Check for
updates

The MIT Press

RESEARCH

Quantifying brain state transition cost via Schrödinger Bridge

Genji Kawakita[1], Shunsuke Kamiya[1], Shuntaro Sasai[2,3], Jun Kitazono[1], and Masafumi Oizumi[1]

[1]Graduate School of Arts and Sciences, University of Tokyo, Tokyo, Japan
[2]Araya Inc., Tokyo, Japan
[3]University of Wisconsin–Madison, Madison, WI, USA

ABSTRACT

Quantifying brain state transition cost is a fundamental problem in systems neuroscience. Previous studies utilized network control theory to measure the cost by considering a neural system as a deterministic dynamical system. However, this approach does not capture the stochasticity of neural systems, which is important for accurately quantifying brain state transition cost. Here, we propose a novel framework based on optimal control in stochastic systems. In our framework, we quantify the transition cost as the Kullback-Leibler divergence from an uncontrolled transition path to the optimally controlled path, which is known as Schrödinger Bridge. To test its utility, we applied this framework to functional magnetic resonance imaging data from the Human Connectome Project and computed the brain state transition cost in cognitive tasks. We demonstrate correspondence between brain state transition cost and the difficulty of tasks. The results suggest that our framework provides a general theoretical tool for investigating cognitive functions from the viewpoint of transition cost.

AUTHOR SUMMARY

In our daily lives, we perform numerous tasks with different kinds and levels of cognitive demand. To successfully perform these tasks, the brain needs to modulate its spontaneous activity to reach an appropriate state for each task. Previous studies utilized optimal control in deterministic systems to measure the cost for the brain state transition. However, no unified framework for quantifying brain state transition cost that takes account of the stochasticity of neural activities has been proposed. Here, we describe a novel framework for measuring brain state transition cost, utilizing the idea of optimal control in stochastic systems. We assessed the utility of our framework for quantifying the cost of transitioning between various cognitive tasks. Our framework can be applied to very diverse settings because of its generality.

INTRODUCTION

The brain is considered a dynamical system that flexibly transitions through various states (Breakspear, 2017; McKenna, McMullen, & Shlesinger, 1994; Vyas, Golub, Sussillo, & Shenoy, 2020). Depending on the properties of a dynamical system (e.g., the biophysical properties of neurons and the connectivity between neurons), some transitions are difficult to realize. Thus,

Dynamical system:
A system whose state changes over time following a certain rule.

Network control theory:
A field of study that examines the control strategies of dynamic networked systems.

Structural connectivity:
Anatomical connections between brain regions.

Optimal control:
The control law for a dynamical system that optimizes a certain cost function.

Kullback-Leibler divergence:
An information-theoretic measure that quantifies the difference between two probability distributions.

Schrödinger Bridge:
The most likely time evolution of a system from an initial to a target probability distribution given the transition probability distribution of the system, which is mathematically equivalent to the optimally controlled path linking the initial and target distribution.

characterizing the dynamical properties of brain state transition would be important for understanding various brain functions (Kringelbach & Deco, 2020), including decision-making (Taghia et al., 2018), motor control (Shenoy, Sahani, & Churchland, 2013), and working memory (Simmering & Perone, 2012), with potential applications in the diagnosis and clinical treatment of disease (Adhikari et al., 2017; Aerts et al., 2020; Deco & Kringelbach, 2014). To date, however, no unified framework for quantifying brain state transition cost from brain activity data has been available.

One promising framework for quantifying brain state transition cost is the network control–theoretic framework (Medaglia, Pasqualetti, Hamilton, Thompson-Schill, & Bassett, 2017; see also Suweis et al., 2019; Tu et al., 2018, for some limitations). Control theory provides useful perspectives for measuring the cost required for controlling a dynamical system to reach a desirable state. Considering the brain as a dynamical system, control-theoretic approaches enable us to quantify the cost of transitioning to a brain state that produces desirable behavior. Recently, the network control-theoretic framework was proposed for study of the control property of the brain by viewing the brain as a networked dynamical system (Bassett & Sporns, 2017; Cornblath et al., 2020; Gu et al., 2017; Gu et al., 2015). Although the framework provides an important new perspective for fundamentally understanding brain state transition, it has two major limitations. First, it does not capture the stochasticity of brain activity, which is ubiquitous in brain activity and is essential for accurately describing brain dynamics (Deco, Rolls, & Romo, 2009; Rieke, 1999; Shadlen & Newsome, 1998). Disregarding stochasticity may result in an inaccurate estimation of transition cost. Second, the model obtained from structural connectivity, which is static over time, may not be able to capture change in the functional dynamics of the brain (Kringelbach & Deco, 2020), such as while performing tasks, for instance. Moreover, it is difficult to model even the resting-state dynamics from structural connectivity (Honey et al., 2009). Recently, alternative models using functional and effective connectivity have been proposed (Deng & Gu, 2020; Szymula, Pasqualetti, Graybiel, Desrochers, & Bassett, 2020), but these models still do not capture stochasticity. Thus, no unified framework able to take account of the key properties of brain dynamics is available.

Here, by employing control-theoretic approaches, we propose a novel framework for measuring brain state transition cost that can account for stochasticity. In our framework, we consider transition from a probability distribution of brain states to another distribution, rather than a transition from one brain state to another brain state (i.e., a point-to-point transition in a state-space) in contrast to a previous work (Cornblath et al., 2020) utilizing network control theory. To transition from an initial distribution to a target distribution, the brain needs to modulate (control) its baseline transition probability. Although there are many possible ways to reach the target distribution, in this study we consider the optimally controlled path only and estimate the lower bound of brain state transition cost. We propose defining the minimum brain state transition cost as the Kullback-Leibler (KL) divergence from the baseline uncontrolled path to the optimally controlled path, that is, the closest path to the original path, with the fixed initial and target distributions. The problem of finding the closest path to the original path connecting the initial and target distribution is known as Schrödinger Bridge problem (Schrödinger, 1931), which has been studied in the fields of stochastic process and optimal transport (Beghi, 1996; Chen, Georgiou, & Pavon, 2016a, 2016b; Dai Pra, 1991; Léonard, 2013).

Here, as proof of concept, we apply the proposed framework based on the Schrödinger Bridge problem to evaluate the cost of task switching (Monsell, 2003), an executive function for moving from one task to another. Specifically, we address two questions. First, is the cost of transitioning to a more difficult task larger? A previous study (Kitzbichler, Henson, Smith,

Nathan, & Bullmore, 2011) reported that performing effortful tasks drives larger reconfiguration of functional brain networks. We therefore hypothesized that transitioning to a more difficult task required a larger cost. Second, is the brain state transition cost asymmetric? Specifically, is the transition cost from an easier task to a more difficult task larger than the cost accompanying the reverse transition?

To address these questions, we apply our framework to functional magnetic resonance imaging (fMRI) data from the Human Connectome Project (HCP; Van Essen et al., 2013). We use fMRI data of n = 937 subjects in the resting state and in seven cognitive tasks. After preprocessing and parcellation, we computed the probability distributions of coarse-grain brain activity patterns for the rest and cognitive tasks (Cornblath et al., 2020; Lynn, Cornblath, Papadopoulos, Bertolero, & Bassett, 2020). We then calculated the brain state transition cost by finding the Schrödinger Bridge, that is, the optimally controlled path (Pavlichin, Quek, & Weissman, 2019) between the initial and target distributions of brain states. We found that the transition cost to a more difficult task carried a larger transition cost. We also observed that the transition cost between an easy and a difficult task is asymmetric. Overall, our findings provide a new perspective on the investigation of brain state transition, which may facilitate our understanding of cognitive functions.

Brain parcellation:
Partitioning the brain into spatially or functionally distinct regions.

RESULTS

Quantification of Brain State Transition Cost From the Schrödinger Bridge Problem

In this study, we propose a novel framework to quantify state transition cost in a stochastic neural system, building on the formulation of the Schrödinger Bridge problem (Beghi, 1996). We consider brain state transition to be the transition from an initial probability distribution of brain states to a target probability distribution. In order to reach the target probability distribution, the brain is assumed to follow some controlled paths. Although there are many possible paths that bridge the initial and target probability distributions, we look for the optimally controlled path that minimizes the Kullback-Leibler (KL) divergence from an uncontrolled to a controlled path. Here, we define brain state transition cost as the minimum KL divergence from an uncontrolled to a controlled path that bridges the initial and target probability distributions (Figure 1).

