nature protocols

https://doi.org/10.1038/s41596-024-01023-w

Protocol

Check for updates

A network control theory pipeline for studying the dynamics of the structural connectome

Linden Parkes[1,2,3,15] ✉, Jason Z. Kim[4,15], Jennifer Stiso[2], Julia K. Brynildsen[2], Matthew Cieslak[3,5,6], Sydney Covitz[3,5,6], Raquel E. Gur[5,6], Ruben C. Gur[5,6], Fabio Pasqualetti[7], Russell T. Shinohara[8,9,10], Dale Zhou[2], Theodore D. Satterthwaite[3,5,6,9] & Dani S. Bassett[2,6,11,12,13,14]

Abstract

Network control theory (NCT) is a simple and powerful tool for studying how network topology informs and constrains the dynamics of a system. Compared to other structure–function coupling approaches, the strength of NCT lies in its capacity to predict the patterns of external control signals that may alter the dynamics of a system in a desired way. An interesting development for NCT in the neuroscience field is its application to study behavior and mental health symptoms. To date, NCT has been validated to study different aspects of the human structural connectome. NCT outputs can be monitored throughout developmental stages to study the effects of connectome topology on neural dynamics and, separately, to test the coherence of empirical datasets with brain function and stimulation. Here, we provide a comprehensive pipeline for applying NCT to structural connectomes by following two procedures. The main procedure focuses on computing the control energy associated with the transitions between specific neural activity states. The second procedure focuses on computing average controllability, which indexes nodes' general capacity to control the dynamics of the system. We provide recommendations for comparing NCT outputs against null network models, and we further support this approach with a Python-based software package called 'network control theory for python'. The procedures in this protocol are appropriate for users with a background in network neuroscience and experience in dynamical systems theory.

Key points

- We present a protocol on how to model the dynamics of neural connectivity states using network control theory (NCT) via a software package written in Python to compute the control energy associated with the transitions between states and the average controllability of the network's dynamics.
- NCT complements biophysical models of neuronal communication and graph-theoretical measures of internodal communication.

Key references

Gu, S. et al. *Neuroimage* **148**, 305–317 (2017): https://doi.org/10.1016/j.neuroimage.2017.01.003

Kim, J. Z. et al. *Nat. Phys.* **14**, 91–98 (2018): https://doi.org/10.1038/nphys4268

Stiso, J. et al. *Cell Rep.* **28**, 2554–2566.e7 (2019): https://doi.org/10.1016/j.celrep.2019.08.008

Karrer, T. M. et al. *J. Neural Eng.* **17**, 026031 (2020): https://doi.org/10.1088/1741-2552/ab6e8b

Parkes, L. et al. *Sci. Adv.* **8**, eadd2185 (2022): https://doi.org/10.1126/sciadv.add2185

A full list of affiliations appears at the end of the paper. ✉e-mail: linden.parkes@rutgers.edu

Introduction

Network neuroscience is principally concerned with studying the connectome[1], the description of whole brain connectivity. The connectome is often encoded as a graph of nodes interconnected by edges that can be defined across multiple scales, species and data modalities[2,3]. In any case, this description of brain connectivity gives rise to complex topology, including hubs, modules, small-worldness and core-periphery structure[4], and understanding how this topology shapes and constrains the brain's rich repertoire of dynamics is a central goal of network neuroscience.

Network control theory (NCT) provides an approach to studying these dynamics and yields insights into the relationship between patterns of neural dynamics and the topology of the underlying structural connectome[5–8]. The application of NCT has revolutionized both the understanding and design of complex networks in contexts as diverse as space and terrestrial exploration and modeling of financial markets, airline networks and fire-control systems. Briefly, NCT assumes that inter-nodal communication follows a linear model of diffusion, in which activity from one set of nodes (i.e., an initial state) spreads across the network over time along a series of fronts[4,9]. Then, upon this dynamical system, NCT models a set of external control signals designed to guide the diffusing activity patterns toward a chosen target state. This choice can be informed by a measurement of activity evoked by behavior, spontaneous activity or the type of brain system. These control signals are found by minimizing the total magnitude of their input over a given time horizon; that is, they are designed to achieve a desired state transition with the lowest amount of control energy. Once modeled, these control signals can be examined to determine to what extent, and how, they were constrained by topology, thus allowing researchers to study how the connectome might be leveraged to control dynamics.

Recently, we have developed and tested the application of NCT to brain network data across multiple contexts, scales and definitions of connectivity[10–22]. Here, we present a protocol for applying NCT to two different structural connectomes: one defined using undirected connectivity estimated in the human brain[23,24] and the other using directed connectivity estimated in the mouse brain[25–27]. Briefly, we detail two common applications of NCT that we—as well as other groups[28–34]—have deployed and that focus on (i) quantifying the amount of energy that is required to complete transitions between specific brain states (Fig. 1) and (ii) modeling regional capacity to control unspecified state transitions (Fig. 2). The former approach is useful for researchers interested in testing whether dynamic state changes can be controlled and induced across the network, while the latter is useful for researchers interested in analyzing topographic maps of control. In addition, we provide recommendations for the visualization of model outputs and discuss the use of null network models to examine which topological features affect model outputs.

Development of the procedure

The methods that underlie NCT are based on the established fields of control theory and dynamical systems theory. Dating back to at least the 19th century[35], control theory is primarily concerned with engineering perturbations to achieve desired behaviors in the states of a system, and specifically the evolution of such states over time. Hence, one of the most natural ways to formulate theories of control is through differential and difference equations that mathematically define the next state of a system given its current state. A common example of a control system is an inverted pendulum on a cart: the system states are the positions and velocities of the cart and pendulum, the differential equations are determined through the governing Newtonian physics and the control task is to perturb the cart so that the cart and pendulum end up in a desired state. For example, one might want to push the cart back and forth in such a way as to stabilize the pendulum so that it remains upright[36].

From one perspective, the inverted pendulum is not unlike the brain, where the system states are the activities of neural units (e.g., brain regions), the differential equations are determined through the diffusion of activity through structural connections between those

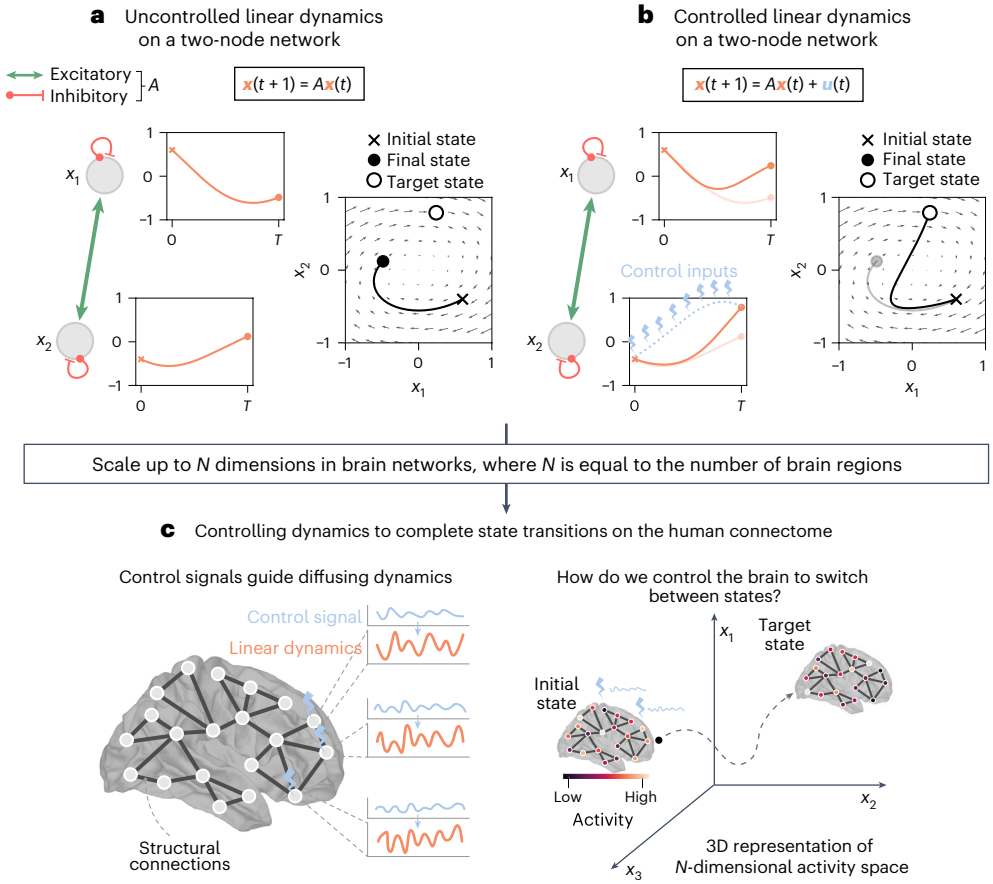

Fig. 1 | Modeling the control energy required to complete a state transition. NCT finds the control signals that, when injected into a networked system, will guide simulated neural activity from an initial state to a target state. Here, we show a two-node toy network (x_1, x_2) that illustrates the difference between neural activity (solid orange lines) in the absence (**a**) and presence (**b**) of a control signal (dashed blue line). **a**, Uncontrolled linear dynamics on a two-node network. Given an initial state in which $x_1 = 0.3$ and $x_1 = -0.2$, as well as coupling between nodes encoded by A, uncontrolled neural activity unfolds as shown on the left. These dynamic trajectories can also be represented in two dimensions as a vector field as shown on the right. Under this uncontrolled regime, the state of the system culminates in $x_1 = -0.24$ and $x_2 = 0.06$ at time T. **b**, Controlled linear dynamics on a two-node network. By contrast, when we introduce a control signal to x_2, the trajectory changes to now culminate in $x_1 = 0.12$ and $x_1 = 0.39$ at time T. Thus, NCT has found the control signal that drove our system from our initial state [$x_1 = 0.3, x_2 = -0.2$] to our desired target state [$x_1 = 0.12, x_2 = 0.39$]. **c**, NCT applied to the human connectome. The above model can be extended to the scale of N brain regions that constitute a whole-brain connectome (left). In doing so, we can model and examine the control signals required to transition the brain between various states of interest (right).

units and the control task is to perturb the brain to steer it to a desired state. There is a rich history of such modeling of the brain as a dynamical system using differential equations, ranging from biophysical models of single neurons[37] to phenomenological[38] and coarse-grained[39] models of neural populations. In tandem, there is a very practical translational need to understand how to control brain dynamics[40] to compensate for abnormal dynamics that may be present in neurological and neuropsychiatric disorders. For example, there is growing interest in using neuroimaging to personalize non-invasive neurostimulation techniques, such as transcranial magnetic stimulation, for depression[41]. NCT is well positioned to assist in these endeavors and may provide insight into the propagating dynamics that these techniques elicit.

Despite the above, the analogy between the brain and an inverted pendulum is insufficient for several reasons. First, the dimensionality and complexity of the brain cannot be reduced to such simple models. Understanding how the topology of the structural connectome gives rise to brain function is a difficult task that has motivated a large body of work in the last two decades.

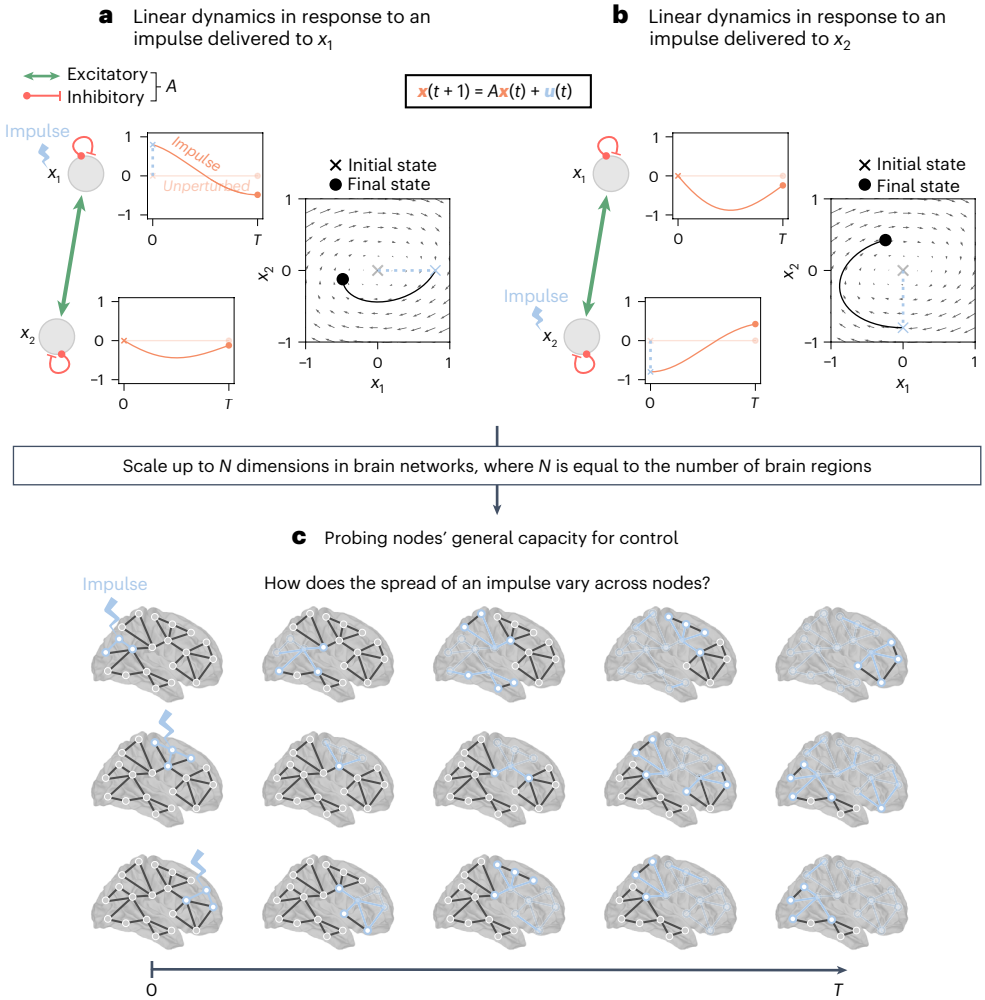

Fig. 2 | Average controllability: modeling the impulse response of the system from each node. NCT can be used to probe regions' general capacity to control dynamics. This is achieved by studying how the system responds to an impulse delivered to each node. Here, we show a two-node toy network (x_1, x_2) coupled by A. Upon this network, we demonstrate how neural activity (solid orange lines) unfolds when an impulse (dashed blue line) is delivered to x_1 (**a**) and x_2 (**b**). **a**, An impulse is delivered to x_1 that sets the initial state of the system to $[x_1 = 0.4, x_2 = 0]$. **b**, An impulse is delivered to x_2 that sets the initial state of the system to $[x_1 = 0, x_2 = -0.4]$. In each case, the impulse response of the system is quantified as the area under the squared curves of the two orange traces. Intuitively, this measurement corresponds to the amount of activity propagated throughout the system over time. We refer to this measure as the average controllability. Thus, for a given time horizon (T), a region with higher average controllability is better able to broadcast an impulse. **c**, Impulse response modeled for the human connectome. The above model can be extended to the scale of N brain regions that constitute a whole-brain connectome. In doing so, we can compare each region's capacity to broadcast an impulse across the whole brain.

