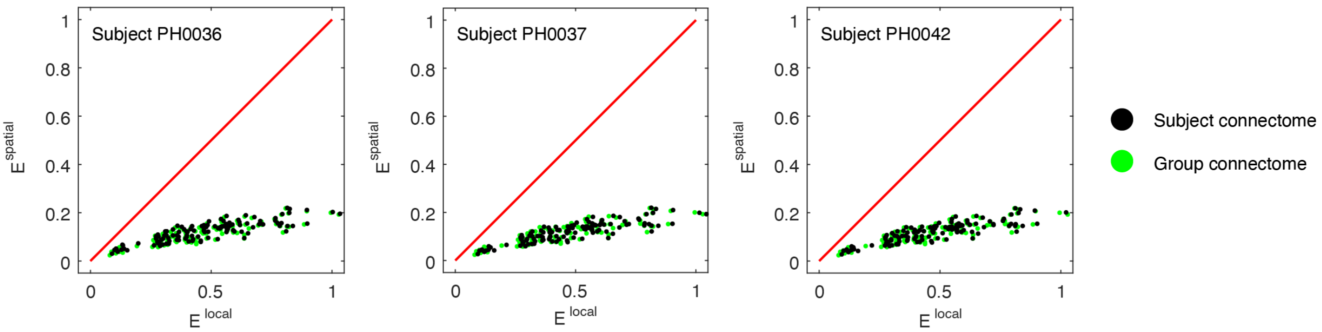

Figure S1. **Replicating main findings using single-subject connectomes.** In the main text, we used a group-representative connectome to estimate control energies. Here, we show that we obtain similar results using connectomes estimated at the subject level. Specifically, we calculate the global control energy for all 121 transitions (11 states to 11 states). In every instance, the spatially informed input strategy resulted in reduced control energy compared to the standard input model.

