The control costs of human brain dynamics

Eric G Ceballos[1],* Andrea I Luppi[1], Gabriel Castrillon[3,4,5], Manish Saggar[2], Bratislav Misic[1], Valentin Riedl[3,4]†
[1]*Montréal Neurological Institute, McGill University, Montréal, QC, Canada*
[2]*Department of Psychiatry and Behavioral Sciences, Stanford University, Stanford, CA, USA*
[3]*Department of Neuroradiology, Klinikum rechts der Isar,*
TUM School of Medicine and Health, Technical University of Munich, Munich, Germany
[4]*Department of Neuroradiology, Uniklinikum Erlangen,*
Friedrich-Alexander-University Erlangen-Nuremberg, Erlangen, Germany
[5]*Research Group in Medical Imaging, SURA Ayudas Diagnósticas, Medellín, Colombia*

Abstract

The human brain is a complex system with high metabolic demands and extensive connectivity that requires control to balance energy consumption and functional efficiency over time. How this control is manifested on a whole-brain scale is largely unexplored, particularly what the associated costs are. Using network control theory, here we introduce a novel concept, time-averaged control energy (TCE), to quantify the cost of controlling human brain dynamics at rest, as measured from functional and diffusion MRI. Importantly, TCE spatially correlates with oxygen metabolism measures from positron emission tomography, providing insight into the bioenergetic footing of resting state control. Examining the temporal dimension of control costs, we find that brain state transitions along a hierarchical axis from sensory to association areas are more efficient in terms of control costs and more frequent within hierarchical groups than between. This inverse correlation between temporal control costs and state visits suggests a mechanism for maintaining functional diversity while minimizing energy expenditure. By unpacking the temporal dimension of control costs, we contribute to the neuroscientific understanding of how the brain governs its functionality while managing energy expenses.

INTRODUCTION

The intricate nature of the human brain, characterized by its extensive connectivity [5, 10, 81] and correspondingly high metabolic demands [58, 65], has led to the assumption that its inherent organization represents a delicate balance between energy consumption and functional efficiency [70]. This notion underscores the role of the nervous system as a regulatory entity [27], orchestrating its functions to fulfill cognitive demands while simultaneously managing energy-related constraints. While prior research has shed light on how the brain manages its energy expenditures at a cellular level [28, 43, 58], the broader exploration of how these mechanisms translate to the regulation of functional costs at a whole-brain scale remains relatively uncharted.

Drawing inspiration from engineering principles, network control theory (NCT) offers a novel perspective on this problem by conceptualizing the brain as a networked control system in order to explain its dynamics [33]. In its most basic form, NCT considers brain dynamics as a composite outcome of a region's connectivity profile and the necessary control inputs to guide neural activity toward a desired state [33, 39, 60] (Fig. 1a). The former aspect delves into the constant anatomical interactions between various brain regions, while the latter presents an adaptable measure to optimally transition between states within the confines of energy constraints (Fig. 1b). These constraints are quantified as control energy, which measures the costs entailed in controlling the brain across diverse states (Fig. 1c).

Control energy has been shown to be a powerful tool in various domains of neuroscience, providing an explanation of how the enhancement of executive functions is supported by structural changes across development [17], explaining the effects of psychedelics on brain dynamics [71], or predicting cortical responses to electrical stimulation [79]. Nevertheless, studies employing NCT traditionally focus on the control energy between single pairs of states without accounting for how such costs organize in time. Furthermore, control energy remains a purely statistical notion [78], dissociated from the actual energetic currency employed by the brain (e.g., glucose and oxygen metabolism), with only recent efforts attempting to establish a connection between control costs and metabolism by showing that relative hemispheric differences in glucose uptake are mirrored by differences in control energy for temporal lobe epilepsy patients [36].

Despite these advances, the costs associated with controlling neurotypical resting state activity as well as their energetic signature remain elusive. In light of these considerations, we propose a novel framework to examine the expenses incurred in regulating resting state dynamics, i.e., activity over time. Building upon the findings of previous reports [36], we further investigate whether theoretically derived control costs manifest in tangible biological metrics of energy expenditure by comparing them to normative measures of energy

* eric.ceballosdominguez@mail.mcgill.ca
† valentin.riedl@tum.de

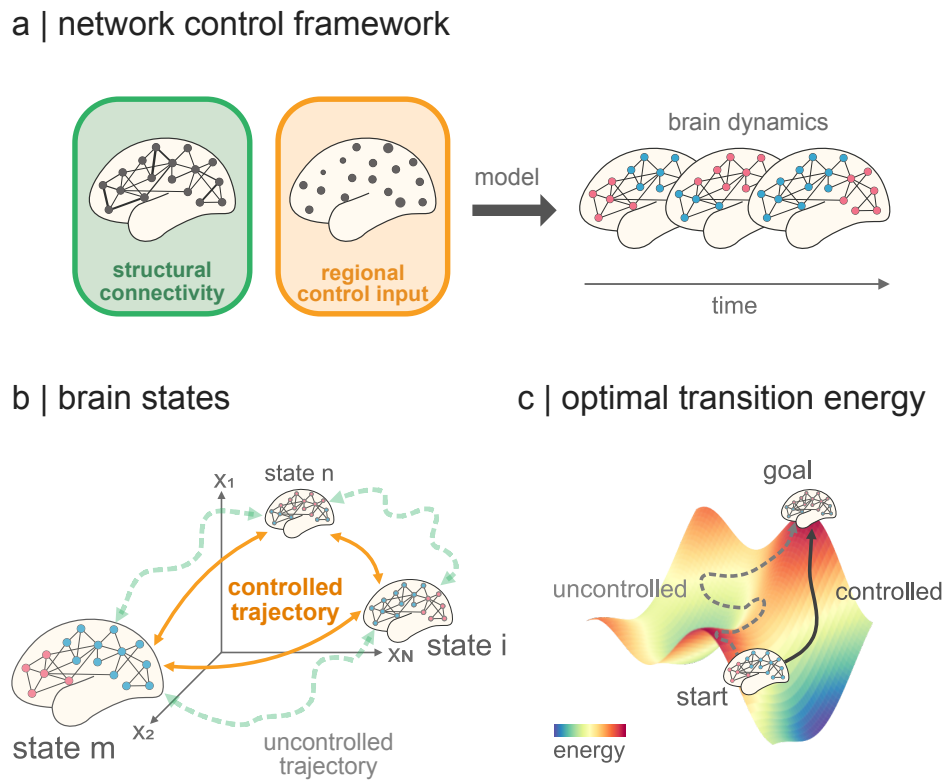

Figure 1. **Network control theory |** (a) The brain as a network of connected regions fluctuates through different modes of activity across time. We can model these transitions throughout time using network control theory. Network control theory is a framework that uses information about the anatomical connections in the brain, which are fixed, and regional control inputs, which are optimised, as bases to model the temporal evolution of brain activity (highlighted by the changing node color patterns on the right). In other words, the changes in brain activity of a particular region are modelled as depending on its connectivity profile and an "activation" that spreads out along such profile, but also spreads in from other regions. Such activation is referred to control input and is fitted to represent an optimal bridge between the brain's current and next state [60]. (b) Before employing a control model, we define what a brain state is, i.e., how maps of brain activity are represented. In previous literature, states have been defined as functional intrinsic networks [15, 17, 36], cytoarchitectonical hierarchy levels [61], or cognitive maps [48], to name a few. These states are represented as single points in an N-dimensional state space, with N being the number of brain regions that we control. We seek to find the shortest possible trajectory between two points in this space and generate activations such that the brain is steered, or "controlled" along a previously optimised trajectory. (c) We define the current and next state of brain activity to be related by the structural connectome of an individual (see *Methods*). The connectome imposes a resistance, or flow, into the state space such that it is easier to make activity changes between states in one direction than another. Think of walking on a hill where the walk uphill will be more strenuous than the downhill walk. In analogy to this, it costs the brain more energy, called "control energy", to control its activity between particular states than others. This asymmetry has been the subject of study in previous literature [48, 61]

metabolism. Moreover, we extend our methodology to investigate organizational principles in brain dynamics through control costs. Ultimately, our efforts seek to establish a foundation for comprehending how the brain governs its functionality and the associated costs entailed in this intricate control.

RESULTS

Here, we present our methodology to study the control costs of human brain dynamics, i.e., activity changes over time, using network control theory. We begin by describing our analysis pipeline and the modelling choices made to simulate the regional energy expenditure based on anatomical and functional properties from diffusion and functional magnetic resonance imaging (MRI) data from $n = 327$ participants in the Human Connectome Project Young Adult dataset [87]. Our work introduces a new measure of control energy over time, time-averaged control energy (TCE), which we relate to actual measurements of energy metabolism based on positron emission tomography (PET). Finally, we conclude our results by investigating the transition costs between functional hi-

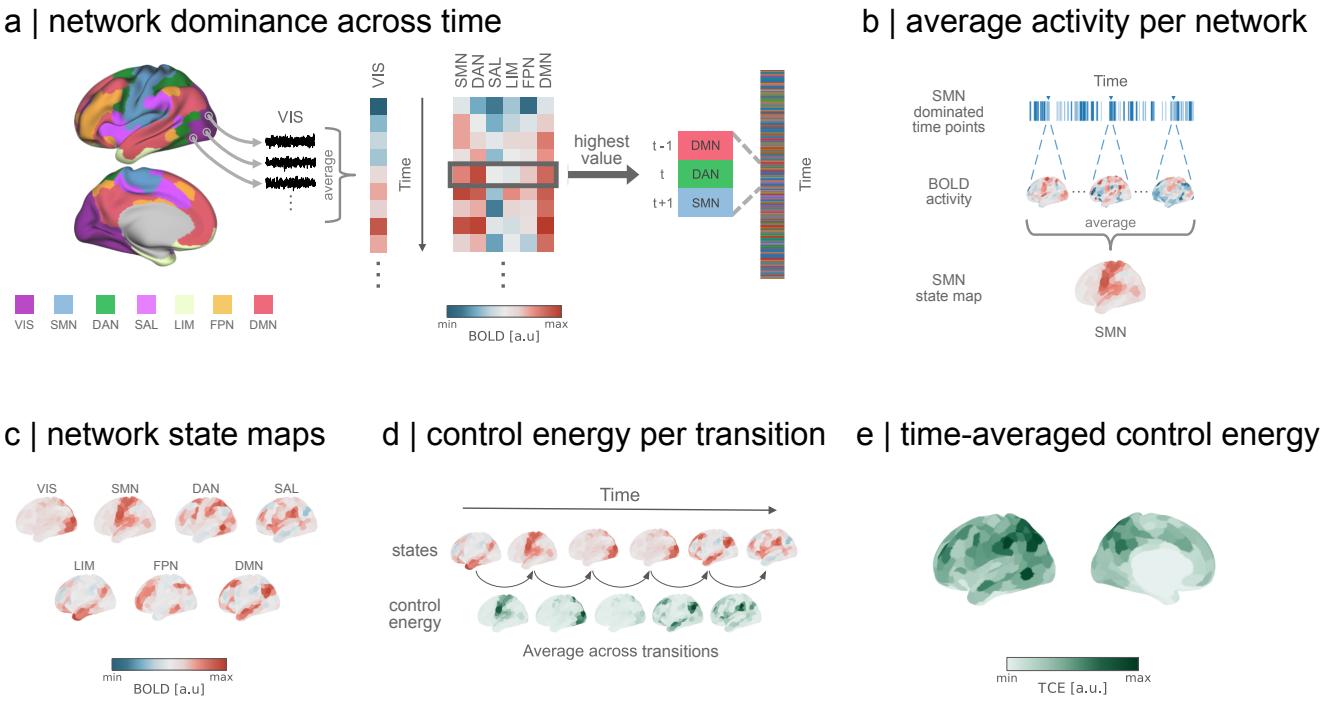

Figure 2. **Estimating the costs of brain dynamics |** (a) We begin by averaging the activity of all regions belonging to one of seven canonical intrinsic networks of Yeo-Krienen [95] in the Schaefer-parcellation [69] and compare the amplitude in each time point. For every time point, the network with the highest amplitude is assigned as the dominating network in that moment, yielding a sequence of network dominance across the time (right). (b) Using the time markers of dominance for each network, BOLD activity maps are averaged to generate a representative state map of activity. (c) We generate seven representative network state maps, reflecting the average magnitude and sign of cortical activity during each network activation. (d) Following the sequence of network dominance, we simulate the optimal control strategy in which the brain transitions between pairs of network states. Thus, each transition yields a map of momentary control costs to move between states. We finally average all control maps across transitions. (e) Our framework results in a map of time-averaged control energy for each individual, where the value of each region represents the amount of control energy integration a region saw throughout the recording on average.
VIS: visual network; SMN: somatomotor network; DAN: dorsal attention network; SAL: salience ventral network; LIM: limbic network; FPN: frontoparietal network; DMN: default mode network; BOLD: blood-oxygen-level-dependent; TCE: time-averaged control energy.

erarchical levels in the brain and find a dynamical rule whereby brain switch frequency within and between levels of the cortical hierarchy is inversely proportional to its control cost.

Simulating the control costs of resting state dynamics

We take a multi-step approach to simulate the control costs of human brain dynamics. In brief, our procedure consists of identifying the dominant networks at each time point (Fig. 2a). This information is used to estimate the average activity across time-points where each network dominates (Fig. 2b) in order to create state maps for each intrinsic network (Fig. 2c). We then use the sequence of network dominance to track the control energy required to transition through the respective states (Fig. 2d). Finally, we average across all transitions, yielding an estimate of the average costs to control brain activity over time (Fig. 2e).

Our analysis begins by delineating intrinsic networks as brain maps. When seeking to model task-related brain states, it is possible to rely on statistical maps of parameter estimates [8] or meta-analytic maps [48]. However, neither of these approaches is feasible for modelling brain activity at rest. Previous studies have taken two approaches to model resting states: On the one hand, states have been defined according to a priori specified parcellations, where areas that fall into a given intrinsic network are represented with ones and all other areas with zeros [17, 36]. On the other, resting states have been derived from clustering techniques and then categorized based on their similarity to binary intrinsic network maps like in the previous approach [15, 71].

