

RESEARCH ARTICLE

Improved Measures of Integrated Information

Max Tegmark^{1,2*}

1 Department of Physics & MIT Kavli Institute, Massachusetts Institute of Technology, Cambridge, Massachusetts, United States of America, **2** Theiss Research, La Jolla, California, United States of America

* tegmark@mit.edu

Abstract

Although there is growing interest in measuring integrated information in computational and cognitive systems, current methods for doing so in practice are computationally unfeasible. Existing and novel integration measures are investigated and classified by various desirable properties. A simple taxonomy of Φ -measures is presented where they are each characterized by their choice of factorization method (5 options), choice of probability distributions to compare (3×4 options) and choice of measure for comparing probability distributions (7 options). When requiring the Φ -measures to satisfy a minimum of attractive properties, these hundreds of options reduce to a mere handful, some of which turn out to be identical. Useful exact and approximate formulas are derived that can be applied to real-world data from laboratory experiments without posing unreasonable computational demands.



OPEN ACCESS

Citation: Tegmark M (2016) Improved Measures of Integrated Information. PLoS Comput Biol 12(11): e1005123. doi:10.1371/journal.pcbi.1005123

Editor: Anil Seth, University of Sussex, UNITED KINGDOM

Received: January 9, 2016

Accepted: August 27, 2016

Published: November 21, 2016

Copyright: © 2016 Max Tegmark. 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.

Data Availability Statement: This paper is a theoretical work and therefore uses no data.

Funding: This research was supported by ARO grant W911NF-15-1-0300. The funder had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The author has declared that no competing interests exist.

Author Summary

How can one determine whether an unresponsive patient is conscious or not? Of all the information processing in your brain that can be measured with modern sensors, which corresponds to information that you are subjectively aware of and which is unconscious? A theory that has garnered much recent attention proposes that the answer involves measuring a quantity called *integration* that quantifies the extent to which information is interconnected into a unified whole rather than split into disconnected parts. Unfortunately, proposed measures of integration are too slow to compute in practice from patient data. In this paper, I explore and classify existing and novel integration measures by various desirable properties, and derive useful exact and approximate formulas that can be applied to real-world data from laboratory experiments without posing unreasonable computational demands. This improves the prospects of making fascinating questions and theories about consciousness experimentally testable.

Introduction

What makes an information-processing system conscious in the sense of having a subjective experience? Although many scientists used to view this topic as beyond the reach of science, the study of Neural Correlates of Consciousness (NCCs) has become quite mainstream in the neuroscience community in recent years—see, e.g., [1, 2]. To move beyond correlation to

causation [3], neuroscientists have begun searching for a theory of consciousness that can predict what physical phenomena cause consciousness (defined as subjective experience [3]) to occur. Dehaene [4] reviews a number of candidate theories currently under active discussion, including the *Nonlinear Ignition* model (NI) [5, 6], the *Global Neuronal Workspace* (GNW) model [7–9] and *Integrated Information Theory* (IIT) [10, 11]. Rapid progress in artificial intelligence is further fueling interest in such theories and how they can be generalized to apply not only to biological systems, but also to engineered systems such as computers and robots and ultimately arbitrary arrangements of elementary particles [12].

Although there is still no consensus on necessary and sufficient conditions for a physical system to be conscious, there is broad agreement that it needs to be able to store and process information in a way that is somehow *integrated*, not consisting of nearly independent parts. As emphasized by Tononi [10], it must be impossible to decompose a conscious system into nearly independent parts—otherwise these parts would feel like two separate conscious entities. While integration as a *necessary* condition for consciousness is rather uncontroversial, IIT goes further and makes the bold and controversial claim that it is also a *sufficient* condition for consciousness, using an elaborate mathematical integration definition [11].

As neuroscience data improves in quantity and quality, it is timely to resolve this controversy by testing the many experimental predictions that IIT makes [11] with state-of-the-art laboratory measurements. Unfortunately, such tests have been hampered by the fact that the integration measure proposed by IIT is computationally infeasible to evaluate for large systems, growing super-exponentially with the system's information content. This has led to the development of various alternative integration measures that are simpler to compute or have other desirable properties. For example, Barrett & Seth [13] proposed an attractive integration measure that is easier to compute from neuroscience data, but whose interpretation is complicated by the fact that it can be negative in some cases [14, 15]. [16] used an integration measure inspired by complexity theory to successfully predict who was conscious in a sample including patients who were awake, in deep sleep, dreaming, sedated and with locked-in syndrome. [17] suggest that state transition entropy correlates with consciousness. Griffith & Koch have proposed defining integration of a system as the synergistic information that its parts have about the future, which appears promising although there does not yet exist a unique formula for it [18]. Even the team behind IIT has updated their integration measure twice through successive refinements of their theory [10, 11]. Despite these definitional and computational challenges, interest in measuring integration is growing, not only in neuroscience but also in other fields, ranging from physics [12] and evolution [19] to the study of collective intelligence in social networks [20].

It is therefore interesting and timely to do a comprehensive investigation of existing and novel integration measures, classifying them by various desirable properties. This is the goal of the present paper, as summarized in Tables 1 and 2. In the Methods section, we investigate general integration measures and their properties. In the Results section, we first present our taxonomy of integration measures, then derive useful formulas for many of these measures that can be applied to the sort of time-series data that is typically measured in laboratory experiments with continuous variables, and finally explore further algorithmic speedups and approximations. We summarize our conclusions in the Discussion session.

Methods

In this section, we present our methods for building a taxonomy of integration measures.

Following Tononi [10], we will use the symbol Φ to denote integrated information. All measures of Φ aim to quantify the extent to which a system is interconnected, yielding $\Phi = 0$ if the

Table 1. Properties of different integration measures. All but the third are desirable properties; capitalized N/Y (no/yes) indicate when an integration measure lacks a desirable property or has an undesirable one. The first four properties are generally agreed to be important, while the second set of four have been argued to be important by some authors. Interpretability refers to the extent to which the measure can be given an information-theoretic interpretation satisfying desirable properties of integration (see text). Computability refers to the feasibility of evaluating the measure in practice (see text).

	$\phi^{2.5}$	$\phi^{2.5'}$	$\phi^{2.5''}$	$\phi^{3.0}$	ϕ^M	ϕ^B	ϕ^{MD}	ϕ_{kk}^M	ϕ^{oak}	ϕ^{opk}	ϕ^{ots}	ϕ^{ofu}	ϕ^{nas}	ϕ^{mas}	ϕ^{xfk}	
Major	Always non-negative	y	y	y	y	N	y	y	y	y	y	y	y	y	y	
	Always finite even for ∞ -dimensional system	N	y	y	N	y	y	y	y	y	y	y	N	y	y	
	Vanishes for deterministic system (drawback)	n	n	n	n	n	n	n	n	n	n	n	n	n	Y	
	Vanishes for separable system	y	y	y	y	y	N	y	y	y	y	N	y	y	y	
Minor	Vanishes for afferent system	y	y	y	y	N	N	N	N	y	N	N	N	y	y	N
	Vanishes for efferent system	y	y	y	y	N	N	N	N	N	y	N	N	N	N	N
	State-dependent	y	y	y	y	N	N	N	y	y	y	N	N	N	N	y
	Based on symmetric probability distance	N	N	N	y	N	N	N	N	N	N	N	N	N	N	N
Intuitively interpretable	2	2	2	2	2	0	2	2	2	2	0	1	0	0	0	
Computationally tractable	1	2	2	0	2	2	2	2	2	2	2	2	1	2	2	

doi:10.1371/journal.pcbi.1005123.t001

system consists of two independent parts, and a larger Φ the more the parts affect each other. Mathematically, all Φ -measures are defined in a two-step process:

1. Given an imaginary cut that partitions the system into two parts, define a measure ϕ of how much these two parts affect each other. Table 2 lists many ϕ -options.
2. Define Φ as the ϕ -value for the “cruellest cut” that minimizes ϕ . A major numerical challenge is that the number of cuts to be minimized over grows super-exponentially with the number of bits in the system. A further challenge in this step is how to best handle cuts splitting the system into parts of unequal size.

Note that our analysis is focused *only* on integration, not on consciousness; besides integration, a true measure of consciousness may involve additional requirements that this paper does not consider. For example, Scott Aaronson has criticized in a widely read blog post the claim that integration is a sufficient condition for consciousness, and IIT discusses postulates including cause-effect power, composition and exclusion [11].

Before delving into the many different options for defining Φ , let us first introduce convenient notation general enough to describe all proposed integration measures, as illustrated in Fig 1.

Interpreting evolution as a Markov process

Consider two random vectors \mathbf{x}_0 and \mathbf{x}_1 whose joint probability distribution is $p(\mathbf{x}_0, \mathbf{x}_1)$. We will interpret them as the state of a time-dependent system $\mathbf{x}(t)$ at two separate times t_0 and t_1 . For example, if these are two vectors of 5 bits each, then p is a table of 2^{10} numbers giving the probability of each possible bit string, while if these are two vectors in 3D space, then p is a function of 6 real continuous variables. We obtain the marginal distribution $p^{(n)}(\mathbf{x}_n)$ for the n^{th} vector, where $n = 0$ or $n = 1$, by summing/integrating p over the other vector.

Below we will often find it convenient to denote these vectors as single indices $i = \mathbf{x}_0$ and $j = \mathbf{x}_1$. For example, this allows us to write the marginal distribution $p_0(\mathbf{x}_0)$ as $\sum_j p_{ij}$, where the sum over j is to be interpreted as summation/integration over all allowed values of \mathbf{x}_1 . We also adopt the notation where replacing an index by a dot means that this index is to be summed/integrated over. This lets us write the marginal distributions $p^{(0)}(\mathbf{x}_0)$ and $p^{(1)}(\mathbf{x}_1)$ as

$$p_i^{(0)} = p_{i\cdot} \quad \text{and} \quad p_j^{(1)} = p_{\cdot j} \tag{1}$$

Table 2. Integration ϕ for different measures. $\mathbf{A} \equiv \mathbf{B}^t \mathbf{C}^{-1}$, $\mathbf{A}_b \equiv \mathbf{B} \mathbf{C}^{-1} = \Sigma \mathbf{A}^t \Sigma_b^{-1}$, $\Sigma \equiv \mathbf{C} - \mathbf{B}^t \mathbf{C}^{-1} \mathbf{B} = \mathbf{C} - \mathbf{A} \mathbf{C} \mathbf{A}^t$ and $\Sigma_b \equiv \mathbf{C} - \mathbf{B} \mathbf{C}^{-1} \mathbf{B}^t = \mathbf{C} - \mathbf{A}_b \mathbf{C} \mathbf{A}_b^t = [\mathbf{C}^{-1} + \mathbf{A}^t \Sigma^{-1} \mathbf{A}]^{-1} = \mathbf{C} - \mathbf{C} \mathbf{A}^t \mathbf{C}^{-1} \mathbf{A} \mathbf{C}$. \mathbf{C} is the data covariance matrix and \mathbf{B} is the cross-covariance between different times as defined by eq (46).