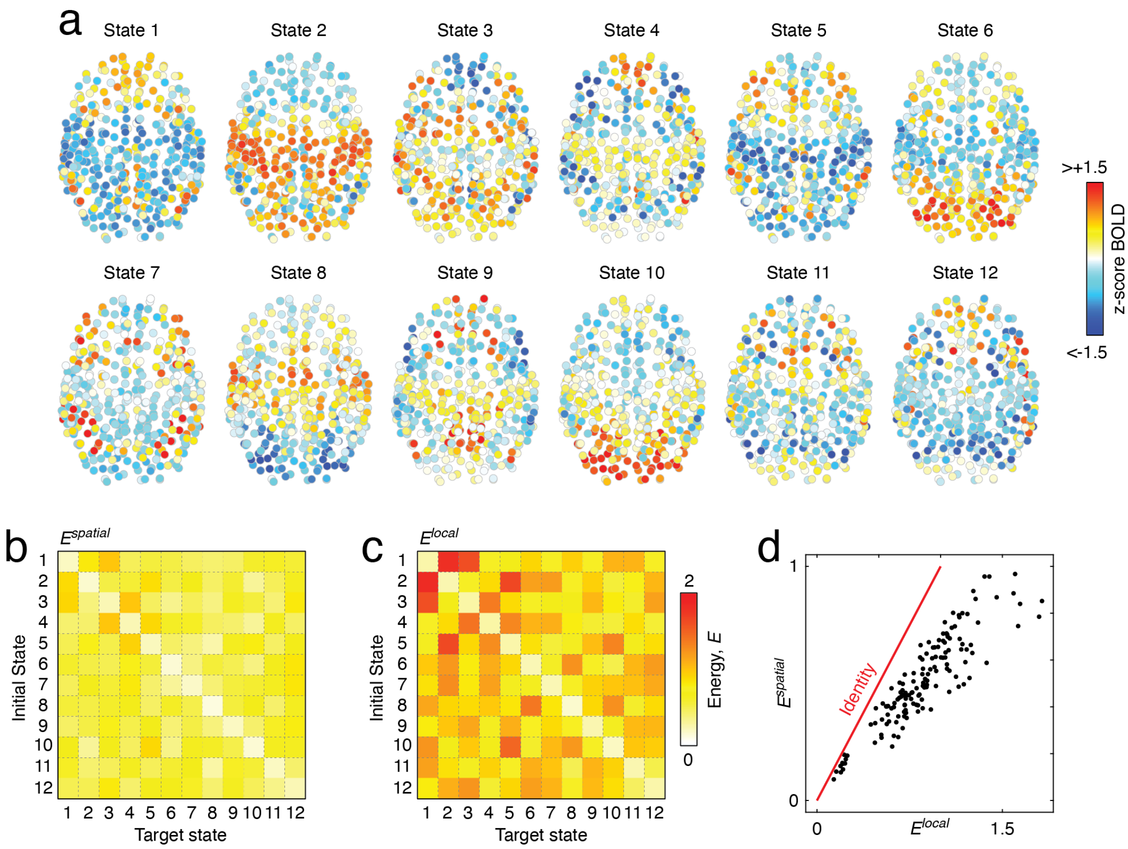

Figure S2. **Independent replication of main finding.** In the main text we showed that spatially diffuse inputs reduce the control energy needed to transition between a set of 11 empirically defined brain states compared to the traditional "local" input strategy. Here, we replicate this finding using a larger and independently acquired and processed dataset. Specifically, we used resting-state and diffusion MRI data from the Human Connectome Project. In general, we followed an analysis pipeline identical to the one described in the main text. We used peak sampling to obtain 6241 peaks in the resting-state activity (`REST1_LR`) from 95 unrelated participants with complete datasets that also met minimum data quality standards. (*a*) After clustering the peak activity patterns, we identified 12 brain states that appeared in at least 50% of the participants. We show the topography of those states here. Panels *b* and *c* show control energies for all pairs of brain state transitions. (*d*) Scatter plot of energies. For all 144 transitions, the spatial input strategy yielded lower energies than the local strategy.

