.1523/ENEURO.0506-19.2020.
119. Minxha, J., Adolphs, R., Rutishauser, U., Mamelak, A., and Fusi, S. (2020). Data for: Flexible Recruitment of Memory-Based Choice Representations by Human Medial-Frontal Cortex (Center for Open Science). https://doi.org/10.17605/OSF.IO/U3KCP.

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental models: Organisms/strains		
C57BL/6J-Tg(Thy1-GCaMP6s)GP4.3	The Jackson Laboratory	RRID: IMSR_JAX:024275
Software and algorithms		
suite2P	Pachitariu et al.[110]	https://github.com/MouseLand/suite2p
ScanImage 2016	Vidrio Technologies	https://archive.scanimage.org/SI2016/index.html
Image preprocessing and motion correction code	Chettih et al.[111]	https://github.com/HarveyLab/Acquisition2P_class
OASIS algorithm for deconvolution of calcium imaging data	Friedrich et al.[112]	https://github.com/zhoupc/OASIS_matlab

METHOD DETAILS

Multi-region recurrent neural networks

Network elements

Network models represent real biological circuits, but they do so with different levels of fidelity.[10,18,20] We constructed Model RNNs that are directly constrained by experimentally-obtained time-series neural data. Each network unit or model neuron, here indexed by i is described by a total current $x_i(t)$, and $r_i(t)$, a nonlinear function ϕ of $x_i(t)$, where $i=1, 2, \ldots, N$ is the total number of units in the network. Each variable $x_i(t)$ obeys the following equation:

$$\tau \frac{dx_i(t)}{dt} = -x_i(t) + \sum_j^N \boldsymbol{J}_{ij} \boldsymbol{r}_j(t) + h_i(t) \quad \text{(Equation 4)}$$

where h_i is the external input to the unit, $\mathbf{J}$ is a heterogeneous matrix of recurrent connections, and τ is the unit's time constant selected based on the data to be modeled (for the datasets in this manuscript, see Table S1 below). The sum of $r_j(t)$ weighted by J_{ij} represents the current received by unit i from all network units, which will be filtered by the leaky integration and give rise to $x_i(t)$, the current going into unit i. We use ϕ=tanh(x), but other saturating nonlinearities, such as sigmoids, have been explored in related work (see Refs. [28,41]).

In Equation 4, we use parentheses $x(t)$ to represent that the variable is continuous in time. In practice, the network equations are integrated using the Euler method and an integration time step, Δt_{RNN}. Note that we allow network integration to interpolate with a finer time step than the sampling rate of the experimental data to be modeled. This allows for smoother dynamics when the experimental data may be relatively sparsely sampled. In such discrete implementations, we use square brackets $x[t]$ to represent that the time is defined in a discrete domain:

$$x_i[t] = x_i[t-1] + \frac{\Delta t}{\tau}\left(-x_i[t-1] + \sum_j^N \boldsymbol{J}_{ij} \boldsymbol{r}_j[t-1] + h_i[t-1]\right) \quad \text{(Equation 5)}$$

Recurrent weights carrying inputs onto a target unit $i=1, 2, \ldots, N$ from its source partner j, J_{ij}, which are the elements of the matrix $\mathbf{J}$, are either fixed or modifiable (plastic), depending on how much structure is introduced into the connectivity of the initially disordered network. We introduce no *a priori* structure in $\mathbf{J}$, allowing all elements to be modified during training. J_{ij} can potentially be modified by a number of different learning algorithms (see below). The $\mathbf{J}$ matrix is arranged such that each block along the diagonal represents the recurrent connections within each brain region being considered, and the off-diagonal blocks, the inter-region projections to and from them. In this way a two-region Model RNN has two blocks on the main diagonal relating each region to itself, and two regions on the off-diagonal relating each region to the other (e.g., Figure S1G). Following this pattern, multi-region Model RNNs can be initialized with $\mathbf{J}$ matrices containing more than 2 blocks.

Typically, the initial, untrained directed interaction matrix $\mathbf{J}_0$ is constructed to be the same size as the number of neurons in the dataset to be modeled, though larger networks in which different subsets of weights are modified in a data-dependent manner have been explored previously in Ref. [28]. Here, each Model RNN unit is matched to one recorded neuron from the respective experimental dataset. The individual weights in $\mathbf{J}_0$ are initially chosen independently and randomly from a Gaussian distribution with mean and variance given by

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

$<\mathbf{J}_0>=0$ and $<\mathbf{J}_0{}^2>_J=g^2/N$. The control parameter g determines the strength of the recurrent connections, and thus whether ($g>1$) or not ($g<1$) the network produces spontaneous activity with non-trivial dynamics.[19,41,51] We set $g=1.5$ here, though in practice we observe qualitatively similar results for a range of values provided g is sufficiently large to facilitate chaotic spontaneous dynamics in the network in the absence of inputs. Ultimately, the elements of **J** will be modified by the training algorithm until the activity of the model RNN's units match the target neural data, and after training, autonomously produce dynamics consistent with the experimental recordings.

Design of external inputs

In biological neural circuits, populations of cells are constantly driven by external and inter-regional inputs that we cannot always observe. To mimic this effect without modeling entire sensory-motor transforms and pathways, and to keep the external drive as general as possible, the external inputs to the each unit i in the Model RNN, denoted by $h_i(t)$, are generated from filtered and spatially delocalized white noise that is frozen, using the equation:

$$\tau_{WN}\frac{dh_i}{dt} = -h_i(t) + h^0\eta(t) \quad \text{(Equation 6)}$$

where η is a random variable drawn from a Gaussian distribution with 0 mean and unit variance, and the parameters h^0 and τ_{WN} control the scale of these inputs and their correlation time, respectively. We use $h^0=1$ and $\tau_{WN}=0.1$ in this paper. There are typically as many different inputs as there are model neurons in the network, with individual model neurons receiving the same input on every simulated trial in cases where there is an explicit trial structure (such as the monkey and human datasets used in this paper).

Model RNN training

The directed interaction weight matrix **J** is trained by a target-based method in order to match the network activity with the experimental data from cellular-resolution neural recordings or imaging. As the target is discrete in time, we take a discrete implementation of the RNN dynamics as in Equation 5, where we use square brackets to note the discrete domain. During training, the activity of individual units in the Model RNN is compared directly to target functions derived from the experimentally-recorded neurons, denoted by $a_i[t]$, to yield error values $e_i[t]$ that will be used to update **J** The error can be computed based on the influx current $x_i[t]$ (''current-based'' training) or the firing rate $r_i[t]=\phi(x_i[t])$ (''rate-based'' training) of individual units.

For ''current-based'' training, the target functions $a_i[t]$ are for the currents into each RNN unit. If the original experimental data is firing rates, it would need to be translated into current-like quantities by inverting the chosen activation function. At each time step, the influx current of individual network units, $x_i[t]$, is treated as though it were being evaluated at steady state based on Equation 5, and then compared with the target functions $a_i[t]$ to yield the error value $e_i[t]$:

$$e_i[t] = x_i[t] - a_i[t] \quad \text{(Equation 7)}$$

The error function is a linear function of the parameter being optimized (**J**), so we can adopt a learning rule similar to Recursive Least Squares (RLS) approaches.[21,28,113] During training, the elements in the directed interaction matrix **J**, J_{ij}, undergo modification at a rate proportional to three factors: i) the error term computed above; and a regularization term consisting of, ii) the ''presynaptic'' or source firing rate of each neuron; and iii) a matrix **P** with N x N elements (defined below in Equation 10).

Training proceeds iteratively as schematized in Figure 2B. At each learning time step t (fixed to a multiple of the integration timestep in the RNN dynamics) for $i=1, 2,..., N$ target units, the corresponding elements of **J** are adjusted from their values at the previous time step ($t-1$), as introduced in Equation 1, according to:

$$\mathbf{J}_{ij}[t] = \mathbf{J}_{ij}[t-1] - \Delta\mathbf{J}_{ij}[t] \quad \text{(Equation 8)}$$

where the update term[21,113] is computed according to:

$$\Delta\mathbf{J}_{ij}[t] = e_i[t]\sum_k^N P_{jk}[t]\mathbf{r}_k[t-1] \quad \text{(Equation 9)}$$

In the above, **P** is defined mathematically as the inverse cross-correlation matrix of the firing rates of units in the network, such that its elements P_{ij} are given by:

$$P_{ij} = \left(C^{-1}\right)_{ij}, \text{where } C_{ij} = \langle \mathbf{r}_i\mathbf{r}_j \rangle \quad \text{(Equation 10)}$$

The matrix **P** tracks the correlations in the firing rate fluctuations across the whole network at every time step, and is computed for all $i=1, 2,\ldots, N$ target units and all $j=1, 2, \ldots, N$ source units. It is not generally necessary to calculate the matrix **P** explicitly. Instead, **P** can be updated iteratively[113] according to (see Figure S10 for details).

$$P[t] = P[t-1] - \frac{P[t-1]\mathbf{r}[t-1]\mathbf{r}'[t-1]P[t-1]}{1+\mathbf{r}'[t-1]P[t-1]\mathbf{r}[t-1]}, \quad \text{(Equation 11)}$$

where $\mathbf{r}[t]$ denotes a vector whose ith element is $r_i[t]$. The matrix **P** is initialized to the identity matrix scaled by a factor P^0 which controls the overall learning rate. In practice, training is most effective when P^0 is set to be 1 to 10 times the overall amplitude of the external inputs (h^0).

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

As an alternative to the "current-based" training, the network can be trained using a "rate-based" learning rule, The target function $a_i[t]$ is now for the firing rate of units. The RNN unit firing rates, $r_i[t]$, are directly compared to the targets to yield the error values:

$$e_i[t] = \boldsymbol{r}_i[t] - a_i[t] \quad \text{(Equation 12)}$$

The rate-based training is similar to the current-based training and takes the same updating rule as in Equation 9, if we replace $e_i[t]$ with the rate-based definition. Therefore, the choice of learning rule (current-based or rate-based as described above) would impact the calculation of $\Delta\mathbf{J}$.

For the interpretation of the rate-based training, the update term $\Delta\mathbf{J}$ to the matrix in this case amounts to a "gradient-like" approach, albeit one with a multiplicative regularizer mediated by $\mathbf{P}$. This multiplicative term makes a summary of the activation of all the units, weighted (regularized) by $\mathbf{P}$, which potentially balances the global activation. This is not present in standard gradient-based training for RNNs. The key distinction between the rate-based and current-base methods is that in rate-based training, we do not need to make any assumptions about the nonlinearity of the biological neuron that we are targeting, while in current-based training such a nonlinearity would need to be explicitly inverted to produce the current-like target functions from rate data. Additional implications of this type of learning rule and a rigorous comparison between the two formalisms, beyond Figure 4 here, is for future work.

Since the learning algorithm updates $\mathbf{J}$ at each learning time step, accurate fits could be observed during training even when the algorithm has not fully converged. Thus, after training for a fixed number of iterations (typically between 1500 and 3000 iterations for model RNNs based on experimental neural datasets), we stopped the training to run the "learned" autonomous dynamical system for a few additional iterations to compute and evaluate the final goodness of fit. We assessed the quality of the fit and convergence using two metrics: 1) the training error (χ^2) between the Model RNN rates and the teacher/target functions (derived from data in the current-based training algorithm, or the time-series directly, in the case of rate-based training algorithm), computed as the mean-squared error $e_i(t)$ along all $i=1, 2, \ldots, N$ target neurons; and 2) the proportion of variance explained (pVar) as one minus the ratio of the average squared difference between the neural data and outputs of the network compared to the variance of the data:

$$pVar = 1 - \frac{\langle e_i[t]^2 \rangle}{\langle (a_i[t] - \bar{a}[t])^2 \rangle} \quad \text{(Equation 13)}$$

where $\bar{a}[t]$ is the mean across all units at each time step and the triangle brackets indicate average across time points and neurons.

Analyzing the Directed Interaction matrix J after training

The directed interaction matrix inferred by the Model RNN quantifies the strength of interactions between the units in the network. These values can be either positive or negative, suggesting excitatory or inhibitory effects on the target neuron, respectively. Since the RNNs we build are extensively constrained by neural dynamics, we find that it is possible to consistently infer similarly distributed matrices, when starting from different random realizations of the initial weight distribution (e.g., Figure S6). Thus the statistical properties of the interaction strengths we derive from data-constrained RNNs can be reliably compared across brain regions, as well as between RNNs trained to match a range of experimental datasets from different species. In this paper, we summarized the statistical properties of such model-derived interaction strengths by computing histograms using the total number of elements of either the full $\mathbf{J}$ matrix or specific submatrices containing the strength and type of interactions within and between individual brain regions. Notably, when analyzing these matrices, we scaled the distributions by the square root of the number of source units to account for differences in population sizes. We also normalized each histogram by the maximum value to facilitate comparison between matrices derived from RNNs of different sizes, and visualized the distributions using a logarithmic scale. These distributions could be further summarized and quantified by metrics such as the median, standard deviation, skewness, or kurtosis.

Computing current sources to specific brain regions

As indicated by Equation 4, the time derivative of the net current going into one unit in the Model RNN, $dx_i(t)/dt$, is proportional to the current it receives from all source units, defined as the product of the corresponding row of the directed interaction matrix and the activity of all source units at the previous time step. We define the total input current the i^{th} target unit received as $I_i(t)$, which is a linear combination of the rates of all the (source) units in the networks::

$$I_i(t) = \boldsymbol{J}_{i1}\boldsymbol{r}_1(t) + \boldsymbol{J}_{i2}\boldsymbol{r}_2(t) + \ldots + \boldsymbol{J}_{iN}\boldsymbol{r}_N(t) \quad \text{(Equation 14)}$$

CURBD adopts this linear decomposition to study brain-wide currents between active neurons across multiple interacting brain regions. In this manuscript, we computed the currents in the target regions using the weights in different submatrices as described here. However, this method can be readily extended to separately infer excitatory (or inhibitory) currents by first setting all of the negative (or positive) values in the $\mathbf{J}$ matrix to be zero and then repeating the summation in Equation 15.

Due to the large number of free parameters in the Model RNN, i.e., order N^2 elements for RNNs with N units, the training algorithm does not necessarily infer the precise entries, element-by-element, in the directed interaction matrix, even when ground truth simulated data originated via low-rank or smoothed connectivity (for details, see Ref. [28]). However, we find consistent and reliable estimates, i.e., recapitulating statistical properties of groups of weights in the directed interaction matrices. Furthermore, after training, the RNNs are able to produce highly consistent dynamics even when starting from different initial conditions. In practice,

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

when taking the dot product of **J** and **r**(t) to compute the currents for CURBD, random element-by-element fluctuations in the individual reconstructed weights between pairs of units are averaged out, but the overall population dynamics are preserved. For this reason, in its current state, CURBD is best applied to infer interactions between source and target brain regions with sufficient numbers of active neurons. Future extensions, e.g., those that incorporate known connectivity between regions[87,96] or additional constraints from data, such as behavioral covariates, could provide reliable current estimates with finer granularity than at the level of individual regions and possibly across different behavioral states.