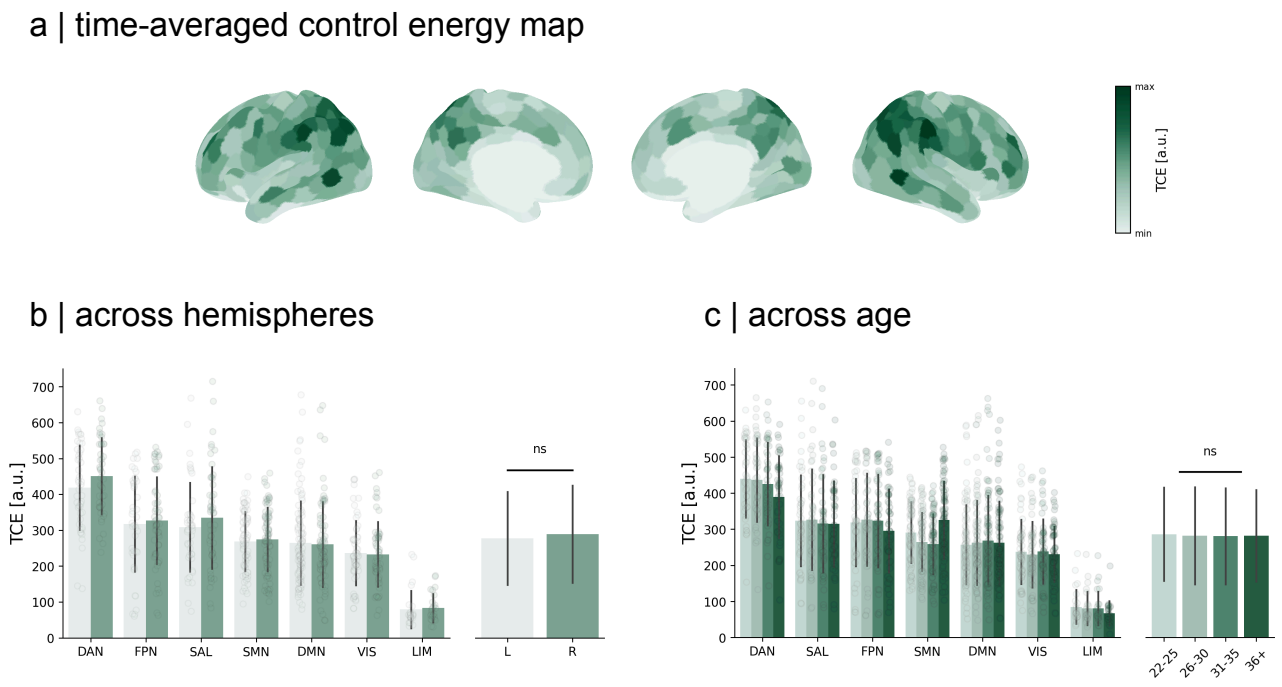

Figure 3. **Time-Averaged Control Energy |** (a) Averaging all TCE maps across individuals, we observe that the costs to control resting state dynamics are regionally heterogeneous. (b) Regions in the dorsal attention network require the most control with limbic regions costing the least. In general, temporal control costs display no preference for a hemisphere, with comparable values evident in both the left and right hemispheres [Mann-Whitney U = 76390, P = 0.27]. (c) When stratifying our analysis by age, we see no statistical difference in temporal control costs across ages groups [ANOVA F(3, 1596) = 0.133, P = 0.94]. Each dot represents the subject-averaged value for a region-of-interest. Bars around the mean indicate standard deviation of the data.

While both approaches offer a respective advantage, they also come with their drawbacks. The former approach profits from the interpretability of previously established brain networks with defined neurocognitive functions. However, such binary, discontinuous state representations are equally defined for all individuals, therefore losing important unique information about the subject. The latter approach circumvents this by using data-driven clustering techniques to derive states from individual data. However, given the possibility of multiple solutions in clustering [23, 37, 41], there is no guarantee that any representation of intrinsic networks are present, thereby complicating the interpretability and replicability of findings across studies. We thus propose a new approach that preserves the intelligibility of intrinsic networks as states, is non-binary, and tailored to individual data.

Our approach relies on an a priori definition of intrinsic networks, such as the canonical intrinsic networks proposed by Yeo-Krienen [95], which are widely-used in the neuroimaging literature. We begin by partitioning the voxel-wise BOLD imaging data into parcel time series representing the time-activity of a region from an intrinsic network. The time series are then z-scored across time and later averaged with other time series from the same network (Fig. 2a). This yields a network-wise average activity for each time point in the recording. Intrinsic networks are then compared at each time point based on their average activity and the network with the highest activity is designated the dominating network of that time point. To control for noise, we discard time points where no network activity surpasses a fixed threshold set at half a standard deviation in average activity. We redo our analysis with a lower threshold and also leave out the network modelling altogether to validate our approach (see *Sensitivity and robustness analysis*)

With the identified time points of dominance for each intrinsic network in hand we proceed to compute the actual state representations of networks. The dominance labels serve as temporal markers to define which time points belong together. We average the activity within such markers to yield a centroid-like representation that resembles an intrinsic network, with a continuous magnitude and sign that is captured by the activity of each individual (Fig. 2b). These state coefficient maps represent each intrinsic network in the a priori defined parcellation, with additional information about regions outside the network that co-fluctuate with it (Fig. 2c). Essen-

tially, we are defining network representations for each individual by tailoring the network weights to their brain activity.

To simulate the control energy between each network coefficient map, we leverage structural information about the anatomical connections of each individual, derived from diffusion MRI. As referenced earlier, network control methods require state maps to represent initial and target positions from where we simulate the control costs to move between them. Such costs are subject to anatomical constrains imposed by the individual's connectome. In other words, we quantify the effort required in each region to change its current activity to a desired state, with other connected regions also contributing to its change [33, 39, 60]. To maintain biological plausibility, we set our simulated path to follow an optimal trajectory that balances the energy required to enable a transition with the proximity to the desired target state [33]. More simply, we simulated transitions that focus on energy minimization while imposing constraints to prevent them from deviating into implausible, non-physiological pathways that may otherwise provide a low-energy pathway.

The results of our modelling of pairwise state transitions manifest in an asymmetric matrix representing the energy cost required to transition the entire brain from one network state to another (Suppl. Fig. S2). The inherent asymmetry in these costs arises from the directional flow imposed by the structural connectivity within the state space. In another context, the states we have modeled could be thought of lying on peaks and valleys, where access is significantly easier in one direction than the other — a phenomenon analogous to the contrast between ascending and descending a hill, with the former requiring more energy.

Having captured the control cost of the transition between all pairs of resting states, we now reconcile this information with the temporal sequence of the dominant networks. This means that we use the network sequence to simulate the energy required for the transitions between states as they occurred during the recording (Fig. 2d). We aggregate the transition energy over all pairs of states in the recording sequence and then normalize the cumulative energy by dividing it by the number of transitions, i.e., we average the overall energy across transitions. This normalization ensures that our measurement remains comparable across different recording durations. Ultimately, our approach culminates in the derivation of a metric we refer to as time-averaged control energy, or TCE (Fig. 2e).

TCE is an estimate of the temporal control costs associated with idiosyncratic intrinsic network activations across time (Fig. 3a). It is a spatially heterogeneous measure but we find that its distribution in left and right hemispheres is statistically indistinguishable [Mann-Whitney U = 76390, P = 0.27](Fig. 3b). Moreover, previous evidence indicates that control energy is modulated by age [17, 80]. We find that stratifying by age (using bins of 22-25, 26-30, 31-35, and 36+, with 247, 527, 418 and 14 individuals, respectively) does not lead to any statistically significant group differences [Levene W = 108, P = 0.955; ANOVA F (3, 1596) = 0.133, P = 0.94] (Fig. 3c). Taken together, we show that the control costs of human brain dynamics are region-specific, consistent between hemispheres, and conserved across ages studied here.

Grounding control costs in metabolism

Previously, network control theory metrics were shown to correlate with physiological markers of energy metabolism [36], thereby establishing an initial bridge between theoretical and tangible indicators of energy. As such, we contextualize our results with neurobiological maps to look for energetic correlates of temporal control costs. To accomplish this, we scrutinize the relationship between temporal control costs and two pivotal substrates of metabolic energy: oxygen and glucose metabolism [13, 65] (Fig. 4).

Our initial investigation involved comparing TCE to the regional variations in cerebral glucose metabolism (CMR_{glc}) across the cortex. To achieve this, we employed a group-average template derived from an internally acquired dataset [14] comprising n = 20 subjects who underwent PET scans, utilizing ^{18}F-Fluorodeoxyglucose (FDG) as a tracer to monitor their glucose uptake (see *Methods*).

Given the inherent spatial smoothness of our PET map, we employ rigorous measures to validate our results by contrasting them against 10,000 null maps that replicate the variogram of the data [12] (see *Methods)*. These null maps randomize the values in our PET map while accounting for the spatial autocorrelation inherent in the recordings, enabling us to distinguish true effects from spatially induced artifacts using a modified P-value (P_{SMASH}) [12, 51]. Our results showed a weak correlation between TCE and CMR_{glc}, which, however, did not exhibit statistical significance when the spatial smoothness of the data set was taken into account [Spearman ρ = 0.144, P = 0.003, P_{SMASH} = 0.347].

Subsequently, we investigate the relationship between TCE and cerebral oxygen uptake within the cortex. Here, we employed a normative map of cerebral oxygen metabolism (CMR_{O2}) derived from [86]. This map was computed from $n = 33$ participants who underwent PET scans to estimate their cerebral oxygen consumption, utilizing ^{15}O-labeled oxygen (see *Methods*). We note that the map is not quantified into cerebral metabolic units; however, we retain the name for oxygen uptake as CMR_{O2}, in accordance with the original authors' nomenclature for their activity concentration.

Given the same technical limitations as with the CMR_{glc} map, we generate 10,000 null maps that preserve the spatial autocorrelation of the original data using the same variogram-matching algorithm [12]. Our results

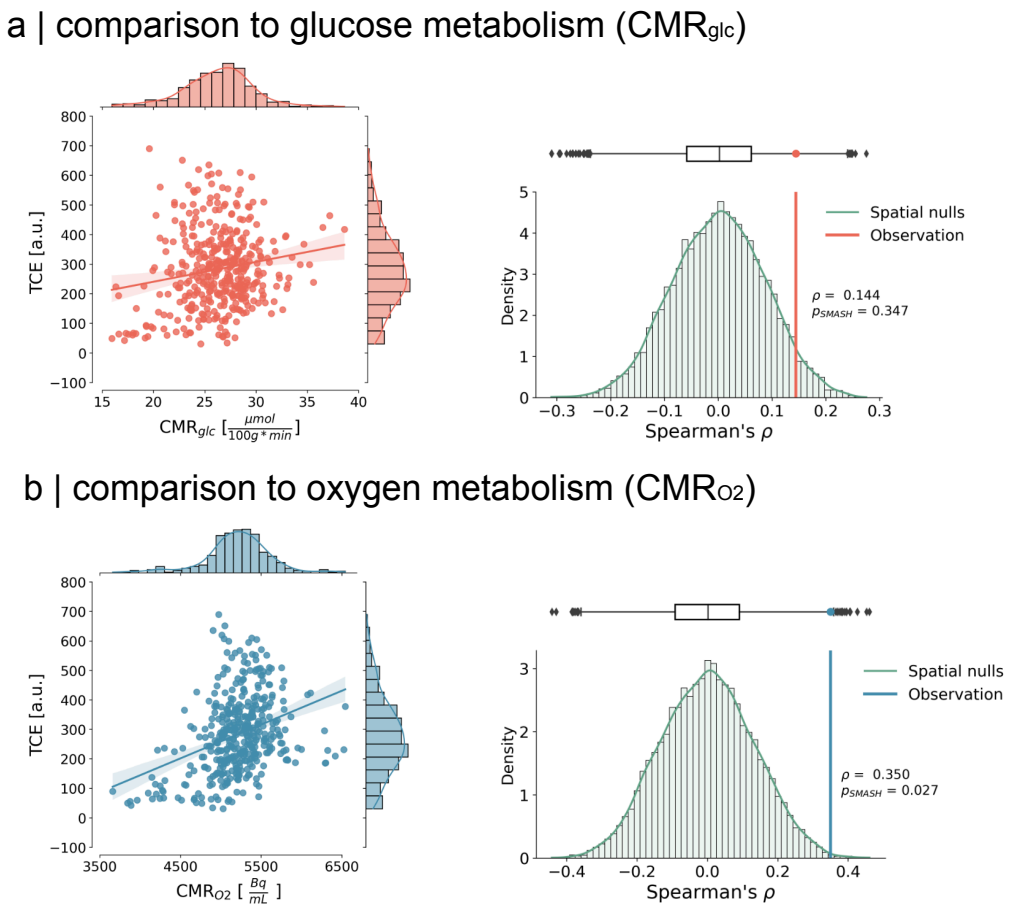

Figure 4. **Relationship to metabolism |** (a) Left: We compared how time-averaged control energy spatially co-varies with glucose metabolism from an internally acquired positron emission tomography (PET) dataset [14]. Glucose uptake has been previously shown to scale with momentary control energy [36]. We, however, do not find a similar significant relationship to control energy when considering its temporal dimension [Spearman $\rho = 0.144$, $P_{SMASH} = 0.347$]. Shades around regression line represent 95% confidence interval with 1000 bootstrap samples. Right: Distribution of spatial nulls. Correlation results accounted for spatial autocorrelation in the data by using variogram-matched nulls [12, 51, 88]. (b) Left: Contrary to glucose metabolism, we found that TCE is spatially similar to a representative cortical distribution of oxygen uptake from PET imaging [86] [Spearman $\rho = 0.35$, $P_{SMASH} = 0.027$]. Shades around regression line represent 95% confidence interval with 1000 bootstrap samples. Right: Distribution of spatial nulls. Correlation results accounted for spatial autocorrelation in the data by using variogram-matched nulls [12, 51, 88].

show that TCE significantly correlates with CMR_{O2}, and that this relationship is significant even after accounting for the spatial autocorrelation in the data [Spearman $\rho = 0.35$, $P < 0.001$, $P_{SMASH} = 0.027$]. Altogether, we find that the regional cortical pattern of temporal control costs resembles the spatial pattern of oxygen uptake, establishing further evidence about the link between biological and control energy.

Temporal organisation of control costs

We now turn to a new perspective on brain states by reframing our analysis from the perspective of functional hierarchical processing as defined by Mesulam [54]. Previous work has shown that control energy differs across cortical hierarchies [61]. In this context, we examine the control costs of switching within and between hierarchical levels across time, ultimately looking at their dynamical organisation in time.