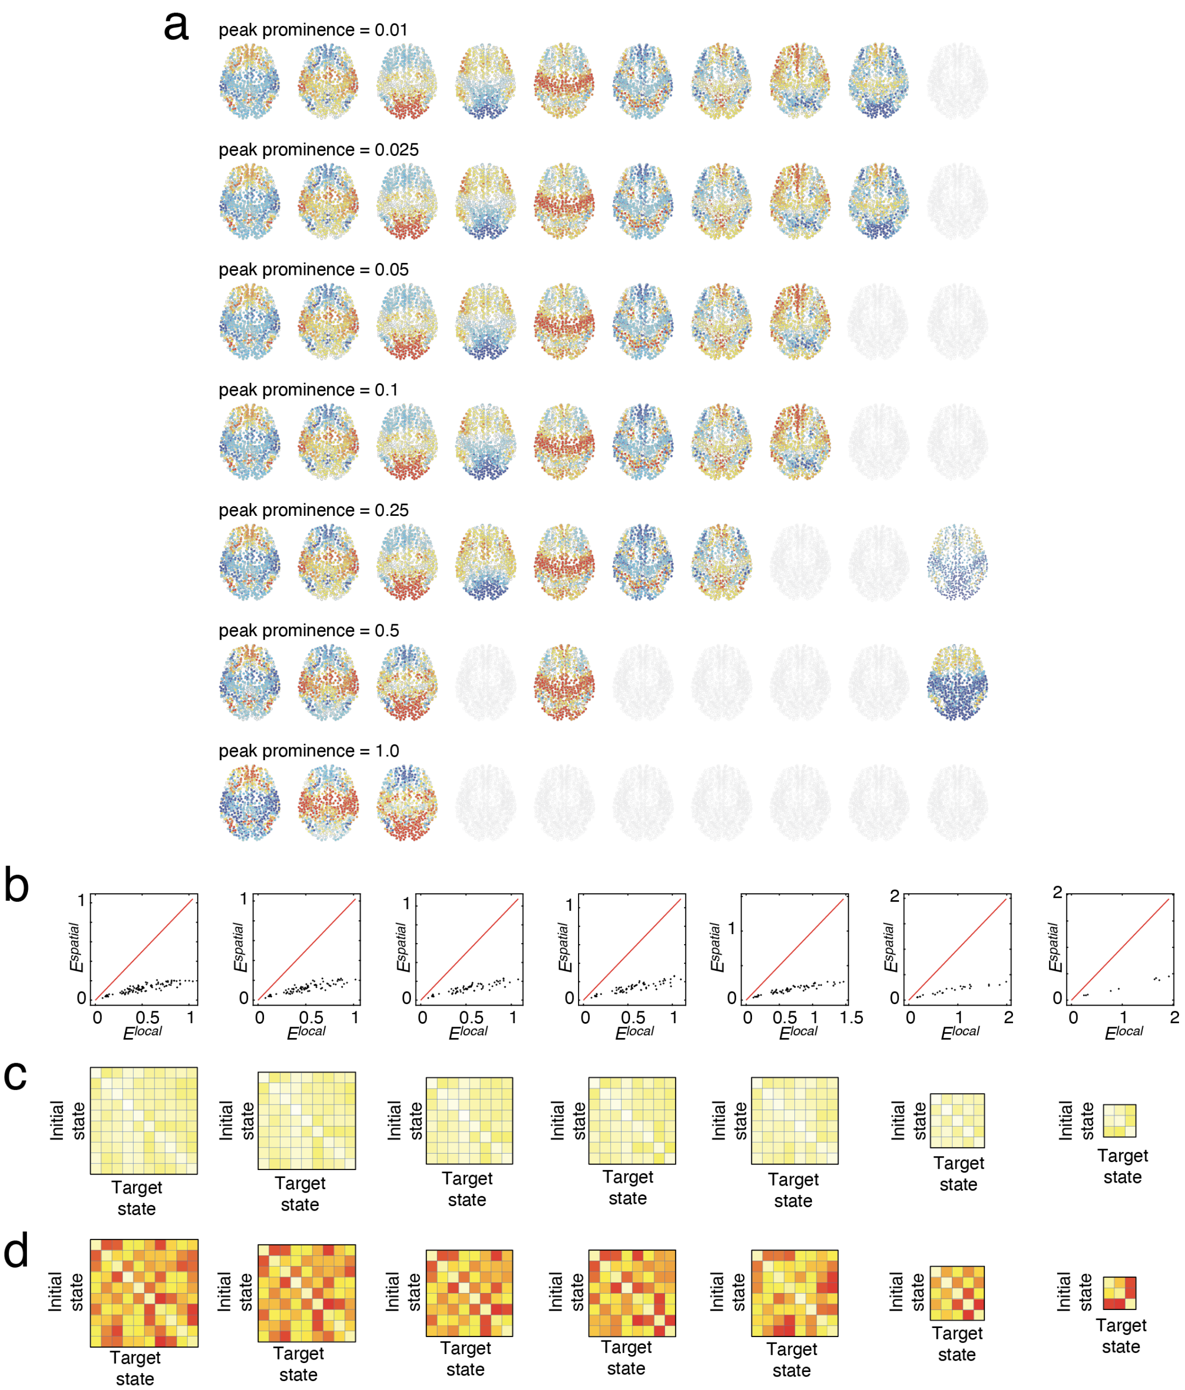

Figure S3. **Alternative state definitions.** In the main text we showed that spatially diffuse inputs reduce the control energy needed to transition between a set of 11 empirically defined brain states compared to the traditional "local" input strategy. The brain states we defined by sampling and clustering whole-brain activation patterns from peaks in the global signal. However, some of these peaks may reflect "noise" – i.e. though technically peaks, the peak may not be prominent. Here, we explore the effect of alternative brain state definitions in which we retain only the most prominent peaks by adjusting the 'prominence" parameter (the drop on both sides of the peak) from 0.01, 0.025, 0.05, 0.1, 0.25, 0.5, and 1. (*a*) Cluster centroids at each prominence value. Note that, as in the main text, we retain only those clusters in which at least 50% of the participants were represented. Cluster centroids have been aligned visually so that columns reflect, approximately, the same clusters across parameter values. Treating the clusters as estimates of "brain states", we calculated all pairwise transitions under both the spatial and local input strategies. (*b*) Scatter plot comparing energies under spatial and local models. As in the main text, on a per-transition basis, the spatial energy was consistently lower than the local energy. Panels *c* and *d* show energy values for the spatial input strategy (panel *c*) and local strategy (panel *d*).

