Connectome as Compiler: Reading Computation from Connectivity and then Writing it Back

Sovesh Mohapatra[1] Dani S. Bassett[1,2,3,4,*]

[1]Department of Bioengineering, School of Engineering and Applied Science, University of Pennsylvania, PA, USA

[2]Departments of Electrical & Systems Engineering and Physics & Astronomy, University of Pennsylvania, PA, USA

[3]Departments of Psychology and Biomedical Engineering, Yale University, CT, USA

[4]Wu Tsai Institute, Yale University, CT, USA

[*]Corresponding author: **dani.bassett@yale.edu**

Abstract

Neural systems use information from earlier inputs to shape their current responses. How recurrent connectivity supports the retention and combination of this information remains difficult to quantify. Here, we introduce **Co**nnectome **R**esponse **A**nalysis and **D**esign (CoRAD), an analytic framework that derives order-specific responses from a connectome and decompiles them into dimensions computationally associated with linear memory and nonlinear mixing. Each dimension is an effective count of the response directions generated at its corresponding order, providing a compact summary of how past inputs shape network activity. CoRAD uses their gradients to guide constrained changes to existing connections. Across a range of networks—from synthetic reservoirs, to human and non-human connectomes—the response dimensions predicted their corresponding simulated memory capacities in held-out data. Their decompositions showed how response contributions were distributed across input histories and cortical regions. Gradient-guided reweighting changed the corresponding held-out capacities and moved one or both dimensions toward specified targets under wiring constraints. Most of the gain found by the constrained searches was recovered using only a small fraction of existing connections, while different reweightings reached similar joint targets. Together, CoRAD provides a quantitative link between recurrent connectivity and the explanation, interpretation, and constrained design of predicted temporal-processing capacities.

1 Introduction

Neural circuits transform streams of input into evolving patterns of activity that support sensation, action, and cognition[1–4]. These transformations unfold through recurrent interactions among neurons and brain regions. More specifically, anatomical connectivity constrains how activity propagates, persists, and coordinates across a system[5–7]. The relationship between anatomical connectivity and function is not one-to-one[8;9]. The same anatomical circuit can express different activity patterns when the signals entering it or its neuromodulatory state change[4;10;11], while markedly different combinations of cellular and synaptic parameters can generate similar network activity[12–14]. A connectome therefore provides essential information about circuit architecture but does not specify the dynamics or functions that the circuit will express[9;15]. A central challenge is to determine which functional capacities are constrained by recurrent connectivity, and how those constraints depend on the system's dynamics, inputs, and operating state[16–18].

One functional capability of particular interest in the field is temporal information processing: a circuit's current state can both *retain* information about earlier inputs and *combine* signals that arrived at different times[19–21]. Reservoir computing provides a controlled framework for studying these two aspects of temporal processing[18;22]: retention and combination. A reservoir transforms a time-dependent input into an evolving state through fixed recurrent dynamics; a trained linear readout then tests which functions of the input history can be recovered from the current state[18;22;23]. Linear memory capacity (a form of *retention*) measures how accurately the readout reconstructs individual past inputs at successive delays, capturing the fading influence of progressively older inputs[21;24]. Here, adjective "linear" in the phrase *linear memory* refers to the target and readout, not necessarily to the recurrent dynamics[21;25]. Linear memory is important because tasks with temporal dependencies require recent inputs to remain accessible, and its capacity provides a common benchmark for comparing recurrent architectures[20;24;26]. We note that linear memory capacity is a property of the specified reservoir model, not a measure of a participant's memory as experienced by themselves or as measured by neuropsychological tests; yet, prior work has shown that simulated memory capacity can statistically track cognitive performance in humans[27]. To complement linear memory as a measure of retention, we also consider nonlinear mixing as a measure of combination. Nonlinear mixing asks whether the same readout can recover nonlinear functions of the input history, such

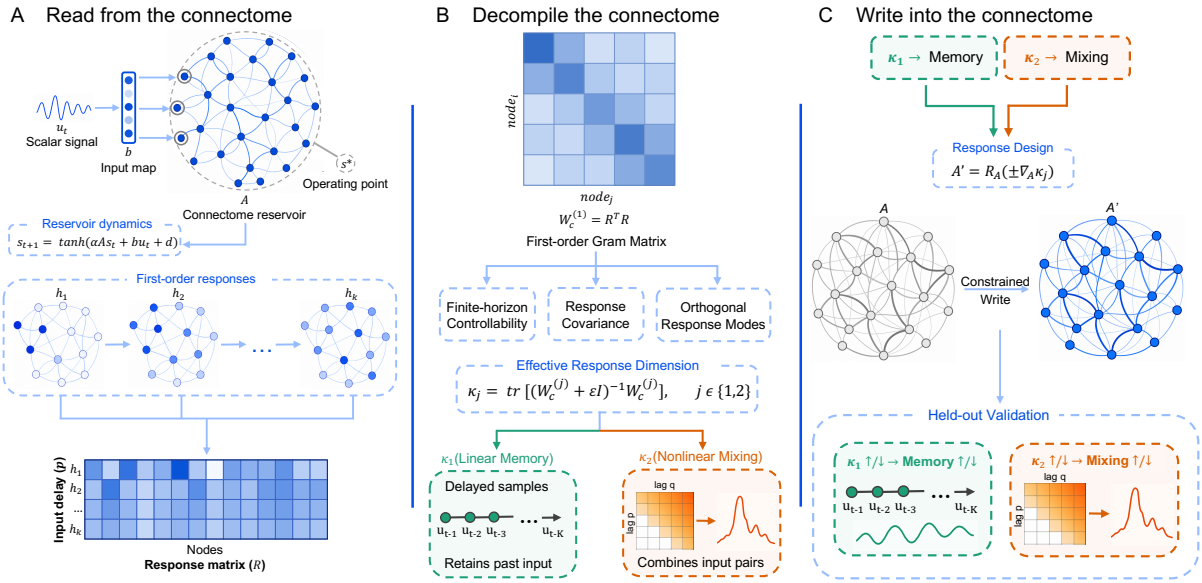

Figure 1: **Connectome Response Analysis and Design (CoRAD).** CoRAD reads input-evoked responses from a connectome, decompiles their geometry into dimensions associated with simulated linear-memory and nonlinear-mixing, and uses those dimensions to guide constrained writes into the wiring. **A. Read from the connectome.** A time-varying signal enters a connectome reservoir through an input map. Linearizing the dynamics about the operating point gives the task-aligned first-order responses $h_p = F^{p-1}g$, for $p = 1, \ldots, K$, stacked by delay as the rows of the response matrix $R \equiv R^{(1)}$. **B. Decompile the connectome.** The node-space Gram matrix $W_c^{(1)} = R^\top R$ describes the geometry of these responses. Under zero-mean temporally white unit-variance input, it is both the finite-horizon controllability Gramian of (F, g) and the covariance of the truncated first-order propagated response, and its eigenvectors define orthogonal response modes. Applying the same effective-dimension summary to $W_c^{(1)}$ and to the corresponding second-order matrix $W_c^{(2)} = R^{(2)\top} R^{(2)}$ gives κ_1 and κ_2, associated with linear memory and nonlinear mixing. **C. Write into the connectome.** The operator $\mathcal{R}_A$ maps positive or negative gradients of either response dimension to a constrained write A'. Held-out simulations test whether designed increases or decreases in κ_1 and κ_2 produce corresponding changes in memory and mixing.

as products of inputs presented at different delays[21;28]. This capability is important when an output depends on interactions among earlier inputs: recovery by a linear readout shows that the reservoir has both retained and nonlinearly transformed those inputs[21;28;29]. With these two tools in hand, we seek to address the question of how recurrent connectivity shapes these functional capacities under fixed dynamical and input conditions[16–18].

Several approaches have linked network structure to dynamical function. Connectome-based reservoir studies use empirical connectomes to constrain recurrent architecture and quantify memory or task performance by simulating reservoir states and fitting readouts[16–18]. Learning studies further link modular organization to task demands: the modular organization of neural activity is associated with task performance[30], and multitask and incremental training can in-

crease the modularity of recurrent connectivity[31]. Analytic approaches characterize recurrent networks directly from their governing equations. Reservoir programming represents states and dynamics in a basis of their inputs, enabling functions to be recovered from trained reservoirs and programmed through output and feedback connections[32]. For linear memory, reconstruction scores across delays sum to total memory capacity[21;24]; for linear recurrent networks with independent inputs, this theoretical total equals the rank of the associated controllability matrix[25;33]. When the controllability matrix is full rank, the total equals the number of units and is therefore identical for all controllable systems of the same size. More detailed response geometry is retained by the controllability Gramian, which quantifies the minimum input energy required for specified state transitions[34;35]. Building on this framework, a computational affordance landscape represents the distribution of these costs across possible goal-directed state transitions[36]. Edge-centric control has also combined structural controllability with functional coupling to identify connection-level perturbations that reduce the energy required for targeted synchronization[37]. These control approaches quantify the energetic accessibility of activity states or synchronization patterns, whereas reservoir memory and mixing concern which functions of the input history can be recovered by a linear readout from the present state[21]. A question still remains whether order-specific response geometry can both predict temporal-processing capacities in unseen networks and provide directions for changing them under structural constraints.

Here, we introduce **Co**nnectome **R**esponse **A**nalysis and **D**esign (CoRAD), an analytic framework for relating recurrent connectivity to linear memory and nonlinear mixing. CoRAD begins with a local response expansion about a specified operating point (Fig. 1A) and defines two separately constructed response dimensions: κ_1 is a ridge-smoothed effective count of first-order response directions, whereas κ_2 is the corresponding count of second-order directions generated by delayed-input products (Fig. 1B). These quantities describe a connectome within specified reservoir dynamics, input map, operating point, response horizon, ridge scale, and spectral normalization; they are not properties of connectivity alone. We test whether κ_1 and κ_2 predict simulated linear-memory and nonlinear-mixing capacities for held-out synthetic reservoirs, human connectomes, and input maps within non-human connectomes. We then examine how these dimensions and capacities change with operating state. Finally, we use their gradients to reweight existing connections while preserving tract-length-weighted wiring cost and spectral radius (Fig. 1C). We test whether these reweightings can increase or decrease either dimension, move it toward specified targets, and produce corresponding changes in held-