We can mathematically formulate brain state transition as follows. Let X_t be a random variable corresponding to a coarse-grained brain state at time t. We consider each brain state to be a discrete number included in the finite set $\mathcal{S}$ = {1, ..., k}, where k is the number of brain states. For instance, $X_t = i$ means that the brain is at the state i at time t. In this study, we used the k-means clustering algorithm to obtain these coarse-grained brain states from high-dimensional brain activity data as described later. Then, let (X_0, ..., X_T) be a time series of brain states that form a first-order Markov chain. We introduce a simplified notation for expressing the time series as $X_0^T = (X_0, ..., X_T)$, where subscript 0 represents the starting time point and superscript T represents the ending time point. We denote the joint probability distribution of the random variables by $q(X_0^T) = q(X_0, ..., X_T)$, which can be expressed using the Markov property as follows:

$$q(X_0^T) = q(X_0)\prod_{t=0}^{T-1}q(X_{t+1}|X_t). \tag{1}$$

Here, we consider a problem of controlling the distribution of brain states to a target distribution, π', at $t = T$ starting from an initial distribution, π, at $t = 0$. The initial distribution, π, is
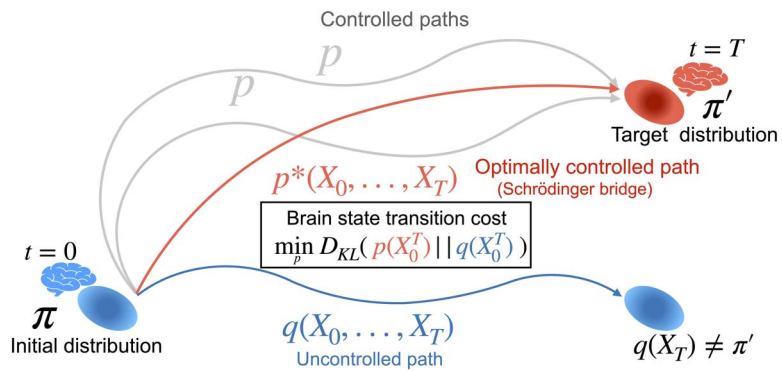
Figure 1. Schematic of brain state transition reframed as the Schrödinger Bridge problem. We consider brain state transition as transition from an initial probability distribution of brain states, π, to a target probability distribution, π'. The brain follows an uncontrolled baseline path, $q(X_0^T)$, which does not lead to the target distribution but to $q(X_T) \neq \pi'$, where $q(X_T)$ represents the probability distribution at $t = 0$, following the uncontrolled path. In order to reach the target distribution, the brain needs to follow a controlled path, $p(X_0^T)$. The brain state transition cost is defined as the minimum Kullback-Leibler divergence between the controlled and uncontrolled paths $p(X_0^T)$ and $q(X_0^T)$, respectively. Optimally controlled path, $p*(X_0^T)$ is equivalent to Schrödinger Bridge.

the same as $q(X_0)$, but the target distribution π' is different from $q(X_T)$, that is, the target distribution, π', cannot be reached by the original dynamics, $q(X_0^T)$. Thus, the transition probability of the brain needs to be modulated by some control input to the system. Although we do not explicitly model control input in this study because we do not model the dynamics of the system (see Chen et al., 2016a, for the model of linear systems with control input), we assume that some control input is implicitly applied to modulate the original dynamics. In this context of controlling the dynamics, we call the original dynamics, $q(X_0^T)$, the "uncontrolled path." In contrast with the uncontrolled path, $q(X_0^T)$, we call a controlled dynamics, $p(X_0^T) = p(X_0, \ldots, X_T)$, a "controlled path," which satisfies the end point constraints, $p(X_0) = \pi$ and $p(X_T) = \pi'$. While there are many possible controlled paths that satisfy the marginal conditions, we consider the problem of finding the optimally controlled path that minimizes some control cost. In this study, we define the control cost as the Kullback-Leibler (KL) divergence between the uncontrolled path and a controlled path (Beghi, 1996; Chen et al., 2016a),

$$D_{KL}\big(p(X_0^T)\|q(X_0^T)\big) = \sum_{X_0, X_1, \ldots, X_T \in \mathcal{S}} p\big(X_0^T\big) \log \frac{p\big(X_0^T\big)}{q\big(X_0^T\big)}. \tag{2}$$

Intuitively, KL divergence measures the difference between two probability distributions. If KL divergence between $p(X_0^T)$ and $q(X_0^T)$ is 0, then we can tell that two paths are equivalent, that is, the system does not change but stays the same. If the KL divergence takes nonzero values, it indicates that the system follows a different path from the uncontrolled path. Using the KL divergence as a transition cost is reasonable since the degree of KL divergence should reflect how different a controlled path is from the uncontrolled path. Here, we define the optimally controlled path, $p*(X_0^T)$, as the minimizer of KL divergence, as shown below:

$$p*\big(X_0^T\big) = \arg \min_p D_{KL}\big(p(X_0^T)\|q(X_0^T)\big). \tag{3}$$

In other words, the optimally controlled path, $p*(X_0^T)$, is the "closest" to the uncontrolled path, $q(X_0^T)$, in terms of KL divergence. Then, using the optimally controlled path, we define the minimum control cost $\mathcal{C}$ as

$$\mathcal{C} \equiv \min_{p(X_0^T)} D_{KL}\big(p(X_0^T)\|q(X_0^T)\big) = D_{KL}\big(p*(X_0^T)\|q(X_0^T)\big). \tag{4}$$

In this study, we propose to use the minimum control cost for quantifying the brain state transition cost. The problem of finding the optimally controlled path from an initial to a target distribution is known to be mathematically equivalent to the Schrödinger Bridge problem, the problem of finding the most likely path linking the initial and target distribution given the transition probability distribution of the system (Beghi, 1996; Chen et al., 2016a).

To solve the minimization problem, we first decompose the KL divergence into two terms, both of which are also KL divergences.

$$
\begin{aligned}
D_{KL}\big(p(X_0^T)\|q(X_0^T)\big) &= \sum_{X_0,...,X_T} p(X_0,...,X_T) \log \frac{p(X_0,...,X_T)}{q(X_0,...,X_T)} \\
&= \sum_{X_0,X_T} p(X_0,X_T) \log \frac{p(X_0,X_T)}{q(X_0,X_T)} \\
&\quad + \sum_{X_0,X_T} p(X_0,X_T) \sum_{X_1,...,X_{T-1}} p(X_1,...,X_{T-1}) \log \frac{p(X_1,...,X_{T-1}|X_0,X_T)}{q(X_1,...,X_{T-1}|X_0,X_T)} \\
&= D_{KL}(p(X_0,X_T)\|q(X_0,X_T)) + D_{KL}\big(p(X_1^{T-1}|X_0,X_T)\|q(X_1^{T-1}|X_0,X_T)\big),
\end{aligned}
\tag{5}
$$

where $X_1^{T-1} = (X_1, ..., X_{T-1})$. The two terms are both nonnegative and we can separately minimize the two terms. The minimum of the second term is obviously 0 when $p*(X_1^{T-1}|X_0, X_T) = q(X_1^{T-1}|X_0, X_T)$. Then, the minimization problem of finding the whole controlled path is reduced to the problem of finding the optimally controlled joint distribution of the end points, $p(X_0, X_T)$ (Beghi, 1996),

$$\min_{p(X_0^T)} D_{KL}\big(p(X_0^T)\|q(X_0^T)\big) = \min_{p(X_0,X_T)} D_{KL}(p(X_0,X_T)\|q(X_0,X_T)). \tag{6}$$

By introducing the useful notation $p(X_0, X_T) = P$, and $q(X_0, X_T) = Q$, where $P_{ij} = p(X_0 = i, X_T = j)$ and $Q_{ij} = q(X_0 = i, X_T = j)$ for ease of computation, we can rewrite the KL divergence as

$$D_{KL}(p(X_0,X_T)\|q(X_0,X_T)) = D_{KL}(P\|Q) = \sum_{i=1}^{k}\sum_{j=1}^{k} P_{ij} \log \frac{P_{ij}}{Q_{ij}}. \tag{7}$$