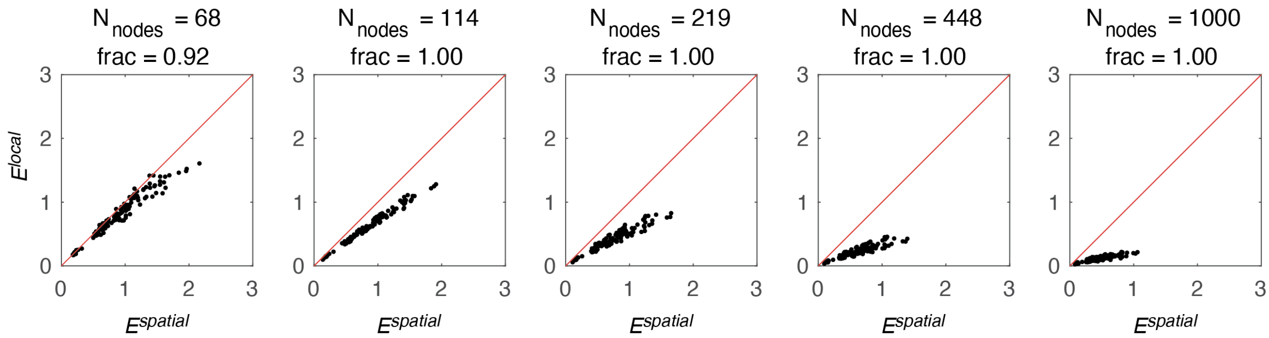

Figure S4. **Alternative parcel definitions.** The "Lausanne" parcellation is multiscale. The five scales contain 68, 114, 219, 448, and 1000 cortical parcels. We repeated our main analysis using connectomes and brain states defined at these other scales. In general, we found that our results replicated. Here, the x and y axes correspond to the energy for each transition under the spatial and local input strategies, respectively. In the title of each plot, "frac" refers to the fraction of transitions in which the spatial model yielded a smaller energy compared to the local model. That is, a value of 1.00 means that all 121 transitions resulted in lower energy under the spatial model.

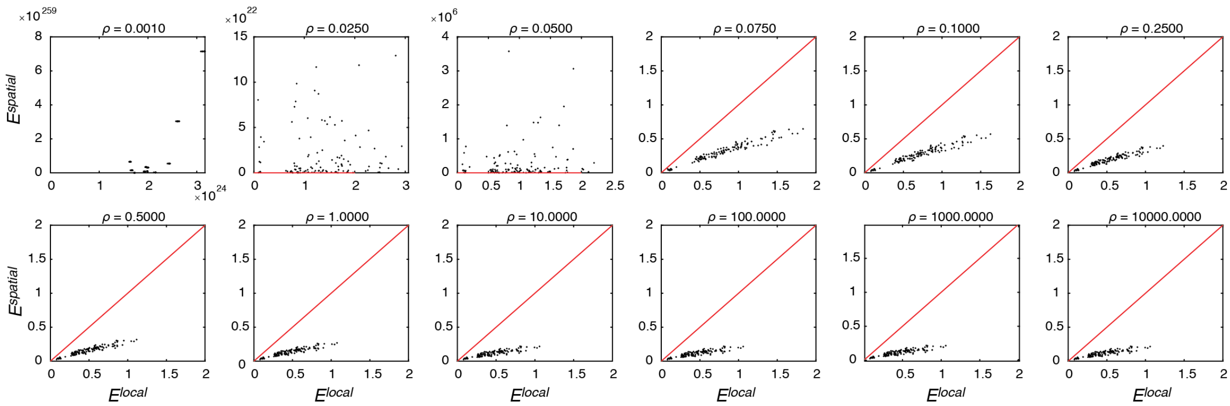

Figure S5. **Effect of varying ρ.** In the optimal control framework, the parameter ρ controls the balance of the two terms in the optimization: larger values of ρ upweight the importance of the control energy relative to the distance from the reference state. Here, we show that, across wide ranges of ρ (spanning roughly six decades, from $\rho = 10^{-1}$ to at least $\rho = 10^{4}$), the spatial input strategy results in reduced control energy compared to the local input strategy. At extremely small values of ρ, we find numerical instabilities, resulting in energies on the order of 10^{100}, and inaccurate estimates of the optimal control signal, making comparisons between the two input strategies difficult.

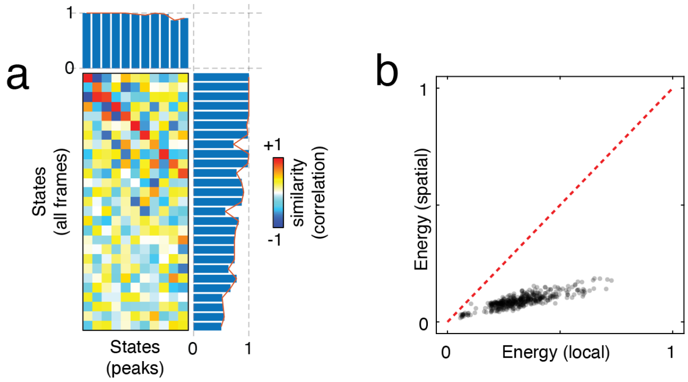

Figure S6. **Replicating main findings using alternative definition of brain states.** In the main text, we identified brain states by clustering activity patterns from peaks in the global activity amplitude (root mean square). Here, we repeated the primary analysis using an alternative definition of brain state. Namely, we clustered all frames, spanning all amplitudes of activity. Using this approach, we found 21 states instead of 11. However, among those 21 states were 11 that closely matched those reported in the main text. (*a*) Correlation (similarity) of states estimated from peak frames with those estimated using all frames. States were matched using the Hungarian algorithm. In the margins, we show on a scale of 0-1, the fraction of participants in which a given state appeared. (*b*) We also calculated the control energy needed to transition between all pairs of states. In line with results reported in the main text, we found that, on a per-transition basis, the energy required using the spatial model was less than what was required using the local model.