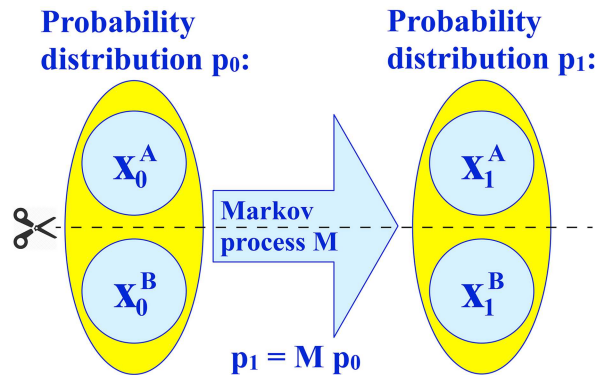
Name	Definition	Formula for Gaussian variables
$\phi^{otu} (\phi^M)$	$I(\mathbf{x}^A, \mathbf{x}^B) - I(\mathbf{x}_0^A, \mathbf{x}_0^B)$	$\frac{1}{2} \log \frac{ \mathbf{T}_A \mathbf{T}_B \mathbf{C} }{ \mathbf{T} \mathbf{C}_A \mathbf{C}_B } = \frac{1}{2} \log \frac{ \Sigma_A \Sigma_B }{ \Sigma }$
ϕ^B	$I(\mathbf{x}_0, \mathbf{x}_1) - I(\mathbf{x}_0^A, \mathbf{x}_1^A) - I(\mathbf{x}_0^B, \mathbf{x}_1^B)$	$\frac{1}{2} \log \frac{ \mathbf{C} ^2 \mathbf{T}_A \mathbf{T}_B }{ \mathbf{T} \mathbf{C}_A ^2 \mathbf{C}_B ^2} = \frac{1}{2} \log \frac{ \mathbf{C} \Sigma_A \Sigma_B }{ \Sigma \mathbf{C}_A \mathbf{C}_B }$
$\phi^{otum} (\phi^{MD})$	$d_{MD}(p, q), \quad q_{ij j'} = \frac{p_{ij} \cdot p_{i'j} p_{j'j'}}{p_{i'j} \cdot p_{j'j'}}$	$\frac{1}{2} \left[\ln \frac{ \hat{\mathbf{C}} }{ \Sigma } + \text{tr} \left(\hat{\mathbf{C}}^{-1} \mathbf{C} - \hat{\Sigma}^{-1} \Sigma - \hat{\mathbf{A}}^t \hat{\Sigma}^{-1} \hat{\mathbf{A}} \mathbf{C} \right) \right]$ for $\beta = 1, \hat{\mathbf{C}} \equiv \hat{\mathbf{A}} \mathbf{C} \hat{\mathbf{A}}^t + \hat{\Sigma}$
ϕ^{ots}	$I(\mathbf{x}^A, \mathbf{x}^B)$	$\frac{1}{2} \log \frac{ \mathbf{T}_A \mathbf{T}_B }{ \mathbf{T} }$
ϕ^{ofs}	$I(\mathbf{x}_1^A, \mathbf{x}_1^B)$	$\frac{1}{2} \log \frac{ \mathbf{C}_A \mathbf{C}_B }{ \mathbf{C} }$
ϕ^{ofu}	$-\sum_{j j'} p_{\cdot j j'} \log \sum_{i j'} \frac{p_{ij} p_{i'j} p_{j'j'}}{p_{i'j} \cdot p_{j'j'}}$	$\frac{1}{2} \left[\ln \frac{ \mathbf{C}_q }{ \mathbf{C} } + \text{tr} \mathbf{C}_q^{-1} \mathbf{C} - n \right], \mathbf{C}_q \equiv \hat{\mathbf{A}} \mathbf{C} \hat{\mathbf{A}}^t + \hat{\Sigma}$
$\phi_{kk'}^{ofk} (\phi_{kk'}^M)$	$\sum_{j j'} \frac{p_{kk'j}}{p_{kk'}} \log \frac{p_{kk'j j'} p_{k'j} p_{k'j'}}{p_{kk'} \cdot p_{k'j} p_{k'j'}}$	$\frac{1}{2} \left[\mathbf{x}_0^t (\hat{\mathbf{A}} - \mathbf{A}) \hat{\Sigma}^{-1} (\hat{\mathbf{A}} - \mathbf{A}) \mathbf{x}_0 + \ln \frac{ \Sigma_A \Sigma_B }{ \Sigma } + \text{tr} \hat{\Sigma}^{-1} \Sigma - n \right]$
$\phi_{kk'}^{oak}$	$\sum_j \frac{p_{kk'j}}{p_{kk'}} \log \frac{p_{kk'j} p_{k'j}}{p_{kk'} \cdot p_{k'j}}$	$\frac{1}{2} \left[\Delta \mathbf{m}^t \tilde{\Sigma}_A^{-1} \Delta \mathbf{m} + \ln \frac{ \tilde{\Sigma}_A }{ \Sigma_A } + \text{tr} \tilde{\Sigma}_A^{-1} \Sigma_A - n_A \right], \Delta \mathbf{m} \equiv (\mathbf{A}_A - \hat{\mathbf{A}}_A) \mathbf{x}_0^A + \mathbf{A}_B \mathbf{x}_0^B$
$\phi_{kk'}^{opk}$	$\sum_j \frac{p_{kk'j}}{p_{kk'}} \log \frac{p_{kk'j} p_{k'j}}{p_{kk'} \cdot p_{k'j}}$	$\frac{1}{2} \left[\Delta \mathbf{m}^t \tilde{\Sigma}_A^{-1} \Delta \mathbf{m} + \ln \frac{ \tilde{\Sigma}_A }{ \Sigma_A } + \text{tr} \tilde{\Sigma}_A^{-1} \tilde{\Sigma}_A - n_A \right], \Delta \mathbf{m} \equiv (\tilde{\mathbf{A}}_A - \hat{\mathbf{A}}_A) \mathbf{x}_0^A + \tilde{\mathbf{A}}_B \mathbf{x}_0^B$
$\phi_{kk'}^{xfk}$	$I(\mathbf{x}_1^A, \mathbf{x}_1^B \mathbf{x}_0 = kk')$	$\frac{1}{2} \log \frac{ \Sigma_A \Sigma_B }{ \Sigma }$
ϕ^{nas}	$-\sum_j p_{\cdot j} \log \sum_{i j} \frac{p_{ij} p_{i'j}}{n_B p_{i'j} \cdot p_{\cdot j}}$	∞
ϕ_k^{nak}	$-\sum_j \frac{p_{kj}}{p_{k\cdot}} \log \sum_{i j} \frac{p_{ki'j} p_{k'j}}{n_B p_{ki'j} \cdot p_{k'j}}$	∞
ϕ^{nps}	$-\sum_{i'} p_{i' \cdot} \log \sum_{j j'} \frac{p_{ij} p_{i'j}}{n_B p_{ij} \cdot p_{i'j}}$	∞
$\phi_k^{npk} (\phi_k^{2.0})$	$-\sum_{i'} \frac{p_{i'k}}{p_{i' \cdot}} \log \sum_{j j'} \frac{p_{ij} p_{i'j} p_{k'j}}{n_B p_{ij} p_{i'j} p_{k'j}}$	∞
ϕ^{mas}	$-\sum_{j j'} p_{\cdot j j'} \log \sum_{i j'} \frac{p_{ij} p_{i'j} p_{j'j'}}{p_{i'j} \cdot p_{j'j'}}$	$\frac{1}{2} \left[\ln \frac{ \mathbf{C}_q }{ \mathbf{C}_A } + \text{tr} \mathbf{C}_q^{-1} \mathbf{C}_A - n_A \right], \mathbf{C}_q \equiv \Sigma_A + \mathbf{A}_A \mathbf{C}_A \mathbf{A}_A^t + \mathbf{A}_{AB} \mathbf{C}_B \mathbf{A}_{AB}^t$
ϕ_k^{mak}	$-\sum_j \frac{p_{kj}}{p_{k\cdot}} \log \sum_{i j} \frac{p_{ki'j} p_{k'j}}{p_{ki'j} \cdot p_{k'j}}$	$\frac{1}{2} \left[\Delta \mathbf{m}^t \tilde{\Sigma}_A^{-1} \Delta \mathbf{m} + \ln \frac{ \tilde{\Sigma}_A }{ \Sigma_A } + \text{tr} \tilde{\Sigma}_A^{-1} \Sigma_A - n_A \right], \Delta \mathbf{m} \equiv (\hat{\mathbf{A}}_A - \mathbf{A}_A) \mathbf{x}_0^A$
ϕ^{mps}	$-\sum_{i'} p_{i' \cdot} \log \sum_{j j'} \frac{p_{ij} p_{i'j} p_{j'j'}}{p_{ij} p_{i'j} p_{j'j'}}$	$\frac{1}{2} \left[\ln \frac{ \mathbf{C}_q }{ \mathbf{C}_A } + \text{tr} \mathbf{C}_q^{-1} \mathbf{C}_A - n_A \right], \mathbf{C}_q \equiv \tilde{\Sigma}_A + \tilde{\mathbf{A}}_A \mathbf{C}_A \tilde{\mathbf{A}}_A^t + \tilde{\mathbf{A}}_{AB} \mathbf{C}_B \tilde{\mathbf{A}}_{AB}^t$
ϕ_k^{mpk}	$-\sum_{i'} \frac{p_{i'k}}{p_{i' \cdot}} \log \sum_{j j'} \frac{p_{ki'j} p_{k'j} p_{j'j'}}{p_{ki'j} p_{k'j} p_{j'j'}}$	$\frac{1}{2} \left[\Delta \mathbf{m}^t \tilde{\Sigma}_A^{-1} \Delta \mathbf{m} + \ln \frac{ \tilde{\Sigma}_A }{ \Sigma_A } + \text{tr} \tilde{\Sigma}_A^{-1} \tilde{\Sigma}_A - n_A \right], \Delta \mathbf{m} \equiv (\hat{\mathbf{A}}_A - \tilde{\mathbf{A}}_A) \mathbf{x}_0^A$
$\phi^{2.5}$	$\min \{ \phi^{nak}, \phi^{npk} \}$	∞
$\phi^{2.5'}$	$\min \{ \phi^{mak}, \phi^{mpk} \}$	
$\phi^{2.5''}$	$\min \{ \phi^{oak}, \phi^{opk} \}$	

doi:10.1371/journal.pcbi.1005123.t002

As illustrated in Fig 1, it is always possible to model this relation between \mathbf{x}_0 and \mathbf{x}_1 as resulting from a Markov process, where \mathbf{x}_1 is causally determined by a combination of \mathbf{x}_0 and random effects. If we write the marginal distributions from eq (1) as vectors $\mathbf{p}^{(0)}$ and $\mathbf{p}^{(1)}$, this Markov process is defined by

$$\mathbf{p}^{(1)} = \mathbf{M} \mathbf{p}^{(0)}, \tag{2}$$

where the Markov matrix M_{ji} specifies the probability that a state i transitions to a state j , and satisfies the conditions $M_{ji} \geq 0$ (non-negative transition probabilities) and $M_{\cdot i} = 1$ (unit column sums, guaranteeing probability conservation). The standard rule for conditional



Φ measures inability to tensor factorize $M = M^A \otimes M^B$

Fig 1. We model the time-evolution of the system state as a Markov process defined by a transition matrix M : when the (possibly unknown) system state evolves from x_0 to x_1 , the corresponding probability distribution evolves from p_0 to $p_1 \equiv M p_0$. All competing definitions of Φ quantify the inability to tensor factorize M , which corresponds to approximating the system as two disconnected parts A and B that do not affect one another.

doi:10.1371/journal.pcbi.1005123.g001

probabilities gives

$$\begin{aligned}
 p_{ij} &= P(x_0 = i \ \& \ x_1 = j) = \\
 &= P(x_0 = i)P(x_1 = j | x_0 = i) = p_i^{(0)}M_{ji},
 \end{aligned}
 \tag{3}$$

which uniquely determines the Markov matrix as

$$M_{ji} = \frac{p_{ij}}{p_i^{(0)}} = \frac{p_{ij}}{p_i},
 \tag{4}$$

which is seen to satisfy the Markov requirements $M_{ji} \geq 0$ and $M_{.i} = 1$.

Note that any system obeying the laws of classical physics can be accurately modeled as a Markov process as long as the time step $\Delta t \equiv t_1 - t_0$ is sufficiently short (defining $x(t)$ as the position in phase space). If the process has “memory” such that the next state depends not only on the current state but also on some finite number of past states, it can reformulated as a standard memoryless Markov process by simply expanding the definition of the state x to include elements of the past.

Also note that although full knowledge of the Markov matrix M completely specifies the dynamics of the system, a person wishing to compute its integration may not know M exactly. If M is not known from having built the system or having examined its inner workings, then passively observing it in action (without active interventions) may not provide enough information to fully reconstruct M [21]. The $n \rightarrow \infty$ limit of continuous variables describes a convenient class of systems where M is relatively easy to determine in practice.

A taxonomy of integration measures

We will now see that this Markov process interpretation allows us create a simple taxonomy of integration measures ϕ that quantify the interaction between two subsystems. The idea is to approximate the Markov process by a *separable* Markov process that does not mix information between subsystems, and to define the integration as a measure of how bad the best such approximation is. Consider the system x as being composed of two subsystems x^A and x^B , so that the elements of the vector x are simply the union of the elements of x^A and x^B , and let us

define the probability distribution

$$p_{ii'jj'} \equiv P(\mathbf{x}_0^A = i \ \& \ \mathbf{x}_0^B = i' \ \& \ \mathbf{x}_1^A = j \ \& \ \mathbf{x}_1^B = j'). \tag{5}$$

(For brevity, we will sometimes refer to this distribution $p_{ii'jj'}$ as simply p below, suppressing the indices, and we will sometimes write \mathbf{x} without indices to refer to the full state at both times.) The Markov matrix of eq (4) then takes the form

$$M_{jj'ii'} = \frac{P_{ii'jj'}}{P_{ii'..}}. \tag{6}$$

The Markov process of eq (2) is separable if the Markov matrix \mathbf{M} is a tensor product $\mathbf{M}^A \otimes \mathbf{M}^B$, i.e., if

$$M_{jj'ii'} = M_{ji}^A M_{j'i'}^B \tag{7}$$

for Markov matrices \mathbf{M}^A and \mathbf{M}^B that determine the evolution of \mathbf{x}^A and \mathbf{x}^B .

If our system is integrated so that \mathbf{M} cannot be factored as in eq (7), we can nonetheless choose to approximate \mathbf{M} by a matrix of the factorizable form $\mathbf{M}^A \otimes \mathbf{M}^B$. If we retain the initial probability distribution $p_{ii'..}$ for \mathbf{x}_0 but replace the correct Markov matrix \mathbf{M} by the separable approximation $\mathbf{M}^A \otimes \mathbf{M}^B$, then eq (6) shows that the probability distribution

$$p_{ii'jj'} = M_{jj'ii'} p_{ii'..} \tag{8}$$

gets replaced by the probability distribution $q_{ii'jj'}$ given by

$$q_{ii'jj'} = M_{ji}^A M_{j'i'}^B p_{ii'..} \tag{9}$$

which is an approximation of $p_{ii'jj'}$. If \mathbf{M} is factorizable (meaning that there is no integration), we can factor \mathbf{M} such that the two probability distributions $q_{ii'jj'}$ and $p_{ii'jj'}$ are equal and, conversely, if the two probability distributions are different, we can use how different they are as an integration measure ϕ .

To define an integration measure ϕ in this spirit, we thus need to make four different choices, which collectively specify it fully and determine where the ϕ -measure belongs in our taxonomy:

1. Choose a recipe defining an approximate factorization $\mathbf{M} \approx \mathbf{M}^A \otimes \mathbf{M}^B$.
2. Choose which probability distributions p and q to compare for exact and approximate \mathbf{M} (the distribution for \mathbf{x} , \mathbf{x}_1 or \mathbf{x}_1^A , say).
3. Choose what to treat as known about $p_{ii'..}$ when computing these probability distributions.
4. Choose a metric for how different the two probability distributions p and q are.

These four options are described in Tables 3, 4 and 5, and we will now explore them in detail.

Options for approximately factoring \mathbf{M}

Table 3 lists five factoring options which all have attractive features, and we will now describe each in turn.

Approximately factoring \mathbf{M} using noising. The first option corresponds to the “noising” method used in IIT [10]: the time evolution of one part of the system (\mathbf{x}^A , say) is determined from the past state \mathbf{x}_0^A alone, treating \mathbf{x}_0^B as random noise with some probability distribution $p^{(B0)}$ that is independent of \mathbf{x}_0^A . In other words, we replace the initial probability distribution $p_{ii'}^{(0)} = p_{ii'..}$ by the separable distribution $p_{ii'}^{(0)} = p_i^{(A0)} p_i^{(B0)}$. We will now see that if we start with

Table 3. Different options for approximate factorizations $M \approx M^A \otimes M^B$ and $\tilde{M} \approx \tilde{M}^A \otimes \tilde{M}^B$. These options correspond to the first superscript in ϕ -measures such as ϕ^{ofuk} . The optimal factorizations maximize the accuracy of the approximate probability distribution that they predict, while the “noising” factorizations are instead defined by treating the input from the other subsystem as random noise, either with uniform distribution (option “n”) or with the observed marginal distribution (option “m”).