We begin by mapping the sequence of dominant networks to their respective levels within the functional hierarchy of the brain (Fig. 5a). Visual and sensorimotor networks are hereafter referred to as "unimodal", a distinction that arises from their specialized processing that assumes specific inputs[49, 54]. In turn, we refer to all other networks as "heteromodal" (associative) as they integrate multiple functional streams [54]. Bi-

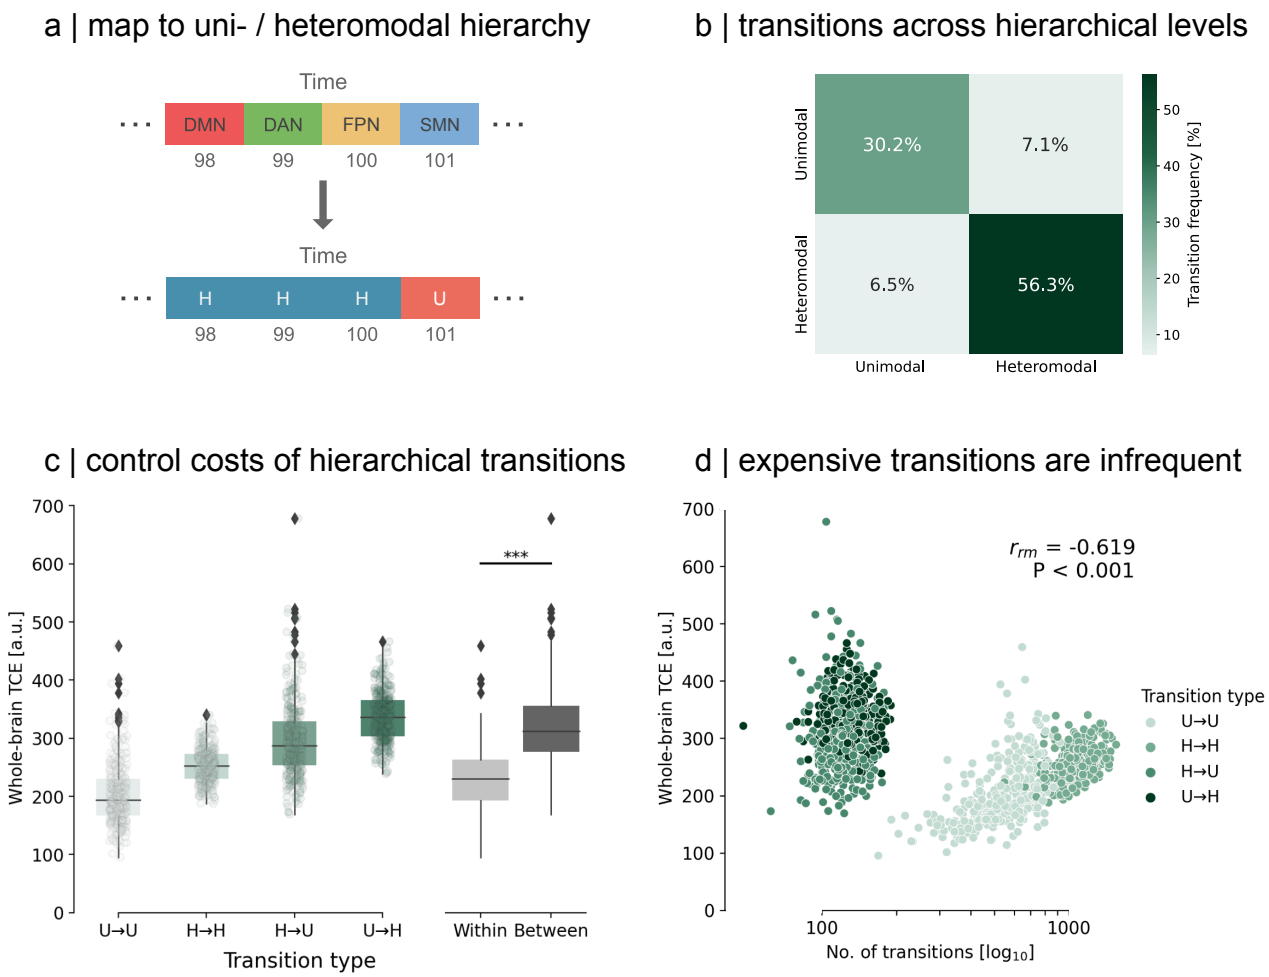

Figure 5. **Temporal organisation of control costs** | (a) We map our sequence of dominating networks to their respective hierarchical level. This means sensory networks are jointly labelled as unimodal networks (U), with the remaining networks defined as heteromodal, association networks (H). (b) Transitions within hierarchical levels represent the majority of transitions (86.5%) at rest. (c) Averaging control costs across the whole brain, we observed that transitions within a hierarchical level required significantly less TCE than between. Asterisks indicate high significance [Mann-Whitney U = 376072, P < 0.001]. Each dot per transition type represents a subject. (d) We relate whole-brain control costs to transition frequency using repeated measures correlation [4] to assess common intra-individual relationships and find that the number of transitions per transition type is negatively correlated to the cost needed to perform such switch [r_{rm} =-0.619, P < 0.001].

modal mapping ensures consistency in our subsequent validation with a different parcellation that defines intrinsic networks differently (Suppl. Fig. S5). Moreover, we show that adopting this new perspective provides a fresh view of the dynamics of the brain, uncovering a relationship between TCE and state switching behavior. We follow the same framework to compute TCE, but now separately average transitions within unimodal networks, i.e., from unimodal to unimodal, and within heteromodal networks, as well as transitions between unimodal and heteromodal levels.

After mapping the sequence of dominating networks, we observe that, on average for individuals, transitions occur predominantly within a hierarchy rather than between hierarchies (Fig. 5b). Analogously, the control costs to transition between hierarchies are more demanding than those within a hierarchical level [Mann-Whitney U = 376072, P < 0.001](Fig. 5c). Using repeated measures correlation [4, see *Methods*] to account for repeated measures of control costs for each participant as well as their individual traits, we examine the commonalities in the relationship between participants' control costs and transition frequency across hierarchies. We find a robust negative correlation throughout the cohort [r_{rm} =-0.619, P < 0.001], indicating that more frequent transitions within a hierarchy are associated with

lower average control energy costs, in stark contrast to infrequent, costly transitions that bridge different hierarchical domains. We show that this relationship is not an artefact of the network partition in our parcellation, but is also evident when using alternative partitions (see *Sensitivity and robustness analysis*). Altogether, our findings suggests a dynamic temporal principle within the brain that strives to minimize its control costs by keeping expensive transitions sparse.

Sensitivity and robustness analysis

To ensure the robustness of our results, we replicate our analyses on an independent dataset [14]. This additional dataset has lower spatial and temporal resolution and is subjected to different preprocessing protocols (see *Methods*). Indeed, even in the face of these differences, our results remain unchanged (Suppl. Fig. S6).

We then repeat our analysis using an alternative cortical parcellation [30]. This approach seeks to validate the reproducibility of our results with a different definition of intrinsic networks. We find that this approach leads to consistent results (Suppl. Fig. S5). Additionally, we investigate the impact of a coarser partition by dividing the cortex into 200 regions instead of the original 400 regions. We find that this modification does not affect our results (Suppl. Fig. S4).

Further, we test the impact of our state modelling step by leaving it out and instead simulate the control costs to transition between time points of raw activity, i.e., our state maps are represented as the BOLD activity at each time point. We form a group-average template like in our main results and compare these results to the temporal signal-to-noise ratio (tSNR) of each region. tSNR is a metric to quantify the quality of a BOLD signal, calculated as the mean of the signal prior to demeaning and divided by the signal standard deviation [92]. We compute the tSNR for each participant run in the HCP dataset, re-scale it from 0 to 1 using min-max normalization, and finally average it across runs.

We find that, when leaving out the state modelling step, regions with high TCE are also regions that exhibit a low temporal signal-to-noise ratio (tSNR) [Spearman ρ = 0.36, P_{spin} = 0.042]. In contrast, using our approach we observe that the group-level results are unlikely related to the group-level regional tSNR [Spearman ρ = 0.234, P_{SMASH} = 0.232]. Following our main approach, we discard time points with average activity below 0.5 but observe that, using a lower threshold of 0 instead, our group-level results again show no significant relationship to regional signal quality [Spearman ρ = 0.274, P_{SMASH} = 0.099]. Moreover, we observed on an individual level that a larger extent of participants' TCE maps were significantly related to their tSNR when leaving out our modelling approach (Fig. S1). Taken together, this suggests that modelling states using our approach mitigates the influence of noise in our analysis.

We conclude by examining whether temporal control costs are related to the occurrence of resting states. This would mean that the patterning of individual TCE maps are a reflection of the most frequently occurring states. To test this, we create a representative map of the occurring intrinsic networks by counting how frequent, i.e., how many time points, each state occurs. Dividing this number by the total number of time points yields an estimate of how much each network was present in the recording, e.g., the visual network dominated 20% of the time. We multiply this ratio with its corresponding state coefficient map, thereby creating time-weighted maps which we can sum together to create a map that represents the frequency of each network and its spatial distribution. Interestingly, we find no significant association between this map and our TCE map (Spearman ρ = 0.27, P_{SMASH} = 0.092; Suppl. Fig. S3). Altogether, we conclude that our findings result from a convergence of characteristic sequences of resting states and individual anatomical constraints, which together form the basis to simulate the control costs of resting state dynamics.

DISCUSSION

In the present study, we introduce a metric for assessing the control costs inherent in human brain dynamics, referred to as time-averaged control energy (TCE). TCE serves as a composite measure quantifying the average control exerted by an individual to sustain dynamic neural activity over time. Notably, our findings reveal a statistically significant spatial resemblance between the distribution of TCE and cortical oxygen metabolism, a major energetic substrate crucial for compensating neural activity [35]. Our results further show that transitions within uni- or heteromodal networks entail distinct temporal control costs as opposed to transitions between modalities, thereby offering insight into the temporal organisation of brain dynamics. Importantly, we provide open-access code to facilitate the replication and extension of our results in subsequent studies.

Temporal control is a confluence of anatomical constraints and state sequences

Network control theory is a powerful tool to investigate the emergence of dynamics from the structural scaffold of the brain [60]. Its application in neuroscience has yielded valuable insights, e.g., into individual differences in psychiatric conditions [48, 62] or developmental changes in youth [17, 80]. Notably, controllability is reported to increase asymptotically and plateau with adolescence [80], potentially explaining the absence of age effects observed in our results.

We find that the distribution of temporal control costs is spatially distinct, with certain brain regions requiring higher TCE than others (Fig. 3). The extent of

engagement in this control dynamic may be influenced by the manner in which control inputs propagate, reminiscent of the role of connectivity profiles in facilitating diverse spreading dynamics across brain regions [56]. This spreading process may parallel the flow of control throughout the connectome [78].

Control costs are further influenced by the specific states the brain traverses. Our results corroborate previous findings [48, 61] indicating that control energy is dependent on both initial and goal states (Fig. 5; Suppl. Fig. S2). This underscores the significance of our state-modelling approach that determines the sequence of network dominance based on intrinsic network activations, thereby establishing the order of state transitions.

Our methodology demonstrates robustness across different partitions of intrinsic networks (Suppl. Fig. S5, S4), as well as separate datasets (Suppl. Fig. S6). Nonetheless, our state-modelling approach trades off robustness for temporal granularity, thereby losing the ability to capture temporal shifts in brain activity. Indeed, brain dynamics are hypothesized to be temporally organised such that high-amplitude modes of activity are counterbalanced by transient 'off' states [67, 90, 99]. We prompt future investigations to experiment with alternative state-modelling methods, like sliding-window approaches, for improved temporal precision. Appropriate tests should, however, be developed in tandem to distill real fluctuations from spurious ones [42, 97].

Finally, we note that our model of network control is linear, therefore assuming that linear dynamics govern brain transition [39]. While nonlinear models of neural dynamics have shown promise in simulating brain activity [9], the use of first-order linear approximations has proven sufficient to model activity at a macroscale [57]. Furthermore, recent studies have showcased how such linear control models are expandable to account for biological heterogeneities in receptor availability [71, 72] or cortical abnormalities in patient populations [48]. Consequently, we see potential for future efforts to expand on our work in order to investigate how temporal control costs are modulated by incorporating additional biological insights or clinical distinctions.

Control costs reflect metabolic features of the brain

Our results show that temporal control costs co-localize with oxygen metabolism, but not glucose metabolism (Fig. 4). A previous study from He et al. [36] showed that control energy and glucose uptake are similarly put out of balance between hemispheres in temporal lobe epilepsy (TLE) patients. The authors, however, only show this relationship in the limbic system and, moreover, their methodology does not consider the costs of temporal switching, but seeks to link relative hemispheric asymmetries in momentary control costs and glucose uptake to explain structural changes in TLE patients. In contrast, our study relates control and metabolism on an absolute whole-brain basis and considers the temporal dimension of control costs.

Moreover, while both glucose and oxygen constitute fundamental components for meeting the energetic demands of neural activity [13, 35, 44, 53, 86, 91], they exhibit distinct spatial patterns. The predominant energy-producing pathways in the brain encompass glycolysis and oxidative metabolism, where glycolysis relies on glucose and oxidative metabolism relies on glucose-derived pyruvate and oxygen [13, 21]. The disparate spatial distribution of glucose and oxygen metabolism is attributed to their dissimilar upscaling capacities for energy generation [20, 22, 96]. Specifically, certain regions are posited to favor glycolysis beyond basal levels of oxygen and glucose consumption at an equable ratio in order to afford greater flexibility in energy demand necessary for supporting human cognition [47, 86]. In light of these considerations, our association of temporal control costs with the spatial distribution of oxygen, rather than glucose metabolism, supports the interpretation that control of human brain dynamics is grounded in baseline metabolic activity constituent of oxidative energy generation.

We note, however, that our conclusion is based on group-level comparisons of control costs and metabolic maps. Further research is warranted to investigate the nuanced relationship between control costs and energy metabolism at an individual level. Advances in MR imaging now enable the simultaneous acquisition of BOLD and oxygen metabolism data from the same subjects [24], providing an avenue for in-depth exploration of this relationship. Finally, we underline that our results are not confounded by the fact that we use BOLD-imaging data, as the BOLD signal reflects blood flow, which represents oxygen availability rather than necessarily indicating oxygen consumption [26].