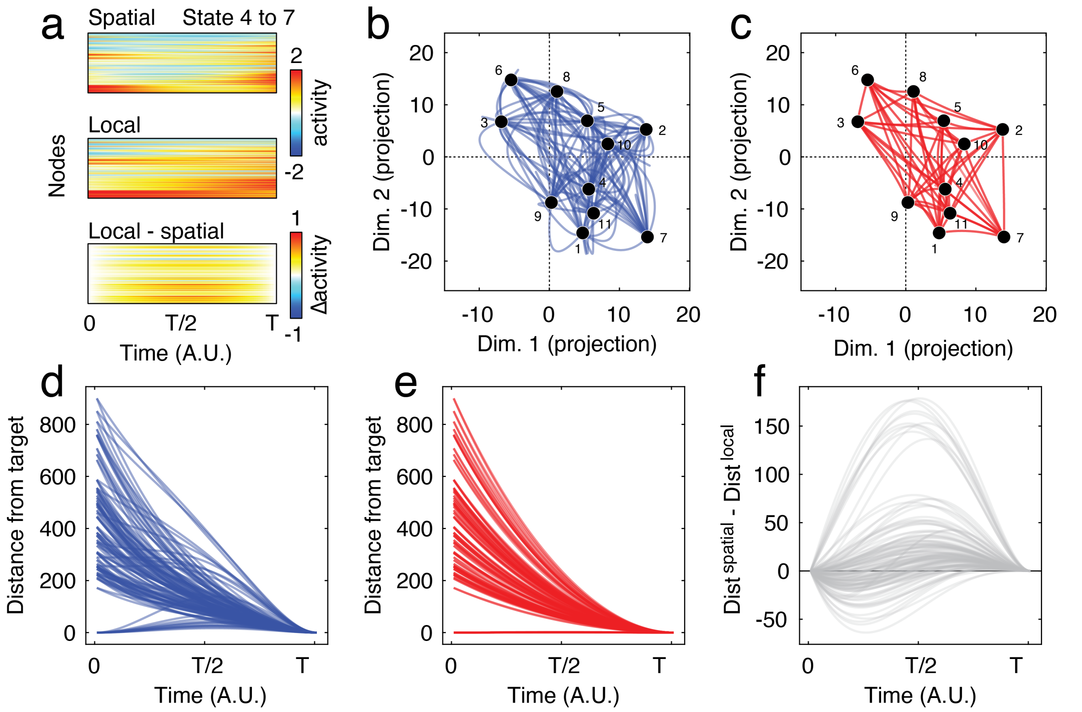

Figure S7. **Brain state trajectories under spatial *versus* local control strategies.** (a) Heatmaps showing nodes' states across an example control task (a transition from state 4 into state 7). Trajectories were estimated using a diffusivity parameter of $\beta = 0.15$. The top plot shows node state changes under the spatial input strategy, the middle shows local, and the bottom plot local minus spatial. In the main text, we showed that state trajectories were less direct under the spatial input strategy. That is, the trajectory was allowed to move further from the target state than compared to the local input strategy. To visualize this, we projected brain states onto two random $N \times 1$ vectors whose elements were sampled from a normal distribution with zero mean and unit variance. In panel b, we show low-dimensional representations of the transitions between all pairs of states under the spatial input strategy. Note that the trajectories tend to be curved. In panel c we show an analogous plot for transition under the local strategy. Note that, in this case, the trajectories tend to be straight lines. This effect is evident when we plot the distance from target for each transition as a function of time. Panels d and e show distances for all 121 transitions (spatial and local control strategies, respectively), while panel f shows the differences in distances. Positive/negative values indicate that the trajectory under the spatial/local input strategy was "loopier."