With these new notations, we can restate the original minimization problem (Equation 3) as

$$P* = \arg\min_{P} D_{KL}(P\|Q) \tag{8}$$

with the following constraints,

$$\sum_{j=1}^{k} P_{ij} = \pi_i, \quad \sum_{i=1}^{k} P_{ij} = \pi_j', \quad \sum_{i=1}^{k}\sum_{j=1}^{k} P_{ij} = 1. \tag{9}$$

To further clarify the mathematical property of the optimization problem, we rewrite the control cost as follows:

$$
\begin{aligned}
D_{KL}(P\|Q) &= -\sum_{i=1}^{k}\sum_{j=1}^{k} P_{ij} \log Q_{ij} + \sum_{i=1}^{k}\sum_{j=1}^{k} P_{ij} \log P_{ij} \\
&= \sum_{i=1}^{k}\sum_{j=1}^{k} C_{ij} P_{ij} - H(P),
\end{aligned}
\tag{10}
$$

where $H(P)$ is the entropy of the joint end point distribution, P. By rewriting the cost in this way, we can regard the problem of minimizing KL divergence as the entropy regularized optimal transport problem with the transportation cost matrix, $C_{ij} = -\log Q_{ij}$ (Amari, Karakida, & Oizumi, 2018; Cuturi, 2013; see e.g., Chen, Georgiou, & Pavon, 2021; De Bortoli, Thornton, Heng, & Doucet, 2021; Léonard, 2013, for the connection between the Schrödinger Bridge problem and optimal transport problem). The existence and uniqueness of the optimal solution, P^*, is guaranteed because this is a strongly convex optimization problem (Amari et al., 2018; Cuturi, 2013).

Here, we explicitly find the solution of the optimization problem using the method of Lagrange multipliers. Let $\mathcal{L}(P, \alpha, \beta)$ be the Lagrangian of Equation 10 with Lagrange multipliers, α_i and β_j.

$$\mathcal{L}(P, \alpha, \beta) = -\sum_{i=1}^{k}\sum_{j=1}^{k} P_{ij}\log Q_{ij} + \sum_{i=1}^{k}\sum_{j=1}^{k} P_{ij}\log P_{ij} - \sum_{i=1}^{k}\sum_{j=1}^{k}(\alpha_i + \beta_j)P_{ij}. \tag{11}$$

Differentiating Equation 11 with respect to P_{ij} yields

$$\frac{\partial \mathcal{L}}{\partial P_{ij}} = -\log Q_{ij} + \log P_{ij} + 1 - (\alpha_i + \beta_j). \tag{12}$$

By setting the partial derivative to 0, we obtain the following optimal solution

$$P_{ij}^* = ce^{\alpha_i + \beta_j}Q_{ij}, \tag{13}$$

where c is the normalization constant,

$$c = \frac{1}{\sum_{i=1}^{k}\sum_{j=1}^{k} e^{\alpha_i + \beta_j}Q_{ij}}. \tag{14}$$

We determine the Lagrange multipliers, α_i and β_j, by the constraints in Equation 9,

$$\sum_{j=1}^{k} P_{ij}^* = \sum_{j=1}^{k} ce^{\alpha_i + \beta_j}Q_{ij} = \pi_i, \tag{15}$$

$$\sum_{i=1}^{k} P_{ij}^* = \sum_{i=1}^{k} ce^{\alpha_i + \beta_j}Q_{ij} = \pi_j'. \tag{16}$$

With some manipulation of the above equations, we obtain

$$\alpha_i = \log(\pi_i) - \log\left(\sum_{j=1}^{k} e^{\beta_j}Q_{ij}\right) - \log c, \tag{17}$$

$$\beta_j = \log(\pi_j') - \log\left(\sum_{i=1}^{k} e^{\alpha_i}Q_{ij}\right) - \log c. \tag{18}$$

These Lagrange multipliers can be numerically determined by iteratively updating α_i and β_j according to the above equations starting from arbitrary initial values. This algorithm is known as the Sinkhorn algorithm (Cuturi, 2013; Sinkhorn, 1967). The implementation of the algorithm is available at https://github.com/oizumi-lab/SB_toolbox (Kawakita & Oizumi, 2021).

Quantification of Brain State Transition Cost in fMRI Data

To test the utility of our proposed method, we applied the Schrödinger Bridge–based framework to real fMRI data. We used resting-state fMRI and task-fMRI (emotion, gambling, language, motor, relational, social, and working memory) from the Human Connectome Project (HCP; Van Essen et al., 2013). We first performed preprocessing of the BOLD signals and parceled them into 100 cortical regions (Schaefer et al., 2018). As shown in Figure 2, we concatenated the pre-processed data of all subjects for all the tasks to obtain $M \times N$ time series data, where M is the number of cortical parcels (100) and N is the total time frames of the concatenated data. We consider a point in $M = 100$ dimensional space as the activity of the whole brain at a particular time frame. In total, there are N points in this high-dimensional space. We applied the k-means clustering algorithm to classify N points into k coarse-grained states. In this section, we show only the results when we set the number of coarse-grained states to $k = 8$ (see Supporting Information for the results with different numbers of coarse-grained states).

To compute the brain state transition cost (Equation 4), we need to obtain initial and target distributions as well as the joint probability distribution for the uncontrolled path. Here, we assume that an uncontrolled path in the brain is the resting-state transition probability. To determine the joint probability distribution for the uncontrolled path, we need to set the value of T, the time steps to reach the target distribution. We computed brain state transition cost with various T and observed that the results did not change qualitatively. Thus, we show here the results with $T = 1$, that is, the next time frame (see Supporting Information S6 for results when $T > 1$). A probability distribution for each task was computed as an empirical probability distribution by counting the number of the occurrences of the coarse-grained states in each task time series data. We estimated the joint probability distribution of the resting state for two consecutive frames by counting transition pairs of the coarse-grained brain states with trajectory bootstrapping (see Methods for more details). From the joint probability distribution, we obtained the transition probability matrix of the resting state. Using these probability distributions and the transition probability matrix of the resting state, we calculated brain state transition cost

Empirical probability distribution: A probability distribution estimated from empirical data.
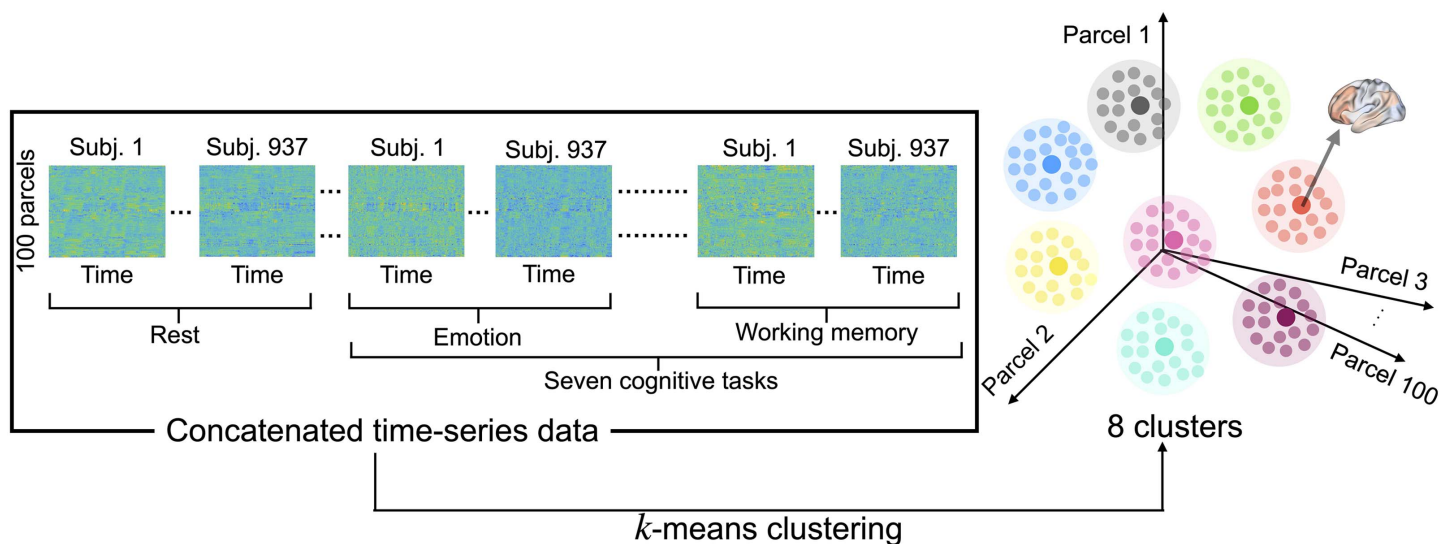