State switching is grounded by control costs

We observed that the costs required to transition within and between hierarchical levels, i.e., unimodal or heteromodal, were significantly different (Fig. 5c,d). Similarly, this phenomenon was mirrored in the frequency with which the brain either transitions out or persists in a hierarchical level (Fig. 5b,d). This means that costly transitions between hierarchical levels were rare, with more efficient transitions within hierarchical levels happening more frequently.

In order to enable information flow, the human brain is posited to switch between modes of integration and segregation over time, where segregation refers to regions clustering together to form functionally distinct modules and integration represents a form of global intercommunication between modules [18, 25, 70]. Building upon this, we propose that the transitions between functional hierarchical levels may reflect an integration-segregation dynamic. In this sense, switches between hi-

erarchical levels might represent modes of communication among functional hierarchies, counterbalanced by switches within hierarchical levels as modes of segregated activity. As previous studies have shown the disruption of integration-segregation dynamics by pathology or altered states of consciousness [45, 46], future works could investigate whether such conditions equally disrupt state switching dynamics.

Furthermore, we hypothesize that the inverse correlation between control costs and state visits represents a mechanism where the brain seeks to maintain functional diversity, while minimizing its energy expenditure. This is similar to the hypothesis put forward by Zalesky, et al. [98], where the authors posit that the brain transitions through costly, intermittent states in order to facilitate global integration while keeping energy demand minimal. As such, we encourage future efforts to further elucidate our understanding of the intricacies governing state switching dynamics and control costs.

In conclusion, our work presents a new perspective of control energy as a measure of the costs to control human brain dynamics. We are able to show that temporal control costs are spatially related to oxygen metabolism as an energy substrate; and uncover an organisational principle that relates temporal dynamics to their costs in the brain. As such, we envision that our work can help advance research on the costs of regulating brain activity.

METHODS

All preprocessed data is available at https://osf.io/nw9zt. The code and additional data used to perform the analyses are available at https://github.com/NeuroenergeticsLab/control_costs.

Data

Human Connectome Project (HCP)

The main data used for this study consisted of resting state time series from functional MRI (fMRI) and structural connectomes from diffusion MRI (dMRI) taken from the Human Connectome Project S900 Young Adult release [87]. Scans from 327 unrelated participants (mean age 28.6 ± 3.73 years, 55% females) were used to ensure that familial factors do not confound our analysis [55]. Informed consent was obtained for all subjects (the protocol was approved by the Washington University Institutional Review Board as part of the HCP). The participants were scanned in the HCP's custom Siemens 3T "Connectome Skyra" scanner, and the acquisition protocol included four 15-minute resting state fMRI sessions and a high angular resolution diffusion imaging (HARDI) sequence. The resting state fMRI data was acquired using a gradient-echo EPI sequence (TR = 720 ms; TE = 33.1 ms; FOV = 208 × 180 mm^2; voxel size = 2 mm^3; number of slices = 72; and number of volumes = 1200). The dMRI data was acquired with a spin-echo EPI sequence (TR = 5520 ms; TE = 89.5 ms; FOV = 210 × 180 mm^2; voxel size = 1.25 mm^3; b-value = three different shells i.e., 1000, 2000, and 3000 s/mm^2; number of diffusion directions = 270; and number of b0 images = 18). Additional information regarding the acquisition protocol can be found under [87].

To process the functional data, each run of each subject's resting state fMRI recording was pre-processed in terms of gradient distortion correction, motion correction, and spatial normalization according to [76] and [29]. Artefacts were then removed using ICA-FIX [68]. Inter-subject registration of cerebral cortex was carried out using areal feature-based alignment and the Multimodal Surface Matching algorithm [66]. The clean time series were then parcellated into 400 and 200 cortical region time series according to the Schaefer functional atlas[69].

To process the diffusion data, structural connectomes were reconstructed from the dMRI data using the MRtrix3 package [84]. Grey matter was parcellated into 400 and 200 cortical regions according to the Schaefer functional atlas [69] and fiber orientation distributions were generated using a multi-shell multi-tissue constrained spherical deconvolution algorithm [19, 38]. The initial tractogram was generated with 40 million streamlines, with a maximum tract length of 250 and a fractional anisotropy cutoff of 0.06. Spherical-deconvolution informed filtering of tractograms (SIFT2) was used to reconstruct whole brain streamlines weighted by cross-section multipliers [74]. More information regarding the individual network reconstructions is available in [59].

TUM Dataset

The replication data consisted of simultaneously measured FDG-PET and fMRI, with subsequent dMRI recordings, all while the participants kept their eyes open. Scans from 20 participants (mean age 34.2 ± 5.99 years, 50% females) were used. Informed consent was obtained for all subjects (the protocol was approved by the Technical University of Munich Review Board). The participants were scanned an integrated PET/MR (3T) Siemens Biograph mMR scanner (Siemens, Erlangen, Germany) and used a 12- channel phase-array head coil for the MRI acquisition. The PET data were collected in list-mode format with an average intravenous bolus injection of 184 MBq (s.d. = 12 MBq) of [18F]FDG. In parallel to the PET measurement, automatic arterial blood samples were taken from the radial artery every second to measure blood radioactivity using a Twilite blood sampler (Swisstrace, Zurich, Switzerland).

The functional MRI data were acquired during a 10 min time interval using a single-shot echo planar imaging sequence (300 volumes; 35 slices; repetition time, TR = 2000 ms; echo time, TE = 30 ms; flip angle, FA =

90°; field of view, FOV = 192 × 192 mm2 ; matrix size = 64 × 64; voxel size = 3 × 3 × 3.6 mm3). Diffusion-weighted images were acquired using a single-shot echo planar imaging sequence (60 slices; 30 non-colinear gradient directions; b-value = 800 s/mm² and one b=0 s/mm² image; TR = 10800 ms, TE = 82 ms; FA = 90°; FOV = 260 x 264 mm²; matrix size = 130 x 132; voxel size = 2 x 2 x 2 mm3). Additional information regarding the acquisition protocol can be found under [14].

To process the functional data, each subject's resting state fMRI recording was pre-processed using the Configurable Pipeline for the Analysis of Connectomes (CPAC v1.4.0) [16]. This included slice-timing correction, motion correction, spatial normalization, quadratic and linear detrending of scanner drift, anatomical CompCor regression of white matter and cerebrospinal fluid activity [6], and subsequent bandpass-filtering (0.01 - 0.1 Hz). More details can be found in [14].

To process the diffusion data, structural connectomes were generated using the MRtrix3_connectome BIDS App [31], which operates principally using tools provided in the MRtrix3 package [84]. This included: DWI denoising [89], Gibbs ringing removal [40], preprocessing [1, 2], bias field correction [85], inter-modal registration [7], brain extraction [75], T1 tissue segmentation [63, 73, 75, 100], spherical deconvolution [38, 83] and probabilistic tractography [82] utilizing Anatomically-Constrained Tractography [73] and dynamic seeding [74]. The resulting fiber track files were subsequently converted into streamline counts by counting the number of streamlines that passed through one of the 400 cortical regions according to the Schaefer functional atlas[69]. In order to compensate for the bias toward longer fibers inherent in the tractography procedure, as well as differences in region size, we normalized the streamline count by the average length of the streamlines and average surface area of the two regions [34]

To process the PET data, the first 45 minutes of the PET acquisition were reconstructed offline using the NiftyPET library [52] based on the ordered subsets expectation maximization (OSEM) algorithm with 14 subsets, 4 iterations, and divided into 33 dynamic frames: 10 x 12 s, 8 x 30 s, 8 x 60 s, 2 x 180 s and 5 x 300 s. The attenuation-correction was based on the T1-derived pseudo-CT images [11]. All reconstructed PET images were motion corrected and spatially smoothed (Gaussian filter, FWHM = 6 mm). The net uptake rate constant (K_i) was calculated using the Patlak plot model [64] based on the last 5 frames of the preprocessed PET images (frames between 20 to 45 minutes) and the arterial input function derived from the preprocessed arterial blood samples. The cerebral metabolic rate of glucose (CMR_{glc}) was calculated by multiplying the K_i map with the concentration of glucose in plasma of every participant, divided by a lumped constant of 0.65 [93]. Finally, the CMR_{glc} maps were partial volume corrected using the GM, WM and CSF masks derived from the T1 images using the iterative Yang method [94] and finally registered to the MNI152NLin6ASym 3mm template through the anatomical image.

Additional data

In addition to the metabolic map of glucose uptake from the TUM dataset, we complemented our analysis with a normative, group-average map of oxygen metabolism (CMR_{O2}) from [86], acquired through neuromaps [50]. The brain map was derived from 33 healthy, right-handed neurologically normal participants (mean age 25.4 ± 2.6 years, 58% females) that were recruited from the Washington University community. The recording was performed using a Siemens model 961 ECAT EXACT HR 47 PET scanner (Siemens/CTI) with 47 slices encompassing an axial field of view of 15 cm. Transverse resolution was 3.8–5.0 mm FWHM and axial resolution was 4.7–5.7 mm full width at half maximum (FWHM). Attenuation data were obtained using 6868GeGa rotating rod sources to enable quantitative reconstruction of subsequent emission scans. Emission data were obtained in the 2D mode (interslice septa extended). The PET data were reconstructed using a ramp filter (≈6 mm FWHM) and then blurred to 12 mm FWHM. Distribution of CMR_{O2} was measured with a 40 second emission scan (derived from a 120 second dynamic scan) after brief inhalation of 60 mCi of [15O]oxygen in room air. More details on the recording parameters and processing can be found in [86].

Network control theory

Network control theory (NCT) is a framework that provides the approach to computing metrics related to the control of brain activity.

Fundamental to NCT is the definition of the brain as a networked system with nodes connected by edges and defined by an adjacency matrix A. Biologically, the matrix A represents the thick bundles of myelinated axonal fibers that run through large-scale connections of brain regions [77]. These fibers are thought to play a critical role in coupling the activity of distant brain regions [3].

By accounting for connections of at the brain-region level, an equation relating brain structure to time-evolving brain activity $\mathbf{x}(t)$ can be formulated as:

$$\dot{\mathbf{x}} = A\mathbf{x}(t), \tag{1}$$

where $\dot{\mathbf{x}}$ and $\mathbf{x}$ are of size N×1 and A N×N, with N is the number of brain regions. In practice, the matrix A is normalized in order to avoid infinite growth, and as such defined as:

$$A_{norm} = \frac{A}{\lambda(A)_{max} + c} - I \tag{2}$$

where $\lambda(A)_{max}$ is the largest eigenvalue of A and c a constant that determines if the system approaches the corresponding mode of A ($c = 0$) or decays towards zero ($c > 0$). Here, c is set to 1. For simplicity, A_{norm} will be referred to as A from now on, assuming that all the following equations make use of a normalized matrix.

Looking at the assumptions made, equation 1 imposes that the temporal evolution of the brain is a *linear* function that is described by its connectivity and its current state in time. In addition to linearity it is important to note that this equation assumes that A is not changing, i.e., the brain is a time-invariant system, and that brain dynamics are noise-free [39, 78].

From here, equation 1 can be extended to account for controlled dynamics that steer the system away from its natural trajectories through an external input. Formally,

$$\dot{\mathbf{x}} = A\mathbf{x}(t) + B\mathbf{u}(t) \tag{3}$$

with $\mathbf{u}(t)$ representing the injected control input into the system at time t and the columns in B representing the total number of voxels or regions of interest to control. Depending on the number of regions m that the modeler decides to control, $\mathbf{u}(t)$ is of size m×1 and, correspondingly, B is of size N×m. Each element in b_{ij} represents the influence that an external input $u_j(t)$ has on region i, where $j = 1, .., m$ and $i = 1, ..., N$. By including controlled dynamics in equation 3 it is now possible to model optimal trajectories between brain states and calculate their optimal control energy.

Optimal control energy

Optimal control energy provides a measure of the controllability of a system in terms of 1) the state trajectory traversed and 2) the work required to reach a state. It further specifies a time horizon T during which the control input $u(t)$ is effective in moving the system from the initial state $x(0) = x_0$ to the goal state $x(T) = x_T$.

Formally, the problem of reaching x_T starting from x_0 while keeping 1) and 2) minimal is as follows:

$$\mathbf{u}(t)^* = \arg\min_{\mathbf{u}} \int_0^T (\Delta\mathbf{x}(t)^\mathsf{T}\Delta\mathbf{x}(t) + \rho\mathbf{u}(t)^\mathsf{T}\mathbf{u}(t))dt \tag{4}$$

with Δ**x**(t) being the distance to the goal state at time t, i.e.,

$$\Delta\mathbf{x}(t) = \mathbf{x}_T - \mathbf{x}(t). \tag{5}$$

Moreover, ρ determines the relative weighting between the costs associated with the length of the state trajectory and input energy. An usual choice is to set $\rho = 1$ as to weight both goals equally [32, 39]. Once equation 4 is minimized, an optimal control energy can be derived from $\mathbf{u}(t)^*$:

$$E^* = \int_0^T \mathbf{u}(t)^{*\mathsf{T}}\mathbf{u}(t)^* dt \tag{6}$$

The optimal energy E^* was computed for each spatial unit (in this case regions from a parcellation). These values therefore represent the quadratic input across model time to optimally move towards a goal state given an initial state.

Readers interested in a more in-depth discussion about the mechanisms of network control theory are invited to read the review from [78], as well as [60]. For more discussions on parameter sensibilities of the model, interested readers are referred to the exhaustive study by [39].

Spatial autocorrelation preserving null model

For all of our hypothesis tests that relied on comparing two brain maps, we employed BrainSMASH [12] as implemented in the neuromaps [50] library. Briefly, the methodology of BrainSMASH entails two steps: 1) random permutation of values within a designated brain map and 2) smoothing and re-scaling to restore the spatial autocorrelation characteristic of the original data.

In the first step, values within the brain map undergo random permutation. Subsequently, the permuted data undergoes a transformation to reintroduce spatial autocorrelation. This transformation is expressed as

$$y = |\beta|^{1/2}x' + |\alpha|^{1/2}z \tag{7}$$

where x' represents the permuted data, $z \sim N(0, 1)$, and α and β are estimated via least-squares optimization, aligning variograms of the original and permuted data.