Code	Factorization method	M_{ij}^A	$M_{ij'}^B$	\tilde{M}_{ij}^A	$\tilde{M}_{ij'}^B$	State-dependent?
n	Noising	$\frac{1}{n_B} \sum_{j'} \frac{P_{ij'k}}{P_{i'j'}}$	$\frac{1}{n_A} \sum_i \frac{P_{ij'k}}{P_{i'j'}}$	$\frac{1}{n_B} \sum_{j'} \frac{P_{ij'j}}{P_{i'j'}}$	$\frac{1}{n_A} \sum_i \frac{P_{ij'j}}{P_{i'j'}}$	N
m	Mild noising	$\sum_{j'} \frac{P_{ij'j} P_{i'j'k}}{P_{i'j'}}$	$\sum_i \frac{P_{ij'j} P_{i'j'k}}{P_{i'j'}}$	$\sum_{j'} \frac{P_{ij'j} P_{i'j'j}}{P_{i'j'}}$	$\sum_i \frac{P_{ij'j} P_{i'j'j}}{P_{i'j'}}$	N
o	Optimal not knowing state \mathbf{x}_0	$\frac{P_{ij.}}{P_{i..}}$	$\frac{P_{i'j'}}{P_{i'..}}$	$\frac{P_{ij.}}{P_{.j}}$	$\frac{P_{i'j'}}{P_{i'j}}$	N
x	Optimal given \mathbf{x}_0	$\frac{P_{kk'j.}}{P_{kk'..}}$	$\frac{P_{kk'j'}}{P_{kk'..}}$	$\frac{P_{kk'j.}}{P_{kk'..}}$	$\frac{P_{kk'j'}}{P_{kk'..}}$	Y
a	Optimal given \mathbf{x}_0 , on average	$\frac{P_{ij.}}{P_{i..}}$	$\frac{P_{i'j'}}{P_{i'..}}$	$\frac{P_{ij.}}{P_{.j}}$	$\frac{P_{i'j'}}{P_{i'j}}$	N

doi:10.1371/journal.pcbi.1005123.t003

Table 4. Different options for which probability distributions p and q to compare, corresponding to the second and third superscripts in ϕ -measures such as ϕ^{ofuk} . The last three columns specify the formula for q for the three conditioning options we consider: when the state \mathbf{x}_0 is unknown (u), has a separable probability distribution (s) and is known (k), respectively.

Code	Comparison option	p	q	Conditioning option		
				u	s	k
				\mathbf{x}_t unknown	\mathbf{x}_t -distribution separable	\mathbf{x}_t known
				$q^u (p_{ij}^{(0)} = p_{i'j'})$	$q^s (p_{ij}^{(0)} = p_{i.} p_{i'j'})$	$q^k (p_{ij}^{(0)} = \delta_{ik} \delta_{i'k'})$
t	Two-time state	$P_{i'j'j}$	$Q_{i'j'j}$	$M_{ij}^A M_{i'j'}^B P_{i'j'j}^{(0)}$	$(M_{ij}^A p_{i.}^{(0)}) (M_{i'j'}^B p_{i'j'}^{(0)})$	$M_{ij}^A \delta_{ik} M_{i'j'}^B \delta_{i'k'}$
f	Future state	$P_{.ij}$	$Q_{.ij}$	$\sum_{i'} M_{ij}^A M_{i'j'}^B P_{i'j'}$	$(\sum_i M_{ij}^A p_{i.}^{(0)}) (\sum_{i'} M_{i'j'}^B p_{i'j'}^{(0)})$	$M_{jk}^A M_{i'k'}^B$
a	Future state of subsystem A	$p_{.j}$	$q_{.j}$	$\sum_i M_{ij}^A p_{i.}^{(0)}$	$\sum_i M_{ij}^A p_{i.}^{(0)}$	M_{jk}^A
p	Past state of subsystem A	$p_{i..}$	$\tilde{q}_{i..}$	$\sum_j \tilde{M}_{ij}^A p_{i'j}^{(1)}$	$\sum_j \tilde{M}_{ij}^A p_{i'j}^{(1)}$	\tilde{M}_{ik}^A

doi:10.1371/journal.pcbi.1005123.t004

Table 5. Different options for measuring the difference d between two probability distributions p and q : Kullback-Leibler divergence d_{KL} , L_1 -norm d_1 , L_2 -norm d_2 , Hilbert-space distance d_H , Shannon-Jensen distance $d_{S,J}$, Earth-Movers distance d_{EM} and Mismatched Decoding distance d_{MD} . These options correspond to the fourth superscript in ϕ -measures such as ϕ^{ofuk} . In the text, we considered options where p and q had one, two or four indices, but in this table, we have for simplicity combined all indices into a single Greek index α .

Code	Metric	Definition	Positivity	Monotonicity	Interpretability	Tractability	Symmetry
k	$d_{KL}(p, q)$	$\sum_{\alpha} p_{\alpha} \log \frac{p_{\alpha}}{q_{\alpha}}$	Y	Y	Y	Y	N
1	$d_1(p, q)$	$\sum_{\alpha} p_{\alpha} - q_{\alpha} $	Y	Y	(Y)	N	Y
2	$d_2(p, q)$	$\left[\sum_{\alpha} (p_{\alpha} - q_{\alpha})^2 \right]^{1/2}$	Y	Y	(N)	(Y)	Y
h	$d_H(p, q)$	$\cos^{-1} \sum_{\alpha} (p_{\alpha} q_{\alpha})^{1/2}$	Y	Y	Y	N	Y
s	$d_{S,J}(p, q)$	$\left[\sum_{\alpha} \left(\frac{p_{\alpha}}{2} \log \frac{2p_{\alpha}}{p_{\alpha} + q_{\alpha}} + \frac{q_{\alpha}}{2} \log \frac{2q_{\alpha}}{p_{\alpha} + q_{\alpha}} \right) \right]^{1/2}$	Y	Y	Y	N	Y
e	$d_{EM}(p, q)$	$\min_{f_{\alpha\beta} \geq 0} \sum_{\alpha\beta} f_{\alpha\beta} d_{\alpha\beta}; f_{\alpha.} = p_{\alpha}, f_{. \beta} = q_{\beta}$	Y	Y	Y	N	Y
m	$d_{MD}(p, q)$	See eqs (24) and (25)	Y	Y	Y	N	N

doi:10.1371/journal.pcbi.1005123.t005

eq (2), *i.e.*, the Markov equation $\mathbf{p}^{(1)} = \mathbf{M}\mathbf{p}^{(0)}$, then this noising prescription gives $\mathbf{p}^{(A1)} = \mathbf{M}^A \mathbf{p}^{(A0)}$ for a particular matrix \mathbf{M}^A . Eq (2) states that

$$p_{jj'}^{(1)} = \sum_{i'} M_{jj'i'} P_{i'}^{(0)}, \tag{10}$$

and substituting the separable “noising” form of $p_{i'}^{(0)}$ from above gives

$$p_j^{(A1)} \equiv p_j^{(1)} = \sum_{i'} M_{j:i'} P_i^{(A0)} p_{i'}^{(B0)} = \sum_i M_{ji}^A P_i^{(A0)}, \tag{11}$$

where we have defined

$$M_{ji}^A \equiv \sum_{i'} M_{j:i'} P_{i'}^{(B0)} = \sum_{i'} \frac{P_{i'i} P_{i'}^{(B0)}}{P_{i'..}}. \tag{12}$$

IIT chooses the noise to have maximum entropy, *i.e.*, a uniform distribution over the n_B possible states of subsystem B [10]:

$$P_{i'}^{(B0)} = \frac{1}{n_B}. \tag{13}$$

Table 3 lists the \mathbf{M}^A -matrix corresponding to this noising choice as well as the analogous \mathbf{M}^B -matrix.

Approximately factoring M using mild noising. One drawback of this choice is that uniform distributions are undefined for continuous variables such as measured voltages, because they cannot be normalized. This means that any ϕ -measure based on this noising factorization is undefined and useless for continuous systems. This problem can be solved by adopting another natural choice for the noise distribution:

$$P_{i'}^{(B0)} = P_{i'..}, \tag{14}$$

i.e., simply the marginal distribution for \mathbf{x}_0^B . We term this option “mild noising”, since the noise is less extreme (its entropy is lower) than with the previous noising option. Table 3 lists the \mathbf{M}^A -matrix corresponding to this mild noising choice as well as the analogous \mathbf{M}^B -matrix.

Optimally factoring M. A drawback of both factorizations that we have considered so far is that they might overestimate integration: there may exist an alternative factorization that is better in the sense of giving a smaller ϕ . The natural way to remedy this problem is to define ϕ by minimizing over all factorizations. This elegantly unifies with the fact that capital Φ is defined by minimizing over all partitions of the system into two parts: we can capture both minimizations by simply saying “minimize over all factorizations”, since the choice of a tensor factorization includes a choice of partition.

In practice, the definition of the optimal factorization depends on what we optimize. We discuss various options below, and identify three particularly natural choices which are listed in Table 3. The first option makes the approximate probability distribution $q_{ii'jj'}$ as similar as possible to $p_{ii'jj'}$, where similarity is quantified by KL-divergence. The second option treats the present state \mathbf{x}_0 as known and makes the conditional probability distribution for the future state \mathbf{x}_1 as similar as possible to the correct distribution. This factorization thus depends on the state and hence on time, whereas all the others we have considered are state-independent. The third option is the factorization that minimizes this state-dependent ϕ *on average*; we will prove below that this factorization is identical to the first option.

In summary, Table 3 lists five factorization options that each have various attractive features; options 3 and 5 turn out to be identical. It is easy to show that if the Markov matrix \mathbf{M} is

factorizable (which means that the probability distribution is separable as $p_{i'j'} = p_{ij}^A p_{i'j'}^B$), then all five factorizations coincide, all giving $M_{ji}^A = p_{ji}^A / p_i^A$ and $M_{j'i'}^B = p_{j'i'}^B / p_{i'}^B$. This means that they will all agree on when $\phi = 0$; otherwise the noising factorizations will yield higher ϕ than an optimized factorization.

Options for which probability distributions to compare

Table 4 lists four options for which probability distributions p and q to compare. Arguably the most natural option is to simply compare the full distributions $p_{i'j'}$ and $q_{i'j'}$ that describe our knowledge of the system at both times (the present state and the future state). Another obvious option is to merely compare the predictions, *i.e.*, the probability distributions $p_{\cdot jj'}$ and $q_{\cdot jj'}$ for the future state. A third interesting option is to compare merely the predictions for one of the two subsystems (which we without loss of generality can take to be subsystem A), thus comparing $p_{\cdot j}$ and $q_{\cdot j}$.

Generally, the less we compare, the easier it is to get a low ϕ -value. To see this, consider a system where A affects B but B has no effect on A . We could, for example, consider A to be photoreceptor cells in your retina and B to be the rest of your brain. Then the second comparison option (“f”) in Table 4 would give $\phi > 0$ because we predict the future of your brain worse if we ignore the information flow from your retina, while the third comparison option (“a”) in the table would give $\phi = 0$ because the rest of your brain does not help predict the future of your retina. In other words, comparison option “a” makes ϕ vanish for *afferent* pathways, where information flows only inward toward the rest of the system.

IIT argues that any good ϕ -measure indeed *should* vanish for afferent pathways, because a system can only be conscious if it can have effects on itself—other systems that it is affected by without affecting will act merely as parts of its unconscious outside world [10]. Analogously, IIT argues that any good ϕ -measure should vanish also for *efferent* pathways, where information flows only outward away from the rest of the system. The argument is that other systems that the conscious system affects without being affected by will again be unconscious, acting merely as unconscious parts of the outside world as far as the conscious system is concerned.

Option “p” in Table 4 has this property of ϕ vanishing for efferent pathways. It is simply the time-reverse of option “a”, quantifying the ability of x_1^A to determine its past cause x_0^A instead of quantifying the ability of x_0^A to determine its future effect x_1^A .

To formalize this, consider that there is nothing in the probability distribution $p_{i'j'}$ that breaks time-reversal symmetry and says that we must interpret causation as going from t_0 to t_1 rather than vice versa. In complete analogy with our formalism above, we can therefore define a time-reversed Markov process \tilde{M} whereby the future determines the past according to the time-reverse of eq (2):

$$\mathbf{p}^{(0)} = \tilde{M}\mathbf{p}^{(1)}, \tag{15}$$

where eqs (6) and (9) get replaced by

$$\tilde{M}_{i'j'} = \frac{p_{i'j'}}{p_{\cdot jj'}} \tag{16}$$

and

$$\tilde{q}_{i'j'} = \tilde{M}_y^A \tilde{M}_{rj'}^B p_{jj'}^{(1)}. \tag{17}$$

This time reversal symmetry doubles the number of q -options we could list in Table 4 to six in total, augmenting $q_{i'j'}$, $q_{\cdot jj'}$ and $q_{\cdot j}$ by $\tilde{q}_{i'j'}$, $\tilde{q}_{\cdot jj'}$ and $\tilde{q}_{\cdot j}$. In the interest of brevity, we have

chosen to only list $\tilde{q}_{..j}$, because of its ability to kill ϕ for efferent pathways—the formulas for the two omitted options are trivially analogous to those listed.

Options for what to treat as known about the current state

Above we listed options for which probabilities p and q to compare to compute ϕ . To complete our specification of these probabilities, we need to choose between various options for our knowledge of the present state; the three rightmost columns of Table 4 correspond to three interesting choices.