Two-region model producing idealized, ground truth, simulated data to validate CURBD

Design of the generator model

We simulated a model that generated idealized ground truth data to test when CURBD approach would be the most effective at disentangling inter-region interactions, and to probe the conditions under which it would fail to perform optimally. We generated two 1000 unit RNNs, each with random connectivity weights drawn from a Gaussian distribution, as described in the initialization procedure for the Model RNN above. One RNN (corresponding to Region A) was driven by an external sinusoidal signal, $S_A[t]$, oscillating at $\pi/8$ Hz, while the second (corresponding to Region B) was driven by another sinusoid, $S_B[t]$, oscillating at $\pi/3$ Hz and phase shifted by $\pi/3$. The two sinusoidal inputs began after two seconds of a simulated "resting state" during which the inputs to the RNNs were set to zero. These external inputs were connected to 33% of the units in their respective RNN with a fixed input weight, picked from a uniform distribution. The two RNNs were recurrently connected, with a varying percentage of neurons in each region (randomly selected) receiving inputs from the other region with a fixed weight of one. We computed the time-series current activity of the i^{th} unit in the two RNNs, $x_{A,\,i}[t]$ *and* $x_{B,\,i}[t]$ for ten seconds of data using the following steps. We initialized the states of the two RNNs to random values between -1 and 1. For each subsequent time step, we computed the change in activity of each RNN unit *i* based on its inputs according to:

$$\Delta x_{A,i}[t] = \boldsymbol{J}_A \boldsymbol{r}_{A,i}[t-1] + w_{rgn} C_{BtoA,i} \boldsymbol{r}_B[t-1] + w_{in} C_{StoA,i} S_A[t-1] \quad \text{(Equation 15)}$$

$$\Delta x_{B,i}[t] = \boldsymbol{J}_B \boldsymbol{r}_{B,i}[t-1] + w_{rgn} C_{AtoB,i} \boldsymbol{r}_A[t-1] + w_{in} C_{StoB,i} S_B \quad \text{(Equation 16)}$$

C_{StoA} and C_{StoB} define binary connectivity vectors describing the connectivity of the external sinusoidal inputs to their respective regions. Similarly, C_{AtoB} and C_{BtoA} are binary vectors describing the connectivity between regions A and B. The fraction of entriess in the above inter-region connectivity vectors set to 1 is defined as p_{rgn}. w_{rgn} and w_{in} are scalars that set the connection weights for the sinusoidal inputs and inter-region connections, respectively. The $\mathbf{J}_A$ and $\mathbf{J}_B$ matrices are initialized to ${g_A}^2/N$ and ${g_B}^2/N$, where the scaling parameters g_A and g_B control how chaotic each RNN is, as described in the Model RNN training section above. The activity of each RNN unit *i* was then computed according to:

$$x_{A,i}[t] = x_{A,i}[t-1] + \frac{\Delta t\left(-x_{A,i}[t-1] + \Delta x_{A,i}[t]\right)}{\tau}, \quad \text{(Equation 17)}$$

$$x_{B,i}[t] = x_{B,i}[t-1] + \frac{\Delta t\left(-x_{B,i}[t-1] + \Delta x_{B,i}[t]\right)}{\tau}, \quad \text{(Equation 18)}$$

Lastly, the activity of each RNN unit *i* was transformed into a firing rate by passing through the nonlinearity, as described above:

$$\boldsymbol{r}_{A,i}[t] = tanh\left(x_{A,i}[t]\right), \quad \text{(Equation 19)}$$

$$\boldsymbol{r}_{B,i}[t] = tanh\left(x_{B,i}[t]\right), \quad \text{(Equation 20)}$$

Checking robustness of CURBD over a range of simulation parameters for the generator model

We repeated the ground truth simulations sweeping over a broad range of parameters applicable to the generator model (Figure S1L): 1) g_A, the dynamical regime of Region A; 2) w_{rgn}, the strength of the weights of the recurrent connections between the networks; and 3) p_{rgn}, the proportion of neurons in Regions A and B receiving input from the other region. This parameter-sweeping process helped us explore how effectively CURBD operates to untangle currents resulting from the external sinusoidal inputs as the properties of the modeled networks change. The remaining parameters were fixed for all of these simulations. Values for all of the parameters we tested are provided in Table S2.

For each combination of parameters in the generator model, we trained a 2000-unit Model RNN to reproduce the activity of the two Regions from the generator model using the algorithm described above. The parameters chosen for this "fit" Model RNN (Table S1) were held fixed to ensure we studied the effect *only* of the network properties as we swept parameters, not of variations in the Model RNN. For each of these Model RNNs, we applied CURBD to Region A to assess how effectively we could isolate the currents from the two external inputs (the sinusoidal input driving Region A and the sinusoidal input driving Region B) from the population. Since each external input was effectively one-dimensional, we first reduced each estimated population-wide source current to a single component using PCA. We then computed how accurately we could infer the external inputs by directly comparing the correlation coefficient (denoted by R^2 in Figure S1K) between the leading PC of each source current and each of the two sinusoidal external

inputs. We repeated this analysis for different combinations of parameters to assess which parameter regimes consistently gave the highest R^2 values.

Three-region ground truth simulation to validate CURBD

Design of the generator model

We designed a second idealized ground truth simulation to validate whether CURBD could effectively infer source currents between different interacting regions even when there are no external inputs driving a particular neural population. We simulated three 1000-unit RNNs using a generator model similar to the two-region model described above (Figure 3A). However, rather than sinusoidal inputs as in the two-region ground truth model, two of the three interconnected RNNs received time-varying patterns of inputs from other model networks. The external inputs driving Region B were provided by a network generating a Gaussian "bump" propagating across the network sequentially, SQ(t). The sequence began 2 seconds after the start of the simulation and ended 4 seconds later, with each sequentially activating unit $i=1, 2,\ldots, N$ behaving according to:

$$SQ_i(t) = e^{-\frac{1}{2\sigma^2}\left(\frac{i-\sigma-Nt}{T}\right)^2} \quad \text{(Equation 21)}$$

where σ denotes the width of the bump across the population (here, 20% of the units), N represents the population size (here, 1000 units), and T represents the total simulation time of twelve seconds. The external input to Region C was provided by another 1000-unit network generating a fixed point, $FP(t)$, for 8 seconds that instantaneously shifted to a new fixed point for an additional 4 seconds. The fixed points were generated by sampling SQ(t) at two different time points (t=2s and t=5s) and holding them at the sampled value of firing rate for the duration of the fixed point. The external inputs were connected to 50% of the units in their respective regions (randomly selected) with a fixed negative weight (inhibitory) for Region B and positive weight (excitatory) for Region C. The third RNN (Region A) received only the recurrent inputs from the other two RNNs, no external drive. The Region A RNN was modeled at a different value of g from the other two networks, yielding distinct dynamics (Table S3). The following update equations (again, discrete implementations with square brackets) governed the interactions of the regions at each time step (with a resolution of $\Delta t=0.01$), with the subsequent activity evolving similarly to the two-region simulation described by Equations 18–21:

$$x_{A,i}[t] = \boldsymbol{J}_A \boldsymbol{r}_{A,i}[t-1] + w_{rgn} C_{BtoA,i} \boldsymbol{r}_B[t-1] + w_{rgn} C_{CtoA,i} \boldsymbol{r}_C[t-1] \quad \text{(Equation 22)}$$

$$\Delta x_{B,i}[t] = \boldsymbol{J}_B \boldsymbol{r}_{B,i}[t-1] + W_{rgn} C_{AtoB,i} \boldsymbol{r}_A[t-1] + w_{rgn} C_{CtoB,i} \boldsymbol{r}_C[t-1] + w_{in} C_{SQtoB,i} SQ[t-1] \quad \text{(Equation 23)}$$

$$\Delta x_{C,i}[t] = \mathbf{J}_C \boldsymbol{r}_{C,i}[t-1] + W_{rgn} C_{AtoC,i} \boldsymbol{r}_A[t-1] + w_{rgn} C_{BtoC,i} \boldsymbol{r}_B[t-1] + W_{in} C_{FPtoC,i} FP[t-1] \quad \text{(Equation 24)}$$

As in the previous simulation, the **J** matrices for each of the three networks were initialized based on the selected g_A, g_B, and g_C parameters.

Description of the Model RNN and CURBD analysis

We trained a 3000-unit Model RNN to match the activity of the three-region generator model using the procedure described above. We found that the Model RNN reproduced the simulated data accurately over a wide range of parameters; for the simulations reported in this paper (Figures 2 and S1), we used the values reported in Table S1. We performed CURBD to infer the nine source currents governing interactions between the three regions. We reduced the dimensionality of the full 1000-unit population of each of the three regions and the source currents using PCA. We chose the leading five dimensions for the following analyses, which sufficed to capture more than 95% of the total variance in each source current, though we observed similar results with other assumed dimensionalities (data not shown).

Since the true connectivity of the network was defined in the simulated dataset, we computed the ground truth currents between each region to isolate the effect of isolated inputs from the source regions on the target region, including how the input activity would propagate through the recurrent connections of the target region. We adapted the update equations defined above to compute the three current sources into one region at each time step, here using Region A as an example:

$$Ix_{GT,AtoA,i}[t] = \boldsymbol{J}_A \boldsymbol{r}_{A,i}[t-1] \quad \text{(Equation 25)}$$

$$Ix_{GT,BtoA,i}[t] = \boldsymbol{J}_A W_{rgn} C_{BtoA,i} \boldsymbol{r}_B[t-1] \quad \text{(Equation 26)}$$

$$Ix_{GT,CtoA,i}[t] = \boldsymbol{J}_A W_{rgn} C_{CtoA,i} \boldsymbol{r}_C[t-1] \quad \text{(Equation 27)}$$

The same process was performed for Regions B and C using similar equations. We performed the same dimensionality reduction analysis on the ground truth currents as on the inferred source currents from CURBD. Since the Model RNN was trained to reproduce time-varying activity from all the units in the multi-region generator model, each inferred source current has the same dimensionality as the ground truth current, and is embedded within the same high-dimensional space of the population activity of the respective simulated region. We could thus directly compare each leading PC using VAF (Figures 3G and S1).

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

$$VAF = 1 - \frac{\sum_{t}^{T} [Ix_{GT}(t) - Ix(t)]^2}{\sum_{t}^{T} [Ix_{GT}(t) - I\bar{x}_{GT}(t)]^2}, \quad \text{(Equation 28)}$$

Comparison of inferred inter-region currents to shared dynamics identified by Canonical Correlation Analysis

We compared the performance of CURBD to an analogous decomposition obtained by canonical correlation analysis (CCA).[42,114] CCA obtains an optimal linear transformation relating the dimensionality-reduced population activity of the source and target regions to identify shared dynamics. In brief, we first took the low-dimensional trajectories of each region and performed a QR decomposition to identify for each region of the resulting **Q**, which provides an orthonormal basis for the column space of the low dimensional trajectories. For any pair of regions, for example Region A and Region B, we performed a singular value decomposition of the inner product of the corresponding **Q** matrices:

$$Q_A' Q_B = USV' \quad \text{(Equation 29)}$$

This process effectively finds new dimensions within the manifold of Region A (denoted by **U**) and Region B (denoted by **V**) that maximize the correlation between the two trajectories. To analyze the shared dynamics between the two regions, we projected the activity of either region onto the corresponding axes. Unlike CURBD, the mapping obtained from CCA (and similar methods of inferring functional connectivity only from the covariance matrix of recorded neural activity) is not directional and is purely correlational. Thus, only one "current" can be obtained for each pair of regions. We compared the VAF by the first component identified by CCA to the first PC of the ground truth currents to assess the effectiveness of this approach (Figure S1E).

Addressing partial sampling issues present in experimental data

We repeated the above simulation to determine whether CURBD is effective when only a fraction of the total multi-regional activity is 'observed' by the Model RNN. This control analysis addresses partial sampling issues present in real data when activity can be experimentally measured from only a relatively small fraction of the total number of neurons in a region. To simulate this scenario and test the efficacy of CURBD in the face of partial sampling issues, we trained Model RNNs to match activity from 5%, 20%, 50%, and 100% of the available neurons in each region (randomly selected) of the ground truth multi-region generator model. We repeated the simulation ten times at each subsampling level to help account for variability in the random sampling of neurons, as well as variability from different random initializations of the **J** matrix. Such variability scales inversely with network size for Gaussian weights; the ten repetitions at 100% sampling thus provide a lower-bound on the variability that would be expected within this model. Unlike the initial Model RNN analysis where every neuron was sampled, in the subsampling case, we can no longer guarantee that the axes should be oriented similarly in PC space and VAF is not a reliable measure of how well the method performed. Thus, here we again employed CCA not to identify shared dynamics between regions, but to compensate for differences in the number of sampled neurons generating the dynamics of a single region,[42] In this application, CCA provides a quantitative "similarity index"-quantified by the canonical correlation of the leading aligned dimension–of the population dynamics between the currents identified by CURBD and the ground truth currents (Figure 3H).

Multi-region calcium fluorescence recordings in larval zebrafish

We analyzed neural recordings from larval zebrafish in a behavioral challenge paradigm.[11]

Experimental setup

All procedures were approved by the Stanford University Institutional Animal Care and Use Committee. The experimental details have been previously described in Ref. [11]. In brief, five larval zebrafish expressing nucleus-localized GCaMP6s — homozygous Tg(elavl3:H2B-GCaMP6s)[53] — were immersed in water and partially head-fixed in agarose (Thermo Fisher Scientific 16250–100), leaving their tails free to attempt swimming movements. The fish were illuminated with 850nm infrared light and tail movements were recorded using a high-speed camera (Allied Vision Technology Manta G-031B) and macro lens (Nikon AF-S DX Micro Nikkor 85 mm f/3.5G ED VR) through both short-pass (Thorlabs FES0900-1) and long-pass (Thorlabs FEL0800-1) filters. The spontaneous neural activity and movements of the fish were recorded for an extended period of time (48+ minutes). In one cohort of 5 fish, gentle but inescapable electrical shocks were applied to the water in the dish[11] which induced stress responses. In a second cohort of 5 fish, no shocks were applied but the recordings were made for the same duration.

Image acquisition and processing

Images were acquired with an Olympus FVMPE multiphoton microscope (Olympus) using fast piezo axial scanning at 1–1.2 Hz to image a volume of 562 x 241 x 76mm on average with 6–8 mm between planes. 920nm light was used to excite the GCaMP6S sensors. Light emission was collected through a 485 nm dichroic mirror and 495–540 nm emission filter. After collection, drift between images was corrected using an affine robust alignment by sparse and low-rank decomposition[115] and regions of interest (ROIs) were identified as the 90th percentile of each plane. These ROIs were pruned based on size and template-matching to nucleus-sized disks to isolate putative neurons. Fluorescent activity (F) was normalized to compute a typical $\Delta F/F$ trace by subtracting and dividing by the fluorescence value of all pixels within each ROI and defining F as the 5th percentile of all values in the session. Cells were

then manually assigned to brain regions based on anatomical landmarks. In this work, we considered only cells identified as part of the telencephalon, habenula, thalamus, or raphe nucleus (Table S4).