Figure 2. Clustering fMRI data. After preprocessing raw fMRI data, we concatenated the preprocessed data of all subjects for all the tasks. We then used k-means clustering to group similar brain activity patterns into eight coarse-grained brain states. Each point in the 100-dimensional state space corresponds to the activity of the whole brain at a particular time frame (see S5 in the Supporting Information for brain maps of the centroids of the eight clusters).
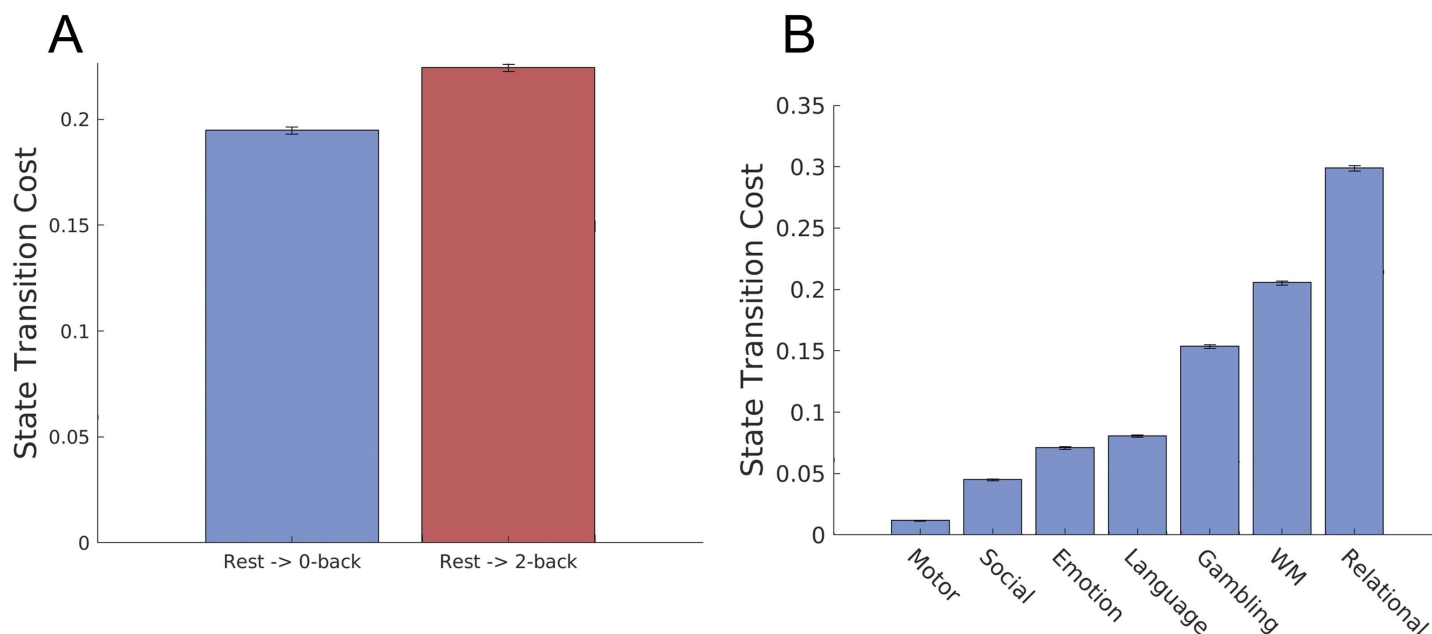
Figure 3. Brain state transition cost from the resting state to tasks. We computed the transition cost from the rest to cognitive tasks in the HCP data, as the minimum Kullback-Leibler divergence from the optimally controlled path to the baseline uncontrolled path. (A) The transition cost to the 2-back task (more difficult) is larger than the transition cost to 0-back task (easier) (one-sided *t* test, $p \ll 0.001$, $t > 60$, $df = 198$). (B) Transition cost from the rest to the seven cognitive tasks in the HCP dataset. Values are averaged over 100 bootstrapping trajectories and error bars indicate one standard deviation estimated with trajectory bootstrapping.

represented by the minimized KL divergence (Equation 4). For instance, when we computed the transition cost from the gambling task to the motor task, we set the initial distribution, π, to be the empirical probability distribution obtained from the gambling task data and the target distribution, π', to be the empirical probability distribution obtained from motor task data. Here, the probability distribution of the uncontrolled path, $Q_{ij} = q(X_0 = i, X_T = j)$, is computed by the product of the initial probability distribution π_i and the transition probability distribution of the resting state $Q_{j|i} = q(X_T = j|X_0 = i)$, $Q_{ij} = Q_{j|i}\pi_i$. With π, π', and Q, we can determine the optimally controlled path, P^*, and then compute the transition cost as $\mathcal{C} = D_{KL}(P^*||Q)$.

We began by testing whether transition cost from rest to a more difficult task is larger. For this purpose, we quantified the transition cost from the distribution at rest to those during 0-back (easier) and 2-back (more difficult) tasks in the working memory (WM) task data. We chose the WM task because the WM task data are the only task data in HCP, wherein subjects perform tasks with objectively different levels of task difficulty. As shown in Figure 3, we found that the transition cost to a 2-back task is larger than that to a 0-back task. This result suggests that our cost metric may capture the level of task difficulty from fMRI data.

N-back task:
A type of working memory task in which a subject reports whether a presented stimulus matches a stimulus shown *n* times ago.

Brain State Transition Cost to Multiple Tasks

To further check the behavior of the proposed metric for transition cost, we then computed brain state transition cost to multiple task distributions in the HCP dataset (emotion, gambling, language, motor, relational, social, and working memory). Note that, unlike the working memory tasks, we cannot objectively compare their task difficulties since these tasks are qualitatively different. Thus, the analysis here is exploratory without any prior hypothesis.

We found that the degree of transition cost to the seven cognitive tasks is significantly different. Figure 3B shows the rank order of transition costs in the seven cognitive tasks. Notably,

transition cost to a motor task was smallest among the seven tasks, whereas the transition cost to a relational task was the largest (see Discussion).

Asymmetry of Brain State Transition Cost

We then investigated whether state transition cost between tasks with different task difficulty was asymmetric. We hypothesized that it would require a larger transition cost to switch from an easier task to a more difficult task. To test this hypothesis, we computed the transition cost between 0-back and 2-back tasks in the working memory task. As shown in Figure 4A, we found that the transition cost from a 0-back task to a 2-back task was larger than the cost accompanying the reverse direction, which agreed with our hypothesis (one-sided t test, $p \ll 0.001$, $t > 80$, $df = 198$). Note that the asymmetry of the brain state transition cost does not result from the asymmetry of KL divergence, because the cost is not solely computed from the end point distributions but with underlying transition probability.

Finally, we examined whether the asymmetric property of brain state transition cost would be observed in other tasks whose task difficulties cannot be objectively compared. Here, we checked whether the following relationship would hold for all the pairs of tasks (note that here we regarded the rest as a task):

If the transition cost from rest to task A is larger than that from rest to task B, then the transition cost from task B to task A is larger than that in the reverse direction,

$$\mathcal{C}(\text{task B} \to \text{task A}) > \mathcal{C}(\text{task A} \to \text{task B}),$$
$$\text{if } \mathcal{C}(\text{rest} \to \text{task A}) > \mathcal{C}(\text{rest} \to \text{task B}),$$
(19)

where $\mathcal{C}(X \to Y)$ represents the brain state transition cost from X to Y, which is quantified by the KL divergence. To evaluate the relationship, we calculated the difference in transition cost between every pair of tasks, which is obtained as follows.