The first option is where the state is unknown, described simply by the probability distribution we have used above:

$$p_{ij}^{(0)} = p_{i'..} \tag{18}$$

This corresponds to us knowing \mathbf{M} , the mechanism by which the state evolves, but not knowing its current state \mathbf{x}_0 . Note that a generic Markov process eventually converges to a unique stationary state $\mathbf{p} = \mathbf{p}^{(0)} = \mathbf{p}^{(1)}$ which, since it satisfies $\mathbf{M}\mathbf{p} = \mathbf{p}$, can be computed directly from \mathbf{M} as the unique eigenvector whose eigenvalue is unity (the only Markov processes that do not converge to a unique steady state are ones where \mathbf{M} has more than one eigenvalue equal to unity; these form a set of measure zero on the set of all Markov processes). This means that if we consider a system that has been evolving for a significantly long time, its full two-time distribution $p_{i'j'}$ is determined by \mathbf{M} alone; conversely, $p_{i'j'}$ determines \mathbf{M} through eq (6). Alternatively, if $p_{i'j'}$ is measured empirically from a time-series \mathbf{x}_t which is then used to compute \mathbf{M} , we can use eq (6) to describe our knowledge of the state at a random time.

A second option is to assume that we know the initial probability distributions for \mathbf{x}_0^A and \mathbf{x}_0^B , but know nothing about any correlations between them. This corresponds to replacing eq (18) by the separable distribution

$$p_{ij}^{(0)} = p_{i..}p_{.j'..}, \tag{19}$$

and can be advantageous for ϕ -measures that would conflate integration with initial correlations between the subsystems.

A third option, advocated by IIT [10], is to treat the current state as known:

$$p_{ij}^{(0)} = \delta_{ik}\delta_{i'k'}, \tag{20}$$

i.e., we know with certainty that the current state $\mathbf{x}_0 = kk'$ for some constants k and k' . IIT argues that this is the correct option from the vantage point of a conscious system which, by definition, knows its own state.

A natural fourth option is a more extreme version of the first: treating the state not merely as unknown, with $p^{(0)}$ given by its ensemble distribution, but completely unknown, with a uniform distribution:

$$p_{ij}^{(0)} = \text{constant}. \tag{21}$$

Although straightforward enough to use in our formulas, we have chosen not to include this option in Table 4 because it is rather inappropriate for most physical systems. For continuous variables such as voltages, it becomes undefined. For brains, such maximum-entropy states never occur: they would have typical neurons firing about half the time, corresponding to much more extreme “on” behavior than during an epileptic seizure. The related option of consistently treating \mathbf{x}_0^A as known but \mathbf{x}_0^B as unknown when predicting \mathbf{x}_1^A (and vice versa when

predicting \mathbf{x}_1^B) corresponds to the noising factorization options described above. For further discussion of this, including so-called “noising at the connection”, see [11, 14, 22].

Finally, please note that if we choose to determine the past rather than the future (the “p”-option from the previous section and Table 4), then all the choices we have described should be applied to $p_{ij}^{(1)}$ rather than $p_{ij}^{(0)}$.

Options for comparing probability distributions

The options in the past three sections uniquely specify two probability distributions \mathbf{p} and \mathbf{q} , and we want the integration ϕ to quantify how different they are from one another:

$$\phi \equiv d(\mathbf{p}, \mathbf{q}) \tag{22}$$

for some distance measure d that is larger the worse \mathbf{q} approximates \mathbf{p} . There are a number of properties that we may consider desirable for d to quantify integration:

1. **Positivity:** $d(\mathbf{p}, \mathbf{q}) \geq 0$, with equality if and only if $\mathbf{p} = \mathbf{q}$.
2. **Monotonicity:** The more different \mathbf{q} is from \mathbf{p} in some intuitive sense, the larger $d(\mathbf{p}, \mathbf{q})$ gets.
3. **Interpretability:** $d(\mathbf{p}, \mathbf{q})$ can be intuitively interpreted, for example in terms of information theory.
4. **Tractability:** $d(\mathbf{p}, \mathbf{q})$ is easy to compute numerically. Ideally, the optimal factorizations can be found analytically rather than through time-consuming numerical minimization.
5. **Symmetry:** $d(\mathbf{p}, \mathbf{q}) = d(\mathbf{q}, \mathbf{p})$.

Any distance measure d meets the mathematical requirements of being a *metric* on the space of probability distributions if it obeys positivity, symmetry and the triangle inequality $d(\mathbf{p}, \mathbf{q}) \leq d(\mathbf{p}, \mathbf{r}) + d(\mathbf{r}, \mathbf{q})$.

Table 5 lists seven interesting probability distribution distance measures $d(\mathbf{p}, \mathbf{q})$ from the literature together with their definitions and properties. All these measures are seen to have the positivity and monotonicity, and all except the first are also symmetric and true metrics. We will now discuss them one by one in greater detail.

The distance d_{KL} is the Kullback-Leibler divergence, and measures how many bits of information are lost when \mathbf{q} is used to approximate \mathbf{p} , in the sense that if you developed an optimal data compression algorithm to compress data drawn from a probability distribution \mathbf{q} , it would on average require $d_{KL}(\mathbf{p}, \mathbf{q})$ more bits to compress data drawn from a probability distribution \mathbf{p} than if the algorithm had been optimized for \mathbf{p} [23]. This has been argued to be the best measure because of its desirable properties related to information geometry [24, 25].

d_1 and d_2 measure the distance between the vectors \mathbf{p} and \mathbf{q} using the L_1 -norm and L_2 -norm, respectively. The former is particularly natural for probability distributions \mathbf{p} , since they all have L_1 norm of unity: $d_1(0, \mathbf{p}) = p = 1$. It is easy to see that $0 \leq d_1(\mathbf{p}, \mathbf{q}) \leq 2$ and $0 \leq d_2(\mathbf{p}, \mathbf{q}) \leq \sqrt{2}$.

The measure d_H is the Hilbert-space distance: if, for each probability distribution, we define a corresponding wavefunction $\psi_i \equiv p_i^{1/2}$, then all wavefunctions lie on a unit hypersphere since they all have unit length: $\langle \psi | \psi \rangle = p = 1$. The distance d_H is simply the angle between two wavefunctions, *i.e.*, the distance along the great circle on the hypersphere that connects the two, so $d_H(\mathbf{p}, \mathbf{q}) \leq \pi/2$. It is also the geodesic distance of the Fisher metric, hence a natural “coordinate free” distance measure on the manifold of all probability distributions.

The measure d_{SJ} is the Shannon-Jensen distance, whose square is defined as the average of the KL-divergences of the two distributions to their average:

$$d_{SJ}(\mathbf{p}, \mathbf{q})^2 \equiv \frac{d_{KL}(\mathbf{p}, [\mathbf{p} + \mathbf{q}]/2) + d_{KL}(\mathbf{q}, [\mathbf{p} + \mathbf{q}]/2)}{2} \tag{23}$$

$$= S[(\mathbf{p} + \mathbf{q})/2] - (S[\mathbf{p}] + S[\mathbf{q}])/2.$$

It is bounded by $0 \leq d_{SJ} \leq 1$, satisfies the triangle inequality and is information-theoretically motivated [26].

The measure d_{EM} is the Earth-Movers distance [27]. If we imagine piles of earth scattered across the space \mathbf{x} , with $p(\mathbf{x})$ specifying the fraction of the earth that is in each location, then d_{EM} is the average distance that you need to move earth to turn the distribution $p(\mathbf{x})$ into $q(\mathbf{x})$. The quantity d_{ij} in the definition in Table 5 specifies the distance between points i and j in this space. For example, if \mathbf{x} is a 3D Euclidean space, this may be chosen to be simply the Euclidean metric, while if \mathbf{x} is a bit string, d_{ij} may be chosen to be the L_1 “Manhattan distance”, *i.e.*, the number of bit flips required to transform one bit string into another. IIT 3.0 argues that the earth mover’s distance d_{EM} is the most appropriate measure d on conceptual grounds (whereas IIT 2.0 was still implemented using d_{KL}). Unfortunately, d_{EM} rates poorly on the tractability criterion. It’s definition involves a linear programming problem which needs to be solved numerically, and even with the fastest algorithms currently available, the computation grows faster than quadratically with the number of system states—which in turn grows exponentially with the number of bits. For continuous variables \mathbf{x} , the number of states and hence the computational time is formally infinite.

The measure d_{MD} is based on “mismatched decoding” as advocated by [15]. The distance measure d_{MD} is defined not for all probability distributions, but for all distributions over two variables, which we can write with two indices as p_{ij} :

$$d_{MD}(\mathbf{p}, \mathbf{q}) \equiv I(\mathbf{p}) - \max_{\beta} I^*(\mathbf{p}, \mathbf{q}, \beta), \tag{24}$$

where

$$I(\mathbf{p}) \equiv -\sum_j p_{.j} \log p_{.j} + \sum_{ij} p_{ij} \log p_{ji}, \tag{25}$$

$$I^*(\mathbf{p}, \mathbf{q}, \beta) \equiv -\sum_j p_{.j} \log \sum_i q_{ji}^{\beta} p_{.i} + \sum_{ij} p_{ij} \log q_{ji}^{\beta},$$

and the conditional distribution $q_{ji} \equiv q_{ij}/q_{.i}$. Here $I(\mathbf{p})$ is simply the mutual information between the two variables, since combining eq (34) with the conditional entropy definition from eq (37) gives the well-known equivalent expression for mutual information

$$I(A, B) = S(A) - S(A|B). \tag{26}$$

$I^*(\mathbf{p}, \mathbf{q}, \beta)$ can be interpreted as the amount of information that one variable predicts about the other if the correct conditional distribution p_{ji} is replaced by a possibly incorrect one q_{ji}^{β} (renormalized to sum to unity) when making the prediction [28]. This renormalization is strictly speaking unnecessary, because it cancels out between the two terms in $I^*(\mathbf{p}, \mathbf{q}, \beta)$. Raising probabilities to positive powers β has the effect of concentrating them (decreasing entropy) if $\beta > 1$ and spreading them more evenly (increasing entropy) if $\beta < 1$. It can be shown that $I^*(\mathbf{p}, \mathbf{q}, \beta) \leq I(\mathbf{p})$ with equality for $\mathbf{q} = \mathbf{p}$ and $\beta = 1$, and that $I^*(\mathbf{p}, \mathbf{q}, \beta) \geq 0$, so one always has $0 \leq d_{MD}(\mathbf{p}, \mathbf{q}) \leq I(\mathbf{p})$ [28]. Mismatched decoding can presumably be further generalized by

replacing the maximization over powers p^β by maximization over arbitrary monotonically increasing functions $f(p)$ that map the unit interval onto itself.

The integration measures of IIT3.0 have a more complex probability comparison that cannot be fully cast in the form of a simple function of $d(\mathbf{p}, \mathbf{q})$: it makes the metric choice $d(\mathbf{p}, \mathbf{q}) = d_{EM}(\mathbf{p}, \mathbf{q})$, but considers not only probability distributions for the whole system and a bipartition, but also for all possible subsets, providing an elaborate interpretation of the results in terms of “conceptual structures” [11].

Results

In this section, we present our taxonomy results.

Optimal factorization with d_{KL}

Our taxonomy of integration measures is determined by four choices: of factorization, variable selection, conditioning and distance measure. Although we have now explored these four choices one at a time, there are important interplays between them that we must examine. First of all, the three optimal factorization options in Table 3 depend on what is being optimized, so let us now explore which of these optimizations are feasible and interesting to perform in practice and let us find out what the corresponding factorizations and ϕ -measures are.

The mathematics problem we wish to solve is

$$\phi \equiv \min_{\mathbf{M}^A, \mathbf{M}^B} d(\mathbf{p}, \mathbf{q}) \tag{27}$$

i.e., minimizing $d(\mathbf{p}, \mathbf{q})$ over \mathbf{M}^A and \mathbf{M}^B given the constraints that \mathbf{M}^A and \mathbf{M}^B are Markov Matrices: $M_{.j}^A = 1$, $M_{.j}^B = 1$, $M_{ij}^A \geq 0$ and $M_{ij}^B \geq 0$. Table 4 specifies the options for how \mathbf{p} and \mathbf{q} are computed and how \mathbf{q} depends on \mathbf{M}^A and \mathbf{M}^B , while Table 5 specifies the options for computing the distance measure d . We enforce the column sum constraints using Lagrange multipliers, minimizing

$$\mathcal{L} \equiv d(\mathbf{p}, \mathbf{q}) - \sum_i \lambda_i (M_{.i}^A - 1) - \sum_{j'} \mu_{j'} (M_{.j'}^B - 1), \tag{28}$$

and need to check afterwards that all elements of \mathbf{M}^A and \mathbf{M}^B come out to be non-negative (we will see that this is indeed the case).

As mentioned, numerical tractability is a key issue for integration measures. This means that it is valuable if the Lagrange minimization can be rapidly solved analytically rather than slowly by numerical means, since this needs to be done separately for large numbers of possible system partitions. There is only one d -option out of the above-mentioned five for which I have been able to solve the optimization over \mathbf{M} -factorizations analytically: the KL-divergence d_{KL} . The runner-up for tractability is d_2 , for which everything can be easily solved analytically except for a final column normalization step, but the resulting formulas are cumbersome and unilluminating, falling foul of the interpretability criterion. Although d_{KL} lacks the symmetry property, it has the above-mentioned positivity, monotonicity and interpretability properties, and we will now show that it also has the tractability property.

Let us begin with the q -options in the upper left corner of Table 4, *i.e.*, comparing the two-time distributions treating the present state as unknown. Substituting eq (9) into the definition

of d_{KL} from Table 5 gives

$$d_{KL}(\mathbf{p}, \mathbf{q}) = \sum_{i'j'} p_{i'j'} \log \frac{P_{i'j'}}{P_{i'..} M_{ji}^A M_{j'j'}^B} = S(\mathbf{x}_0) - S(\mathbf{x}) - \sum_{ij} p_{ij} \log M_{ji}^A - \sum_{j'j} p_{.j'j} \log M_{j'j'}^B, \tag{29}$$

where the entropy for a random variable \mathbf{x} with probability distribution \mathbf{p} is given by Shannon’s formula [29]

$$S(\mathbf{x}) = - \sum_i p_i \log p_i. \tag{30}$$

To avoid a profusion of notation, we will often write as the argument of S a random variable rather than its probability distribution. For convenience, we will take all logarithms to be in base 2 for discrete distributions (so that entropies are measured in units of bits) and in base e for continuous Gaussian distributions (so that equations get simpler). In the latter case, where the entropy is based on the natural logarithm, entropy is measured in “nits” or “nats” which equal $1/\ln 2 \approx 1.44$ bits.

