

REVIEW

Energy landscape analysis based on the Ising model: Tutorial review

Naoki Masuda^{1,2*}, Saiful Islam², Si Thu Aung¹, Takamitsu Watanabe³

1 Department of Mathematics, State University of New York at Buffalo, Buffalo, New York, United States of America, **2** Institute for Artificial Intelligence and Data Science, State University of New York at Buffalo, Buffalo, New York, United States of America, **3** International Research Centre for Neurointelligence, The University of Tokyo, Tokyo, Japan

* naokimas@gmail.com

Abstract

We review a class of energy landscape analysis method that uses the Ising model and takes multivariate time series data as input. The method allows one to capture dynamics of the data as trajectories of a ball from one basin to a different basin to yet another, constrained on the energy landscape specified by the estimated Ising model. While this energy landscape analysis has mostly been applied to functional magnetic resonance imaging (fMRI) data from the brain for historical reasons, there are emerging applications outside fMRI data and neuroscience. To inform such applications in various research fields, this review paper provides a detailed tutorial on each step of the analysis, terminologies, concepts underlying the method, and validation, as well as recent developments of extended and related methods.

OPEN ACCESS

Citation: Masuda N, Islam S, Aung S T, Watanabe T. (2025) Energy landscape analysis based on the Ising model: Tutorial review. *PLoS Complex Sys* 2(4): e0000039. <https://doi.org/10.1371/journal.pcsy.0000039>

Editor: Hocine Cherifi, Université de Bourgogne: Universite de Bourgogne, FRANCE

Published: May 9, 2025

Copyright: © 2025 Masuda et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: TW and NM acknowledge support from the Japan Science and Technology Agency (JST) Moonshot R&D (under grant no. JPMJMS2021). TW also acknowledges support from Grant-in-aid for Research Activity from Japan Society for Promotion of Sciences (23H04217), JST FOREST Program (24012854), Institute of AI and Beyond of UTokyo, and Cross-ministerial Strategic Innovation Promotion Program (SIP) on “Integrated Health Care System” (JPJ012425). NM also acknowledges support from National Institute of General Medical Sciences (under grant no. R01GM148973), the National Science

1. Introduction

Energy landscape is an established concept, catered with computational methods, in protein folding research [1–4]. With this dynamic view, we comprehend protein folding as an ensemble of locally stable protein conformations that are stochastically switched between depending on how stable each conformation is and how likely one conformation changes to another. According to the principles of statistical physics, a conformation with a low energy occurs with a high probability, and transitions between conformations occur with a high frequency if the energy barrier separating them is low (see [Figs 1 and 2](#) for schematics). The purpose of this review paper is to explain a family of data analysis methods for constructing such energy landscapes and dynamics on them for general multivariate time series data for which we do not know any underlying kinetic, chemical, or other laws.

The type of energy landscape analysis (ELA) that this article deals with, which we simply refer to as ELA in the following text, was initially proposed in [5,6] for functional magnetic resonance imaging (fMRI) data recorded from the human brain. It uses the Ising spin system model. The model fitting part of this ELA is shared with earlier studies, in particular in neuroscience [7–13]. Related to this fact, the ELA has predominantly been used in neuroscience, especially for human fMRI data [14–35]. Other applications in neuroscience include electroencephalogram (EEG) [36–38], magnetoencephalography (MEG) [39], and structural

Foundation (under grant no. 2204936), and JSPS KAKENHI (under grant nos. JP 23H03414, 24K14840, and 24K030130). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

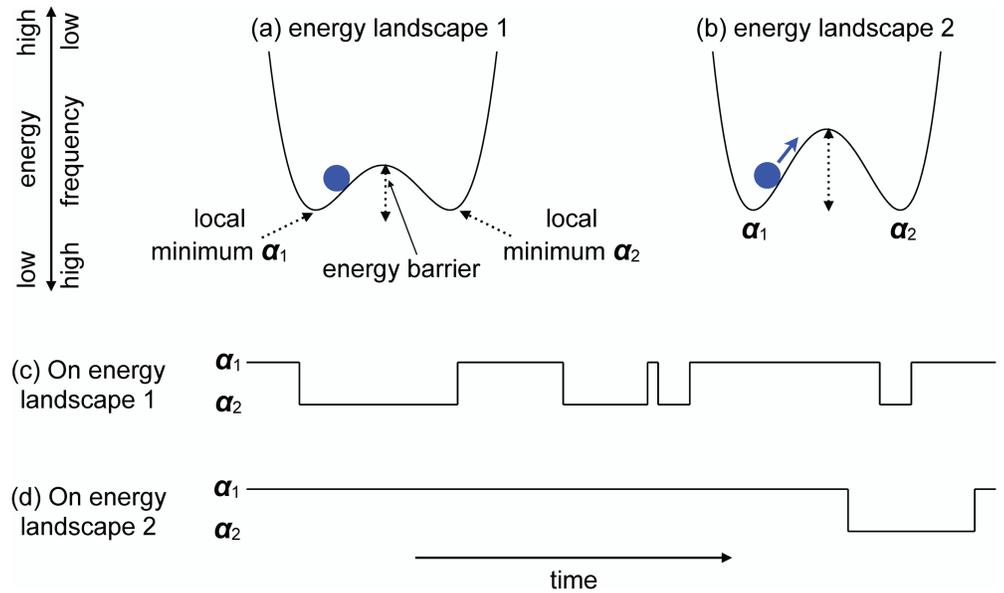


Fig 1. One-dimensional schematic of energy landscapes. (a) An energy landscape with two local minima of energy, α_1 and α_2 . They are separated by an energy barrier. The filled circle, or “ball,” represents the state of the system, which tends to go downhill but can also go uphill with a smaller rate. The position of the ball along the horizontal axis specifies the system’s state, which we call the activity pattern. (b) An energy landscape with a higher energy barrier between two local minima. (c) Sample time course of the system’s dynamics on the energy landscape shown in (a). (d) Sample time course on the energy landscape shown in (b). Note that switching between α_1 and α_2 is rarer in (d) than (c) because α_1 and α_2 are separated by a higher energy barrier in (b) than (a). In fact, the system’s state (i.e., filled circle) can take activity patterns different from the two local minima (e.g., a pattern close to but different from α_1). However, we only show stochastic transition dynamics between α_1 and α_2 in (c) and (d) for simplicity.

<https://doi.org/10.1371/journal.pcsy.0000039.g001>

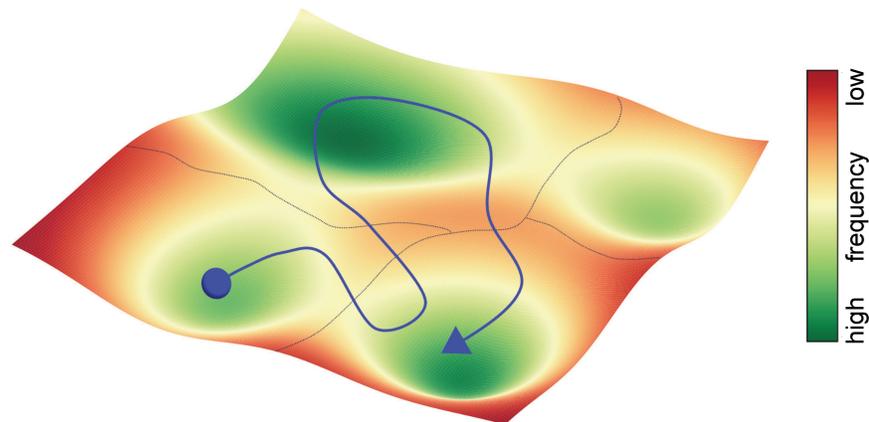


Fig 2. Two-dimensional schematic of an energy landscape. Four local minima are present. The dotted lines show the borders between the basin of different local minima. The “ball” represents the state of the system, and its position on the plane (ignoring the vertical coordinate) represents the activity pattern.

<https://doi.org/10.1371/journal.pcsy.0000039.g002>

brain network [40] data. However, the ELA does not impose any neuroscientific assumptions on the data and therefore applicable to any multivariate time series data as long as the data are long enough relative to the number of variables. In fact, the ELA has seen applications to

simulated data of cardiac fibrillation [41], time series of a combination of multimodal observables from patients in rheumatoid arthritis (such as the rheumatoid factor and swollen joint count) [42], Raman spectroscopy data [43], microbiome data [44–47], and mixture of microbiome data and physiological signals from calves [48]. ELA crucially takes a system's dynamics point of view in the sense that the method tracks the dynamics of the system's state, which is defined by the values of all variables at each time t . Such a system's state is abstracted into the filled ball in Figs 1(a), 1(b), and 2. As these figures show, ELA aims to capture the given multivariate time series data as the single abstracted system's state that transits among discrete states, similar to a Markov chain. Furthermore, the present ELA is based on the computation and statistical physics theory of the Ising spin system model. There are various other energy landscape analyses in addition to those in protein folding studies (e.g. [49–53]). In this review paper, we exclusively focus on the ELA based on the Ising model.

Our intention to write this tutorial-prone survey paper is to guide potential users of ELA in various research fields. At the same time, probably owing to open-source code (reviewed in Section 4), including some with a graphical interface, one can now easily carry out ELA without knowing details. This situation is beneficial for boosting applications of the method. However, there are recommended practices and pitfalls in ELA, which we will advocate in this paper. Furthermore, a detailed step-by-step tutorial for non-specialists is expected to help users understand intuitions, physics theory, and logic behind the scene. Therefore, although we published a shorter review in 2017 [54], here we provide a detailed tutorial, which includes discussion of recommended practices and pitfalls, advantages and disadvantages of ELA, and recent methodological developments, followed by some example applications.

2. Workflow of the energy landscape analysis

2.1. Overview

We outline a workflow of ELA in Fig 3. First, the input data are multivariate time series with N variables in discrete time (Fig 3(a)). We denote the time series by $\{\mathbf{x}(1), \dots, \mathbf{x}(t_{\max})\}$, where $\mathbf{x}(t)$ (with $t \in \{1, \dots, t_{\max}\}$) is the N -dimensional time series at time t , and t_{\max} is the number of observations, or the length of the time series. We denote the N -dimensional vector $\mathbf{x}(t)$ by

$$\mathbf{x}(t) = (x_1(t), \dots, x_N(t)), \quad (1)$$

where $x_i(t) \in \mathbb{R}$ is the signal of the i th variable at time t . We usually assume N of the order of 10; we discuss the choice of N in Section 3.3.

Second, we binarize each $x_i(t)$, $i \in \{1, \dots, N\}$, $t \in \{1, \dots, t_{\max}\}$. See Fig 3(b) for an example of binarized multivariate time series. We denote the obtained binarized multivariate time series by $\{\boldsymbol{\sigma}(1), \dots, \boldsymbol{\sigma}(t_{\max})\}$, where the system's binarized state, which we call the activity pattern, at time t is written as

$$\boldsymbol{\sigma}(t) = (\sigma_1(t), \dots, \sigma_N(t)), \quad (2)$$

where $\sigma_i(t) \in \{-1, 1\}$. Another common convention is $\sigma_i(t) \in \{0, 1\}$. In either convention, 1 corresponds to high activity of the i th variable at time t ; the other value (i.e., -1 or 0 depending on the coding scheme) corresponds to low activity. Because the linear transformation $f(x) = (x + 1)/2$ maps $\{-1, 1\}$ to $\{0, 1\}$, i.e., $f(-1) = 0$ and $f(1) = 1$, the choice of either convention does not essentially matter for the following analyses. See [54] for more discussion of the two conventions.

Third, we estimate the Ising model, i.e., infer the model parameters from the binarized multivariate time series data $\{\boldsymbol{\sigma}(1), \dots, \boldsymbol{\sigma}(t_{\max})\}$ (Fig 3(c) and 3(d)). The Ising model is also

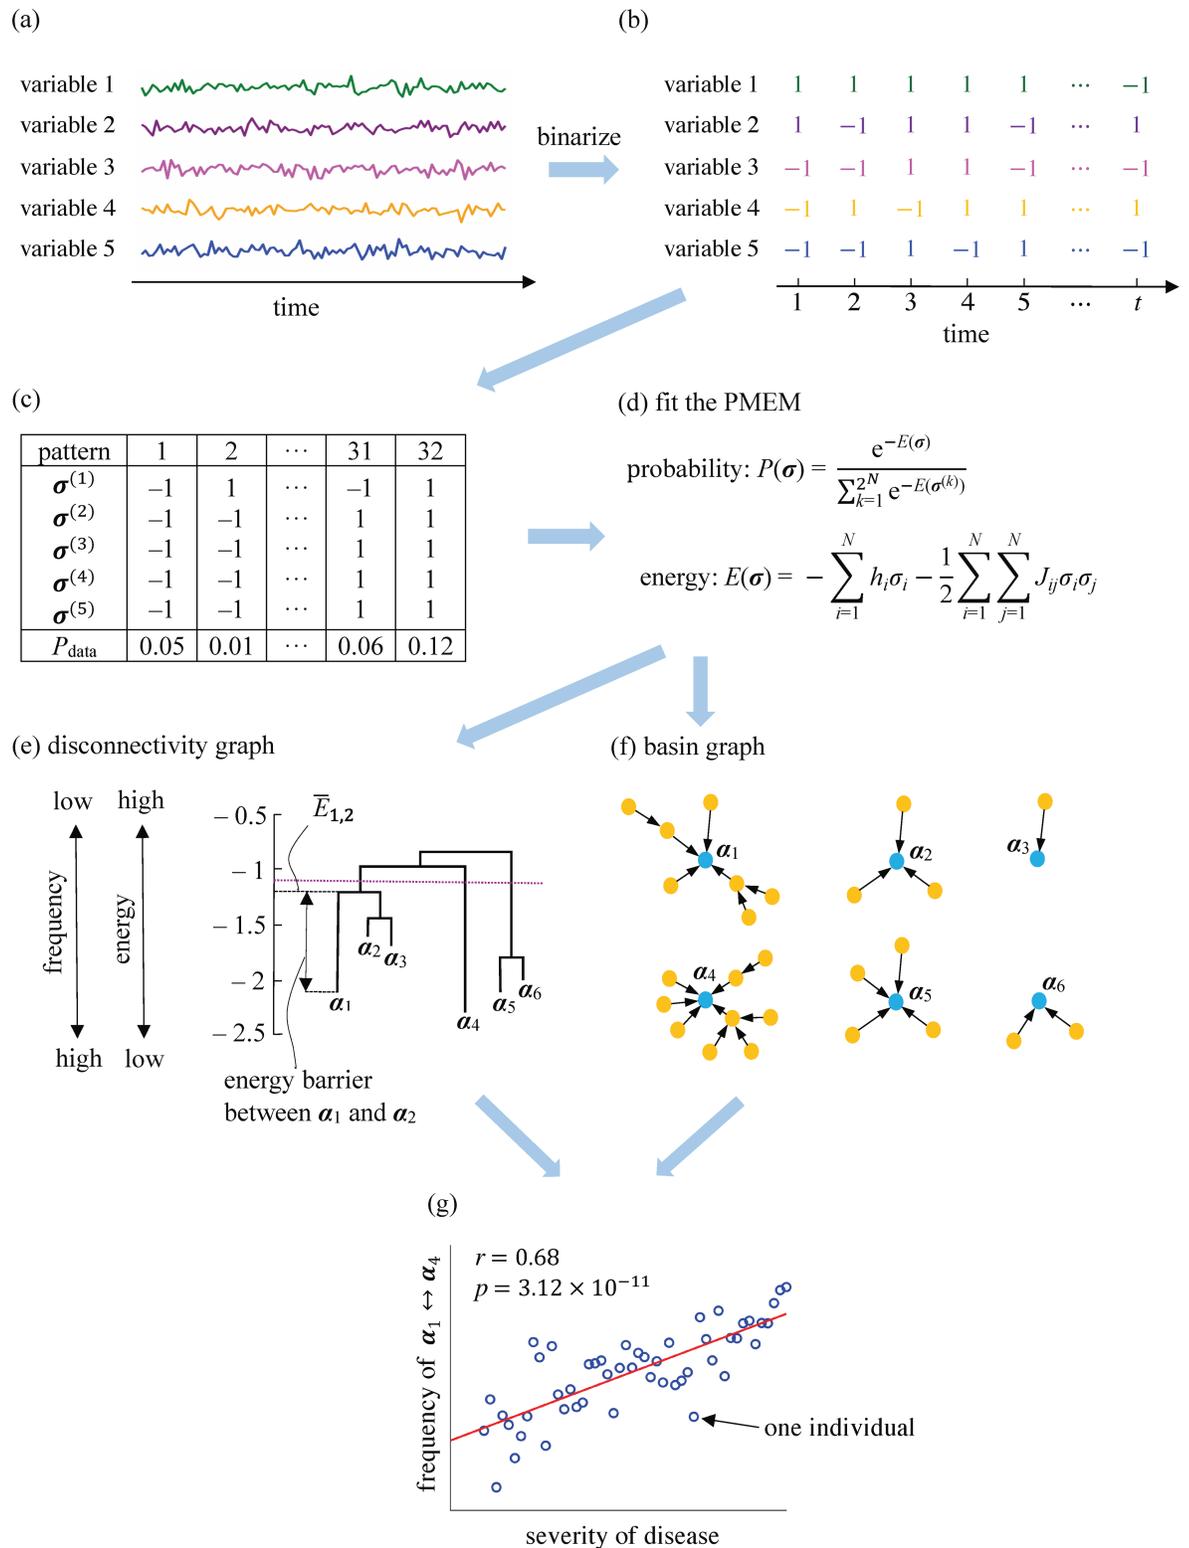


Fig 3. Steps of the ELA. (a) The input data is multivariate time series with N variables. (b) Binarized time series. (c) Fraction of times, $P(\sigma)$, with which each of the 2^N activity patterns, σ , occurs in the binarized data shown in (b). (d) Fitting of the PMEM to the distribution of σ shown in (c). (e) Disconnectivity graph derived from the estimated PMEM. It shows relationships between the activity patterns that are local minima. Each $\alpha_i, i \in \{1, \dots, 6\}$, represents a local minimum. The double-headed arrow indicates the height of the energy barrier between local minima α_1 and α_2 when the activity pattern tries to travel from α_1 to α_2 . (f) Basin graph derived from the estimated PMEM. (g) Downstream analysis, which may use a feature computed from the estimated disconnectivity graph or basin graph.