Model RNN fitting

We used the 2-region larval zebrafish dataset to compare CURBD output using both current-based and rate-based learning rules. We used the 3-region larval zebrafish dataset to study the mechanisms underlying behavioral state transitions. For each zebrafish, we trained Model RNNs with 2 regions (Fish cA: 4372 units; Fish cB: 2937 units; Fish cC: 1578 units; Fish cD: 4303 units; Fish cE: 2171 units) or 3 regions (Fish cA: 5432 units; Fish cB: 3861 units; Fish cC: 1678 units; Fish cD: 5063 units; Fish cE: 2333 units; Fish sA: 4704 units; Fish sB: 3908 units; Fish sC: 1784 units; Fish SD: 2177 units; Fish SE: 3857 units) to reproduce the time-series Ca^{2+} data from the multi-regional activity. We used identical model parameters for all fits to all fish (Table S1). Before training each model, we randomly initialized the **J** matrix and white noise inputs as described above. For the 2 region models, in order to most accurately compare the learning rules, we fixed these initial conditions and trained the Model RNN to fit the experimental data using one of the two learning rules. We then repeated this procedure five additional times to assess the reproducibility over runs. Each data sample contained 6 minutes of recordings from the control cohort of fish. For the 3-region models, we used rate-based training and divided the 48 minute dataset into 8 epochs of 6 minutes each. We sequentially trained Model RNNs on each epoch starting from an initial random **J** matrix in epoch 1, and using the resultant **J** post-training as the initial guess for the subsequent epoch (a process called *chaining*). We then concatenated the current estimates identified within each epoch to reconstruct the interactions in the full 48 minute dataset.

Analysis of model output to compare learning rules

We compared the similarity of the resulting **J** matrices after rate-based and current-based learning by computing the two-dimensional correlation of the matrices (Figure S2). As a reference, we compared the correlations obtained across different training runs with the same learning rule as well as after shuffling each matrix. We then used CURBD to compute the current sources driving the telencephalon populations. We compared the currents obtained by each learning rule by using PCA to compute the leading 10 components and computing the vector norm at each time point within this reduced dimensional space (Figure 4I). We lastly computed the similarity within this low-dimensional space using CCA as described above (Figure 4J).

Analysis of inter-region currents during the behavioral challenge paradigm

We analyzed the behavior of the two cohorts of fish. We identified timestamps of each tail whipping movement using the high speed cameras monitoring the fish. We then convolved this sequence of timestamps with a Gaussian kernel (60s kernel width) to produce an estimate of the continual movement rate. We then averaged this movement rate across fish in each cohort (Figure 5I). To study the large-scale temporal dynamics of each current source, we decomposed the regional inputs to the Habenula and then performed PCA separately to reduce the activity in the space of habenular neurons to a single leading dimension (Figure S3). We then averaged this trajectory across fish in each cohort to compare the temporal modulation of the dominant patterns of inputs from each region to habenula (Figure 5J). Lastly, we analyzed the directed interaction weights to look for possible underlying mechanisms for the current dynamics. In each epoch we computed the distribution of weights in the sub-matrix capturing inputs from the raphe nucleus to the habenula. We computed this distribution for each fish of the shocked cohort in the 6 minute baseline epoch (pre-shocks), the first epoch of shocks, and the final epoch of the experiment (after the transition to passivity) and averaged these distributions across fish (Figure 5K).

Multi-region calcium fluorescence recordings in mice

We analyzed neural recordings from mice during spontaneous running in the dark.[14]

Surgery

All experimental procedures were approved by the Harvard Medical School Institutional Animal Care and Use Committee and were performed in compliance with the Guide for the Care and Use of Laboratory Animals. Two female mice expressing GCaMP6s (C57BL/6J-Tg(Thy1-GCaMP6s)GP4.3, The Jackson Laboratory, stock 024275) were implanted with cranial windows over the cortical surface. Mice were 3–5 months old at the time of surgery, and given an injection of dexamethasone (3 μg per g body weight) 4–8 h before the surgery. Mice were anesthetized with isoflurane (1%–2% in air). A cranial window surgery was performed to either fit a 'crystal skull' curved window (LabMaker UG) exposing the dorsal surface of both cortical hemispheres,[46] or to fit a stack of custom laser-cut quartz glass coverslips (three coverslips with #1 thickness each (Electron Microscopy Sciences), cut to a 'D'-shape with maximum dimensions of 5.5 mm medial-lateral and 7.7 mm anterior-posterior, and glued together with UV-curable optical adhesive (Norland Optics NOA 65), exposing the left cortical hemisphere. The dura was removed before sealing the window using dental cement (Parkell). A custom titanium headplate was affixed to the skull using dental cement mixed with carbon powder (Sigma-Aldrich) to prevent light contamination. A custom aluminum ring was affixed on top of the headplate using dental cement. During imaging, this ring interfaced with a black rubber balloon enclosing the microscope objective for light-shielding.

Imaging and behavior setup

Data were collected using a large field of view two-photon microscope assembled as described in Ref. [58]. In brief, the system contained a combination of a fast resonant scan mirror and several large galvanometric scan mirrors allowing for especially large scan angles. Paired with a remote focusing unit to rapidly move the focus depth, this setup enabled random access imaging in a field of view of 5-mm diameter with 1 mm depth. The setup was assembled on a vertically mounted breadboard whose XYZ positions and rotation were controlled electronically via a gantry system (Thorlabs). Thus, to position the imaging objective with regards to the

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

mouse, the position and rotation of the entire microscope were adjusted while the position of the mouse remained fixed. Mice were head-fixed and placed on an air-suspended 8-inch diameter Styrofoam spherical treadmill that enabled spontaneous running. Using two optical sensors (ADNS-9800, Avago Technologies), we tracked the treadmill velocity, which was translated into pitch, roll, and yaw velocity using custom code on a Teensy microcontroller (PJCR) as a readout of the mouse's running speed and direction. Individual recording sessions lasted from 45–60 minutes. Mice were extensively acclimated to head-fixation and running on the treadmill before data collection. We recorded behavioral and neural activity while mice spontaneously ran on the ball. The room was kept in complete darkness throughout the experiment. We defined running bouts as periods when the ball movement speed crossed a fixed threshold set to be the 90th percentile of the running speeds throughout the session.

Image acquisition

The excitation wavelength was 920 nm, and the average power at the sample was 60–70 mW. The microscope was controlled by ScanImage 2016 (Vidrio Technologies). We targeted four distinct regions in the left cortical hemisphere: primary visual cortex (V1), secondary motor cortex (M2), posterior parietal cortex (PPC), and retrosplenial cortex (RSC). These regions were targeted based on retinotopic mapping (see below). In each region, we acquired images in layer 2/3 from two planes spaced 50 μm in depth, at 5.36 Hz per plane at a resolution of 512 x 512 pixels (600 μm x 600 μm).

Retinotopic mapping for selecting Ca2+ imaging locations

We performed retinotopic mapping in the mice used for calcium imaging experiments as previously described in Ref. [13]. Mice were lightly anesthetized with isoflurane (0.7%–1.2% in air). GCaMP fluorescence was imaged using a tandem-lens macroscope where excitation light (455 nm LED, Thorlabs) was filtered (469 nm with 35 nm bandwidth, Thorlabs) and reflected onto the brain through a camera lens (NIKKOR AI-S FX 50 mm f/1.2, Nikon) focused 400 μm below the brain surface. GCaMP emission light was collected using the same lens, filtered (525 nm with 39 nm bandwidth, Thorlabs), and imaged with another camera lens (SY85MAE-N 85 mm F1.4, Samyang) and a CMOS camera at 60 Hz (ace acA1920-155um, Basler). These images were synchronized to visual stimuli presented on a gamma-corrected 27 inch IPS LCD monitor (MG279Q, Asus). The monitor was centered in front of the mouse's right eye at an angle of 30 degrees from the mouse's midline. The visual stimulus, a spherically corrected black and white checkered moving bar[116] (12.5 degree width, 10 deg/s speed), was presented in 7 blocks, each consisting of 10 repeats of 4 movement directions (up, down, forward, backward). To produce retinotopic maps, we calculated the temporal Fourier transform at each pixel of the imaging data and extracted the phase at the stimulus frequency.[117] These phase images were averaged across repetitions for a given movement direction and smoothed with a Gaussian filter (25 μm s.d.). Lastly, we calculated field sign maps by computing the sine of the angle between the gradients of the average horizontal and vertical retinotopic maps.

For each retinotopic mapping session, we also acquired an image of the superficial brain vasculature pattern under the same field of view. We then acquired a similar brain vasculature image under the large field of view two-photon microscope. These two reference images were manually aligned and used to directly locate V1 and PPC locations for two-photon imaging. The location for RSC imaging was positioned adjacent to the midline and about 300 μm anterior of the PPC location. The location for M2 imaging was positioned one millimeter anterior of the RSC location.

Pre-processing of imaging data

We used custom code to correct for motion artifacts, as described in Ref. [111]. In brief, motion correction was implemented as a sum of shifts on three distinct temporal scales: sub-frame, full-frame, and minutes- to hour-long warping. After motion correction, ROIs were extracted using Suite2P.[110] Afterwards, somatic sources were identified with a custom two-layer convolutional network in MATLAB trained on manually annotated labels to classify ROIs as neural somata, processes, or other.[111] Only somatic sources were used. This yielded large populations from neurons from each of the four targeted regions in Mouse A and Mouse B (Table S5).

After identifying individual neurons, we computed average fluorescence in each ROI and converted this value into a normalized change in fluorescence ($\Delta F/F$). We corrected the numerator of the $\Delta F/F$ calculation for neuropil by subtracting a scaled version of the neuropil signal estimated per neuron during source extraction:

$$F_{corr} = F - 0.7F_{neuropil} \quad \text{(Equation 30)}$$

We estimated the baseline fluorescence (F_{base}) of this trace as the 8th percentile of fluorescence within a 60-s window and subtracted this baseline to get the numerator:

$$\Delta F = F_{corr} - F_{base} \quad \text{(Equation 31)}$$

We divided this by the baseline (again 8th percentile of 60s window) of the raw fluorescence signal to get $\Delta F/F$. We deconvolved the $\Delta F/F$ trace per neuron using the constrained AR-1 OASIS method44. We initialized the decay constants at two seconds and then optimized separately for each neuron. To fit the Model RNN, we temporally smoothed the sparse deconvolved spike estimates using a Gaussian kernel with four times the width of the sampling rate. We applied the same filter to the behavioral signals (pitch, roll, and yaw of the ball) to preserve the temporal relationship with the neural activity. For visualization of the neural population activity in the heatmaps of Figures 6 and S4, we scaled each neuron by the mean of the total activity.

CURBD analysis to infer source currents from mouse data

For each mouse, we trained Model RNNs (Mouse A: 2787 units; Mouse B, Session 1: 999 units; Mouse B, Session 2: 1444 units) to reproduce the time-series Ca^{2+} data from the four regions. We used identical parameters for each Model RNN (Table S1).

Please cite this article in press as: Perich et al., Inferring brain-wide interactions using data-constrained recurrent neural network models, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.07.016

We applied CURBD to infer the sixteen source currents comprising the multi-region population activity. We first assessed how much unique explanatory power each source current had in the total V1 population. We developed a partial coefficient of determination analysis to quantify this as follows. We subtracted each source current one-by-one from each V1 neuron and computed the sum-squared error of this difference and the recorded neural data. We then computed the sum-squared error of the full Model RNN fit compared to the recorded neural data. We defined the unique variance explained by the source current according to:

$$VAF_{unique} = 1 - \frac{\sum_{t}^{T} \{a(t) - [\boldsymbol{r}(t) - I(t)]\}^2}{\sum_{t}^{T} [a(t) - \boldsymbol{r}(t)]^2} \quad \text{(Equation 32)}$$

where *I(t)* denotes the partial source current that is being evaluated. Effectively, this computes the variance that cannot be explained by any of the three remaining source currents. Importantly, this metric can be computed at individual time points. We normalized each calculation by the sum of the four unique variances at each time point to give a proportion of unique variance explained by each source current. For cleaner presentation, we smoothed these normalized traces with a Gaussian kernel of width 500 ms (Figure 6H).

We then reduced the dimensionality of all sixteen source currents using PCA, selecting a 5-dimensional manifold which sufficed to explain more than 80% of the total variance in all source currents. We trained Wiener cascade filters, a type of linear-nonlinear decoder,[118] to predict the running speed using the five-dimensional activity of each source current at each time step as well as the most recent 5 time steps of history. To perform cross validation, we randomly withheld 20% of time steps (the test set) and trained the decoders using the remaining 80% of the data. We quantified the performance of each decoder output on the left-out test set of time steps using VAF, as described above. We repeated this process for 100 iterations, randomly leaving out 20% of time steps for the test set on each iteration, and averaged across all iterations for the final decoder performance (Figures 6K and 6L).

Multi-region electrophysiology recordings in monkeys

We analyzed neural recordings from macaque monkeys during Pavlovian conditioning.[15]

Behavioral task

All procedures were reviewed and approved by the Icahn School of Medicine Animal Care and Use Committee. For detailed descriptions of the experimental setup and protocol, see Ref. [15], where these data were previously reported. In brief, two rhesus macaque monkeys (*Macaca mulatta*; Monkey D: female, 5.6 kg; Monkey H: Male, 11.0 kg) were trained to sit in a custom primate chair with their head restrained and fixate on a computer monitor for four seconds, before performing a Pavlovian conditioning task for liquid rewards. They fixated on a neutral gray square for 800–1000ms. They were then presented with one of three visual conditioned stimuli for 500–600 ms on each trial corresponding to three different reward outcomes: no reward, water (0.5 mL), and juice (0.5 mL). An additional trial type occurred with equal frequency in which no conditioned stimuli was presented, and the gray square persisted throughout the trial. On all trials, a small (0.1 mL) water reward was given two seconds after the stimulus onset. Conditioned stimuli varied between monkeys and consisted of gray shapes, covering 1.10° of visual angle for Monkey D and 2.45° for Monkey H. We trained the Model RNNs described below using all four trial types to utilize as much training data as possible, though in this paper we only analyzed the CURBD output for the three conditioned stimuli.

Surgical procedures and neural recordings

After training, each monkey was implanted with a titanium head restraint device followed by a plastic recording chamber over the exposed cranium of the left frontal lobe. During the behavioral experiments, tungsten microelectrodes (FHC, Inc. or Alpha Omega, 0.5–1.5 M at 1 KHz) or 16-channel multi-contact linear arrays (Neuronexus Vector array) advanced by an 8-channel micromanipulator (NAN instruments, Nazareth, Israel) were attached to the recording chamber and inserted into the brain. The targeted brain regions were located using stereotaxic coordinates and verified by T_1-weighted MRI imaging with the electrodes implanted. Recordings from subcallosal ACC were made on the medial surface of the brain ventral corpus callosum. Amygdala recordings were made between 22 and 18.5 mm anterior to the interaural plane. Rostromedial striatum recordings were made in the anterior medial segment corresponding to the zone where subcallosal and basal amygdala projections overlap. Spikes from putative single neurons were captured online using a Plexon Multichannel Acquisition Processor and later isolated with Plexon Offline Sorter. The small number of neurons recorded in each experimental session were then pooled into a pseudopopulation. First, the spike trains for each neuron on each trial were converted to an estimated firing rate. The firing rates were aligned on the stimulus presentation for each trial, then averaged across all trials of each stimulus type for that session in each monkey to give substantially large pseudopopulations (Table S6). As with the mouse dataset, neural activity was visualized with a heatmap after scaling each neuron's firing rate by its mean activity (Figures 7B and S6).