Substituting eq (29) into eq (28) and requiring vanishing derivatives with respect to M_{ij}^A , $M_{j'j}^B$, λ_j and $\mu_{j'}$ shows that the solution to our minimization problem is

$$M_{ji}^A = \frac{P_{ij.}}{P_{i...}}, \quad M_{j'j}^B = \frac{P_{.j'j}}{P_{.j'..}}. \tag{31}$$

We recognize these equations as simply the Markov matrix estimator from eq (4) applied separately to subsystems A and B after marginalizing over the other system. Substituting this back into eq (9) gives

$$q_{i'j'} = \frac{P_{i'..} P_{ij.} P_{.j'j}}{P_{i...} P_{.j'..}}. \tag{32}$$

Although the full probability distributions q and p typically differ, eq (32) implies that three marginal distributions are identical: $q_{i.j.} = p_{i.j.}$, $q_{.j'j} = p_{.j'j}$ and $q_{ij'..} = p_{ij'..}$.

Substituting eq (32) back into the definition of d_{KL} gives the extremely simple result that the integration is

$$\begin{aligned} \phi^{otuk}(\mathbf{p}) &= \sum_{i'j'} p_{i'j'} \log \frac{P_{i'j'} P_{i...} P_{.j'..}}{P_{i'..} P_{ij.} P_{.j'j}} \\ &= I(\mathbf{x}^A, \mathbf{x}^B) - I(\mathbf{x}_0^A, \mathbf{x}_0^B), \end{aligned} \tag{33}$$

where the mutual information between two random variables is given in terms of entropies by the standard definition

$$I(\mathbf{x}^A, \mathbf{x}^B) \equiv S(\mathbf{x}^A) + S(\mathbf{x}^B) - S(\mathbf{x}). \tag{34}$$

Since we will be deriving a large number of different ϕ -measures that we do not wish to conflate with one another, we superscript each one with four code letters denoting the four taxonomical choices that define it. These letter codes are

1. factorization: n/m/o/x/a
2. comparison: t/f/a/p
3. conditioning: u/s/k

4. measure: k/1/2/h/s/e/m

and are defined in Tables 3, 4 and 5. For example, the integration measure ϕ^{otuk} from eq (33) denotes optimized (o) factorization comparing the two-time (t) probability distributions with the current state unknown (u) and KL-divergence (k). Almost all measures discussed below will involve the k-measure (KL-divergence), so when this is the case we will typically drop this last index k to avoid a confusing profusion of indices, for example writing $\phi^{otuk} = \phi^{otu}$. For brevity, we will also define $\phi^M \equiv \phi^{otu}$, since we will be referring to this “Markov measure” ϕ^{otu} many times below.

Although we derived this optimal factorization by comparing the two-time distribution (option t) for an unknown state (option u), an analogous calculation leads to the exact same optimal factorization for the options a+u, s+f and a+s. The option t+s is undefined and the option f+u gives messy equations I have been unable to solve analytically. It is therefore reasonable to view eq (31) as *the* optimal factorization when the state is unknown (option o), and for the remainder of this paper, we will simply define the o-option as using the factorization given by eq (31).

Note that our result in eq (33) involves a time-asymmetry, singling out t_0 rather than t_1 in the second term. This is because we chose to interpret our Markov process as operating *forward* in time, determining the state at t_1 from the state at t_0 . As we discussed in the previous section, we could equally well have done the opposite, using the Markov process \tilde{M} operating *backward* in time, which would have yielded the alternative integration measure

$$\phi^{\tilde{o}iu}(\mathbf{p}) = I(\mathbf{x}^A, \mathbf{x}^B) - I(\mathbf{x}_1^A, \mathbf{x}_1^B). \tag{35}$$

In practice, one usually estimates all statistical properties from a time-series that is assumed to be stationary. This means that $I(\mathbf{x}_0^A, \mathbf{x}_0^B) = I(\mathbf{x}_1^A, \mathbf{x}_1^B)$, so that the these two integration measures become identical.

Comparison with the Ay/Barrett/Seth integration measures

In the paper [13] where Barrett & Seth proposed their easier-to-compute integration measure ϕ^B (see below), they also mentioned an alternative measure that they termed $\tilde{\phi}_E$, defined by

$$\tilde{\phi}_E \equiv S(\mathbf{x}_0^A | \mathbf{x}_1^A) + S(\mathbf{x}_0^B | \mathbf{x}_1^B) - S(\mathbf{x}_0 | \mathbf{x}_1), \tag{36}$$

where the *conditional entropy* of two variables A and B is defined by

$$S(A|B) \equiv S(A, B) - S(B). \tag{37}$$

This measure had been introduced earlier by Ay [30, 31] in a context unrelated to IIT, under the name “stochastic interaction”, and was further discussed in [15, 32]. Applying eqs (37) and (34) to eq (36) shows that

$$\begin{aligned} \tilde{\phi}_E &= S(\mathbf{x}^A) - S(\mathbf{x}_1^A) + S(\mathbf{x}^B) - S(\mathbf{x}_1^B) - S(\mathbf{x}) + S(\mathbf{x}_1) \\ &= I(\mathbf{x}^A, \mathbf{x}^B) - I(\mathbf{x}_1^A, \mathbf{x}_1^B) = \phi^{\tilde{o}iu}, \end{aligned} \tag{38}$$

i.e., that $\tilde{\phi}_E$ is identical to the time-reversed Markov measure $\phi^{\tilde{o}iu}$. This equivalence provides another convenient interpretation of $\phi^{\tilde{o}iu}$: as the average KL-divergence between (i) the probability distribution of the past state \mathbf{x}_0 given the present state \mathbf{x}_1 and (ii) the product of these conditional distributions for the two subsystems.

It is also interesting to compare our result in eq (33) with the popular integration measure

$$\phi^B(\mathbf{p}) = I(\mathbf{x}_0, \mathbf{x}_1) - I(\mathbf{x}_0^A, \mathbf{x}_1^A) - I(\mathbf{x}_0^B, \mathbf{x}_1^B) \tag{39}$$

proposed by Barrett & Seth [13]. The intuition behind this definition is to take the amount of information that a system predicts about its future and subtract of the information predicted by both of its subsystems. Unfortunately, the result can sometimes go negative [14, 15], violating the desirable positivity property and making the ϕ_B difficult to interpret. Consider the simple example of two independent bits that never change. If they start out perfectly correlated, then they will remain perfectly correlated, giving $I(\mathbf{x}_0, \mathbf{x}_1) = I(\mathbf{x}_0^A, \mathbf{x}_1^A) = I(\mathbf{x}_0^B, \mathbf{x}_1^B) = 1$ and integrated information $\phi^B(\mathbf{p}) = -1$.

By substituting eq (34) into eqs (33) and (39), we find that

$$\phi^B(\mathbf{p}) = \phi^M(\mathbf{p}) - I(\mathbf{x}_1^A, \mathbf{x}_1^B). \tag{40}$$

In other words, we can make the Barrett-Seth measure non-negative by adding back any final mutual information between the two subsystems. When this is done, it becomes the integration measure we derived, therefore having a simple information-theoretic interpretation: it is the KL-divergence between the actual probability distribution \mathbf{p} and the best separable approximation, which is guaranteed to be non-negative.

Comparison with the mismatched decoding integration measure

The measure ϕ^M is also closely related to the *mismatched decoding* measure ϕ^{MD} introduced in [15]. ϕ^{MD} makes the same taxonomical choices “otu” as ϕ^M for the first three options: optimal factorization (o), comparing full two-time distributions (t), and treating the past state as unknown (u). However, it uses probability distance measure “m” (mismatched decoding d_{MD}) instead of KL-divergence. We can therefore write this measure in our notation as $\phi^{otum} = d_{MD}(p, q)$, where q is the optimal factorization given by eq (32). Whether this factorization is also optimal in the sense of minimizing $d_{MD}(p, q)$ is not obvious.

The measure ϕ^M (or more specifically its time-reverse ϕ^{otuk}) has been criticized in [15, 25] for being able to exceed the mutual information $I(\mathbf{x}_0, \mathbf{x}_1)$ between the past and present: for example, if a two-bit system evolves from “00” to either “00” or “11” with equal probability, then $\phi^M = I(\mathbf{x}^A, \mathbf{x}^B) - I(\mathbf{x}_0^A, \mathbf{x}_0^B) = 1 - 0 = 1$ bit, even though $I(\mathbf{x}_0, \mathbf{x}_1) = 0$. This means that ϕ^M counts as a contribution to integration also correlated random noise added to both subsystems. It is debatable whether this should count as integration: the “con” argument is that no information flows between the subsystems, while the “pro” argument is that the two subsystems get linked by shared information flowing into both of them.

Both ϕ^M and ϕ^{MD} have intuitive bounds: $0 \leq \phi^M \leq I(\mathbf{x}^A, \mathbf{x}^B)$ and $0 \leq \phi^{MD} \leq I(\mathbf{x}_0, \mathbf{x}_1)$; these upper bounds correspond to the total mutual information across space and time, respectively.

Optimal state-dependent factorization

Let us now turn to factorization option “x”, optimized knowing the current state. Consider some conscious observer (perhaps the system itself) who knows nothing about the system except its dynamics (encoded in \mathbf{M}) and its state at the present instant, encoded in $\mathbf{x}_0 = kk'$. What can this observer say about the system state at earlier and later times? How integrated will this observer feel that the system is? To answer this question, we simply want to find the best approximate factorization of the conditional future state $M_{jj'kk'}$ (or the past state $M_{kk'ii'}$), where k and k' are known constants.

To gain intuition for this, let us temporarily write this conditional distribution as $p_{i'}$, suppressing the known parameters kk' for simplicity. Given an arbitrary bivariate probability distribution $p_{i'}$, what is best separable approximation $q_{i'} \equiv a_i b_{i'}$ in the sense that it minimizes $d_{KL}(\mathbf{p}, \mathbf{q})$? By minimizing $d_{KL}(\mathbf{p}, \mathbf{q})$ using Lagrange multipliers, one easily obtains the long-known result that $a_i = p_{i\cdot}$, $b_{i'} = p_{\cdot i'}$ and $d_{KL}(\mathbf{p}, \mathbf{q}) = I$, the mutual information of \mathbf{p} . In other words, even if we had never heard of marginal distributions or mutual information, we could derive them all from d_{KL} : the best factorization simply uses the marginal distributions, and the mutual information of a bivariate distribution is simply the KL-measure of how non-separable it is.

This means that the optimal factorization given k and k' is simply the one giving the marginal conditional distributions

$$M_{ji}^A = \frac{P_{kk'j\cdot}}{P_{kk'\cdot\cdot}}, \quad M_{j'i'}^B = \frac{P_{kk'j'i'}}{P_{kk'\cdot\cdot}} \tag{41}$$

and the corresponding integration is simply

$$\phi^{xtkk} = I(\mathbf{x}_1^A, \mathbf{x}_1^B | \mathbf{x}_0) \tag{42}$$

ϕ^{xtkk} is identical. We can alternatively obtain this result directly from eq (33) by noting that the $I(\mathbf{x}_0^A, \mathbf{x}_0^B)$ -term vanishes now that the state \mathbf{x}_0 is known.

This result highlights a striking and arguably undesirable feature of measures based on the x-factorization option: they vanish for any deterministic system! If the system is deterministic and the present state \mathbf{x}_0 is known, then the future state \mathbf{x}_1 is also known, so all entropies in eq (42) vanish and we obtain $\phi = 0$. With ϕ -measures based on x-factorization, the only source of integration is therefore correlated noise generated by the system.

Minimizing integration on average

Let us now turn to our final factorization option, “a”, where we pick the state-independent factorization that minimizes integration on average. Given the present state $\mathbf{x}_0 = kk'$, let us compare the exact and approximate future probability distributions

$$\begin{aligned} p_{j'i'} &= P(\mathbf{x}_1 = j'i' | \mathbf{x}_0 = kk') = \mathbf{M}_{kk'j'i'} \\ q_{j'i'} &= P(\mathbf{x}_1^A = j | \mathbf{x}_0^A = k) P(\mathbf{x}_1^B = i' | \mathbf{x}_0^B = k') = \mathbf{M}_{kj}^A \mathbf{M}_{k'i'}^B \end{aligned} \tag{43}$$

by computing their KL-divergence $\phi = d_{KL}(p, q)$. The answer clearly depends on the present state kk' , and we saw in the previous section what happens when we minimize separately for each state kk' . Let us now instead average $d_{KL}(p, q)$ over all current states and find the state-independent factorization that minimizes this average:

$$\begin{aligned} \langle d_{KL}(p, q) \rangle &= \sum_{kk'} P(\mathbf{x}_0 = kk') d_{KL}(p, q) |_{\mathbf{x}_0 = kk'} \\ &= \sum_{kk'} p_{kk'\cdot\cdot} \sum_{j'i'} \mathbf{M}_{kk'j'i'} \log \frac{\mathbf{M}_{kk'j'i'}}{\mathbf{M}_{kj}^A \mathbf{M}_{k'i'}^B} \end{aligned} \tag{44}$$

Substituting eq (6) shows that this expression is identical to that from eq (29), so minimizing it gives the exact same optimal factors \mathbf{M}^A and \mathbf{M}^B and the exact same minimum ϕ . The comparison option “t” gives the same result as well, so in conclusion, although they appear quite different from their definitions, the factorization options “o” and “a” are in fact identical.