<https://doi.org/10.1371/journal.pcsy.0000039.g003>

referred to the pairwise maximum entropy model (PMEM), which we use throughout this paper, and the Boltzmann machine. There are various inference algorithms. The estimated Ising model is typically a maximum likelihood estimator for the data under some constraints. In fact, the procedures up to this point, i.e., fitting the Ising model to multivariate time series and other multivariate data, is not a new idea. See [11,12,55–58] for reviews.

The following steps are unique to the ELA. Fourth, we estimate the so-called disconnectivity graph [59], of which a hypothetical example is shown in Fig 3(e). The bottom of each valley schematically shown in Figs 1 and 2 corresponds to the tip of each branch hanging down vertically in the disconnectivity graph (see Fig 3(e)). It is called a local minimum of energy (local minimum for short). The disconnectivity graph represents all the local minima and their relationships in terms of the ease of transition for the estimated PMEM and is a primary product of the ELA. Another main product of the ELA is the basin graph, shown in Fig 3(f). The basin graph encodes the width of the valley of each local minimum schematically shown in Fig 1 or its area schematically shown in Fig 2. It also contains other information.

Fifth, because the disconnectivity graph and basin graph are graphical representations of the energy landscape, we quantify them into numbers or vectors, which we collectively call features of the energy landscape. These features provide insights into the given multivariate time series data or a system generating the data.

Sixth, one often wants to compare the energy landscapes obtained from a group of people or an individual against the energy landscapes obtained from a different group or individual. Such a comparison is made in terms of features of the energy landscape. For example, we examine whether or not larger values of a particular feature from the energy landscape estimated for different human individuals are associated with the severity of a disease (see Fig 3(g)). In the following sections, we explain each step in detail.

2.2. Fitting the PMEM

We start by binarizing the original time series (see Fig 3(a) and (b)). For each $i \in \{1, \dots, N\}$ and $t \in \{1, \dots, t_{\max}\}$, we set $\sigma_i(t) = -1$ if $x_i(t) < \theta_i$ and $\sigma_i(t) = 1$ if $x_i(t) \geq \theta_i$, where θ_i is an arbitrary threshold. A common choice of θ_i is the 50 percentile of $\{x_i(1), \dots, x_i(t_{\max})\}$ such that $\sigma_i(t) = -1$ and $\sigma_i(t) = 1$ with probability 0.5 each. If the distribution of $x_i(t)$ is not extremely skewed, one can also set θ_i to the average of the time series, i.e., $\sum_{t=1}^{t_{\max}} x_i(t)/t_{\max}$, to obtain roughly 50% of $\sigma_i(t) = -1$ and $\sigma_i(t) = 1$ each.

Because $\sigma(t)$ is an N -dimensional binary vector, there are 2^N possible activity patterns in total, which we enumerate as $\sigma^{(1)}, \dots, \sigma^{(2^N)}$. This is depicted in Fig 3(c). For example, one obtains $\sigma^{(1)} = (-1, -1, \dots, -1)$, $\sigma^{(2)} = (1, -1, \dots, -1)$, and $\sigma^{(2^N)} = (1, 1, \dots, 1)$. We then count the fraction of each activity pattern σ , denoted by $P_{\text{data}}(\sigma)$, in the given data. Note that σ is any of $\sigma^{(1)}, \dots, \sigma^{(2^N)}$ and that $P_{\text{data}}(\sigma)$ is simply the number of times that σ appears in the data that is divided by t_{\max} . Because $P_{\text{data}}(\sigma)$ is a fraction, we obtain

$$\sum_{k=1}^{2^N} P_{\text{data}}(\sigma^{(k)}) = 1. \tag{3}$$

The PMEM aims to fit a particular form of the distribution, denoted by $P_{\text{model}}(\sigma)$ and explained in the following text, to the empirical distribution, $P_{\text{data}}(\sigma)$. With this methodological choice, we are drastically ignoring the temporal order of the given data. For example, the estimated PMEM will be the same as that for the original time series even if we re-arrange the original multivariate time series in the reversed order (i.e., $\{\sigma(t_{\max}), \sigma(t_{\max} - 1), \dots, \sigma(1)\}$) or shuffle the time stamp of $\{\sigma(t)\}$ randomly before the analysis. Therefore, the ELA is in a

sense not a method for analyzing time series. It is a method to describe the distribution of the given data. However, ELA still allows us to infer and interpret the *dynamics* of the data; e.g., the system's state (i.e., the ball in Figs 1 and 2) stays near a local minimum α_1 for some time, then transits to a different local minimum α_2 , and comes back to α_1 . This is because the PMEM implies a particular equilibrium stochastic dynamics. In other words, once a PMEM is estimated, it dictates the probability or rate of the move from any activity pattern σ to another, which we explain in Section 3.2.3.

The mean activity of the i th variable for the data is given by

$$\langle \sigma_i \rangle_{\text{data}} \equiv \frac{1}{t_{\max}} \sum_{t=1}^{t_{\max}} \sigma_i(t). \tag{4}$$

The mean pairwise joint activity for the i th and j th variables for the data is given by

$$\langle \sigma_i \sigma_j \rangle_{\text{data}} \equiv \frac{1}{t_{\max}} \sum_{t=1}^{t_{\max}} \sigma_i(t) \sigma_j(t). \tag{5}$$

The PMEM maximizes the entropy of the distribution of activity patterns, $P(\sigma)$, under the condition that $\langle \sigma_i \rangle$ and $\langle \sigma_i \sigma_j \rangle$ (with $1 \leq i < j \leq N$) are the same between the estimated Ising model and the given data for all $i, j \in \{1, \dots, N\}$. As a result of entropy maximization, one obtains (see Fig 3(d))

$$P(\sigma) = \frac{e^{-E(\sigma)}}{\sum_{k=1}^{2^N} e^{-E(\sigma^{(k)})}}, \tag{6}$$

where $E(\sigma)$ represents the energy of activity pattern σ and is given by

$$E(\sigma) = - \sum_{i=1}^N h_i \sigma_i - \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N J_{ij} \sigma_i \sigma_j. \tag{7}$$

Equation (6) is a Boltzmann distribution (also called Gibbs distribution) in statistical mechanics without the temperature parameter, or equivalently, the temperature normalized to 1. Equation (6) implies that a high probability of the activity pattern $P(\sigma)$ corresponds to a low energy $E(\sigma)$ and vice versa. It should be noted that the energy here is a purely mathematical construct and does not imply any physical energy. In Eq (7), the fitting parameter h_i represents the tendency that the i th variable is active (i.e., $\sigma_i = 1$), and J_{ij} quantifies the sign and strength of the pairwise interaction between the i th and j th variables. The PMEM assumes that $J_{ij} = J_{ji}$. Furthermore, the $J_{ij} \sigma_i \sigma_j$ when $i = j$ is reduced to a constant $J_{ii} \sigma_i^2 = J_{ii}$ because σ_i^2 with $\sigma_i \in \{-1, 1\}$ is always equal to 1. Because $P(\sigma)$ in Eq (6) does not change by the addition of a constant term to $E(\sigma)$, one can simplify Eq (7) to

$$E(\sigma) = - \sum_{i=1}^N h_i \sigma_i - \sum_{i=1}^N \sum_{j=1}^{i-1} J_{ij} \sigma_i \sigma_j. \tag{8}$$

We denote the expected single-variable activity and pairwise activity from the PMEM by $\langle \sigma_i \rangle_{\text{model}}$ and $\langle \sigma_i \sigma_j \rangle_{\text{model}}$, respectively. They are given by

$$\langle \sigma_i \rangle_{\text{model}} = \sum_{k=1}^{2^N} \sigma_i^{(k)} P(\sigma^{(k)}) \tag{9}$$

and

$$\langle \sigma_i \sigma_j \rangle_{\text{model}} = \sum_{k=1}^{2^N} \sigma_i^{(k)} \sigma_j^{(k)} P(\boldsymbol{\sigma}^{(k)}), \tag{10}$$

where $\sigma_i^{(k)} \in \{-1, 1\}$ is the binarized activity of the i th variable in the k th activity pattern $\boldsymbol{\sigma}^{(k)}$.

A popular and conceptually easy method to infer h_i and J_{ij} is to iteratively adjusting them using a gradient ascent algorithm; it is an ascent algorithm because it gradually increases the likelihood of the data, $\mathcal{L} = \prod_{t=1}^{t_{\max}} P(\boldsymbol{\sigma}(t))$ along the ascent direction. The iteration scheme is given by

$$h_i^{\text{new}} = h_i^{\text{old}} + \frac{\epsilon}{t_{\max}} \frac{\partial \mathcal{L}}{\partial h_i} = h_i^{\text{old}} + \epsilon (\langle \sigma_i \rangle_{\text{data}} - \langle \sigma_i \rangle_{\text{model}}) \tag{11}$$

and

$$J_{ij}^{\text{new}} = J_{ij}^{\text{old}} + \frac{\epsilon}{t_{\max}} \frac{\partial \mathcal{L}}{\partial J_{ij}} = J_{ij}^{\text{old}} + \epsilon (\langle \sigma_i \sigma_j \rangle_{\text{data}} - \langle \sigma_i \sigma_j \rangle_{\text{model}}), \tag{12}$$

where superscripts new and old represent the values after and before a single updating step, respectively, and ϵ is the learning rate. The factor $1/t_{\max}$ on the right-hand sides of Eqs (11) and (12) is for normalization. A slightly different iteration scheme called the iterative scaling algorithm [60], in which the right-hand side of Eqs (11) and (12) is replaced by $\epsilon \operatorname{sgn}(\langle \sigma_i \rangle) \log(\langle \sigma_i \rangle_{\text{data}} / \langle \sigma_i \rangle_{\text{model}})$ and $\epsilon \operatorname{sgn}(\langle \sigma_i \sigma_j \rangle) \log(\langle \sigma_i \sigma_j \rangle_{\text{data}} / \langle \sigma_i \sigma_j \rangle_{\text{model}})$, respectively, is also common [7,11,13]. With either iterative algorithm, the convergence (i.e., $h_i^{\text{new}} = h_i^{\text{old}}$ and $J_{ij}^{\text{new}} = J_{ij}^{\text{old}}$ for all i,j) implies that $\langle \sigma_i \rangle_{\text{model}} = \langle \sigma_i \rangle_{\text{data}}$ and $\langle \sigma_i \sigma_j \rangle_{\text{model}} = \langle \sigma_i \sigma_j \rangle_{\text{data}}$ for all i,j .

One can verify that Eq (6) is concave in terms of each h_i and J_{ij} by checking that the Hessian of $\log \mathcal{L}$ is a type of sign-flipped covariance matrix, which is negative semi-definite. The concavity guarantees that the gradient ascent scheme always yields the unique maximum likelihood estimator. However, a single updating step requires reference to $P(\boldsymbol{\sigma})$ for all the 2^N activity patterns for calculating $\langle \sigma_i \rangle_{\text{model}}$ and $\langle \sigma_i \sigma_j \rangle_{\text{model}}$. Because $P(\boldsymbol{\sigma})$ changes in every iteration step, the gradient ascent method is computationally costly for large N . Our experience is that, without dedicated hardware, the gradient ascent method would be slow even with $N = 15$ or 20 . To circumvent this problem, there are various approximate algorithms (e.g., pseudo-likelihood maximization) for estimating the PMEM [54–56,61,62].

2.3. Accuracy of fit

If the PMEM poorly fits to the data, ELA is probably not a good choice. Therefore, while many papers skip measuring an accuracy of fit, it is important to measure it.

A least stringent test, which is often used, is to examine the correlation between $\langle \sigma_i \rangle_{\text{model}}$ and $\langle \sigma_i \rangle_{\text{data}}$ across the i values and that between $\langle \sigma_i \sigma_j \rangle_{\text{model}}$ and $\langle \sigma_i \sigma_j \rangle_{\text{data}}$ across (i,j) pairs. A good fit of the PMEM to the data should yield $\langle \sigma_i \rangle_{\text{model}} \approx \langle \sigma_i \rangle_{\text{data}}$ and $\langle \sigma_i \sigma_j \rangle_{\text{model}} \approx \langle \sigma_i \sigma_j \rangle_{\text{data}}$, where \approx represents “approximately equal to”. However, these relationships are not enough for the following reason. If we run a gradient ascent algorithm to fit the PMEM (Section 2.2), it will converge due to the concavity of the problem for any data while convergence may take a long time if N is relatively large. Once the algorithm sufficiently converges, one obtains $\langle \sigma_i \rangle_{\text{model}} \approx \langle \sigma_i \rangle_{\text{data}}$ and $\langle \sigma_i \sigma_j \rangle_{\text{model}} \approx \langle \sigma_i \sigma_j \rangle_{\text{data}}$ even if the data are a poor fit to the PMEM.

Therefore, one often plots $P(\boldsymbol{\sigma})$ against $P_{\text{data}}(\boldsymbol{\sigma})$ over all the 2^N activity patterns $\boldsymbol{\sigma}$. If $P(\boldsymbol{\sigma})$ and $P_{\text{data}}(\boldsymbol{\sigma})$ are highly correlated, the accuracy of the fit is considered to be high [5–10,12,13,17].

Another commonly used accuracy measure [8,11,54] is

$$r_D = \frac{D_1 - D_2}{D_1}, \tag{13}$$

where

$$D_\ell = \sum_{k=1}^{2^N} P_{\text{data}}(\boldsymbol{\sigma}^{(k)}) \log_2 \frac{P_{\text{data}}(\boldsymbol{\sigma}^{(k)})}{P_\ell(\boldsymbol{\sigma}^{(k)})} \tag{14}$$

is the Kullback-Leibler divergence between the probability distribution of the activity pattern in the ℓ th-order ($\ell = 1, 2$) maximum entropy model, $P_\ell(\boldsymbol{\sigma})$, and the distribution of the activity pattern in the original data, $P_{\text{data}}(\boldsymbol{\sigma})$. Distribution $P_2(\boldsymbol{\sigma})$ is equivalent to $P(\boldsymbol{\sigma})$ given by Eqs (6) and (7). It is called second-order because the highest-order function in terms of σ_i appearing in the model is of second order (i.e., $J_{ij}\sigma_i\sigma_j$). The first-order, or independent, maximum entropy model (i.e., $\ell = 1$) is given by Eq (6) but without interaction terms. In other words, we force $J_{ij} = 0 \forall i, j$ in Eq (7) and set each h_i so that $\langle \sigma_i \rangle_{\text{model}} = \langle \sigma_i \rangle_{\text{data}}$ is satisfied. Quantity r_D ranges between 0 and 1 for the maximum likelihood estimator of the PMEM. We obtain $r_D = 1$ when the PMEM perfectly fits the distribution of the activity pattern for the given data. In contrast, $r_D = 0$ when the PMEM does not fit the data any better than the first-order model. For the maximum likelihood estimator, r_D is equal to another often employed accuracy measure given by $(S_1 - S_2)/(S_1 - S_{\text{data}})$, where $S_\ell = -\sum_{k=1}^{2^N} P_\ell(\boldsymbol{\sigma}^{(k)}) \log P_\ell(\boldsymbol{\sigma}^{(k)})$ is the Shannon entropy of distribution P_ℓ [9,11].

The accuracy of fit would be small when N is large or t_{max} is small. In fact, the ELA is a data-hungry method. Therefore, it is important to check and report the value of r_D or a different accuracy of fit measure. We recommend the value of r_D above 0.85, or ideally 0.9. This is because, although these guideline values do not have any theoretical ground, an influential early study of the PMEM applied to neural spiking data realized $r_D \approx 0.9$ [7].

2.4. Local minima of energy and the network of activity patterns

Once we have estimated the PMEM, the next task is to construct and draw the so-called disconnectivity graph. To understand the meaning and construction of the disconnectivity graph, we need to introduce two concepts: local minimum of energy and network of activity patterns.

Intuitively, a local minimum of energy is the bottom of a basin of the energy landscape shown in Figs 1 and 2. However, the 2^N activity patterns do not lie in a one- or two-dimensional space as depicted in these figures. To precisely define the local minimum, let us consider Fig 4, in which there are $N = 5$ variables and an activity pattern $\boldsymbol{\sigma}^{(19)} = (-1, 1, -1, -1, 1)$ is neighbored by five other activity patterns. Each of these neighboring activity patterns is different from $\boldsymbol{\sigma}^{(19)}$ in the sign of just one σ_i . In other words, these five activity patterns are in Hamming distance 1 from $\boldsymbol{\sigma}^{(19)}$. There are five activity patterns neighboring $\boldsymbol{\sigma}^{(19)}$ because one can flip any one entry of $\boldsymbol{\sigma}^{(19)}$ to obtain a neighboring activity pattern. Fig 4 shows that $\boldsymbol{\sigma}^{(19)}$ has a lower energy than all the five neighbors. Therefore, $\boldsymbol{\sigma}^{(19)}$ appears more frequently than each of its neighboring activity patterns in the data. When this is the case, we say that $\boldsymbol{\sigma}^{(19)}$ is a local minimum. If we allow the activity pattern to vary over time, the system's state is more likely to transit from a neighbor of $\boldsymbol{\sigma}^{(19)}$, say $\boldsymbol{\sigma}^{(17)}$, to $\boldsymbol{\sigma}^{(19)}$ than vice versa because $\boldsymbol{\sigma}^{(19)}$ is more frequent than $\boldsymbol{\sigma}^{(17)}$. In this sense, $\boldsymbol{\sigma}^{(19)}$ is locally attractive.