CURBD analysis of monkey dataset

For each monkey, we trained Model RNNs (Monkey D: 343 units; Monkey H: 199 units) to match the pseudopopulation data for all four conditions. We used identical parameters for the Model RNNs for both monkeys (Table S1). In the previous simulations and the mouse dataset, the Model RNN learned a single dynamical system that reproduces the neural data based on one initial condition. However, here we have four different initial conditions corresponding to the four trial types, and we seek to learn a single dynamical system that reproduces all of them. To achieve this, we concatenated the time-series data from the four conditions and reset the

state of the Model RNN to match the real neural data at the first time point of each new condition. We repeated each Model RNN fit an additional four times, yielding five runs in total, each starting from a different randomly initialized matrix $\boldsymbol{J_0}$ each time. We performed CURBD using each Model RNN to infer the nine source currents comprising the full multi-region activity. We then quantified the magnitude of current arriving to each source region from each target region by summing the absolute value of the source currents at each time step. We averaged this across the five runs in each monkey to assess the consistency of our solutions (Figures 7F and 7G).

We performed a systematic subsampling analysis to assess whether applying CURBD to pseudopopulation data would be reliable even if different numbers and types of neurons were recorded experimentally. We randomly subsampled between 50% and 90% of the available neurons to create new, smaller pseudopopulations for each monkey. We used CCA (similar to the description in X above) to compute a similarity metric between the currents inferred by CURBD from each subsampled population and the currents originally inferred by CURBD using all of the available neurons for each monkey (Figure S7).

Multi-region electrophysiology recordings in humans

We analyzed intra-operative neural recordings from human participants performing a long-term memory task.[16]

Behavioral task

The institutional review boards of Cedars-Sinai Medical Center and the California Institute of Technology approved all protocols. Detailed descriptions of the experimental procedures are described in Ref. [16], where these data were previously reported. All datasets are available in a public repository.[119] In brief, recordings were made from thirteen adult participants being evaluated for surgical treatment of drug-resistant epilepsy that provided informed consent and volunteered for this study. Of these thirteen participants, eight did not have a sufficiently large number of neurons to create a population for CURBD and were thus excluded. Our final analyses focused on five participants (P44, P51, P56, P57, and P58 from the original manuscript). The participants were seated in a chair facing a screen and reported decisions using either button presses or eye movements. They each performed eight forty-trial blocks that alternated between two tasks. In the categorization task, participants classified pseudorandomly presented images as belonging to one of four target categories (human faces, monkey faces, fruits, or cars) with a "yes" or "no" response. In the memory task, participants were shown an image and asked "Have you seen this image before, yes or no?" to which they responded "yes" or "no". In the first block, all images were necessarily novel (40 unique images). In all subsequent blocks, the participants viewed 20 new images that were randomly intermixed with 20 familiar images. The 20 repeated images remained the same throughout the remainder of the session. We ignored trials where the participant provided an incorrect response (e.g. mistakenly identifying a novel image as familiar). This gave sixteen different conditions: four blocks of trials for each of two different tasks, each with correct "yes" and "no" responses. Participant task performance tended to improve throughout the session. Thus, we primarily focused our analysis on blocks 4, 6, and 8 (the final three memory blocks) when performance was highest.

Neural recordings

As participants performed this task, we recorded bilaterally from the amygdala, hippocampus, pre-supplementary motor area (pre-SMA), and dorsal anterior cingulate cortex (dACC) of each participant using microwires embedded in a hybrid electrode.[69] Electrode locations were confirmed using post-operative MRI or CT scans. We identified putative single neurons using a semi-automated spike sorting procedure. For the purposes of the CURBD analyses, we pooled recordings from each region from either hemisphere to ensure that we had sufficiently large populations for the dimensionality-reduction analysis. We estimated the instantaneous firing rate of each recorded neuron by convolving the spike train with a Gaussian kernel of width 150 ms. The relatively large width was necessary to accurately estimate firing rates for many low-firing neurons in the hippocampus and amygdala, and we opted to use a uniform width for all neurons. We created pseudopopulations by aligning each trial on the time of stimulus presentation and averaging across all trials for each task and correct response type (Table S7). In each trial, we kept two seconds before the stimulus presentation and three seconds after the stimulus presentation. As with the previous datasets, neural activity was visualized with a heatmap after scaling each neuron's firing rate by its mean activity (Figures 8F and S8).

CURBD analysis of human dataset

We trained Model RNNs based on the spiking pseudopopulation data from all sixteen conditions for each participant. As with the monkeys, we compensated for the discontinuities at trial boundaries by resetting the state of the Model RNN at the start of each condition. We used the same Model RNN parameters for all participants to ensure consistency (Table S1). We applied CURBD to infer the nine source currents comprising the multi-region interactions in the dataset. We reduced the dimensionality of each source current, as well as the full population activity of each region, using PCA. Since the stimulus was presented two seconds after the start of the trials, we defined the first two seconds to be the 'resting state' (Figure 8G). We then computed the Mahalanobis distance of each source current at each time step. This gave a time-varying estimate of how much the population activity or source currents responded to the stimulus. For each participant, we averaged across all five training runs to obtain the most reliable estimate of the current dynamics. We then averaged across all participants to show the group effect (Figure 8H).

Neuron, Volume *114*

Supplemental information

Inferring brain-wide interactions using data-constrained recurrent neural network models

Matthew G. Perich, Charlotte Arlt, Sofia Soares, Siyan Zhou, Manuel Beiran, Aaron S. Andalman, Tyler Benster, Megan E. Young, Clayton P. Mosher, Juri Minxha, Eugene Carter, Ueli Rutishauser, Peter H. Rudebeck, Christopher D. Harvey, Karl Deisseroth, and Kanaka Rajan

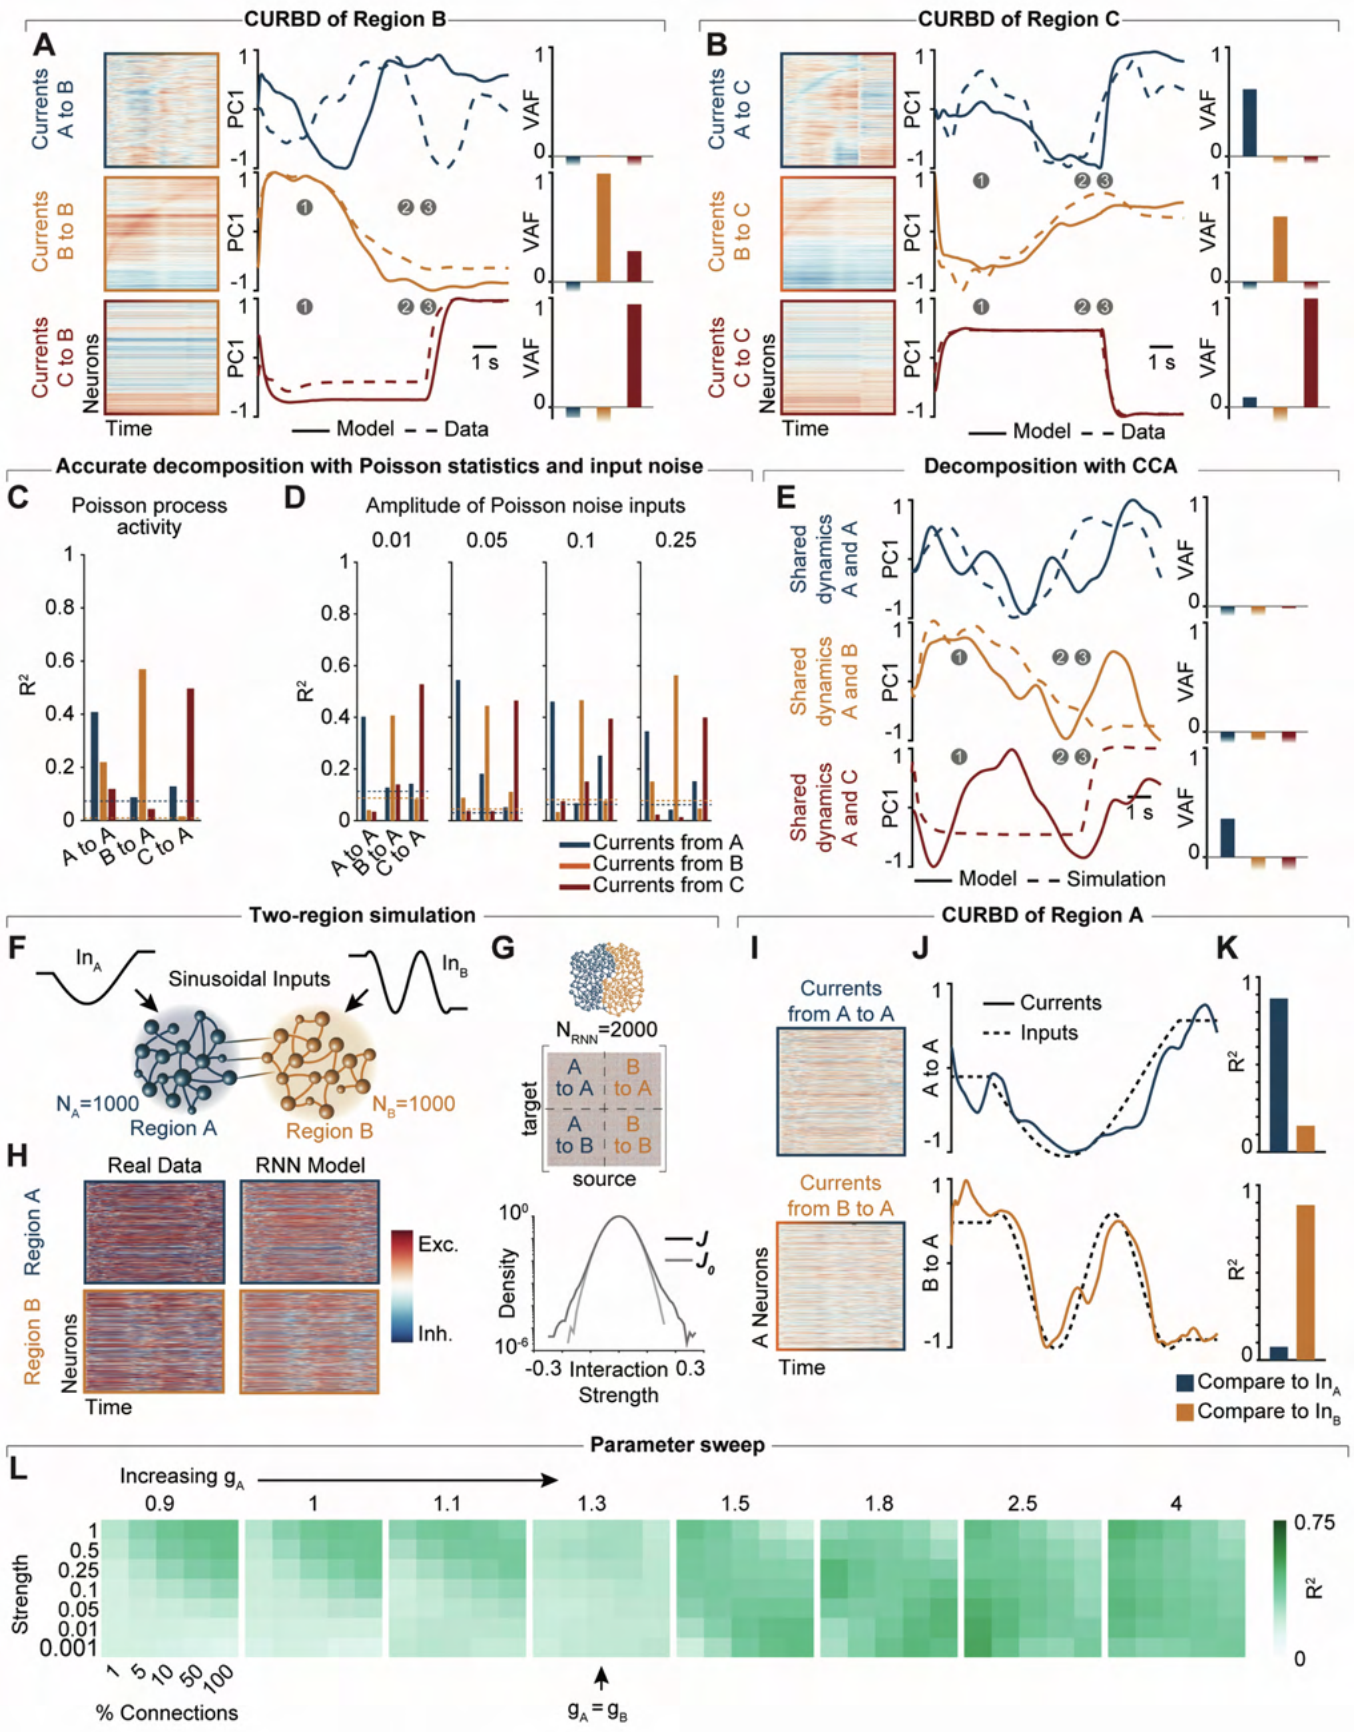

Figure S1. Supporting data for three-region ground truth simulations, related to Figure 3.

(A) Analysis of source currents within Region B. Data presented as in Figure 3E. Currents from Region B and Region C are accurately reconstructed, though currents from Region A are missed, presumably due to the lack of strong external drive to Region A and the similar intrinsic dynamics between the three regions ($VAF_{AtoB}<0$; $VAF_{BtoB}=0.98$; $VAF_{CtoB}=0.94$).

(B) Analysis of source currents within Region C. Data presented as in Figure 3E. All three source currents are accurately reconstructed ($VAF_{AtoC}=0.61$; $VAF_{BtoC}=0.60$; $VAF_{CtoC}=0.99$).