This research has revealed that structure–function coupling is not one to one; instead it varies spatially across the cortex[42–46] and is stronger when indirect structural pathways are accounted for under multiple models of network communication[47,48]. Second is the distributed nature of brain states for human function. Although some brain regions may be thought of as supporting specific functions (e.g., the fusiform face area), carrying out complex human functions typically requires the recruitment of a network of brain regions to a distributed brain state[49]. Finally, biology imposes relatively tight operating constraints. To support complex human functions, the brain needs to optimize for efficient signaling while balancing the need to minimize wiring cost within the spatial constraints of the cranial cavity. Hence, there is a need

to express the unique complexities and constraints of controlling brain structure–function coupling in the quantitative formalism of dynamics and control.

NCT emerges as a flexible framework to address this need that is methodologically based in optimal control theory[50] and can accommodate a wide range of theoretical and experimental hypotheses and constraints about structure–function coupling through a consistent mathematical framework[20,51,52]. Because NCT posits a model of neural dynamics at the level of individual neural interactions, it allows us to probe the role of the complex structural connectome on brain function at the level of those interactions[14,22,53]. In addition, because NCT parameterizes which regions to control and how, as well as the precise patterns of initial and target neural activity, it can answer questions ranging from the importance of a single region for propagating activity[10] to the cost of transitioning between specific brain states[11]. Hence, the development of NCT has largely served to provide a simple, first-order biophysical model with the flexibility and power to study more advanced hypotheses of brain function.

The modeling framework of NCT comprises N nodes (e.g., neurons or brain regions) and m inputs and stipulates that the state of each node, $x_i(t)$, evolves in time as a weighted sum of the state of all upstream nodes, $x_j(t)$, and any inputs, $u_m(t)$. If the evolution of the system can be framed in terms of discrete states, in which the activity of upstream nodes determines the state of downstream nodes at the next discrete point in time, then the model takes the form of a difference equation:

$$\boldsymbol{x}(t+1) = A\boldsymbol{x}(t) + B\boldsymbol{u}(t)\,, \tag{1}$$

where $\boldsymbol{x}(t) = [x_1(t), x_2(t), \cdots, x_n(t)]^{\mathrm{T}}$ is the vector of neural states, A is the $N{\times}N$ connectome that comprises N nodes along the rows and columns interconnected by N^2 edges, $\boldsymbol{u}(t) = [u_1(t), u_2(t), \ldots, u_m(t)]^{\mathrm{T}}$ is the vector of independent control signals and B is the $N{\times}m$ matrix that quantifies how each input affects the nodes. If instead the evolution of the system can be framed in terms of rates, in which the activity of upstream nodes affects the continuous rate at which the state of downstream nodes change, then the model takes the form of the differential equation:

$$\frac{d}{dt}\boldsymbol{x}(t) = A\boldsymbol{x}(t) + B\boldsymbol{u}(t)\,. \tag{2}$$

Although these two models appear similar because the right side of the equations are identical, their definition, properties and behavior differ substantially. In turn, the interpretation of the model parameters and outputs can vary dramatically between them. We discuss several implications of this difference in Experimental design.

In the Procedure, we discuss two common operationalizations of NCT that can be derived from either of these models. The first (Procedure 1) uses a time-varying perturbation, $\boldsymbol{u}(t)$, to drive the neural activity, $\boldsymbol{x}(t)$, from an initial state, $\boldsymbol{x}_0$, to a target state, $\boldsymbol{x}_f$, given a balance of constraints on the magnitude of both the neural states and the perturbations. The magnitude of these perturbations is summarized as the control energy, which we interpret as the amount of effort that the model system must exert to complete a given state transition. The second (Procedure 2) is average controllability, which is the magnitude of the neural activity, $\boldsymbol{x}(t)$, in response to an impulse stimulus delivered to a single node; a node with higher average controllability is better able to leverage graph topology to spread an impulse throughout the system. Note that average controllability is only one example of a node-level NCT metric that falls within the broader category of controllability statistics. This category encompasses NCT outputs that describe different ways in which the nodes of the system may control its dynamics. Although we have used other controllability statistics in our previous work (e.g., modal controllability; see Understanding the influence of topology below), we focus on average controllability in this protocol because of its simple intuitive nature and broad appeal.

We focus on these two operationalizations, control energy and average controllability, because they encompass two common sets of questions about the brain. The first set of questions stems from advances in neuroimaging that allow us to empirically measure neural

states via functional MRI (fMRI), electrophysiology and calcium imaging[54]. Given these state-level empirical data, a natural question is "how does the brain reach or switch between these states using regimens of internal or external control?" Optimal control theory provides a powerful and flexible set of tools to explore these questions under various constraints and at different spatiotemporal scales. For example, in human neuroimaging, it is common to study different brain activity states and their relation to cognitive function. By leveraging control energy, NCT enables researchers to study how the topology of the brain's underlying structural connectivity enables it to transition between activity states that are relevant to (i) different cognitive functions[52] or (ii) different conditions within a single cognitive domain[55]. The second set of questions stems from empirical evidence demonstrating that individual and groups of brain regions (i) may be important for enabling specific functions, such as visual processing[56], motor processing[57] and cognition[58,59]; (ii) may be important for supporting critical functional processes in the brain, such as segregation and integration[60,61]; and (iii) may be disproportionately affected by disease processes[62,63]. Given these data, a natural question to ask is "what is the contribution of these sets of regions to the control of brain activity?" Average controllability measures the magnitude of propagation of stimulation along neural tracts. In turn, average controllability provides a coarse-grained understanding of individual regions' control over brain activity, insofar as it probes their general importance to guiding diverse ensembles of state transitions. This general nature stands in contrast to control energy, which requires that specific state transitions be defined a priori.

Applications

The analysis of brain data using a network representation is increasingly popular in neuroscience, and researchers have used a wide range of connectomic data to perform NCT analysis. For example, the availability of multimodal neuroimaging data in large cohorts accompanied by clinical and cognitive data[23,24,64–66], as well as indices of neurobiology not measurable in vivo (e.g., high-resolution histology[67] and gene expression[68]), enable researchers to validate NCT against brain function and biology, as well as examine differences between individuals. Indeed, we have applied NCT in our research with a view toward achieving these goals. Here, we briefly review selections of this work to show how our protocol may be applied to study the brain. Specifically, we discuss how model outputs from NCT link to network topology; explain differences between individuals in mental health symptoms, cognition and age; predict the effects of neurostimulation; explain switching between functional task states; and link to neuroanatomy. Note that unless otherwise specified, all the specific studies described below used NCT to examine the structural connectome.

Understanding the influence of topology

In our early work, we began by contextualizing nodal controllability statistics against what we know about connectome topology from graph theory. Specifically, Gu et al.[10] examined how nodal control properties—specifically average controllability and modal controllability—correlated with nodes' strength (the sum of a node's edge weights). Gu et al.[10] found that nodes' strength correlated strongly positively and negatively with average and modal controllability, respectively. These relationships were conserved across both humans and macaques. Collectively, these results indicate that a node's local topological importance predicts its capacity to control the dynamics of a system.

We have also examined how connectome topology influences the control energy associated with specific state transitions. Betzel et al.[52] found that nodes' topological importance predicted their capacity to facilitate transitions between eight canonical brain states (seven resting-state cortical networks as well as a subcortical network[69]). Specifically, Betzel et al. found that target states that intersected with the brain's rich club[61], a set of highly interconnected nodes that form the connectome's core, exhibited low transition energy. This result demonstrates that the rich club is well positioned in the network to act as an efficient target state to which a diverse set of initial states can transition with low control energy. Thus, the topology of the human connectome may be optimized to guide dynamics toward the rich club, bolstering the idea that these nodes support functional integration[60,70,71].

Given these advances in understanding how connectome topology contributes to control, we subsequently analyzed what the underlying control equations could tell us about network topology. Starting from the NCT equations, Kim et al.[53] derived the features of network architecture that were the most important for determining control energy. Kim et al.[53] discovered that a strong and diverse set of connections from stimulated nodes to unstimulated nodes were the leading-order contributors to the control cost. Using this discovery, the authors reduced the cost of controlling connectomes in *Drosophila*, mice and humans by virtually resecting edges and developed a method to meaningfully compare the control cost between different species and connectomes. These results provide simple and quantitative knowledge about the most important features of topology according to NCT.

Individual differences

Although the strong correlation between average controllability and strength reported by Gu et al.[10] may seem to imply redundancy between nodal controllability statistics and measures from graph theory incorporating weighted degree, we note that this correlation was spatial (i.e., across the brain) rather than between subjects. In subsequent work examining individual differences, Parkes et al.[21] compared the capacity of average controllability and strength to predict psychosis spectrum symptoms using out-of-sample testing. Parkes et al.[21] found that average controllability significantly outperformed strength in this predictive task and demonstrated that this improved performance was concentrated in higher-order default mode cortex[72]. These results show that although high average controllability may depend upon high strength, there exists unique inter-individual variation between the metrics and that this variance in average controllability couples more tightly to mental health symptoms. Furthermore, these results support the use of NCT in population-based predictive modeling studies.

We have also shown that average controllability exhibits robust developmental and sex effects. Average controllability increases between the ages of 8 and 22 years[12] and is higher in females in the cortex but higher in males in the subcortex[19]. Furthermore, Tang et al.[12] showed that age effects were strongest in nodes with higher controllability, underscoring the developmental importance of nodes that are well positioned in the network to control dynamics. When examining control energy, Cui et al.[73] demonstrated that the amount of energy required to activate the fronto-parietal system, a brain network thought to support executive function[74], was negatively correlated with both age and executive function in the same sample. This result suggests that the developmental emergence of executive function is associated with increased efficiency of neural signaling within the human connectome.

Predicting stimulation effects

An application of NCT that has clear translational impact is modeling the relationship between brain structure and function. To this end, we have examined whether NCT can predict the brain's functional response to neurostimulation from its structural connectome. For example, in patients with epilepsy, Stiso et al.[14] found that NCT was able to predict electrophysiological neuronal responses (measured with electrocorticography) after direct electrical stimulation. This result shows that our model, wherein neural activity is simulated upon the structural connectome, explains variance in experimentally manipulated empirical changes in brain state.

We have also examined NCT in the context of noninvasive neurostimulation techniques. In a pair of studies, Medaglia et al.[16,17] delivered transcranial magnetic stimulation to the left inferior frontal gyrus between repeated sessions of a set of language tasks. Across both studies, the authors found that NCT metrics extracted from the left inferior frontal gyrus explained variance in changes to task performance before and after transcranial magnetic stimulation. These results demonstrate that NCT can be used to probe the network mechanisms that underpin how neurostimulation elicits changes in behavior.

Modeling switches between functional brain states

In addition to predicting the effects of neurostimulation, NCT can be used to investigate how the topology of the structural connectome supports transitions between empirically observed

functional brain states. Our group has studied this process using brain states derived from fMRI. Cornblath et al.[20] clustered resting-state fMRI (rs-fMRI) into brain states representing instantaneous co-activations among canonical brain networks and used NCT to model the energy required to transition between those states. Using a series of null network models, Cornblath et al.[20] found that the topology of the structural connectome was wired to support efficient switching between brain states. This result demonstrates that the topology of the connectome is optimized to support dynamic fluctuations in resting-state activity.

Subsequent work by Braun et al.[55] examined transitions between brain states elicited by a working memory task. Braun et al.[55] found that transitioning from a 0-back brain state to the more cognitively demanding 2-back brain state required more energy than the reverse transition, demonstrating an asymmetry in control energy. Braun et al.[55] also found that this energy asymmetry was more pronounced in patients with schizophrenia than in healthy control subjects. Thus, while connectome topology may be set up to enable low-cost fluctuations in resting state[20], activating cognitively demanding brain states may require more control effort. Furthermore, this increased control effort appears to scale with within-task differences in cognitive demand and is further elevated in psychopathology.

Biologically informed NCT

Neuroscience is increasingly moving toward a multi-scale approach that seeks to understand how features of the brain observed at one scale link to properties observed at another, and vice versa[3,75–83]. Recently, we have applied this multi-scale approach to NCT by examining how dynamics within the model are influenced by variations to regions' cellular composition. Specifically, we examined how regions' profiles of cytoarchitecture affected the energy associated with state transitions that spanned the cortical hierarchy[22] (i.e., the sensory-fugal axis[84]). We found that state transitions traversing bottom up along the cortical hierarchy of cytoarchitecture required lower control energy to complete compared to their top-down counterparts, and we observed that nodes' position along this hierarchy predicted their importance in facilitating these transitions. This result shows that spatial variations in cortical microstructure constrain macroscopic connectome topology; this effect is consistent with work from neuroanatomy that describes a precise relationship between regions' profiles of cytoarchitecture and their extrinsic connectivity[85].

In recent work from outside our group, Luppi et al.[34] characterized how the control energy associated with a large set of activity maps derived from NeuroSynth[86] related to cognition. In addition, the authors examined how these transition energies varied when they used a broad range of neurotransmitter density maps to modify the control weights. This work ties together switching between functional brain states and biologically informed connectome analysis to provide the field with a comprehensive 'look-up table' of how regions' diverse biology affects control energy.

Alternative methods

We consider NCT with respect to other models that also seek to understand how communication unfolds within a structurally interconnected complex system. For neuroscientists, NCT complements both more complex biophysical models of dynamics and graph-theoretical measures of inter-nodal communication. Although both of these approaches model communication, they differ in their biological plausibility and complexity. Biophysical models aim to capture neuronal communication by distilling the various biophysical processes necessary for functional activity into separate model parameters. These parameters are tuned to simulate biologically plausible nonlinear dynamics within and between neurons at multiple scales. For example, at the scale of single neurons, the Hodgkin-Huxley model is concerned with modeling neuronal spiking activity[87] and is based on parameterizing the flow of sodium and potassium ions across the cell membrane. At the next scale up, mean-field models focus on the collective activity patterns of co-located populations of neurons[88,89]. Coupling multiple mean-field models together—where each model represents distinct neuronal populations—enables researchers to study how nonlinear dynamics emerge from brain structure at the macroscale[88].

In turn, this approach gives rises to a wide range of complex dynamical behaviors, including synchronized oscillators[38], learning[90,91], large-scale traveling brain waves[92] and structure–function coupling[93–95]. Broadly, NCT trades biophysical accuracy and the complexity of specific model behaviors for more power in designing and studying stimuli. For example, in lieu of studying state transitions that emerge from different models of associative memory[96] and context integration[97], NCT allows us to design specific stimuli to transition the model system to states that are known to be important for memory and cognitive control under specific constraints[52].

By contrast, graph-theoretic approaches instantiate relatively simple models of inter-nodal communication that rely on assumptions such as shortest-path routing, spatial proximity, random walks and diffusion processes[4,9]. Although these assumptions are an oversimplification of brain dynamics and are thus less biologically plausible, their simplicity confers greater analytic tractability and scalability, which are both desirable features when studying the human brain. This benefit compounds when the goal of a given study is to examine inter-individual differences, wherein dynamical models may be fit to thousands of participants. As such, despite their relative simplicity, graph-theoretic approaches have deepened our insights into large-scale brain organization[58,59,98–101], improved our understanding of the link between the brain and mental disorders[62,102–104] and helped elucidate the link between structure and function[43,47,48,105–107].