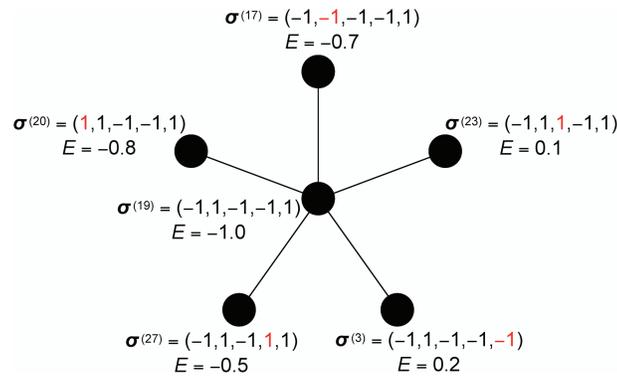


Fig 4. A schematic of the local minimum of energy. A system of $N = 5$ variables is assumed. Activity pattern $\sigma^{(19)}$ is a local minimum because it has a lower energy than all its neighbors. The neighbors are the activity patterns with Hamming distance 1 from $\sigma^{(19)}$. The numbering scheme for the activity patterns is consistent with the one used in Fig 3(c).

<https://doi.org/10.1371/journal.pcsy.0000039.g004>

However, such dynamics are by definition stochastic. Therefore, $\sigma^{(19)}$ being a local minimum implies that the system’s state arriving at $\sigma^{(19)}$ will still escape it albeit at a lower rate. In empirical data, we commonly see transitions out of local minima.

In fact, any activity pattern has N neighbors no matter whether or not it is a local minimum. The 2^N activity patterns form a “network of activity patterns” shown in Fig 5. This network is called the hypercube. One can verify that each node of this network, corresponding to an activity pattern, has exactly N neighbors. Using the terminology of network science (or graph theory), there are 2^N nodes, and the degree of each node (i.e., the number of edges the node has) is equal to N . Furthermore, in Fig 5, each node has ten nodes in distance (i.e., minimum number of hops) 2, ten nodes in distance 3, five nodes in distance 4, and one node in distance 5. The single node in distance 5 is the one in which all the N values of σ_i are flipped. For example, the node in distance 5 from $\sigma^{(19)} = (-1, 1, -1, -1, 1)$ is $\sigma^{(14)} = (1, -1, 1, 1, -1)$. In general, there are $\binom{N}{d}$ nodes in distance d from any node, where $\binom{N}{d}$ represents the binomial coefficient. The network of activity patterns represents admissible routes of the system’s state, $\sigma(t)$, when $\sigma(t)$ is allowed to change in one entry, i , at one time.

In empirical data, $\sigma(t)$ and $\sigma(t + 1)$ are often different in more than one i unless the time resolution of the observation is high. For example, in Fig 3(b), $\sigma(1) = (1, 1, -1, -1, -1)$ and $\sigma(2) = (1, -1, -1, 1, -1)$ are different in two variables, i.e., $i = 2$ and $i = 4$. If $\sigma(t)$ and $\sigma(t + 1)$ are indeed different in more than one i for some t , the trajectory is not a discrete-time walk on the hypercube because the walker can move from a node to its non-neighboring node in one time step. However, even in this situation, the assumption that only a single entry of $\sigma(t)$ can flip at a time is still plausible by considering that we are observing a continuous-time walk in discrete time. This is the case if synchronous threshold crossing, which would simultaneously flip multiple $\sigma_i(t)$, is absent for the resolution of discrete times $t = 1, 2, \dots$ at which one observes the continuous-time walk.

To discuss the interpretation of $\sigma(t)$ as a discrete-time observation of a continuous-time walk, we postulate an underlying N -dimensional dynamics in continuous time, $\mathbf{x}(t')$, $t' \in \mathbb{R}$, and consider that we observe it in discrete times $t \in \{1, 2, \dots\}$, obtaining $\sigma(t)$. Discrete time t does not have to be the same as t' . In Fig 6, the t shown in blue is identical with t' . In contrast, the t shown in red is sampled three times less frequently than the t shown in blue. Recall that $\sigma_i(t)$ flips if and only if $x_i(t)$ crosses the threshold θ_i , no matter whether the time is

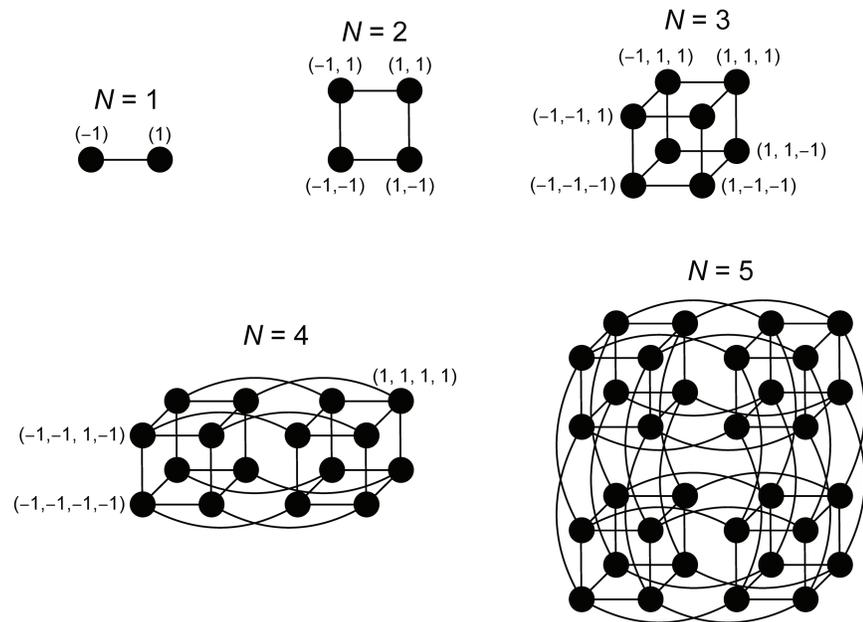


Fig 5. Networks of activity patterns for $N \in \{1, \dots, 5\}$. They are hypercubes having 2^N nodes. Each node represents an activity pattern and has N neighbors.

<https://doi.org/10.1371/journal.pcsy.0000039.g005>

continuous or discrete. If different $x_i(t)$'s simultaneously cross their thresholds, then multiple $\sigma_i(t)$'s flip at the same time. In this case, $\{\sigma(1), \sigma(2), \dots\}$ is not a nearest-neighbor walk on the hypercube. On the contrary, if such a synchronous threshold crossing is absent and the time resolution for sampling $\sigma(t)$ is high enough, then $\sigma(t)$ flips at most one bit at a time. In this case, $\{\sigma(1), \sigma(2), \dots\}$ is a nearest-neighbor walk. In other words, the walker stays at one node of the hypercube for some discrete times, then moves to one of its neighbors, and so on, producing a discrete-time walk on the hypercube, without jumps from a node to a non-neighboring node. For example, the walk shown in blue in Fig 6(a) is a nearest-neighbor walk. The figure indicates that $x_j(t')$ changes from positive to negative first, followed by the same change in $x_i(t')$. Then, $x_k(t')$ changes from negative to positive. Then, $x_i(t')$, $x_k(t')$, and $x_j(t')$ turn from positive to negative in this order. The discrete-time walk $\sigma(t)$ shown in blue captures all these events including their chronology. In contrast, the walk shown in red in Fig 6(a) is not a nearest-neighbor walk because the time resolution for observing $\sigma(t)$ is too low. For example, both $\sigma_1(t)$ and $\sigma_2(t)$ change between $t = 1$ and $t = 2$.

In the PMEM fitted to empirical data, we often find moves from a node on the hypercube to its non-nearest-neighbor, as shown in Fig 3(b) and as $\sigma(t)$ shown in red in Fig 6(a). We justify this situation by regarding such $\sigma(t)$ as an observation of originally continuously varying $\mathbf{x}(t')$ at a relatively low time resolution, if there is no approximately simultaneous threshold crossing event.

Fig 6(b) shows a case with approximately simultaneous threshold crossing events in the time series $\{\mathbf{x}(t'); t' \in \mathbb{R}\}$, marked by the green and yellow circles. In this case, even if the time resolution for observing $\sigma(t)$ is high, $\{\sigma(1), \sigma(2), \dots\}$ is not a nearest-neighbor walk. For example, the walker shown in blue in Fig 6(b) moves to a non-neighboring node between $t = 2$ and $t = 3$ and between $t = 7$ and $t = 8$, which are two intervals in which approximately simultaneous threshold crossing occurs. In fact, a move to a non-neighboring node also occurs

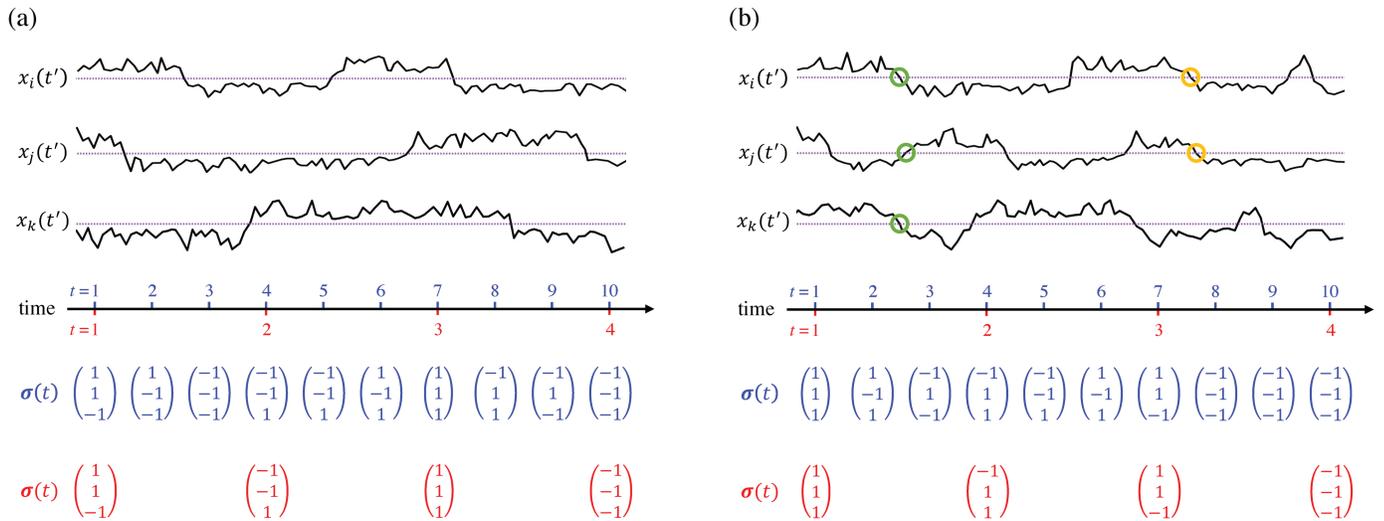


Fig 6. Time series of activity pattern as a discrete-time observation of a continuous-time walk. (a) Absence and (b) presence of synchronous threshold crossing. The horizontal dotted lines represent the threshold for binarization, θ_i . We show the case $t = t'$ and $t = 3t' - 2$ in blue and red, respectively. In (b), the green circles show an approximately simultaneous threshold crossing of $x_i(t')$, $x_j(t')$, and $x_k(t')$. The yellow circles show an approximately simultaneous threshold crossing of $x_i(t')$ and $x_j(t')$.

<https://doi.org/10.1371/journal.pcsy.0000039.g006>

between $t = 6$ and $t = 7$, which contains a less stringent simultaneous threshold crossing of $x_j(t')$ and $x_k(t')$. We also note that the same $\sigma(t)$, shown in blue, fails to capture the two threshold crossings of $x_k(t')$ between $t' = t = 8$ and $t' = t = 9$ and two threshold crossings of $x_i(t')$ between $t' = t = 9$ and $t' = t = 10$. In this situation, it is difficult to justify the use of the PMEM and hence ELA.

Approximately simultaneous threshold crossing is a relative concept. If we observe $\mathbf{x}(t')$, $t' \in \mathbb{R}$, at large regular intervals to produce $\sigma(t)$, then $x_i(t')$ with multiple i 's would cross their thresholds, perhaps multiple times for some i 's, between adjacent discrete times (i.e., between t and $t + 1$). This is a case of simultaneous threshold crossing. In the extreme case, if the observation occurs with excessively long time intervals, then $\mathbf{x}(t)$ and $\mathbf{x}(t + 1)$ would be completely uncorrelated, and so are $\sigma(t)$ and $\sigma(t + 1)$. Then, the trajectory $\{\sigma(1), \sigma(2), \dots\}$ will be composed of independent samples drawn from the distribution $P(\sigma)$, completely ignoring the hypercube structure. Given this discussion, we could try to justify the use of the PMEM and ELA even in the case of $\sigma(t)$ shown in blue and red in Fig 6(b). If we use a higher sampling rate, such as $t = 100t'$, such that $t \in \{1, 2, \dots\}$ corresponds to $t' \in \{0.01, 0.02, \dots\}$, then $\{\sigma(1), \sigma(2), \dots\}$ will be a nearest-neighbor walk unless there are perfectly simultaneous threshold crossing events in $\mathbf{x}(t')$, $t' \in \mathbb{R}$. However, excessively synchronous threshold crossing of $x_i(t')$'s would violate the detailed balance condition underlying the random walk assumption (Section 3.2.3). Therefore, we do not recommend that users rely on random-walk views when strongly synchronous dynamics among the variables is known or expected.

The network of activity patterns should not be confused with networks that variables of the multivariate time series form. One can form a network of the N variables by, for example, naively laying an edge (i, j) if and only if the Pearson correlation coefficient between the two time series $\{x_i(1), \dots, x_i(t_{\max})\}$ and $\{x_j(1), \dots, x_j(t_{\max})\}$ is sufficiently large [63]. In such a correlation network, a node is a variable of the multivariate time series, and there are N nodes. There may also be a structural network generating $\{\mathbf{x}(t)\}$ (e.g., a neuronal network with neurons as nodes, or a social network with human individuals as nodes). In contrast,

in the network of activity patterns, a node is an activity pattern, there are 2^N nodes, and the network structure is always the hypercube.

2.5. Disconnectivity graph

The disconnectivity graph is a first-hand and main summary of the energy landscape. The disconnectivity graph is a tool from chemical physics and represents the relationship between different local minima of energy [59,65]. For expository purposes, we show a hypothetical disconnectivity graph in Fig 3(e); see Fig 7(a) for the disconnectivity graph for a real fMRI data set from the human brain.

There may be multiple local minima in the network of activity patterns. In the hypothetical network of activity patterns for $N = 5$ variables shown in Fig 8(a), we indicate the local minima by red circles. To construct the disconnectivity graph, we need to determine the energy barrier separating each pair of local minima. This task would be easy if the energy landscape were one-dimensional, as shown in Fig 1(a) and 1(b) because there is only one shortest (i.e., non-redundant) path connecting the two local minima. In fact, this task is more involved because the actual energy landscape of the PMEM is defined on the hypercube. Let us consider possible paths between each pair of local minima. Partly because such a path does not have to be a shortest path (whose length is equal to the Hamming distance between the two local minima), there are many paths unless N is small. The non-black lines in Fig 8(a) show three among many paths connecting local minima α_1 and α_2 . Because the Hamming distance

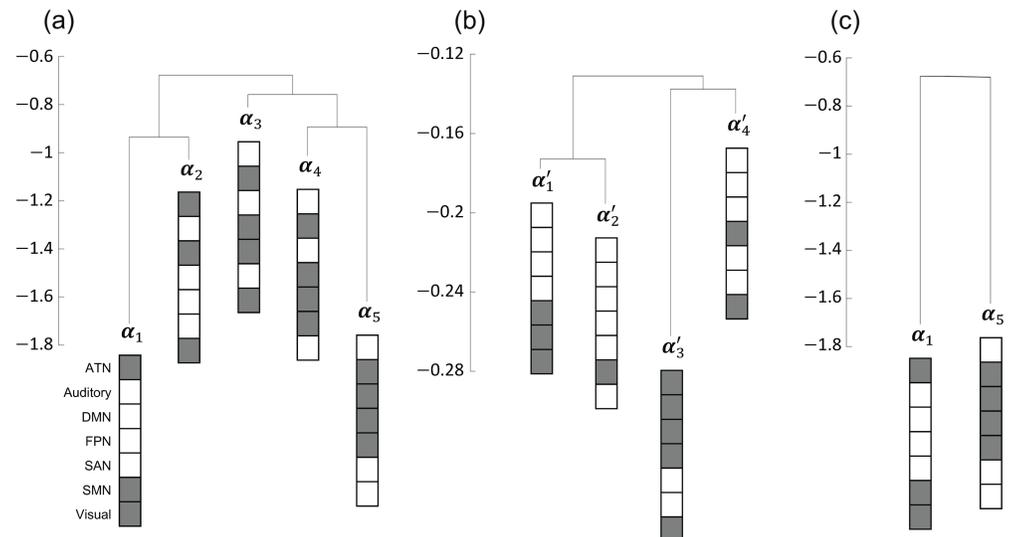


Fig 7. Disconnectivity graphs. (a) Disconnectivity graph from fMRI data obtained from the brain of healthy human participants. We used the same data as the one used for Fig A1(a) in [32]. It is a concatenation of the first four fMRI recording sessions (resulting in $t_{\max} = 818 \times 4 = 3,272$) from a single participant in the Midnight Scan Club data set [64]. We used the first participant (i.e., MSC01). We used the so-called whole-brain network with $N = 7$ nodes; the N variables are the attention network (ATN), auditory network, default mode network (DMN), fronto-parietal network (FPN), salience network (SAN), somatosensory and motor network (SMN), and visual network [14]. See [32] for more details and the procedures for preprocessing the fMRI data. The activity pattern at each local minimum is indicated below the local minimum; a gray square represents $\sigma_i = -1$, and a white square represents $\sigma_i = 1$. The accuracy of fit $r_D = 0.915$. (b) Disconnectivity graph for randomly generated data with the same N and t_{\max} as those used in (a). (c) Disconnectivity graph obtained via pruning of (a). Local minima α_1 and α_5 in the disconnectivity graph shown in (a) have survived the significance test, whereas the other three local minima have not.