(C) We repeated the 3-region RNN simulation by simulating spiking networks with a Poisson process applied to the RNN unit activations. We then smoothed the putative spike counts to reconstruct firing rates and performed CURBD. CURBD accurately recovered the ground truth interactions with the added Poisson statistics.
(D) We drove the 3-region RNN simulation with simulated distractor networks using new inputs of Poisson processes applied to all units with random weights. CURBD accurately recovered the ground truth interactions even at high amounts of Poisson input noise.
(E) Comparison of ground truth current inputs and shared dynamics identified by CCA. CCA finds a single space capturing shared dynamics between each region, with a linear transformation (provided by the weight matrices *w*) relating each source and target region. However it does not provide a directional estimate of interactions. The shared dynamics plots show the aligned trajectories between pairs of regions projected onto the leading two aligned components.Unlike CURBD, the shared dynamics identified by CCA do not accurately match the ground truth current dynamics ($VAF_{AandA}<0$; $VAF_{BandA}<0$; $VAF_{CandA}<0$).
(F) We simulated two interconnected RNNs representing distinct brain regions. Each was driven by a sinusoid of different frequencies.
(G-H) We fit a Model RNN directly to the time-series data of the two regions to perform CURBD. From the Model RNN we obtained a matrix describing the directed interactions within and between each of the two regions.
(I) We applied CURBD to obtain the currents driving each neuron in Region A from other Region A neurons (top) and from region B (bottom).
(J) We performed PCA to identify the dominant component of each source current. The currents from Region A resembled the low-frequency sinusoid driving Region A, while the currents from Region B matched the higher-frequency sinusoid driving Region B.
(K) We computed R^2 values comparing the first PC of each source current to the two sinusoidal inputs.
(L) Reconstruction accuracy of B to A currents for different simulation parameter values. We explored three key simulation parameters: i) the amount of chaos (g parameter; see Methods) from overdamped (g<1) to strongly chaotic (g>1.5); ii) the strength of the external inputs driving the system from very weak (0.001) to very strong (1); iii) the sparsity of inter-region connections from very sparse (1%) to full-rank (100%). Each heatmap shows the strength of inter-region connections against the percent of neurons receiving inter-region connections, and heatmaps going left to right show increasing g_A. For low values of g_A corresponding to damped dynamics, the inputs can only be reconstructed with strong connectivity. When both regions have similar dynamics (g_A = 1.3) the currents cannot be accurately demixed. The optimal regime occurs when g_A and g_B have different dynamics, with a tradeoff between sparsity and strength of the inter-region connectivity.

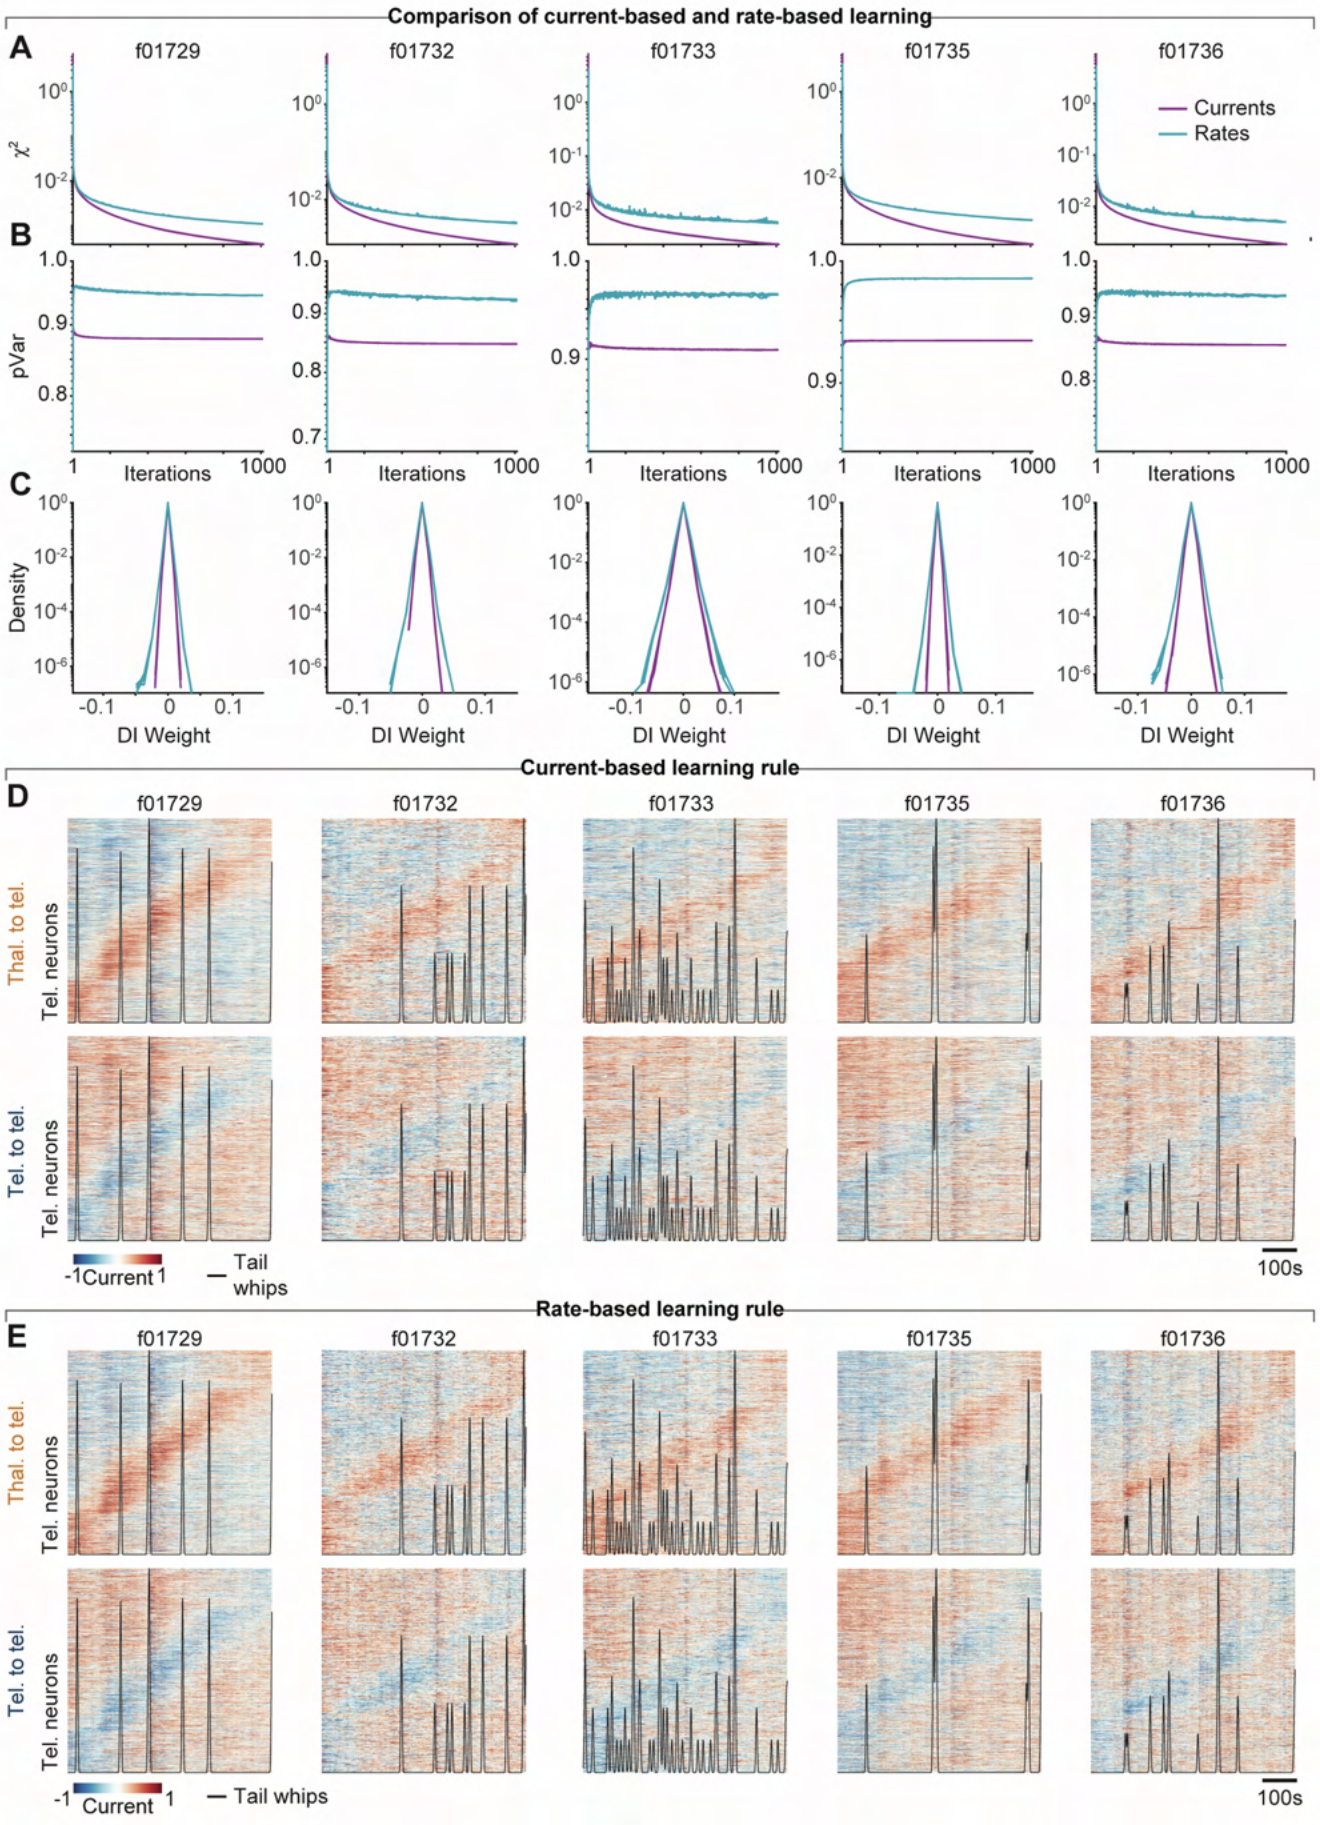
Comparison of current-based and rate-based learning
A
B
C
f01729
f01732
f01733
f01735
f01736
Currents
Rates
χ^2
pVar
Iterations
Density
DI Weight
Current-based learning rule
D
Thal. to tel.
Tel. neurons
Tel. to tel.
-1Current 1
Tail whips
100s
Rate-based learning rule
E

Figure S2. Supporting data for comparing current-based and rate-based Model RNNs in the larval zebrafish dataset, related to Figure 4.

(A) Model error over 1000 training iterations for the Model RNN for the five fish for current-based learning (magenta) and rate-based learning (cyan).

(B) Proportion of variance explained by the Model RNN for 1000 training iterations for the two learning rules.

(C) Final distributions of directed interaction weights for the trained **J** matrices using the two learning rules.

(D) CURBD on the inputs to telencephalon neurons from thalamus (top) and recurrent connections from telencephalon (bottom) using the current-based learning rule. Heatmaps show inhibitory (blue) and excitatory (red) driving each telencephalon neuron over time. Neurons are sorted by the time of peak activation of excitatory thalamic inputs, and the same sort is used for the telencephalon inputs. Black overlaid lines show tail whipping movements.

(E) Data presented as in Panel A for CURBD using the rate-based learning rule.

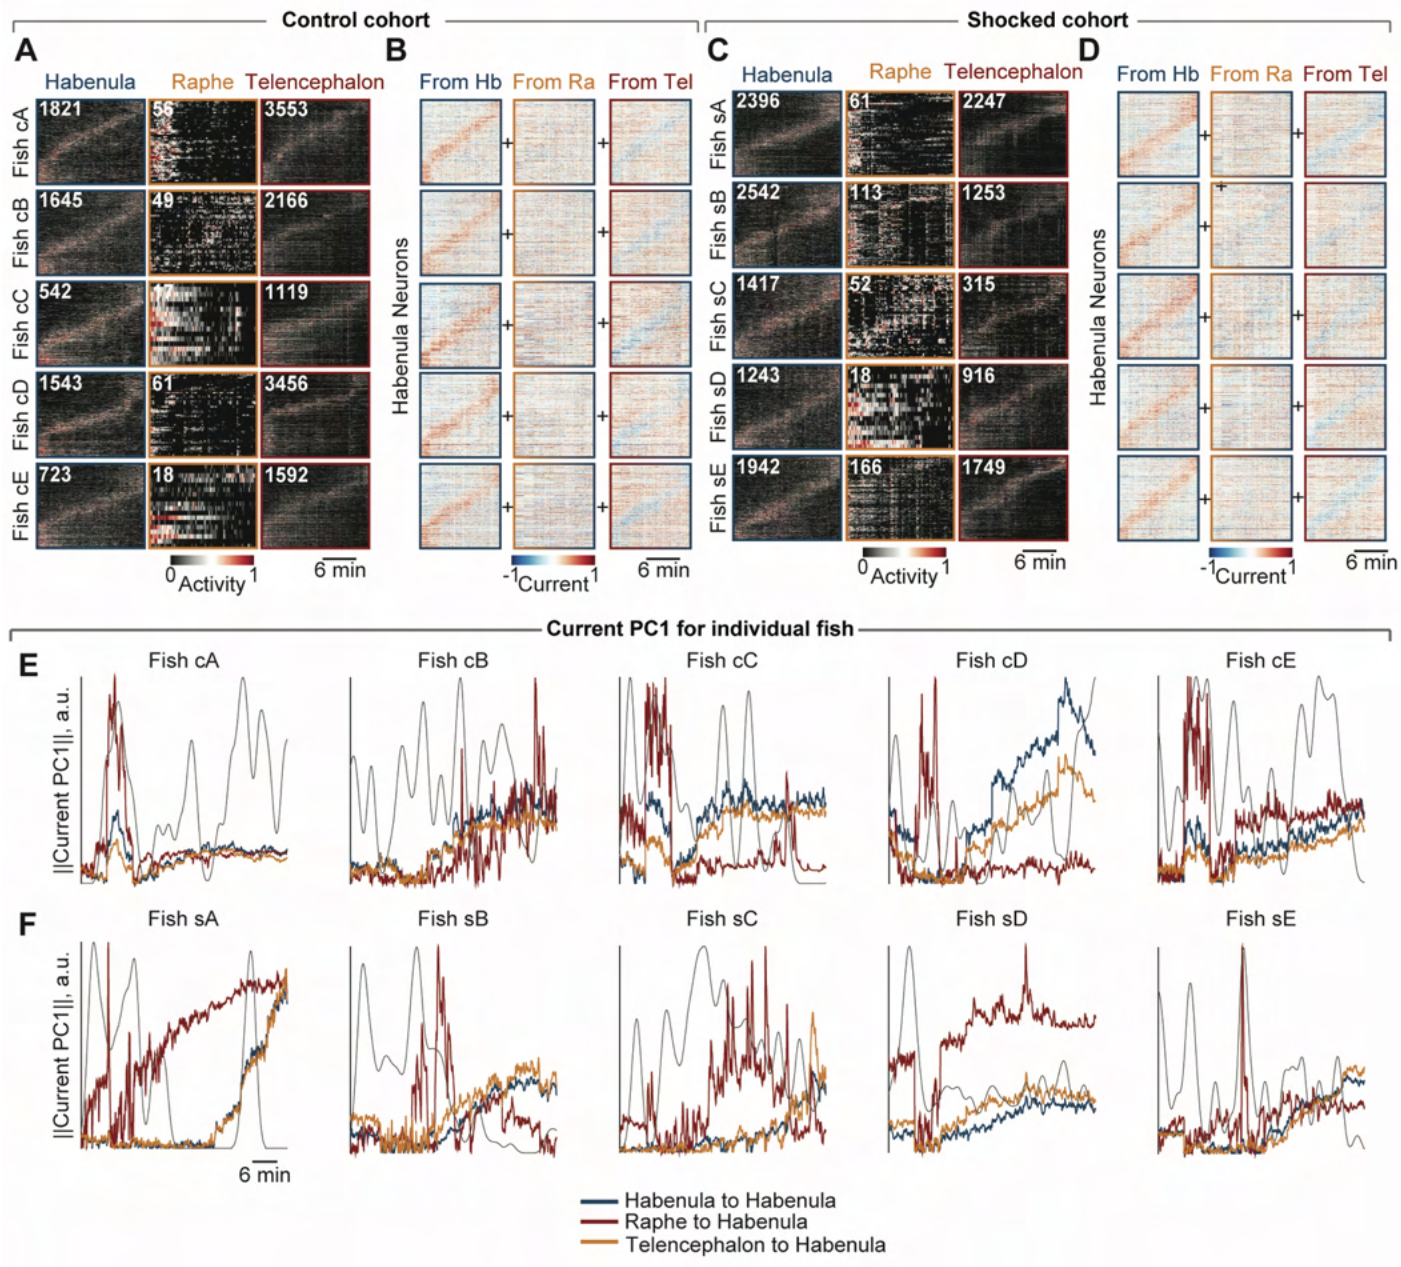

Figure S3. Supporting data for the three-region larval zebrafish dataset, related to Figure 5.