Repeated measures correlation

Repeated measures correlation (r_{rm}) is a statistical technique developed by Bakdash et al. [4] to examine the common association between variables within each subject, focusing on accounting for within-individual associations in paired measures. This method is particularly valuable when dealing with repeated measures on the same participants — in our case four control cost estimates for each hierarchical transition. Because this method only assesses the common slope across participants, we can investigate whether there is a consistent relationship trend throughout measurements while automatically correcting for individual confounding factors such as the participant's gender or age.

Acknowledgments

We thank Filip Milisav, Justine Hansen, Vincent Bazinet, Asa Farahani, Zhen-Qi Liu, Moohebat Pourmajidian, Laura Suarez, and Markus Ploner for helpful discussion. EGC acknowledges support from the Molson

Foundation, the IFI fellowship of the German Academic Exchange Service and the Elite Network of Bavaria. AIL was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), [funding reference number 202209BPF-489453-401636, Banting Postdoctoral Fellowship]. VR has been funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (ERC Starting Grant, ID 759659).

[1] Andersson, J. L. R., Skare, S., and Ashburner, J. (2003). How to correct susceptibility distortions in spin-echo echo-planar images: Application to diffusion tensor imaging. *NeuroImage*, 20(2):870–888.

[2] Andersson, J. L. R. and Sotiropoulos, S. N. (2016). An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging. *NeuroImage*, 125:1063–1078.

[3] Avena-Koenigsberger, A., Mišić, B., Hawkins, R. X. D., Griffa, A., Hagmann, P., Goñi, J., and Sporns, O. (2017). Path ensembles and a tradeoff between communication efficiency and resilience in the human connectome. *Brain Structure and Function*, 222(1):603–618.

[4] Bakdash, J. Z. and Marusich, L. R. (2017). Repeated Measures Correlation. *Frontiers in Psychology*, 8.

[5] Bassett, D. S. and Sporns, O. (2017). Network neuroscience. *Nature Neuroscience*, 20(3):353–364.

[6] Behzadi, Y., Restom, K., Liau, J., and Liu, T. T. (2007). A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *NeuroImage*, 37(1):90–101.

[7] Bhushan, C., Haldar, J. P., Choi, S., Joshi, A. A., Shattuck, D. W., and Leahy, R. M. (2015). Co-registration and distortion correction of diffusion and anatomical images based on inverse contrast normalization. *NeuroImage*, 115:269–280.

[8] Braun, U., Harneit, A., Pergola, G., Menara, T., Schäfer, A., Betzel, R. F., Zang, Z., Schweiger, J. I., Zhang, X., Schwarz, K., Chen, J., Blasi, G., Bertolino, A., Durstewitz, D., Pasqualetti, F., Schwarz, E., Meyer-Lindenberg, A., Bassett, D. S., and Tost, H. (2021). Brain network dynamics during working memory are modulated by dopamine and diminished in schizophrenia. *Nature Communications*, 12(1):3478.

[9] Breakspear, M. (2017). Dynamic models of large-scale brain activity. *Nature Neuroscience*, 20(3):340–352.

[10] Bullmore, E. and Sporns, O. (2012). The economy of brain network organization. *Nature Reviews Neuroscience*, 13(5):336–349.

[11] Burgos, N., Cardoso, M. J., Thielemans, K., Modat, M., Pedemonte, S., Dickson, J., Barnes, A., Ahmed, R., Mahoney, C. J., Schott, J. M., Duncan, J. S., Atkinson, D., Arridge, S. R., Hutton, B. F., and Ourselin, S. (2014). Attenuation Correction Synthesis for Hybrid PET-MR Scanners: Application to Brain Studies. *IEEE Transactions on Medical Imaging*, 33(12):2332–2341.

[12] Burt, J. B., Helmer, M., Shinn, M., Anticevic, A., and Murray, J. D. (2020). Generative modeling of brain maps with spatial autocorrelation. *NeuroImage*, 220:117038.

[13] Camandola, S. and Mattson, M. P. (2017). Brain metabolism in health, aging, and neurodegeneration. *The EMBO Journal*, 36(11):1474–1492.

[14] Castrillon, G., Epp, S., Bose, A., Fraticelli, L., Hechler, A., Belenya, R., Ranft, A., Yakushev, I., Utz, L., Sundar, L., Rauschecker, J. P., Preibisch, C., Kurcyus, K., and Riedl, V. (2023). An energy costly architecture of neuromodulators for human brain evolution and cognition. *Science Advances*, 9(50):eadi7632.

[15] Cornblath, E. J., Ashourvan, A., Kim, J. Z., Betzel, R. F., Ciric, R., Adebimpe, A., Baum, G. L., He, X., Ruparel, K., Moore, T. M., Gur, R. C., Gur, R. E., Shinohara, R. T., Roalf, D. R., Satterthwaite, T. D., and Bassett, D. S. (2020). Temporal sequences of brain activity at rest are constrained by white matter structure and modulated by cognitive demands. *Communications Biology*, 3(1):1–12.

[16] Craddock, C., Sikka, S., Cheung, B., Khanuja, R., Ghosh, S. S., Yan, C., Li, Q., Lurie, D., Vogelstein, J., Burns, R., et al. (2013). Towards automated analysis of connectomes: The configurable pipeline for the analysis of connectomes (c-pac). *Front Neuroinform*, 42(10.3389).

[17] Cui, Z., Stiso, J., Baum, G. L., Kim, J. Z., Roalf, D. R., Betzel, R. F., Gu, S., Lu, Z., Xia, C. H., He, X., Ciric, R., Oathes, D. J., Moore, T. M., Shinohara, R. T., Ruparel, K., Davatzikos, C., Pasqualetti, F., Gur, R. E., Gur, R. C., Bassett, D. S., and Satterthwaite, T. D. (2020). Optimization of energy state transition trajectory supports the development of executive function during youth. *eLife*, 9:e53060.

[18] Deco, G., Tononi, G., Boly, M., and Kringelbach, M. L. (2015). Rethinking segregation and integration: Contributions of whole-brain modelling. *Nature Reviews Neuroscience*, 16(7):430–439.

[19] Dhollander, T., Raffelt, D., and Connelly, A. (2016). Unsupervised 3-tissue response function estimation from single-shell or multi-shell diffusion mr data without a co-registered t1 image. In *ISMRM workshop on breaking the barriers of diffusion MRI*, volume 5. Lisbon, Portugal.

[20] Díaz-García, C. M., Mongeon, R., Lahmann, C., Koveal, D., Zucker, H., and Yellen, G. (2017). Neuronal Stimulation Triggers Neuronal Glycolysis and Not Lactate Uptake. *Cell Metabolism*, 26(2):361–374.e4.

[21] Dienel, G. A. (2012). Fueling and Imaging Brain Activation. *ASN Neuro*, 4(5):AN20120021.

[22] Dienel, G. A. (2019). Brain Glucose Metabolism: Integration of Energetics with Function. *Physiological Reviews*, 99(1):949–1045.

[23] Eickhoff, S. B., Thirion, B., Varoquaux, G., and Bzdok, D. (2015). Connectivity-based parcellation: Critique and implications. *Human Brain Mapping*, 36(12):4771–4792.

[24] Epp, S. M., Castrillón, G., Yuan, B., Andrews-Hanna, J., Preibisch, C., and Riedl, V. (2023). Two distinct modes of hemodynamic responses in the human brain.

[25] Finc, K., Bonna, K., He, X., Lydon-Staley, D. M., Kühn, S., Duch, W., and Bassett, D. S. (2020). Dynamic reconfiguration of functional brain networks during working memory training. *Nature Communications*, 11(1):2435.

[26] Fox, P. T., Raichle, M. E., Mintun, M. A., and Dence, C. (1988). Nonoxidative Glucose Consumption During Focal Physiologic Neural Activity. *Science*, 241(4864):462–464.

[27] Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138.

[28] Gautron, L., Elmquist, J. K., and Williams, K. W. (2015). Neural Control of Energy Balance: Translating Circuits to Therapies. *Cell*, 161(1):133–145.

[29] Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J. R., Van Essen, D. C., and Jenkinson, M. (2013). The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, 80:105–124.

[30] Gordon, E. M., Laumann, T. O., Adeyemo, B., Huckins, J. F., Kelley, W. M., and Petersen, S. E. (2016). Generation and Evaluation of a Cortical Area Parcellation from Resting-State Correlations. *Cerebral Cortex*, 26(1):288–303.

[31] Gorgolewski, K. J., Alfaro-Almagro, F., Auer, T., Bellec, P., Capotă, M., Chakravarty, M. M., Churchill, N. W., Cohen, A. L., Craddock, R. C., Devenyi, G. A., Eklund, A., Esteban, O., Flandin, G., Ghosh, S. S., Guntupalli, J. S., Jenkinson, M., Keshavan, A., Kiar, G., Liem, F., Raamana, P. R., Raffelt, D., Steele, C. J., Quirion, P.-O., Smith, R. E., Strother, S. C., Varoquaux, G., Wang, Y., Yarkoni, T., and Poldrack, R. A. (2017). BIDS apps: Improving ease of use, accessibility, and reproducibility of neuroimaging data analysis methods. *PLOS Computational Biology*, 13(3):e1005209.

[32] Gu, S., Betzel, R. F., Mattar, M. G., Cieslak, M., Delio, P. R., Grafton, S. T., Pasqualetti, F., and Bassett, D. S. (2017). Optimal trajectories of brain state transitions. *NeuroImage*, 148:305–317.

[33] Gu, S., Pasqualetti, F., Cieslak, M., Telesford, Q. K., Yu, A. B., Kahn, A. E., Medaglia, J. D., Vettel, J. M., Miller, M. B., Grafton, S. T., and Bassett, D. S. (2015). Controllability of structural brain networks. *Nature Communications*, 6(1):8414.

[34] Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C. J., Wedeen, V. J., and Sporns, O. (2008). Mapping the Structural Core of Human Cerebral Cortex. *PLOS Biology*, 6(7):e159.

[35] Harris, J. J., Jolivet, R., and Attwell, D. (2012). Synaptic Energy Use and Supply. *Neuron*, 75(5):762–777.

[36] He, X., Caciagli, L., Parkes, L., Stiso, J., Karrer, T. M., Kim, J. Z., Lu, Z., Menara, T., Pasqualetti, F., Sperling, M. R., Tracy, J. I., and Bassett, D. S. (2022). Uncovering the biological basis of control energy: Structural and metabolic correlates of energy inefficiency in temporal lobe epilepsy. *Science Advances*, 8(45):eabn2293.

[37] Jain, A. K. (2010). Data clustering: 50 years beyond K-means. *Pattern Recognition Letters*, 31(8):651–666.

[38] Jeurissen, B., Tournier, J.-D., Dhollander, T., Connelly, A., and Sijbers, J. (2014). Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data. *NeuroImage*, 103:411–426.

[39] Karrer, T. M., Kim, J. Z., Stiso, J., Kahn, A. E., Pasqualetti, F., Habel, U., and Bassett, D. S. (2020). A practical guide to methodological considerations in the controllability of structural brain networks. *Journal of Neural Engineering*, 17(2):026031.

[40] Kellner, E., Dhital, B., Kiselev, V. G., and Reisert, M. (2016). Gibbs-ringing artifact removal based on local subvoxel-shifts. *Magnetic Resonance in Medicine*, 76(5):1574–1581.

[41] Kleinberg, J. (2002). An Impossibility Theorem for Clustering. In *Advances in Neural Information Processing Systems*, volume 15. MIT Press.

[42] Leonardi, N. and Van De Ville, D. (2015). On spurious and real fluctuations of dynamic functional connectivity during rest. *NeuroImage*, 104:430–436.

[43] Levin, B. E. (2006). Metabolic sensing neurons and the control of energy homeostasis. *Physiology & Behavior*, 89(4):486–489.

[44] Li, S. and Sheng, Z.-H. (2022). Energy matters: Presynaptic metabolism and the maintenance of synaptic transmission. *Nature Reviews Neuroscience*, 23(1):4–22.

[45] Lord, L.-D., Stevner, A. B., Deco, G., and Kringelbach, M. L. (2017). Understanding principles of integration and segregation using whole-brain computational connectomics: Implications for neuropsychiatric disorders. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 375(2096):20160283.

[46] Luppi, A. I., Craig, M. M., Pappas, I., Finoia, P., Williams, G. B., Allanson, J., Pickard, J. D., Owen, A. M., Naci, L., Menon, D. K., and Stamatakis, E. A. (2019). Consciousness-specific dynamic interactions of brain integration and functional diversity. *Nature Communications*, 10(1):4616.

[47] Luppi, A. I., Rosas, F. E., Noonan, M. P., Mediano, P. A. M., Kringelbach, M. L., Carhart-Harris, R. L., Stamatakis, E. A., Vernon, A. C., and Turkheimer, F. E. (2022). Oxygen and the Spark of Human Brain Evolution: Complex Interactions of Metabolism and Cortical Expansion across Development and Evolution. *The Neuroscientist*, page 10738584221138032.

[48] Luppi, A. I., Singleton, S. P., Hansen, J. Y., Bzdok, D., Kuceyeski, A., Betzel, R. F., and Misic, B. (2023). Transitions between cognitive topographies: Contributions of network structure, neuromodulation, and disease.

[49] Margulies, D. S., Ghosh, S. S., Goulas, A., Falkiewicz, M., Huntenburg, J. M., Langs, G., Bezgin, G., Eickhoff, S. B., Castellanos, F. X., Petrides, M., Jefferies, E., and Smallwood, J. (2016). Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proceedings of the National Academy of Sciences*, 113(44):12574–12579.

[50] Markello, R. D., Hansen, J. Y., Liu, Z.-Q., Bazinet, V., Shafiei, G., Suárez, L. E., Blostein, N., Seidlitz, J., Baillet, S., Satterthwaite, T. D., Chakravarty, M. M., Raznahan, A., and Misic, B. (2022). Neuromaps: Structural and functional interpretation of brain maps. *Nature Methods*, 19(11):1472–1479.

[51] Markello, R. D. and Misic, B. (2021). Comparing spatial null models for brain maps. *NeuroImage*, 236:118052.

[52] Markiewicz, P. J., Ehrhardt, M. J., Erlandsson, K., Noonan, P. J., Barnes, A., Schott, J. M., Atkinson, D., Arridge, S. R., Hutton, B. F., and Ourselin, S. (2018). NiftyPET: A High-throughput Software Platform for High Quantitative Accuracy and Precision PET Imaging and Analysis. *Neuroinformatics*, 16(1):95–115.

[53] Mergenthaler, P., Lindauer, U., Dienel, G. A., and Meisel, A. (2013). Sugar for the brain: The role of glucose in physiological and pathological brain function. *Trends in*

Neurosciences, 36(10):587–597.