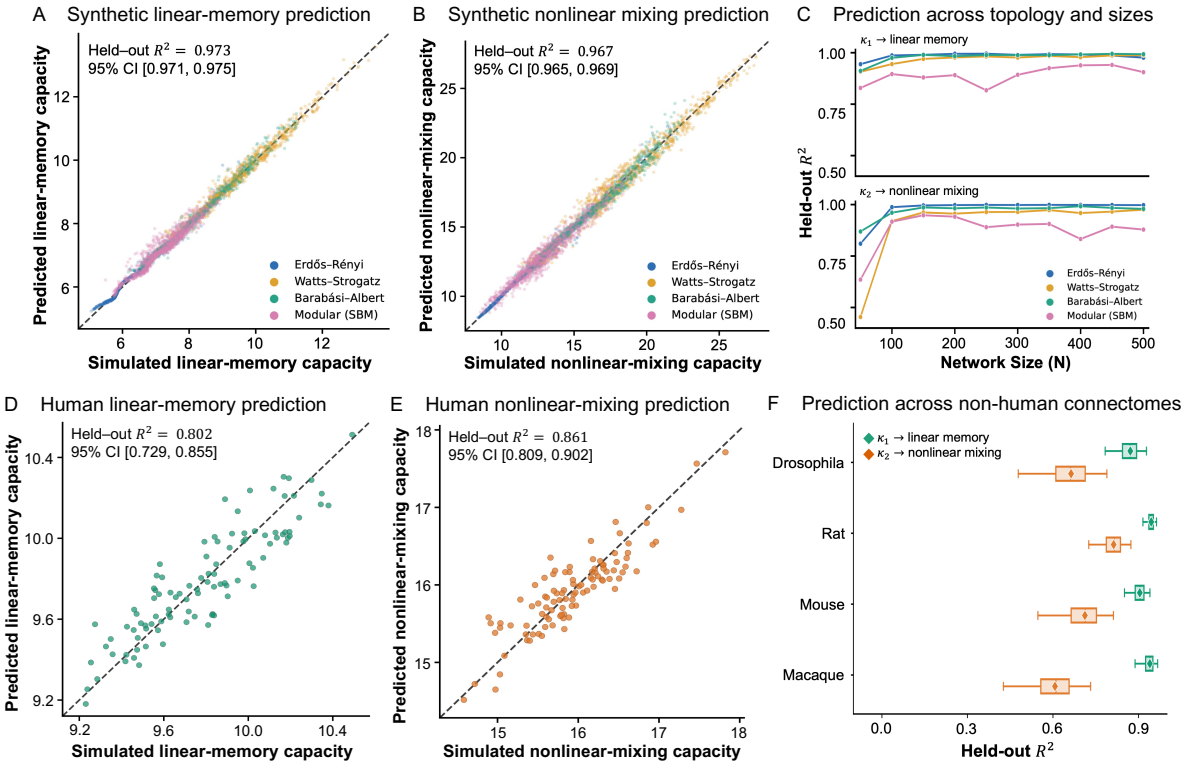

Figure 2: **Response dimensions predict simulated memory and mixing capacities.** First- and second-order response dimensions predict their matched simulated capacities in held-out synthetic reservoirs, human connectomes, and input maps within non-human connectomes. **A. Synthetic linear-memory prediction.** κ_1 predicts simulated linear-memory capacity across 4,000 synthetic reservoirs (held-out $R^2 = 0.973$). **B. Synthetic nonlinear mixing prediction.** κ_2 predicts simulated nonlinear-mixing capacity in the same reservoirs (held-out $R^2 = 0.967$). **C. Prediction across topology and sizes.** Held-out R^2 exceeds zero in all 40 topology-by-size groups for both capacities. **D. Human linear-memory prediction.** κ_1 predicts simulated linear-memory capacity across 100 HCP-YA connectomes (held-out $R^2 = 0.802$). **E. Human nonlinear-mixing prediction.** κ_2 predicts simulated nonlinear-mixing capacity across the same connectomes (held-out $R^2 = 0.861$). **F. Prediction across non-human connectomes.** Within one fixed connectome from each of four species, both dimensions predict their matched capacities across 60 Gaussian input maps that distribute the same temporal input across all nodes. Diamonds show observed held-out R^2, boxes show the interquartile range of 4,000 input-map bootstrap estimates, and whiskers show 95% intervals.

out simulations of memory and mixing. Together, these analyses link response geometry to the prediction and controlled modification of simulated linear-memory and nonlinear-mixing capacities within a specified recurrent-network model.

2 Results

2.1 Response dimensions predict simulated memory and mixing capacities

We first asked whether response dimensions, κ_1 and κ_2, computed without running the task simulations, could predict their matched capacities under fixed reservoir conditions. We generated 4,000 reservoirs spanning Erdős–Rényi, Watts–Strogatz, Barabási–Albert, and modular stochastic block-model topologies, with ten network sizes from 50 to 500 nodes. Linear-memory capacity summed the held-out squared correlations with which linear readouts reconstructed earlier inputs, whereas nonlinear-mixing capacity used the same score for products of inputs at pairs of delays. Within each topology and size, we fitted linear regression models to predict each capacity from its matched response dimension using training reservoirs and evaluated them on held-out reservoirs. κ_1 predicted linear-memory capacity with a held-out $R^2 = 0.973$ (95% CI $[0.971, 0.975]$), and κ_2 predicted nonlinear-mixing capacity with a held-out $R^2 = 0.967$ (95% CI $[0.965, 0.969]$; Fig. 2A,B). Across the 40 topology-by-size groups, held-out R^2 ranged from 0.818 to 0.995 for linear-memory prediction and from 0.454 to 0.999 for nonlinear-mixing prediction (Fig. 2C). Thus, each response dimension predicted variation in its matched simulated capacity within every tested topology and network size.

Beyond the synthetic reservoirs, κ_1 and κ_2 predicted their matched simulated capacities in human and non-human connectomes. Across 100 participants from the Human Connectome Project Young Adult (HCP-YA) cohort, participant-held-out R^2 was 0.802 when predicting linear-memory capacity from κ_1 and 0.861 when predicting nonlinear-mixing capacity from κ_2 (95% CIs $[0.729, 0.855]$ and $[0.809, 0.902]$, respectively; Fig. 2D,E). For the non-human analysis, each connectome was held fixed while only its input map varied. For each of the four non-human species (Drosophila, rat, mouse, and macaque), we generated 60 maps that distributed the same temporal input stream across all nodes through independently sampled Gaussian input weights with the same scaling. Models fitted on one set of input maps predicted the capacities of held-out maps, with R^2 ranging from 0.871 to 0.945 for linear memory and from 0.607 to 0.813 for nonlinear mixing (Fig. 2F). The corresponding associations between each response dimension and its matched simulated capacity were positive across the synthetic, HCP-YA, and non-human analyses (Fig. S1). The same held-out analysis in 100 participants from the Lifespan Human Connectome Project in Aging (HCP-Aging) cohort yielded $R^2 = 0.710$ for linear memory and $R^2 = 0.808$ for nonlinear mixing (95% CIs $[0.582, 0.794]$ and $[0.721, 0.876]$;

Fig. S5). These human analyses concern simulated reservoir capacities of participant-specific connectomes, not participants' behavioral memory performance.

The response dimensions are derived from local responses around a specified operating point, so their predictive accuracy may depend on the model and input conditions. Stronger input drive moves the nonlinear reservoir farther from that point, where effects omitted from the local expansion can become more influential. Input gain is the dimensionless weight multiplying the external signal. We used 0.30 as the primary input gain because, for inputs in $[-1, 1]$, it confines the direct input contribution to $[-0.30, 0.30]$, a range over which the tanh activation remains close to linear. An input-gain sweep from 0.03 to 1.00 tested prediction under weaker and stronger drive. Held-out R^2 remained positive throughout this range, varying from 0.581 to 0.972 for linear memory and from 0.669 to 0.873 for nonlinear mixing (Fig. S2A). Sensitivity analyses also examined three other choices. Recurrent gain is the dimensionless multiplier on recurrent coupling, response ridge controls how strongly weak response directions contribute, and response horizon specifies how much input history is included. We varied all three quantities over the ranges shown in Fig. S2B–D. All held-out R^2 point estimates remained positive across the tested settings, although their magnitude varied. Prediction also remained accurate across spatial input maps. In HCP-YA, we repeated the analysis with 12 prespecified maps, each delivering the same temporal input to a different set of 40 of the 400 cortical parcels. Held-out R^2 ranged from 0.698 to 0.837 for linear memory and from 0.790 to 0.869 for nonlinear mixing (Fig. S3). Order specificity was tested using capacities summed over 12 prespecified mixtures of delayed inputs or delayed-input products. At the prespecified operating state for each order, κ_1 predicted linear-memory capacity with $R^2 = 0.793$, compared with $R^2 = 0.196$ using κ_2. Conversely, κ_2 predicted nonlinear-mixing capacity with $R^2 = 0.868$, compared with $R^2 = 0.358$ using κ_1 (Fig. S4A). For both capacities, the response dimension of the corresponding order also yielded lower participant-level mean squared error ($p = 4.0 \times 10^{-5}$; Fig. S4B). Finally, we compared the response dimensions with a combined model of seven conventional graph descriptors, mean strength, strength variability, density, algebraic connectivity, weighted clustering, spectral gap, and modularity, selected to summarize common weighted and topological properties of the same connectomes. The combined graph model yielded held-out $R^2 = 0.622$ for linear memory and $R^2 = 0.354$ for nonlinear mixing, compared with 0.802 and 0.861 for the response dimensions (Fig. S6A,B). Models based on the response dimensions also had lower mean participant-level squared error for linear memory ($p = 2.0 \times 10^{-4}$) and nonlinear mixing

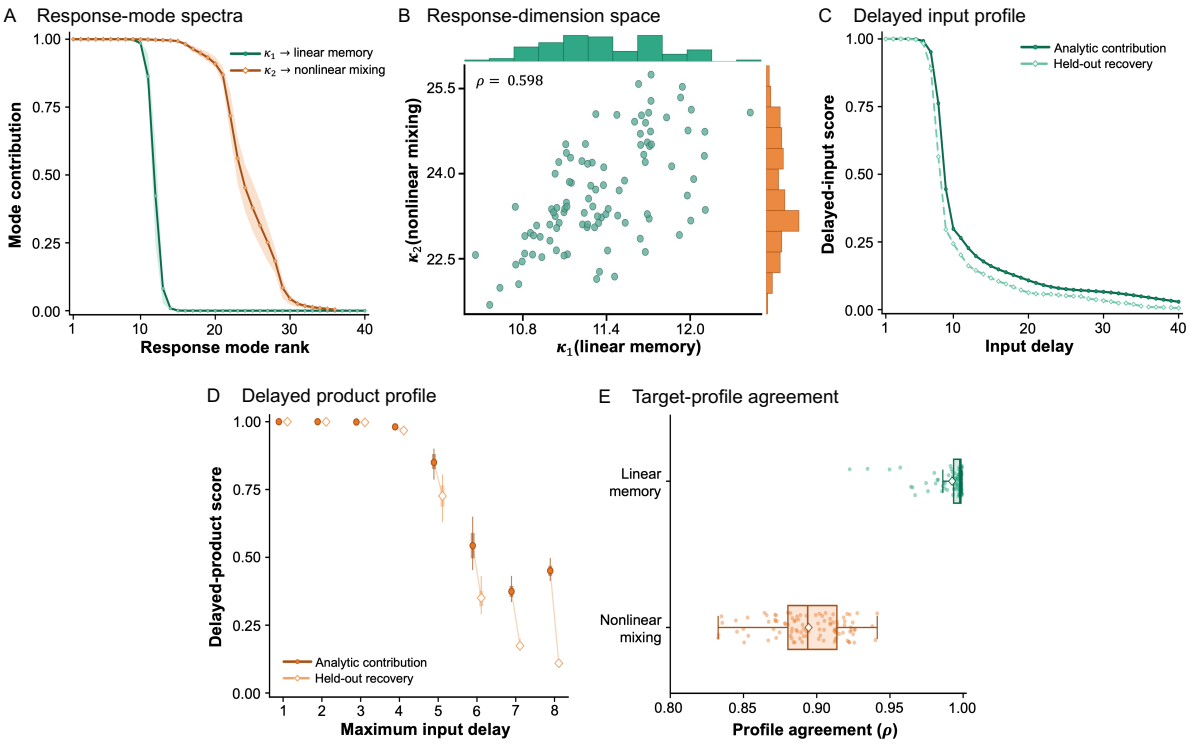

Figure 3: **Response dimensions decompose into recoverable memory and mixing profiles.** Response-mode spectra and target-resolved profiles show how contributions to κ_1 and κ_2 are distributed and correspond to held-out recovery across 100 HCP-YA connectomes. **A. Response-mode spectra.** Ranked first- and second-order mode contributions reveal distinct spectra for κ_1 and κ_2. **B. Response-dimension space.** κ_1 and κ_2 are moderately associated across connectomes ($\rho = 0.598$). **C. Delayed input profile.** Analytic first-order contributions and held-out recovery are shown across 40 input delays. **D. Delayed product profile.** The 36 second-order targets comprise eight centered squared-input terms and 28 products of inputs at different delays spanning lags 1–8. Analytic contributions and held-out recovery are averaged within each target's maximum input delay, giving the eight plotted positions. **E. Target-profile agreement.** Within each connectome, ρ compares analytic contributions with held-out recovery across all targets. Curves in **A**, **C**, and **D** show cohort means; bands in **A** show interquartile ranges, bands in **C** show 95% participant-bootstrap intervals, and bars in **D** show interquartile and 5th–95th percentile ranges. In **E**, points represent participants, boxes show interquartile ranges, center lines show medians, whiskers extend to 1.5× the interquartile range, and diamonds show means.