We consider NCT as situated between these two modeling approaches. As discussed in Development of the procedure, NCT is essentially a model of two parts, dynamics and control. For the former component, NCT models dynamics according to a diffusion process; thus, like graph theory, NCT makes simplifying assumptions of inter-nodal communication, which confers the advantages of analytic tractability and scalability. However, the second component, control, adds an additional layer of model parameterization that allows researchers to probe how the system might behave under different contexts (e.g., in response to task manipulations, cognitive control or neurostimulation protocols). This added flexibility brings NCT closer to biophysical modeling, insofar as they both seek to understand how the dynamics of a system respond to external perturbation. Indeed, we have shown that NCT can be used to predict changes in the dynamics of coupled Wilson-Cowan oscillators after simulated stimulation[18], suggesting that NCT can explain some of the behaviors engendered by nonlinear biophysical models. We have also shown that the underlying diffusion model is able to predict empirical neural responses measured with fMRI and electroencephalography[108], suggesting that NCT is well positioned to explain the types of data that are typically acquired to study the human brain.

Limitations and ongoing development

NCT can flexibly accommodate many scientific questions and generate concise knowledge from a simple model. However, NCT also has several limitations for the study of high-dimensional complex systems, such as numerical stability of algorithms, validation against empirical data, approximations of complex interactions and interpretation of model parameters.

Numerical stability of optimization

One limitation is the numerical stability of Procedure 1 under certain parameter conditions, which arises from ill-conditioned matrices that are built while solving for the control signals. This issue occurs most frequently when using a relatively small control set—a small m in the $N \times m$ matrix B—to control a network with large N. It is intuitive that precisely controlling the initial and target states of the whole brain from only a few nodes is difficult. In light of this limitation, it is crucial that the researcher carefully study the generated trajectories of the neural activity to ensure that the desired initial and target states are reached and that the numerical integrator does not generate a warning of numerically ill-conditioned matrices (see the Troubleshooting section). In the event that the control set must be small for the purposes of the research question, one solution may be to extend the control set by heavily weighting the desired control nodes and lightly weighting the remaining nodes[14]. Another option is to use Procedure 2 to study the average controllability of the control set.

Validation of NCT against empirical data

A second limitation is the validation of the model at the level of individual neural states. Phrased another way: given a connectome, A, and stimulations $\boldsymbol{u}(t)$ delivered to brain regions B starting at neural activations $\boldsymbol{x}(0)$, does experimentally measured neural activity agree with the simulated trajectory $\boldsymbol{x}(t)$? The challenges associated with addressing this sort of question extend far beyond NCT and to a significant portion of systems and network neuroscience. Microstate validation between neural structure and activity is most evident in small systems of neural circuits[109], but how to perform similar validations for large-scale systems such as the human brain remains an open area of research. Challenges include (i) the multiple possible scales of constructing brain networks[2,110], (ii) differing measures of inter-areal connectivity[111,112], (iii) multiple definitions of simulated neural activity[88,113–115] and (iv) the diverse spatial and temporal resolutions at which we can record whole-brain activity[54]. Along this active area of research, we have demonstrated that linear models outperform nonlinear and kernel-based models in both one-step prediction and model complexity for both fMRI and electroencephalography data[108], as well as correspondence between control energy and local metabolism[116]. We have also shown that NCT is able to explain variance in neuronal responses recorded using electrocorticography data[14], but more work is needed to fully validate NCT against empirical data.

Linear dynamics

A third limitation is the assumption of linear dynamics, which enables the calculation of powerful measures such as optimal control trajectories but hinders the biophysical realism of the framework. More sophisticated nonlinear models capture complex dynamics from individual neurons[87] to neural populations[39], thereby enabling the study of fine-grained experimental behavior[117] and complex nonlinear phenomena[118]. These models make fewer simplifying assumptions to capture nonlinear behaviors of biological systems such as complex memory landscapes[119]. Although prior work has shown that linear models outperform several classes of nonlinear models in describing and predicting brain-wide neural activity[108], extensions of NCT to nonlinear systems will enable greater flexibility to accommodate and explore the impact of nonlinear biophysical constraints. Although the theory of nonlinear control is an active area of research[120], there are immediate applications of NCT to nonlinear systems and many exciting potential extensions of NCT to capture more biophysical realism.

Broadly speaking, the linear dynamics assumed by NCT can be thought of as being valid for a nonlinear system within small deviations of an operating state[7]. Hence, the most immediate application of NCT to nonlinear systems is to linearize the model about an operating point, such as the upright position of an inverted pendulum[121]. Along these lines, the next immediate generalization to NCT is to linear time-varying systems[50], in which the model is linearized not about a point, but about a trajectory. Although methods to implement control for linearized and time-varying systems are well established in the control community, a biophysically meaningful implementation and interpretation of the parameters—namely $A(t)$ and $B(t)$—remains an area of active work[122]. Another approximation that is particularly relevant for high-dimensional neural systems is at the limit of weakly coupled oscillators[123,124], whereby a high-dimensional system of oscillators with weak interactions can be reduced to a low-dimensional phase-response curve, allowing for the potential linearization of the system about phase-locked states.

In addition to linearizing dynamics about points and trajectories, NCT can also meaningfully be applied to nonlinear dynamical systems that can be made linear through a nonlinear change of variables. One such example is by using finite-dimensional Koopman subspaces, which allow for the recasting of nonlinear systems with single fixed points as higher-dimensional linear systems[125], and closely related methods in dynamic mode decomposition[126]. Furthermore, advances in nonlinear control enable us to probe important coarse-grained questions such as the control set necessary to push nonlinear systems between attracting states[127]. Other control strategies take advantage of the ability of nonlinear systems to access states that lie outside of their linearization[128].

Experimental design

The adjacency matrix

The goal of an NCT analysis, as it is conceptualized in this protocol, is to understand how the topology of the structural connectome supports and constrains spreading dynamics and to what extent those dynamics can be controlled. Thus, core to this analysis is the acquisition of one or more structural connectomes from a model organism. For input to our protocol, we represent a given connectome as an adjacency matrix, A. In A, the N nodes of the system are stored on the rows and columns, and the $N \times N$ edge values are stored in the entries (for a discussion of size limits for A, see Supplementary Information, Theoretical and practical limitations on system size). Both the nodes and the edges of A can be defined in numerous ways. For example, the nodes of the system may be defined as single neurons in organisms such as *Caenorhabditis elegans*[51,129] or as brain regions of varying size and definition in organisms such as the mouse[77], *Drosophila*[53], macaque[130] and human[4]. The edges of A may be defined as either the directed or undirected connectivity between nodes. In humans, we typically extract structural connectomes from diffusion-weighted imaging (DWI) sequences obtained using MRI. Tractography algorithms are applied to DWI scans to model the white matter pathways intersecting pairs of brain regions, which are then used to populate connectome edges with the number of those pathways (e.g., the streamline count)[4]. This example constitutes a weighted undirected connectome upon which NCT can be conducted. Critically, our model assumes that A_{ij} encodes the strength of diffusion of activity along the edge connecting node j to node i. In other words, our model assumes that the columns of A store the source nodes (i.e., projections from node j), while the rows store the target nodes (i.e., projections to node i). Although this distinction is irrelevant for undirected connectomes in which $A_{ij} = A_{ji}$, it is crucial for directed connectomes, and researchers must ensure that their directed A matrix conforms to the above assumptions. Relatedly, although A may be dense or sparse, special care must be taken if A is not one connected component (i.e., node i cannot reach every node j through a path of any length). Disconnected components will not be able to influence each other, and control signals cannot propagate between them, such that every disconnected component must receive its own subset of inputs. Overall, NCT can be flexibly applied to directed and undirected connectomes, dense and sparse connectomes and human and non-human connectomes.

The best choice of connectome will depend upon researchers' goals and available data. If researchers are primarily interested in studying a specific transition or set of transitions, then a group-averaged connectome is preferable. In this case, a group-averaged connectome will allow researchers to closely study how topology constrains their transitions of interest while minimizing the sources of noise present in subject-specific connectomes (see below). Furthermore, if the data are available, a directed structural connectome is preferable to an undirected structural connectome, because the former will yield state transitions with greater biological plausibility. On the other hand, if researchers are interested in studying individual differences, then it will be necessary to use subject-specific connectomes. In this case, as with other techniques, replication across multiple datasets is desirable to ensure generalizability of findings. We focus on group-averaged structural connectomes in this protocol and provide examples of applying NCT to both directed and undirected edge weights.

Given that connectomes are central to the application of NCT, any artefacts present in the connectomes will be reflected in model outputs. For example, connectomes populated by DWI estimates of connectivity are known to contain false positives and false negatives, which may be partly mitigated by the use of thresholding techniques[111,131]. In-scanner head motion is well known to spuriously affect these estimates of connectivity as well[132,133]. Finally, connectome reconstruction is affected by acquisition parameters[134] and benefits from parameters that result in higher signal-to-noise ratio and angular resolution. Thus, the accurate generation and rigorous quality control of connectomes are both crucial considerations for experimental design. For human connectomes, we recommend that researchers consult the extant literature on the acquisition, processing and quality control of DWI scans[111,134–136] (see also https://qsiprep.readthedocs.io/).

Another consideration for connectome estimation is the brain parcellation used to define system nodes. If, as mentioned above, structural connectivity is determined by streamline count, then variations in the size of regions across the parcellation will bias connectome edge weights; larger brain regions will intersect with more white matter pathways and thus show higher connectivity to the remaining regions. As with any analysis of graph topology, this bias will affect the outputs of NCT; for example, larger regions may show higher average controllability just by virtue of being more directly connected to the system. It is for this reason that we recommend that researchers reproduce their results using several different parcellation definitions and resolutions. Doing so ensures that their results are not driven by a specific parcellation choice.

Defining a time system

Once an adjacency matrix (A) has been defined, the next decision is whether to model the linear dynamical system in discrete or continuous time. As discussed in Development of the procedure, in a discrete-time system, the states of the system, $\boldsymbol{x}(t)$, evolve forward in time according to a set of discrete steps ($\boldsymbol{x}(t) \rightarrow \boldsymbol{x}(t+1)$). In a continuous-time system, the states of the system are continuously changing in time ($\boldsymbol{x}(t) \rightarrow \boldsymbol{x}(t) + \dot{\boldsymbol{x}}(t)\,dt$). The choice of time system will depend upon the research question and affects all subsequent analyses because of differences in the mathematical implementation of NCT under each system. We refer the reader to Karrer et al.[8], Kim and Bassett[7], Hespanha[50] and other texts in linear systems theory for extended discussion.

One guiding principle for choosing a time system can be seen in the following simple example. Consider a 1D discrete-time system $\boldsymbol{x}(t+1) = -\boldsymbol{x}(t)$ starting at an initial condition of $\boldsymbol{x}(0) = 1$. This system evolves by jumping between 1 and −1 without visiting any intermediate value. By contrast, the state of a continuous-time system can transition between 1 and −1 only by smoothly visiting all values in between. The former process is reminiscent of 'all-or-nothing' phenomena such as neuron spiking, whereas the latter process is closer to population-level average activity in macroscopic connectomes. As such, when studying the macroscopic human connectome, we generally recommend the use of continuous-time systems. Note that one seemingly reasonable justification for using a discrete-time system would be to model fMRI activity, which is sampled at discrete time intervals on the order of ~1 s. However, care must be taken with this line of reasoning because the discretization of a linear continuous-time system at some sampling interval Δt produces an equivalent discrete-time system matrix with non-negative eigenvalues through the conversion $A_{\text{discrete}} = e^{A_{\text{continuous}}\Delta t}$. This conversion exponentiates all eigenvalues such that they are non-negative. As such, using an adjacency matrix with negative eigenvalues in a discrete-time system cannot truly be seen as a discrete sampling of a linear time-invariant continuous-time system. In general, we recommend continuous-time systems as the default choice. If researchers choose to use discrete-time systems, then we suggest that they replicate their primary results using continuous-time systems.

Normalizing the adjacency matrix

Once a time system is chosen, A needs to be normalized before analysis. If a discrete-time system is chosen, we normalize A according to the following equation:

$$A_{\text{norm}} = \frac{A}{|\lambda(A)|_{\text{max}} + c}.$$

Here, $|\lambda(A)|_{\text{max}}$ denotes the largest absolute eigenvalue of the system. In addition, c is a user-defined input parameter that determines the rate of decay of system dynamics. We set $c = 1$ by default, which ensures that all modes of the system decay and thus that activity goes to zero over time (note that this is true of any positive c value). This normalization ensures that internal dynamics decay in a manner that is necessary for the stabilization of the system. Specifically, the largest absolute value of a matrix's eigenvalues is called the 'spectral radius', and this normalization ensures that the spectral radius is <1: a condition known as Schur stability.

Intuitively, a discrete-time system given by Eq. 1 with no input (i.e., $\boldsymbol{u}(t) = 0$) will evolve as $\boldsymbol{x}(n) = A^n\boldsymbol{x}(0)$, and the most unstable eigenmode of the system will evolve as $\lambda(A)^n_{\max}$. To ensure that this mode does not grow infinitely with n, it must have a magnitude <1.

If a continuous-time system is chosen, we normalize A according to the following equation:

$$A_{\text{norm}} = \frac{A}{|\lambda(A)|_{\max} + c} - I.$$

Here, I denotes the identity matrix of size $N \times N$. As above, we normalize such that the spectral radius is <1, but we take the additional step of subtracting the identity. This step exists because a continuous-time system given by Eq. 2 with no input will evolve as $\boldsymbol{x}(t) = A^{At}\boldsymbol{x}(0)$, and the eigenmodes of the system will evolve as $e^{\lambda_i t}$. Hence, for the system to decay, all λ_i must have a negative real component, which is achieved through the subtraction of I.

Apart from the above approaches, there are alternatives to normalizing A. One could divide different edges of A by different amounts instead of uniformly by $|\lambda(A)|_{\max} + c$. One could also subtract the diagonals of A by some other matrix besides the identity (for continuous-time systems). If there exists sufficient empirical evidence to use more specific normalization parameters, then they may be used. For example, the timescales of regions' neuronal dynamics are known to vary across the brain[137–140], and this variance could be incorporated into normalization by subtracting non-uniform values from the diagonal of A. In this case, for a continuous-time system, subtracting larger values from the diagonal of A will yield dynamics that decay more quickly, akin to relatively fast neuronal dynamics. What we present here is a simple method for stabilizing A in the absence of additional empirical constraints. Key properties that are preserved after our normalization approach are (i) the rank ordering of the eigenvalues of A and (ii) the eigenspaces (eigenvectors) of those modes; this preservation is a direct result of uniform normalization by a scalar ($|\lambda(A)|_{\max} + c$) and subtraction of the identity matrix. Thus, this minimal normalization approach preserves the necessary properties to run NCT analysis while also maintaining comparability between studies.

Defining brain states

Beyond the core requirement of a connectome, the flexibility of NCT makes it applicable to a broad range of experimental designs (see Applications); the most critical component is that researchers have hypotheses that pertain to studying the control of brain dynamics. However, in the case of Procedure 1 (control energy), in which researchers will study the control signals, $\boldsymbol{u}(t)$, there are some additional considerations. Specifically, to analyze state transitions, researchers need to provide a pair of brain states relevant to their hypotheses. Providing these states allows NCT to find the control signals, $\boldsymbol{u}(t)$, that are required to transition between them and to summarize those control signals as control energy.