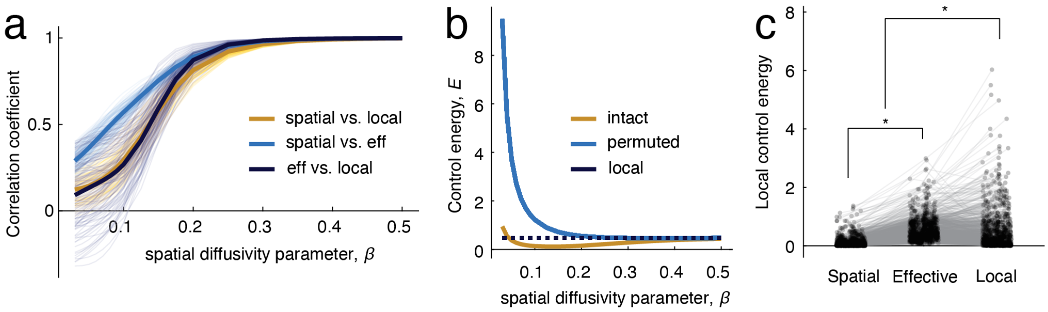

Figure S8. **Supplementary analyses.** (a) Similarity of control energy vectors as a function of the diffusivity parameter, β. Each curve corresponds to one of the possible 121 state transitions. (b) In the main text we show that the spatial input strategy leads to a reduction in the control energy. We hypothesized that this is because it, effectively, takes advantage of spatial correlations in structural connectivity and brain activations (states). That is, because nearby nodes tend to have similar state values and similar connectivity patterns, weaker but spatially diffuse inputs can have the same effect as the strong but spatially localized inputs. To test this, we simply randomly permuted the rows/columns of the input matrix, thereby destroying any spatial relationships but otherwise preserving exactly its total weight and global structure. We found that when we destroy spatial relationships the control energy increases dramatically. The gold curve is identical to the curve shown in Fig. 3f. The dark blue curve is the energy after destroying spatial structure. (c) Regional energies under the spatial input strategy, after calculating the effective energy under the spatial model, and under the local model. We find significant differences between all three distributions.

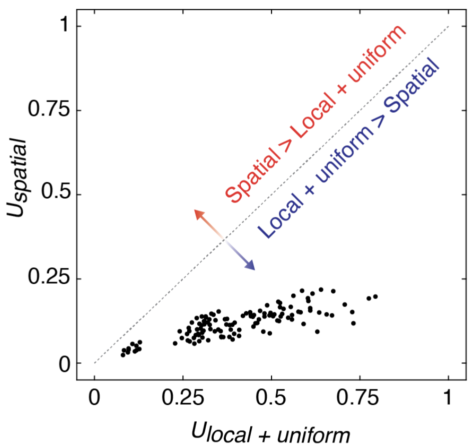

Figure S9. **Control energies under null model.** The spatial model effectively adds extra nonzero elements to the input matrix. To ensure that the reduction in control energy under the spatial model is not driven exclusively by the added weight, we compared the control energy for all 121 transitions against the energies obtained under the following null model. For column i, we calculate $w_i = \sum_{j \neq i} B_{ij}^{spatial}$. The value of w_i varies from node to node and corresponds to the additional weight added to column i under the spatial model. Next, we initialize a null input matrix as the identity matrix, i.e. $B_{ii} = 1$ for all i. We then set all off-diagonal elements, B_{ij} with $i \neq j$, equal to $\frac{w_i}{n-1}$. That is, we uniformly distribute the added weight across all non-identity elements within each column. This resulting input matrix has total weight equal to that of the spatial model, its column sums are identical to that of the spatial model, and it preserves the identity line (a feature that both the spatial and local models both exhibit). Next, using this matrix, we calculated the control energy associated with each of the 121 transitions, comparing it to the corresponding energy under the spatial model. We found that the energy under the uniform model was greater than that of the spatial model for every transition. This observation suggests that the observed reduction in energy under the spatial model is not driven solely by the weight added to the input matrix.