($p = 2.0 \times 10^{-5}$; Fig. S6C). Together, these analyses show that the response dimensions predict their corresponding simulated capacities across the tested reservoir and input conditions and more accurately than the tested graph descriptors.

2.2 Response dimensions decompose into mode- and target-specific profiles

The response dimensions predict total simulated linear-memory and nonlinear-mixing capacities. To resolve their internal geometry, we decomposed each aggregate dimension across individual response modes, defined as orthogonal directions in the corresponding first- or second-order

response space. Directions that are strong relative to the response ridge contribute values near one, whereas weak directions contribute values near zero; these contributions sum to the corresponding response dimension. Across 100 HCP-YA connectomes, the first-order spectrum fell sharply near rank 11, yielding a mean $\kappa_1 = 11.36$ across 40 modes. The second-order spectrum declined more gradually across higher ranks, yielding a mean $\kappa_2 = 23.65$ across 36 modes (Fig. 3A). Participant rankings in κ_1 and κ_2 were moderately associated ($\rho = 0.598$; Fig. 3B), indicating that connectomes ranked highly by one dimension were not necessarily ranked highly by the other.

The mode spectra describe how strongly orthogonal response directions contribute. A target-level decomposition linked each response dimension to specific functions of the input history. For κ_1, the targets were the 40 individual inputs $u_{t-1}, \ldots, u_{t-40}$. For κ_2, they were 36 degree-two functions spanning delays 1–8: eight centered squared-input terms and 28 products of inputs at different delays. Eight delays were chosen because they yield 36 unique degree-two targets, keeping the second-order basis close in size to the 40-target first-order basis; sensitivity to shorter and longer horizons was tested in Fig. S2D. Each target has an analytic contribution between zero and one, and the contributions sum to the corresponding response dimension. We compared each analytic contribution with the squared correlation between the target and the output of a linear readout evaluated on held-out time points. For linear memory, analytic contributions and held-out recovery were high for recent inputs and declined with increasing delay (Fig. 3C). For nonlinear mixing, both analytic contributions and held-out recovery declined as maximum input delay increased (Fig. 3D). Individual results for all eight squared-input and 28 cross-lag product targets are shown in Fig. S7.

Within each connectome, we used ρ to test whether targets with larger analytic contributions also had higher held-out recovery. The mean participant-level ρ was 0.992 for linear memory and 0.894 for nonlinear mixing (Fig. 3E). The held-out recovery profiles were also reproducible across the two independent input streams, with mean agreement of $\rho = 0.988$ for linear memory and $\rho = 0.982$ for nonlinear mixing (Fig. S8). Thus, under the specified reservoir conditions, the aggregate response dimensions can be decomposed into target-level profiles whose ordering corresponds closely to held-out target recovery.

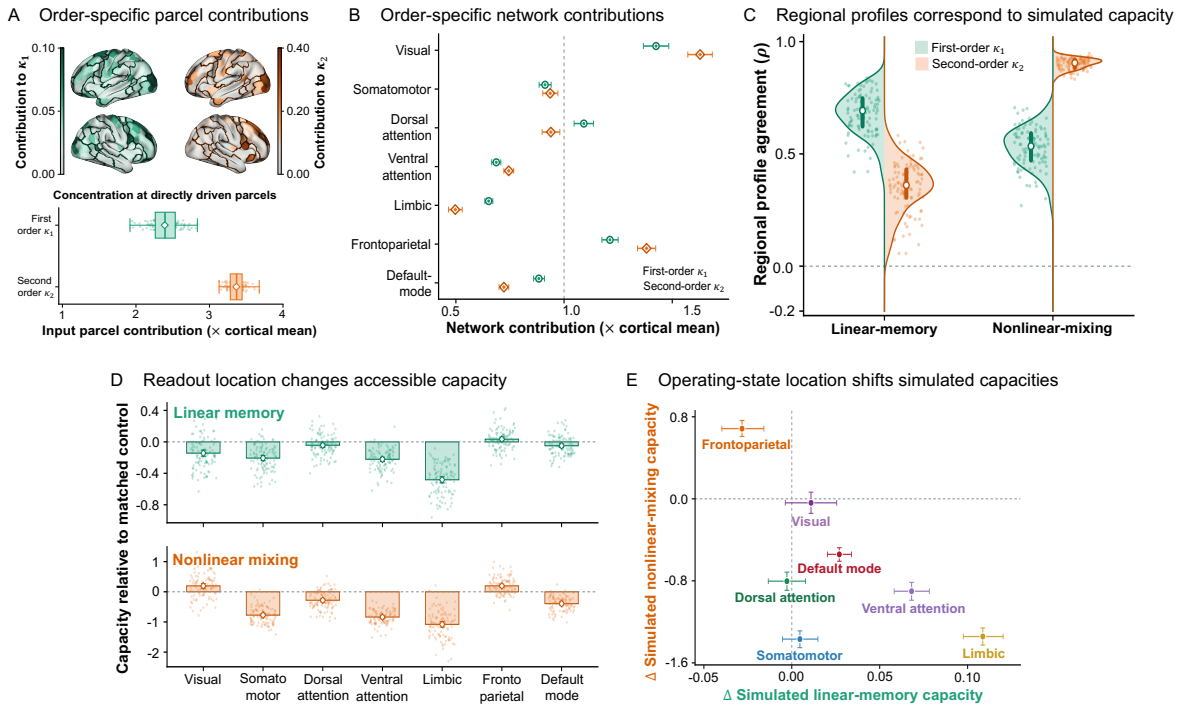

Figure 4: **Regional response profiles correspond to simulated memory and mixing capacities.** Regional decompositions relate response geometry to capacities accessible from cortical readouts and show how operating-state location alters full-state capacity across 100 HCP-YA connectomes. **A. Order-specific parcel contributions.** Cortical maps show cohort-mean parcel contributions to κ_1 and κ_2; the 40 parcels receiving temporal input contribute 2.39 and 3.37 times the cortical mean, respectively. **B. Order-specific network contributions.** Contributions averaged within the seven Yeo networks reveal distinct first- and second-order regional profiles. **C. Regional profiles correspond to simulated capacity.** Across 84 regional readout conditions per participant, comprising seven networks, six support realizations, and network and matched readouts, restricted κ_1 and κ_2 correspond to their matched capacities with mean $\rho = 0.681$ and 0.901, respectively. **D. Readout location changes accessible capacity.** Readouts restricted to 26 parcels within each Yeo network recover different capacities relative to matched 26-parcel readouts outside that network. **E. Operating-state location shifts simulated capacities.** The same time-independent bias with root-mean-square magnitude 0.09 is applied to 26 parcels within each Yeo network or to matched parcels outside it; plotted differences are network placement minus matched outside-network placement. Error bars in **B**, **D**, and **E** show 95% participant-bootstrap intervals. Dashed reference lines mark the cortical mean (1×) in **A** and **B**, $\rho = 0$ in **C**, and equality with the matched outside-network condition in **D** and **E**.

2.3 Regional response profiles correspond to accessible simulated capacities

A parcel-level decomposition characterized the spatial organization of the response dimensions. For each parcel, we calculated a nonnegative contribution to the ridge-smoothed count of response directions; these contributions sum exactly to κ_1 or κ_2. Across 100 HCP-YA connectomes, contributions extended throughout the cortex but were concentrated in the 40 parcels that received the temporal input. The mean contribution per directly driven parcel was 2.39

times the cortical mean for κ_1 and 3.37 times the cortical mean for κ_2 (Fig. 4A). Network contributions varied across the seven Yeo networks at both response orders ($p = 2.0 \times 10^{-5}$ for each order). Visual and frontoparietal contributions exceeded the cortical mean at both orders. Dorsal-attention contributions exceeded the mean at first order but not at second order, while somatomotor, ventral-attention, limbic, and default-mode contributions remained below the mean at both orders (Fig. 4B).

Regional restriction tested whether these response profiles corresponded to the simulated capacities accessible from the same locations. To compare networks without confounding their different sizes, we used 26 parcels from each Yeo network, matching the size of the smallest network, and paired each selection with 26 matched parcels outside that network. Both the analytic response matrix and the trained linear readout were restricted to the same parcels, while the reservoir dynamics and resulting activity remained unchanged. This produced 84 regional readout conditions per participant. Within each participant, the mean correlation between restricted κ_1 and linear-memory capacity was $\rho = 0.681$, and that between restricted κ_2 and nonlinear-mixing capacity was $\rho = 0.901$ (Fig. 4C). Participant-held-out models predicted regional linear-memory capacity with $R^2 = 0.663$ and nonlinear-mixing capacity with $R^2 = 0.802$ (95% CIs $[0.639, 0.686]$ and $[0.789, 0.814]$, respectively; Fig. S9A,B). For each of three independent temporal input streams, the response dimension of the corresponding order yielded higher participant-held-out R^2 for its matched capacity than did the response dimension of the other order (Fig. S9C). The difference between network and matched outside-network readouts varied across the seven networks for both capacities ($p = 2.0 \times 10^{-5}$ for each capacity). Relative to matched outside-network readouts, network-restricted readouts recovered less linear-memory capacity in six networks and more in the frontoparietal network. For nonlinear mixing, network-restricted readouts recovered more capacity in the visual and frontoparietal networks and less in the other five networks (Fig. 4D).