Brain states can be defined in a number of ways. The simplest approach is to define each brain state as a binary vector, in which nodes that are within a given state are assigned an arbitrary constant value (e.g., 1), and any remaining nodes are assigned a value of 0. In this setup, NCT is tasked with transitioning the brain between actuating different sets of nodes to a constant arbitrary level of neural activity. An alternative approach is to allow brain states to represent a variable pattern of activity. As mentioned above, Cornblath et al.[20] modeled the energy required to transition between brain states derived from clustering of fMRI data, while Braun et al.[55] used task-activation maps extracted from an fMRI contrast. These approaches allowed the authors to generate state vectors that encode non-zero activity across all nodes of the system. The choice to study binary or non-binary brain states depends on researchers' goals and available data. For example, using non-binary brain states confers the advantage of studying empirical neural activity, which may lead to more biologically plausible state transitions. By contrast, binary brain states are typically more easily defined because they do not depend on the presence of functional data. In addition, binary states confer the advantage of studying more focused transitions on the connectome because they involve guiding activity from one focal point on the network to another, simulating system-to-system communication. Thus, generally, if researchers wish to maximize the biological plausibility of their state transitions,

then we recommend that they use non-binary brain states derived from empirical neural data (e.g., fMRI). If instead researchers wish to study transitions from one specific point on the network to another (e.g., between canonical brain systems[52] or between specific node types[22]), then we argue that binary brain states are the better choice. In either case, if researchers choose to study subject-specific states (which is possible for both binary and non-binary states), then they must take particular care to ensure that their results are not confounded by issues of state definition (see below). In this protocol, we illustrate examples using both binary (see Procedure 1) and non-binary (see Supplementary Information, Variations to Procedure 1) brain states, both defined at the group level.

Differences in brain states' magnitude will affect control energy, potentially necessitating the normalization of state magnitude. For example, if researchers are examining transitions between patterns of brain activity (e.g., using functional data as in refs. 20,55), then differences between states' mean activity will affect control energy; assuming a common initial state, target states with higher activity will require more energy to transition to compared to target states with lower activity. This effect generalizes to binary brain states as well. In this case, differences in state size (i.e., the number of regions in each state) constitute differences in state magnitude; transitioning to larger target states will require more energy. If there are differences in state magnitude, we recommend normalizing states before computing control energy (see Step 3 in Procedure 1). Note that the need for this normalization will depend upon researchers' analyses. For example, if researchers are studying individual differences in the energy associated with a single transition, then normalization may not be necessary as long as state definition is consistent across subjects. What is critical is that researchers consider what comparisons they want to make and whether variations in state definition would confound those comparisons.

Defining a control set

In addition to brain states, for Procedure 1, researchers also need to designate a control set; these are the nodes that NCT will use to complete state transitions. As discussed in Development of the procedure, the $N \times m$ control set defines the extent to which the nodes of the system can affect changes in its dynamics. In turn, the definition of B determines the dimensions of $\boldsymbol{u}(t)$; the greater the number of control nodes, the more independent control signals will be generated. In our work, we have often deployed a uniform full control set, which means that all of the nodes of the system are designated as controllers (full), and all are given equivalent control over dynamics (uniform). In this case, $m = N$. Intuitively, this approach assumes that the entire brain is being controlled—either internally or externally—when completing a state transition. However, depending on a researcher's hypotheses, this assumption may not be appropriate. Instead, researchers may want to define only a subset of nodes as controllers (e.g., refs. 52,53) or assign variable weights to control nodes (e.g., refs. 22,32–34), or both. Note that assigning variable control weights serves to give some nodes more control over system dynamics than others. In any case, it is critical that researchers check whether their designated control set was able to complete the associated state transition (see Procedure 1, Step 5); successful completion of a state transition is not guaranteed in the model, and completion is less likely when transitions are driven by a small control set.

Expertise needed to implement the protocol

We provide open-source and broadly accessible tools that implement optimal control and average controllability in a Python-based software package called 'network control theory for python' (nctpy). In nctpy, we provide a flexible implementation that enables researchers to make model assumptions that best fit their research question. As a result, while a full understanding of linear systems and optimal control theory are not required, the researcher must have enough expertise to make key modeling decisions that best represent the data (see above).

The first piece of expertise needed is to understand the differences between (and implications of) discrete-time systems and continuous-time systems. As discussed above, this difference is not merely a conceptual one, because discrete-time systems display a fundamentally different set of behaviors than continuous-time systems. That is, a discrete-time system is not simply a temporally coarse-grained version of a continuous-time system. Instead, each system exhibits different dynamics.

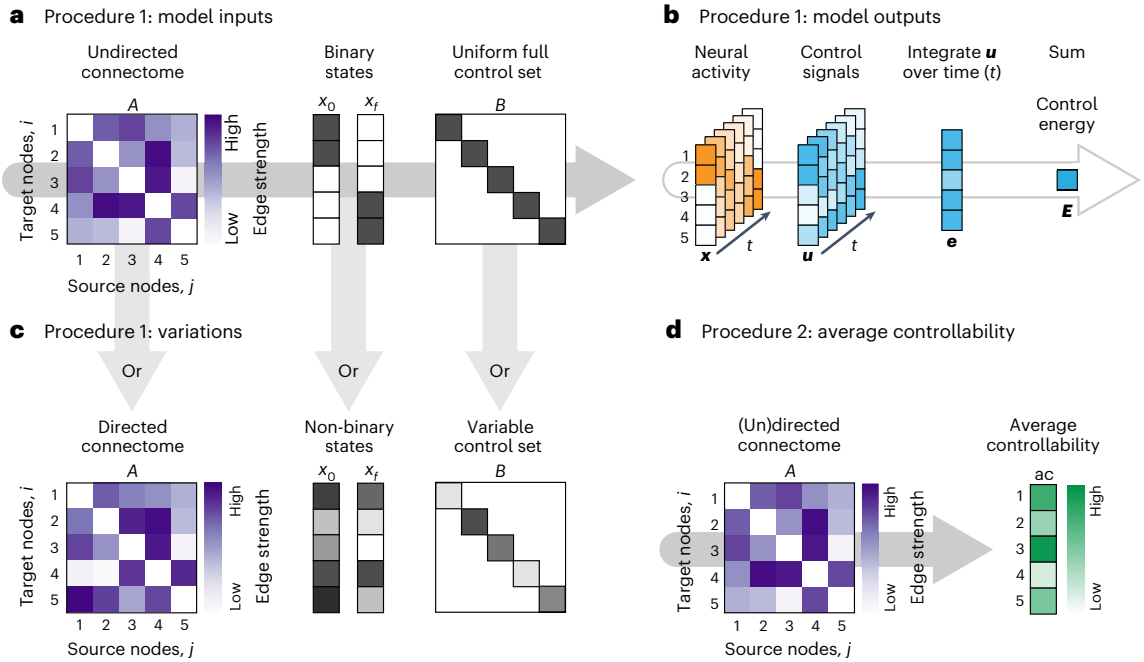

Fig. 3 | Schematic representation of the NCT protocol. Our protocol is split into two procedures. Primarily, our protocol focuses on modeling the control energy associated with user-defined control tasks. We refer to this part of our protocol as Procedure 1 (**a**–**c**). Procedure 1 will be of interest to researchers who seek to study specific state transitions. We also outline a brief protocol for estimating nodes' average controllability. We refer to this part of our protocol as Procedure 2 (**d**). Procedure 2 will be of interest to researchers who want to examine nodes' general capacity to control system dynamics. **a**, Inputs required for Procedure 1. To compute control energy, researchers must provide a structural connectome (A), an initial state ($\boldsymbol{x}_0$) and a target state ($\boldsymbol{x}_f$) and must also define a control set (B). **b**, Model outputs from Procedure 1. Given these inputs, our protocol will output the state trajectory (neural activity, $\boldsymbol{x}(t)$) and the control signals ($\boldsymbol{u}(t)$). Once inspected, the control signals can be integrated over time to obtain node-level energy ($\boldsymbol{e}$), which in turn can be summed over nodes to get the control energy ($\boldsymbol{E}$). **c**, Variations to Procedure 1. Procedure 1 can handle a diverse range of inputs, including but not limited to undirected and directed connectomes (left), binary and non-binary brain states (middle) and control sets with uniform or variable weights (right). **d**, Procedure 2: average controllability. Procedure 2 requires only a structural connectome (A) as input and will return the average controllability of each node. This metric quantifies the impulse response of the system from a given node. Higher average controllability indicates that a node is better positioned in the network to propagate dynamics.

The second piece of expertise needed is to understand the nature of Procedure 1 (control energy) and Procedure 2 (average controllability) to interpret the outputs. Procedure 1 solves an optimization problem. Specifically, we first provide a model of the dynamics (i.e., A, B), the initial and target states and some optimization parameters. Then, we solve for the control signals, $\boldsymbol{u}(t)$, that minimize the cost. Hence, all interpretations of $\boldsymbol{u}(t)$ should be made with the understanding that they were determined by the user-defined optimization parameters. Procedure 2 does not solve an optimization problem and thus does not receive any optimization parameters. Rather, it measures the magnitude of the neural states over time as a result of an impulse stimulation. Because Procedure 1 and Procedure 2 use the same dynamics but output different quantities through different means, more expertise in linear systems and optimal control is needed to meaningfully compare and contrast the two pathways.

Overview of the procedure

As discussed above, we split our protocol into two procedures (Fig. 3). The primary procedure of our protocol focuses on computing control energy (Procedure 1). This procedure is illustrated in Fig. 3a–c. Briefly, Fig. 3a outlines the inputs required to run Procedure 1, Fig. 3b outlines the corresponding model outputs and Fig. 3c outlines some of the variations to model inputs that we have discussed thus far. Note that example implementations of all these variations are presented in the Supplementary Information (see Variations to Procedure 1). The second procedure focuses on computing nodes' average controllability (Procedure 2; Fig. 3d). Both of these procedures are underpinned by the same set of core steps, which are (i) defining a

time system and (ii) normalizing the adjacency matrix. After these core steps, the protocol bifurcates, and most of the subsequent steps focus on computing control energy. Broadly, computing control energy requires (i) defining a control task, (ii) computing control signals and (iii) summarizing those control signals into control energy. By contrast, computing average controllability is done in a single step after the two core steps.

Materials

Equipment

- A computer with Python (tested on version 3.9) and nctpy installed alongside its dependencies. This protocol has been tested on Mac OS running on Intel Core i5/i7/i9 processors as well as on Apple Silicon. We have also tested this protocol on Linux Ubuntu running on Intel processors. Random access memory (RAM) requirements will vary depending on researchers' data and analyses, but we recommend ≥16 GB. Finally, we recommend installing nctpy inside a virtual environment managed by Anaconda (https://www.anaconda.com/). The following core dependencies are required to run nctpy:
  - numpy (https://numpy.org/), tested on version 1.24.3
  - scipy (https://scipy.org/), tested on version 1.10.1
  - tqdm (https://github.com/tqdm/tqdm), tested on version 4.65

  In addition, there are some functions in `nctpy.plotting` and `nctpy.utils` that require the following:
  - statsmodels (https://www.statsmodels.org/), tested on version 0.13.5
  - matplotlib (https://matplotlib.org/), tested on version 3.7.1
  - seaborn (https://seaborn.pydata.org/), tested on version 0.12.2
  - nibabel (https://nipy.org/nibabel/), tested on version 5.1
  - nilearn (https://nilearn.github.io/), tested on version 0.10.1

  Finally, the following optional packages were used to run the analyses illustrated in this protocols paper:
  - (Optional) pandas (https://pandas.pydata.org/), tested on version 1.5.3
  - (Optional) scikit-learn (https://scikit-learn.org/), tested on version 1.2.2

  See https://github.com/LindenParkesLab/nctpy for more details. Creating a Python environment using Anaconda and installing the above dependencies should take ≤30 min.

Input data

- Adjacency matrix, A (required for Procedure 1 and Procedure 2)
- Brain states, $\boldsymbol{x}_0$ and $\boldsymbol{x}_f$ (required for Procedure 1)
- Control set, B (required for Procedure 1)

Example dataset

Here, we primarily used undirected structural connectomes derived from DWI performed on the human brain. We obtained these connectomes from the Philadelphia Neurodevelopmental Cohort (PNC)[23,24], a community-based study of brain development in youths aged 8–22 years. The neuroimaging sample of the PNC consists of 1,601 participants. From this original sample, we retained 253 typically developing participants who had no medical co-morbidity or radiological abnormalities and who were not taking psychoactive medications at the time of assessment. In addition, these participants' T1-weighted, DWI and rs-fMRI scans all passed stringent quality control procedures[136,141,142].

- Structural connectome reconstruction was performed using QSIprep 0.14.2[135], which is based on Nipype 1.6.1[143]. Connectomes were extracted using the 200-node variant of the Schaefer parcellation[110], ordered according to seven canonical brain systems[69]. The strength of inter-regional connectivity was summarized using the number of streamlines that intersected

each pair of parcels. Connectomes were averaged over subjects. This group-averaged connectome was thresholded by retaining the edges that were present in ≥60% of participants' connectomes[131]. This process resulted in a final connectome with 98% edge density.

- rs-fMRI was also obtained from the same 253 PNC participants[23]. These data were used to generate empirical brain activity states to perform non-binary state transitions (see Supplementary Information, Variations to Procedure 1). The eXtensible Connectivity Pipeline[142,144] was used to post-process the outputs of fMRIPrep version 20.2.3[145]. The eXtensible Connectivity Pipeline was built with Nipype 1.7.0[143]. Processed rs-fMRI time series were extracted from the same 200-node parcellation mentioned above[110].
  We also studied a directed structural connectome obtained from the Allen Mouse Brain Connectivity Atlas[25,27] (see Supplementary Information, Variations to Procedure 1).
- Whole-brain structural connectomes were constructed with 2×10^5 voxels at a spatial resolution of 100 µm (see refs. 25,27 for more details). Voxels were assigned to regions (coarse structures) according to a 3-D Allen Mouse Brain Reference Atlas. Isocortex was further divided into six systems (auditory, lateral, medial, prefrontal, somatomotor and visual) on the basis of prior work that applied community detection to identify stable modules[26]. Connection strengths were modeled for all source and target voxels by using data from 428 anterograde tracing experiments in wild-type C57BL/6J mice[27]. Normalized connection strengths were obtained by dividing the connection strengths by the source and target region sizes. Here, we retained only the 43 isocortical regions. This process resulted in a fully connected directed structural connectome.