<https://doi.org/10.1371/journal.pcsy.0000039.g007>

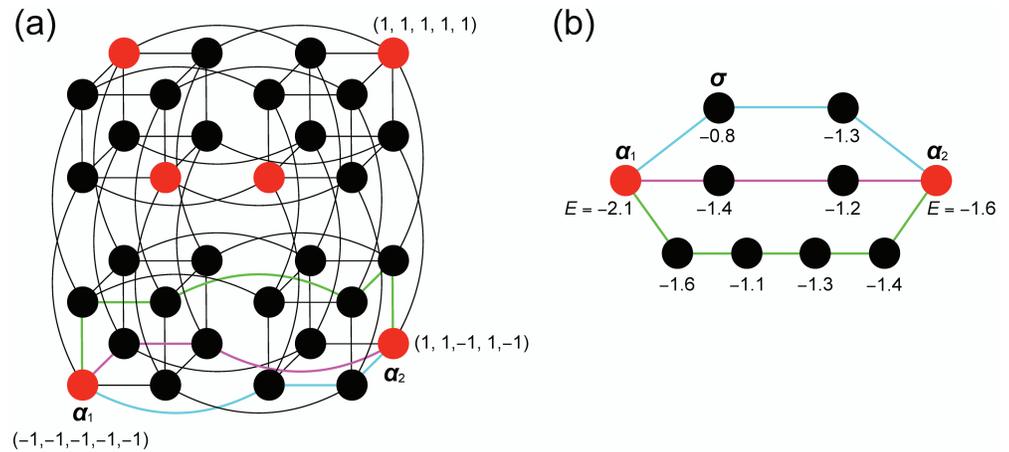


Fig 8. Local minima and paths between them. (a) The network of activity patterns in the case of $N = 5$ variables. The red circles represent local minima. (b) Three paths between local minima α_1 and α_2 in the network of activity patterns shown in (a).

<https://doi.org/10.1371/journal.pcsy.0000039.g008>

between α_1 and α_2 is three, any path has at least length 3. For example, the path shown in green has length 5.

Then, to determine the ease with which the system's state can travel from α_1 to α_2 and vice versa, let us tabulate all the paths between them and the energy value of each node on the path. One can omit walks visiting the same node more than once because such redundant walks do not make the travels between α_1 and α_2 easier than the corresponding shortest path. Fig 8(b) shows the three paths shown in color in Fig 8(a) together with the energy on the nodes. The energy value of each node on the path dictates how much height of the hill the system's state needs to climb to go from one end to the other. This is because, by definition, a high-energy activity pattern is a rare activity pattern in the data or the estimated PMEM. For example, in the path shown in skyblue, the highest hill to be overcome is activity pattern σ , which has energy $E = -0.8$. The magenta path makes the travel easier with the largest energy of $E = -1.2$. Although the green path is longer than the skyblue path, the green path has the largest energy of $E = -1.1$, which is smaller than that for the skyblue path. As this example shows, longer paths may be easier to go through than shorter paths. If all the other paths between α_1 and α_2 have an activity pattern with $E > -1.2$, we conclude that the magenta path is the easiest path between α_1 and α_2 and record its peak height, denoted by $\bar{E}_{1,2}$ ($= -1.2$ in the current example). Computationally, we do not need to actually enumerate all the possible paths because there are efficient algorithms to realize the same computation [6,32,54,66].

How easy it is for the system's state to travel between α_i and α_j depends on the energy of α_i and α_j as well as on $\bar{E}_{i,j}$. If $E(\alpha_i)$ is much smaller than $\bar{E}_{i,j}$, then it is difficult for a trajectory starting at α_i to reach $\bar{E}_{i,j}$ even if $\bar{E}_{i,j}$ is small. We define the energy barrier between α_i and α_j from the viewpoint of α_i by $\bar{E}_{i,j} - E(\alpha_i)$. The energy barrier between the same pair of local minima from the viewpoint of α_j , which is equal to $\bar{E}_{i,j} - E(\alpha_j)$, is different from that from the viewpoint of α_i in general. For example, Fig 8(b) indicates that $\bar{E}_{1,2} - E(\alpha_1) = -1.2 - (-2.1) = 0.9$ and $\bar{E}_{1,2} - E(\alpha_2) = -1.2 - (-1.6) = 0.4$.

We carry out the same computation between each pair of local minima to obtain all $\bar{E}_{i,j}$. Once this is done, one can visualize all the local minima and energy barrier between each pair of them in a dendrogram, which is the disconnectivity graph, as shown in Figs 3(e) and 7(a).

The vertical axis of the disconnectivity graph represents the energy. A leaf, which vertically hangs down in the upside-down tree, represents a local minimum. The vertical position of the tip of the leaf represents the energy of the local minimum. Observe in Fig 3(e) that $E(\alpha_1) = -2.1$ and $E(\alpha_2) = -1.6$, which are consistent with Fig 8(b). We also observe that α_4 is the global minimum activity pattern in terms of the energy. Therefore, α_4 is the most frequent in the estimated PMEM and also likely so in the given data. The horizontal axis of the disconnectivity graph does not bear any meaning; different local minima are horizontally spread out only for visualization purposes. We also observe in Fig 3(e) that the branch of α_1 and that of α_2 joins at $\bar{E}_{1,2} = -1.2$, which is also consistent with Fig 8(b). In this manner, the disconnectivity graph represents all the local minima with their energy and the minimal height of the hill that any trajectory between each pair of local minima has to overcome.

The disconnectivity graph shown in Fig 3(e) suggests grouping of the six local minima into three groups, one composed of α_1 , α_2 , and α_3 , one composed solely of α_4 , and the other composed of α_5 and α_6 . Alternatively, one may want to avoid grouping α_1 together with α_2 and α_3 . One obtains such a grouping by cutting the disconnectivity graph into subtrees by drawing an arbitrary horizontal line into the disconnectivity graph (e.g., at $E = -1.1$, as shown by the dotted line in Fig 3(e)). For example, the authors of [6] similarly grouped 10 local minima into three groups (shown in Fig 1(e) of [6]). This procedure is the same as determination of the number of clusters in hierarchical clustering and necessarily arbitrary. Many studies relied on visual inspections or arbitrary criteria. We will explain a more principled method of grouping together some local minima into one major state in Section 3.2.2.

2.6. Basin graph

The disconnectivity graph only contains the local minimum activity patterns. How should we interpret activity patterns that are not local minima? The low-dimensional schematics shown in Figs 1 and 2 suggest that activity patterns that are not local minima belong to the basin of a local minimum. In other words, if we release a ball at the position of an arbitrary activity pattern, σ , on the energy landscape, and if the ball only goes downhill (which is contrary to the actual stochastic dynamics implied by the PMEM though), the ball eventually reaches one of the local minimum α and stops there. In this case, we say that σ belongs to the basin of α . In reality, visual understanding, such as in Figs 1 and 2, is not easy to obtain because the network of activity patterns is a hypercube. However, the definition of the basin remains the same. We can tell the basin to which an arbitrary activity pattern, σ , belongs as follows. We start by identifying the neighbor σ' of σ that attains the lowest energy among all the N neighbors of σ (see Fig 9(a) for a schematic). Unless σ is itself a local minimum, σ' has a lower energy than σ , which we assume here. If there is a tie, we break it arbitrarily to set σ' . The direction from σ to σ' is the steepest descent direction from σ . Therefore, we move from σ and σ' in search of the bottom of the basin. Then, at σ' , we similarly identify the neighboring activity pattern having the lowest energy, denoted by σ'' . If σ'' 's energy is lower than σ' 's energy, then we move to σ'' . Otherwise, σ' has an energy that is lower than all its neighbors such that σ' is a local minimum. In the latter case, we stop the journey at σ' and conclude that σ belongs to the basin of σ' . If this is not the case, we repeat the procedure until we reach a local minimum, and we call that σ belongs to the basin of the finally visited activity pattern, which is a local minimum. In Fig 9(a), the steepest descent path starting at σ stops at σ'' . Therefore, σ'' is the local minimum associated with σ . In this manner, we can identify to which basin each activity pattern belongs.

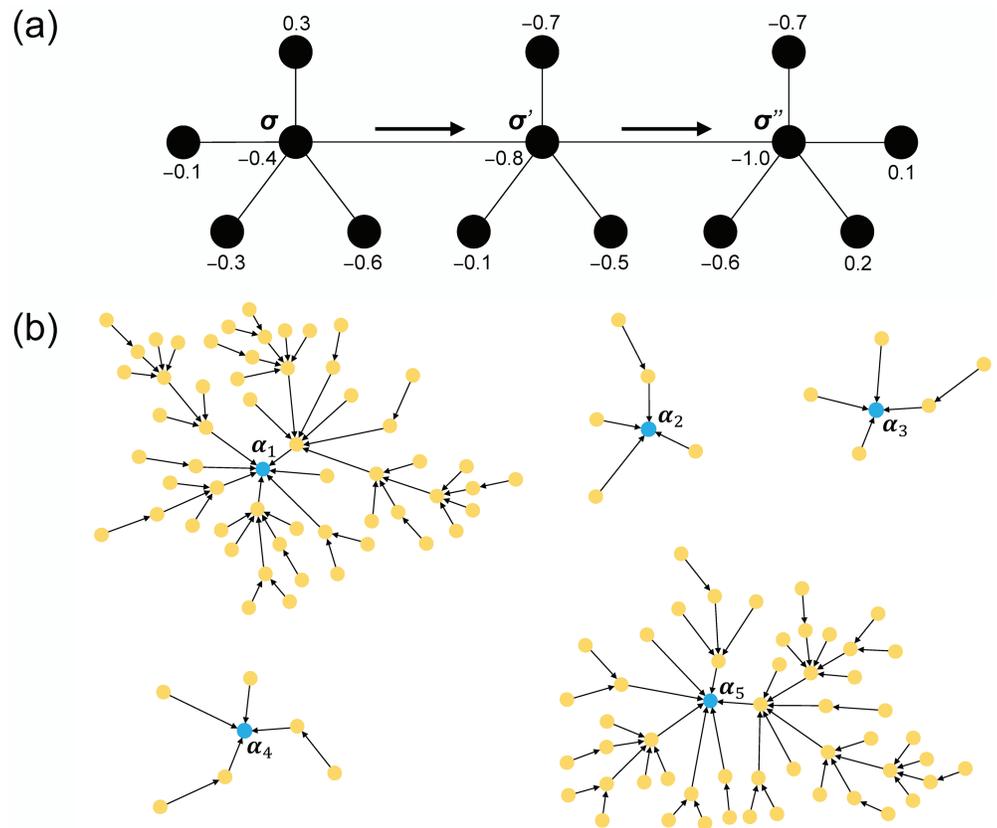


Fig 9. Basin graph. (a) The steepest descent path from σ to its associated local minimum σ'' . A node represents an activity pattern. The numbers attached to each activity pattern represents its energy value. We set $N = 5$. In fact, σ and σ'' share two neighbors because they are distance 2 apart from each other. One such neighbor is σ' . The other shared neighbor is an activity pattern with $E = -0.6$. For expository purposes, we showed this second neighbor as separate nodes in this figure. (b) An empirical basin graph with $N = 7$ corresponding to the disconnectivity graph shown in Fig 7(a).

<https://doi.org/10.1371/journal.pcsy.0000039.g009>

What we call the basin graph tells us the basin of each activity pattern and more; see Fig 3(f) for a schematic and Fig 9(b) for a real example. A node of the basin graph is an activity pattern. Therefore, there are 2^N nodes; in Figs 3(f) and 9(b), there are $2^5 = 32$ and $2^7 = 128$ nodes, respectively. The edge of the basin graph represents the steepest descent move from an activity pattern to its neighbor. Differently from the network of activity patterns, which is also a network among the 2^N activity pattern nodes, the edge of the basin graph is directed, and each node has at most one out-going edge. If a node is not a local minimum, it has one out-going edge (i.e., the node's out-degree is equal to 1); see the orange nodes in Figs 3(f) and 9(b). If a node is a local minimum, it does not have any out-going edge; see the blue nodes in the same figures. Any node has between 0 to N in-coming edges; its number is called the in-degree of the node. The basin graph is composed of so-called weakly connected components (i.e., islands), each of which contains just one local minimum. This fact aids the interpretation of the basin graph. The number of nodes in the connected components is equal to the size of the basin of the corresponding local minimum. Comparison between Figs 3(e) and 3(f) and between Figs 7(a) and 9(b) suggests that an local minimum with a smaller energy value

(i.e., therefore, high frequency of appearance) tends to have a large basin. However, while this association is often found in empirical data, it is not guaranteed.

Some activity patterns may be located near the border between different basins. Such an activity pattern roughly, but not exactly, corresponds to those on the top of a hill separating two basins (e.g., the local maximum in Fig 1 or a dotted line in Fig 2). Precisely speaking, if activity pattern σ belonging to the basin of α_i has a neighboring activity pattern σ' belonging to the basin of a different local minimum, then we say that σ and σ' are at the boundary between the two basins.

2.7. Features of energy landscapes

The disconnectivity graph and basin graph are visualizations of the energy landscape. They are useful for us to intuitively understand the dynamics generating the multivariate time series data and differences between different data or between different groups of data (e.g., between a group of healthy individuals and a group of those with disease) in terms of the energy landscape. However, we need to quantify such differences to draw conclusions. We refer to such a quantity as a feature of the energy landscape. A primary usage of a feature is to look for association between the feature and a property of interest in the research domain. For example, a schizophrenia researcher may be interested in seeking correlation between the severity of schizophrenia and a feature of the energy landscape. As another example, one may want to measure the effect of treatment in terms of the change in the value of a feature, especially if the feature is known to correlate with the severity of the disease.

A feature is also useful for assessing reliability of the ELA. If one measures a multivariate time series data of length t_{\max} twice from the same system (e.g., same person in two separate days), run an ELA, and calculate a feature from the two energy landscapes, the value of the feature should be roughly the same between the two energy landscapes. Otherwise, one cannot rely on the ELA. Such a test, which should be done statistically, also requires a feature of energy landscape. Popular features of energy landscapes are as follows:

- The number of local minima [5,16], denoted by n_m . Because a local minimum may not be significant in the sense that it is separated from a different local minimum with a small energy barrier (see Section 3.2.2 for a procedure to remove insignificant local minima), one may want to only count major local minima, i.e., significantly deep local minima.
- The frequency with which a local minimum α appears in the data [16,17]. It is the number of times t at which $\sigma(t) = \alpha$, divided by t_{\max} . Because the dynamics on the energy landscape are noisy by definition, it is more common to calculate the frequency with which the basin of a local minimum or that of a major local minimum is visited [6,14,16,21,22,26,37].
- The basin size [5,6,14–16]. It is the number of activity patterns belonging to the basin of a local minimum α , that can be divided by 2^N depending on the preference. It is equal to the number of nodes in the connected component of the basin graph to which α belongs. The basin size tends to correlate with $E(\alpha)$ [5] and the frequency of appearance of the basin [37].
- The energy barrier between local minima α_i and α_j . It can be quantified by $\min\{\bar{E}_{i,j} - E(\alpha_i), \bar{E}_{i,j} - E(\alpha_j)\}$, where we recall that $\bar{E}_{i,j}$ is the smallest energy value that any path between α_i and α_j needs to overcome [5,15,16]. The energy barrier quantifies the difficulty with which the system's state transits from α_i to α_j or vice versa. If the energy barrier is high, the transition occurs at a low rate in both directions. Sometimes, the unsymmetrized,

or directed, version of the energy barrier, i.e., $\bar{E}_{i,j} - E(\alpha_i)$ or $\bar{E}_{i,j} - E(\alpha_j)$, but not the minimum of the two, is used [15,37]. See Section 3.2.3 for quantitative relationships between the height of the energy barrier and the state transition rate.

A variant of this feature is an average over different local minima. In [32], the authors used the average of branch length over all the major local minima in the pruned disconnectivity graph. By the pruned disconnectivity graph, we mean one created by sequentially removing insignificant local minima, to keep only the major local minima (see Section 3.2.2 for the procedure). The branch length for local minimum α_i is $\min_{j \text{ such that } j \neq i} \{\bar{E}_{i,j} - E(\alpha_i)\}$, where $\bar{E}_{i,j}$ now refers to the smallest energy at which the i th and j th branches of the pruned disconnectivity graph merge. For example, in Fig 3(e), suppose that the pruning process discards α_2 , α_3 , and α_6 , leaving α_1 , α_4 , and α_5 as major local minima. This pruning is equivalent to setting the threshold energy to the one shown by the dotted horizontal line in the figure to define three major states, each of which is represented by a major local minimum. Then, the branch length for α_1 , α_4 , and α_5 are equal to $\bar{E}_{1,4} - E(\alpha_1)$, $\bar{E}_{1,4} - E(\alpha_4)$, and $\bar{E}_{1,5} - E(\alpha_5)$, respectively. The last one is equal to $\bar{E}_{4,5} - E(\alpha_5)$ because α_1 and α_4 merge earlier than α_5 does as we climb up the disconnectivity graph from a low energy value. The average branch length used in [32] is the average of these three branch lengths and represents how distinct the different major local minima are from each other.