The restricted-readout analysis changed which regional states were observed while leaving the reservoir activity unchanged. A complementary analysis changed the operating state itself without altering connectivity. A bias is a constant additive term in the reservoir update; unlike the temporal input, it does not vary over time. Applying a bias to a parcel shifts its baseline state and local response gain. We applied the same bias values to 26 parcels within each Yeo network or to matched parcels outside it, and measured both capacities from the full reservoir state. The network-minus-control difference therefore isolates bias location rather than the presence of a

bias. The simulated capacity differences varied across the seven networks ($p = 2.0 \times 10^{-5}$ for both capacities; Fig. 4E). Relative to their matched outside-network placements, bias placement within the ventral-attention, limbic, and default-mode networks yielded higher simulated linear-memory capacity ($\Delta = 0.068$, 0.109, and 0.027, respectively) but lower simulated nonlinear-mixing capacity ($\Delta = -0.901$, -1.343, and -0.542, respectively). Frontoparietal placement yielded lower simulated linear-memory capacity ($\Delta = -0.028$) but higher simulated nonlinear-mixing capacity ($\Delta = 0.685$). Matching controls more closely on continuous node strength reduced their mean strength mismatch from 0.039 to 0.006 and preserved the nonlinear-mixing direction for all seven networks and the linear-memory direction for five of seven networks (Fig. S10). Together, these results show that, within the specified reservoir model, the parcels available to the readout changed accessible simulated capacity, while the location of a constant bias changed full-state capacity.

2.4 Response-gradient reweighting changes simulated memory and mixing capacities

The response dimensions predict simulated capacities and decompose their target-level and regional organization. Constrained reweighting tested whether their gradients identify connection-weight changes that alter those same held-out capacities. For each connectome, we reweighted existing positive connections by 0.1% of $\|A\|_F$ along the positive and negative gradients while preserving the original support, tract-length-weighted wiring cost, and unit spectral radius. The κ_1 and κ_2 reweightings were evaluated at the zero-bias and fixed biased operating points, respectively. Across 100 HCP-YA connectomes, reweighting along $+\nabla_A \kappa_1$ increased simulated linear-memory capacity by $\Delta = 0.010$ (95% CI $[0.008, 0.011]$), whereas reweighting along $-\nabla_A \kappa_1$ decreased it by $\Delta = -0.009$ ($[-0.011, -0.008]$). Similarly, reweighting along $+\nabla_A \kappa_2$ increased simulated nonlinear-mixing capacity by $\Delta = 0.017$ ($[0.016, 0.018]$), whereas reweighting along $-\nabla_A \kappa_2$ decreased it by $\Delta = -0.017$ ($[-0.018, -0.016]$; Fig. 5A). Each direction also changed the other capacity, but these mean changes were smaller than those for the corresponding capacity. The analytic response dimensions themselves moved in the requested direction in every connectome, with mean changes of 0.016 and -0.016 for κ_1 and 0.029 and -0.029 for κ_2 (Fig. S11A). On average, the directed capacity change exceeded 97% of eight magnitude- and constraint-matched random reweightings for linear memory and 100% for nonlinear mixing (Fig. S11B). The total capacity changes were distributed across multiple delayed-input and delayed-product

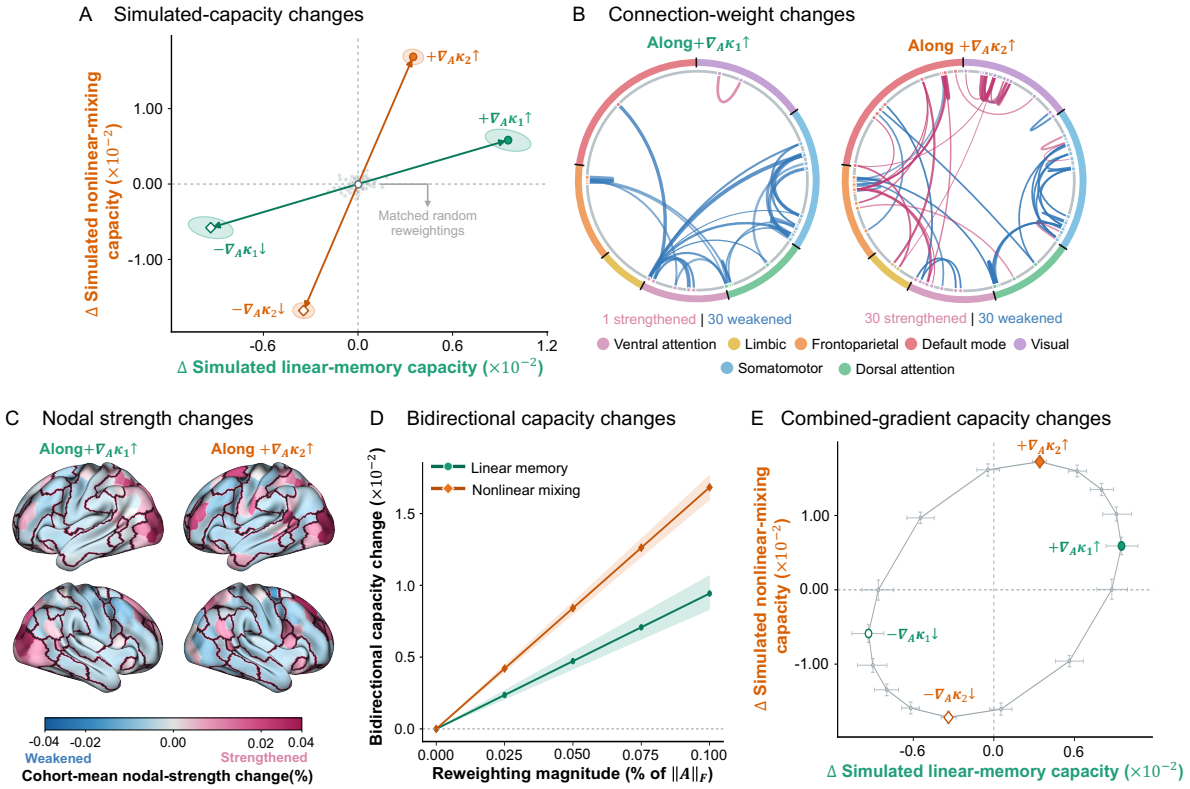

Figure 5: **Response-gradient reweighting changes simulated memory and mixing capacities.** Across 100 HCP-YA connectomes, constrained reweightings along the gradients of κ_1 and κ_2 change their matched held-out capacities and span distinct local directions in capacity space. **A. Simulated-capacity changes.** Changes in both capacities are shown after reweighting by 0.1% of $\|A\|_F$ along $\pm\nabla_A\kappa_1$ and $\pm\nabla_A\kappa_2$; grey symbols show participant means across eight matched random reweightings. **B. Connection-weight changes.** Connectograms show the largest cohort-consistent strengthened and weakened existing connections after reweighting along each positive-gradient direction. **C. Nodal strength changes.** Cortical maps show cohort-mean percentage changes in recurrent nodal strength after the same positive-gradient reweightings. **D. Bidirectional capacity changes.** Sign-aligned changes from the positive- and negative-gradient directions are averaged within each participant across reweighting magnitudes from zero to 0.1% of $\|A\|_F$. **E. Combined-gradient capacity changes.** Sixteen equal-magnitude mixtures of the two gradient directions trace cohort-mean capacity changes; the four pure-gradient directions are highlighted. In **A**, ellipses show 95% participant-bootstrap regions; in **D**, bands show 95% participant-bootstrap intervals; and in **E**, error bars show coordinate-wise 95% participant-bootstrap intervals. Connections in **B** are present in at least 80 participants and change with the same sign in at least 75% of participants carrying that connection; up to 30 per response order and sign are shown.

scores (Fig. S12).

Connection- and parcel-level summaries characterized how the positive-gradient reweightings were distributed across the connectome. For each connection, we calculated the log-weight change, $\log(A'_{ij}/A_{ij})$, with positive values denoting strengthening and negative values denoting weakening. We retained connections that were present in at least 80 participants and changed with the same sign in at least 75% of the participants in whom they were present. Under

these criteria, one strengthened and 29,018 weakened connections were retained along $+\nabla_A \kappa_1$, while 45 strengthened and 42,638 weakened connections were retained along $+\nabla_A \kappa_2$ (Fig. 5B). Parcel-level effects were summarized using recurrent nodal strength, defined as the total recurrent connection weight incident to a parcel. After reweighting along $+\nabla_A \kappa_1$, cohort-mean nodal-strength changes ranged from -0.009% to 0.039% across the 400 parcels. Changes along $+\nabla_A \kappa_2$ ranged from -0.015% to 0.112% (Fig. 5C). The two response objectives therefore produced different distributions of connection-weight and nodal-strength changes while retaining the original set of connections.

A magnitude sweep from zero to 0.1% of $\|A\|_F$ tested whether capacity changes increased with perturbation size. At each magnitude, we sign-aligned and averaged the positive- and negative-gradient capacity changes within each participant. Bidirectional changes increased across the tested magnitudes, reaching $\Delta = 0.009$ for linear memory (95% CI $[0.008, 0.011]$) and $\Delta = 0.017$ for nonlinear mixing ($[0.016, 0.018]$) at 0.1% (Fig. 5D). Expressed relative to each participant's baseline capacity, the mean bidirectional changes were 0.095% for linear memory and 0.094% for nonlinear mixing, similar in scale to the 0.1% connection-weight perturbation. The change was monotonic across the five magnitudes in 97 participants along $+\nabla_A \kappa_1$, 97 along $-\nabla_A \kappa_1$, and all 100 participants along both κ_2 directions. Finally, we evaluated 16 equal-magnitude mixtures of the two gradient directions. The resulting capacity changes formed a two-dimensional trajectory rather than collapsing onto a single line (Fig. 5E). A local capacity-gain matrix summarizing the two effects had a positive determinant in 95 of 100 connectomes. Its mean normalized span was 0.709 (95% CI $[0.645, 0.770]$), where zero indicates collinear capacity changes and one indicates orthogonal changes. The two response gradients therefore provide distinct local directions for changing simulated linear-memory and nonlinear-mixing capacities.

2.5 Constrained reweighting targets and increases response dimensions

The response gradients provide distinct local directions for changing simulated capacities. Repeated constrained reweighting tested whether the corresponding response dimensions could be moved toward specified values. For each connectome and response order, a target controller recalculated the gradient after every update and selected the direction that reduced the remaining target error while preserving the original set of positive connections, tract-length-weighted wiring cost, and unit spectral radius. Each κ_j is a dimensionless effective count, so a change