The full taxonomy

Now that we have derived the explicit form of all our factorization options, we can complete our integration measure classification. Our taxonomy is determined by four choices: of factorization (n/m/o/x/a), variable selection (t/f/a/p), conditioning (u/s/k) and distance measure (k/1/2/h/s/e/m). Although this nominally gives $5 \times 4 \times 3 \times 7 = 420$ different integration measures, most of these options turn out to be zero, undefined or identical to other options. For noising factorizations (factorization options n and m), subsystem B is randomized, so the only well-defined options are ϕ^{nas^*} , ϕ^{nak^*} , ϕ^{nps^*} , ϕ^{npx^*} , ϕ^{mas^*} , ϕ^{mak^*} , ϕ^{mps^*} and ϕ^{mpk^*} , where $*$ denotes any option for the distance measure. For o -factorization, we find that $\phi^{oau^*} = \phi^{oas^*} = \phi^{opu^*} = \phi^{ops^*} = 0$ and $\phi^{otk^*} = \phi^{ofk^*}$. For x -factorization, $\phi^{xt^{**}}$ is undefined and one easily shows that $\phi^{xak^*} = \phi^{xpk^*} = 0$, $\phi^{xau^*} = \phi^{xas^*}$ and $\phi^{xpu^*} = \phi^{xps^*}$. We interpret k -conditioning as \mathbf{x}_0 being known for o -factorization and as \mathbf{x}_0^A being known for noising factorizations, since the reverse options vanish and are undefined, respectively.

Whereas there are strong interactions between the factorization, variable selection and conditioning, we can freely choose any of the 7 distance measures independently of the other choices without changing whether ϕ vanishes or is well-defined. We consider the option k (KL-divergence) by default below since it results in the simplest and most intuitive formulas; the formulas for the other options are straightforward to derive by combining Tables 3, 4 and 5. This leaves us with only the 21 separate options shown in Table 2 to consider. To provide intuition for these formulas, let us recapitulate key definitions in words:

- ϕ^M is the KL-divergence of the two-state probability distribution and the best separable approximation.
- ϕ^{MD} is a measure of how much less information the present gives about the past if factorized dynamics is assumed.
- $\phi_{kk'}^M$ is the KL divergence between (i) the future of the whole given the specific present state of the whole, and (ii) the product of this for the parts calculated separately.
- ϕ^{oak} is the KL divergence between (i) the distribution for the future state of subset A given the current state of A and (ii) the distribution for the future state of subset A given the current state of the whole system.
- ϕ^{opk} is ϕ^{oak} swapping “future” for “past”.
- The subsequent ones are versions from above with different factorizations applied.

Which integration measures are best?

Table 1 summarizes the desirable and undesirable traits for each of these integration measures, showing that merely a handful lack any major drawbacks. Let us now rate the various options in more detail.

For the choice of **probability distance measure** (k/1/2/h/s/e/m), option “e” (the Earth-Mover’s distance d_{EM} used in $\phi^{3.0}$ [11]) remains an attractive candidate for discrete distributions with small number of bits, but is otherwise computationally unfeasible as we discussed above. All options in Table 1 except $\phi^{3.0}$ and ϕ^{MD} therefore use option “k” (the KL-divergence). Note that whether it is an advantage for the probability distance measure to be symmetric (as advocated in [11]) depends on the interpretational context. For example, there is nothing asymmetric about the mutual information that ends up defining ϕ^M in Table 2.

For the choice of **factorization** (n/m/o/x/a), we can quickly dispense with option “a” (for being identical to “o”) and option “x” (because it has the highly undesirable property of always

vanishing for deterministic systems). Which of the remaining options (n/m/o) is preferable depends on other choices. If one wishes to use a distance measure other than the KL-divergence, then the noising options “n” or “m” are computationally preferable, since the optimal factorization “o” can no longer be found analytically. Otherwise, “m” is arguably inferior to “o” because it is no simpler to evaluate and can overestimate the integration as described above. If one has a philosophical preference for the factorization depending only on the mechanism \mathbf{M} and not on any other information about state probabilities, then “n” is the only choice. If one wishes to consider continuous systems, on the other hand, “n” is undefined. In summary, the best factorizations are therefore “o” and “n”, depending on preferences. In practice, numerical experiments show that “n”, “m” and “o” usually give quite similar ϕ -values for a wide range of \mathbf{M} -matrices and probability distributions, so the choice between the three is a relatively minor one.

Turning now to the choice **variable selection** and **conditioning**, Table 1 shows that many of the otherwise well-defined integration measures from Table 2 have serious flaws.

Neither ϕ^{ots} and ϕ^{ofs} are guaranteed to vanish for separable systems, which means that we cannot in good conscience interpret them as measures of integration. Numerical experiments show that ϕ^{nas} , ϕ^{nps} , ϕ^{mas} and ϕ^{nps} tend to be extremely small in practice (ϕ^{mas} is plotted in Fig 2). This is because they differ little from the corresponding measures using optimal factorization (ϕ^{oas} and ϕ^{ops}), which always vanish. In other words, they are not really measures of

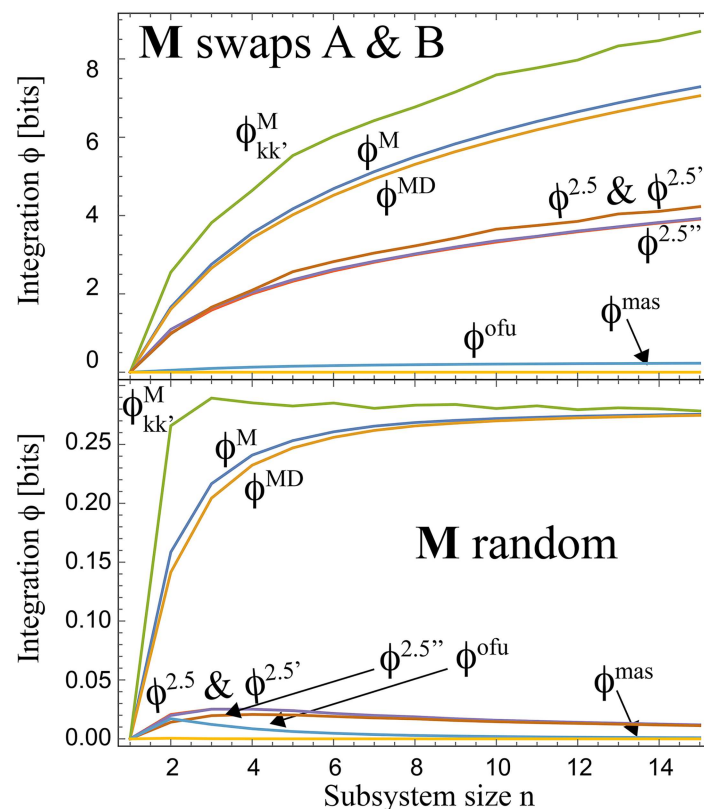


Fig 2. Numerical comparison of different integration measures, averaged over 3,000 random trials. In the bottom panel, all elements of p are independently drawn from a uniform distribution and normalized to sum to unity. In the top panel, only $p^{(0)}$ is randomly generated, and \mathbf{M} is defined so as to swap the two subsystems, i.e., $M_{ij} = \delta_{ij} \delta_{i'j'}$.

doi:10.1371/journal.pcbi.1005123.g002

integration, merely measures of how suboptimal the factorizations “ n ” and “ m ” are. For brevity, we have included merely three of these six flawed measures in Table 1.

Fig 2 shows that ϕ^{ofu} also tends to be much smaller than some other integration measures. We can intuitively understand this by recalling that $\phi^{\text{oua}} = 0$, which means that optimal factorization lets us predict the future marginal distributions for A and B perfectly. Since ϕ^{ofu} quantifies the inability of optimal factorization to predict the full future distribution, we expect that it will at most be of the order of $I(\mathbf{x}_1^A, \mathbf{x}_1^B)$, the extent to which this distribution is not separable (determined by its marginal distributions). For randomly generated probability distributions generated as in Fig 2, one can show that $I(\mathbf{x}_1^A, \mathbf{x}_1^B) \rightarrow 1 - 1/2 \ln 2 \approx 0.28$ bits in the limit where $n \rightarrow \infty$, and numerical experiments indicate that ϕ^{ofu} is never much larger than this value for any p .

Dispensing with flawed/problematic ϕ -measures narrows our list of remaining top candidates to merely nine: ϕ^{otu} , ϕ^{otum} , ϕ^{ofk} , ϕ^{oak} , ϕ^{opk} , ϕ^{nak} , ϕ^{npk} , ϕ^{mak} and ϕ^{mpk} . Moreover, the last six can be elegantly combined into merely three even better ones. As we discussed above, they have the advantage that they vanish for either afferent or efferent systems.

By following the prescription of [10] and taking the minimum of two such complementary measures, we can construct an even better one that vanishes for both afferent and efferent systems. All three of these improved measures are listed in Table 2. The first is $\phi^{2.5} \equiv \min\{\phi^{\text{nak}}, \phi^{\text{npk}}\}$. We denote it “2.5” because it combines attractive features of both IIT2.0 and IIT3.0: it starts with the ϕ^{npk} , which is precisely the IIT2.0 measure, and improves it by taking the minimum of cause/effect integration in the spirit of IIT3.0 (but retaining the KL-divergence of IIT2.0 instead of the harder-to-compute Earth-mover’s distance of IIT3.0). The second is $\phi^{2.5'} \equiv \min\{\phi^{\text{mak}}, \phi^{\text{mpk}}\}$, which has the advantage of remaining defined even for continuous variables. The third is $\phi^{2.5''} \equiv \min\{\phi^{\text{oak}}, \phi^{\text{opk}}\}$, which uses the optimal factorization.

How large can ϕ get?

In summary, our taxonomy of ϕ -measures produces merely a handful of truly attractive options: $\phi^{2.5}$, $\phi^{2.5'}$, $\phi^{2.5''}$, $\phi^{3.0}$, ϕ^{MD} , ϕ^M and $\phi_{kk'}^M$. Fig 2 shows examples of what they evaluate to numerically. The lower panel shows that for randomly generated probability distributions, none of them exceed $1 - 1/2 \ln 2 \approx 0.28$ bits on average, which as mentioned above is the mutual information in a random bivariate distribution. However, $\phi^{2.5}$, $\phi^{2.5'}$, $\phi^{2.5''}$, ϕ^M , ϕ^{MD} and $\phi_{kk'}^M$ can get arbitrarily large for some systems, as illustrated in the top panel, growing logarithmically with the size n of the subsystems A and B. In other words, the maximum integration is of the order of the number of subsystem bits. For the example shown where the dynamics merely swaps the two subsystems, we obtain $\phi^{2.5} = \log_2 n$, because noising gives $M^A = 1/n$, $q = 1/n^2$ and p is a Kronecker δ . ϕ^M , ϕ^{MD} and $\phi_{kk'}^M$ are seen to give about twice the integration for this example.

Note that although this dynamics \mathbf{M} that merely swaps the subsystems has such a large ϕ -value only for this particular cut that separates the systems being swapped. Consider, for example, a system of four bits labeled 1, 2, 3 and 4, where the dynamics swaps 1 with 3 and 2 with 4. There is a different cut where $\phi = 0$: simply define the new subsystems A' and B' to be the first and second halves of the A and B-systems, i.e., $A' = 1, 3$ and $B' = 2, 4$. The swapping is now carried out internally within A' and B', revealing that there is no integration and upper-case $\Phi = 0$.

However, there are plenty of systems for which even the true integration Φ grows like the number of subsystem bits, $\log_2 n$. A simple example accomplishing this (in the spirit of the random coding example in [12]) is when the n^4 probabilities $p_{i'j'}$ are all set to zero except for a

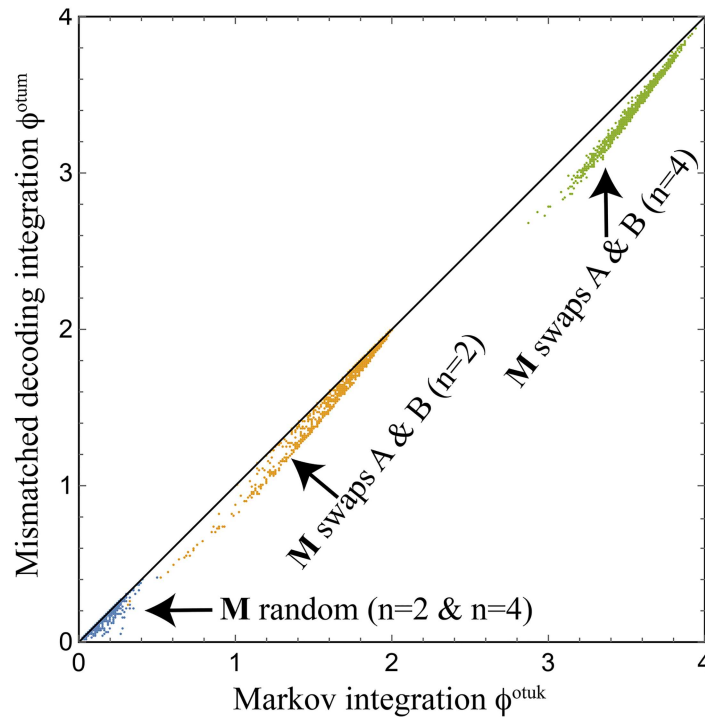


Fig 3. Numerical comparison of the two measures ϕ^{otuk} and ϕ^{ofum} for 3,000 random trials, generated the same way as in Fig 2. The two measures are seen to be rather similar for these examples, and to satisfy the inequality $\phi^{ofum} \leq \phi^{otuk}$.

doi:10.1371/journal.pcbi.1005123.g003

randomly selected subset of n^2 of them that are set to $1/n^2$. Now $\phi^M \sim \log_2 n$ even when minimized over all bipartitions of the $2\log_2 n$ bits in the system. For this example, we have $S(\mathbf{x}) = \log_2 n^2 = 2 \log_2 n$. The marginal distributions for $\mathbf{x}^A, \mathbf{x}^B, \mathbf{x}_0^A$ and \mathbf{x}_0^B are all rather uniform, with entropy on average less than a bit from the value for a uniform distribution, giving $S(\mathbf{x}^A) \sim S(\mathbf{x}^B) \sim \log_2 n^2, S(\mathbf{x}_0^A) \sim S(\mathbf{x}_0^B) \sim \log_2 n, I(\mathbf{x}^A, \mathbf{x}^B) = S(\mathbf{x}^A) + S(\mathbf{x}^B) - S(\mathbf{x}) \sim 2 \log_2 n, I(\mathbf{x}_0^A, \mathbf{x}_0^B) = S(\mathbf{x}_0^A) + S(\mathbf{x}_0^B) - S(\mathbf{x}_0) \sim 0$ and therefore $\phi^M = I(\mathbf{x}^A, \mathbf{x}^B) - I(\mathbf{x}_0^A, \mathbf{x}_0^B) \sim 2 \log_2 n \sim \log_2 n$.