- The frequency of transitions between basins [6,14,15,17]. This is the number of transitions from the basin of local minimum α_i to that of local minimum α_j . To explain the idea, we note that the activity pattern at each time $t \in \{1, \dots, t_{\max}\}$, i.e., $\sigma(t)$, belongs to the basin of an i th local minimum, which we write as $s(t) = i$, where each $s(t) \in \{1, \dots, n_m\}$. Then, the time series of the basin is given by $\{s(1), \dots, s(t_{\max})\}$. The normalized frequency of transitions from the i th basin (i.e., basin of local minimum α_i) to the j th basin is the count of times with $s(t) = i$ and $s(t + 1) = j$ over $t \in \{1, \dots, t_{\max} - 1\}$, divided by $t_{\max} - 1$. Another way to normalize the count is to divide it by the number of times at which $s(t) = i$ over $t \in \{1, \dots, t_{\max} - 1\}$. The latter quantity is equal to the empirical probability of transitions from $s(t) = i$ to $s(t + 1) = j$ conditioned that $s(t) = i$.

There are some variants of the frequency of transitions. First, one sometimes measures the number of transitions from the local minimum α_i , not from its basin, to α_j [17,25]. A second variant is the frequency of indirect transitions from i (either the local minimum or its basin) to j through k [6,14,26,27,29,37]. In this count, we exclude the transitions from i to j that do not go through k . We call this type of transitions “indirect”. Those that do not go through another local minimum or its basin are called direct transitions. Third, one may want to count the number of iterations between two major local minima [6]. This quantity is basically the sum of the number of transitions from i to j and that from j to i .

- Average dwell time of an i th basin [15,16]. This is the duration for which $\sigma(t)$ resides in the i th basin (i.e., $s(t) = i$) before leaving it. Because the dwell time is generally different for different visits to the i th basin in a single data set, we average the dwell time of the i th basin over the number of the visits to the i th basin. In one report, the average dwell time was strongly correlated with the basin size [15].
- The average number of time steps for moving from the i th basin to the j th basin to another, forming a path [6].
- Various quantities that can be directly calculated from the estimated J_{ij} or h_i values. For example, a definition of the node strength (i.e., weighted degree) of the i th variable is $\sum_{j=1}^N |J_{ij}|$ [16]. This class of features only requires the estimated PMEM, not further steps of the ELA.

One often needs to pool data from many different individual data, such as different participants in the experiment, and then estimate a single group PMEM with a high accuracy. Even in this situation, one can measure some of these features for each individual despite that only one energy landscape for the entire group of individuals is available. One can do this by using the group energy landscape and still computing these features for time series $\{\mathbf{x}(1), \dots, \mathbf{x}(t_{\max})\}$ for each individual. For example, one can calculate the dwell time of the i th basin for a single individual because the group energy landscape gives α_i and its basin that are common to all the individuals. In contrast, in the same situation, one cannot use features that are directly measured for the estimated energy landscape, not for $\{\mathbf{x}(1), \dots, \mathbf{x}(t_{\max})\}$ for each individual. Such features include n_m and the energy barrier.

3. Further methods

In this section, we discuss advanced ELA methods, their backgrounds, and validation methods.

3.1. Phase diagram methods

An advantage of carrying out ELA using the Ising model is that one can potentially benefit from statistical mechanical theory of the Ising spin-system model, which is a mature research field. The notions of local minima of energy and energy barrier have been investigated in spin-system theory for decades. One can also attempt to use other knowledge of spin-system theory for analyzing and interpreting multivariate time series data. One such method is to determine the phase of the multivariate time series data [19].

The statistical mechanical theory exploited here is that of prototypical spin systems, such as the Sherrington-Kirkpatrick (SK) model [67]. The SK model is a type of Ising model and defined by drawing each J_{ij} ($= J_{ji}; 1 \leq i \neq j \leq N$) independently from the Gaussian distribution with a given mean and standard deviation. The SK model, as well as other spin system models, shows different phases depending on the parameter values (see Fig 10 for a schematic). In the paramagnetic phase, the spins (i.e., σ_i s) tend to be random and orderless in the sense that the

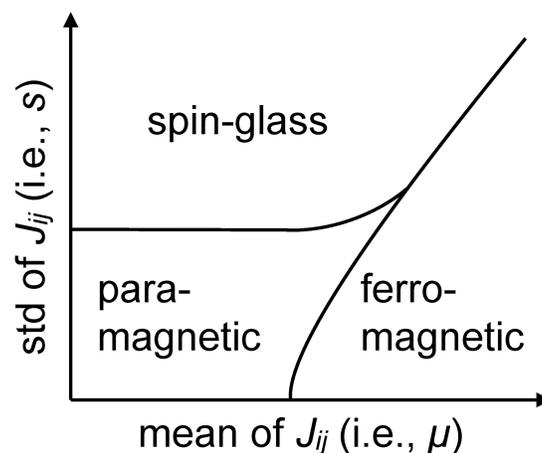


Fig 10. Schematic of the phase diagram of the SK model. A more common parameter in physics literature is the temperature. However, in this figure, we chose the average and standard deviation of J_{ij} because their interpretation is easier in our context.

<https://doi.org/10.1371/journal.pcsy.0000039.g010>

mean activity

$$m = \frac{1}{N} \sum_{i=1}^N \langle \sigma_i \rangle_{\text{model}} \quad (15)$$

and the so-called spin-glass order parameter defined by

$$q = \frac{1}{N} \sum_{i=1}^N (\langle \sigma_i \rangle_{\text{model}})^2 \quad (16)$$

are both close to 0. Theoretically, m and q converge to 0 as $N \rightarrow \infty$. The spin-glass order parameter q represents the degree of local (i.e., for each variable) magnetization, or the local stickiness of the activity. In an energy landscape, the paramagnetic phase corresponds to a situation in which local minima are shallow and not distinguishable from those existing in completely random spin systems. Therefore, if we run a significance test for the local minima in this situation (see Section 3.2.2 for a method), local minima in the paramagnetic phase should be classified to be insignificant.

In the ferromagnetic phase, the spins strongly tend to align in one direction, let say, $\sigma_i = 1$. It is defined by $m \neq 0$ and $q > 0$. For example, $m > 0$ if a majority of σ_i s tends to be 1. A strongest case of ferromagnetic phase corresponds to an energy landscape in which $\sigma = (1, \dots, 1)$ is by far the deepest local (and global) minimum. However, in ELA, we usually impose that $\langle \sigma_i \rangle_{\text{data}}$ is close to 0 by choosing a threshold on $x_i(t)$ as such. Therefore, $\langle \sigma_i \rangle_{\text{model}}$ is also expected to be close to 0, and m substantially far from 0 is not likely to occur. Instead, in practice, we often have an energy landscape in which $\sigma = (1, \dots, 1)$ and $\sigma = (-1, \dots, -1)$ are the only most dominant local minima [17,25,27,34].

The other phase is the spin-glass phase, which has attracted a lot of statistical physics research. The spin-glass phase corresponds to rugged energy landscapes (and therefore, complex state-transition dynamics) [67] and chaotic dynamics [68–70]. It is defined by $m = 0$ and $q > 0$ (in the $N \rightarrow \infty$ limit). In our ELA, the obtained energy landscape would have relatively many local minima whose branches are not too short to be ignorable.

Criticality, or phase transition, i.e., the boundary between (typically) the paramagnetic phase and the spin-glass phase in a parameter space (see Fig 10), is of special interest due to rich dynamical behavior of the spin system near there. On these grounds, Ezaki et al. developed a data-driven method to locate a given multivariate time series data as a point in the phase diagram to tell which phase the given data resides in, and quantify how close the data is from the boundary with a different phase [19].

The first step of the method is to fit the Ising model to the given data, as in the ordinary ELA. Because the above discussion holds true in principle in the limit of $N \rightarrow \infty$, we need to estimate a relatively large Ising model. If N is small, phase transition lines would be blurred. We used $N = 264$ in [19]. Then, we need to give up an exact estimation method and resort to an approximate model estimation method such as pseudo-likelihood maximization. A large N leads to a low accuracy of fit, which is an issue to be investigated in future research.

Next, to draw a phase diagram from the given multivariate time series data, we fix $\mathbf{h} \equiv (h_1, \dots, h_N)$ at that for the estimated PMEM by following a convention [67,71]. Under this constraint, we vary the mean of $\mathbf{J} \equiv (J_{12}, J_{13}, \dots, J_{N-1,N})$, denoted by μ , and the standard deviation of \mathbf{J} , denoted by s , to linearly transform \mathbf{J} as follows:

$$J_{ij} = (\hat{J}_{ij} - \hat{\mu}) \frac{s}{\hat{s}} + \mu, \quad (17)$$

where $\hat{\mu}$, for example, is the value of μ for the originally estimated PMEM. With Eq (17), if we arbitrarily choose the values of μ and s , then we obtain J , which specifies one PMEM. For each J specified by a (μ, s) pair, we perform Monte Carlo simulations to produce $\{\sigma(1), \sigma(2), \dots\}$ and calculate observables such as m and q . In this manner, we can generate a phase diagram for each observable in terms of μ and s , as in Fig 10.

With this method, energy landscapes estimated for fMRI data from healthy human participants were in the paramagnetic phase but close to the boundary between the paramagnetic and spin-glass phases [19]. It was also shown that the participants with higher intelligence quotient scores were, albeit weakly, closer to the phase boundary, consistent with the so-called critical brain hypothesis [19].

3.2. Validation methods and good practices

Because ELA fits a model (i.e., PMEM) to data, it is imperative to validate the methods on the data. The accuracy of fit introduced in Section 2.3 is one validation method. In this section, we explain other validation methods.

3.2.1. Threshold variation. The choice of the threshold, θ_i , for reducing each i th time series $\{x_i(1), \dots, x_i(t_{\max})\}$ into a binarized time series $\{\sigma_i(1), \dots, \sigma_i(t_{\max})\}$ is an arbitrary parameter. We recommend that one tests the robustness of the results, such as the accuracy of fit, the disconnectivity graph at the level of major local minima, and association between a feature of the an energy landscape and, e.g., a cognitive score for a human participant in the experiment, for a range of values of θ_i . For example, in some studies, it was verified that the accuracy of fit remained large for a range of θ_i , $i \in \{1, \dots, N\}$, in which the frequency of $\sigma_i = 1$ (as opposed to $\sigma_i = -1$) varied between approximately 30% and 70% [13,17,72].

3.2.2. Randomization. In Fig 3, local minimum α_2 is shallow in the sense that the energy barrier between α_2 and α_3 is small. Therefore, the dwell time in α_2 would be small. Therefore, although α_2 is technically a local minimum, it may be better to ignore it in practice. Only the local minima that have sufficiently high energy barriers with others are trustable. We already discussed this topic; by choosing a threshold energy level, shown by the dotted line in Fig 3(e), one can classify the local minima into groups. However, how should we choose the threshold? How can one tell whether or not a local minimum is too shallow to be seriously considered? This decision has been rather arbitrary in most studies using the ELA. The root of the problem is basically the same as an often asked problem for clustering algorithms for data: how to determine the number of clusters?

For ELA, we developed a statistical test to tell which local minima are significant [32]. Specifically, to generate a completely random disconnectivity graph to be contrasted with the disconnectivity graph from the given data, we generate a completely random multivariate time series of the same size (i.e., same N and t_{\max}) as that of the original time series data by assigning -1 or 1 with probability 0.5 each, independently to each of $\sigma_i(t)$, $i \in \{1, \dots, N\}$, $t \in \{1, \dots, t_{\max}\}$. For the generated random multivariate time series, we estimate the PMEM and then the disconnectivity graph. A completely random time series should not have any meaningful structure including local minima. Therefore, we regard that the branch length in the generated random disconnectivity graph (see Section 2.7 for the definition of the branch length) for any α is meaningless. To operationalize this idea, we record the maximal branch length among all the local minima in the random disconnectivity graph.

We compute the maximal branch length for a sufficiently large number, c , of completely random multivariate time series data. We set $c = 10^2$ in [32] and use $c = 10^3$ in the following running example in this section. We denote the mean and standard deviation of the maximal branch length over the c random disconnectivity graphs by $\langle L \rangle$ and $\text{std}(L)$, respectively.

Then, we remove the local minimum with the shortest branch length in the original disconnectivity graph as being insignificant if its branch is shorter than $\langle L \rangle + 2 \times \text{std}(L)$. For example, Fig 7(b) shows the disconnectivity graph for a random multivariate binary time series having the same $N (= 7)$ and $t_{\max} (= 3, 272)$ as those for the original data whose disconnectivity graph is shown in Fig 7(a). Due to the randomization, the disconnectivity graph in Fig 7(b) has shorter branch lengths than the disconnectivity graph in Fig 7(a) in general. The maximal branch length in Fig 7(b) is that of α'_3 and equal to $L = 0.125$. Based on $c = 10^3$ random disconnectivity graphs, we obtained $\langle L \rangle = 0.0902$ and $\text{std}(L) = 0.0471$. Therefore, $\langle L \rangle + 2 \times \text{std}(L) = 0.184$. The shortest branch of the empirical disconnectivity graph shown in Fig 7(a), which is for α_3 , is shorter than 0.184. Therefore, we regard α_3 as an insignificant local minimum and prune it away.

If we have removed a local minimum according to this procedure, then we recompute the branch length of the local minima affected by the removal. In Fig 7(a), removal of the branch of α_3 does not affect the branch length of any other local minima, especially because α_4 and α_5 merge before they merge α_3 . Therefore, there is no need to recompute any branch length. After the possible recomputation, we repeat the above procedure. In other words, if the shortest branch is shorter than $\langle L \rangle + 2 \times \text{std}(L)$, we remove the branch and recompute the affected branch lengths. For example, in Fig 7(a), we next remove the branch of α_2 because it is shorter than the threshold (i.e., 0.184). This removal changes the branch length for α_1 from 0.814 to 1.082. Next, we remove the branch of α_4 , which increases the branch length for α_5 from 0.761 to 0.902.

We repeat the procedure until the branches of all the surviving local minima are longer than the threshold in the pruned disconnectivity graph. We refer to the surviving local minima as major local minima. For example, for Fig 7(a), the final, pruned disconnectivity graph, shown in Fig 7(c), only consists of α_1 and α_5 , which are the two major local minima. This pruning algorithm also tells us how to associate each removed local minimum and its basin to a major state. In Fig 7(c), α_1, α_2 , and their basins constitute a major state represented by α_1 . Similarly, $\alpha_3, \alpha_4, \alpha_5$, and their basins constitute a major state represented by α_5 .

There are two remarks. First, the threshold $\langle L \rangle + 2 \times \text{std}(L)$ is arbitrary and probably too generous. We test the significance of the branch length at most n_m times, where we recall that n_m is the number of local minima. Therefore, one may want to set a more conservative threshold, such as $\langle L \rangle + 3 \times \text{std}(L)$ or one that explicitly carries out Bonferroni correction, to avoid a potential multiple comparison problem. Second, one may want to use a different method to generate random multivariate time series of the same size (i.e., $N \times t_{\max}$). In fact, it is a common practice to randomize the given time series, creating surrogate time series, then measure a quantity to characterize time series, q , and compare the value of q with that for the original time series to assess whether the q value for the original time series is statistically meaningful [73,74]. We can implement such a procedure here, too. However, it should be noted that time stamp shuffling (i.e., randomly permute the time stamps of $\{\sigma(1), \dots, \sigma(t_{\max})\}$) does not work because the PMEM ignores the time order of the data. The PMEM estimated for the original data and that estimated for the data after time stamp shuffling are exactly the same.

3.2.3. Random-walk simulations. Once the PMEM is estimated, one can generate random-walk trajectories on the hypercube to emulate the dynamics of the system's state with a high time resolution (because each step of the walk is only to the nearest neighbor on the hypercube by definition) and for arbitrary long time [6,14,15]. We consider discrete-time random walks for simplicity and construct them such that the moves from activity pattern σ to σ'

and vice versa respect the detailed balance condition, i.e.,

$$\frac{T_{\sigma \rightarrow \sigma'}}{T_{\sigma' \rightarrow \sigma}} = e^{E(\sigma) - E(\sigma')}, \quad (18)$$

where $T_{\sigma \rightarrow \sigma'}$ is the transition probability of the random walk from σ to σ' . This detailed balance condition is a quantitative formulation of the already made claim that the movement of the state uphill on the energy landscape is not much likely to occur, especially when the energy barrier to be overcome to go to a different activity pattern is high. In this way, moves from a high energy activity pattern to a low energy one are more likely than vice versa. However, the random walk is a stochastic process, so moves from a lower energy activity pattern to a higher one do occur with some probability. Note that Eq (18) is consistent with the stationary density given by Eq (6).

Random-walk simulations are useful for validation of the ELA. For example, we can calculate the transition probability from a major state represented by a major local minimum α to an adjacent major state represented by major local minimum α' , denoted by $T_{\alpha \rightarrow \alpha'}^{(\text{basin})}$, from both random-walk simulations and the given data. If these two numbers are drastically different, the ELA may not be a suitable method for the data. Such a discrepancy, if it exists, may be because the transition probability in the given data is more complex than what the detailed balance condition dictates. For example, if the probability of the move $\sigma \rightarrow \sigma'$ depends on where the trajectory comes from to σ , i.e., the last activity pattern before σ , it should not be fitted by the PMEM in principle. This is because the PMEM, hence ELA, does not assume such a dependence. Alternatively, if the given data is too short, we may not observe sufficiently many transitions from the basin of α to that of α' such that the statistical fluctuation in $T_{\alpha \rightarrow \alpha'}^{(\text{basin})}$ estimated for the given data is too large.