In all of the below code, we assume the existence of a Python environment with nctpy installed alongside its dependencies. First, we import all the functions we need to run our protocol:

```
# import
import os
import numpy as np
import pandas as pd
import scipy as sp
from scipy import stats
from scipy.spatial import distance
from sklearn.cluster import KMeans
from tqdm import tqdm
# import plotting libraries
import matplotlib.pyplot as plt
import seaborn as sns
from nilearn import datasets
from nilearn import plotting
# import nctpy functions
from nctpy.energies import integrate_u, get_control_inputs
from nctpy.pipelines import ComputeControlEnergy,
ComputeOptimizedControlEnergy
from nctpy.metrics import ave_control
from nctpy.utils import matrix_normalization, convert_states_str2int,
normalize_state, normalize_weights, get_null_p, get_fdr_p
from nctpy.plotting import roi_to_vtx, null_plot, surface_plot, add_
module_lines
from null_models.geomsurr import geomsurr
```

Note that depending on their goals, researchers may need only a subset of this import call. Next, we load a structural connectome as our adjacency matrix:

```
# directory where data is stored
datadir = '/path/to/data'
```

```
adjacency_file = 'structural_connectome.npy'
# load adjacency matrix
adjacency = np.load(os.path.join(datadir, adjacency_file))
n_nodes = adjacency.shape[0]
print(adjacency.shape)
(200, 200)
# check for self-connections
print(np.any(np.diag(adjacency) > 0))
True
# get density including self connections
density = np.count_nonzero(np.triu(adjacency, k=0)) / (n_nodes**2 / 2)
print(density)
0.9768
```

The above code demonstrates that our connectome comprises 200 nodes, includes self-connections (i.e., $A_{ii} > 0$) and has an edge density of 98%. See Supplementary Fig. 1 for control energy plotted as a function of edge density.

Procedure 1

Core steps

● **TIMING <5 s**

1. *Define a time system*. Determine whether to model the linear dynamical system in discrete or continuous time:

```
# determine time system. Note, delete the line below that is not needed.
system = "discrete"
# or
system = "continuous"
```

2. *Normalize the adjacency matrix*. Once a time system has been determined, normalize the adjacency matrix, A:

```
# normalize adjacency matrix
adjacency_norm = matrix_normalization(A=adjacency, system=system, c=1)
```

Irrespective of the chosen time system, the above step outputs `adjacency_norm`, which contains the structural connectome as a normalized adjacency matrix that is ready for NCT analysis. See Supplementary Fig. 2 for control energy plotted as a function of c (as well as other user-defined input parameters outlined in subsequent sections), which demonstrates that control energy remains stable with increasing c. In all of the code and results shown below, `adjacency_norm` was produced for a continuous-time system.

Define a control task

● **TIMING <5 s**

3. *Define a control task*. Define a control task that comprises an initial state, $\boldsymbol{x}_0$, a target state, $\boldsymbol{x}_f$, and a control set, B. Here, we illustrate an example control task that involves transitioning between a pair of binary brain states controlled by a uniform full control set. Our `adjacency_norm` is ordered according to seven canonical brain systems[110]. We leverage this grouping to define a state transition between the visual system and the default mode network. To begin, we set up a vector, `states`, that stores integer values denoting to which brain system each node belongs. That is, `states == 0` represents nodes that belong to system 1, `states == 1` represents nodes that belong to system 2, etc.

(A) **Create `states` from a list of strings that groups nodes into the aforementioned canonical brain systems (this file can be found here):**

```
# load node-to-system mapping
system_labels = list(
 np.loadtxt(os.path.join(datadir, "pnc_schaefer200_system_labels.txt"),
dtype=str)
)
print(len(system_labels))
200
print(system_labels[:20])
['Vis', 'Vis', 'Vis', 'Vis', 'Vis', 'Vis', 'Vis', 'Vis', 'Vis', 'Vis',
'Vis', 'Vis', 'Vis', 'Vis', 'SomMot', 'SomMot', 'SomMot', 'SomMot',
'SomMot', 'SomMot']
```

(B) **Use `convert_states_str2int` to convert this list of strings:**

```
# use list of system names to create states
states, state_labels = convert_states_str2int(system_labels)
print(type(state_labels), len(state_labels), state_labels)
<class 'list'> 7 ['Cont', 'Default', 'DorsAttn', 'Limbic',
'SalVentAttn', 'SomMot', 'Vis']
print(type(states), states.shape, states)
<class 'numpy.ndarray'> (200,) [6 6 6 6 6 6 6 6 6 6 6 6 6 6 5 5 5
5 5 5 5 5 5 5 5 5 5 5
5 5 2 2 2 2 2 2 2 2 2 2 2 2 4 4 4 4 4 4 4 4 4 4 3 3 3 3 3 3
0 0 0 0 0 0 0 0 0 0 0 0
0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 6 6 6 6 6
6 6 6 6 6 6 6 6 6 6 5
5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 2 2 2 2 2 2 2 2 2 2 2 4 4
4 4 4 4 4 4 4 4 3 3
3 3 3 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1]
```

(i) As can be seen from the above `print` commands, our `system_labels` variable comprises a string variable for every node in our system that denotes to which brain system that node belongs. `convert_states_str2int` takes that list of strings and returns an array of integers, `states`, with a corresponding list of labels, `state_labels`.

(C) **Extract x_0 and x_f by using the integers that correspond to the visual system (`'Vis'`) and the default mode system (`'Default'`):**

```
# extract initial state
initial_state = states == state_labels.index('Vis')
# extract target state
target_state = states == state_labels.index('Default')
```

(i) `initial_state` and `target_state` will be Boolean vectors, wherein `True` encodes the nodes that belong to a given state.

(D) **Normalize state magnitude:**

```
# normalize state magnitude
initial_state = normalize_state(initial_state)
target_state = normalize_state(target_state)
```

(i) This process will convert `initial_state` and `target_state` from Boolean entries to floating point numbers that have been normalized using the Euclidean norm of the vector. This normalization constrains state magnitude to a unit sphere (see Experimental design).

(E) **Finally, define a control set. Unlike the initial and target states, the control set is encoded along the diagonal of an *N*×*N* matrix, *B*. To define a uniform full control set, use the identity matrix:**

```
# specify a uniform full control set: all nodes are control nodes
# and all control nodes are assigned equal control weight
control_set = np.eye(n_nodes)
```

Compute control energy

● **TIMING ~1 min**

4. *Compute control signals and state trajectory*. After the definition of a control task, the next step is to find the control signals, $\boldsymbol{u}(t)$, that drive the system to transition between $\boldsymbol{x}_0$ and $\boldsymbol{x}_f$. $\boldsymbol{u}(t)$ will be an $m \times T$ matrix of m time-varying signals injected into the control nodes over a specified time horizon T. Here, because of our use of a full control set, $m = N$. Critically, injecting these control signals into a system whose initial state is encoded by $\boldsymbol{x}_0$ should result in a system whose final state is encoded by $\boldsymbol{x}_f$ at time T. Alongside the control signals, we also extract the state trajectory, $\boldsymbol{x}(t)$. The state trajectory, which will be an $N \times T$ matrix, is the time-varying pattern of simulated neural activity that unfolds as the system traverses between $\boldsymbol{x}_0$ and $\boldsymbol{x}_f$.

(A) **Compute $\boldsymbol{u}(t)$ and $\boldsymbol{x}(t)$:**

```
# set parameters
time_horizon = 1 # time horizon (T)
rho = 1 # mixing parameter for state trajectory constraint
trajectory_constraints = np.eye(n_nodes) # nodes in state
trajectory to be constrained
# get the state trajectory, x(t), and the control signals, u(t)
state_trajectory, control_signals, numerical_error = get_control_
inputs(
 A_norm=adjacency_norm,
T=time_horizon,
B=control_set,
x0=initial_state,
xf=target_state,
system=system,
 rho=rho,
 S=trajectory_constraints,
)
```

(i) By default, we set `time_horizon=1`. Note that this value is arbitrary and does not correspond to any real-world time units (e.g., seconds). Importantly, `get_control_inputs` uss a cost function that includes both the magnitude of the control signals and the magnitude of the state trajectory. The input parameter `rho` allows researchers to tune the mixture of these two costs while finding the input $\boldsymbol{u}(t)$ that achieves the state transition. Specifically, `rho=1` places equal cost over the magnitude of the control signals and the state trajectory. Reducing `rho` below 1 increases the extent to which the state trajectory adds to the cost function alongside the control signals. Conversely, increasing `rho` beyond 1 reduces the state trajectory contribution, thus increasing the relative prioritization of the control signals. Lastly, `S` takes in an $N \times N$ matrix whose diagonal elements define which nodes' activity will be constrained in the state trajectory. In summary, `S` designates which nodes' neural activity will be constrained while `rho` determines

by how much it will be constrained, relative to the control signals. Here, by setting `rho=1` and `S=np.eye(n_nodes)`, we are implementing what we refer to as 'optimal control'[11]. If `S` is instead set to an $N \times N$ matrix of zeros, then the state trajectory is completely unconstrained, a setup we refer to as 'minimum control'[20,22]. In this case, `rho` is ignored.

(ii) Alternatively, researchers may choose to constrain only a subset of the state trajectory by defining partial constraint sets. See Supplementary Fig. 2 for control energy plotted as a function of different combinations of `time_horizon` and `rho`. Also see here for a notebook outlining different use cases of `get_control_inputs`. However, to avoid issues associated with researcher degrees of freedom, we recommend that users implement either optimal control, by setting `rho=1` paired with their chosen `S` matrix, or minimum control, by setting `S=np.zeros((n_nodes, n_nodes))`.

(iii) In addition to `state_trajectory` and `control_signals, get_control_inputs` also outputs `numerical_error`, which stores two forms of numerical error. The first error is the inversion error, which measures the conditioning of the optimization problem. If this error is small, then solving for the control signals was well conditioned. The second error is the reconstruction error, which is a measure of the distance between $\boldsymbol{x}_f$ and $\boldsymbol{x}(T)$. If this error is small, then the state transition was successfully completed; that is, the neural activity at the end of the simulation was equivalent to the neural activity encoded by $\boldsymbol{x}_f$. We consider errors $<10^{-8}$ as adequately small.

(iv) Check the inversion and reconstruction errors:

```
# print errors
thr = 1e-8
# the first numerical error corresponds to the inversion error
print(
 "inversion error = {:.2E} (<{:.2E}={:})".
format(numerical_error[0], thr, numerical_error[0] < thr
 )
)
inversion error = 1.36E-15 (<1.00E-08=True)
# the second numerical error corresponds to the reconstruction
error
print(
 "reconstruction error = {:.2E} (<{:.2E}={:})".
format(numerical_error[1], thr, numerical_error[1] < thr
 )
)
reconstruction error = 5.16E-14 (<1.00E-08=True)
```

5. *Visualize state trajectory and control signals*. Once $\boldsymbol{x}(t)$ and $\boldsymbol{u}(t)$ have been derived, they should be visualized before computing control energy. Visualization provides intuition regarding how the model is behaving and is helpful for confirming that the state transition was completed successfully (see Fig. 4 and Box 1).
◆ **TROUBLESHOOTING**

6. *Compute control energy*. The final step is to summarize the control signals into control energy. This is done by numerically integrating the control signals over time. In this step, we use Simpson's rule—an extension of the trapezoidal rule that fits a polynomial through neighboring sets of points—to achieve this integration, yielding a vector of node-level energy.
(A) **Compute node-level control energy:**

```
# integrate control signals to get control energy
node_energy = integrate_u(control_signals)
```

```
print(node_energy.shape)
(200,)
print(np.round(node_energy[:5], 2))
[21.13 37.65 23.55 21.55 28.34]
```

(B) **Summarize node-level energies to produce a single estimate of control energy:**

```
# summarize nodal energy
energy = np.sum(node_energy)
print(np.round(energy, 2))
2604.71
```

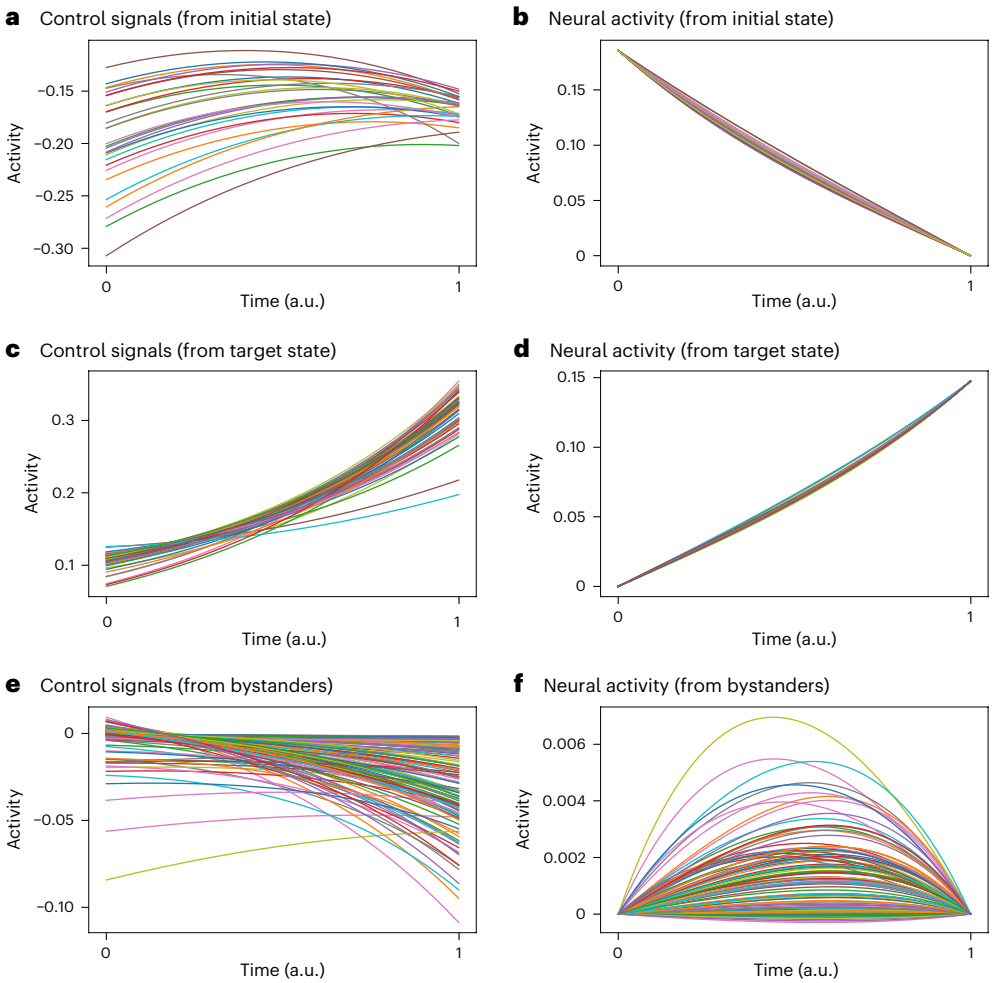

Fig. 4 | Visualize the control signals and the state trajectory. For a given state transition, the control signals ($\boldsymbol{u}(t)$, left column) and the state trajectory ($\boldsymbol{x}(t)$, right column) should be visualized. Here, because of our use of binary brain states, we group this visualization by nodes in the initial state ($\boldsymbol{x}_0$, top row), the target state ($\boldsymbol{x}_f$, middle row) and the remaining nodes (bystanders, bottom row). Across all subplots, differently colored line plots represent different system nodes (brain regions) within each group. This plot provides intuition on model behavior by showing the kinds of control signals that drive specific changes in neural activity. The top row shows that the model drives negative time-varying control signals into the nodes of the initial state (**a**), which drives their activity to 0 over time (**b**). The middle row shows that the model drives positive time-varying control signals into the nodes of the target state (**c**), which drives their activity from 0 to ~0.15 over time (**d**). The bottom row shows that the model drives diverse time-varying control signals into the bystander nodes (**e**), which in turn guide changes in these regions' activity (**f**).