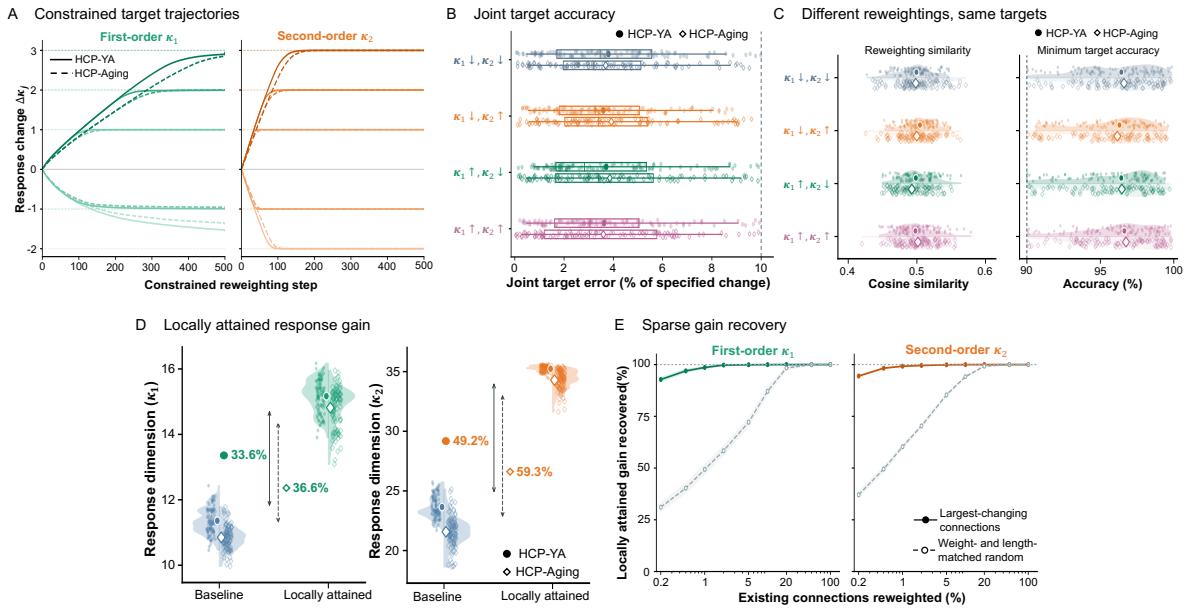

Figure 6: **Constrained reweighting targets and increases response dimensions.** Across 100 HCP-YA and 100 HCP-Aging connectomes, constrained searches target κ_1 and κ_2, find distinct reweightings with similar response endpoints, and concentrate locally attained gains in a small fraction of existing connections. **A. Constrained target trajectories.** Cohort-mean trajectories move toward changes of -2, -1, $+1$, $+2$, and $+3$ from baseline; the first 500 of up to 1,000 reweighting steps are shown, and dotted lines mark the targets. **B. Joint target accuracy.** Four joint targets combine increases or decreases in both dimensions, with each component set to one-third of its ten-step response change; joint target error is the larger component-wise percentage error. **C. Different reweightings, same targets.** For each joint target, cosine similarity compares two alternative connection-weight changes, and minimum accuracy is evaluated across the standard solution, both alternatives, and both dimensions. **D. Locally attained response gain.** At the endpoints of positive-gradient searches, mean participant-level gain in κ_1 is 33.6% in HCP-YA and 36.6% in HCP-Aging, while mean gain in κ_2 is 49.2% and 59.3%, respectively. **E. Sparse gain recovery.** In HCP-YA, the largest-changing 0.2% of existing connections recover 92.8% of the locally attained κ_1 gain and 94.5% of the κ_2 gain, compared with 31.1% and 37.1% for weight- and length-matched random subsets; HCP-Aging results are shown in Fig. S14. In **A**–**D**, solid curves or filled circles denote HCP-YA and dashed curves or open diamonds denote HCP-Aging; pale marks in **B**–**D** represent participants. Boxes show interquartile ranges and medians, whiskers span the 5th–95th percentiles, and half-violins show participant distributions. Curves in **E** show cohort means, bands show 95% participant-bootstrap intervals, and the dashed line marks complete recovery.

of one represents one unit in the effective number of response directions. We specified changes of -2, -1, $+1$, $+2$, and $+3$ from baseline and considered a target attained when its error fell below 0.01 within 1,000 steps. Across 100 HCP-YA and 100 HCP-Aging connectomes, every participant attained all five κ_2 targets and the $+1$ and $+2$ targets for κ_1. Most participants also attained $\Delta\kappa_1 = +3$ (93 HCP-YA and 94 HCP-Aging) and $\Delta\kappa_1 = -1$ (97 and 88), whereas $\Delta\kappa_1 = -2$ was attained in 29 and 17 participants, respectively (Fig. 6A). The controller selected reweightings using only the analytic response dimensions. Across three held-out input

streams, linear-memory capacity changed in the corresponding direction in 96–99 HCP-YA participants and 95–97 HCP-Aging participants, whereas nonlinear-mixing capacity did so in every participant (Fig. S13).

Joint targeting extended this test to both response dimensions simultaneously. We formed four participant-scaled joint targets by combining a decrease or increase in each dimension, with each change set to one-third of the response change produced by ten constrained gradient steps in the corresponding direction. Across 100 HCP-YA and 100 HCP-Aging connectomes, all 800 participant–target combinations reached both response dimensions within the 10% error threshold in the allowed 20 iterations. Across the standard joint-controller solutions, mean joint target error was 3.72%, and the largest error was 9.98% (Fig. 6B). To test non-uniqueness, we constructed two additional constraint-preserving reweightings for each participant and target. Together with the standard solution, all 800 participant–target sets contained three distinct reweightings that reached the specified response values. The two alternatives had a median cosine similarity of 0.499, showing that their connection-weight changes were not identical, while the lowest accuracy across all three reweightings and both response dimensions was 90.0% (Fig. 6C). Thus, under the specified reservoir model and constraints, a target pair of response dimensions does not uniquely determine how the existing connections must be reweighted.

Finally, repeated positive-gradient searches from each original connectome tested how much each response dimension could be increased under the same constraints and how broadly the attained gain was distributed across existing connections. Each search ran for up to 1,000 constrained steps. Mean participant-level gain in κ_1 was 33.6% in HCP-YA and 36.6% in HCP-Aging (95% CIs $[32.6, 34.7]$ and $[35.6, 37.6]$, respectively). Mean gain in κ_2 was 49.2% and 59.3% ($[48.2, 50.2]$ and $[57.9, 60.7]$, respectively), and every participant in both cohorts increased both response dimensions. Every search reached the 1,000-step limit while continuing to improve. These endpoints therefore represent locally attained gains and lower bounds on the attainable response, rather than local or global optima (Fig. 6D). To test whether these gains required changes across most existing connections, we ranked connections by the magnitude of their change in the full search and repeated the constrained search while allowing only nested subsets of the highest-ranked connections to vary. In HCP-YA, the largest-changing 0.2% of connections recovered 92.8% of the locally attained κ_1 gain and 94.5% of the κ_2 gain, compared with 31.1% and 37.1% for weight- and length-matched random subsets (Fig. 6E). The same analysis in HCP-Aging recovered 88.9% and 91.9% of the respective gains, compared with 26.4% and 33.8% for

matched random subsets (Fig. S14). Together, these results show that constrained response design can move both dimensions toward specified values through multiple connection-level solutions, while concentrating much of the locally attained gain in a small fraction of existing connections.

3 Discussion

We have introduced Connectome Response Analysis and Design (CoRAD), a framework for predicting, decomposing, and modifying simulated linear-memory and nonlinear-mixing capacities under specified reservoir conditions. The order-specific response dimensions predicted their matched capacities across held-out synthetic reservoirs, human connectomes, and input maps within non-human connectomes. Decomposing these dimensions across response modes, input-history targets, and cortical parcels revealed how capacity is distributed and how it depends on readout location and operating state. Constrained reweighting along the response gradients produced corresponding changes in the matched simulated capacities and moved both dimensions toward individual and joint targets. Different connection-weight changes reached similar response targets, while a small fraction of existing connections recovered much of the locally attained gain. Together, these findings show response geometry as a link between recurrent connectivity and the analysis and design of model-defined temporal-processing capacities.

The response dimensions retain graded information about response geometry that algebraic rank alone discards. For linear recurrent networks with independent inputs, rank determines total linear memory but cannot distinguish full-rank networks of equal size[25;33]. By weighting response modes relative to the response ridge, κ_1 and κ_2 distinguish networks through their response spectra. CoRAD therefore complements control-theoretic uses of the Gramian to quantify the energy required to reach target states and computational affordance landscapes[34–36]: it instead quantifies which specified functions of the input history remain recoverable by a linear readout[21]. It also complements analytic reservoir programming through output and feedback connections[32]; after calibration within a model setting, CoRAD predicts matched capacities in held-out recurrent networks and differentiates the response dimensions with respect to recurrent connection weights.

The regional analyses emphasize that a fixed connectome does not have a single simulated-capacity profile. Restricting the readout changed which capacities were accessible from the

same reservoir activity, while relocating a constant bias changed the capacities expressed by the same connectivity. This dependence on observation and operating state is consistent with evidence that the same anatomical circuit can express different dynamics when its dynamical regime or neuromodulatory conditions change[10;11;16]. The cortical profiles should therefore be interpreted as properties of the specified reservoir configuration, not as evidence that those cortical systems implement behavioral memory or nonlinear mixing in vivo. The response-design analyses addressed a complementary intervention question: whether constrained changes to existing connection weights could shift these model-based capacities. Reweighting along the response gradients produced corresponding changes in the matched held-out capacities, providing a model-internal intervention test of the link between response geometry and simulated capacity. Mixed gradients spanned a local two-dimensional range of linear-memory and nonlinear-mixing changes, and different admissible reweightings reached similar joint targets. This non-uniqueness is consistent with circuit degeneracy, in which different cellular and synaptic parameter combinations can generate similar network activity[12–14]. Much of the locally attained gain was also recovered by reweighting a small fraction of existing connections, indicating concentrated leverage within the optimization but not yet identifying biological control points. Response-gradient design may therefore provide a useful strategy for engineering physical or neuromorphic reservoirs implemented in manipulable devices and substrates[38].

Several limitations define the scope of these conclusions. First, κ_1 and κ_2 depend on the chosen state-update model, input map, operating point, response horizon, ridge scale, and spectral normalization; κ_2 additionally summarizes the fixed set of second-order targets studied here. Their predictions remained robust across the tested gains, ridge scales, response horizons, and input maps. Different state-update rules, temporally correlated inputs, and response bases beyond second order remain to be examined. Second, the human results describe simulated capacities of connectomes reconstructed using diffusion tractography[39;40]; they do not measure participants' behavioral memory. Each non-human species was represented by a single connectome. Larger datasets linking model-derived capacities to behavior and containing multiple connectomes per species will be needed to establish biological relevance and cross-species generality. The regional analyses also used a fixed adult group atlas; repeating them with data-driven, cohort- and development-specific functional parcellations would test how strongly the spatial profiles depend on the adopted regional definition[41]. Third, response design reweights a mathematical network under specified constraints; it neither identifies a biological mecha-

nism for producing those changes nor proves that the locally attained endpoints are optimal. Future work should test CoRAD in experimentally manipulable neural or physical reservoirs, incorporate temporally structured and multivariate inputs, and extend the response basis to higher-order and task-specific functions. Design constraints could also be expanded to include cell types, plasticity rules, metabolic costs, and other limits on biologically realizable changes. Together, these directions will determine how far connectomes can be read to characterize their responses, decompiled to identify the organization of simulated capacities, and written to produce targeted changes in constrained network models.

4 Methods

4.1 Connectome data and preprocessing

We analyzed structural connectomes from two human cohorts: 100 unrelated participants in the Human Connectome Project Young Adult (HCP-YA) S1200 release[42] and 100 participants in the Lifespan Human Connectome Project in Aging (HCP-Aging)[43;44]. The HCP-YA cohort comprised 54 female and 46 male participants, with age-band midpoints from 23.5 to 38.0 years (mean 29.4 years). HCP-YA connectomes were reconstructed in MRtrix3 using multi-shell multi-tissue constrained spherical deconvolution, anatomically constrained probabilistic tractography, and SIFT2 streamline weighting[39;40]; HCP-Aging connectomes used the same SIFT2 and inverse-node-volume weighting. In both cohorts, we retained the 400 cortical parcels of the Schaefer-400 atlas from the 4S456 parcellation[45]. The weight between parcels i and j was the SIFT2-weighted streamline count multiplied by $2/(V_i + V_j)$, where V_i and V_j are parcel volumes. Regional analyses used the seven Yeo cortical networks[46].