[54] Mesulam, M. M. (1998). From sensation to cognition. *Brain*, 121(6):1013–1052.

[55] Miranda-Dominguez, O., Feczko, E., Grayson, D. S., Walum, H., Nigg, J. T., and Fair, D. A. (2018). Heritability of the human connectome: A connectotyping study. *Network Neuroscience*, 02(02):175–199.

[56] Mišić, B., Betzel, R. F., Nematzadeh, A., Goñi, J., Griffa, A., Hagmann, P., Flammini, A., Ahn, Y.-Y., and Sporns, O. (2015). Cooperative and Competitive Spreading Dynamics on the Human Connectome. *Neuron*, 86(6):1518–1529.

[57] Nozari, E., Bertolero, M. A., Stiso, J., Caciagli, L., Cornblath, E. J., He, X., Mahadevan, A. S., Pappas, G. J., and Bassett, D. S. (2023). Macroscopic resting-state brain dynamics are best described by linear models. *Nature Biomedical Engineering*, pages 1–17.

[58] Padamsey, Z. and Rochefort, N. L. (2023). Paying the brain's energy bill. *Current Opinion in Neurobiology*, 78:102668.

[59] Park, B.-y., Vos de Wael, R., Paquola, C., Larivière, S., Benkarim, O., Royer, J., Tavakol, S., Cruces, R. R., Li, Q., Valk, S. L., Margulies, D. S., Mišić, B., Bzdok, D., Smallwood, J., and Bernhardt, B. C. (2021). Signal diffusion along connectome gradients and inter-hub routing differentially contribute to dynamic human brain function. *NeuroImage*, 224:117429.

[60] Parkes, L., Kim, J. Z., Stiso, J., Brynildsen, J. K., Cieslak, M., Covitz, S., Gur, R. E., Gur, R. C., Pasqualetti, F., Shinohara, R. T., Zhou, D., Satterthwaite, T. D., and Bassett, D. S. (2023). Using network control theory to study the dynamics of the structural connectome.

[61] Parkes, L., Kim, J. Z., Stiso, J., Calkins, M. E., Cieslak, M., Gur, R. E., Gur, R. C., Moore, T. M., Ouellet, M., Roalf, D. R., Shinohara, R. T., Wolf, D. H., Satterthwaite, T. D., and Bassett, D. S. (2022). Asymmetric signaling across the hierarchy of cytoarchitecture within the human connectome. *Science Advances*, 8(50):eadd2185.

[62] Parkes, L., Moore, T. M., Calkins, M. E., Cieslak, M., Roalf, D. R., Wolf, D. H., Gur, R. C., Gur, R. E., Satterthwaite, T. D., and Bassett, D. S. (2021). Network Controllability in Transmodal Cortex Predicts Positive Psychosis Spectrum Symptoms. *Biological Psychiatry*, 90(6):409–418.

[63] Patenaude, B., Smith, S. M., Kennedy, D. N., and Jenkinson, M. (2011). A Bayesian model of shape and appearance for subcortical brain segmentation. *NeuroImage*, 56(3):907–922.

[64] Patlak, C. S. and Blasberg, R. G. (1985). Graphical Evaluation of Blood-to-Brain Transfer Constants from Multiple-Time Uptake Data. Generalizations. *Journal of Cerebral Blood Flow & Metabolism*, 5(4):584–590.

[65] Raichle, M. E. and Gusnard, D. A. (2002). Appraising the brain's energy budget. *Proceedings of the National Academy of Sciences*, 99(16):10237–10239.

[66] Robinson, E. C., Jbabdi, S., Glasser, M. F., Andersson, J., Burgess, G. C., Harms, M. P., Smith, S. M., Van Essen, D. C., and Jenkinson, M. (2014). MSM: A new flexible framework for Multimodal Surface Matching. *NeuroImage*, 100:414–426.

[67] Saggar, M., Shine, J. M., Liégeois, R., Dosenbach, N. U. F., and Fair, D. (2022). Precision dynamical mapping using topological data analysis reveals a hub-like transition state at rest. *Nature Communications*, 13(1):4791.

[68] Salimi-Khorshidi, G., Douaud, G., Beckmann, C. F., Glasser, M. F., Griffanti, L., and Smith, S. M. (2014). Automatic denoising of functional MRI data: Combining independent component analysis and hierarchical fusion of classifiers. *NeuroImage*, 90:449–468.

[69] Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X.-N., Holmes, A. J., Eickhoff, S. B., and Yeo, B. T. T. (2018). Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. *Cerebral Cortex*, 28(9):3095–3114.

[70] Shine, J. M. and Poldrack, R. A. (2018). Principles of dynamic network reconfiguration across diverse brain states. *NeuroImage*, 180:396–405.

[71] Singleton, S. P., Luppi, A. I., Carhart-Harris, R. L., Cruzat, J., Roseman, L., Nutt, D. J., Deco, G., Kringelbach, M. L., Stamatakis, E. A., and Kuceyeski, A. (2022). Receptor-informed network control theory links LSD and psilocybin to a flattening of the brain's control energy landscape. *Nature Communications*, 13(1):5812.

[72] Singleton, S. P., Timmermann, C., Luppi, A. I., Eckernäs, E., Roseman, L., Carhart-Harris, R. L., and Kuceyeski, A. (2023). Time-resolved network control analysis links reduced control energy under DMT with the serotonin 2a receptor, signal diversity, and subjective experience.

[73] Smith, R. E., Tournier, J.-D., Calamante, F., and Connelly, A. (2012). Anatomically-constrained tractography: Improved diffusion MRI streamlines tractography through effective use of anatomical information. *NeuroImage*, 62(3):1924–1938.

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

[75] Smith, S. M. (2002). Fast robust automated brain extraction. *Human Brain Mapping*, 17(3):143–155.

[76] Smith, S. M., Beckmann, C. F., Andersson, J., Auerbach, E. J., Bijsterbosch, J., Douaud, G., Duff, E., Feinberg, D. A., Griffanti, L., Harms, M. P., Kelly, M., Laumann, T., Miller, K. L., Moeller, S., Petersen, S., Power, J., Salimi-Khorshidi, G., Snyder, A. Z., Vu, A. T., Woolrich, M. W., Xu, J., Yacoub, E., Uğurbil, K., Van Essen, D. C., and Glasser, M. F. (2013). Resting-state fMRI in the Human Connectome Project. *NeuroImage*, 80:144–168.

[77] Sporns, O. (2011). The human connectome: A complex network. *Annals of the New York Academy of Sciences*, 1224(1):109–125.

[78] Srivastava, P., Nozari, E., Kim, J. Z., Ju, H., Zhou, D., Becker, C., Pasqualetti, F., Pappas, G. J., and Bassett, D. S. (2020). Models of communication and control for brain networks: Distinctions, convergence, and future outlook. *Network Neuroscience*, 4(4):1122–1159.

[79] Stiso, J., Khambhati, A. N., Menara, T., Kahn, A. E., Stein, J. M., Das, S. R., Gorniak, R., Tracy, J., Litt, B., Davis, K. A., Pasqualetti, F., Lucas, T. H., and Bassett, D. S. (2019). White Matter Network Architecture Guides Direct Electrical Stimulation through Optimal State Transitions. *Cell Reports*, 28(10):2554–2566.e7.

[80] Tang, E., Giusti, C., Baum, G. L., Gu, S., Pollock, E., Kahn, A. E., Roalf, D. R., Moore, T. M., Ruparel, K., Gur, R. C., Gur, R. E., Satterthwaite, T. D., and Bassett, D. S. (2017). Developmental increases in white matter network controllability support a growing diversity of brain dynamics. *Nature Communications*, 8(1):1252.

[81] Thiebaut de Schotten, M. and Forkel, S. J. (2022). The emergent properties of the connected brain. *Science*, 378(6619):505–510.

[82] Tournier, J.-D., Calamante, F., Connelly, A., et al. (2011). Effect of step size on probabilistic streamlines: implications for the interpretation of connectivity analysis. In *Proc Intl Soc Mag Reson Med*, volume 19, page 2019.

[83] Tournier, J. D., Calamante, F., Gadian, D. G., and Connelly, A. (2004). Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. *NeuroImage*, 23(3):1176–1185.

[84] Tournier, J.-D., Smith, R., Raffelt, D., Tabbara, R., Dhollander, T., Pietsch, M., Christiaens, D., Jeurissen, B., Yeh, C.-H., and Connelly, A. (2019). MRtrix3: A fast, flexible and open software framework for medical image processing and visualisation. *NeuroImage*, 202:116137.

[85] Tustison, N. J., Avants, B. B., Cook, P. A., Zheng, Y., Egan, A., Yushkevich, P. A., and Gee, J. C. (2010). N4ITK: Improved N3 Bias Correction. *IEEE Transactions on Medical Imaging*, 29(6):1310–1320.

[86] Vaishnavi, S. N., Vlassenko, A. G., Rundle, M. M., Snyder, A. Z., Mintun, M. A., and Raichle, M. E. (2010). Regional aerobic glycolysis in the human brain. *Proceedings of the National Academy of Sciences*, 107(41):17757–17762.

[87] Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., and Ugurbil, K. (2013). The WU-Minn Human Connectome Project: An overview. *NeuroImage*, 80:62–79.

[88] Váša, F. and Mišić, B. (2022). Null models in network neuroscience. *Nature Reviews Neuroscience*, 23(8):493–504.

[89] Veraart, J., Novikov, D. S., Christiaens, D., Ades-aron, B., Sijbers, J., and Fieremans, E. (2016). Denoising of diffusion MRI using random matrix theory. *NeuroImage*, 142:394–406.

[90] Vidaurre, D., Smith, S. M., and Woolrich, M. W. (2017). Brain network dynamics are hierarchically organized in time. *Proceedings of the National Academy of Sciences*, 114(48):12827–12832.

[91] Watts, M. E., Pocock, R., and Claudianos, C. (2018). Brain Energy and Oxygen Metabolism: Emerging Role in Normal Function and Disease. *Frontiers in Molecular Neuroscience*, 11.

[92] Welvaert, M. and Rosseel, Y. (2013). On the Definition of Signal-To-Noise Ratio and Contrast-To-Noise Ratio for fMRI Data. *PLOS ONE*, 8(11):e77089.

[93] Wu, H.-M., Bergsneider, M., Glenn, T. C., Yeh, E., Hovda, D. A., Phelps, M. E., and Huang, S.-C. (2003). Measurement of the Global Lumped Constant for 2-Deoxy-2-[18F]Fluoro-D-Glucose in Normal Human Brain Using [15O]Water and 2-Deoxy-2-[18F]Fluoro-D-Glucose Positron Emission Tomography Imaging: A Method with Validation Based on Multiple Methodologies. *Molecular Imaging & Biology*, 5(1):32–41.

[94] Yang, J., Hu, C., Guo, N., Dutta, J., Vaina, L. M., Johnson, K. A., Sepulcre, J., Fakhri, G. E., and Li, Q. (2017). Partial volume correction for PET quantification and its impact on brain network in Alzheimer's disease. *Scientific Reports*, 7(1):13035.

[95] Yeo, B. T. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., Roffman, J. L., Smoller, J. W., Zöllei, L., Polimeni, J. R., Fischl, B., Liu, H., and Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, 106(3):1125–1165.

[96] Zala, D., Hinckelmann, M.-V., Yu, H., Lyra da Cunha, M. M., Liot, G., Cordelières, F. P., Marco, S., and Saudou, F. (2013). Vesicular Glycolysis Provides On-Board Energy for Fast Axonal Transport. *Cell*, 152(3):479–491.

[97] Zalesky, A. and Breakspear, M. (2015). Towards a statistical test for functional connectivity dynamics. *NeuroImage*, 114:466–470.

[98] Zalesky, A., Fornito, A., Cocchi, L., Gollo, L. L., and Breakspear, M. (2014). Time-resolved resting-state brain networks. *Proceedings of the National Academy of Sciences*, 111(28):10341–10346.

[99] Zamani Esfahlani, F., Jo, Y., Faskowitz, J., Byrge, L., Kennedy, D. P., Sporns, O., and Betzel, R. F. (2020). High-amplitude cofluctuations in cortical activity drive functional connectivity. *Proceedings of the National Academy of Sciences*, 117(45):28393–28401.

[100] Zhang, Y., Brady, M., and Smith, S. (2001). Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging*, 20(1):45–57.