Fig 3 shows that the measures ϕ^M and ϕ^{MD} can sometimes be quite similar: they give numerically similar values for the 3,000 random examples shown. Moreover, they appear to satisfy the inequality $\phi^{ofum} \leq \phi^{otuk}$. Further examination shows that for these these random examples, the β -complication in eq (24) makes essentially no perceptible difference in practice, in the sense that the computation of ϕ^{MD} can be accurately accelerated by setting $\beta = 1$ rather than minimizing over it. However, [15] shows that there are real-world cases where β is far from unity and also where $\phi^{MD} \ll \phi^M$, particularly when noise correlations dominate over causal correlations. To understand this, consider the extreme case of two perfectly correlated bits that are independently randomized by both time 0 and time 1, so that $\mathbf{x}_0^A = \mathbf{x}_0^B$ and $\mathbf{x}_1^A = \mathbf{x}_1^B$, with no correlation between the two times. Then $\phi^{MD} = 0$ whereas $\phi^M = I(\mathbf{x}^A, \mathbf{x}^B) - I(\mathbf{x}_0^A, \mathbf{x}_0^B) = 2 - 1 = 1$, which is arguably undesirable.

The $n \rightarrow \infty$ limit of continuous variables

All our previous results are fully general, applying regardless of whether the variables are discrete (such as bits that equal zero or one) or continuous (such as voltages or other variables

measured in fMRI, EEG, MEG or electrophysiology studies). We can view the latter as the $n \rightarrow \infty$ limit of the former, since a single real number can be represented as an infinite string of bits. In this section, we will focus on the continuous case and see how our previous formulas can be greatly simplified by assuming Gaussianity. We therefore replace i, i', j and j' in all our formulas by $\mathbf{x}_0^A, \mathbf{x}_0^B, \mathbf{x}_1^A$ and \mathbf{x}_1^B , respectively, and replace all sums by integrals.

How Gaussianity gives linearity. To make things tractable, we will make one strong but very useful assumption: that \mathbf{x} has a Gaussian distribution. The most general d -dimensional multivariate Gaussian distribution is parametrized by its mean vector $\mathbf{m} \equiv \langle \mathbf{x} \rangle$ and covariance matrix $\mathbf{T} \equiv \langle \mathbf{x}\mathbf{x}^t \rangle - \mathbf{m}\mathbf{m}^t$ and takes the form

$$g[\mathbf{x}; \mathbf{m}, \mathbf{T}] \equiv \frac{1}{(2\pi)^{d/2} |\mathbf{T}|^{1/2}} e^{-\frac{1}{2}(\mathbf{x}-\mathbf{m})^t \mathbf{T}^{-1}(\mathbf{x}-\mathbf{m})}, \tag{45}$$

so we are making the assumption that there is some \mathbf{m} and \mathbf{T} such that $p(\mathbf{x}) = g(\mathbf{x}; \mathbf{m}, \mathbf{T})$. Let us write \mathbf{m} and \mathbf{T} as

$$\mathbf{m} = \begin{pmatrix} \mathbf{m}_0 \\ \mathbf{m}_1 \end{pmatrix}, \quad \mathbf{T} = \begin{pmatrix} \mathbf{C}_0 & \mathbf{B} \\ \mathbf{B}^t & \mathbf{C}_1 \end{pmatrix}, \tag{46}$$

where \mathbf{m}_i and \mathbf{C}_i are the mean and covariance of \mathbf{x}_i , respectively.

Interpreting the sum in the denominator of eq (6) as an integral and evaluating it gives

$$\mathbf{M}(\mathbf{x}_1, \mathbf{x}_0) = g[\mathbf{x}_1; \mathbf{m}_1 + \mathbf{A}(\mathbf{x}_0 - \mathbf{m}_0), \Sigma], \tag{47}$$

where

$$\mathbf{A} \equiv \mathbf{B}^t \mathbf{C}_0^{-1}, \tag{48}$$

$$\Sigma \equiv \mathbf{C}_1 - \mathbf{B}^t \mathbf{C}_0^{-1} \mathbf{B} = \mathbf{C}_1 - \mathbf{A} \mathbf{C}_0 \mathbf{A}^t. \tag{49}$$

The following well-known matrix identities are useful in the derivation of this and other matrix results in this paper:

$$\begin{vmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{vmatrix} = |\mathbf{A}\mathbf{D} - \mathbf{A}\mathbf{C}\mathbf{A}^{-1}\mathbf{B}|, \tag{50}$$

$$\begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{pmatrix}^{-1} = \begin{pmatrix} [\mathbf{A} - \mathbf{B}\mathbf{D}^{-1}\mathbf{C}]^{-1} & -\mathbf{A}^{-1}\mathbf{B}[\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}]^{-1} \\ [\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}]^{-1}\mathbf{C}\mathbf{A}^{-1} & [\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}]^{-1} \end{pmatrix}, \tag{51}$$

$$[\mathbf{A} + \mathbf{B}\mathbf{D}^{-1}\mathbf{C}]^{-1} = \mathbf{A}^{-1} - \mathbf{A}^{-1}\mathbf{B}[\mathbf{D} + \mathbf{C}\mathbf{A}^{-1}\mathbf{B}]\mathbf{C}\mathbf{A}^{-1}. \tag{52}$$

Eq (47) encodes the well-known result that the conditional distribution $\mathbf{x}_1 | \mathbf{x}_0$ for Gaussian variables is Gaussian with mean $\mathbf{m}_1 + \mathbf{B}\mathbf{C}_0^{-1}(\mathbf{x}_0 - \mathbf{m}_0)$ and covariance matrix $\mathbf{C}_1 - \mathbf{B}^t \mathbf{C}_0^{-1} \mathbf{B}$. These equations embody a remarkable simplicity that we can exploit. First of all, the covariance matrix Σ is independent of \mathbf{x}_0 , which allows us to interpret \mathbf{x}_1 as simply a function of \mathbf{x}_0 plus a random noise vector \mathbf{n} that is independent of \mathbf{x}_0 . Second, this function is affine, involving simply a linear term plus a constant. In other words, we can write

$$\mathbf{x}_1 = \mathbf{m}_1 + \mathbf{A}(\mathbf{x}_0 - \mathbf{m}_0) + \mathbf{n}, \tag{53}$$

where the noise vector \mathbf{n} satisfies

$$\langle \mathbf{n} \rangle = 0, \quad \langle \mathbf{n}\mathbf{x}^t \rangle = 0, \quad \langle \mathbf{n}\mathbf{n}^t \rangle = \Sigma. \quad (54)$$

It is worth reflecting on how remarkable this is, since it is easy to overlook. The future state \mathbf{x}_1 of a system can depend on the present state \mathbf{x}_0 in some arbitrarily complicated non-linear way. Moreover, for a generic Markov process, the scatter of \mathbf{x}_1 around its mean will depend strongly on \mathbf{x}_0 . Yet as long as all probability distributions are Gaussian, which is often a useful approximation for laboratory data, both of these complications vanish and we are left with the simple linear dynamics of eq (53).

Autoregressive processes. Let us now briefly review the formalism of so-called autoregressive processes and how it relates to our problem at hand. A simple special case of the above is where the random process is *stationary*, i.e., where the statistical properties are independent of time. This implies that $\mathbf{m}_i = \mathbf{m}$ and $\mathbf{C}_i = \mathbf{C}$ for some \mathbf{m} and \mathbf{C} that are independent of i . For a stationary process, it is convenient to redefine new zero-mean variables $\mathbf{x}'_i \equiv \mathbf{x}_i - \mathbf{m}$. Dropping the prime for simplicity, this allows us to rewrite eq (53) as

$$\mathbf{x}_{i+1} = \mathbf{A}\mathbf{x}_i + \mathbf{n}_i, \quad (55)$$

where the noise vectors \mathbf{n}_i have vanishing mean and vanishing correlations between different times, i.e., $\langle \mathbf{n}_i, \mathbf{n}_j^t \rangle = \delta_{ij}\Sigma$. The covariance matrix between vectors at two subsequent times is therefore

$$\langle \mathbf{x}\mathbf{x}^t \rangle \equiv \begin{pmatrix} \mathbf{C} & \mathbf{C}\mathbf{A}^t \\ \mathbf{A}\mathbf{C} & \mathbf{A}\mathbf{C}\mathbf{A}^t + \Sigma \end{pmatrix}, \quad \mathbf{x} \equiv \begin{pmatrix} \mathbf{x}_0 \\ \mathbf{x}_1 \end{pmatrix}. \quad (56)$$

Even if the random process is not stationary initially, it will eventually converge to a stationary state where covariance is time-independent as long as all eigenvalues of \mathbf{A} have magnitude below unity, so that memory of the past gets exponentially damped over time. Once the covariance has become time-independent, eq (56) implies that $\mathbf{C} = \mathbf{A}\mathbf{C}\mathbf{A}^t + \Sigma$. This is known as the Lyapunov equation, and is readily solved by special-purpose techniques or, rapidly enough, by simply iterating it to convergence. If we write the covariance matrix $\langle \mathbf{x}\mathbf{x}^t \rangle$ measured from actual time series data as

$$\mathbf{T} \equiv \langle \mathbf{x}\mathbf{x}^t \rangle = \begin{pmatrix} \mathbf{C} & \mathbf{B} \\ \mathbf{B}^t & \mathbf{C} \end{pmatrix}, \quad (57)$$

then equating it with eq (56) lets us compute the matrices we need from the data:

$$\mathbf{A} = \mathbf{B}^t\mathbf{C}^{-1}, \quad (58)$$

$$\Sigma = \mathbf{C} - \mathbf{A}\mathbf{C}\mathbf{A}^t = \mathbf{C} - \mathbf{B}^t\mathbf{C}^{-1}\mathbf{B}. \quad (59)$$

These equations hold regardless of whether the probability distributions are Gaussian or not. If the noise \mathbf{n} is Gaussian, then all distributions will be Gaussian in the steady state, so this is an alternative way of deriving eqs (67) and (49) (without the subscripts).

We saw above how we can equally well interpret our system as a Markov process operating backward in time, where the future causes the past. Repeating the above derivation for this case, we can write

$$\mathbf{x}_{i-1} = \tilde{\mathbf{A}}\mathbf{x}_i + \mathbf{n}_i, \quad (60)$$

where

$$\tilde{\mathbf{A}} = \mathbf{B}\mathbf{C}^{-1} = \Sigma\mathbf{A}^t\tilde{\Sigma}^{-1}, \quad (61)$$

$$\begin{aligned} \tilde{\Sigma} &= \mathbf{C} - \mathbf{B}\mathbf{C}^{-1}\mathbf{B}^t = \mathbf{C} - \tilde{\mathbf{A}}\mathbf{C}\tilde{\mathbf{A}}^t, \\ &= [\mathbf{C}^{-1} + \mathbf{A}^t\Sigma^{-1}\mathbf{A}]^{-1} = \mathbf{C} - \mathbf{C}\mathbf{A}^t\mathbf{C}^{-1}\mathbf{A}\mathbf{C}. \end{aligned} \quad (62)$$

Although the matrices Σ and $\tilde{\Sigma}$ are different, it is easy to prove that their determinants are identical, which means that the conditional entropy is the same both forward and backward in time.

Optimal factorization. In summary, a Markov process $\mathbf{p}_1 = \mathbf{M}\mathbf{p}$ can be described much more simply when all probability distributions are Gaussian: instead of keeping track of the infinite-dimensional Markov matrix \mathbf{M} or the infinite-dimensional rank-4 tensor \mathbf{p} (both of which have as indices the four continuous variables $\mathbf{x}_0^A, \mathbf{x}_0^B, \mathbf{x}_1^A, \mathbf{x}_1^B$), we merely need to keep track of the $2n \times 2n$ covariance matrix \mathbf{T} , from which we can compute and quantify the deterministic and stochastic parts of the dynamics as the matrices \mathbf{A} and Σ , respectively.