$$\text{Diff}(A, B) = \mathcal{C}(\text{task A} \to \text{task B}) - \mathcal{C}(\text{task B} \to \text{task A}).$$
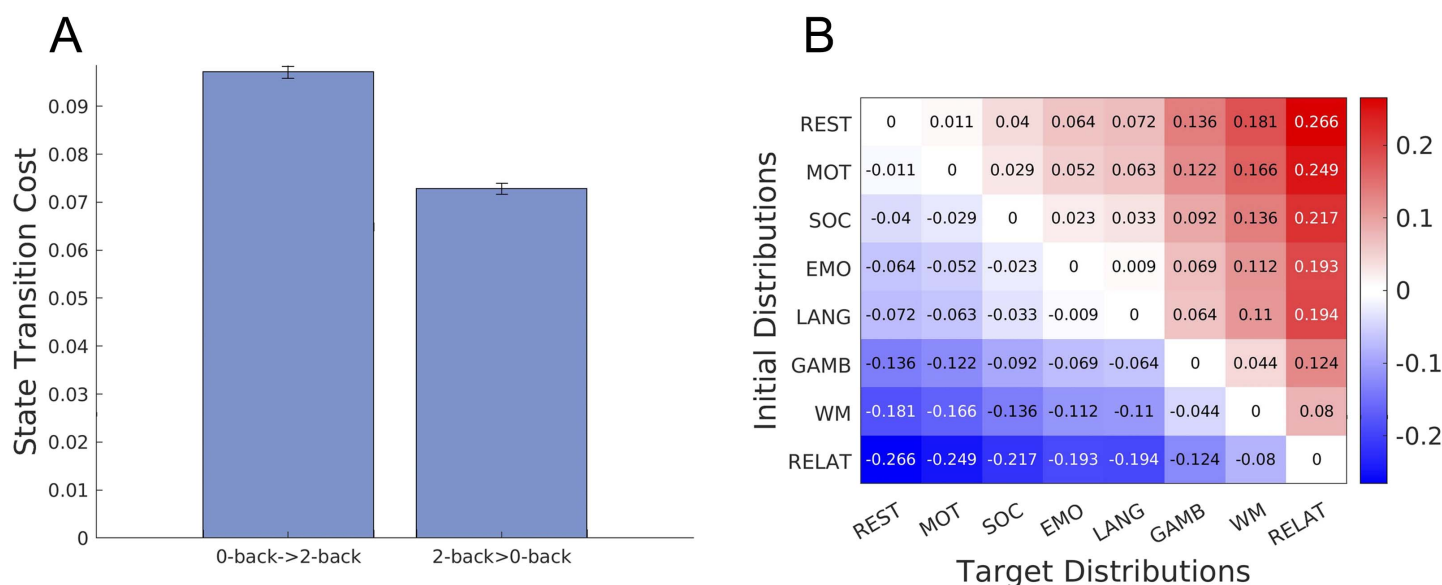(20)

Figure 4. Brain state transition cost between task states. (A) Brain state transition cost between the 0-back and 2-back tasks in the working memory task. Values are averaged over 100 bootstrapping trajectories and error bars indicate one standard deviation estimated using trajectory bootstrapping. (B) Asymmetry of brain state transition costs for the rest and seven cognitive tasks. Each element in the matrix represents a difference in transition cost between tasks.

The result is summarized in the matrix in Figure 4B, wherein entries (tasks) are arranged in ascending order by transition cost from rest, that is, the first row (column) corresponds to the task with the smallest transition cost from rest and the last row (column) corresponds to the task with the largest transition cost from rest. The (i, j) entry of the matrix represents Diff $(task_i, task_j) = \mathcal{C}(task\ i \rightarrow task\ j) - \mathcal{C}(task\ j \rightarrow task\ i)$. We observed that every entry in the upper (lower) triangular parts was positive (negative). This means that the relationship represented in Equation 19 holds for every pair of tasks in the dataset. That is, the transition cost is asymmetric between tasks with different degrees of transition cost.

DISCUSSION

In this study, we propose a novel framework for quantifying brain state transition cost in stochastic neural systems by framing brain state transition as the Schrödinger Bridge problem (SBP). This framework resolves the problem of previous methods that cannot take account of the inherent stochastisity of neural systems (Daunizeau, Stephan, & Friston, 2012; Deco et al., 2009) while still utilizing principled control-theoretic approaches. Under this framework, we assumed that the brain follows the resting-state activity as the baseline uncontrolled dynamics, and transitions to other distributions of brain state by modulating the baseline dynamics. Transition cost is measured as the minimum KL divergence from the uncontrolled path to the controlled path with the fixed end point probability distributions. We tested the utility of our framework by applying it to fMRI data from the Human Connectome Project release. The results indicated that the transition cost metric proposed in our framework may be useful for elucidating the characteristics of transition cost between tasks with different task difficulties.

Correspondence Between Transition Cost and Cognitive Demands

In the present study, we aimed to examine the relationship between the degree of brain state transition cost and task difficulty as proof of concept. We refer to task difficulty as objectively quantifiable task difficulty (e.g., 0-back and 2-back) only, not as a subjectively experienced task difficulty, which could vary among subjects. As for the objective task difficulty, we observed that the transition cost to a 2-back task (a more difficult task) is larger than that to a 0-back task (an easier task). Further studies using different types of tasks with various levels of difficulty are needed to determine the generality of this result.

On the other hand, we did not deal with subjectively experienced task difficulty or cognitive demands, as the dataset does not contain subjective reporting on the cognitive demand of each task. Nevertheless, we quantified the transition cost to the seven qualitatively different tasks, whose task difficulty or cognitive demand cannot be objectively quantified. Although it is unclear whether the observed order of transition cost correlates with subjective cognitive demand (Figure 3B), one may at least consider it reasonable that transition cost to a motor task is substantially smaller than that to a relational task. This is because performing a motor task only requires subjects to move a part of their body (e.g., right hand or tongue), whereas performing a relational task requires processing multiple visual stimuli and reasoning their relationships, which appears significantly more demanding than performing a motor task. Although we could not further examine whether the degree of transition cost correlates with the degree of cognitive demand, investigating this relation in more detail would be an interesting future work. We expect that while there may be a rough correlation between transition cost and cognitive demand, there can never be one-to-one correspondence, as many factors affect subjective evaluation of cognitive demand (Frömer, Lin, Dean Wolf, Inzlicht, & Shenhav, 2021; Kool, McGuire, Rosen, & Botvinick, 2010; McGuire & Botvinick, 2010; Rosenbaum & Bui, 2019). It would be intriguing

to investigate the difference in transition cost and cognitive demand and in what cases these behave similarly or differently. The direction of study proposed in the present study might be an important step toward bridging cognitive demand and brain activity.

Relation to Previous Theoretical Work

In the present study, we considered the brain dynamics as a discrete stochastic process by coarse-graining brain activity patterns, which reduced the computational cost. On the other hand, previous studies using network control-theoretic framework (Gu et al., 2015) employed a linear continuous process. Our framework can also be extended to a linear continuous stochastic process because the Schrödinger Bridge problem is not limited to a discrete process but has also been studied in continuous settings as well (Chen et al., 2016b; Dai Pra, 1991; Léonard, 2013). We therefore expected that we could directly fit the high-dimensional neural recording data with a continuous model (e.g., stochastic differential equations) as recently implemented (Nozari et al., 2020) and carry out a similar analysis as the present study. Developing a continuous version of the present framework may allow us to gain more insights into brain state transition cost.

Both the discrete stochastic process (model-free dynamics) utilized in the present study and linear dynamical models (Chen et al., 2016a; Gu et al., 2015) possess pros and cons for the application in the analysis of neural activity. In linear dynamical models, because the control input is explicitly modeled, the biophysical meaning of the control input is clear. By taking advantage of the high interpretability of the control input, linear dynamical models, for example, can provide insights into the contribution of each brain region to the control of the whole system (Gu et al., 2015). However, linear dynamical models are not suitable for the analysis of neural activity that is highly nonlinear. The discrete stochastic process used in the present study can be applied to nonlinear neural activity although the biophysical interpretation of control input is unclear. It is imperative to select an approach that fits the purpose of the study and the property of the data. By choosing appropriate models, we can compute brain state transition cost in various types of data (e.g., fMRI, EEG, ECoG).

Similarly to our framework, some recent works have utilized information theoretic measures to quantify cognitive costs (Zénon, Solopchuk, & Pezzulo, 2019) and connectivity changes between brain states (Amico, Arenas, & Goñi, 2019). While these studies compute only distance or divergence between the two distributions of brain state, our framework takes account of the underlying baseline activity of the brain. We employ this approach because including this baseline spontaneous activity provides a more accurate transition cost measure from the viewpoint of dynamical system theory.

Physical Interpretation of the Brain State Transition Cost

The KL divergence–based control cost proposed in this study may seem to be a distant concept from the conventional control cost, namely the time integral of squared input (Gu et al., 2017), in a linear deterministic model. However, it was shown in a previous study that the KL divergence cost in a stochastic linear model is analytically computed as the expectation of the time integral of squared input (Beghi, 1996; Chen et al., 2016a). In this sense, the KL divergence cost is tightly connected to the conventional control cost in a linear system.