Human adjacency matrices were symmetrized, assigned zero diagonals, clipped to nonnegative weights, scaled by their largest edge weight, and normalized to unit spectral radius, yielding recurrent matrices A with $\rho(A) = 1$. Tract-length matrices were symmetrized, assigned zero diagonals, clipped to nonnegative values, and scaled by their largest entry to give L.

The non-human analyses used one directed connectome from each of four species: Drosophila ($N = 49$)[47], rat ($N = 73$)[48], mouse ($N = 112$)[49], and macaque ($N = 82$)[50]. We oriented each matrix so that A_{ij} denotes propagation from node j to node i, set its diagonal to zero, scaled it by its largest absolute edge weight, and normalized it by the modulus of its leading eigenvalue.

4.2 Reservoir dynamics, inputs, and operating states

We modeled each connectome as the fixed recurrent matrix A of a nonlinear reservoir[18;22;23], with state $s_t \in \mathbb{R}^N$ evolving as

$$s_{t+1} = \tanh(\alpha A s_t + b\, u_t + d)\,, \tag{1}$$

where tanh was applied elementwise, u_t was drawn independently from Uniform$[-1, 1]$, b was the spatial input map, and d was a time-independent bias. The bias changed the operating state without carrying temporal information. Because $\rho(A) = 1$, the dimensionless recurrent gain α scaled the strength and persistence of recurrent propagation; we set $\alpha = 0.9$ unless otherwise stated. The primary human input map contained 40 randomly selected parcels, with $b_i = 0.30$ at selected parcels and zero elsewhere. Hence $b_i u_t \in [-0.30, 0.30]$, where tanh remains close to linear.

At $u_t = 0$, the operating point s^* satisfies

$$s^* = \tanh(\alpha A s^* + d)\,. \tag{2}$$

We obtained s^* by fixed-point iteration from zero until convergence. First-order prediction and linear-memory analyses used $d = 0$, giving $s^* = 0$. Primary human second-order and nonlinear-mixing analyses used $d = 0.30\,\xi$, where $\xi \in \mathbb{R}^{400}$ was a fixed standard-normal vector shared across the two human cohorts. Synthetic and non-human analyses used fixed dimension-matched standard-normal vectors. Because $\tanh''(0) = 0$, a nonzero operating point was required for a nonzero quadratic response. Figure 4A–D evaluated both response orders at the biased operating point to compare their spatial profiles under a common state.

4.3 First- and second-order response dimensions

Let $M = \alpha A$. The first and second derivatives of the activation function at the operating point are

$$g_1 = 1 - s^{*2}, \qquad g_2 = -2s^* \odot (1 - s^{*2}). \tag{3}$$

These derivatives define the local transition matrix and direct input response,

$$F = \mathrm{diag}(g_1)M, \qquad g = g_1 \odot b. \tag{4}$$

For a response horizon $K_1 = 40$, the first-order response matrix stacks the propagated input responses by delay,

$$R^{(1)} = \left[g, Fg, \ldots, F^{K_1-1}g\right]^\top, \qquad K_1 = 40. \tag{5}$$

Row p therefore corresponds to the target u_{t-p}, including the direct response g at $p = 1$. The node-space Gram matrix $W_c^{(1)} = R^{(1)\top}R^{(1)}$ is the finite-horizon controllability Gramian of (F, g)[34]. Under temporally white unit-variance input, it is also the covariance of the truncated first-order propagated response. Supplementary Proposition 1 states and proves both identities.

We obtained the second-order responses from the quadratic terms of the same local expansion. Define $a_1 = b$ and $a_p = MF^{p-2}g$ for $p \geq 2$. For each unordered lag pair $1 \leq p \leq q \leq K_2$, with $K_2 = 8$,

$$q_{pq} = c_{pq}\, g_2 \odot a_p \odot a_q + \mathbf{1}_{p>1} F q_{p-1,q-1}, \qquad c_{pq} = \begin{cases} \frac{1}{2}, & p = q, \\ 1, & p < q. \end{cases} \tag{6}$$

where the factor $1/2$ is the diagonal Taylor coefficient. The corresponding targets are $u_{t-p}u_{t-q}$ for $p < q$ and $u_{t-p}^2 - \mathbb{E}[u^2]$ for $p = q$. Under $u \sim \text{Uniform}[-1, 1]$, these targets have variances $1/9$ and $4/45$, respectively. We formed $R^{(2)} \in \mathbb{R}^{36\times N}$ by stacking the vectors $q_{pq}^\top$ after scaling each by the square root of its target variance, and defined $W_c^{(2)} = R^{(2)\top}R^{(2)}$.

For either response order $j \in \{1, 2\}$, the response dimension was

$$\kappa_j = \operatorname{tr}\left[W_c^{(j)}(W_c^{(j)} + \epsilon I_N)^{-1}\right] = \sum_r \frac{\lambda_{jr}}{\lambda_{jr} + \epsilon}, \qquad \epsilon = 10^{-10}, \tag{7}$$

where λ_{jr} are the eigenvalues of $W_c^{(j)}$. Each eigenvalue contributes between zero and one, making κ_j a ridge-smoothed effective count of response directions. Thus $0 \leq \kappa_1 \leq 40$ and $0 \leq \kappa_2 \leq 36$. The first-order construction is exact for the stated local linearized system, whereas the second-order construction is a finite degree-two approximation over the 36 delayed-product targets.

4.4 Mode-, target-, and parcel-level decompositions

The contribution of response mode r was

$$m_{jr} = \frac{\lambda_{jr}}{\lambda_{jr} + \epsilon}, \tag{8}$$

with $\sum_r m_{jr} = \kappa_j$. The contribution of input-history target k was

$$\tau_{jk} = \left[R^{(j)} (W_c^{(j)} + \epsilon I_N)^{-1} R^{(j)\top} \right]_{kk}. \tag{9}$$

Each τ_{jk} lies between zero and one, and $\sum_k \tau_{jk} = \kappa_j$. We compared these analytic contributions with the held-out squared-correlation recovery scores defined below. Profile agreement was the within-participant Spearman correlation across the 40 first-order or 36 second-order targets. For the second-order delay profile, targets sharing the same maximum delay were averaged within participant. Recovery profiles were averaged across two independent temporal input streams, and their between-stream agreement was assessed separately.

The contribution of parcel i was the corresponding diagonal entry of the node-space ridge leverage matrix,

$$\ell_{ji} = \left[W_c^{(j)} (W_c^{(j)} + \epsilon I_N)^{-1} \right]_{ii}. \tag{10}$$

These contributions are nonnegative and satisfy $\sum_i \ell_{ji} = \kappa_j$. Network-level enrichment was the mean parcel contribution within a Yeo network divided by the cortical mean contribution.

4.5 Simulated linear-memory and nonlinear-mixing capacities

Reservoir simulations began from the zero state and used 4,500 input values, with the first 500 states discarded. Rows with complete target histories were divided chronologically into 70% training and 30% held-out test sets. Each target was mean-centered using the training data and decoded from the reservoir state by ridge regression without an intercept and with penalty 10^{-6}. Its held-out recovery score was

$$c_k = \mathrm{corr}^2(y_k, \widehat{y}_k). \tag{11}$$

Following standard reservoir-capacity definitions[21;24], linear-memory capacity was the sum across 40 delayed inputs,

$$C_1 = \sum_{p=1}^{40} \mathrm{corr}^2(u_{t-p}, \widehat{u}_{t-p}). \tag{12}$$

Nonlinear-mixing capacity was the sum across all 36 unordered delayed products,

$$C_2 = \sum_{1 \le p \le q \le 8} \mathrm{corr}^2\big(u_{t-p}u_{t-q}, \widehat{u_{t-p}u_{t-q}}\big). \tag{13}$$

For $p = q$, mean-centering makes the squared-input target equivalent under correlation scoring to the centered square used in the analytic basis. All networks within a comparison received the same input stream.

4.6 Synthetic reservoirs and non-human input maps

The synthetic ensemble comprised 4,000 weighted undirected reservoirs: 100 realizations for each combination of four topology families and ten sizes $N = 50, 100, \ldots, 500$. The families were Erdős–Rényi (ER)[51], Watts–Strogatz (WS)[52], Barabási–Albert (BA)[53], and stochastic block models (SBM)[54]. ER connection probabilities were sampled from $[0.05, 0.40]$. WS neighborhood degree varied from 4 to $\min(30, \lfloor N/3 \rfloor)$ over even values, with rewiring probability sampled from $[0.05, 0.50]$. BA attachment parameters were sampled from integers below $\min(15, \lfloor N/4 \rfloor)$, with a minimum value of 2. SBM networks contained two to five near-equal communities, with between- and within-community probabilities sampled from $[0.02, 0.10]$ and $[0.30, 0.70]$, respectively. Existing edges received independent Uniform$[0.1, 1]$ weights before preprocessing. Each reservoir received input with weight 0.30 at 36 nodes, and the input nodes were held fixed across realizations of the same size.

For each non-human connectome, we generated 60 spatial input maps,

$$b = 0.30\,\xi/\sqrt{N}, \qquad \xi \sim \mathcal{N}(0, I). \tag{14}$$

Each map distributed the same scalar temporal input across all nodes through independently sampled signed weights. The temporal input stream was fixed across maps within a connectome. First-order simulations used 4,000 inputs and a 400-state washout; second-order simulations used 4,500 inputs and a 500-state washout. Both used the chronological split and readout ridge defined above, and the response dimensions and capacities were recomputed for every map. This analysis evaluated variation across input maps within each fixed connectome.

4.7 Held-out prediction and sensitivity analyses

Each prediction model used one response dimension as a scalar predictor in a linear regression with an intercept fitted only to training units. Synthetic prediction used five-fold cross-validation repeated 20 times, with a separate model fitted within each topology-by-size group. HCP-YA and HCP-Aging prediction used ten-fold participant cross-validation repeated 20

times, and non-human prediction used ten-fold input-map cross-validation repeated 20 times. Out-of-fold predictions and training-fold-mean reference predictions were averaged across repeats before scoring

$$R^2_{\text{held-out}} = 1 - \frac{\sum_i (y_i - \widehat{y}_i)^2}{\sum_i (y_i - \widehat{y}_{0i})^2}, \tag{15}$$

where $\widehat{y}_{0i}$ is the corresponding training-fold outcome mean. Calibration, feature standardization, and hyperparameter selection were confined to the training folds.

Raw associations used Spearman correlation. Sensitivity analyses repeated HCP-YA prediction across input gains from 0.03 to 1.00, recurrent gains from 0.60 to 0.98, response ridges from 10^{-14} to 10^{-4}, and response horizons of 10–80 delays for linear memory and 4–10 lags for nonlinear mixing. In the horizon analysis, each capacity was summed over the target window used to construct its response dimension. Input-map sensitivity used 12 distinct maps, each selecting 40 HCP-YA parcels while holding the remaining conditions fixed.

Response-order specificity was tested using 12 equal-variance orthogonal mixtures of delayed inputs as linear targets and 12 mixtures of the complete degree-two basis as nonlinear targets. Capacity scores were averaged across two independent input streams. The corresponding mappings, $\kappa_1 \to C_1$ and $\kappa_2 \to C_2$, were compared with the two cross-order mappings using identical participant folds. Participant-level squared-error differences were tested by paired sign flips with Holm adjustment across the two capacities.