A practical method to run random-walk simulations is a Markov chain Monte Carlo method with the Metropolis-Hastings algorithm [75] (used in, e.g., [6]). In each discrete time step, the random walker selects a neighbor, σ' , of the currently visited node, σ , with equal probability $1/N$. The walker actually moves from σ to σ' with probability $T_{\sigma \rightarrow \sigma'} = \min\{1, e^{E(\sigma) - E(\sigma')}\}$. Otherwise, the walker does not move. We remark that, if one prefers, main features of the random walk such as the dwell time and transition probabilities, including those between basins, can be theoretically calculated without running simulations [76]. This is because the complete transition probability matrix of the random walk on the hypercube with 2^N nodes is available once the PMEM is estimated.

3.2.4. Test-retest reliability. If ELA is reliable, the energy landscape constructed from the data collected in one day should be similar to that on another day unless the system has sufficiently changed between the two days. Otherwise, we cannot trust the constructed energy landscapes even if the accuracy of the fit values are high. We recently developed a statistical test of test-retest reliability for experimental designs where data are collected in multiple sessions from each participant of the experiment and over multiple participants [32]. The main idea is to compare two energy landscapes constructed from two disjoint sets of sessions from the same participant, called the within-participant comparison, and two energy landscapes constructed from different participants, called the between-participant comparison. If ELA is reliable, the within-participant comparison should yield a smaller discrepancy score than the between-participant comparison, which we operationalized. Because one session is often too short to yield a high accuracy of fit, we built the above procedure to also work when one estimates each energy landscape from data pooled across sessions. In other words, a within-participant comparison is made between two energy landscapes from the same participant, with each energy landscape being built from \tilde{m} sessions from the same participant.

Then, we need at least $2\tilde{m}$ sessions of data from the same participant. The obtained within-participant comparison result is to be contrasted with the between-participant comparison, which consists in comparing the energy landscape constructed from \tilde{m} sessions from \tilde{m} participants, one session from each participant, and that constructed from \tilde{m} sessions from other \tilde{m} participants.

3.3. Challenges

We recognize two main difficulties in the original ELA.

First, ELA is data-hungry. If we aim at maintaining a high accuracy of fit of the PMEM, we roughly need to increase the data length, t_{\max} , twice as we add one variable (i.e., as we increase N by 1). Therefore, t_{\max} needed increases exponentially in terms of N , i.e., in proportion to 2^N [54]. Then, if we seriously take the accuracy, one cannot go beyond approximately $N = 10$ or $N = 15$, depending on t_{\max} of the data, while ignoring the fitting accuracy to conduct a large-scale ELA (e.g., see Section 3.1) is a methodological choice.

For this reason, in neuroimaging data analysis, in which the ELA has been most intensively used, it is customary to pool data from different participants in the experiment to enhance the size of the data. A typical approach is to pool data across participants in a control group and run an ELA, estimating just one energy landscape. Then, we do the same for a set of participants in a treatment group (such as those receiving the same diagnosis) and compare the results between the two groups. This procedure ameliorates the data length issue at a cost: the data from different participants in either group may not be as homogeneous as that from a single participant. If the participants within a group are too heterogeneous, then estimating a common ELA from the pooled data may not be justified.

Note that, when pooling data, we may need to use different threshold values, θ_i , for binarizing $x_i(t)$ for different participants. For example, assume that one participant, denoted by p_1 , tends to have a large $x_i(t)$. We also assume that a different participant classified to the same group, denoted by p_2 , tends to have a small $x_i(t)$. Then, if we use θ_i shared by everybody in the group, $x_i(t)$ from p_1 tends to yield $\sigma_i(t) = 1$ and $x_i(t)$ from p_2 tends to yield $\sigma_i(t) = -1$. However, this is not what ELA intends to do; ELA intends to analyze the system's dynamics for each participant under the assumption that the dynamics are similar among the participants in the same group. A common method to evade this situation is to set θ_i separately for different participants such that, for example, the fraction of $\sigma_i(t) = -1$ and that of $\sigma_i(t) = 1$ are approximately the same (i.e., 50%) for each participant (e.g., [17,32]).

Second, binarization implies loss of information contained in the original multivariate time series data. If the original time series data for each variable, $\{x_i(1), \dots, x_i(t_{\max})\}$, shows bimodality, binarization is better justified. However, bimodal distributions of a time series are not frequently found in many domains. Despite the apparent loss of information by binarization, the ELA has been successful in many studies. The reason for this is unknown.

In the next two sections, we introduce two methods that expand the original ELA to address these problems.

3.4. Variational Bayes approximation method

Consider a collection of multivariate time series data from different participants in the same experiment. We assume that these participants belong to the same group (e.g., people with a particular psychiatric disorder). It is typical that the data from a single participant is not long enough to be able to estimate the PMEM with a high accuracy. If we pool the data across all the participants in the group, the combined data may be sufficiently long, but we may suffer from heterogeneity in the data among the different participants, as we discussed in

Section 3.3. Because the PMEM implicitly assumes that observed data points, no matter whether they come from the same participant or different participants, are drawn independently from the same distribution, such a heterogeneity, which may be better modeled by a participant-dependent distribution of data, is undesirable. In this situation, a variational Bayes approximation method [20,21] enables us to estimate the PMEM, and hence the energy landscape, separately for each participant while benefiting from some commonality that the participants in the same group have. Core steps of this Bayesian method are as follows.

Assume that there are D multivariate time series data (such as those from D participants) belonging to the group. This method uses a prior distribution of the PMEM parameters specific to each time series data in the group (referred to as “series” for short in the remainder of this section), denoted by

$$\theta_{D'} = (\theta_{D'1}, \dots, \theta_{D'M}) \equiv (h_1, h_2, \dots, h_N, J_{12}, J_{13}, \dots, J_{N-1,N}) \in \mathbb{R}^M \tag{19}$$

for the D' th series, where $D' \in \{1, \dots, D\}$, and $M \equiv N(N + 1)/2$ is the number of parameters of the PMEM. We set the prior distribution of

$$\Theta = [\theta_1, \dots, \theta_D] \tag{20}$$

to

$$p(\Theta|\boldsymbol{\eta}, \tilde{\boldsymbol{\alpha}}) = \prod_{D'=1}^D \prod_{M'=1}^M p(\theta_{D'M'}|\mathcal{N}(\eta_{M'}, 1/\tilde{\alpha}_{M'})), \tag{21}$$

where $p(x|\mathcal{N}(\mu, \sigma^2))$ represents the probability density of x obeying the one-dimensional normal distribution with mean μ and variance σ^2 . Here, $\boldsymbol{\eta} = (\eta_1, \dots, \eta_M)^\top \in \mathbb{R}^M$ is the prior mean vector, $\tilde{\boldsymbol{\alpha}} = (\tilde{\alpha}_1, \dots, \tilde{\alpha}_M)^\top \in \mathbb{R}_+^M$ is the prior precision vector, and $^\top$ represents transposition. Equation (21) implies that the data from the D series are mutually independent.

It is intractable to derive the posterior distribution of Θ because the normal distribution is not a conjugate prior for the Boltzmann distribution. Therefore, we use a variational approximation to the posterior distribution [77] using the normal distribution as follows:

$$q(\Theta|\mathcal{S}, \boldsymbol{\eta}, \tilde{\boldsymbol{\alpha}}) = \prod_{D'=1}^D \prod_{M'=1}^M p(\theta_{D'M'}|\mathcal{N}(\mu_{D'M'}, 1/\beta_{D'M'})), \tag{22}$$

where q is the posterior distribution, and \mathcal{S} is the set of the data from all the D series.

We denote the posterior mean vector and the posterior precision vector for the D' th series by $\boldsymbol{\mu}_{D'} = (\mu_{D'1}, \dots, \mu_{D'M})^\top \in \mathbb{R}^M$ and $\boldsymbol{\beta}_{D'} = (\beta_{D'1}, \dots, \beta_{D'M})^\top \in \mathbb{R}_+^M$, respectively. One obtains the variational approximate solution for q by optimizing the evidence lower bound (ELBO), which is also called the free energy [20,21]. By maximizing the free energy with respect to q , we calculate the posterior quantities for each D' th series, $\boldsymbol{\mu}_{D'}$ and $\boldsymbol{\beta}_{D'}$, from the prior mean vector $\boldsymbol{\eta} = (\eta_1, \dots, \eta_M)^\top$ and the prior precision vector $\tilde{\boldsymbol{\alpha}} = (\tilde{\alpha}_1, \dots, \tilde{\alpha}_M)^\top$ as follows:

$$\boldsymbol{\mu}_{D'} = \boldsymbol{\eta} + t_{\max} \mathbb{A}_{\boldsymbol{\eta}, \tilde{\boldsymbol{\alpha}}}^{-1} (\langle \bar{\boldsymbol{\sigma}} \rangle_{\text{data}, D'} - \langle \bar{\boldsymbol{\sigma}} \rangle_{\boldsymbol{\eta}}), \tag{23}$$

$$\boldsymbol{\beta}_{D'} = \tilde{\boldsymbol{\alpha}} + t_{\max} \mathbf{C}_{\boldsymbol{\eta}}, \tag{24}$$

where

$$\mathbb{A}_{\boldsymbol{\eta}, \tilde{\boldsymbol{\alpha}}} = \text{diag}(\tilde{\boldsymbol{\alpha}}) + t_{\max} \mathbf{C}_{\boldsymbol{\eta}}, \tag{25}$$

is an $M \times M$ matrix, and $\text{diag}(\cdot)$ represents the diagonal matrix whose entries are given by the arguments. In Eq (23),

$$\langle \bar{\sigma} \rangle_{\text{data}, D'} \equiv (\langle \sigma_1 \rangle_{\text{data}, D'}, \dots, \langle \sigma_N \rangle_{\text{data}, D'}, \langle \sigma_1 \sigma_2 \rangle_{\text{data}, D'}, \langle \sigma_1 \sigma_3 \rangle_{\text{data}, D'}, \dots, \langle \sigma_{N-1} \sigma_N \rangle_{\text{data}, D'})^T \quad (26)$$

is the vector composed of the mean activity of the individual variables and the mean pairwise joint activities of the data; vector $\langle \bar{\sigma} \rangle_{\eta}$ is the model mean of

$$\bar{\sigma} \equiv (\sigma_1, \dots, \sigma_N, \sigma_1 \sigma_2, \sigma_1 \sigma_3, \dots, \sigma_{N-1} \sigma_N)^T \quad (27)$$

when the PMEM parameters $(h_1, \dots, h_N, J_{12}, J_{13}, \dots, J_{N-1, N})$ are given by η ; matrix C_{η} is the $M \times M$ covariance matrix of $\bar{\sigma}$ when the PMEM parameters are η . In Eq (24), c_{η} is the vector composed of the diagonal entries of C_{η} . In other words, the i th entry of c_{η} is the variance of the i th entry of $\bar{\sigma}$ given PMEM parameters η .

Now, we fix q and maximize the free energy with respect to η and $\tilde{\alpha}$ to obtain the updating rules for η and $\tilde{\alpha}$ as follows:

$$\eta_{M'} = \frac{1}{D} \sum_{D'=1}^D \mu_{D'M'}, \quad (28)$$

$$\tilde{\alpha}_{M'} = \left[\frac{1}{D} \sum_{D'=1}^D \left\{ (\mu_{D'M'} - \eta_{M'})^2 + \frac{1}{\beta_{D'M'}} \right\} \right]^{-1}, \quad (29)$$

where $M' \in \{1, \dots, M\}$.

Thus, we have updated the posterior distribution $\theta_{D'M'} \sim \mathcal{N}(\mu_{D'M'}, 1/\beta_{D'M'})$, $D' \in \{1, \dots, D\}$, $M' \in \{1, \dots, M\}$ using the prior distribution $\theta_{D'M'} \sim \mathcal{N}(\eta_{M'}, 1/\tilde{\alpha}_{M'})$, and then updated the prior distribution using the obtained posterior distribution. These operations constitute one cycle of the algorithm. The steps of the variational Bayes approximation method are as follows:

1. Initialize the hyperparameters by independently drawing each $\eta_{M'}$ (with $M' \in \{1, \dots, M\}$) from the normal distribution with mean 0 and standard deviation 0.1. We also set the first N entries of the prior precision vector $\tilde{\alpha}$, corresponding to h_i , $i \in \{1, \dots, N\}$, to 6, and set the remaining $M-N$ entries of $\tilde{\alpha}$, corresponding to J_{ij} (with $1 \leq i < j \leq N$), to 30. These values are used in [20,21] and their accompanying code.
2. Calculate the posterior mean vector, $\mu_{D'}$, and the posterior precision vector, $\beta_{D'}$, for each $D' \in \{1, \dots, D\}$ using Eqs (23) and (24).
3. Update the prior mean vector, η , and the prior precision vector, α , using Eqs (28) and (29).
4. If $\left| \frac{\text{ELBO}(\text{iter})}{\text{ELBO}(\text{iter}-1)} - 1 \right| < 10^{-8}$, we stop iterating steps 2 and 3. Otherwise, we return to step 2. Here, $\text{ELBO}(\text{iter})$ represents the ELBO value after ‘iter’ iterations.

3.5. Clustering and Markovian approaches

An attempt to mitigate possible loss of information by binarizing the originally continuous signal is to dichotomize the original signal into a larger number of discrete values. However, this idea does not help much. For example, using ternary activity for each variable (i.e., $\sigma_i(t) \in \{-1, 0, 1\}$) did not enhance the accuracy of fit in one study [13]. This is presumably because, by introducing a larger number of discrete activity levels for each

variable, we have less samples (i.e., time points in the case of time series) in each activity pattern. If there are three activity levels per variable, there are 3^N possible activity patterns as opposed to 2^N . If many of the 3^N possible activity patterns are never visited by the data or their counts are close to 0, the estimation of the corresponding PMEM may not be accurate.

Another strategy is to fit a different model than the PMEM to try to realize the same goal as the ELA. The main goals of ELA are to identify a small number of major states from the multivariate time series data in a data-driven manner and to interpret the time series as a journey from one identified major state to another and further.

One simple alternative to the ELA along this idea is to cluster the t_{\max} data points in \mathbb{R}^N using any clustering algorithm and regard each cluster as a major state [78–82]. In the EEG data analysis for the brain, such clustering analysis, or more sophisticated versions of this, is known as microstate analysis [83,84]. Application of various conventional clustering methods used in microstate analysis to fMRI data has shown that, in terms of various features, the test-retest reliability of conventional clustering methods is not particularly behind the ELA on the same data sets [82]. This result is encouraging because the clustering methods avoid dichotomization of data. Furthermore, one can carry out many (albeit not all) measurements of energy landscapes also for clustering analysis. For example, one can count the transitions from one cluster to another and compute the dwell time in each cluster. If transitions from cluster C to cluster C' and vice versa are considerably rarer than typical cluster-to-cluster transitions, one can make an analogue to ELA to regard that C and C' are separated by a high energy barrier.

Hidden Markov models (HMMs) applied to multivariate time series share the goal of capturing the given data as transitory dynamics among discrete states. They see applications to, e.g., neuroimaging data [80,85,86]. The rate of transition from one state to another particular state and the dwell time in each state, for example, are features the original time series data as in the case of the conventional clustering algorithms. A key difference between the HMM and the conventional clustering is that the HMM assumes that the state transition rate depends on the currently visited state, whereas the conventional clustering approach does not. The HMM explicitly estimates the state transition rate for each pair of states. In contrast, in the clustering approach, one only estimates discrete clusters, or states, without looking at dynamics of the time series data. Then, one calculates the state transition rate as the normalized number of the state transitions (e.g., from state i to j) present in the data.

The ELA, or PMEM, is closer to the conventional clustering approach than to the HMM in this aspect. In estimating the PMEM from data, only the equilibrium distribution of the data is fitted, similar to conventional clustering. With the PMEM, one does not fit the state transition rate from each particular state to another. The state transition rates in the case of ELA are determined by the detailed balance condition (i.e., Eq (18)) and the structure of the estimated disconnectivity graph.

It may sound that the HMM is better than conventional clustering and the ELA because the HMM allows the explicit estimation of the individual state-transition rates dependent on the current state. However, the HMM usually has more parameters to be estimated, due to the state-dependent state-transition rate parameters, at least than clustering analyses. Therefore, if the given data are relatively short or the number of variables is large, it may be better to give up the HMM to run for conventional clustering or ELA. In addition, poor temporal resolution and strong autocorrelation of the data may also make the estimation of state-transition rates in the HMM difficult.

4. Software

We are aware of the following open-resource code packages for ELA and the related methods introduced in this article.

The energy landscape analysis toolbox (ELAT), accompanying a review paper [54], is a MATLAB package with a graphical user interface [87]. It can estimate the PMEM from data, compute the accuracy of fit, basin graphs in two and three dimensions (the three-dimensional version shows the energy value of each activity pattern on the vertical axis), disconnectivity graph, list of activity patterns at each local minimum (each being an N -dimensional binary vector, as in the gray-white vectors shown in Fig 7), time series of the activity pattern, time series of the basin ID to which the activity pattern at each time belongs, and some features of the constructed energy landscape. The code accompanying an MEG study [39] is provided in MATLAB [88]. It uses ELAT and provides analyses complementary to ELAT such as calculation of the best threshold for binarization and a bootstrapping for statistical analysis of energy values. A Python implementation of ELAT is available [89] as an accompaniment of [43]. Code for carrying out ELA is also available by other authors in R [90] and Mathematica [91], as an accompaniment of [44]. The code accompanying an EEG study [36] is also written in R and capable of estimating the PMEM by pseudo-likelihood maximization and carrying out ELA [92].