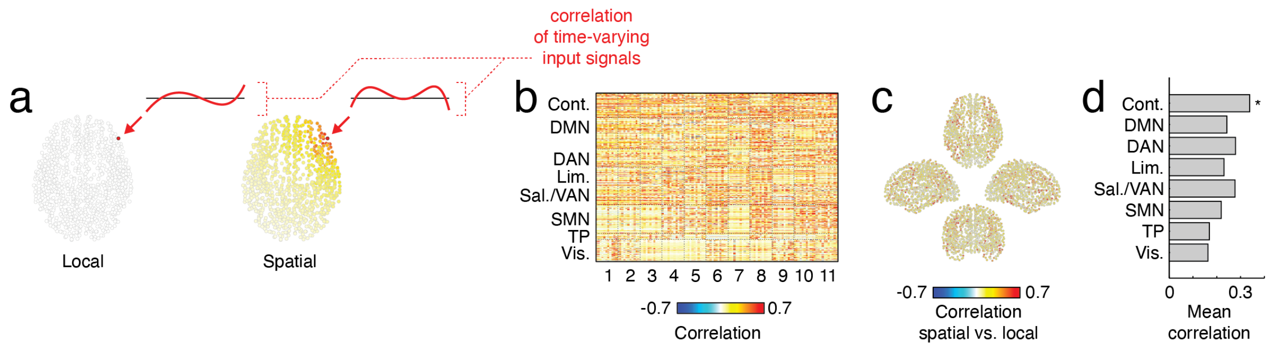

Figure S10. **Correlated input signals.** In the main text we described a procedure in which we calculated the correlation between regional input signals under the spatial and local input strategies. Here, we elaborate on that procedure. (*a*) Schematic showing input signals being delivered to the same node under different input strategies. Each element in the heatmap shown in *b* corresponds to a correlation coefficient. Rows correspond to nodes and columns correspond to transitions. The columns are ordered by target state. In panel *c* we show, in anatomical space, the mean correlation projected into anatomical space. In panel *d* we show mean correlation at the level of brain systems. The correlation in the control network was the only system whose mean correlation exceeded chance (spin tests).

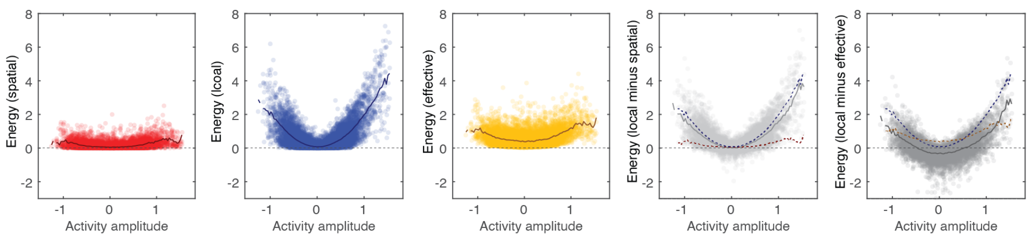

Figure S11. **Regions that drive differences between spatial and local input strategies.** In the main text we showed that the spatial input strategy reduced control energies compared to the local input strategy. Here, we show that these global reductions are driven by regions whose activity in the target state is extreme (highly active, positively or negatively). The left-most plot (red) depicts regional energies for all control tasks as a function of their activity. The next two plots are analogous to the first; they show (in blue and yellow) regional energies under the local strategy and their effective energy under the spatial strategy. The next panel shows the difference in energy – local minus spatial. It shows that the biggest differences are at the extremes; regions with high activity require proportionally less input under the spatial strategy than under the local strategy. Regions with low levels of activity are comparable between the two input strategies. The right-most panel is analogous to the previous, but shows the difference between local and effective energies. Note that, again, the energy of regions with extreme levels of activity correspond to the biggest gap. However, regions with relatively low levels of activity now receive slightly more energetic inputs compared to the local control strategy.

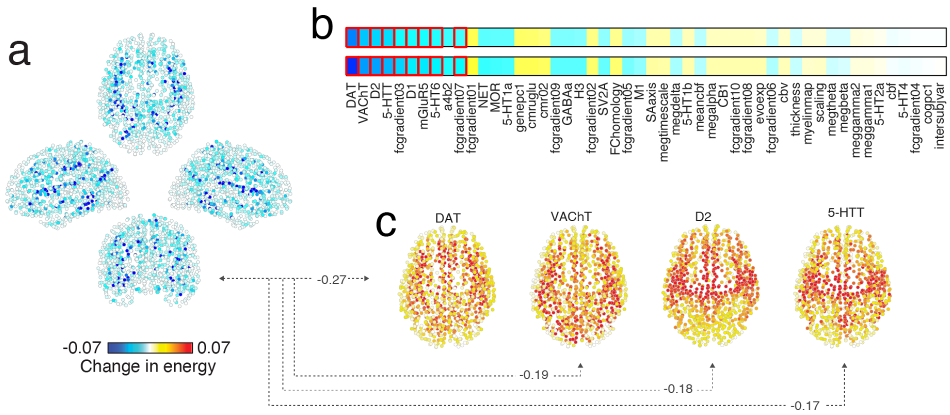

Figure S12. **Regions whose energy increased under the spatial input strategy.** In the main text we noted that a small subset of brain regions exhibited consistent increases in energy under the spatial input strategy compared to the local. (*a*) Spatial distribution of increases. (*b*) Correlation of the brain map from *a* with the annotations. Red outlines indicate statistically significant correlations. We show the annotation maps associated with the strongest statistically significant correlations in panel *c*.