(A) Heatmaps of all recorded cells for the 3 regions in each individual fish of the control cohort. Data presented as in Figure 5B.

(B) Heatmaps of current inputs to habenular neurons for each fish of the control cohort. Data presented as in Figure 5H.

(C) Same as Panel A for the shocked cohort. Panel duplicated from Figure 5B.

(D) Same as Panel B for the individual fish of the shocked cohort. Top row (Fish sA) is reproduced from Figure 5H.

(E-F) The leading Principal Component across Habenular neurons capturing the dominant pattern of inputs to Habenula from each region. Plots show each individual fish from the control (Panel E) and shocked (Panel F) cohorts.

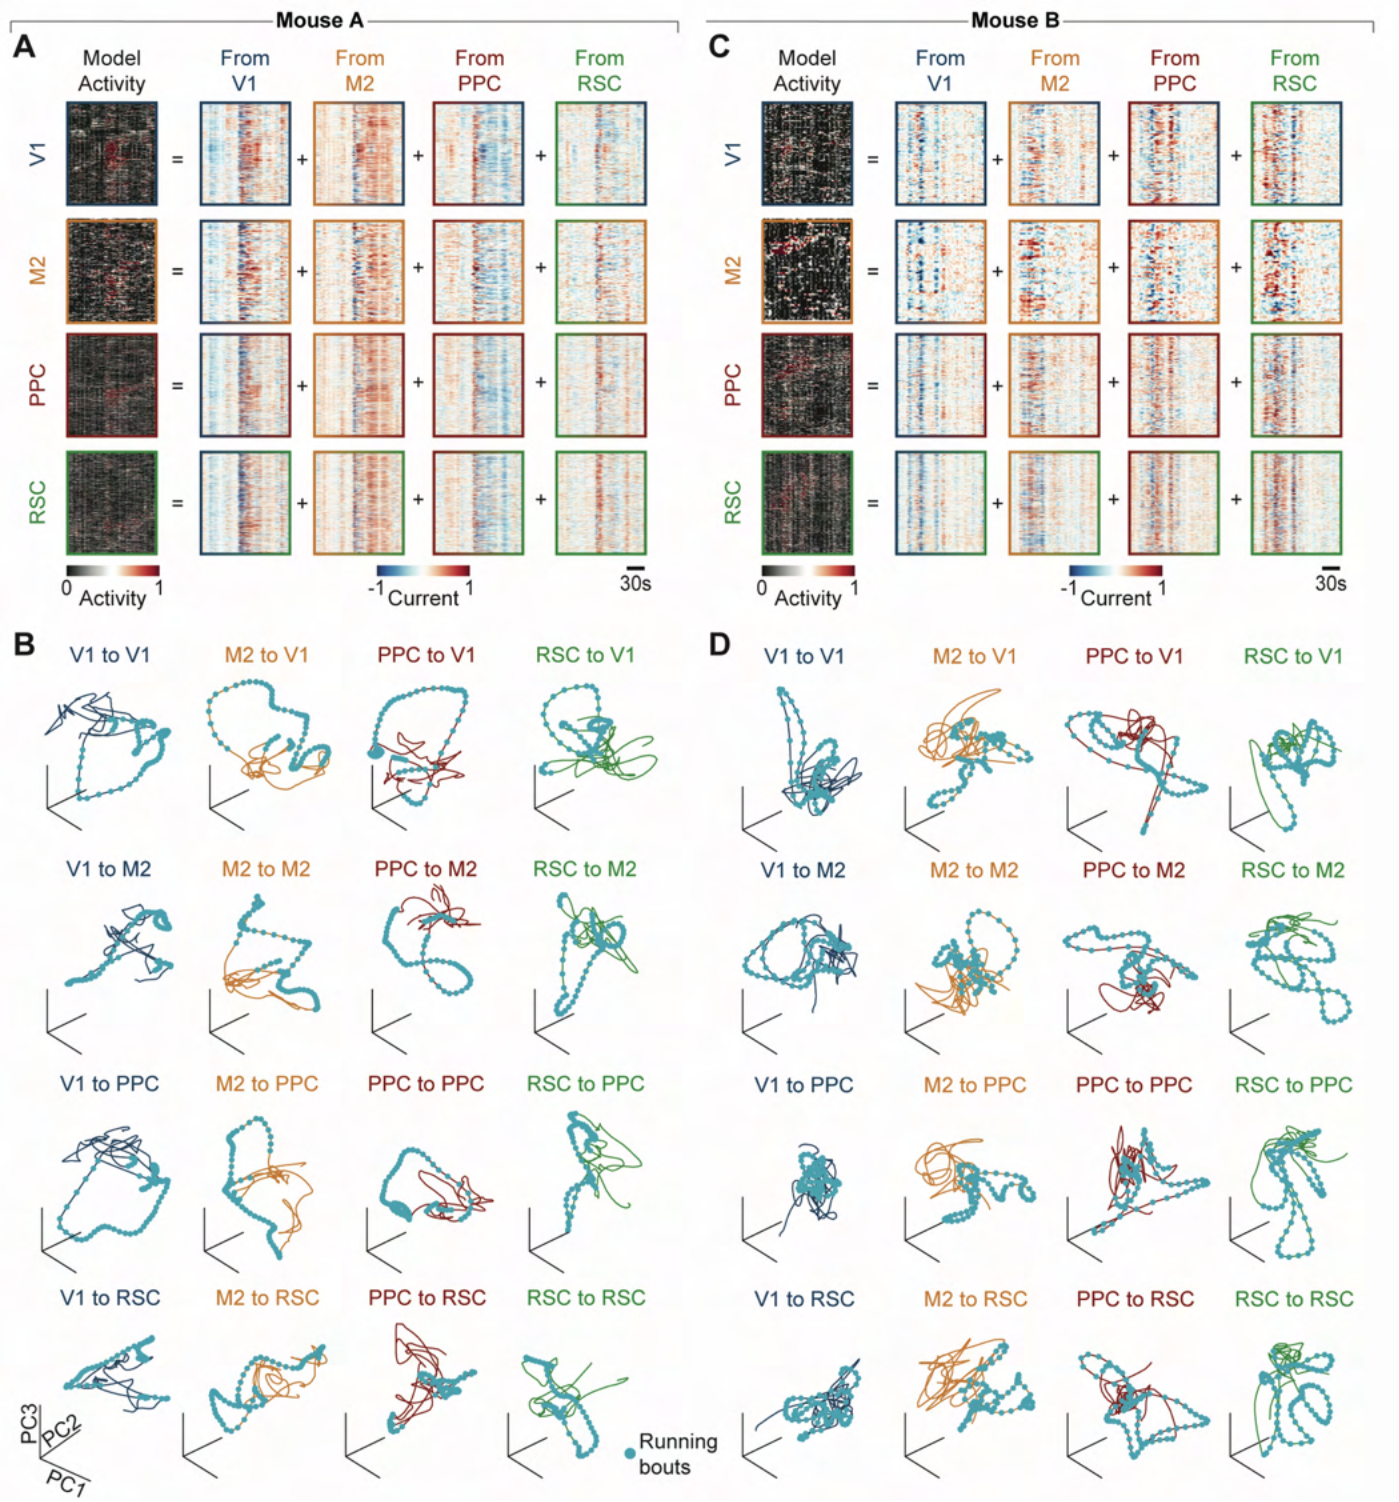

Figure S4. Supporting data for the multi-region mouse dataset, related to Figure 6.

(A) Model RNN activity and source current activity for Mouse A. Figures reproduced from Figure 6G.

(B) Current trajectories in the first three PCs for all sixteen source currents from Mouse A. The V1 source currents (top row) are reproduced from Figure 6H.

(C) Data presented as in Panel A for Mouse B.

(D) Data presented as in Panel B for Mouse B.

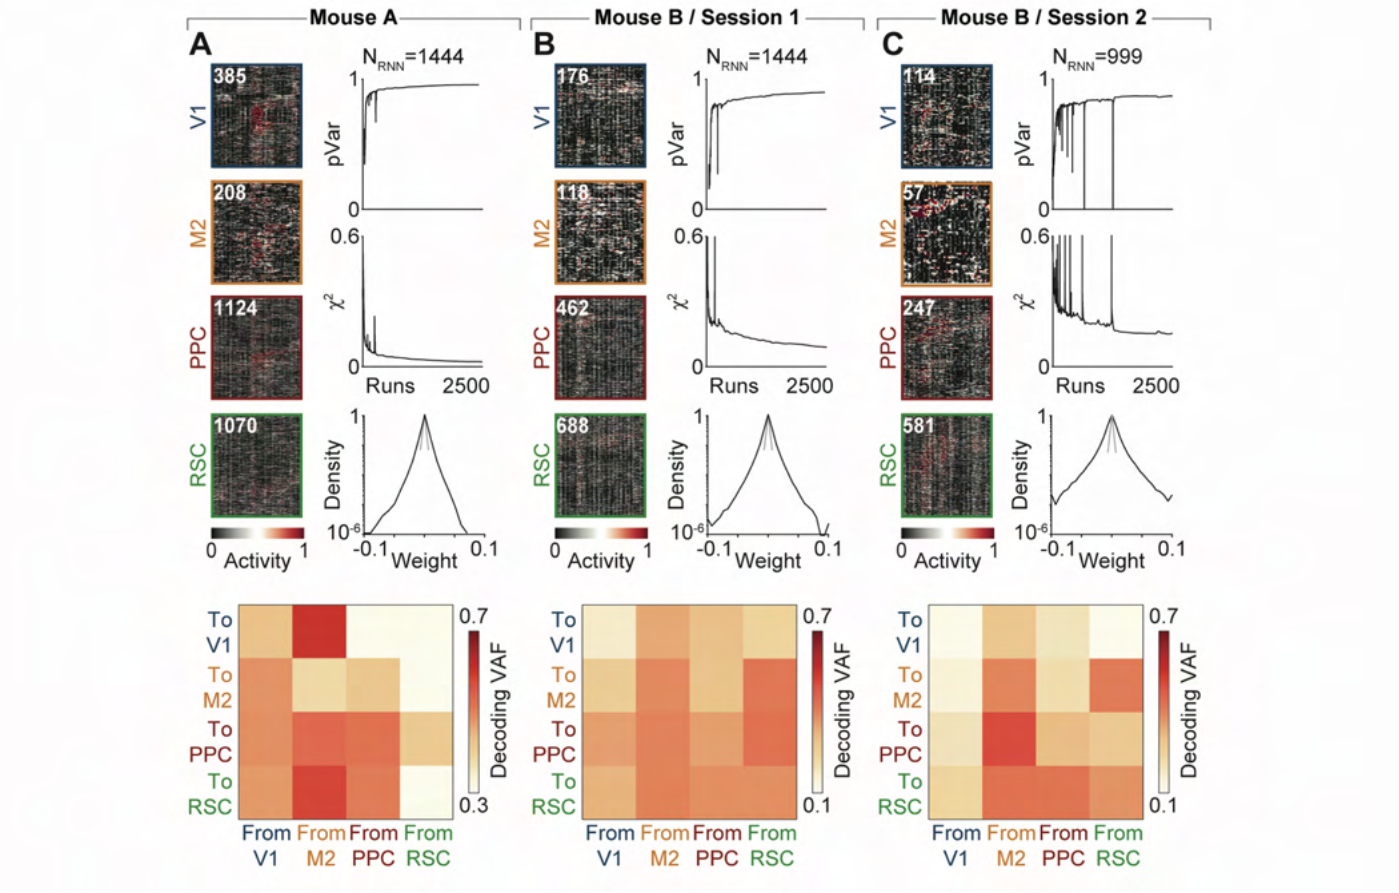

Figure S5. All recording sessions for the mouse dataset, related to Figure 6.

(A) (Top) Model RNN output (left) and training performance (right) for the session from Mouse A. (Bottom) Decoding performance for the sixteen source currents. All data are reproduced from Figure 6.

(B) (Top) Model RNN output (left) and training performance (right) for Session 1 from Mouse B. (Bottom) Decoding performance for the sixteen source currents. Portions are reproduced from Figure 6.

(C) (Top) Model RNN output (left) and training performance (right) for Session 2 from Mouse B. (Bottom) Decoding performance for the sixteen source currents.