The code accompanying [20] is written in R and provides PMEM estimation by the standard maximum likelihood and Bayesian methods [93]. The code accompanying [32] is provided mostly in MATLAB and partly in R [94]. It allows us to estimate the standard PMEM, calculate significant local minima, run the variational Bayes method, and run test-retest reliability analyses. The code accompanying [19] is provided in MATLAB [95]. It also allows us to estimate the standard PMEM and provides part of code for the phase diagram methods explained in Section 3.1. Code for clustering-based estimation of discrete states as an analogue of ELA (see Section 3.5) is available [96] as an accompaniment of [82].

5. Applications

In this section, we showcase some applications of ELA.

In our first application, we studied energy landscapes of fMRI data obtained from healthy humans undergoing a visual bistable perception task [6]. We used $N = 7$ biologically informed locations in the brain, called regions of interest (ROI), as variables for constructing the PMEM. We constructed a single energy landscape from the fMRI data during the task. For the data pooled across the participants, we found ten local maxima and three heuristically determined major local maxima. In general, by inspecting the activity pattern at each major local minimum, one can see which ROI is relatively active (i.e., $\sigma_i(t) = 1$) or inactive (i.e., $\sigma_i(t) = -1$), based on which one can biologically characterize the major local minima. One major state was an aggregation of five local minima and their basins, and we named this major state the frontal-area state because the ROIs in the frontal area of the brain were active at the major local minimum defining this major state. Likewise, the visual-area state was composed of two other local minima and their basins. The visual-area state is complementary to the frontal-area state in that the ROIs active in the former are inactive in the latter and vice versa. The third major state, named the intermediate state, contained three local minima. These three major states remained the same for the PMEM separately estimated for the data from each single participant. In the bistable visual perception task, the participants were asked to watch a set of moving dots on the screen, which allows bistable perception (i.e., either of the two directions of the apparently moving sphere composed of dots). The participants reported whenever their subjective perception spontaneously switched from one direction to the other.

The main result was that participants with a larger basin size of the frontal-area state tended to have a shorter duration of the percept (i.e., how long the participant perceives one direction of motion of the visual stimulus before it subjectively switches to the other direction). In contrast, a larger basin size of the visual-area state tended to yield a longer duration of the percept. In this manner, we found an association between behavior (i.e., the median duration of the percept) and brain dynamics revealed by ELA (i.e., basin sizes), as schematically shown in Fig 3(g).

ELA also succeeded in identifying atypical brain mechanisms underlying prevalent neuropsychiatric conditions. One study on autism spectrum disorder (ASD) using fMRI data found that both high-functioning ASD adults and control individuals exhibited similar energy landscapes with the same two major states, each composed of two local minima that were the same between the ASD and control groups. The ASD adults showed atypically reduced indirect transitions between the two major states that went through one of the two local minima not belonging to a major state [14]. In particular, this specific transition frequency was so closely correlated with the severity of the ASD symptoms that it could be used for accurate diagnosis of autism. In a similar manner, another resting-state fMRI study showed that attention deficit hyperactivity disorder (ADHD) is characterized by atypical increases in overall transition probabilities between basins of local minima [29].

Multistability is often present in the composition of multispecies ecosystems, and different compositions may have different impacts on biodiversity, resistance to infection, and other functions. Suzuki et al. was among the first to apply the ELA based on the Ising model to outside neuroscience [44]. They analyzed gut microbiota data obtained from feces of the same six mice over the course of 72 weeks. They collected 18 data points per mouse, totaling $18 \times 6 = 108$ data points in each week in which the measurement was made. They tracked the changes in the energy landscape on $N = 8$ variables, which were the binarized abundance of bacterial species, as the mice aged. They found that the energy at the local minima gradually decreased (i.e., the local minima became more frequent) over the age and that the energy barrier between the two local minima with the smallest energy also decreased, making the transition between them easier for older mice.

6. Outlook

We have provided a tutorial of ELA based on the Ising model, with an emphasis on concepts and steps that are often confusing to users. We conclude this article by stating some possible directions for future research.

The problem of the data length, t_{\max} , is inherent in the ELA because t_{\max} required for a given accuracy seems to scale with 2^N [54]. Given this situation, effectiveness of the Bayesian methods (Section 3.4), which tactically combines the strength of pooling of the data across many time series to earn an effective t_{\max} and enable estimation of the Ising model for each time series, should be further investigated. Possible research includes applications to other data sets and developments of other variants of Bayes approximation methods. In a broader context, transfer learning for the PMEM [97,98] may help mitigate the data paucity problem. In particular, pooling of different data sets (e.g., similar data obtained from experiments carried out by different research groups) as well as pooling of different time series data within one data set may worth pursuing.

We provided a list of existing applications of the ELA in the introduction section. While ELA has been progressively used in other domains, most applications so far have been to neuroscience data, especially fMRI data, presumably reflecting the history of the method. However, we emphasize that the ELA is completely domain-general and can be applicable to any

multivariate time series data as long as there are at least about seven variables (i.e., $N \geq 7$) and t_{\max} is large enough to justify the choice of N . Because of the binarization step, ELA may be particularly useful when each variable shows a bimodal distribution over time, such as expression of some genes [99–102]. In fact, the fitting of the Ising model to neuronal spiking data [7–12], which preceded the ELA based on the Ising model, benefited from the observation that a neuron's state is approximately binary in the sense that it either spikes or it does not at any given time. With this remark in mind, applications of ELA to unforeseen types of multivariate time series data would be exciting.

In Section 3.1, we introduced a phase diagram analysis, which enables us to qualitatively tell properties of the given data set. If the system generating the observed time series data is changing on a slow timescale, which may be due to the environmental change (i.e., gradual pollution of an ecosystem) or the internal changes of the system (e.g., aging of the participant), the energy landscape of the system may transit from one phase to another (e.g., from the paramagnetic to spin-glass phase) on a slow timescale. A related challenge is to find early warning signals for large regime shifts of the system, which may occur as a result of bifurcations of the dynamical system underlying the data or for different reasons. Popular early warning signals include the sample variance and the decay rate of the sample autocorrelation function obtained from the observed time series, which would increase as the system approaches a bifurcation point [103–107]. A large regime shift that one wants to anticipate may correspond to a mass extinction of species in an ecosystem or a rapid progression of a disease in a human individual. From an energy landscape perspective, these methods usually assume that a large regime shift occurs when a local minimum in which the system's state is currently sitting is lost via the diminishment of the energy barrier between the currently visited basin and a neighboring basin; see Fig 1 in [103], Figs 2 and 3 in [104], and Box 1 in [106] for schematics. Then, the regime shift corresponds to the movement of the system's state from the current basin to the neighboring basin. Cross-fertilization of ELA and early warning signal research may be interesting.

ELA is not only a powerful tool for inferring and interpreting brain state dynamics but also for modifying neural dynamics and relevant behaviors. A recent study has developed a non-invasive closed-loop neural stimulation system using transcranial magnetic stimulation (TMS) and demonstrated that one can modify the height of a specific energy barrier and control the dynamics of perceptual consciousness in humans [37]. Also see [38]. Such intervention methods should be useful when we know desirable versus undesirable energy landscapes or when we can measure the effect of changing the energy landscape. While the TMS is considered to locally change h_i and hence σ_i for nodes i in the region of the brain affected by the TMS, other types of intervention, especially those that changes the connectivity between specific pairs of nodes, may be possible depending on the system [108]. Simulation studies are expected to be useful for understanding specificity of intervention [16,18,108]. Further studies along this line, including applications outside neuroimaging, are warranted.

Acknowledgements

We acknowledge Pitambar Khanra for technical information.

Author contributions

Conceptualization: Naoki Masuda.

Data curation: Saiful Islam, Si Thu Aung, Takamitsu Watanabe.

Funding acquisition: Naoki Masuda, Takamitsu Watanabe.

Investigation: Naoki Masuda, Saiful Islam, Si Thu Aung.

Methodology: Naoki Masuda.

Project administration: Naoki Masuda.

Supervision: Naoki Masuda.

Validation: Naoki Masuda, Saiful Islam, Si Thu Aung.

Visualization: Naoki Masuda, Saiful Islam, Si Thu Aung.

Writing – original draft: Naoki Masuda, Takamitsu Watanabe.

Writing – review & editing: Naoki Masuda, Saiful Islam, Si Thu Aung, Takamitsu Watanabe.

References

1. Onuchic JN, Luthey-Schulten Z, Wolynes PG. Theory of protein folding: the energy landscape perspective. *Annu Rev Phys Chem.* 1997;48:545–600. <https://doi.org/10.1146/annurev.physchem.48.1.545> PMID: 9348663
2. Wales DJ. Energy landscapes. Cambridge, UK: Cambridge University Press; 2003.
3. Wales DJ. The energy landscape as a unifying theme in molecular science. *Philos Trans A Math Phys Eng Sci.* 2005;363(1827):357–77. <https://doi.org/10.1098/rsta.2004.1497> PMID: 15664888
4. Röder K, Wales DJ. The energy landscape perspective: encoding structure and function for biomolecules. *Front Mol Biosci.* 2022;9:820792. <https://doi.org/10.3389/fmolb.2022.820792>
5. Watanabe T, Hirose S, Wada H, Imai Y, Machida T, Shirouzu I, et al. Energy landscapes of resting-state brain networks. *Front Neuroinform.* 2014;8:12. <https://doi.org/10.3389/fninf.2014.00012> PMID: 24611044
6. Watanabe T, Masuda N, Megumi F, Kanai R, Rees G. Energy landscape and dynamics of brain activity during human bistable perception. *Nat Commun.* 2014;5:4765. <https://doi.org/10.1038/ncomms5765> PMID: 25163855
7. Schneidman E, Berry MJ, Segev R, Bialek W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature.* 2006; 440(7087):1007–12. <https://doi.org/10.1038/nature04701> PMID: 16625187
8. Shlens J, Field GD, Gauthier JL, Grivich MI, Petrusca D, Sher A, et al. The structure of multi-neuron firing patterns in primate retina. *J Neurosci.* 2006;26(32):8254–66. <https://doi.org/10.1523/jneurosci.1282-06.2006> PMID: 16899720
9. Tang A, Jackson D, Hobbs J, Chen W, Smith JL, Patel H, et al. A maximum entropy model applied to spatial and temporal correlations from cortical networks in vitro. *J Neurosci.* 2008;28(2):505–18. <https://doi.org/10.1523/jneurosci.3359-07.2008> PMID: 18184793
10. Yu S, Huang D, Singer W, Nikolić D. A small world of neuronal synchrony. *Cereb Cortex.* 2008;18(12):2891–01. <https://doi.org/10.1093/cercor/bhn047> PMID: 18400792
11. Yeh FC, Tang A, Hobbs JP, Hottowy P, Dabrowski W, Sher A, et al. Maximum entropy approaches to living neural networks. *Entropy.* 2010;12(1):89–106. <http://dx.doi.org/10.3390/e12010089>
12. Mora T, Bialek W. Are biological systems poised at criticality? *J Stat Phys.* 2011;144(2):268–302. <https://doi.org/10.1007/s10955-011-0229-4>
13. Watanabe T, Hirose S, Wada H, Imai Y, Machida T, Shirouzu I, et al. A pairwise maximum entropy model accurately describes resting-state human brain networks. *Nat Commun.* 2013;4:1370. <https://doi.org/10.1038/ncomms2388> PMID: 23340410
14. Watanabe T, Rees G. Brain network dynamics in high-functioning individuals with autism. *Nat Commun.* 2017;8:16048. <https://doi.org/10.1038/ncomms16048> PMID: 28677689
15. Ashourvan A, Gu S, Mattar MG, Vettel JM, Bassett DS. The energy landscape underpinning module dynamics in the human brain connectome. *NeuroImage.* 2017;157:364–80. <https://doi.org/10.1016/j.neuroimage.2017.05.067> PMID: 28602945
16. Kang J, Pae C, Park HJ. Energy landscape analysis of the subcortical brain network unravels system properties beneath resting state dynamics. *NeuroImage.* 2017;149:153–64. <https://doi.org/10.1016/j.neuroimage.2017.01.075> PMID: 28159684

17. Ezaki T, Sakaki M, Watanabe T, Masuda N. Age-related changes in the ease of dynamical transitions in human brain activity. *Hum Brain Mapp.* 2018;39(6):2673–88. <https://doi.org/10.1002/hbm.24033> PMID: 29524289
18. Kang J, Pae C, Park HJ. Graph-theoretical analysis for energy landscape reveals the organization of state transitions in the resting-state human cerebral cortex. *PLoS ONE.* 2019;14(10):e0222161. <https://doi.org/10.1371/journal.pone.0222161> PMID: 31498822
19. Ezaki T, Fonseca dos Reis E, Watanabe T, Sakaki M, Masuda N. Closer to critical resting-state neural dynamics in individuals with higher fluid intelligence. *Commun Biol.* 2020;3:52. <https://doi.org/10.1038/s42003-020-0774-y> PMID: 32015402
20. Jeong SO, Kang J, Pae C, Eo J, Park SM, Son J, et al. Empirical Bayes estimation of pairwise maximum entropy model for nonlinear brain state dynamics. *NeuroImage.* 2021;244:118618. <https://doi.org/10.1016/j.neuroimage.2021.118618> PMID: 34571159
21. Kang J, Jeong SO, Pae C, Park HJ. Bayesian estimation of maximum entropy model for individualized energy landscape analysis of brain state dynamics. *Hum Brain Mapp.* 2021;42(10):3411–28. <https://doi.org/10.1002/hbm.25442> PMID: 33934421
22. Regonia PR, Takamura M, Nakano T, Ichikawa N, Fermin A, Okada G, et al. Modeling heterogeneous brain dynamics of depression and melancholia using energy landscape analysis. *Front Psychiatry.* 2021;12:780997. <https://doi.org/10.3389/fpsy.2021.780997> PMID: 34899435
23. Yamashita A, Rothlein D, Kucyi A, Valera EM, Esterman M. Brain state-based detection of attentional fluctuations and their modulation. *NeuroImage.* 2021;236:118072. <https://doi.org/10.1016/j.neuroimage.2021.118072> PMID: 33882346
24. Yamashita A, Rothlein D, Kucyi A, Valera EM, Germine L, Wilmer J, et al. Variable rather than extreme slow reaction times distinguish brain states during sustained attention. *Sci Rep.* 2021;11:14883. <https://doi.org/10.1038/s41598-021-94161-0> PMID: 34290318
25. Kondo HM, Terashima H, Ezaki T, Kochiyama T, Kihara K, Kawahara JI. Dynamic transitions between brain states predict auditory attentional fluctuations. *Front Neurosci.* 2022;16:816735. <https://doi.org/10.3389/fnins.2022.816735> PMID: 35368290
26. Fan L, Li C, Huang ZG, Zhao J, Wu X, Liu T, et al. The longitudinal neural dynamics changes of whole brain connectome during natural recovery from poststroke aphasia. *NeuroImage Clin.* 2022;36:103190. <https://doi.org/10.1016/j.nicl.2022.103190> PMID: 36174256
27. Li Y, An S, Zhou T, Su C, Zhang S, Li C, et al. Triple-network analysis of Alzheimer's disease based on the energy landscape. *Front Neurosci.* 2023;17:1171549. <https://doi.org/10.3389/fnins.2023.1171549> PMID: 37287802
28. Theis N, Bahuguna J, Rubin JE, Banerjee SS, Muldoon B, Prasad KM. Energy of functional brain states correlates with cognition in adolescent-onset schizophrenia and healthy persons. *Human Brain Mapp.* 2025;46:e70129. <https://doi.org/10.1002/hbm.70129> PMID: 39777939
29. Watanabe D, Watanabe T. Distinct frontoparietal brain dynamics underlying the co-occurrence of autism and ADHD. *eNeuro.* 2023; 10(5):ENEURO.0146-23.2023. <https://doi.org/10.1523/eneuro.0146-23.2023> PMID: 37414550
30. Hosaka Y, Hieda T, Hayashi K, Jimura K, Matsui T. Linear models replicate the energy landscape and dynamics of resting-state brain activity. *bioRxiv.* preprint. 2024 [posted Sep 3, 2024; cited Nov 12, 2024] Available from: <https://doi.org/10.1101/2024.05.21.595246> doi: 10.1101/2024.05.21.595246
31. Ishida T, Yamada S, Yasuda K, Uenishi S, Tamaki A, Tabata M, et al. Aberrant brain dynamics of large-scale functional networks across schizophrenia and mood disorder. *NeuroImage Clin.* 2024;41:103574. <https://doi.org/10.1016/j.nicl.2024.103574> PMID: 38346380
32. Khanra P, Nakuci J, Muldoon S, Watanabe T, Masuda N. Reliability of energy landscape analysis of resting-state functional MRI data. *Eur J Neurosci.* 2024;60(12):4265–90. <https://doi.org/10.1111/ejn.16390> PMID: 38837814
33. Kindler J, Ishida T, Michel C, Klaassen AL, Stübke M, Zimmermann N, et al. Aberrant brain dynamics in individuals with clinical high risk of psychosis. *Schizophr Bull Open.* 2024;5(1):sgae002. <https://doi.org/10.1093/schizbullopen/sgae002> PMID: 38605980
34. Miyata J, Sasamoto A, Ezaki T, Isobe M, Kochiyama T, Masuda N, et al. Associations of conservatism and jumping to conclusions biases with aberrant salience and default mode network. *Psychiatry Clin Neurosci.* 2024;78(4):322–31. <https://doi.org/10.1111/pcn.13652> PMID: 38414202
35. Xing L, Guo Z, Long Z. Energy landscape analysis of brain network dynamics in Alzheimer's disease. *Front Aging Neurosci.* 2024;16:1375091. <https://doi.org/10.3389/fnagi.2024.1375091> PMID: 38813531
36. Klepl D, He F, Wu M, De Marco M, Blackburn DJ, Sarrigiannis PG. Characterising Alzheimer's disease with EEG-based energy landscape analysis. *IEEE J Biomed Health Inform.* 2022;26(3):992–1000. <https://doi.org/10.1109/jbhi.2021.3105397> PMID: 34406951