The brain state transition cost proposed in the present study has a clear information theoretic meaning, that is, the KL divergence between the optimally controlled path and the uncontrolled path. However, an explicit physical interpretation has yet to be elucidated. A natural choice of control cost from the viewpoint of physics would be the work needed to realize a controlled path from an initial distribution to a target distribution (Chen, Georgiou, & Tannenbaum, 2020;

Horowitz, Zhou, & England, 2017). Minimizing the work is equivalent to minimizing entropy production (work dissipation; Chen et al., 2020; Horowitz et al., 2017). Previous works investigated the optimal control that minimized entropy production (or equivalently the work) and showed that the entropy production is lower bounded by the square of the Wasserstein distance (Chen et al., 2020; Nakazato & Ito, 2021). Interestingly, entropy production is given by the Kullback-Leibler divergence between two probabilities of forward and backward processes (Kawai, Parrondo, & Van den Broeck, 2007). Thus, one may consider that there should be some connection between the Schrödinger Bridge–type information theoretic control cost proposed in this study and the physical cost, work, or entropy production. It would be interesting to clarify the relationship between these different types of control cost.

Brain State Transition and Reconfiguration of Functional Connectivity

Functional connectivity:
Statistical dependencies (often simply referred to as correlation) between brain regions estimated from brain activity data.

The brain state transition cost computed in our framework may be related to the reconfiguration of functional connectivity between tasks. Numerous functional neuroimaging studies have reported the alterations in functional connectivity from the resting-state connectivity during task performance (Cole, Bassett, Power, Braver, & Petersen, 2014; Cole, Ito, Cocuzza, & Sanchez-Romero, 2021; Davison et al., 2015; Spadone et al., 2015; Stitt et al., 2017). In our present study, we showed that transitioning to more difficult tasks carries a larger transition cost. This seems to be consistent with Kitzbichler's work, which demonstrated that larger cognitive demand induces a more global alteration in brain activity (Kitzbichler et al., 2011). It may be the case that our framework captures the cost associated with the degree of reorganization of functional connectivity between different tasks.

Implications to Task Switching

The Schrödinger Bridge–based framework we proposed in this study may provide a new perspective for studying task switching from brain activity data. One of the most important and replicated findings in the task-switching paradigm is the observation of switch cost (Monsell, 2003): Switching to a new task takes a higher reaction time and error rate than repeating the same task. Various hypotheses have been proposed to explain the source of switch cost (Jersild, 1927; Koch, Gade, Schuch, & Philipp, 2010), including reconfiguration of the mental set for performing tasks. However, few studies have quantified the switch cost from brain activity. A recent work suggests that task switching involves the reconfiguration of brain-wide functional connectivity (Daws et al., 2020). Our framework may be used as a quantitative method for measuring switch cost. Subsequent investigations should study the relationship between switch cost and brain state transition cost by measuring brain activity while a subject is performing a task-switching experiment.

Comparison Between the Optimally Controlled Path and the Empirical Path

We investigated only the optimally controlled path, not an empirical path because of the limitation of fMRI recording. To quantify the efficiency of brain state transition, it would be interesting to compare empirical and optimal paths. This may provide insight into individual differences in the performance of task switching. However, the fMRI data from the Human Connectome Project does not include recordings in which subjects perform and switch between multiple tasks, and we were therefore unable to compute an empirical transition path between initial and target distributions of brain states. Even if the dataset contained such data, fMRI would not capture rapid transitions between tasks because the time resolution of the fMRI data is not sufficiently high (TR = 0.72 s in the HCP dataset). Computing empirical transition paths will require the use of recording data with better temporal resolution, such as EEG, MEG, or ECoG. In this regard, our theoretical framework is applicable to other types of recording data besides fMRI.

METHODS

fMRI Data Acquisition and Preprocessing

The 3T functional magnetic resonance imaging (fMRI) data of 937 subjects were obtained from the Washington University–Minnesota Consortium Human Connectome Project (HCP; Van Essen et al., 2013). Every subject provided a written informed consent to the Human Connectome Project consortium, following the HCP protocol. We used minimally preprocessed fMRI data at resting state and seven cognitive task states (emotion, gambling, language, motor, relational, social, and working memory). We selected these 937 subjects as they contain complete data for all the tasks. We then performed denoising by estimating nuisance regressors and subtracting them from the signal at every vertex (Satterthwaite et al., 2013). For this, we used 36 nuisance regressors and spike regressors introduced in a previous study, consisting of (1–6) six motion parameters, (7) a white matter time series, (8) a cerebrospinal fluid time series, (9) a global signal time series, (10–18) temporal derivatives of (1–9), and (19–36) quadratic terms for (1–18). Following a previous study (Satterthwaite et al., 2013), the spike regressors were computed with 1.5-mm movement as a spike identification threshold. After regressing these nuisance time courses, we also applied a band-pass filter (0.01–0.69 Hz) to the data, in which the upper bound of the filter corresponds to the Nyquist frequency of the time series. We then applied a parcellation proposed in (Schaefer et al., 2018) to divide the cortex into 100 brain regions, which reduced the complexity of the following analysis.

Clustering BOLD Signals

In order to model brain dynamics as a discrete stochastic process, we coarse-grained brain activity patterns using the *k*-means clustering algorithm. While there are numerous unsupervised clustering algorithms, we chose the *k*-means clustering because of its effective fit with the dynamics of neural activities (Cornblath et al., 2020). We used cosine similarity as a distance measure between centroids of clusters, which is commonly used in high-dimensional data. As described in the previous studies (Cornblath et al., 2020; Lynn et al., 2020), we concatenated preprocessed BOLD signals of all the subjects during the resting state and seven cognitive tasks. We obtained an $M \times N$ matrix, X, where M is the number of cortical parcels (100), and N is the number of task types times the number of time frames times the number of subjects. In order to prevent the variability in data size across tasks from affecting the clustering results, we used the same number of time frames for each task data. We used a different number of time frames depending on whether we divided the working memory task data into 0-back and 2-back tasks. When we divided the working memory task data (Figures 3A and 4A), we obtain 148 time frames that included either 0-back or 2-back task blocks. Accordingly, we used only the first 148 time frames in the other tasks. When we did not divide the working memory task data (Figures 3B and 4B), we used only the first 176 discrete measurements in each task since the emotion task—the task with the shortest measurement—was recorded for 176 time frames.

We determined the number of clusters using a similar procedure to that in previous studies (Cornblath et al., 2020; Lynn et al., 2020). We first computed the percentage variance explained by the number of clusters varying from $k = 2$ to $k = 12$. We observed that the explained variance plateaued around 75% after $k = 5$ (S4a in Supporting Information). We then examined whether all the coarse-grained states would appear in every subject during each task session (S4b in Supporting Information). We found that when we set the number of clusters to be greater than $k = 8$, some coarse-grained states did not appear in the data of some subjects. For these two reasons, we selected the number of clusters to be 8. While we chose $k = 8$, we were able to

reproduce the major results with $k > 4$ (S1–S3 in Supporting Information), which indicates the robustness of our results regardless of the number of clusters.

Estimating Transition Probabilities and Probability Distributions of Coarse-Grained States Using Trajectory Bootstrapping

We are limited by the finite length of the time series data for estimating brain state transition cost, which is calculated using probability distributions and transition probabilities. To ensure the accuracy of estimated quantities, we applied trajectory bootstrapping (Battle et al., 2016; Lynn et al., 2020) to calculate error bars on the estimated quantities. After we classified brain activity data with k-means clustering, we estimated the joint probability distribution matrix, $M_{ij} = q(X_\tau = i, X_{\tau+1} = j)$, of coarse-grained states at time $t = \tau$ and $t = \tau + 1$ for the resting state. To obtain the matrix, we first created a list of transitions in concatenated time series data of the resting state, in accordance with previous work (Lynn et al., 2020):

$$I = \begin{pmatrix} i_1 & i_2 \\ i_2 & i_3 \\ . & . \\ . & . \\ . & . \\ i_{L-1} & i_L \end{pmatrix},$$ (21)

where i_l is the coarse-grained state at lth frame of the time series, and L is the length of the concatenated time series (L = the number of time frames × the number of subjects). We sampled a pair of transitions from the list for L times to fill in the matrix, M, and normalized it to be a joint probability matrix. Although the transition list is a concatenated time series across subjects, we excluded pairs of transitions that took place across subjects; we only sampled pairs within the same subject. By normalizing each row of the matrix, M, to be 1, we constructed a transition probability matrix for the resting state. Similarly, we computed the probability distributions of coarse-grained brain state using trajectory bootstrapping. From the concatenated time series data of each task, we sampled a coarse-grained brain state L times. We counted the number of occurrences of each coarse-grained brain state and normalized it to 1 to obtain the probability distribution for each task. While we calculated a probability distribution for each task including the rest, we computed transition probability only for the resting state for obtaining the uncontrolled path, Q. We followed this process 100 times and computed the error bars on the estimated quantities in this study using the 100 bootstrap trajectories.