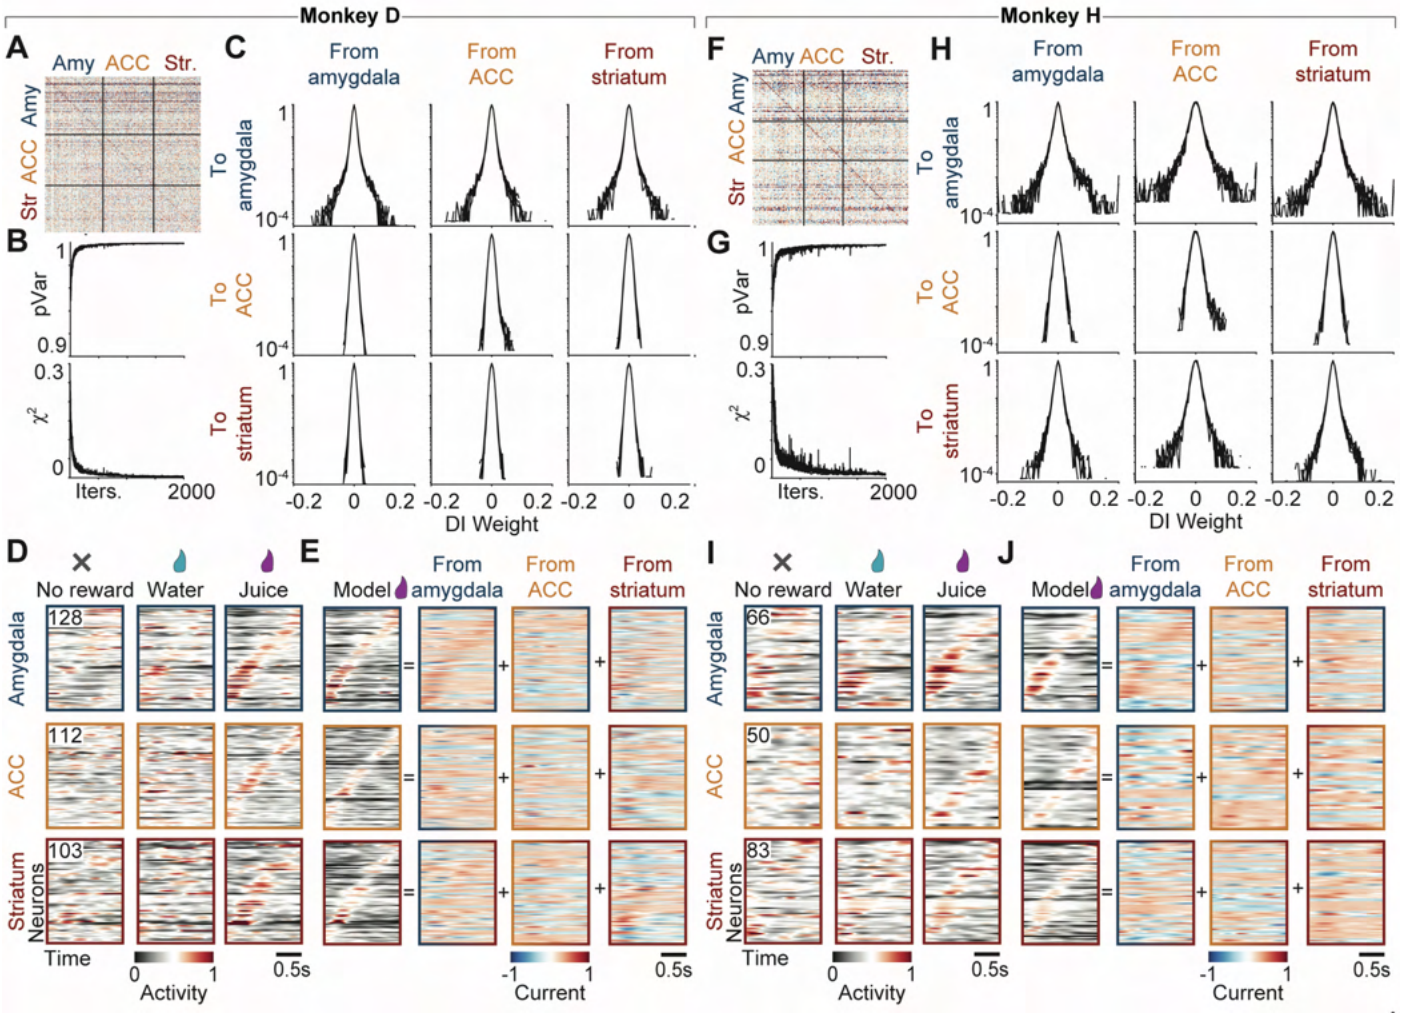

Figure S6. Supporting data for the multi-region macaque electrophysiology dataset, related to Figure 7.

(A) Connectivity matrix for an example Model RNN fit to data from Monkey D. Neurons are ordered by region, starting with amygdala (Amy, blue), subcallosal anterior cingulate cortex (ACC, yellow), and rostromedial striatum (Str, red).

(B) (Top) Proportion of variance explained (pVar) in the neural population as a function of training runs. Training results for five different random initializations are plotted to highlight consistency. (Bottom) Model error (χ^2) for the five initializations shown above.

(C) Distribution of weights in each submatrix used for CURBD. Each column corresponds to a source region, and each row to a target region. All five initializations are plotted to illustrate consistency.

(D) Trial-averaged firing rates for the amygdala (top), subcallosal ACC (middle), and striatum (bottom) comprising the pseudopopulation dataset for Monkey D. Left plot shows data from the unconditioned stimulus (left), water stimulus (middle), and juice stimulus (right). All trials are aligned on presentation of the stimulus. Neurons in each region are sorted according to their time of peak activity in the Juice condition.

(E) CURBD decomposition of activity in each region for the Juice trials. Left plots show the full Model RNN activity. The remaining plots show the inferred source currents to each target region (rows) from all source regions (columns).

(F-J) Data for Monkey H presented as in Panels A-E.

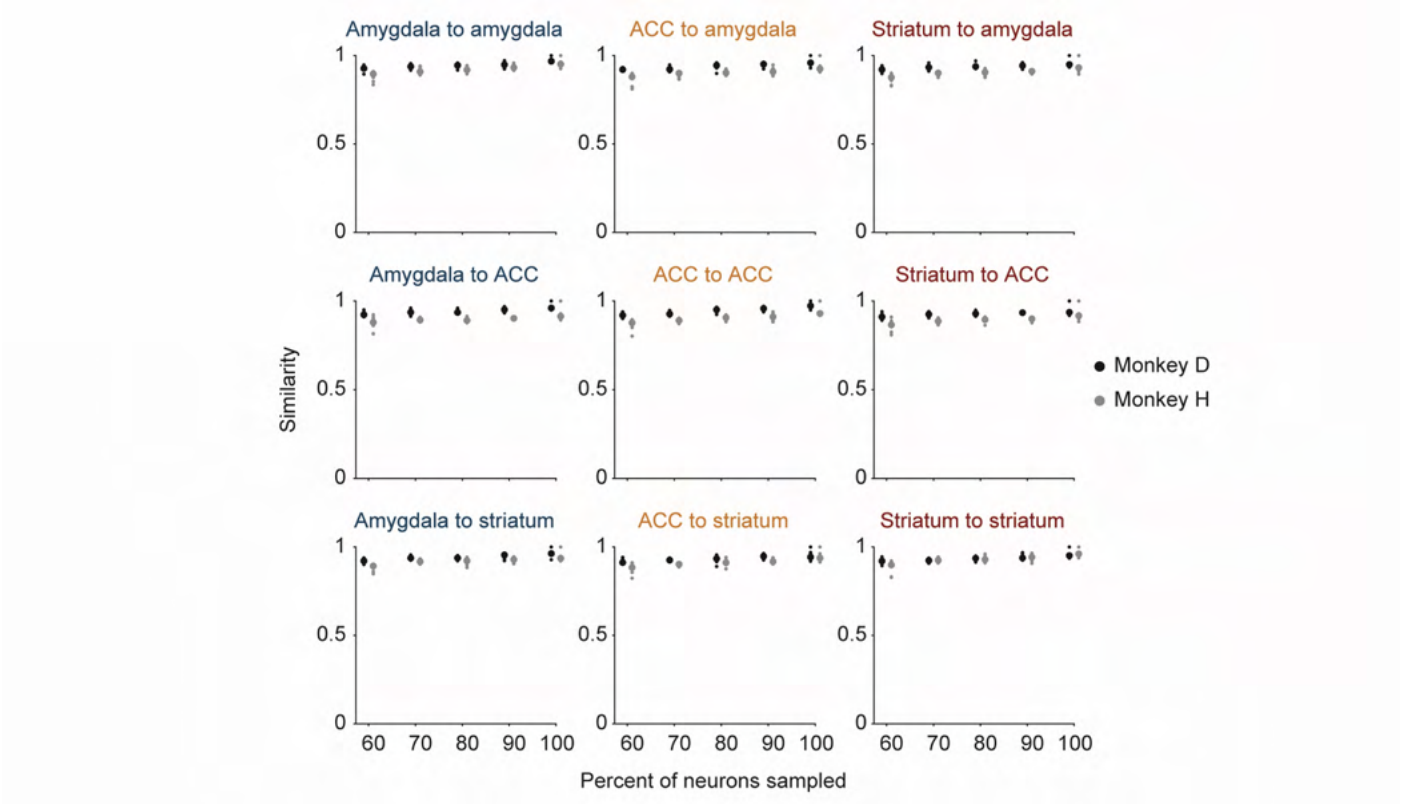

Figure S7. Consistent identification of current dynamics with random subsamples of recorded neurons, related to Figure 7.

Mean canonical correlation in a 20-dimensional space identified by PCA for each source current in Monkey D (black) and Monkey H (gray). Small dots indicate the results of 10 random subsamples of the total neural population at each percentage level. Large circles indicate the median across iterations.

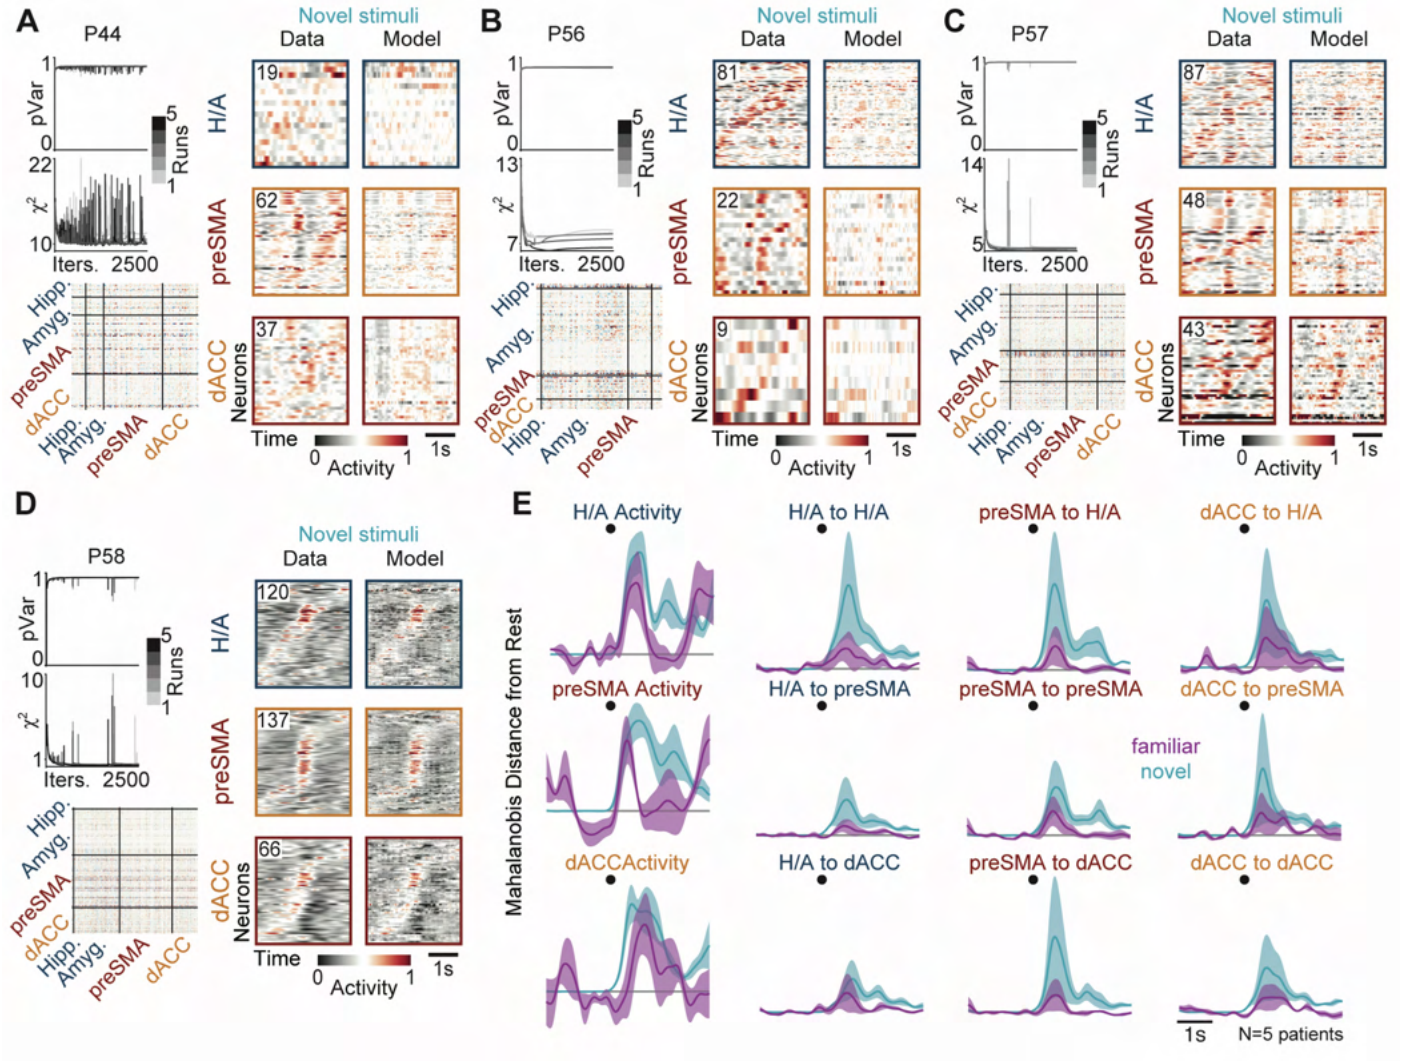

Figure S8. Supporting data for the multi-region human electrophysiology dataset, related to Figure 8.

(A) Model RNN summary for P44. (Top left) Model RNN training performance (pVar and χ^2) for five runs starting from different random initializations of the J matrix. (Bottom left) Example J matrix for one run. (Right) Neural activity from recorded neurons (Data) and the Model RNN units.

(B) Data presented as in Panel A for P56.

(C) Data presented as in Panel A for P57.

(D) Data presented as in Panel A for P58.

(E) Mahalanobis distance from rest for all sixteen source currents on the familiar stimuli trials (magenta) and novel stimuli trials (cyan). Lines show mean and standard error across all five participants. Black dot indicates time of stimulus presentation. Top row is reproduced from Figure 8H.