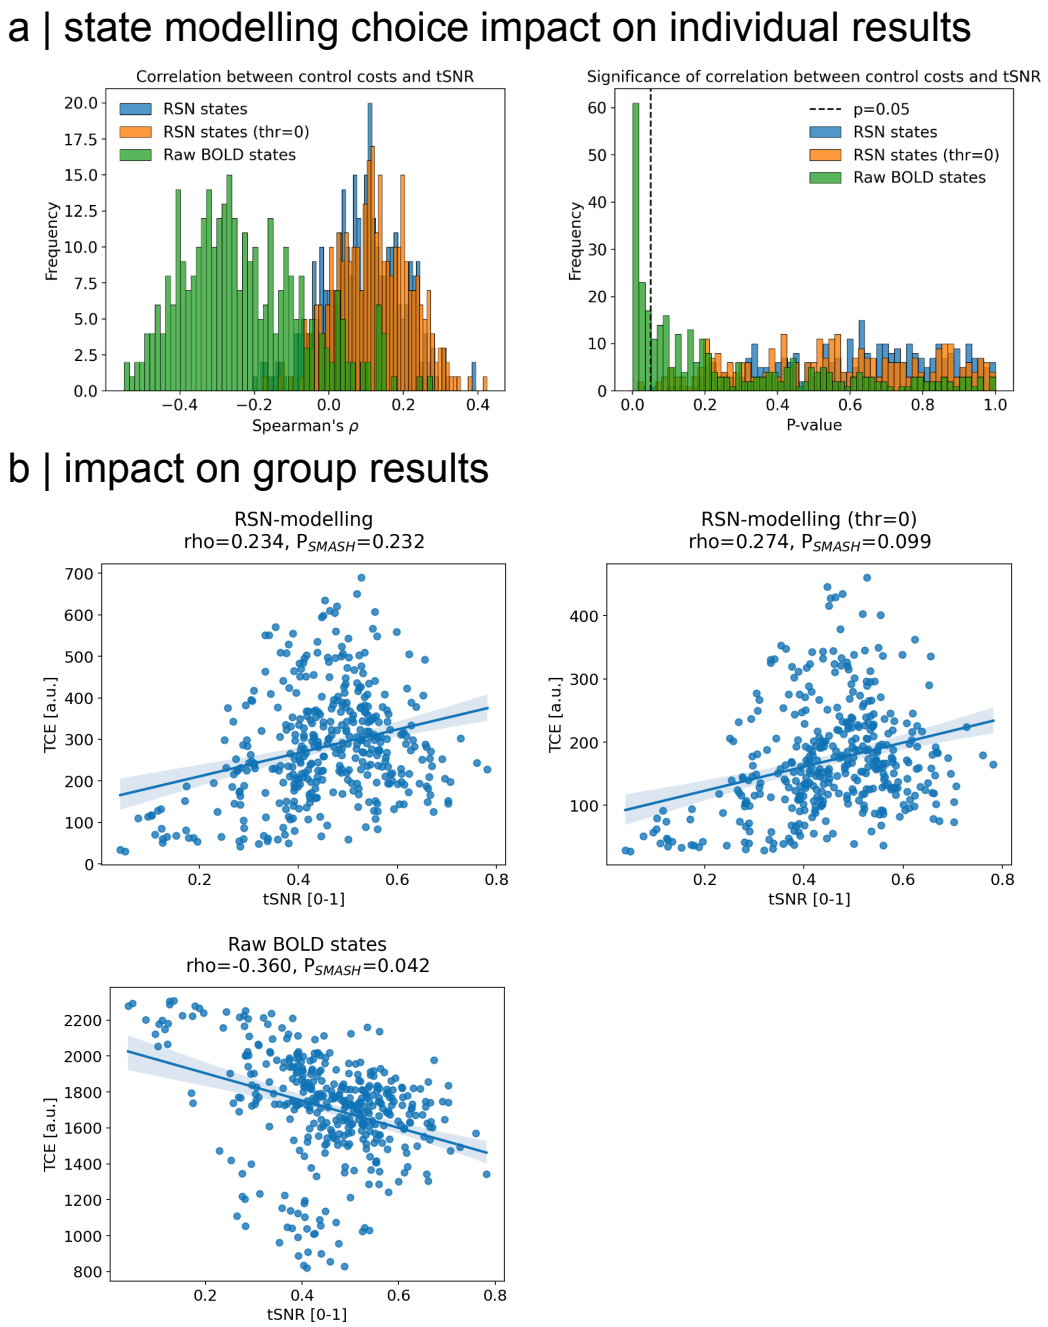

Figure S1. **Intrinsic network modelling validation |** (a) The temporal signal-to-noise ratio (tSNR) of each individual's recording was estimated by dividing their mean signal prior to demeaning with their signal variability [92]. This metric was compared to their control costs from three modelling approaches: 1) our reported approach to use intrinsic networks as states, where only time points with an average network activity above 0.5 are selected; 2) same as 1) but with an activity threshold of 0; 3) raw BOLD activity are used as states. The first two approaches result in TCE maps that correlate similarly to their respective tSNR (left), with only few subjects showing significance in their associations(right). Using the third approach, however, results in control costs being negatively correlated to tSNR (left), and more individual's control maps being significantly related to the tSNR of their recording (right). (b) We see similar results on a group-level, where we average all control maps and tSNR across subjects. In both state modelling approaches, we observe that control costs and tSNR are not significantly associated, however, the lower threshold variant results are statistically more unlikely to be discernible [Spearman ρ = 0.234, P_{SMASH} = 0.232 for thr=0.5 and Spearman ρ = 0.274, P_{SMASH} = 0.099 for thr=0]. Critically, the group-average map of TCE significantly correlates with tSNR when opting for no state modelling [Spearman ρ = -0.36, P_{SMASH} = 0.042].

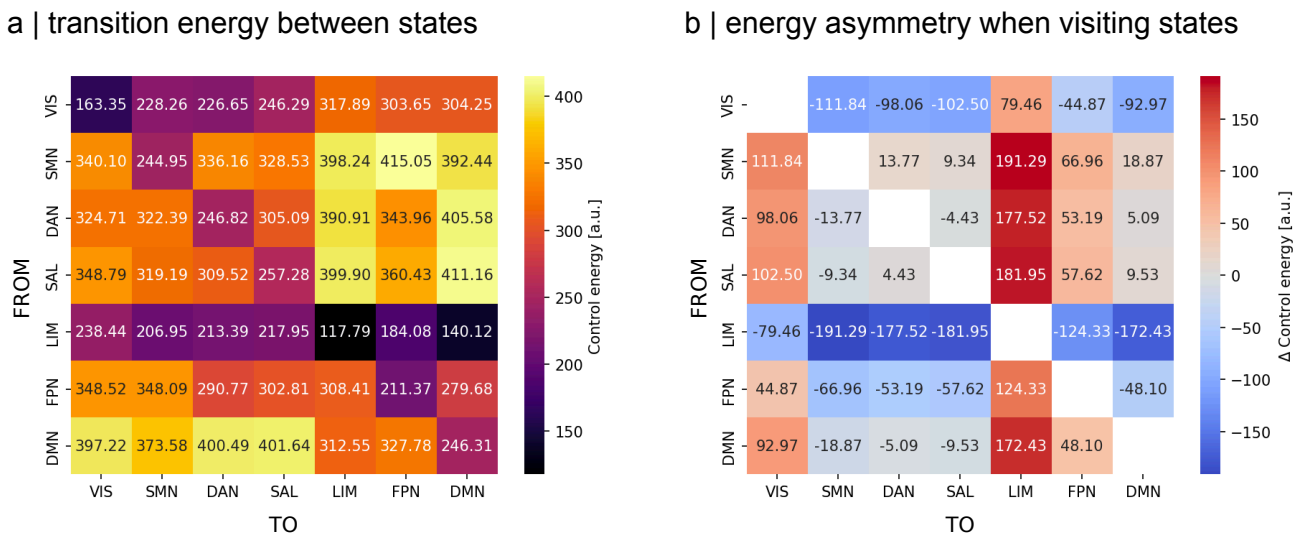

Figure S2. **Transition energy matrix |** (a) We simulate the control energy to transition between all of pairs of intrinsic networks. Similar to previous work [48, 61], we find that the costs to transition into a state are dissimilar to the costs of leaving it. (b) We highlight this asymmetry by subtracting the transition energy matrix with its transpose, thereby accentuating the difference between going into and leaving a state.

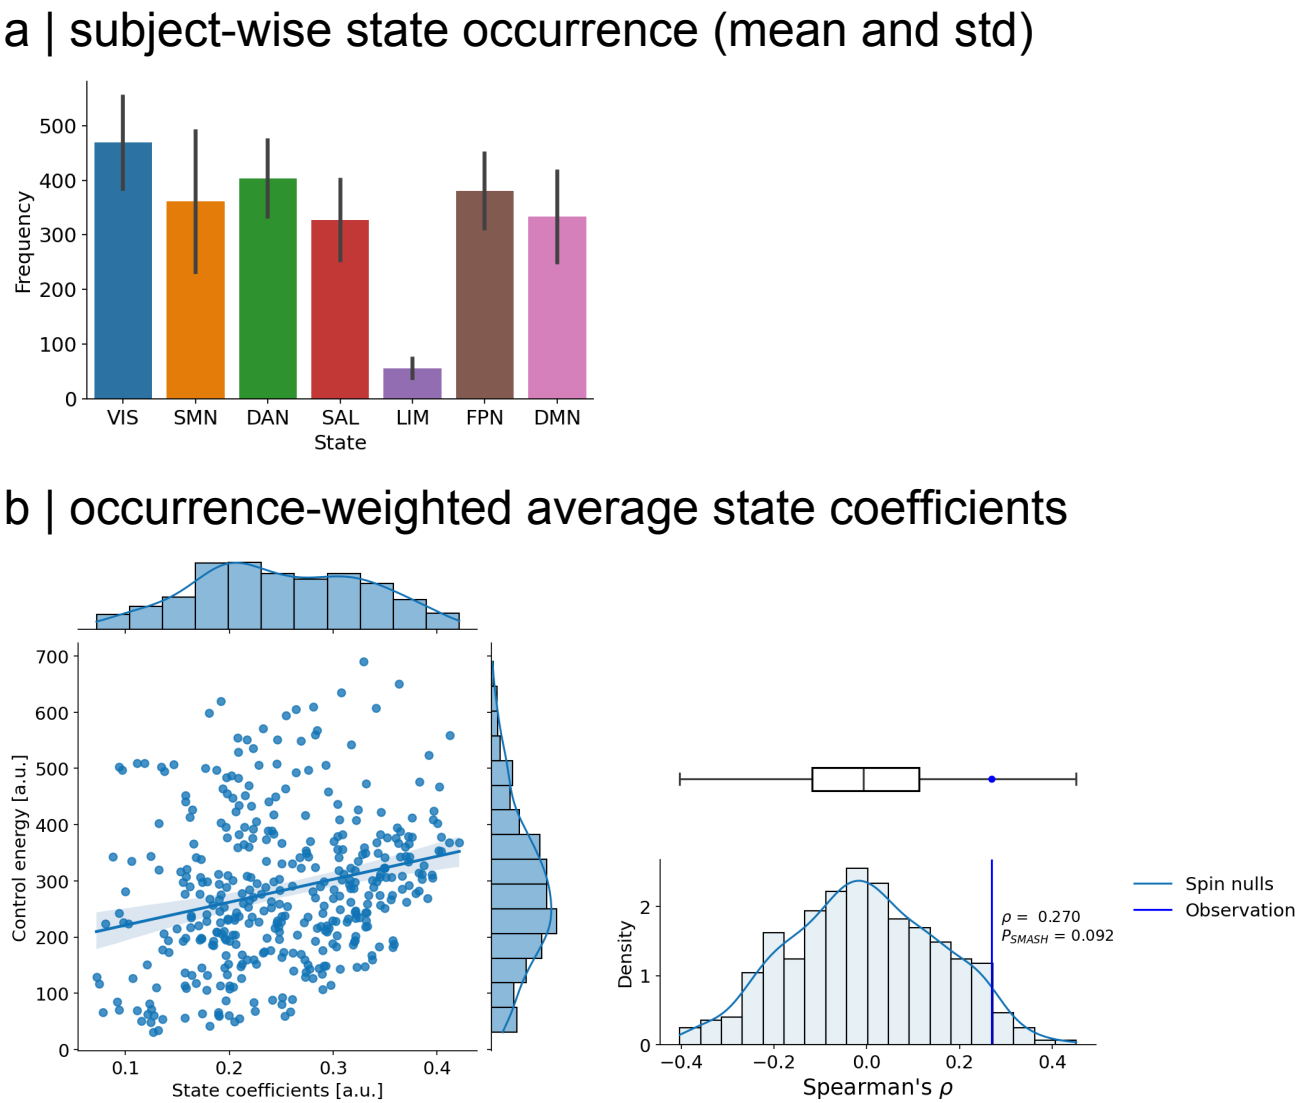

Figure S3. **Control costs vs. state coefficients |** (a) State occurrences are counted for each individual. We report the average and standard deviation of each state distribution. (b) Left: Weighting each network map by their occurrence and then summing all maps yields an occurrence-weighted state coefficient map. This is averaged across all individuals and compared to the group-level TCE map. Right: TCE and state coefficients are not significantly correlated when accounting for spatial autocorrelation [Spearman $\rho = 0.27$, $P_{SMASH} = 0.092$].