CODE AVAILABILITY

The code for computing the transition cost based on optimal control for stochastic systems is available at https://github.com/oizumi-lab/SB_toolbox.

SUPPORTING INFORMATION

Supporting Information (available at https://doi.org/10.1162/netn_a_00213) includes the following supplementary figures. S1: Robustness with different numbers of clusters. S2: Order of the degrees of transition costs from the resting state. S3: Asymmetry of brain transition cost. S4: Criteria for determining the number of clusters for k-means clustering algorithm. S5: Figures of brain maps. S6: Brain state transition cost when the time horizon, T, is set $T > 1$.

AUTHOR CONTRIBUTIONS

FUNDING INFORMATION

REFERENCES

Adhikari, M. H., Hacker, C. D., Siegel, J. S., Griffa, A., Hagmann, P., Deco, G., & Corbetta, M. (2017). Decreased integration and information capacity in stroke measured by whole brain models of resting state activity. *Brain*, *140*(4), 1068–1085. https://doi.org/10.1093/brain/awx021, PubMed: 28334882

Aerts, H., Schirner, M., Dhollander, T., Jeurissen, B., Achten, E., Van Roost, D., … Marinazzo, D. (2020). Modeling brain dynamics after tumor resection using the virtual brain. *NeuroImage*, *213*, 116738. https://doi.org/10.1016/j.neuroimage.2020.116738, PubMed: 32194282

Amari, S.-I., Karakida, R., & Oizumi, M. (2018). Information geometry connecting Wasserstein distance and Kullback–Leibler divergence via the entropy-relaxed transportation problem. *Information Geometry*, *1*(1), 13–37. https://doi.org/10.1007/s41884-018-0002-8

Amico, E., Arenas, A., & Goñi, J. (2019). Centralized and distributed cognitive task processing in the human connectome. *Network Neuroscience*, *3*(2), 455–474. https://doi.org/10.1162/netn_a_00072, PubMed: 30793091

Bassett, D. S., & Sporns, O. (2017). Network neuroscience. *Nature Neuroscience*, *20*(3), 353–364. https://doi.org/10.1038/nn.4502, PubMed: 28230844

Battle, C., Broedersz, C. P., Fakhri, N., Geyer, V. F., Howard, J., Schmidt, C. F., & MacKintosh, F. C. (2016). Broken detailed balance at mesoscopic scales in active biological systems. *Science*, *352*(6285), 604–607. https://doi.org/10.1126/science.aac8167, PubMed: 27126047

Beghi, A. (1996). On the relative entropy of discrete-time markov processes with given end-point densities. *IEEE Transactions on Information Theory*, *42*(5), 1529–1535. https://doi.org/10.1109/18.532893

Breakspear, M. (2017). Dynamic models of large-scale brain activity. *Nature Neuroscience*, *20*(3), 340–352. https://doi.org/10.1038/nn.4497, PubMed: 28230845

Chen, Y., Georgiou, T. T., & Pavon, M. (2016a). On the relation between optimal transport and Schrödinger bridges: A stochastic

control viewpoint. *Journal of Optimization Theory and Applications*, *169*(2), 671–691. https://doi.org/10.1007/s10957-015-0803-z

Chen, Y., Georgiou, T. T., & Pavon, M. (2016b). Optimal steering of a linear stochastic system to a final probability distribution, part I. *IEEE Transactions on Automatic Control*, *61*(5), 1158–1169. https://doi.org/10.1109/TAC.2015.2457784

Chen, Y., Georgiou, T. T., & Pavon, M. (2021). Stochastic control liaisons: Richard Sinkhorn meets Gaspard Monge on a Schrödinger Bridge. *SIAM Review*, *63*(2), 249–313. https://doi.org/10.1137/20M1339982

Chen, Y., Georgiou, T. T., & Tannenbaum, A. (2020). Stochastic control and nonequilibrium thermodynamics: Fundamental limits. *IEEE Transactions on Automatic Control*, *65*(7), 2979–2991. https://doi.org/10.1109/TAC.2019.2939625, PubMed: 33746240

Cole, M. W., Bassett, D. S., Power, J. D., Braver, T. S., & Petersen, S. E. (2014). Intrinsic and task-evoked network architectures of the human brain. *Neuron*, *83*(1), 238–251. https://doi.org/10.1016/j.neuron.2014.05.014, PubMed: 24991964

Cole, M. W., Ito, T., Cocuzza, C., & Sanchez-Romero, R. (2021). The functional relevance of task-state functional connectivity. *Journal of Neuroscience*, *41*(12), 2684–2702. https://doi.org/10.1523/JNEUROSCI.1713-20.2021, PubMed: 33542083

Cornblath, E. J., Ashourvan, A., Kim, J. Z., Betzel, R. F., Ciric, R., Adebimpe, A., … 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), 261. https://doi.org/10.1038/s42003-020-0961-x, PubMed: 32444827

Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in Neural Information Processing Systems*, *26*, 2292–2300.

Dai Pra, P. (1991). A stochastic control approach to reciprocal diffusion processes. *Applied Mathematics and Optimization*, *23*(1), 313–329. https://doi.org/10.1007/BF01442404

Daunizeau, J., Stephan, K. E., & Friston, K. J. (2012). Stochastic dynamic causal modelling of fMRI data: Should we care about neural noise? *NeuroImage*, *62*(1), 464–481. https://doi.org/10.1016/j.neuroimage.2012.04.061, PubMed: 22579726

Davison, E. N., Schlesinger, K. J., Bassett, D. S., Lynall, M.-E., Miller, M. B., Grafton, S. T., & Carlson, J. M. (2015). Brain network adaptability across task states. *PLoS Computational Biology*, *11*(1), e1004029. https://doi.org/10.1371/journal.pcbi.1004029, PubMed: 25569227

Daws, R. E., Scott, G., Soreq, E., Leech, R., Hellyer, P. J., & Hampshire, A. (2020). Optimisation of brain states and behavioural strategies when learning complex tasks.

De Bortoli, V., Thornton, J., Heng, J., & Doucet, A. (2021). Diffusion Schrödinger bridge with applications to score-based generative modeling.

Deco, G., & Kringelbach, M. L. (2014). Great expectations: Using whole-brain computational connectomics for understanding neuropsychiatric disorders. *Neuron*, *84*(5), 892–905. https://doi.org/10.1016/j.neuron.2014.08.034, PubMed: 25475184

Deco, G., Rolls, E. T., & Romo, R. (2009). Stochastic dynamics as a principle of brain function. *Progress in Neurobiology*, *88*(1), 1–16. https://doi.org/10.1016/j.pneurobio.2009.01.006, PubMed: 19428958

Deng, S., & Gu, S. (2020). Controllability analysis of functional brain networks.

Frömer, R., Lin, H., Dean Wolf, C. K., Inzlicht, M., & Shenhav, A. (2021). Expectations of reward and efficacy guide cognitive control allocation. *Nature Communications*, *12*(1), 1030. https://doi.org/10.1038/s41467-021-21315-z, PubMed: 33589626

Gu, S., Betzel, R. F., Mattar, M. G., Cieslak, M., Delio, P. R., Grafton, S. T., ... Bassett, D. S. (2017). Optimal trajectories of brain state transitions. *NeuroImage*, *148*, 305–317. https://doi.org/10.1016/j.neuroimage.2017.01.003, PubMed: 28088484

Gu, S., Pasqualetti, F., Cieslak, M., Telesford, Q. K., Yu, A. B., Kahn, A. E., ... Bassett, D. S. (2015). Controllability of structural brain networks. *Nature Communications*, *6*, 8414. https://doi.org/10.1038/ncomms9414, PubMed: 26423222

Honey, C. J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J. P., Meuli, R., & Hagmann, P. (2009). Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences*, *106*(6), 2035–2040. https://doi.org/10.1073/pnas.0811168106, PubMed: 19188601

Horowitz, J. M., Zhou, K., & England, J. L. (2017). Minimum energetic cost to maintain a target nonequilibrium state. *Physical Review E*, *95*(4–1), 042102. https://doi.org/10.1103/PhysRevE.95.042102, PubMed: 28505816

Jersild, A. T. (1927). Mental set and shift. *Archives of Psychology*, *14*, 89.