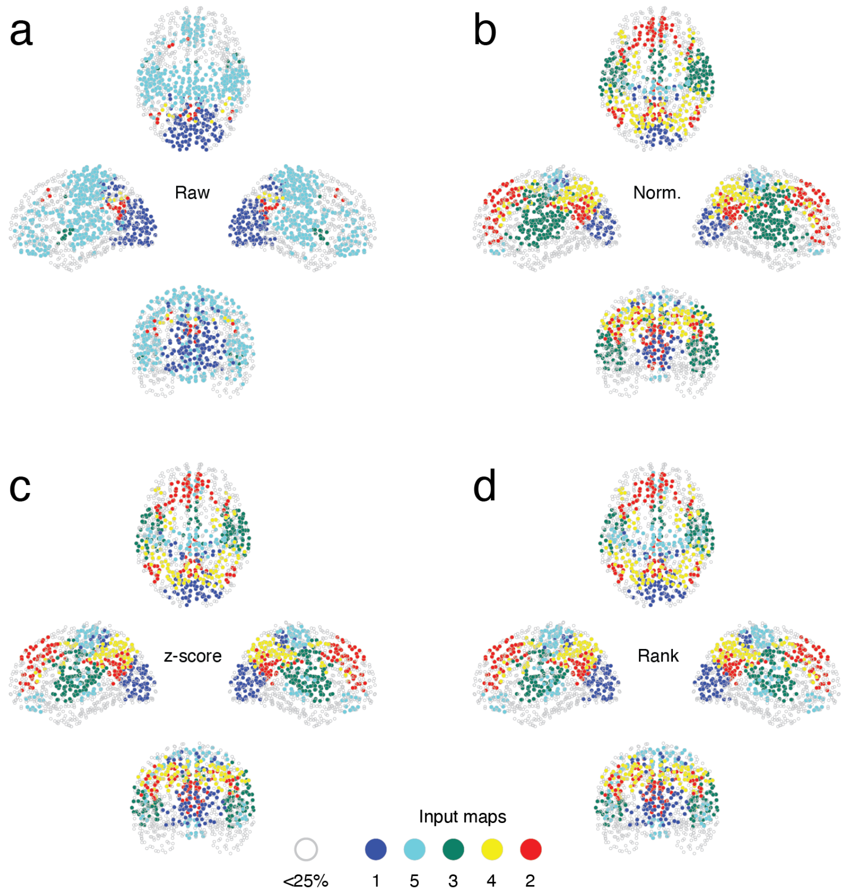

Figure S13. **Dominance of input maps.** In the main text we derived a series of brainwide input maps. Each map is associated with a specific input (control) signal; its elements correspond to how much of that signal should be delivered to each node. In the main text we focused on the top five maps, each of which appeared across many different control tasks. We also calculated for every node i its maximum input weight – that is, $I_i = \max\{m_i^1, m_i^2, m_i^3, m_i^4, m_i^5\}$. Here we show a series of "dominance maps". That is, the index corresponding to the maximum input weight. We show four versions of the dominance maps, each of which we threshold so that the lowest quartile of maximum input weights are "greyed out". (a) Dominance map using the raw (untransformed) input maps. (b) Dominance map after rescaling each map to the interval $[0, 1]$ by subtracting the smallest input weight and dividing by the maximum value. (c) Dominance map after z-scoring the input maps. (d) Dominance map after rank-transforming each input map.

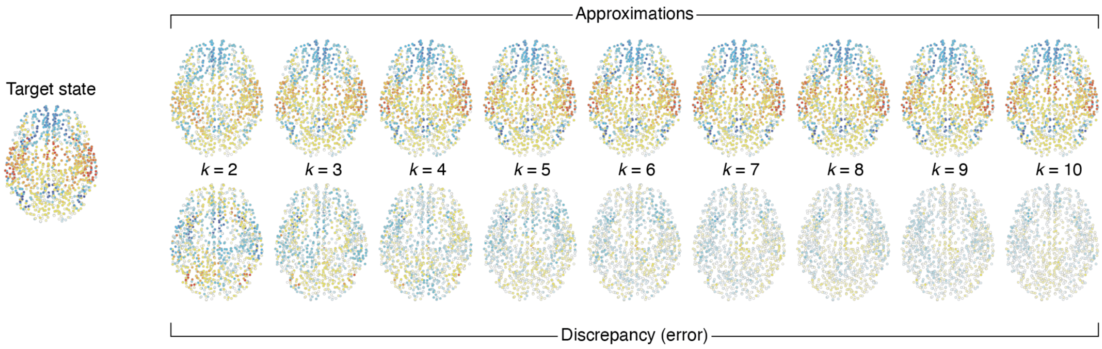

Figure S14. **Example of approximation algorithm.** On the left we show an example target state (state 1). The top row shows example approximations of the target state using k inputs instead of the N inputs usually required. The bottom row shows the error (approximation minus true state) for each k at a regional level.