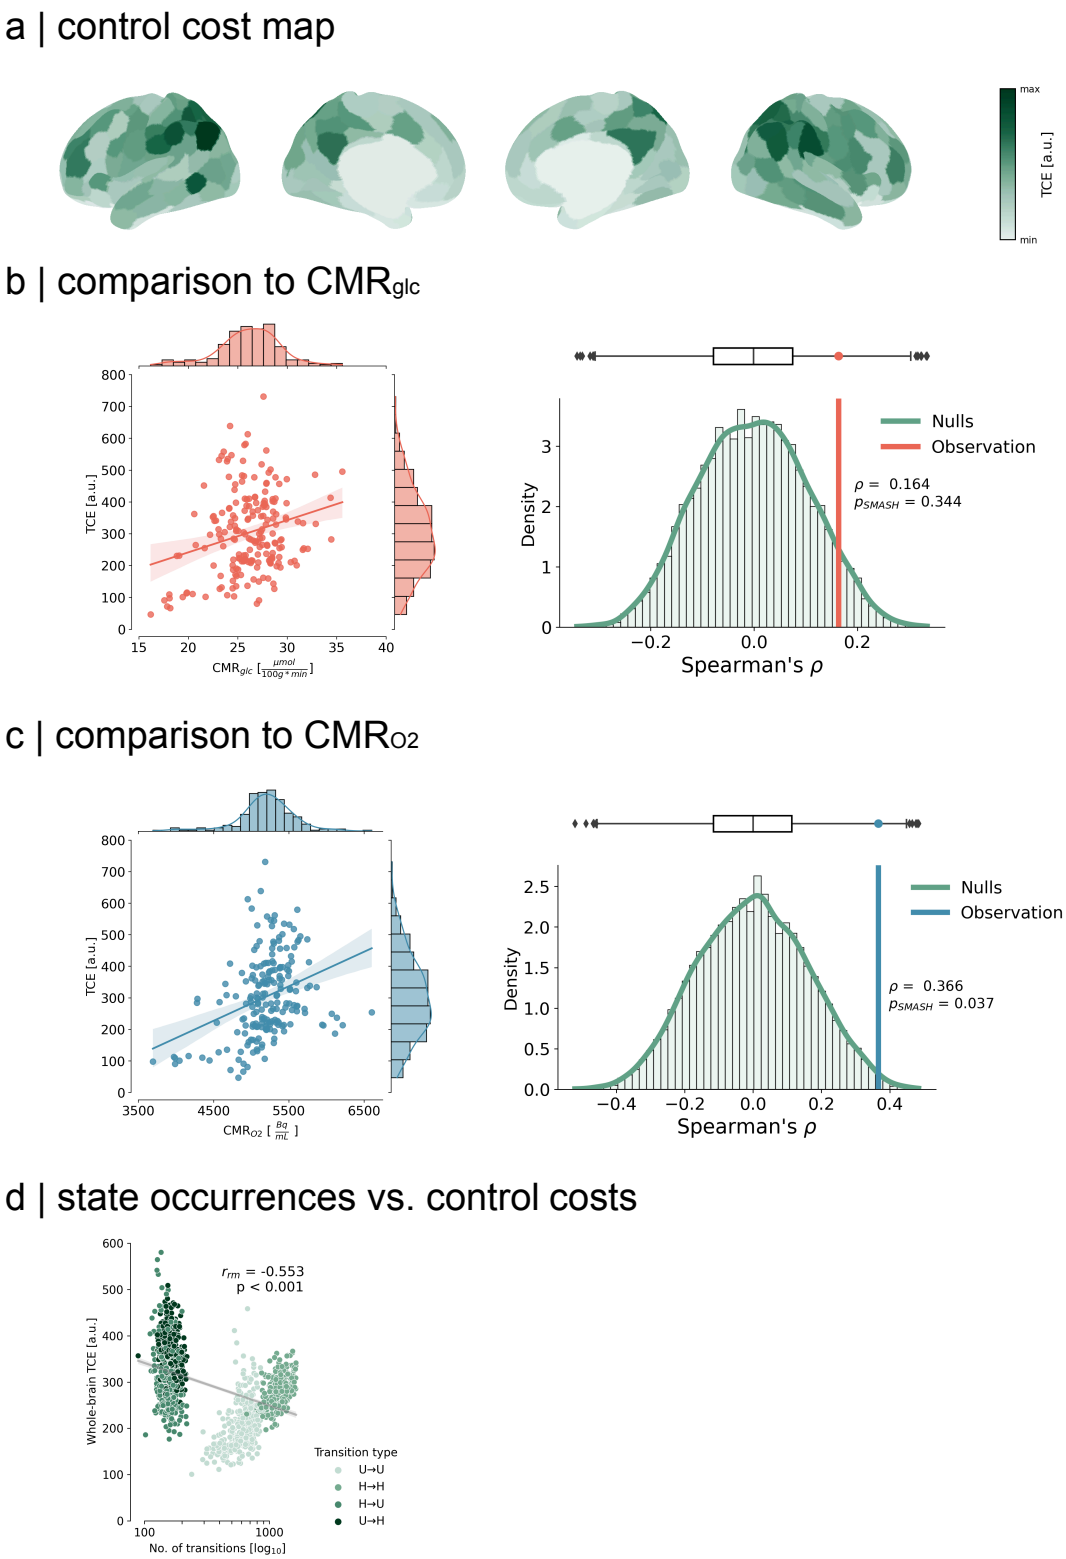

Figure S4. **Replication with Schaefer 200 parcellation|** (a) TCE map across both hemispheres. (b) No significant relationship to CMR_{glc} [Spearman ρ = 0.164, P_{SMASH} = 0.344]. (c) Significant relationship to CMR_{O2} [Spearman ρ = 0.366, P_{SMASH} = 0.037]. (d) Whole-brain TCE is inversely related to the number of transitions across hierarchies [r_{rm} = -0.553, P < 0.001].

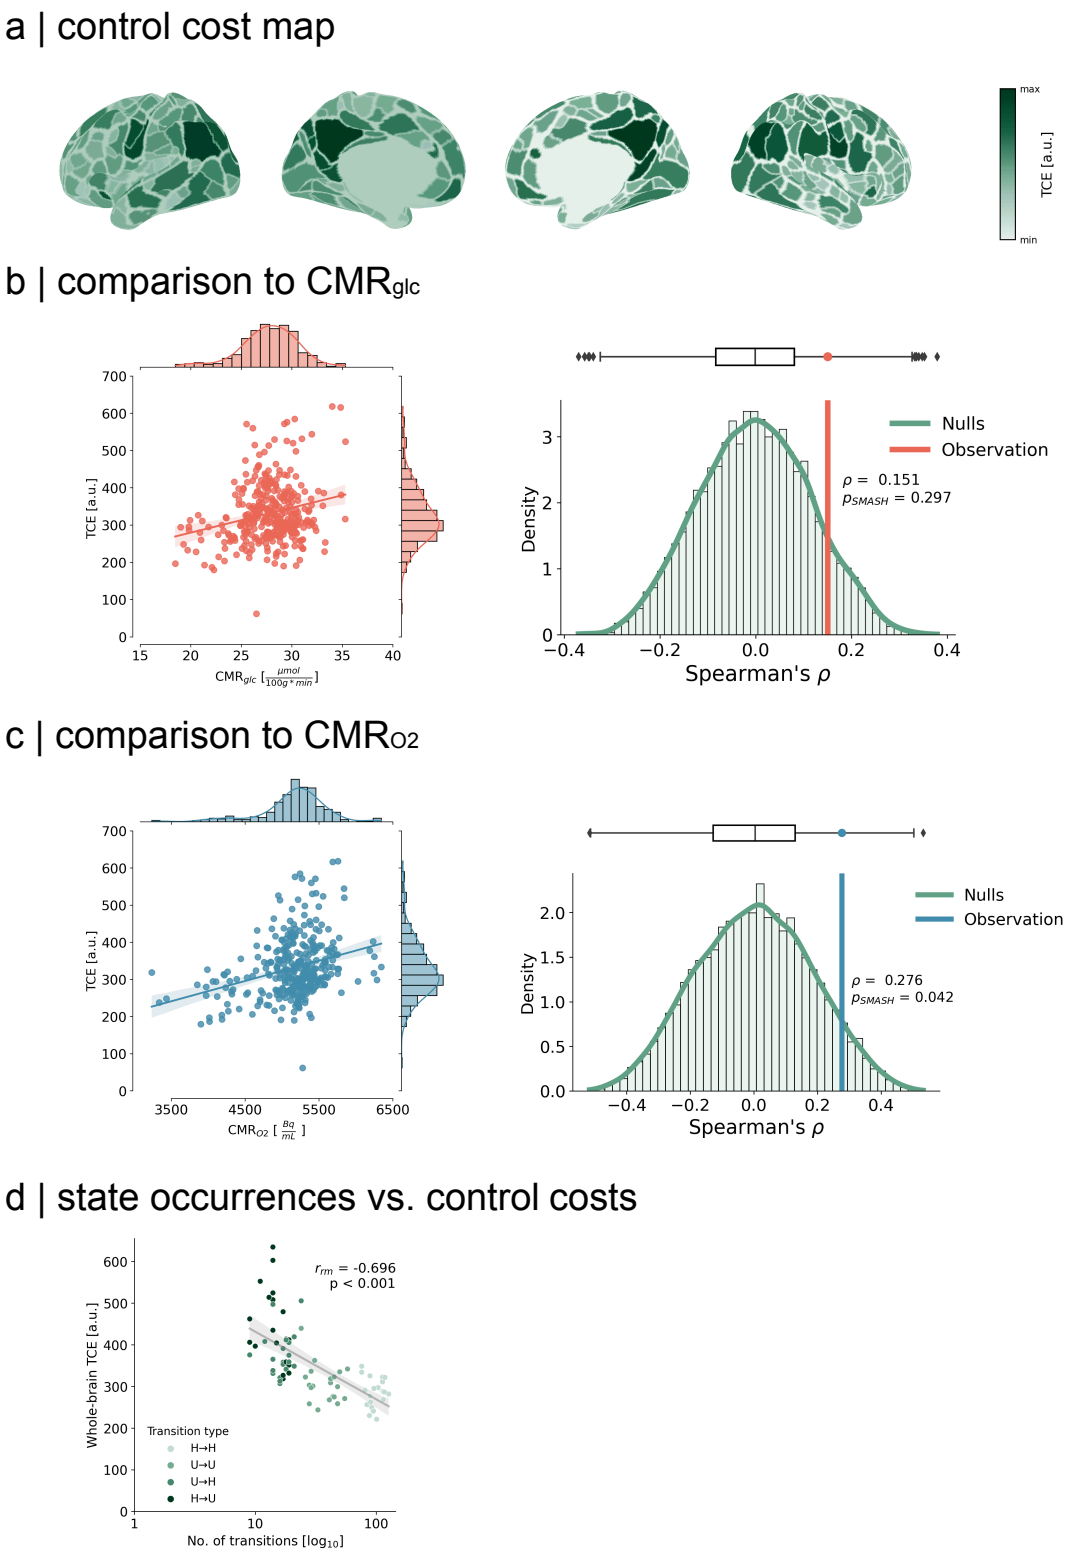

Figure S5. **Replication with Gordon parcellation** | (a) TCE map across both hemispheres. (b) No significant relationship to CMR_{glc} [Spearman ρ = 0.151, P_{SMASH} = 0.297]. (c) Significant relationship to CMR_{O2} [Spearman ρ = 0.276, P_{SMASH} = 0.042]. (d) Whole-brain TCE is inversely related to the number of transitions across hierarchies [r_{rm} = -0.696, P < 0.001].

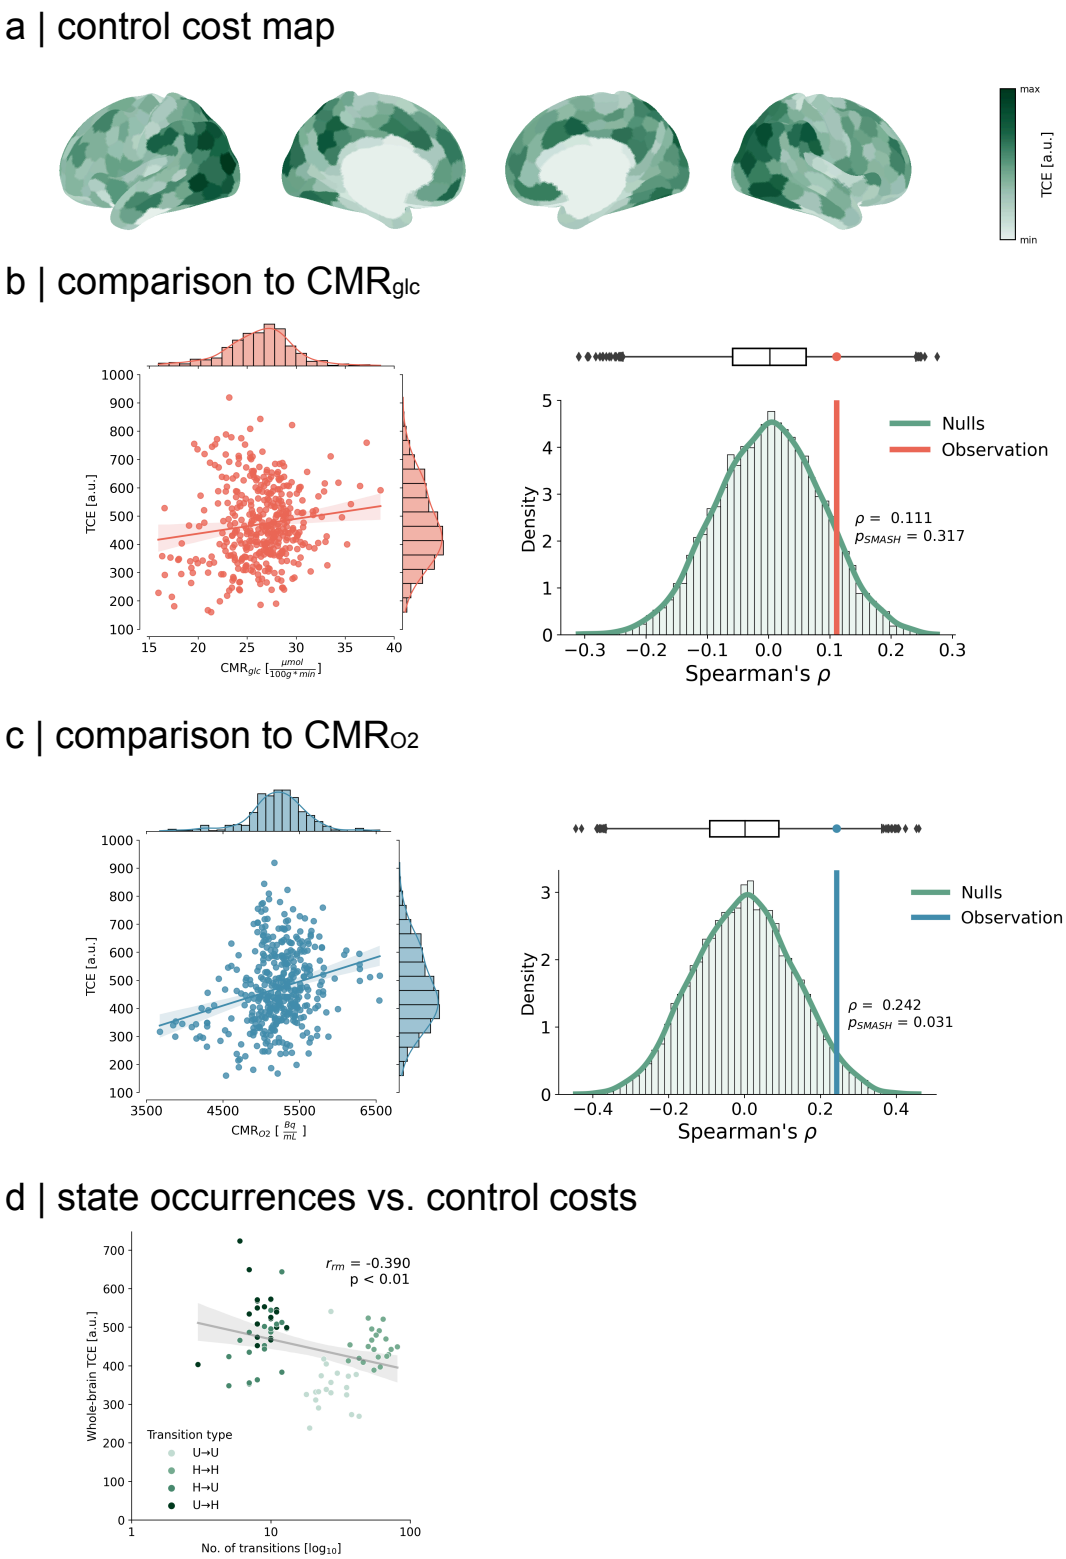

Figure S6. **Replication with separate dataset [14]** | (a) TCE map across both hemispheres. (b) No significant relationship to CMR_{glc} [Spearman ρ = 0.111, P_{SMASH} = 0.317]. (c) Significant relationship to CMR_{O2} [Spearman ρ = 0.242, P_{SMASH} = 0.031]. (d) Whole-brain TCE is inversely related to the number of transitions across hierarchies [r_{rm} = -0.39, P < 0.01].