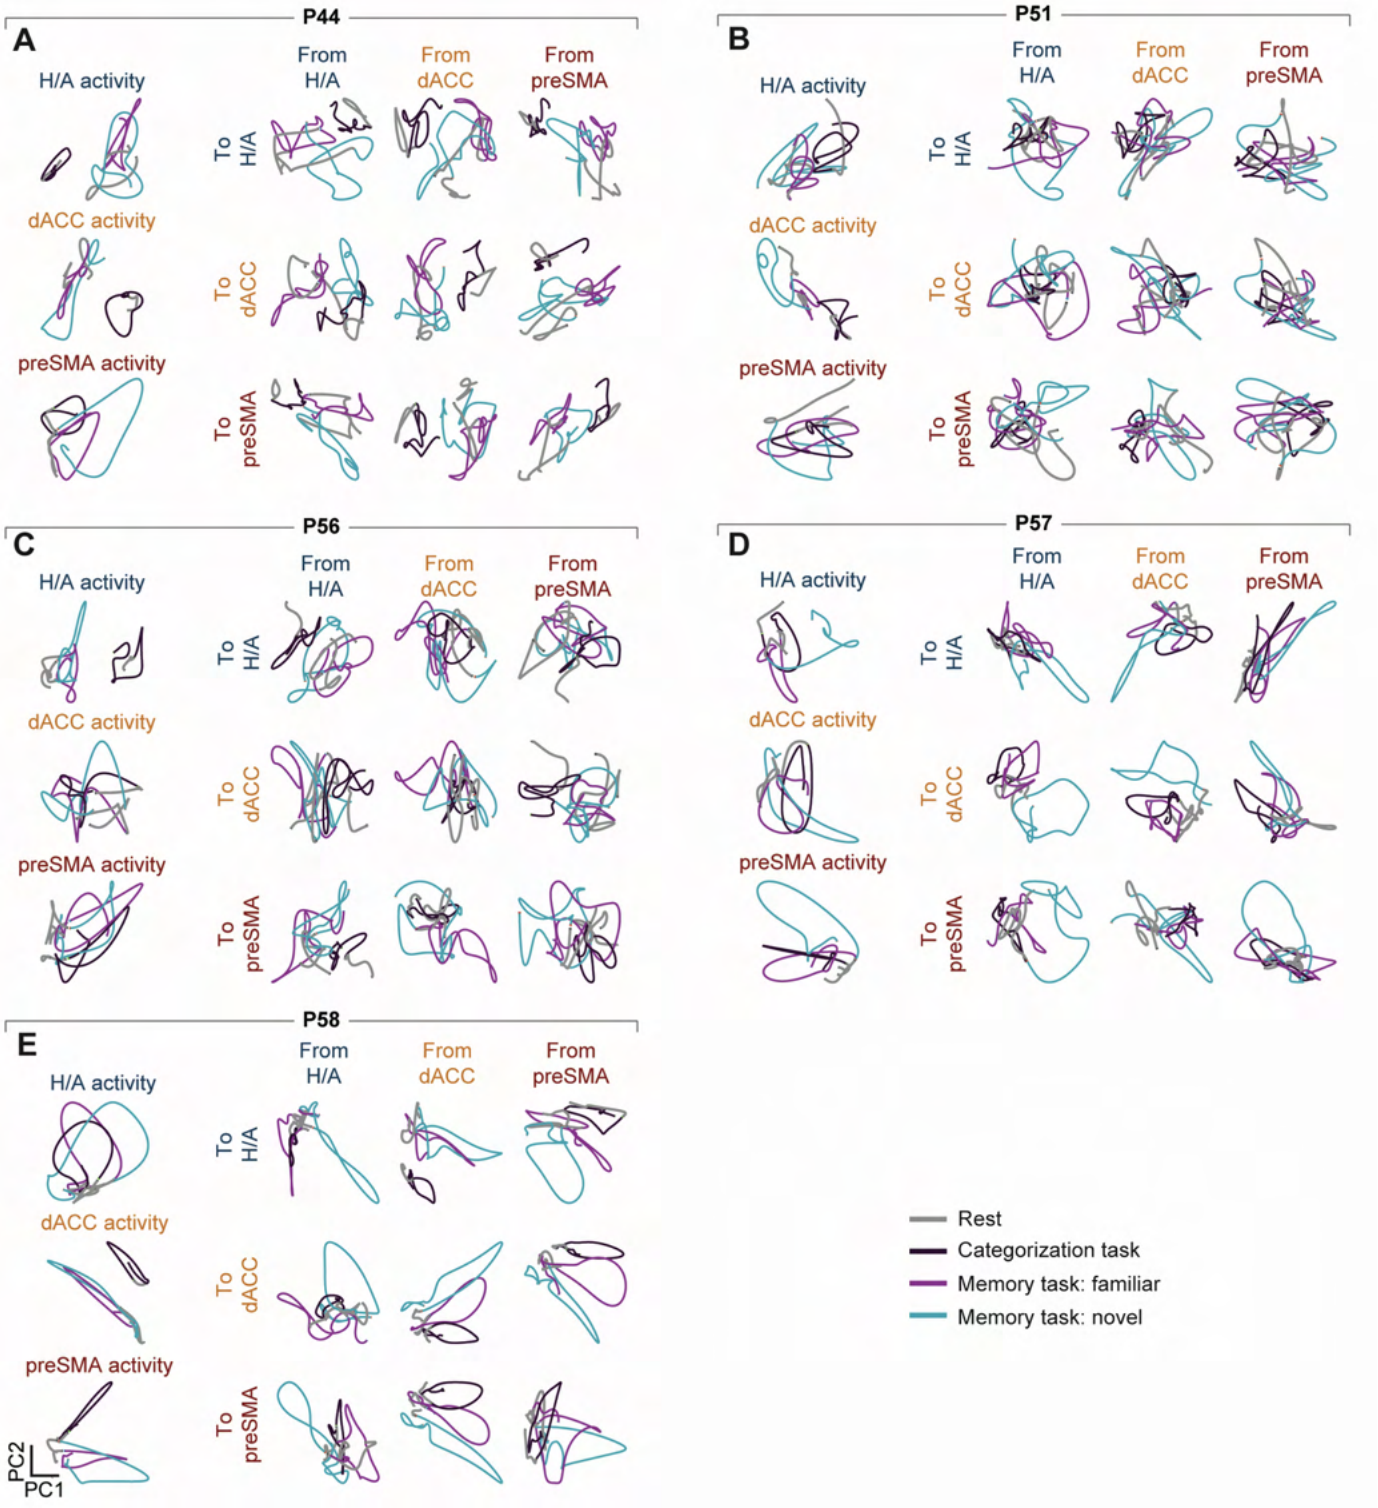

Figure S9. Comparison of current dynamics across tasks for the multi-region human electrophysiology dataset, related to Figure 8.

(A) Neural and current dynamics for both tasks in P44. Each subplot shows the first two PCs of the full population activity of the three regions as well as the nine source currents during the categorization task (purple) and novel (cyan) and familiar (magenta) stimuli during the memory task. Gray shows activity at rest before the stimuli.

(B-E) Neural and current dynamics for both tasks in P51 (Panel B), P56 (Panel C), P57 (Panel D), P58 (Panel E).

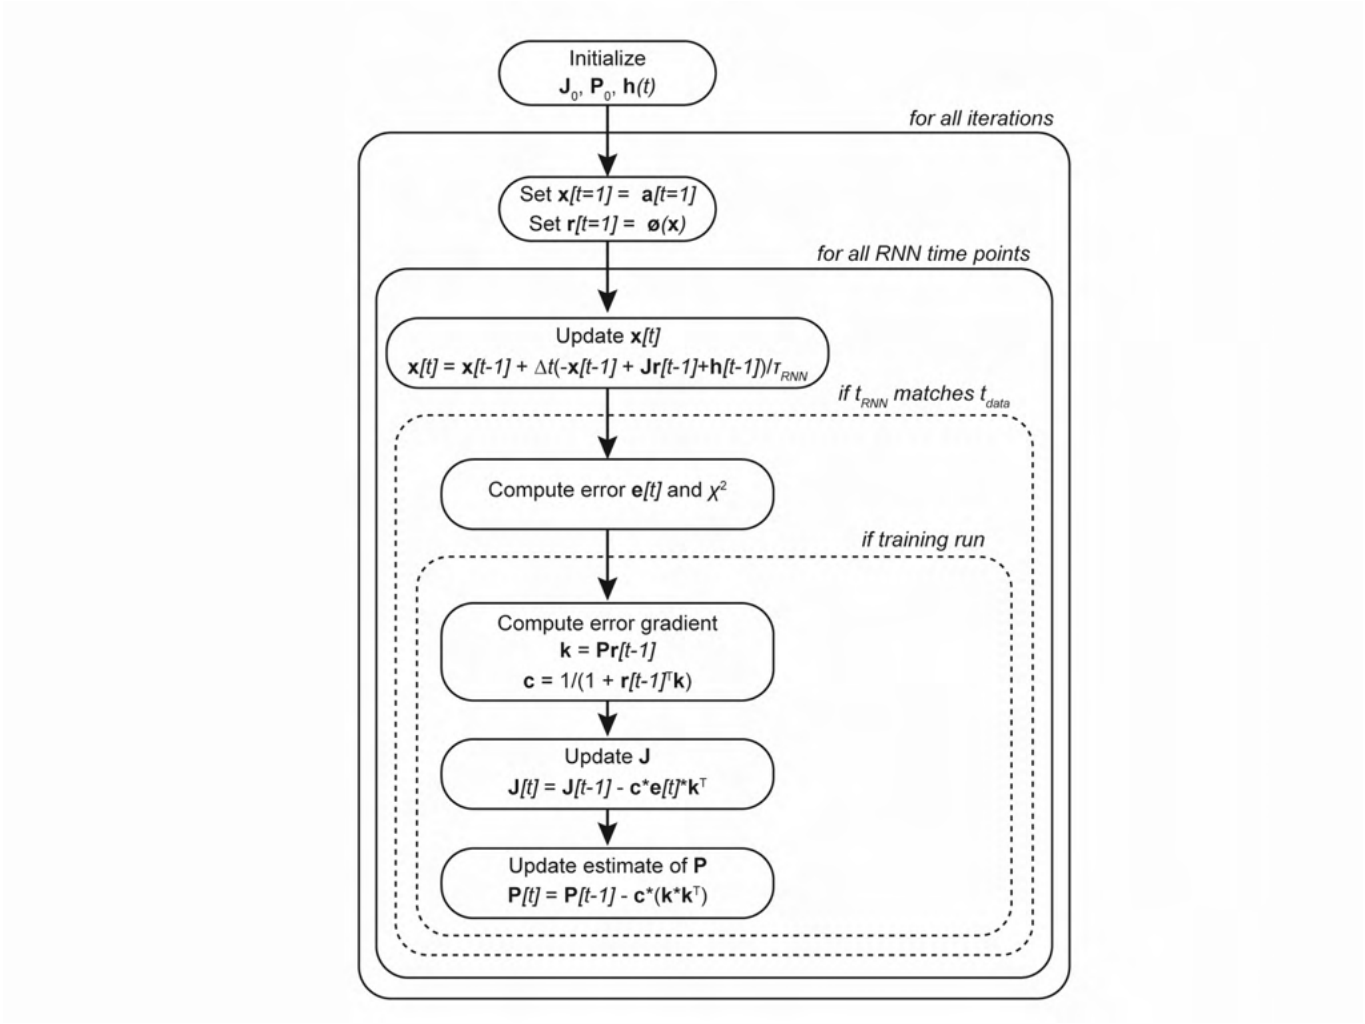

Figure S10. Flowchart of algorithmic implementation of learning rules, related to STAR Methods.
This flowchart indicates the algorithm to iteratively estimate the **P** matrix used during learning.

Table S1. Model RNN training parameters for all datasets, related to STAR Methods

Parameter	Description	2-region simulation	3-region simulation	Zebrafish dataset	Mouse dataset	Monkey dataset	Human dataset
g	Chaotic spontaneous activity	1.5	1.5	1.3	1.5	1.5	1.5
τ	Time constant for model units	0.1	0.1	2.5	0.3	0.001	0.0075
P_0	Learning rate	1.0	1.0	1.0	1.0	1.0	1.0
τ_{WN}	Time constant of filtered white noise	0.1	0.1	0.1	0.1	0.1	0.1
w_{WN}	White noise input weight	0.01	0.01	0.01	0.001	0.0001	0.001
$n_{iterations}$	Number of training iterations	100	500	1000	2500	1500	2500
Δt	Data time step size (s)	0.01	0.01	0.92 - 0.97	0.1866	0.01	0.01
Δt_{RNN}	Model RNN step size (s)	0.001	0.001	0.23 - 0.24	0.0467	0.0005	0.001

Table S2. Two-region generator model parameters, related to STAR Methods.

Parameter	Description	Value(s)
g_A	Region A chaos	[0.9, 1.0, 1.1, 1.3, 1.5, 1.8, 2.5]
g_B	Region B chaos	1.5
τ	Time constant of model units	0.1
T	Simulation time	10

Parameter	Description	Value(s)
w_{rgn}	Inter-region weight	[0.001, 0.01, 0.05, 0.1, 0.25, 0.5, 1.0]
p_{rgn}	Inter-region connection proportion	[0.01, 0.05, 0.1, 0.25, 0.5, 1.0]
w_{in}	External input weight	1.0
Δt	Simulation time step size	0.01

Table S3. Three-region generator model parameters, related to STAR Methods

Parameter	Description	Value
g_A	Region A chaos	1.8
g_B	Region B chaos	1.5
g_C	Region C chaos	1.5
τ_{true}	True decay constant	0.1
T	Simulation time	12

Parameter	Description	Value
w_{rgn}	Inter-region connection weight	0.01
p_{rgn}	Fraction of inter-region connections	0.01
w_{in}	External input weight	1.0
σ	Width of sequential and fixed point-bumps (number of units)	200
Δt	Simulation time step	0.01

Table S4. Simultaneous neuron yield for the selected from the larval zebrafish dataset, related to STAR Methods

Brain Region	Fish cA	Fish cB	Fish cC	Fish cD	Fish cE	Fish sA	Fish sB	Fish sC	Fish sD	Fish sE
Telenc.	3553	2166	1119	3456	1592	2247	1253	315	916	1749
Thalamus	819	771	459	847	579	*N/A*	*N/A*	*N/A*	*N/A*	*N/A*
Habenula	1821	1645	542	1543	723	2396	2542	1417	1243	1942
Raphe	56	49	17	61	18	61	113	52	18	166
Total	6249	4631	2137	5907	2912	4704	3908	1784	2177	3857

Table S5. Simultaneous neuron yield for the mouse dataset, related to STAR Methods

Brain Region	Mouse A	Mouse B	
		Session 1	*Session 2*
V1	385	114	176
M2	208	57	118
PPC	1124	247	462
RSC	1070	581	688
Total	2787	999	1444

Table S6. Pseudopopulation sizes for the monkey dataset, related to STAR Methods

Brain Region	Monkey D	Monkey H
Amygdala	128	66
Subcallosal ACC	112	50
Striatum	103	83
Total	343	199

Table S7. Pseudopopulation sizes for the human dataset, reported as: Left Hemisphere / Right Hemisphere (Total), related to STAR Methods

Brain Region	P44 (two sessions)	P51 (five sessions)	P56 (three sessions)	P57 (three sessions)	P58 (three sessions)
Hippocampus	0 / 15 (15)	25 / 38 (63)	5 / 0 (5)	17 / 0 (17)	5 / 0 (5)
Amygdala	3 / 1 (4)	2 / 22 (24)	33 / 43 (76)	36 / 34 (70)	54 / 61 (115)
preSMA	62 / 0 (62)	16 / 7 (23)	8 / 14 (22)	48 / 0 (48)	65 / 72 (137)
dACC	37 / 0 (37)	72 / 17 (89)	0 / 9 (9)	43 / 0 (43)	7 / 59 (66)
Total	102 / 16 (118)	115 / 84 (199)	46 / 66 (112)	144 / 34 (178)	131 / 192 (323)