The graph benchmark used seven descriptors of each HCP-YA adjacency: mean strength, standard deviation of nodal strength, density, weighted-Laplacian algebraic connectivity, mean weighted clustering, adjacency spectral gap, and Newman modularity. Single-descriptor models used the same linear calibration as the response models. The combined model standardized all descriptors within each training fold and used ridge regression, with its penalty selected by nested cross-validation. Response and graph models used identical outer folds.

4.8 Regional response and operating-state analyses

Regional analyses used the seven Yeo networks in the 100 HCP-YA connectomes. To compare networks at equal size, each regional parcel set contained 26 parcels, matching the size of the smallest network. For each network, we sampled six parcel sets and paired each with 26 parcels outside that network, matching counts by hemisphere, membership in the fixed 40-parcel temporal-input map, and cohort-mean nodal-strength tertile. This yielded $7 \times 6 \times 2 = 84$

regional conditions per participant.

For each analytic condition, we selected the corresponding parcel columns of $R^{(1)}$ or $R^{(2)}$ before applying Eq. (7). For simulated regional capacity, reservoir dynamics and full state trajectories were unchanged across parcel sets, while the linear readout used only the same 26 parcels. Capacities were averaged across three independent temporal input streams. Regional prediction excluded all 84 conditions from the evaluated participant when fitting the prediction model. Spatial profile agreement was the within-participant Spearman correlation across conditions, and network-minus-outside capacity differences were averaged across the six paired parcel sets.

Operating-state analyses tested whether the location of a constant bias changed full-state capacity while the temporal input map b remained fixed. For each Yeo network, the bias was placed on either 26 parcels in that network or a matched parcel set outside it. Each placement used 26 standard-normal values scaled to root-mean-square magnitude 0.09, with the same value multiset assigned to the paired network and outside-network sets. Six paired placements were evaluated per network using readouts fitted to the full reservoir state. Reported shifts were network placement minus outside-network placement, isolating bias location rather than the presence of a bias.

A sensitivity analysis replaced strength-tertile matching with one-to-one outside-network assignments that preserved hemisphere and temporal-input-map membership while minimizing differences in cohort-mean nodal strength. The bias values were unchanged; only the outside-network parcel sets differed.

4.9 Constrained response-gradient reweighting

Response design reweighted only positive existing connections in the symmetric human connectomes. We parameterized their weights in log coordinates, preserving positivity and the original edge set. For κ_1 at $d = 0$, automatic differentiation through Eq. (5) gave $\nabla_A \kappa_1$. The gradient of κ_2 also included the dependence of the operating point on A. If $r = \partial\kappa_2/\partial s^*$ denotes the derivative with A fixed, the adjoint z satisfies

$$(I - F)^\top z = r. \tag{16}$$

The implicit fixed-point contribution $\alpha(g_1 \odot z)s^{*\top}$ was added to the partial adjacency derivative before symmetrization. Both gradients were assigned zero diagonals.

Tract-length-weighted wiring cost was

$$C(A) = \langle A, L \rangle_F = \sum_{i,j} A_{ij} L_{ij}. \tag{17}$$

In log coordinates, the steepest-ascent direction is $A \odot \nabla_A \kappa_j$ on the original edge set. We projected this direction onto the tangent space orthogonal to the constraint normals of $C(A)$ and the leading eigenvalue. Finite reweightings exponentiated the supported log weights and were corrected along these normals to restore $C(A') = C(A)$ and $\rho(A') = 1$; a scalar step set the requested distance $\|A' - A\|_F$. We denote this constraint-preserving retraction by $\mathcal{R}_A$. It preserved symmetry, a zero diagonal, the original positive edge set, exact wiring cost, and unit spectral radius.

Local positive and negative reweightings used $\|A' - A\|_F = 0.001\|A\|_F$. For each response order, eight Gaussian log-weight directions were retracted under the same constraints and magnitude to form random controls. The main signed-gradient, random-control, magnitude, and target-resolved capacity analyses used one common temporal input stream. Signed-gradient endpoints were repeated on two additional streams, and combined-gradient capacities were averaged across all three. Connection-level change was $\log(A'_{ij}/A_{ij})$. Displayed connections were present in at least 80 participants and changed with the same sign in at least 75% of those participants; up to the 30 largest absolute cohort-median changes were shown per group. Nodal strength was the sum of all incident recurrent weights.

The magnitude analysis repeated positive and negative reweightings at five equally spaced distances from zero to $0.001\|A\|_F$. For each participant and response order, the bidirectional capacity change was

$$\Delta C_{\mathrm{bi}} = \frac{\Delta C(+\nabla_A \kappa_j) - \Delta C(-\nabla_A \kappa_j)}{2}. \tag{18}$$

For the combined-gradient analysis, let d_1 and d_2 denote the unit projected log-weight directions for κ_1 and κ_2. We mixed them as $\cos\theta\, d_1 + \sin\theta\, d_2$ at 16 equally spaced angles $\theta \in [0, 2\pi)$, reprojected each mixture, and applied the same 0.1% retraction. First angular harmonics of the two capacity changes defined a 2×2 local gain matrix. Its normalized span was the absolute determinant divided by the product of its column norms, with zero indicating collinear effects and one indicating orthogonal effects.

4.10 Target-directed and repeated constrained reweighting

The target controller recomputed the relevant response gradient after each accepted reweighting and selected its sign to reduce the remaining target error. Positive and negative paths began at the baseline connectome and visited targets from nearest to farthest. Candidate steps used the retraction $\mathcal{R}_A$, were limited to 0.001 times the Frobenius norm of the current adjacency, and were shortened when they failed to reduce the error or crossed the target. A response target was considered attained when its absolute error was at most 0.01. We evaluated $\Delta\kappa_j \in \{-2, -1, 1, 2, 3\}$. Searches used at most 1,000 accepted steps, of which the first 500 are displayed in Fig. 6A.

For the simulated-capacity validation in Fig. S13, positive and negative response excursions were defined by ten constrained gradient steps. Targets were placed at one-third and two-thirds of each excursion. The corresponding simulated capacity was evaluated at the reached endpoints on three independent temporal input streams and averaged within participant.

Joint control used four sign combinations formed by increasing or decreasing each response dimension by one-third of its participant-specific ten-step excursion. At each iteration, the two log-weight gradients were scaled by these excursions and combined by a damped least-squares update that reduced the larger normalized component error. Searches used at most 20 iterations. A joint target was considered attained when both component errors were no larger than the smaller of 0.01 response units and 10% of the specified change.

To test whether a joint response target uniquely determined the connection-weight changes, we generated random log-weight directions orthogonal to both response gradients, the standard adjacency change, and the two constraint normals. After an equal-distance retraction, the joint controller removed the residual target error. Two alternative reweightings were constructed for each participant and target. Distinctness was measured by cosine similarity between adjacency-change vectors, and target accuracy was the minimum percentage accuracy across the standard and alternative reweightings and both response dimensions.

Locally attained headroom was estimated by repeatedly following the positive response gradient from the baseline connectome for at most 1,000 accepted steps, with backtracking used to preserve ascent. Every search was still increasing at the step limit, so its endpoint was treated as a lower bound rather than an optimum.

For sparse gain recovery, existing undirected connections were ranked by their absolute weight change at the unrestricted endpoint. Constrained searches then modified only the top

0.2, 0.5, 1, 2, 5, 10, 20, 50, or 100% of connections. Weight- and tract-length-matched random subsets provided controls at each fraction. Recovery was

$$100 \times \frac{\kappa_{\text{restricted}} - \kappa_{\text{empirical}}}{\kappa_{\text{unrestricted}} - \kappa_{\text{empirical}}}. \tag{19}$$

The procedure was applied independently in HCP-YA and HCP-Aging.

4.11 Statistical analysis

All reported confidence intervals are 95% percentile-bootstrap intervals. Resampling followed the independent unit of each analysis: a synthetic reservoir within its topology–size group, a participant, or an input map within a fixed non-human connectome. Prediction intervals were calculated from fixed out-of-fold predictions and therefore quantify variation across the evaluated observations rather than uncertainty from refitting the complete prediction procedure.

The associations were quantified with Spearman's ρ; Pearson correlations were used only where stated for capacity scores. Paired participant comparisons used sign-flip randomization, and variation across the seven Yeo networks was tested by permuting network labels within participant. Holm adjustment controlled multiplicity across paired response-order or capacity comparisons and across the two response-versus-graph comparisons. Network-specific follow-up tests used the Benjamini–Hochberg false-discovery-rate procedure. Monte Carlo randomization values were calculated as $(1 + n_{\text{extreme}})/(1 + n_{\text{draws}})$.

Citation diversity statement

Recent work in several fields of science has identified a bias in citation practices such that papers from women and other minority scholars are under-cited relative to the number of such papers in the field[55–63]. Here we sought to proactively consider choosing references that reflect the diversity of the field in thought, form of contribution, gender, race, ethnicity, and other factors. First, we obtained the predicted gender of the first and last author of each reference by using databases that store the probability of a first name being carried by a woman[59;64]. By this measure (and excluding self-citations to the first and last authors of our current paper), our references contain 7.15% woman (first)/woman (last), 6.11% man/woman, 19.52% woman/man, and 67.22% man/man. This method is limited in that a) names, pronouns, and social media profiles used to construct the databases may not, in every case, be indicative of gender identity

and b) it cannot account for intersex, non-binary, or transgender people. Second, we obtained predicted racial/ethnic category of the first and last author of each reference by databases that store the probability of a first and last name being carried by an author of color[65;66]. By this measure (and excluding self-citations), our references contain 12.75% author of color (first)/author of color (last), 19.06% white author/author of color, 22.32% author of color/white author, and 45.87% white author/white author. This method is limited in that a) names and Florida Voter Data used to make the predictions may not be indicative of racial/ethnic identity, and b) it cannot account for Indigenous and mixed-race authors, or those who may face differential biases due to the ambiguous racialization or ethnicization of their names. We look forward to future work that could help us to better understand how to support equitable practices in science.

Data availability

The Human Connectome Project Young Adult S1200 data analysed in this study are publicly available from the WU-Minn HCP Consortium through the ConnectomeDB repository (`https://db.humanconnectome.org`) under the HCP Open Access Data Use Terms. The HCP-Aging 2.0 Release data are available from the NIMH Data Archive (`https://nda.nih.gov`; DOI: 10.15154/1520707) under the NDA Data Use Certification. The Schaefer 400-region parcellation is publicly available[45]. The Drosophila, rat, and mouse connectomes[47–49] are available from the processed connectivity data distributed with the conn2res toolbox (`https://doi.org/10.5281/zenodo.10205004`)[18], and the macaque connectome[50] is available from Zenodo (`https://doi.org/10.5281/zenodo.1464105`). Source data are provided with this paper.

Code availability

All codes and statistics required to reproduce the figures are made available at `https://www.soveshmohapatra.com/research/connectome-compiler`.

References

[1] Vyas, S., Golub, M. D., Sussillo, D. & Shenoy, K. V. Computation through neural population dynamics. *Annual Review of Neuroscience* **43**, 249–275 (2020).

[2] Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).

[3] Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical control of arm movements: A dynamical systems perspective. *Annual Review of Neuroscience* **36**, 337–359 (2013).