BOX 1

Checking completion of the state transition

As noted in Experimental design, completion of a state transition is not guaranteed by the model, and an incomplete transition may necessitate revising either the control set (e.g., to provide more control over the system if a partial control set was used) or the time horizon (e.g., to provide more time for the model to complete the transition). We suggest the following simple plot:

```
f, ax = plt.subplots(3, 2, figsize=(7, 7))
# plot control signals for initial state
ax[0, 0].plot(control_signals[:, initial_state != 0], linewidth=0.75)
ax[0, 0].set_title("A | control signals, x0")
# plot state trajectory for initial state
ax[0, 1].plot(state_trajectory[:, initial_state != 0], linewidth=0.75)
ax[0, 1].set_title("B | neural activity, x0")
# plot control signals for target state
ax[1, 0].plot(control_signals[:, target_state != 0], linewidth=0.75)
ax[1, 0].set_title("C | control signals, xf")
# plot state trajectory for target state
ax[1, 1].plot(state_trajectory[:, target_state != 0], linewidth=0.75)
ax[1, 1].set_title("D | neural activity, xf")
# plot control signals for bystanders
ax[2, 0].plot(
 control_signals[:, np.logical_and(initial_state == 0, target_state == 0)],
 linewidth=0.75,
)
ax[2, 0].set_title("E | control signals, bystanders")
# plot state trajectory for bystanders
ax[2, 1].plot(
 state_trajectory[:, np.logical_and(initial_state == 0, target_state == 0)],
 linewidth=0.75,
)
ax[2, 1].set_title("F | neural activity, bystanders")
for cax in ax.reshape(-1):
cax.set_ylabel("activity")
cax.set_xlabel("time (a.u.)")
cax.set_xticks([0, state_trajectory.shape[0]])
cax.set_xticklabels([0, time_horizon])
f.tight_layout()
plt.show()
```

Figure 4 shows the control signals (left column) alongside the state trajectory (i.e., neural activity; right column) separately for nodes within the initial state (top row) and the target state (middle row), as well as the bystanders (bottom row). Note that we define bystanders as nodes that are outside both the initial and target states. We choose this division of nodes because it provides several simple intuitions about model behavior. First, we can see that the model drives negative time-varying control signals into the nodes of the initial state (Fig. 4a), which drives their activity to 0 over time (Fig. 4b). Second, we can see that the model drives positive time-varying control signals into the nodes of the target state (Fig. 4c), which drives their activity from 0 to ~0.15 over time (Fig. 4d). Note that ~0.15 represents the maximum neural activity after state normalization for the states presented here; this maximum activity may vary depending on state definition. Finally, we can see that diverse time-varying control signals are injected into the bystander nodes (Fig. 4e), which in turn guide changes in these regions' activity (Fig. 4f). In other words, Fig. 4 shows that the model performs a combination of 'turning off' the initial state, 'turning on' the target state and guiding diffusing activity toward the target state via the bystanders. Figure 4 also provides a simple visual way to check whether the state transition completed successfully; at the end of the simulation,

(continued from previous page)
it is apparent that activity in the target state is maximal whereas activity in the initial state and bystanders is 0, which accords with our definition of $\boldsymbol{x}_f$. This behavior explains the low reconstruction error reported for this transition (see the main text). In addition, this plot allows researchers to visualize how model behavior varies under different control sets (Supplementary Figs. 3–6) and time horizons (Supplementary Figs. 7–9). Note that while we view Fig. 4 as the simplest way to plot initial model outputs, it is only one of many options. Researchers may choose to plot $\boldsymbol{x}(t)$ and $\boldsymbol{u}(t)$ as heatmaps or on the brain's surface, which would facilitate visualization of spatial patterning (see Supplementary Information, Variations to Procedure 1).

▲ **CRITICAL** There are multiple variations to Steps 1–6 that researchers may wish to consider depending on their research goals. In the Supplementary Information, we illustrate a selection of these variants that are likely to be of broad interest to the field of neuroscience (Fig. 3c). These variations include studying non-binary brain states, implementing partial and non-uniform control sets and examining directed structural connectomes (see Supplementary Information, Variations to Procedure 1).

▲ **CRITICAL STEP** Steps 1–6 above outline how to extract the control energy for a single control task, which we defined as completing a state transition between the visual system and the default mode system by using control signals delivered to all system nodes (to view the above steps in a single notebook, see https://github.com/LindenParkesLab/nctpy/blob/f69ec009d70a46cb019da7c59a0d00b3e254731a/scripts/path_a_control_energy_binary.ipynb). Alternatively, researchers may want to examine many control tasks within the context of a single study. Thus, in nctpy, we include a Python class called `ComputeControlEnergy` that wraps all of the above steps (excluding Step 5) and applies them over a list of control tasks (see Supplementary Information, Wrapping Procedure 1 for ease of use).

Procedure 2

● **TIMING** <1–20 s

▲ **CRITICAL** Procedure 1 is the primary component of our protocol. Implementing these steps assumes that researchers are interested in studying a specific set of state transitions defined in accordance with their research questions and hypotheses. In the absence of such hypotheses, researchers may instead wish to examine nodes' general capacity to control a broad range of unspecified state transitions. To support these types of hypotheses, we present a complementary set of steps that yields estimates of average controllability, in which higher values indicate that a region is better positioned in the network to control dynamics:

1. Compute average controllability (discrete time, <1 s for 200 nodes; continuous time, 10–20 s for 200 nodes).

```
# compute average controllability
average_controllability = ave_control(A_norm=adjacency_norm,
system=system)
```

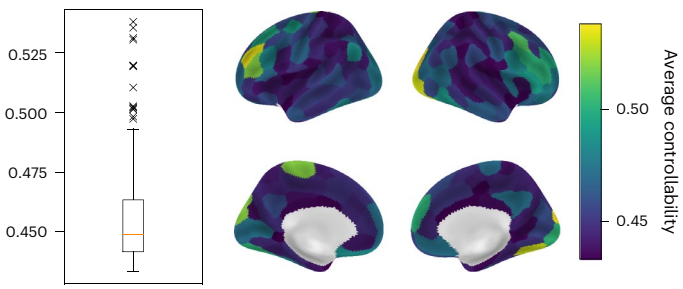

Fig. 5 | Average controllability. Each system node receives an impulse of equal magnitude. Nodes with higher average controllability are able to broadcast that impulse throughout the system to a greater extent than nodes with lower average controllability. Thus, nodes with high average controllability are better positioned within the network to control dynamics. Distribution of average controllability values are displayed using a box plot (left) and are projected onto the cortical surface (right). In the box plot, the orange line represents the median, the box spans the middle 50% of the data, the whiskers span 1.5 times the interquartile range on either side and the 'x' marks represent outliers beyond this limit.

`average_controllability` will be a vector containing the average controllability of each node of the system.

2. Visualize average controllability. Because average controllability is a regional metric, we can simply plot its distribution of values on the surface of the cortex (Fig. 5).

Troubleshooting

Troubleshooting advice can be found in Table 1.

Table 1 | Troubleshooting table for Procedure 1

Step	Problem	Possible reason	Solution
5	Inversion error out of bounds	Ill-conditioned optimization problem	If using a partial control set, try using a full control set
			If using a full control set, try reducing the size of the system (i.e., reduce the number of nodes in *A*)
			Check whether *A* has disconnected components (i.e., groups of nodes that cannot communicate with each other)
	Reconstruction error out of bounds	Failure of the state transition to complete (i.e., $\boldsymbol{x}_f$ not equal to $\boldsymbol{x}(T)$)	Increase time horizon, *T*
			If using a partial control set, try using a full control set
			If using a full control set, try reducing the size of the system (i.e., reduce the number of nodes in *A*)
			Check whether *A* has disconnected components (i.e., groups of nodes that cannot communicate with each other)

Timing

As noted throughout the Procedure, the timing of each step is relatively short.

Procedure 1

Steps 1–2, normalize the adjacency matrix: 1–5 s
Step 3, define a control task: 1–5 s
Step 4–6, compute control energy: ~1 min

Procedure 2

Step 1, compute average controllability: <1–20 s

We note two clarifications to the above timing information. First, these time estimates are only for a single execution of each step as shown in the protocol. In reality, these steps will probably need to be executed many times over to achieve researchers' goals. For example, a given study may need to compute control energy for multiple control tasks across multiple subjects, which will increase total run time. This run time will increase further if null network models are used (see Supplementary Information, Null network models), wherein each step may be run thousands of times for a single control task. However, in these instances, protocol steps can be trivially parallelized using high-performance computing infrastructure, which will reduce total run time. Second, timing will vary as a function of researchers' data specifications. For example, in this protocol, we performed analysis on a structural connectome comprising 200 nodes. Increasing parcellation resolution will increase run time.

Anticipated results

The final outputs of our protocol will depend on whether researchers choose to follow Procedure 1 or Procedure 2. For the former, the output will be one estimate of control energy per control task, or one estimate per brain region per task if energy was not summarized across regions. This value will be positive and can be thought of as the amount of effort the model has to exert to complete a specific control task; higher energy corresponds to greater effort. For the latter, the output will be one estimate of average controllability for each brain region, a regional map of control over system dynamics (Fig. 5). These regional values will also be positive. Greater average controllability indicates that regions are better positioned within the network's topology to broadcast an impulse and, as such, may better orchestrate control of brain dynamics.

What can researchers do with these outputs? The answers to this question are diverse and depend heavily on the researchers' goals. As we discussed in Applications, we have used NCT to investigate a range of research questions that spanned from examining the influence of topology[10,52,53], to predicting state transitions observed in functional data[14,55], to studying individual differences, including psychosis symptoms[21], executive function[73] and sex effects[19]. Providing detailed guidance on each of these applications is beyond the scope of this protocol. However, to support this protocol, we outline the use of null network models as an initial analysis that we believe is an essential step irrespective of researchers' study goals (see Supplementary Information, Null network models).

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

The PNC data are publicly available in the Database of Genotypes and Phenotypes under accession number phs00607.v3.p2 (https://www.ncbi.nlm.nih.gov/projects/gap/cgi-bin/study.cgi?study_id=phs000607.v3.p2).

Code availability

All analysis code is freely available at https://github.com/LindenParkesLab/nctpy/.

Received: 23 August 2023; Accepted: 16 May 2024;
Published online: 29 July 2024

References

1. Bassett, D. S. & Sporns, O. Network neuroscience. *Nat. Neurosci.* **20**, 353–364 (2017).
2. Bassett, D. S., Zurn, P. & Gold, J. I. On the nature and use of models in network neuroscience. *Nat. Rev. Neurosci.* **19**, 566–578 (2018).
3. Betzel, R. F. & Bassett, D. S. Multi-scale brain networks. *Neuroimage* **160**, 73–83 (2017).
4. Fornito, A., Zalesky, A. & Bullmore, E. T. *Fundamentals of Brain Network Analysis*. (Elsevier/Academic Press, 2016).
5. Menara, T., Katewa, V., Bassett, D. S. & Pasqualetti, F. The structured controllability radius of symmetric (brain) networks. In *2018 Annual American Control Conference (ACC)* 2802–2807 (IEEE, Milwaukee, Wisconsin, USA, 2018).
6. Pasqualetti, F., Zampieri, S. & Bullo, F. Controllability metrics, limitations and algorithms for complex networks. *IEEE Trans. Control Netw. Syst.* **1**, 40–52 (2014).
7. Kim, J. Z. & Bassett, D. S. Linear dynamics and control of brain networks. In *Neural Engineering* (ed. He, B.) 497–518 (Springer, (2020).
8. Karrer, T. M. et al. A practical guide to methodological considerations in the controllability of structural brain networks. *J. Neural Eng.* **17**, 026031 (2020).
9. Seguin, C., Sporns, O. & Zalesky, A. Brain network communication: concepts, models and applications. *Nat. Rev. Neurosci.* **24**, 557–574 (2023).
10. Gu, S. et al. Controllability of structural brain networks. *Nat. Commun.* **6**, 8414 (2015).
11. Gu, S. et al. Optimal trajectories of brain state transitions. *Neuroimage* **148**, 305–317 (2017).
12. Tang, E. et al. Developmental increases in white matter network controllability support a growing diversity of brain dynamics. *Nat. Commun.* **8**, 1252 (2017).
13. Tang, E. et al. Control of brain network dynamics across diverse scales of space and time. *Phys. Rev. E* **101**, 062301 (2020).
14. Stiso, J. et al. White matter network architecture guides direct electrical stimulation through optimal state transitions. *Cell Rep.* **28**, 2554–2566.e7 (2019).
15. Scheid, B. H. et al. Time-evolving controllability of effective connectivity networks during seizure progression. *Proc. Natl Acad. Sci. USA* **118**, e2006436118 (2021).
16. Medaglia, J. D. et al. Network controllability in the inferior frontal gyrus relates to controlled language variability and susceptibility to TMS. *J. Neurosci.* **38**, 6399–6410 (2018).
17. Medaglia, J. D. et al. Language tasks and the network control role of the left inferior frontal gyrus. *eneuro* **8**, ENEURO.0382-20.2021 (2021).
18. Muldoon, S. F. et al. Stimulation-based control of dynamic brain networks. *PLOS Comput. Biol.* **12**, e1005076 (2016).