37. Watanabe T. Causal roles of prefrontal cortex during spontaneous perceptual switching are determined by brain state dynamics. *eLife*. 2021;10:e69079. <https://doi.org/10.7554/eLife.69079> PMID: 34713803
38. Gupta D, Du X, Summerfelt A, Hong LE, Choa FS. Brain connectivity signature extractions from TMS invoked EEGs. *Sensors*. 2023;23(10):4078. <https://doi.org/10.3390/s23084078> PMID: 37112420
39. Krzemiński D, Masuda N, Hamandi K, Singh KD, Routley B, Zhang J. Energy landscape of resting magnetoencephalography reveals fronto-parietal network impairments in epilepsy. *Netw Neurosci*. 2020; 4(2):374–396. https://doi.org/10.1162/netn_a_00125 PMID: 32537532
40. Gu S, Cieslak M, Baird B, Muldoon SF, Grafton ST, Pasqualetti F, et al. The energy landscape of neurophysiological activity implicit in brain network structure. *Sci Rep*. 2018;8:2507. <https://doi.org/10.1038/s41598-018-20123-8> PMID: 29410486
41. Song E. Energy landscape analysis of cardiac fibrillation wave dynamics using pairwise maximum entropy model. *arXiv preprint*. arXiv:1809.09798. 2018 [cited Nov 12, 2024]. Available from: <https://arxiv.org/abs/1809.09798>
42. Yamamoto K, Sakaguchi M, Onishi A, Yokoyama S, Matsui Y, Yamamoto W, et al. Energy landscape analysis and time-series clustering analysis of patient state multistability related to rheumatoid arthritis drug treatment: the KURAMA cohort study. *PLoS ONE*. 2024;19(4):e0302308. <https://doi.org/10.1371/journal.pone.0302308> PMID: 38709812
43. Yonezawa S, Haruki T, Koizumi K, Taketani A, Oshima Y, Oku M, et al. Establishing monoclonal gammopathy of undetermined significance as an independent pre-disease state of multiple myeloma using Raman spectroscopy, dynamical network biomarker theory, and energy landscape analysis. *Int J Mol Sci*. 2024;25(3):1570. <https://doi.org/10.3390/ijms25031570> PMID: 38338848
44. Suzuki K, Nakaoka S, Fukuda S, Masuya H. Energy landscape analysis elucidates the multistability of ecological communities across environmental gradients. *Ecol Monogr*. 2021;91(2):e01469. <http://dx.doi.org/10.1002/ecm.1469>
45. Fujita H, Yoshida S, Suzuki K, Toju H. Alternative stable states of microbiome structure and soil ecosystem functions. *Environ Microbiome*. 2025;20:28. <https://doi.org/10.1186/s40793-025-00688-4> PMID: 40050988
46. Fujita H, Ushio M, Suzuki K, Abe MS, Yamamichi M, Iwayama K, et al. Alternative stable states, nonlinear behavior, and predictability of microbiome dynamics. *Microbiome*. 2023;11:63. <https://doi.org/10.1186/s40168-023-01474-5> PMID: 36978146
47. Miyamoto H, Kawachi N, Kurotani A, Moriya S, Suda W, Suzuki K, et al. Computational estimation of sediment symbiotic bacterial structures of seagrasses overgrowing downstream of onshore aquaculture. *Environ Res*. 2023;219:115130. <https://doi.org/10.1016/j.envres.2022.115130> PMID: 36563976
48. Okada S, Inabu Y, Miyamoto H, Suzuki K, Kato T, Kurotani A, et al. Estimation of silent phenotypes of calf antibiotic dysbiosis. *Sci Rep*. 2023;13:6359. <https://doi.org/10.1038/s41598-023-33444-0> PMID: 37076584
49. Kim KY, Wang J. Potential energy landscape and robustness of a gene regulatory network: toggle switch. *PLoS Comput Biol*. 2007;3(3):e60. <https://doi.org/10.1371/journal.pcbi.0030060> PMID: 17397255
50. Marvel SA, Strogatz SH, Kleinberg JM. Energy landscape of social balance. *Phys Rev Lett*. 2009;103(19):198701. <https://doi.org/10.1103/physrevlett.103.198701> PMID: 20365960
51. Lv C, Li X, Li F, Li T. Energy landscape reveals that the budding yeast cell cycle is a robust and adaptive multi-stage process. *PLoS Comput Biol*. 2015;11(4):e1004156. <https://doi.org/10.1371/journal.pcbi.1004156> PMID: 25794282
52. Prada-Gracia D, Gómez-Gardeñes J, Echenique P, Falo F. Exploring the free energy landscape: from dynamics to networks and back. *PLoS Comput Biol*. 2009;5(6):e1000415. <https://doi.org/10.1371/journal.pcbi.1000415> PMID: 19557191
53. Andersson E, Sjö M, Kaji K, Olariu V. CELLoGeNe—an energy landscape framework for logical networks controlling cell decisions. *iScience*. 2022;25(1):104743. <https://doi.org/10.1016/j.isci.2022.104743> PMID: 35942105
54. Ezaki T, Watanabe T, Ohzeki M, Masuda N. Energy landscape analysis of neuroimaging data. *Philos Trans R Soc A*. 2017;375(2096):20160287. <https://doi.org/10.1098/rsta.2016.0287> PMID: 28507232
55. Stein RR, Marks DS, Sander C. Inferring pairwise interactions from biological data using maximum-entropy probability models. *PLoS Comput Biol*. 2015;11(4):e1004182. <https://doi.org/10.1371/journal.pcbi.1004182> PMID: 26225866
56. Nguyen HC, Zecchina R, Berg J. Inverse statistical problems: from the inverse Ising problem to data science. *Adv Phys*. 2017;66(3):197–261. <http://dx.doi.org/10.1080/00018732.2017.1341604>

57. Carleo G, Cirac I, Cranmer K, Daudet L, Schuld M, Tishby N, et al. Machine learning and the physical sciences. *Rev Mod Phys.* 2019;91(4):045002. <http://dx.doi.org/10.1103/RevModPhys.91.045002>
58. Mehta P, Bukov M, Wang CH, Day AGR, Richardson C, Fisher CK, et al. A high-bias, low-variance introduction to Machine Learning for physicists. *Phys Rep.* 2019;810:1–124. <https://doi.org/10.1016/j.physrep.2019.03.001> PMID: 31404441
59. Becker OM, Karplus M. The topology of multidimensional potential energy surfaces: theory and application to peptide structure and kinetics. *J Chem Phys.* 1997;106(4):1495–517.
60. Darroch JN, Ratcliff D. Generalized iterative scaling for log-linear models. *Ann Math Stat.* 1972;43(5):1470–80.
61. Besag J. Statistical analysis of non-lattice data. *J R Stat Soc Ser D.* 1975;24(3):179–95. <https://doi.org/10.2307/2987782>
62. Roudi Y, Aurell E, Hertz JA. Statistical physics of pairwise probability models. *Front Comput Neurosci.* 2009;3:22. <https://doi.org/10.3389/neuro.10.022.2009> PMID: 19949460
63. Masuda N, Boyd ZM, Garlaschelli D, Mucha PJ. Correlation networks: interdisciplinary approaches beyond thresholding. *arXiv preprint. arXiv:2311.09536.* 2023 [cited Nov 12, 2024]. Available from: <https://arxiv.org/abs/2311.09536>
64. Gordon EM, Laumann TO, Gilmore AW, Newbold DJ, Greene DJ, Berg JJ, et al. Precision functional mapping of individual human brains. *Neuron.* 2017;95(4):791–807. <https://doi.org/10.1016/j.neuron.2017.07.011> PMID: 28757305
65. Wales DJ, Miller MA, Walsh TR. Archetypal energy landscapes. *Nature.* 1998; 394(6695):758–60. <http://dx.doi.org/10.1038/29487>
66. Zhou Q. Random walk over basins of attraction to construct Ising energy landscapes. *Phys Rev Lett.* 2011;106(18):180602. <https://doi.org/10.1103/physrevlett.106.180602> PMID: 21635078
67. Fischer KH, Hertz JA. *Spin glasses.* Cambridge, UK: Cambridge University Press; 1991.
68. Bray AJ, Moore MA. Chaotic nature of the spin-glass phase. *Phys Rev Lett.* 1987;58(1):57–60. <https://doi.org/10.1103/physrevlett.58.57> PMID: 10034289
69. Rizzo T, Crisanti A. Chaos in temperature in the Sherrington-Kirkpatrick model. *Phys Rev Lett.* 2003;90(13):137201. <https://doi.org/10.1103/physrevlett.90.137201> PMID: 12689320
70. Aspelmeier T. Free-energy fluctuations and chaos in the Sherrington-Kirkpatrick model. *Phys Rev Lett.* 2008;100(11):117205. <https://doi.org/10.1103/physrevlett.100.117205> PMID: 18517821
71. Tkačik G, Mora T, Marre O, Amodei D, Palmer SE, Berry MJ II, et al. Thermodynamics and signatures of criticality in a network of neurons. *Proc Natl Acad Sci U S A.* 2015;112(37):11508–13. <https://doi.org/10.1073/pnas.1514188112> PMID: 26330611
72. Watanabe T, Kan S, Koike T, Misaki M, Konishi S, Miyauchi S, et al. Network-dependent modulation of brain activity during sleep. *NeuroImage.* 2014;98:1–10. <https://doi.org/10.1016/j.neuroimage.2014.04.079> PMID: 24814208
73. Theiler J, Eubank S, Longtin A, Galdrikian B, Farmer JD. Testing for nonlinearity in time series: the method of surrogate data. *Phys D.* 1992;58(1-4):77–94. [https://doi.org/10.1016/0167-2789\(92\)90102-S](https://doi.org/10.1016/0167-2789(92)90102-S)
74. Schreiber T, Schmitz A. Surrogate time series. *Phys D.* 2000;142(3-4):346–82. [https://doi.org/10.1016/S0167-2789\(00\)00043-9](https://doi.org/10.1016/S0167-2789(00)00043-9)
75. Walter JC, Barkema GT. An introduction to Monte Carlo methods. *Phys A.* 2015;418:78–87. <https://doi.org/10.1016/j.physa.2014.06.014>
76. Masuda N, Porter MA, Lambiotte R. Random walks and diffusion on networks. *Phys Rep.* 2017;716-717:1–58. <https://doi.org/10.1016/j.physrep.2017.07.007>
77. Bishop CM. *Pattern recognition and machine learning.* New York, NY: Springer; 2006.
78. Liu X, Chang C, Duyn JH. Decomposition of spontaneous brain activity into distinct fMRI co-activation patterns. *Front Syst Neurosci.* 2013;7:101. <https://doi.org/10.3389/fnsys.2013.00101> PMID: 24550788
79. Liu X, Zhang N, Chang C, Duyn JH. Co-activation patterns in resting-state fMRI signals. *NeuroImage.* 2018;180:485–94. <https://doi.org/10.1016/j.neuroimage.2018.01.041> PMID: 29355767
80. Ezaki T, Himeno Y, Watanabe T, Masuda N. Modelling state-transition dynamics in resting-state brain signals by the hidden Markov and Gaussian mixture models. *Eur J Neurosci.* 2021;54(8):5404–16. <https://doi.org/10.1111/ejn.15386> PMID: 34250639
81. Paakki JJ, Rahko JS, Kotila A, Mattila ML, Miettunen H, Hurtig TM, et al. Co-activation pattern alterations in autism spectrum disorder—a volume-wise hierarchical clustering fMRI study. *Brain Behav.* 2021;11(8):e02174. <https://doi.org/10.1002/brb3.2174> PMID: 33998178

82. Islam S, Khanra P, Nakuci J, Muldoon SF, Watanabe T, Masuda N. State-transition dynamics of resting-state functional magnetic resonance imaging data: model comparison and test-to-retest analysis. *BMC Neurosci.* 2024;25:14. <https://doi.org/10.1186/s12868-024-00854-3> PMID: 38438838
83. Khanna A, Pascual-Leone A, Michel CM, Farzan F. Microstates in resting-state EEG: current status and future directions. *Neurosci Biobehav Rev.* 2015;49:105–13. <https://doi.org/10.1016/j.neubiorev.2014.12.010> PMID: 25526823
84. Michel CM, Koenig T. EEG microstates as a tool for studying the temporal dynamics of whole-brain neuronal networks: a review. *NeuroImage.* 2018;180:577–93. <https://doi.org/10.1016/j.neuroimage.2017.11.062>
85. Baker AP, Brookes MJ, Rezek IA, Smith SM, Behrens T, Probert Smith PJ, et al. Fast transient networks in spontaneous human brain activity. *eLife.* 2014; 3:e01867. <https://doi.org/10.7554/elife.01867> PMID: 24668169
86. Tait L, Zhang J. +microstate: a MATLAB toolbox for brain microstate analysis in sensor and cortical EEG/MEG. *NeuroImage.* 2022;258:119346. <https://doi.org/10.1016/j.neuroimage.2022.119346> PMID: 35660463
87. Ezaki T. Energy Landscape Analysis Toolkit (ELAT) for MATLAB. GitHub. 2017 [cited Nov 12, 2024]. <https://github.com/tkEzaki/energy-landscape-analysis>.
88. Krzemiński D. Energy landscape scripts. GitHub. 2020 [cited Nov 11, 2024]. https://github.com/dokato/energy_landscape.
89. Oku M. About elapy. GitHub. 2024 [cited Nov 11, 2024]. <https://github.com/okumakito/elapy>.
90. Suzuki K. rELA—an R package for energy landscape analysis. GitHub. 2024 [cited May 30, 2024]. <https://github.com/kecosz/rELA>.
91. Suzuki K. ela. GitHub. 2024 [cited May 30, 2024]. <https://github.com/kecosz/ela>.
92. Klepl D. Alzheimer-s-disease-Energy-landscape. GitHub. 2022 [cited Aug 30, 2024]. <https://github.com/dominikklepl/AD-energy-landscape/tree/master/scripts>.
93. monetyonsei. VEM_MEM. GitHub. 2021 [cited Aug 30, 2024]. https://github.com/monetyonsei/paper/tree/main/VEM_MEM.
94. Khanra P. Energy landscape analysis with the conventional likelihood maximization and variational Bayes approximation methods. GitHub. 2024 [cited May 10, 2024]. https://github.com/pitambarkhanra/energy_landscape_analysis.
95. Ezaki T. MATLAB code for phase diagram analysis. GitHub. 2020 [cited Nov 12, 2024]. https://github.com/tkEzaki/phase_diagram_analysis.
96. Islam S. State-transition dynamics of resting state fMRI data. GitHub. 2024 [cited Nov 11, 2024]. https://github.com/sislam99/fmri_state_transition_dynamics.
97. Yosinski J, Clune J, Bengio Y, Lipson H. How transferable are features in deep neural networks? *Adv Neural Inf Process Syst.* 2014;27:3320–8.
98. Zen R, My L, Tan R, Hébert F, Gattobigio M, Miniatura C, et al. Transfer learning for scalability of neural-network quantum states. *Phys Rev E.* 2020;101(5):053301. <https://doi.org/10.1103/physreve.101.053301> PMID: 32575207
99. Gormley M, Tozeren A. Expression profiles of switch-like genes accurately classify tissue and infectious disease phenotypes in model-based classification. *BMC Bioinform.* 2008;9:486. <https://doi.org/10.1186/1471-2105-9-486> PMID: 19014681
100. Eldar A, Elowitz MB. Functional roles for noise in genetic circuits. *Nature.* 2010;467(7312):167–73. <https://doi.org/10.1038/nature09326> PMID: 20829787
101. Mason CC, Hanson RL, Ossowski V, Bian L, Baier LJ, Krakoff J, et al. Bimodal distribution of RNA expression levels in human skeletal muscle tissue. *BMC Genomics.* 2011;12:98. <https://doi.org/10.1186/1471-2164-12-98> PMID: 21299892
102. Aqil A, Li Y, Wang Z, Islam S, Russel M, Kallak TK, et al. Switch-like gene expression modulates disease susceptibility. *bioRxiv. preprint.* 2024. doi:10.1101/2024.08.24.609537.
103. Scheffer M, Bascompte J, Brock WA, Brovkin V, Carpenter SR, Dakos V, et al. Early-warning signals for critical transitions. *Nature.* 2009;461(7260):53–59. <http://dx.doi.org/10.1038/nature08227>
104. Scheffer M, Carpenter SR, Lenton TM, Bascompte J, Brock W, Dakos V, et al. Anticipating critical transitions. *Science.* 2012;338(6105):344–8. <https://doi.org/10.1126/science.1225244> PMID: 23087241
105. Boettiger C, Ross N, Hastings A. Early warning signals: the charted and uncharted territories. *Theor Ecol.* 2013;6(3):255–64. <https://doi.org/10.1007/s12080-013-0192-6>

106. Dakos V, Carpenter SR, van Nes EH, Scheffer M. Resilience indicators: prospects and limitations for early warnings of regime shifts. *Philos Trans R Soc B*. 2015;370(1659):20130263. <http://dx.doi.org/10.1098/rstb.2013.0263>
107. Southall E, Brett TS, Tildesley MJ, Dyson L. Early warning signals of infectious disease transitions: a review. *J R Soc Interface*. 2021;18(181):20210555. <http://dx.doi.org/10.1098/rsif.2021.0555>
108. Park HJ, Kang J. A computational framework for controlling the self-restorative brain based on the free energy and degeneracy principles. *Front Comput Neurosci*. 2021;15:590019. <https://doi.org/10.3389/fncom.2021.590019>