Kawai, R., Parrondo, J. M. R., & Van den Broeck, C. (2007). Dissipation: The phase-space perspective. *Physical Review Letters*, *98*(8), 080602. https://doi.org/10.1103/PhysRevLett.98.080602, PubMed: 17359081

Kawakita, G., & Oizumi, M. (2021). Schrödinger's Bridge toolbox, GitHub, https://github.com/oizumi-lab/SB_toolbox

Kitzbichler, M. G., Henson, R. N. A., Smith, M. L., Nathan, P. J., & Bullmore, E. T. (2011). Cognitive effort drives workspace configuration of human brain functional networks. *Journal of*

Neuroscience, *31*(22), 8259–8270. https://doi.org/10.1523/JNEUROSCI.0440-11.2011, PubMed: 21632947

Koch, I., Gade, M., Schuch, S., & Philipp, A. M. (2010). The role of inhibition in task switching: A review. *Psychonomic Bulletin and Review*, *17*(1), 1–14. https://doi.org/10.3758/PBR.17.1.1, PubMed: 20081154

Kool, W., McGuire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision making and the avoidance of cognitive demand. *Journal of Experimental Psychology: General*, *139*(4), 665–682. https://doi.org/10.1037/a0020198, PubMed: 20853993

Kringelbach, M. L., & Deco, G. (2020). Brain states and transitions: Insights from computational neuroscience. *Cell Reports*, *32*(10), 108128. https://doi.org/10.1016/j.celrep.2020.108128, PubMed: 32905760

Léonard, C. (2013). A survey of the Schrödinger problem and some of its connections with optimal transport.

Lynn, C. W., Cornblath, E. J., Papadopoulos, L., Bertolero, M. A., & Bassett, D. S. (2020). Non-equilibrium dynamics and entropy production in the human brain.

McGuire, J. T., & Botvinick, M. M. (2010). Prefrontal cortex, cognitive control, and the registration of decision costs. *Proceedings of the National Academy of Sciences*, *107*(17), 7922–7926. https://doi.org/10.1073/pnas.0910662107, PubMed: 20385798

McKenna, T. M., McMullen, T. A., & Shlesinger, M. F. (1994). The brain as a dynamic physical system. *Neuroscience*, *60*(3), 587–605. https://doi.org/10.1016/0306-4522(94)90489-8

Medaglia, J. D., Pasqualetti, F., Hamilton, R. H., Thompson-Schill, S. L., & Bassett, D. S. (2017). Brain and cognitive reserve: Translation via network control theory. *Neuroscience and Biobehavioral Reviews*, *75*, 53–64. https://doi.org/10.1016/j.neubiorev.2017.01.016, PubMed: 28104411

Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, *7*(3), 134–140. https://doi.org/10.1016/S1364-6613(03)00028-7

Nakazato, M., & Ito, S. (2021). Geometrical aspects of entropy production in stochastic thermodynamics based on Wasserstein distance.

Nozari, E., Stiso, J., Caciagli, L., Cornblath, E. J., He, X., Bertolero, M. A., ... Bassett, D. S. (2020). Is the brain macroscopically linear? A system identification of resting state dynamics.

Pavlichin, D. S., Quek, Y., & Weissman, T. (2019). Minimum power to maintain a nonequilibrium distribution of a Markov chain.

Rieke, F. (1999). *Spikes: Exploring the neural code*. MIT Press.

Rosenbaum, D. A., & Bui, B. V. (2019). Does task sustainability provide a unified measure of subjective task difficulty? *Psychonomic Bulletin Review*, *26*(6), 1980–1987. https://doi.org/10.3758/s13423-019-01631-8, PubMed: 31240614

Satterthwaite, T. D., Elliott, M. A., Gerraty, R. T., Ruparel, K., Loughead, J., Calkins, M. E., ... Wolf, D. H. (2013). An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. *NeuroImage*, *64*, 240–256. https://doi.org/10.1016/j.neuroimage.2012.08.052, PubMed: 22926292

Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X.-N., Holmes, A. J., ... Yeo, B. T. T. (2018). Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cerebral Cortex*, *28*(9), 3095–3114. https://doi.org/10.1093/cercor/bhx179, PubMed: 28981612

Schrödinger, E. (1931). Uber die umkehrung der naturgesetze. sitz. ber. der preuss. *Akad. Wissen, Berlin Phys. Math*, *144*.

Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *Journal of Neuroscience*, *18*(10), 3870–3896. https://doi.org/10.1523/JNEUROSCI.18-10-03870.1998, PubMed: 9570816

Shenoy, K. V., Sahani, M., & Churchland, M. M. (2013). Cortical control of arm movements: A dynamical systems perspective. *Annual Review of Neuroscience*, *36*, 337–359. https://doi.org/10.1146/annurev-neuro-062111-150509, PubMed: 23725001

Simmering, V. R., & Perone, S. (2012). Working memory capacity as a dynamic process. *Frontiers in Psychology*, *3*, 567. https://doi.org/10.3389/fpsyg.2012.00567, PubMed: 23335902

Sinkhorn, R. (1967). Diagonal equivalence to matrices with prescribed row and column sums. *The American Mathematical Monthly*, *74*(4), 402–405. https://doi.org/10.2307/2314570

Spadone, S., Della Penna, S., Sestieri, C., Betti, V., Tosoni, A., Perrucci, M. G., ... Corbetta, M. (2015). Dynamic reorganization of human resting-state networks during visuospatial attention. *Proceedings of the National Academy of Sciences*, *112*(26), 8112–8117. https://doi.org/10.1073/pnas.1415439112, PubMed: 26080395

Stitt, I., Hollensteiner, K. J., Galindo-Leon, E., Pieper, F., Fiedler, E., Stieglitz, T., ... Engel, A. K. (2017). Dynamic reconfiguration of cortical functional connectivity across brain states. *Scientific Reports*, *7*(1), 8797. https://doi.org/10.1038/s41598-017-08050-6, PubMed: 28821753

Suweis, S., Tu, C., Rocha, R. P., Zampieri, S., Zorzi, M., & Corbetta, M. (2019). Brain controllability: Not a slam dunk yet. *NeuroImage*, *200*, 552–555. https://doi.org/10.1016/j.neuroimage.2019.07.012, PubMed: 31291605

Szymula, K. P., Pasqualetti, F., Graybiel, A. M., Desrochers, T. M., & Bassett, D. S. (2020). Habit learning supported by efficiently controlled network dynamics in naive macaque monkeys.

Taghia, J., Cai, W., Ryali, S., Kochalka, J., Nicholas, J., Chen, T., & Menon, V. (2018). Uncovering hidden brain state dynamics that regulate performance and decision-making during cognition. *Nature Communications*, *9*(1), 2505. https://doi.org/10.1038/s41467-018-04723-6, PubMed: 29950686

Tu, C., Rocha, R. P., Corbetta, M., Zampieri, S., Zorzi, M., & Suweis, S. (2018). Warnings and caveats in brain controllability. *NeuroImage*, *176*, 83–91. https://doi.org/10.1016/j.neuroimage.2018.04.010, PubMed: 29654874

Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., Ugurbil, K., & WU-Minn HCP Consortium. (2013). The WU-Minn Human Connectome Project: An overview. *NeuroImage*, *80*, 62–79. https://doi.org/10.1016/j.neuroimage.2013.05.041, PubMed: 23684880

Vyas, S., Golub, M. D., Sussillo, D., & Shenoy, K. V. (2020). Computation through neural population dynamics. *Annual Review Neuroscience*, *43*, 249–275. https://doi.org/10.1146/annurev-neuro-092619-094115, PubMed: 32640928

Zénon, A., Solopchuk, O., & Pezzulo, G. (2019). An information-theoretic perspective on the costs of cognition. *Neuropsychologia*, *123*, 5–18. https://doi.org/10.1016/j.neuropsychologia.2018.09.013, PubMed: 30268880