19. Cornblath, E. J. et al. Sex differences in network controllability as a predictor of executive function in youth. *Neuroimage* **188**, 122–134 (2019).
20. Cornblath, E. J. et al. Temporal sequences of brain activity at rest are constrained by white matter structure and modulated by cognitive demands. *Commun. Biol.* **3**, 261 (2020).
21. Parkes, L. et al. Network controllability in transmodal cortex predicts psychosis spectrum symptoms. *Biol. Psychiatry* **89**, S370–S371 (2021).
22. Parkes, L. et al. Asymmetric signaling across the hierarchy of cytoarchitecture within the human connectome. *Sci. Adv.* **8**, eadd2185 (2022).
23. Satterthwaite, T. D. et al. Neuroimaging of the Philadelphia Neurodevelopmental Cohort. *Neuroimage* **86**, 544–553 (2014).
24. Satterthwaite, T. D. et al. The Philadelphia Neurodevelopmental Cohort: a publicly available resource for the study of normal and abnormal brain development in youth. *Neuroimage* **124**, 1115–1119 (2016).
25. Oh, S. W. et al. A mesoscale connectome of the mouse brain. *Nature* **508**, 207–214 (2014).
26. Harris, J. A. et al. Hierarchical organization of cortical and thalamic connectivity. *Nature* **575**, 195–202 (2019).
27. Knox, J. E. et al. High-resolution data-driven model of the mouse connectome. *Netw. Neurosci.* **3**, 217–236 (2019).
28. Chiêm, B., Crevecoeur, F. & Delvenne, J.-C. Structure-informed functional connectivity driven by identifiable and state-specific control regions. *Netw. Neurosci.* **5**, 591–613 (2021).
29. Jeganathan, J. et al. Fronto-limbic dysconnectivity leads to impaired brain network controllability in young people with bipolar disorder and those at high genetic risk. *NeuroImage Clin.* **19**, 71–81 (2018).
30. Kenett, Y. N. et al. Driving the brain towards creativity and intelligence: a network control theory analysis. *Neuropsychologia* **118**, 79–90 (2018).
31. Yuan, J., Ji, S., Luo, L., Lv, J. & Liu, T. Control energy assessment of spatial interactions among macro-scale brain networks. *Hum. Brain Mapp.* **43**, 2181–2203 (2022).
32. Singleton, S. P. et al. Receptor-informed network control theory links LSD and psilocybin to a flattening of the brain's control energy landscape. *Nat. Commun.* **13**, 5812 (2022).
33. Singleton, S. P. et al. Time-resolved network control analysis links reduced control energy under dmt with the serotonin 2a receptor, signal diversity, and subjective experience. Preprint at https://www.biorxiv.org/content/10.1101/2023.05.11.540409v1 (2023).
34. Luppi, A. I. et al. Transitions between cognitive topographies: contributions of network structure, neuromodulation, and disease. Preprint at https://www.biorxiv.org/content/10.1101/2023.03.16.532981v1 (2023).
35. Maxwell, J. C. On governors. *Proc. R. Soc. Lond.* **16**, 270–283 (1867).
36. Grasser, F., D'Arrigo, A., Colombi, S. & Rufer, A. C. JOE: a mobile, inverted pendulum. *IEEE Trans. Ind. Electron.* **49**, 107–114 (2002).
37. Hodgkin, A. L. & Huxley, A. F. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.* **117**, 500–544 (1952).
38. Papadopoulos, L., Kim, J. Z., Kurths, J. & Bassett, D. S. Development of structural correlations and synchronization from adaptive rewiring in networks of Kuramoto oscillators. *Chaos* **27**, 073115 (2017).
39. Wilson, H. R. & Cowan, J. D. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys. J.* **12**, 1–24 (1972).
40. Schiff, S. J. et al. Controlling chaos in the brain. *Nature* **370**, 615–620 (1994).
41. Cash, R. F. H. et al. Using brain imaging to improve spatial targeting of transcranial magnetic stimulation for depression. *Biol. Psychiatry* **90**, 689–700 (2021).
42. Suárez, L. E., Markello, R. D., Betzel, R. F. & Misic, B. Linking structure and function in macroscale brain networks. *Trends Cogn. Sci.* **24**, 302–315 (2020).
43. Vázquez-Rodríguez, B. et al. Gradients of structure–function tethering across neocortex. *Proc. Natl Acad. Sci. USA* **116**, 21219–21227 (2019).
44. Baum, G. L. et al. Development of structure–function coupling in human brain networks during youth. *Proc. Natl Acad. Sci. USA* **117**, 771–778 (2020).
45. Preti, M. G. & Van De Ville, D. Decoupling of brain function from structure reveals regional behavioral specialization in humans. *Nat. Commun.* **10**, 4747 (2019).
46. Luo, N. et al. Structural brain architectures match intrinsic functional networks and vary across domains: a study from 15 000+ individuals. *Cereb. Cortex* **30**, 5460–5470 (2020).
47. Seguin, C., Tian, Y. & Zalesky, A. Network communication models improve the behavioral and functional predictive utility of the human structural connectome. *Netw. Neurosci.* **4**, 980–1006 (2020).
48. Betzel, R. F., Faskowitz, J., Mišić, B., Sporns, O. & Seguin, C. Multi-policy models of interregional communication in the human connectome. Preprint at https://www.biorxiv.org/content/10.1101/2022.05.08.490752v1 (2022).
49. Fox, P. T. & Friston, K. J. Distributed processing; distributed functions? *Neuroimage* **61**, 407–426 (2012).
50. Hespanha, J. P. *Linear Systems Theory: Second Edition* (Princeton University Press, 2018).
51. Yan, G. et al. Network control principles predict neuron function in the *Caenorhabditis elegans* connectome. *Nature* **550**, 519–523 (2017).
52. Betzel, R. F., Gu, S., Medaglia, J. D., Pasqualetti, F. & Bassett, D. S. Optimally controlling the human connectome: the role of network topology. *Sci. Rep.* **6**, 30770 (2016).
53. Kim, J. Z. et al. Role of graph architecture in controlling dynamical networks with applications to neural systems. *Nat. Phys.* **14**, 91–98 (2018).
54. Sejnowski, T. J., Churchland, P. S. & Movshon, J. A. Putting big data to good use in neuroscience. *Nat. Neurosci.* **17**, 1440–1441 (2014).
55. Braun, U. et al. Brain network dynamics during working memory are modulated by dopamine and diminished in schizophrenia. *Nat. Commun.* **12**, 3478 (2021).
56. Tong, F. Primary visual cortex and visual awareness. *Nat. Rev. Neurosci.* **4**, 219–229 (2003).
57. Gordon, E. M. et al. A somato-cognitive action network alternates with effector regions in motor cortex. *Nature* **617**, 351–359 (2023).
58. Bertolero, M. A., Yeo, B. T. T. & D'Esposito, M. The modular and integrative functional architecture of the human brain. *Proc. Natl Acad. Sci. USA* **112**, E6798–E6807 (2015).
59. Bertolero, M. A., Yeo, B. T. T., Bassett, D. S. & D'Esposito, M. A mechanistic model of connector hubs, modularity and cognition. *Nat. Hum. Behav.* **2**, 765–777 (2018).
60. van den Heuvel, M. P. & Sporns, O. Network hubs in the human brain. *Trends Cogn. Sci.* **17**, 683–696 (2013).
61. van den Heuvel, M. P. & Sporns, O. Rich-club organization of the human connectome. *J. Neurosci.* **31**, 15775–15786 (2011).
62. Fornito, A., Zalesky, A. & Breakspear, M. The connectomics of brain disorders. *Nat. Rev. Neurosci.* **16**, 159–172 (2015).
63. Crossley, N. A. et al. The hubs of the human connectome are generally implicated in the anatomy of brain disorders. *Brain* **137**, 2382–2395 (2014).
64. Van Essen, D. C. et al. The WU-Minn Human Connectome Project: an overview. *Neuroimage* **80**, 62–79 (2013).
65. Casey, B. J. et al. The Adolescent Brain Cognitive Development (ABCD) study: imaging acquisition across 21 sites. *Dev. Cogn. Neurosci.* **32**, 43–54 (2018).
66. Alexander, L. M. et al. An open resource for transdiagnostic research in pediatric mental health and learning disorders. *Sci. Data* **4**, 170181 (2017).
67. Amunts, K. et al. BigBrain: an ultrahigh-resolution 3D human brain model. *Science* **340**, 1472–1475 (2013).
68. Hawrylycz, M. J. et al. An anatomically comprehensive atlas of the adult human brain transcriptome. *Nature* **489**, 391–399 (2012).
69. Thomas Yeo, B. T. et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.* **106**, 1125–1165 (2011).
70. de Reus, M. A. & van den Heuvel, M. P. Simulated rich club lesioning in brain networks: a scaffold for communication and integration? *Front. Hum. Neurosci.* **8**, 647 (2014).
71. van den Heuvel, M. P. & Sporns, O. An anatomical substrate for integration among functional networks in human cortex. *J. Neurosci.* **33**, 14489–14500 (2013).
72. Margulies, D. S. et al. Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proc. Natl Acad. Sci. USA* **113**, 12574–12579 (2016).
73. Cui, Z. et al. Optimization of energy state transition trajectory supports the development of executive function during youth. *eLife* **9**, e53060 (2020).
74. Niendam, T. A. et al. Meta-analytic evidence for a superordinate cognitive control network subserving diverse executive functions. *Cogn. Affect. Behav. Neurosci.* **12**, 241–268 (2012).
75. Parkes, L., Fulcher, B. D., Yücel, M. & Fornito, A. Transcriptional signatures of connectomic subregions of the human striatum. *Genes Brain Behav.* **16**, 647–663 (2017).
76. Fulcher, B. D., Murray, J. D., Zerbi, V. & Wang, X.-J. Multimodal gradients across mouse cortex. *Proc. Natl Acad. Sci. USA* **116**, 4689–4695 (2019).
77. Fulcher, B. D. & Fornito, A. A transcriptional signature of hub connectivity in the mouse connectome. *Proc. Natl Acad. Sci. USA* **113**, 1435–1440 (2016).
78. Larivière, S. et al. Microstructure-informed connectomics: enriching large-scale descriptions of healthy and diseased brains. *Brain Connect.* **9**, 113–127 (2019).
79. Arnatkevičiūtė, A., Fulcher, B. D. & Fornito, A. A practical guide to linking brain-wide gene expression and neuroimaging data. *Neuroimage* **189**, 353–367 (2019).
80. Arnatkeviciute, A. et al. Genetic influences on hub connectivity of the human connectome. *Nat. Commun.* **12**, 4237 (2021).
81. Arnatkevičiūtė, A., Fulcher, B. D., Pocock, R. & Fornito, A. Hub connectivity, neuronal diversity, and gene expression in the *Caenorhabditis elegans* connectome. *PLoS Comput. Biol.* **14**, e1005989 (2018).
82. Anderson, K. M. et al. Gene expression links functional networks across cortex and striatum. *Nat. Commun.* **9**, 1428 (2018).
83. Anderson, K. M. et al. Convergent molecular, cellular, and cortical neuroimaging signatures of major depressive disorder. *Proc. Natl Acad. Sci. USA* **117**, 25138–25149 (2020).
84. Paquola, C. et al. Microstructural and functional gradients are increasingly dissociated in transmodal cortices. *PLoS Biol.* **17**, e3000284 (2019).
85. García-Cabezas, M, Á., Zikopoulos, B. & Barbas, H. The Structural Model: a theory linking connections, plasticity, pathology, development and evolution of the cerebral cortex. *Brain Struct. Funct.* **224**, 985–1008 (2019).
86. Yarkoni, T., Poldrack, R. A., Nichols, T. E., Van Essen, D. C. & Wager, T. D. Large-scale automated synthesis of human functional neuroimaging data. *Nat. Methods* **8**, 665–670 (2011).
87. Hodgkin, A. L. & Huxley, A. F. Propagation of electrical signals along giant nerve fibres. *Proc. R. Soc. Lond. B Biol. Sci.* **140**, 177–183 (1952).
88. Breakspear, M. Dynamic models of large-scale brain activity. *Nat. Neurosci.* **20**, 340–352 (2017).
89. Shine, J. M. et al. Computational models link cellular mechanisms of neuromodulation to large-scale neural dynamics. *Nat. Neurosci.* **24**, 765–776 (2021).
90. Lu, Z. & Bassett, D. S. Invertible generalized synchronization: a putative mechanism for implicit learning in neural systems. *Chaos* **30**, 063133 (2020).

91. Suárez, L. E., Richards, B. A., Lajoie, G. & Misic, B. Learning function from structure in neuromorphic networks. *Nat. Mach. Intell.* **3**, 771–786 (2021).
92. Roberts, J. A. et al. Metastable brain waves. *Nat. Commun.* **10**, 1056 (2019).
93. Demirtaş, M. et al. Hierarchical heterogeneity across human cortex shapes large-scale neural dynamics. *Neuron* **101**, 1181–1194.e13 (2019).
94. Deco, G. et al. Resting-state functional connectivity emerges from structurally and dynamically shaped slow linear fluctuations. *J. Neurosci.* **33**, 11239–11252 (2013).
95. Deco, G. et al. Dynamical consequences of regional heterogeneity in the brain's transcriptional landscape. *Sci. Adv.* **7**, eabf4752 (2021).
96. Monasson, R. & Rosay, S. Transitions between spatial attractors in place-cell models. *Phys. Rev. Lett.* **115**, 098101 (2015).
97. Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).
98. Fornito, A., Zalesky, A. & Breakspear, M. Graph analysis of the human connectome: promise, progress, and pitfalls. *Neuroimage* **80**, 426–444 (2013).
99. Bertolero, M. A. & Bassett, D. S. On the nature of explanations offered by network science: a perspective from and for practicing neuroscientists. *Top. Cogn. Sci.* **12**, 1272–1293 (2020).
100. Vázquez-Rodríguez, B., Liu, Z.-Q., Hagmann, P. & Misic, B. Signal propagation via cortical hierarchies. *Netw. Neurosci.* **4**, 1072–1090 (2020).
101. Bazinet, V., Vos de Wael, R., Hagmann, P., Bernhardt, B. C. & Misic, B. Multiscale communication in cortico-cortical networks. *Neuroimage* **243**, 118546 (2021).
102. Fornito, A., Zalesky, A., Pantelis, C. & Bullmore, E. T. Schizophrenia, neuroimaging and connectomics. *Neuroimage* **62**, 2296–2314 (2012).
103. Bassett, D. S. et al. Hierarchical organization of human cortical networks in health and schizophrenia. *J. Neurosci.* **28**, 9239–9248 (2008).
104. Bassett, D. S., Nelson, B. G., Mueller, B. A., Camchong, J. & Lim, K. O. Altered resting state complexity in schizophrenia. *Neuroimage* **59**, 2196–2207 (2012).
105. Seguin, C., Razi, A. & Zalesky, A. Inferring neural signalling directionality from undirected structural connectomes. *Nat. Commun.* **10**, 4289 (2019).
106. Seguin, C., Mansour, L. S., Sporns, O., Zalesky, A. & Calamante, F. Network communication models narrow the gap between the modular organization of structural and functional brain networks. *Neuroimage* **257**, 119323 (2022).
107. Mišić, B. et al. Cooperative and competitive spreading dynamics on the human connectome. *Neuron* **86**, 1518–1529 (2015).
108. Nozari, E. et al. Macroscopic resting-state brain dynamics are best described by linear models. *Nat. Biomed. Eng.* **8**, 68–84 (2023).
109. Bargmann, C. I. & Marder, E. From the connectome to brain function. *Nat. Methods* **10**, 483–490 (2013).
110. Schaefer, A. et al. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cereb. Cortex* **28**, 3095–3114 (2018).
111. Oldham, S. et al. The efficacy of different preprocessing steps in reducing motion-related confounds in diffusion MRI connectomics. *Neuroimage* **222**, 117252 (2020).
112. Sarwar, T., Ramamohanarao, K. & Zalesky, A. Mapping connectomes with diffusion MRI: deterministic or probabilistic tractography? *Magn. Reson. Med.* **81**, 1368–1384 (2019).
113. Robinson, P. A. et al. Eigenmodes of brain activity: neural field theory predictions and comparison with experiment. *Neuroimage* **142**, 79–98 (2016).
114. Deco, G., Jirsa, V. K., Robinson, P. A., Breakspear, M. & Friston, K. The dynamic brain: from spiking neurons to neural masses and cortical fields. *PLoS Comput. Biol.* **4**, e1000092 (2008).
115. Shenoy, K. V. & Kao, J. C. Measurement, manipulation and modeling of brain-wide neural population dynamics. *Nat. Commun.* **12**, 633 (2021).
116. He, X. et al. Uncovering the biological basis of control energy: structural and metabolic correlates of energy inefficiency in temporal lobe epilepsy. *Sci. Adv.* **8**, eabn2293 (2022).
117. McCormick, D. A., Shu, Y. & Yu, Y. Hodgkin and Huxley model—still standing? *Nature* **445**, E1–E2 (2007).
118. Durstewitz, D., Seamans, J. K. & Sejnowski, T. J. Neurocomputational models of working memory. *Nat. Neurosci.* **3**, 1184–1191 (2000).
119. Hu, S. G. et al. Associative memory realized by a reconfigurable memristive Hopfield neural network. *Nat. Commun.* **6**, 7522 (2015).
120. Vaidyanathan, S. & Volos, C., eds. *Advances and Applications in Nonlinear Control Systems*. Vol. **635** (Springer International Publishing, Cham, Switzerland, 2016).
121. Vinodh Kumar, E. & Jerome, J. Robust LQR controller design for stabilizing and trajectory tracking of inverted pendulum. *Procedia Eng.* **64**, 169–178 (2013).
122. Bastos, A. M. & Schoffelen, J.-M. A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. *Front. Syst. Neurosci.* **9**, 175 (2016).
123. Smeal, R. M., Ermentrout, G. B. & White, J. A. Phase-response curves and synchronized neural networks. *Philos. Trans. R. Soc. B Biol. Sci.* **365**, 2407–2422 (2010).
124. Park, Y. & Ermentrout, B. Weakly coupled oscillators in a slowly varying world. *J. Comput. Neurosci.* **40**, 269–281 (2016).
125. Brunton, S. L., Brunton, B. W., Proctor, J. L. & Kutz, J. N. Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. *PLoS One* **11**, e0150171 (2016).
126. Proctor, J. L., Brunton, S. L. & Kutz, J. N. Dynamic mode decomposition with control. *SIAM J. Appl. Dyn. Syst.* **15**, 142–161 (2016).
127. Zañudo, J. G. T., Yang, G. & Albert, R. Structure-based control of complex networks with nonlinear dynamics. *Proc. Natl Acad. Sci. USA* **114**, 7234–7239 (2017).
128. Haynes, G. W. & Hermes, H. Nonlinear controllability via lie theory. *SIAM J. Control* **8**, 450–460 (1970).
129. Towlson, E. K. et al. *Caenorhabditis elegans* and the network control framework—FAQs. *Philos. Trans. R. Soc. B Biol. Sci.* **373**, 20170372 (2018).
130. Felleman, D. J. & Van Essen, D. C. Distributed hierarchical processing in the primate cerebral cortex. *Cereb. Cortex* **1**, 1–47 (1991).
131. de Reus, M. A. & van den Heuvel, M. P. Estimating false positives and negatives in brain networks. *Neuroimage* **70**, 402–409 (2013).
132. Yendiki, A., Koldewyn, K., Kakunoori, S., Kanwisher, N. & Fischl, B. Spurious group differences due to head motion in a diffusion MRI study. *Neuroimage* **88**, 79–90 (2014).
133. Baum, G. L. et al. The impact of in-scanner head motion on structural connectivity derived from diffusion MRI. *Neuroimage* **173**, 275–286 (2018).
134. Sotiropoulos, S. N. & Zalesky, A. Building connectomes using diffusion MRI: why, how and but. *NMR Biomed.* **32**, e3752 (2019).
135. Cieslak, M. et al. QSIPrep: an integrative platform for preprocessing and reconstructing diffusion MRI data. *Nat. Methods* **18**, 775–778 (2021).
136. Roalf, D. R. et al. The impact of quality assurance assessment on diffusion tensor imaging outcomes in a large-scale population-based cohort. *Neuroimage* **125**, 903–919 (2016).
137. Fallon, J. et al. Timescales of spontaneous fMRI fluctuations relate to structural connectivity in the brain. *Netw. Neurosci.* **4**, 788–806 (2020).
138. Shafiei, G. et al. Topographic gradients of intrinsic dynamics across neocortex. *eLife* **9**, e62116 (2020).
139. Wolff, A. et al. Intrinsic neural timescales: temporal integration and segregation. *Trends Cogn. Sci.* **26**, 159–173 (2022).
140. Sydnor, V. J. et al. Intrinsic activity development unfolds along a sensorimotor–association cortical axis in youth. *Nat. Neurosci.* **26**, 638–649 (2023).
141. Rosen, A. F. G. et al. Quantitative assessment of structural image quality. *Neuroimage* **169**, 407–418 (2018).
142. Ciric, R. et al. Mitigating head motion artifact in functional connectivity MRI. *Nat. Protoc.* **13**, 2801–2826 (2018).
143. Gorgolewski, K. et al. Nipype: a flexible, lightweight and extensible neuroimaging data processing framework in Python. *Front. Neuroinform.* **5**, 13 (2011).
144. Satterthwaite, T. D. et al. An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. *Neuroimage* **64**, 240–256 (2013).
145. Esteban, O. et al. fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nat. Methods* **16**, 111–116 (2019).