[4] Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).

[5] Deco, G., Jirsa, V. K. & McIntosh, A. R. Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nature Reviews Neuroscience* **12**, 43–56 (2011).

[6] Honey, C. J. *et al.* Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences* **106**, 2035–2040 (2009).

[7] Mišić, B. *et al.* Cooperative and competitive spreading dynamics on the human connectome. *Neuron* **86**, 1518–1529 (2015).

[8] Park, H.-J. & Friston, K. Structural and functional brain networks: From connections to cognition. *Science* **342**, 1238411 (2013).

[9] Bargmann, C. I. & Marder, E. From the connectome to brain function. *Nature Methods* **10**, 483–490 (2013).

[10] Marder, E. Neuromodulation of neuronal circuits: Back to the future. *Neuron* **76**, 1–11 (2012).

[11] Nadim, F. & Bucher, D. Neuromodulation of neurons and synapses. *Current Opinion in Neurobiology* **29**, 48–56 (2014).

[12] Prinz, A. A., Bucher, D. & Marder, E. Similar network activity from disparate circuit parameters. *Nature Neuroscience* **7**, 1345–1352 (2004).

[13] Marder, E. & Taylor, A. L. Multiple models to capture the variability in biological neurons and networks. *Nature Neuroscience* **14**, 133–138 (2011).

[14] Deistler, M., Macke, J. H. & Gonçalves, P. J. Energy-efficient network activity from disparate circuit parameters. *Proceedings of the National Academy of Sciences* **119**, e2207632119 (2022).

[15] Suárez, L. E., Markello, R. D., Betzel, R. F. & Misic, B. Linking structure and function in macroscale brain networks. *Trends in Cognitive Sciences* **24**, 302–315 (2020).

[16] Suárez, L. E., Richards, B. A., Lajoie, G. & Misic, B. Learning function from structure in neuromorphic networks. *Nature Machine Intelligence* **3**, 771–786 (2021).

[17] Damicelli, F., Hilgetag, C. C. & Goulas, A. Brain connectivity meets reservoir computing. *PLOS Computational Biology* **18**, e1010639 (2022).

[18] Suárez, L. E. *et al.* Connectome-based reservoir computing with the conn2res toolbox. *Nature Communications* **15**, 656 (2024).

[19] Maass, W., Natschläger, T. & Markram, H. Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation* **14**, 2531–2560 (2002).

[20] Ganguli, S., Huh, D. & Sompolinsky, H. Memory traces in dynamical systems. *Proceedings of the National Academy of Sciences* **105**, 18970–18975 (2008).

[21] Dambre, J., Verstraeten, D., Schrauwen, B. & Massar, S. Information processing capacity of dynamical systems. *Scientific Reports* **2**, 514 (2012).

[22] Lukoševičius, M. & Jaeger, H. Reservoir computing approaches to recurrent neural network training. *Computer Science Review* **3**, 127–149 (2009).

[23] Jaeger, H. The "echo state" approach to analysing and training recurrent neural networks. GMD Report 148, GMD Forschungszentrum Informationstechnik, Sankt Augustin, Germany (2001).

[24] Jaeger, H. Short term memory in echo state networks. GMD Report 152, GMD Forschungszentrum Informationstechnik, Sankt Augustin, Germany (2001).

[25] Gonon, L., Grigoryeva, L. & Ortega, J.-P. Memory and forecasting capacities of nonlinear recurrent networks. *Physica D: Nonlinear Phenomena* **414**, 132721 (2020).

[26] White, O. L., Lee, D. D. & Sompolinsky, H. Short-term memory in orthogonal neural networks. *Physical Review Letters* **92**, 148102 (2004).

[27] Mijalkov, M. *et al.* Computational memory capacity predicts aging and cognitive decline. *Nature Communications* **16**, 2748 (2025).

[28] Schulte to Brinke, T., Dick, M., Duarte, R. & Morrison, A. A refined information processing capacity metric allows an in-depth analysis of memory and nonlinearity trade-offs in neurocomputational systems. *Scientific Reports* **13**, 10517 (2023).

[29] Inubushi, M. & Yoshimura, K. Reservoir computing beyond memory-nonlinearity trade-off. *Scientific Reports* **7**, 10199 (2017).

[30] Gu, S., Mattar, M. G., Tang, H. & Pan, G. Emergence and reconfiguration of modular structure for artificial neural networks during continual familiarity detection. *Science Advances* **10**, eadm8430 (2024).

[31] Gu, S. *et al.* Task-structured modularity emerges in artificial networks and aligns with brain architecture. *Research Square* (2025). Preprint.

[32] Kim, J. Z. & Bassett, D. S. A neural machine code and programming framework for the reservoir computer. *Nature Machine Intelligence* **5**, 622–630 (2023).

[33] Ballarin, G., Grigoryeva, L. & Ortega, J.-P. Memory of recurrent networks: Do we compute it right? *Journal of Machine Learning Research* **25**, 1–38 (2024). URL `https://jmlr.org/papers/v25/23-0568.html`.

[34] Pasqualetti, F., Zampieri, S. & Bullo, F. Controllability metrics, limitations and algorithms for complex networks. *IEEE Transactions on Control of Network Systems* **1**, 40–52 (2014).

[35] Gu, S. *et al.* Controllability of structural brain networks. *Nature Communications* **6**, 8414 (2015).

[36] Kulkarni, S. S., Kim, J. Z., Fotiadis, P., Pasqualetti, F. & Bassett, D. S. Quantifying the cost of network computations to unpack structure–function relationships in the brain. *arXiv preprint arXiv:2607.29537* (2026). `2607.29537`.

[37] Mohapatra, S. *et al.* QUIET: Quantifying underutilized influential edges for targeted synchronization. *arXiv preprint arXiv:2606.11091* (2026). `2606.11091`.

[38] Tanaka, G. *et al.* Recent advances in physical reservoir computing: A review. *Neural Networks* **115**, 100–123 (2019).

[39] Tournier, J.-D. *et al.* MRtrix3: A fast, flexible and open software framework for medical image processing and visualisation. *NeuroImage* **202**, 116137 (2019).

[40] Smith, R. E., Tournier, J.-D., Calamante, F. & Connelly, A. SIFT2: Enabling dense quantitative assessment of brain white matter connectivity using streamlines tractography. *NeuroImage* **119**, 338–351 (2015).

[41] Mohapatra, S. *et al.* TReND: Transformer derived features and regularized NMF for neonatal functional network delineation. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2025*, vol. 15971 of *Lecture Notes in Computer Science*, 660–669 (Springer Nature Switzerland, 2025).

[42] Van Essen, D. C. *et al.* The WU-Minn Human Connectome Project: an overview. *NeuroImage* **80**, 62–79 (2013).

[43] Bookheimer, S. Y. *et al.* The Lifespan Human Connectome Project in aging: an overview. *NeuroImage* **185**, 335–348 (2019).

[44] Harms, M. P. *et al.* Extending the Human Connectome Project across ages: Imaging protocols for the Lifespan Development and Aging projects. *NeuroImage* **183**, 972–984 (2018).

[45] Schaefer, A. *et al.* Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cerebral Cortex* **28**, 3095–3114 (2018).

[46] Yeo, B. T. T. *et al.* The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology* **106**, 1125–1165 (2011).

[47] Chiang, A.-S. *et al.* Three-dimensional reconstruction of brain-wide wiring networks in Drosophila at single-cell resolution. *Current Biology* **21**, 1–11 (2011).

[48] Bota, M., Sporns, O. & Swanson, L. W. Architecture of the cerebral cortical association connectome underlying cognition. *Proceedings of the National Academy of Sciences* **112**, E2093–E2101 (2015).

[49] Rubinov, M., Ypma, R. J. F., Watson, C. & Bullmore, E. T. Wiring cost and topological participation of the mouse brain connectome. *Proceedings of the National Academy of Sciences* **112**, 10032–10037 (2015).

[50] Shen, K. *et al.* A macaque connectome for large-scale network simulations in TheVirtualBrain. *Scientific Data* **6**, 123 (2019).

[51] Erdős, P. & Rényi, A. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences* **5**, 17–61 (1960). URL `https://real.mtak.hu/200959/`.

[52] Watts, D. J. & Strogatz, S. H. Collective dynamics of 'small-world' networks. *Nature* **393**, 440–442 (1998).

[53] Barabási, A.-L. & Albert, R. Emergence of scaling in random networks. *Science* **286**, 509–512 (1999).

[54] Holland, P. W., Laskey, K. B. & Leinhardt, S. Stochastic blockmodels: First steps. *Social Networks* **5**, 109–137 (1983).

[55] Mitchell, S. M., Lange, S. & Brus, H. Gendered citation patterns in international relations journals. *International Studies Perspectives* **14**, 485–492 (2013).

[56] Maliniak, D., Powers, R. & Walter, B. F. The gender citation gap in international relations. *International Organization* **67**, 889–922 (2013).

[57] Caplar, N., Tacchella, S. & Birrer, S. Quantitative evaluation of gender bias in astronomical publications from citation counts. *Nature Astronomy* **1**, 0141 (2017).

[58] Dion, M. L., Sumner, J. L. & Mitchell, S. M. Gendered citation patterns across political science and social science methodology fields. *Political Analysis* **26**, 312–327 (2018).

[59] Dworkin, J. D. *et al.* The extent and drivers of gender imbalance in neuroscience reference lists. *Nature Neuroscience* **23**, 918–926 (2020).

[60] Bertolero, M. A. *et al.* Racial and ethnic imbalance in neuroscience reference lists and intersections with gender. *bioRxiv* (2020).

[61] Wang, X. *et al.* Gendered citation practices in the field of communication. *Annals of the International Communication Association* **45**, 134–153 (2021).

[62] Chatterjee, P. & Werner, R. M. Gender disparity in citations in high-impact journal articles. *JAMA Network Open* **4**, e2114509 (2021).

[63] Fulvio, J. M., Akinnola, I. & Postle, B. R. Gender (im)balance in citation practices in cognitive neuroscience. *Journal of Cognitive Neuroscience* **33**, 3–7 (2021).

[64] Zhou, D. *et al.* dalejn/cleanBib: v1.1.3. https://doi.org/10.5281/zenodo.7375227 (2022).

[65] Ambekar, A., Ward, C., Mohammed, J., Male, S. & Skiena, S. Name-ethnicity classification from open sources. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 49–58 (ACM, 2009).

[66] Chintalapati, R., Laohaprapanon, S. & Sood, G. Predicting race and ethnicity from the sequence of characters in a name. *arXiv preprint arXiv:1805.02109* (2018). 1805.02109.

Acknowledgments

Data were provided in part by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 National Institutes of Health (NIH) Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University. Research reported in this publication was supported by the National Institute on Aging of the National Institutes of Health under Award Number U01AG052564 and by funds provided by the McDonnell Center for Systems Neuroscience at Washington University in St. Louis. The HCP-Aging 2.0 Release data used in this report came from DOI: 10.15154/1520707. S.M. and D.S.B. acknowledge support from the Army Research Office MURI program (through grant number W911NF2410228).

Contribution

Conceptualization: S.M., D.S.B. Methodology: S.M., D.S.B. Data Analyses, Writing – Original Draft: S.M. Writing – Review & Editing: S.M., D.S.B. Supervision: D.S.B.

Competing interests

The authors declare no competing interests.