Acknowledgements

The authors acknowledge the following funding bodies: National Institute of Mental Health grant R00MH127296 (to L.P.) (the content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health); NARSAD Young Investigator Grant 28995 from the Brain & Behavior Research Foundation (to L.P.); National Science Foundation grant DGE-1321851 (to J.Z.K.); National Institute of Mental Health grants R21MH106799, R01MH113550 and RF1MH116920 (to D.S.B. and T.D.S.); the Swartz Foundation (to D.S.B.); the John D. and Catherine T. MacArthur Foundation (to D.S.B.); National Institute of Mental Health grants R01MH120482 (to T.D.S.), R01MH107703 (to T.D.S.), R01MH112847 (to T.D.S. and R.T.S.), R37MH125829 (to T.D.S.), R01EB022573 (to T.D.S.), R01MH107235 (to R.C.G.) and R01MH119219 (to R.C.G. and R.E.G.); the Penn-CHOP Lifespan Brain Institute; and National Institute of Mental Health grants RC2MH089983 and RC2MH089924 (Philadelphia Neurodevelopmental Cohort).

Author contributions

Conceptualization was the responsibility of L.P., J.Z.K., T.D.S. and D.S.B. Methodology was the responsibility of L.P., J.Z.K., J.S. and D.S.B. Software was the responsibility of L.P., J.Z.K. and J.S. Formal analysis was performed by L.P. and J.Z.K. Visualization was performed by L.P. and J.Z.K. Data curation was done by J.K.B., M.C., S.C., R.E.G., R.C.G., R.T.S., D.Z. and T.D.S. The original draft was written by L.P. and J.Z.K. Reviewing and editing of the writing was done by L.P., J.Z.K, J.S., J.K.B., M.C., S.C., R.E.G., R.C.G., F.P., R.T.S., D.Z., T.D.S and D.S.B.

Competing interests

R.T.S. receives consulting compensation from Octave Bioscience and compensation for reviewer duties from the American Medical Association. All other authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41596-024-01023-w.

Correspondence and requests for materials should be addressed to Linden Parkes.

Peer review information *Nature Protocols* thanks Masafumi Oizumi, Lucina Uddin and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Related links

Key references using this protocol

Gu, S. et al. *Neuroimage* **148**, 305–317 (2017): https://doi.org/10.1016/j.neuroimage.2017.01.003
Kim, J. Z. et al. *Nat. Phys*. **14**, 91–98 (2018): https://doi.org/10.1038/nphys4268
Stiso, J. et al. *Cell Rep*. **28**, 2554–2566.e7 (2019): https://doi.org/10.1016/j.celrep.2019.08.008
Karrer, T. M. et al. *J. Neural Eng*. **17**, 026031 (2020): https://doi.org/10.1088/1741-2552/ab6e8b
Parkes, L. et al. *Sci. Adv*. **8**, eadd2185 (2022): https://doi.org/10.1126/sciadv.add2185

[1]Department of Psychiatry, Rutgers University, Piscataway, NJ, USA. [2]Department of Bioengineering, University of Pennsylvania, Philadelphia, PA, USA. [3]Lifespan Informatics and Neuroimaging Center (PennLINC), Department of Psychiatry, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA, USA. [4]Department of Physics, Cornell University, Ithaca, NY, USA. [5]Penn/CHOP Lifespan Brain Institute, Perelman School of Medicine, Children's Hospital of Philadelphia Research Institute, Philadelphia, PA, USA. [6]Department of Psychiatry, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA, USA. [7]Department of Mechanical Engineering, University of California, Riverside, Riverside, CA, USA. [8]Department of Biostatistics, Epidemiology, and Informatics, Perelman School of Medicine, Philadelphia, PA, USA. [9]Center for Biomedical Image Computation and Analytics, University of Pennsylvania, Philadelphia, PA, USA. [10]Penn Statistics in Imaging and Visualization Endeavor (PennSIVE), Center for Clinical Epidemiology and Biostatistics, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA, USA. [11]Department of Neurology, Perelman School of Medicine, Philadelphia, PA, USA. [12]Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA, USA. [13]Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA, USA. [14]Santa Fe Institute, Santa Fe, NM, USA. [15]These authors contributed equally: Linden Parkes, Jason Z. Kim.

nature portfolio

Corresponding author(s): DBPR NP-P230449B

Last updated by author(s): Apr 10, 2024

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our Editorial Policies and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed	
☐	☒	The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☒	☐	A description of all covariates tested
☐	☒	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☒	☐	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒	Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	n/a
Data analysis	All analysis code is freely available at https://github.com/LindenParkesLab/nctpy/

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our policy

The PNC data are publicly available in the Database of Genotypes and Phenotypes: accession number: phs00607.v3.p2; https://www.ncbi.nlm.nih.gov/projects/gap/cgi-bin/study.cgi?study_id=phs000607.v3.p2

Human research participants

Policy information about studies involving human research participants and Sex and Gender in Research.

Reporting on sex and gender	N/A
Population characteristics	N/A
Recruitment	N/A
Ethics oversight	N/A

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	253
Data exclusions	N/A
Replication	N/A
Randomization	N/A
Blinding	N/A

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☒	☐ Animals and other organisms
☒	☐ Clinical data
☒	☐ Dual use research of concern

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☐	☒ MRI-based neuroimaging

Magnetic resonance imaging

Experimental design

Design type	diffusion weighted and resting-state and MRI
Design specifications	120 timepoints of resting state data per subject
Behavioral performance measures	N/A

Acquisition

Imaging type(s): functional, diffusion

Field strength: 3T

Sequence & imaging parameters: MRI data were acquired on a 3 Tesla Siemens Tim Trio scanner with a 32-channel head coil at the Hospital of the University of Pennsylvania. Diffusion weighted imaging (DWI) scans were acquired via a twice-refocused spin-echo (TRSE) single-shot echo-planar imaging (EPI) sequence (TR=8100 ms, TE=82 ms, FOV=240mm2/240mm2; Matrix=RL: 128, AP: 128, Slices: 70, in-plane resolution of 1.875 mm2; slice thickness=2 mm, gap=0; flip angle=90□/180□/180□, 71 volumes, GRAPPA factor=3, bandwidth=2170 Hz/pixel, PE direction=AP). The sequence utilized a four-lobed diffusion encoding gradient scheme combined with a 90-180-180 spin-echo sequence designed to minimize eddy-current artifacts. The sequence consisted of 64 diffusion-weighted directions with b=1000 s/mm2 and 7 interspersed scans where b=0 s/mm2. The imaging volume was prescribed in axial orientation and covered the entire brain.
In addition to the DWI scan, a B0 map of the main magnetic field was derived from a double-echo, gradient-recalled echo (GRE) sequence, allowing for the estimation and correction of field distortions. Prior to DWI acquisition, a 5-min magnetization-prepared, rapid acquisition gradient-echo T1-weighted (MPRAGE) image (TR=1810 ms, TE=3.51 ms, FOV=180 x 240 mm, matrix 256 x 192, effective voxel resolution of 0.94 x 0.94 x 1 mm) was acquired for each participant.
Finally, approximately 6 minutes of rs-fMRI data was acquired using a blood oxygen level-dependent (BOLD-weighted) sequence (TR=3000 ms; TE=32 ms; FoV=192 x 192 mm; resolution 3 mm isotropic; 124 volumes).

Area of acquisition: whole brain

Diffusion MRI ☒ Used ☐ Not used

Parameters: MRI data were acquired on a 3 Tesla Siemens Tim Trio scanner with a 32-channel head coil at the Hospital of the University of Pennsylvania. Diffusion weighted imaging (DWI) scans were acquired via a twice-refocused spin-echo (TRSE) single-shot echo-planar imaging (EPI) sequence (TR=8100 ms, TE=82 ms, FOV=240mm2/240mm2; Matrix=RL: 128, AP: 128, Slices: 70, in-plane resolution of 1.875 mm2; slice thickness=2 mm, gap=0; flip angle=90□/180□/180□, 71 volumes, GRAPPA factor=3, bandwidth=2170 Hz/pixel, PE direction=AP). The sequence utilized a four-lobed diffusion encoding gradient scheme combined with a 90-180-180 spin-echo sequence designed to minimize eddy-current artifacts. The sequence consisted of 64 diffusion-weighted directions with b=1000 s/mm2 and 7 interspersed scans where b=0 s/mm2. The imaging volume was prescribed in axial orientation and covered the entire brain.
In addition to the DWI scan, a B0 map of the main magnetic field was derived from a double-echo, gradient-recalled echo (GRE) sequence, allowing for the estimation and correction of field distortions. Prior to DWI acquisition, a 5-min magnetization-prepared, rapid acquisition gradient-echo T1-weighted (MPRAGE) image (TR=1810 ms, TE=3.51 ms, FOV=180 x 240 mm, matrix 256 x 192, effective voxel resolution of 0.94 x 0.94 x 1 mm) was acquired for each participant.
Finally, approximately 6 minutes of rs-fMRI data was acquired using a blood oxygen level-dependent (BOLD-weighted) sequence (TR=3000 ms; TE=32 ms; FoV=192 x 192 mm; resolution 3 mm isotropic; 124 volumes).

Preprocessing

Preprocessing software: - Structural connectome reconstruction was performed using QSIprep 0.14.2135, which is based on Nipype 1.6.1143. Connectomes were extracted using the 200-node variant of the Schaefer parcellation, ordered according to 7 canonical brain systems. The strength of inter-regional connectivity was summarized using the number of streamlines that intersected each pair of parcels. Connectomes were averaged over subjects. This group-averaged connectome was thresholded by retaining the edges that were present in at least 60% of participants' connectomes. This process resulted in a final connectome with 98% edge density.
- rs-fMRI was also obtained from the same 253 PNC participants. These data were used to generate empirical brain activity states in order to perform non-binary state transitions (see Variations to Procedure 1 in the Supplementary Information). The eXtensible Connectivity Pipeline (XCP-D) was used to post-process the outputs of fMRIPrep version 20.2.3145. XCP was built with Nipype 1.7.0143. Processed rs-fMRI time series were extracted from the same 200-node parcellation mentioned above.

Normalization: see above

Normalization template: see above

Noise and artifact removal: see above

Volume censoring: see above

Statistical modeling & inference

Model type and settings: N/A

Effect(s) tested: N/A

Specify type of analysis: ☒ Whole brain ☐ ROI-based ☐ Both

Statistic type for inference (See Eklund et al. 2016): N/A

Correction N/A

Models & analysis

n/a	Involved in the study
☒	☐ Functional and/or effective connectivity
☐	☒ Graph analysis
☒	☐ Multivariate modeling or predictive analysis

Graph analysis undirected diffusion-weighted structural connectome comprising 200 nodes averaged over 250 subjects