Let us now translate the rest of our results from our integration taxonomy into this simpler formalism. To separate out the effects occurring within and between the subsystems A and B, let us name the corresponding blocks of the \mathbf{A} -matrix and the matrix $\mathbf{T} \equiv \langle \mathbf{x}\mathbf{x}^t \rangle$ from eq (46) as follows:

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}_A & \mathbf{A}_{AB} \\ \mathbf{A}_{BA} & \mathbf{A}_B \end{pmatrix}, \quad (63)$$

$$\mathbf{T} = \begin{pmatrix} \mathbf{C}_A & \mathbf{C}_{AB} & \mathbf{B}_A & \mathbf{B}_{AB} \\ \mathbf{C}_{AB}^t & \mathbf{C}_B & \mathbf{B}_{BA} & \mathbf{B}_B \\ \mathbf{B}_A^t & \mathbf{B}_{BA}^t & \mathbf{C}_A & \mathbf{C}_{AB} \\ \mathbf{B}_{AB}^t & \mathbf{B}_B^t & \mathbf{C}_{AB}^t & \mathbf{C}_B \end{pmatrix}, \quad \mathbf{x} = \begin{pmatrix} \mathbf{x}_0^A \\ \mathbf{x}_0^B \\ \mathbf{x}_1^A \\ \mathbf{x}_1^B \end{pmatrix}. \quad (64)$$

Analogously to how eq (6) gave us eq (47), eq (31) now gives the optimal factorization

$$\mathbf{M}^A(\mathbf{x}_1^A, \mathbf{x}_0^A) = g[\mathbf{x}_1^A; \hat{\mathbf{A}}_A\mathbf{x}_0^A, \hat{\Sigma}_A], \quad (65)$$

$$\mathbf{M}^B(\mathbf{x}_1^B, \mathbf{x}_0^B) = g[\mathbf{x}_1^B; \hat{\mathbf{A}}_B\mathbf{x}_0^B, \hat{\Sigma}_B], \quad (66)$$

where

$$\hat{\mathbf{A}}_A \equiv \mathbf{B}_A^t\mathbf{C}_A^{-1}, \quad \hat{\Sigma}_A \equiv \mathbf{C}_A - \mathbf{B}_A^t\mathbf{C}_A^{-1}\mathbf{B}_A, \quad (67)$$

$$\hat{\mathbf{A}}_B \equiv \mathbf{B}_B^t\mathbf{C}_B^{-1}, \quad \hat{\Sigma}_B \equiv \mathbf{C}_B - \mathbf{B}_B^t\mathbf{C}_B^{-1}\mathbf{B}_B. \quad (68)$$

In other words, the “o”-factorization approximates $\mathbf{x}_1 = \mathbf{A}\mathbf{x}_0 + \mathbf{n}$ by

$$\hat{\mathbf{x}}_1 \equiv \begin{pmatrix} \hat{\mathbf{x}}_1^A \\ \hat{\mathbf{x}}_1^B \end{pmatrix} \equiv \hat{\mathbf{A}}\mathbf{x}_0 + \hat{\mathbf{n}}, \quad \hat{\mathbf{A}} \equiv \begin{pmatrix} \hat{\mathbf{A}}_A & 0 \\ 0 & \hat{\mathbf{A}}_B \end{pmatrix}, \quad (69)$$

where the noise vector $\hat{\mathbf{n}}$ has zero mean and covariance matrix

$$\hat{\Sigma} \equiv \begin{pmatrix} \hat{\Sigma}_A & 0 \\ 0 & \hat{\Sigma}_B \end{pmatrix}. \tag{70}$$

We see that tensor factorization in the previous section now corresponds to the matrices \mathbf{A} and Σ being block-diagonal.

Noising factorization. Eq (55) tells us that

$$\begin{pmatrix} \mathbf{x}_1^A \\ \mathbf{x}_1^B \end{pmatrix} = \mathbf{A}\mathbf{x}_0 + \mathbf{n} = \begin{pmatrix} \mathbf{A}_A\mathbf{x}_0^A + \mathbf{A}_{AB}\mathbf{x}_0^B + \mathbf{n}^A \\ \mathbf{A}_B\mathbf{x}_0^B + \mathbf{A}_{BA}\mathbf{x}_0^A + \mathbf{n}^B \end{pmatrix}. \tag{71}$$

The idea with noising is to take the terms $\mathbf{A}_{AB}\mathbf{x}_0^B$ and $\mathbf{A}_{BA}\mathbf{x}_0^A$ and reinterpret them not as signal but as noise, with zero mean and uncorrelated with anything else. The noising option “n” is unfortunately undefined for this continuous-variable case, because it says to use a uniform distribution for these noised versions of \mathbf{x}_0^A and \mathbf{x}_0^B , which has infinite variance and hence gives, e.g., $\langle \mathbf{x}_0^B \mathbf{x}_0^{Bt} \rangle = \infty$ when \mathbf{x}_0^B is noised. The mild noising option “m”, however, remains well-defined, saying to use the actual distributions for these noised versions of \mathbf{x}_0^A and \mathbf{x}_0^B , hence giving $\langle \mathbf{x}_0^A \mathbf{x}_0^{At} \rangle = \mathbf{C}_A$ and $\langle \mathbf{x}_0^B \mathbf{x}_0^{Bt} \rangle = \mathbf{C}_B$ when these variables are noised.

Computing the first and second moments of eq (71) therefore tells us that “m”-factorization approximates $\mathbf{x}_1 = \mathbf{A}\mathbf{x}_0 + \mathbf{n}$ by

$$\bar{\mathbf{x}}_1 \equiv \begin{pmatrix} \bar{\mathbf{x}}_1^A \\ \bar{\mathbf{x}}_1^B \end{pmatrix} \equiv \bar{\mathbf{A}}\mathbf{x}_0 + \bar{\mathbf{n}}, \quad \bar{\mathbf{A}} \equiv \begin{pmatrix} \mathbf{A}_A & 0 \\ 0 & \mathbf{A}_B \end{pmatrix}, \tag{72}$$

where the noise vector $\bar{\mathbf{n}}$ has zero mean and covariance matrix

$$\bar{\Sigma} \equiv \begin{pmatrix} \Sigma_A + \mathbf{A}_{AB}\mathbf{C}_B\mathbf{A}_{AB}^t & 0 \\ 0 & \Sigma_B + \mathbf{A}_{BA}\mathbf{C}_A\mathbf{A}_{BA}^t \end{pmatrix}. \tag{73}$$

Note that in contrast to the “o”-factorization of eq (69), the “m”-factorization has no tildes on the \mathbf{A}_A and \mathbf{A}_B -matrices in eq (72).

Results. We now have all the tools we need to derive the Gaussian versions of the ϕ -formulas in Table 2. Starting with eq (34), interpreting the sum in eq (30) as an integral and performing it when p is the Gaussian distribution of eq (45) gives the well-known formula

$$I(\mathbf{x}_A, \mathbf{x}_B) = \frac{1}{2} \log \frac{|\mathbf{T}_A| |\mathbf{T}_B|}{|\mathbf{T}|} \tag{74}$$

for the mutual information between two multivariate Gaussian random variables. This immediately gives the five matrix formulas for ϕ^M , ϕ^B , ϕ^{ots} , ϕ^{ofs} and ϕ^{xfk} in the right column of Table 2. The second version listed for ϕ^B is also given in [13].

Starting with the KL-divergence definition $d_{\text{KL}}(p, q) \equiv \sum_i p_i \log \frac{p_i}{q_i}$ from Table 5, we again interpret the sum as an integral and use eq (45). This gives the well-known formula

$$D_{\text{KL}}(f_p, f_q) = \frac{1}{2} \left[\Delta \mathbf{m}^t \mathbf{C}_q^{-1} \Delta \mathbf{m} + \text{tr} \mathbf{C}_q^{-1} \mathbf{C}_p + \ln \frac{|\mathbf{C}_q|}{|\mathbf{C}_p|} - n \right] \tag{75}$$

for the KL-divergence between two Gaussian probability distributions f_p and f_q with means \mathbf{m}_i

and covariance matrices C_i ($i = p, q$), where $\Delta \mathbf{m} \equiv \mathbf{m}_p - \mathbf{m}_q$. The first term in eq (75) thus represents the mismatch between the means and the remainder (which is also guaranteed to be nonnegative) represents the mismatch between the covariances.

For ϕ^{ofu} , the future distribution $p(\mathbf{x}_1)$ with mean zero and covariance matrix \mathbf{C} is approximated by the distribution $q(\mathbf{x}_1)$ that has mean zero and covariance matrix $\hat{\mathbf{A}}\hat{\mathbf{C}}\hat{\mathbf{A}}^t + \hat{\Sigma}$, which follows from eqs (69) and (70). Substituting these means and covariance matrices into eq (75) gives the matrix formula for ϕ^{ofu} in Table 2. For ϕ^{mas} , both means again vanish, but now the future distribution $p(\mathbf{x}_1^A)$ has covariance matrix \mathbf{C}_A while the approximation $q(\mathbf{x}_1^A)$ has covariance matrix $\Sigma_A + \mathbf{A}_A\mathbf{C}_A\mathbf{A}_A^t + \mathbf{A}_{AB}\mathbf{C}_B\mathbf{A}_{AB}^t$, which follows from eqs (72) and (73).

For the remaining options in Table 2, i.e., ϕ^{ofk} , ϕ^{oak} , ϕ^{opk} , ϕ^{mak} and ϕ^{mpk} , the means do not vanish, since they reflect information about the known state. For ϕ^{ofk} , the future distribution $p(\mathbf{x}_1)$ with mean $\mathbf{A}\mathbf{x}_0$ and covariance matrix Σ is approximated by the distribution $q(\mathbf{x}_1)$ that has mean $\hat{\mathbf{A}}\mathbf{x}_0$ and covariance matrix $\hat{\Sigma}$, so eq (75) gives the matrix formula for ϕ^{ofk} in the table. For ϕ^{oak} , the future distribution $p(\mathbf{x}_1^A)$ has mean $\mathbf{A}_A\mathbf{x}_0^A + \mathbf{A}_B\mathbf{x}_0^B$ and covariance matrix Σ , while the approximation $q(\mathbf{x}_1^A)$ has mean $\hat{\mathbf{A}}_A\mathbf{x}_0^A$ and covariance matrix $\hat{\Sigma}_A$. Finally, for ϕ^{mak} , the future distribution $p(\mathbf{x}_1^A)$ with mean $\hat{\mathbf{A}}_A\mathbf{x}_0^A$ and covariance matrix $\hat{\Sigma}_A$ is approximated by $q(\mathbf{x}_1^A)$ with mean $\bar{\mathbf{A}}_A\mathbf{x}_0^A$ and covariance matrix $\bar{\Sigma}_A$. The time-reversed measures ϕ^{opk} , ϕ^{mps} and ϕ^{mpk} are identical to ϕ^{oak} , ϕ^{mas} and ϕ^{mak} , but with \mathbf{A} and Σ replaced by their time-reversed versions $\tilde{\mathbf{A}}$ and $\tilde{\Sigma}$ from eq (61).

Substituting the above Gaussian formulas into eqs (24) and (25) gives

$$\phi^{MD} = \frac{\ln \frac{|\hat{\mathbf{C}}|}{|\hat{\Sigma}|} + \text{tr} \left[\hat{\mathbf{C}}^{-1}\mathbf{C} - \hat{\Sigma}^{-1}\Sigma - \hat{\mathbf{A}}^t\hat{\Sigma}^{-1}\hat{\mathbf{A}}\mathbf{C} \right]}{2}, \tag{76}$$

where $\hat{\mathbf{A}} \equiv \mathbf{A} - \hat{\mathbf{A}}$ and $\hat{\mathbf{C}} \equiv \hat{\mathbf{A}}\hat{\mathbf{C}}\hat{\mathbf{A}}^t + \hat{\Sigma}$ for the simple but important case $\beta = 1$.

Graph-theory approximation to make computations feasible

The problem. The ϕ -formulas for discrete variables in the left column of Table 2 require working with the $n \times n$ matrix \mathbf{M} , where $n = 2^b$ for a system of b bits. In other words, the time to evaluate ϕ for a given cut grows exponentially with the system size b , which becomes computationally prohibitive even for modest system sizes such as 100 bits—let alone the set of neurons in the human brain with $b \sim 10^{11}$. Even 300 bits give n greater than the number of particles in our universe.

When the system state is described not by bits but continuous variables (such as voltages or other variables measured in fMRI, EEG, MEG or electrophysiology studies), things get even worse, since representing even a single variable requires an infinite number of bits. However, [13] pointed out that the Gaussian approximation radically simplifies things, and we saw in The $n \rightarrow \infty$ limit of continuous variables how ϕ can then be computed dramatically faster. Not only does the infinity problem go away for most measures in Table 2, but the formulas in the right column are exponentially faster to evaluate than those in the left column even when each bit is replaced by a separate real number! This is because if there are b real numbers, the $n \times n$ matrix \mathbf{T} has $n = 2b$, not $n = 2^b$. This means that ϕ can now be computed in polynomial time, more specifically $O(b^3)$ time, since the slowest matrix operations in Table 2 scale as $O(n^3)$.

Unfortunately, even after this exponential speedup in computing ϕ , computing the upper-case version Φ is still exponentially slow. This is because Φ is the minimum of ϕ over the exponentially many ways of splitting the system into two parts. Even if we limit ourselves to symmetric bipartitions, the number of ways to split an even number n elements into two parts of

size $n/2$ is

$$\binom{n}{n/2} = \frac{n!}{(n/2)!^2} \approx \sqrt{\frac{2}{\pi n}} 2^n, \tag{77}$$

where we have used Stirling’s approximation $n! \approx \sqrt{2\pi n}(n/e)^n$. In other words, examining all symmetric bipartitions is pretty much as exponentially painful as examining all 2^n bipartitions, because most bipartitions are close to symmetric.

An approximate solution. Being able to compute Φ approximately is clearly better than not being able to compute it at all. In this spirit, let us explore an approximation that exponentially accelerates the computation of Φ . Starting with the linear dynamics $\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x} + \mathbf{n}$ from eq (55), let us motivate our approximation by considering the case where the noise is \mathbf{n} uncorrelated (where Σ is diagonal) so that it introduces no correlations between the two systems, regardless of the cut. This means that the only source of integration can be the \mathbf{A} -matrix transferring information between the two subsystems. Let us visualize this information flow as a directed graph (Fig 4, bottom), where each node represents a variable i and each edge represents a non-zero element A_{ij} , i.e., non-zero information flow from element j to element i . If this graph consists of two disconnected parts A and B of equal size, as in the lower right corner of Fig 4, then we clearly have $\Phi = 0$, since there is no information flow and hence no integration between these two parts. In other words, if we permute the elements so that all elements of A precede all elements of B, the matrix \mathbf{A} becomes block-diagonal (Fig 4, middle right), for which all integration measures in the right column of Table 2 will give $\phi = 0$.

Note that before the elements were permuted (Fig 4, top right), this fact that $\phi = 0$ was less obvious. Moreover, examining all $n!$ permutations (or all $\binom{n}{n/2}$ symmetric bipartitions) would have been an enormously inefficient way of finding that best bipartition for which ϕ vanishes. In contrast, finding the connected components of a graph is quite simple, as is evident from staring at Fig 4, with complexity between $O(n)$ and $O(n^2)$. This means that if we know that $\Phi = 0$, then we can find the best bipartition (“cruellest cut”) easily, in polynomial time.

Let us now define an approximation taking advantage of this idea: *replace all unimportant elements $|A_{ij}| < \epsilon$ by zero, and adjust ϵ so that the largest connected component has size as close as possible to $n/2$* . Letting this largest connected component define our approximation of the best bipartition, we now compute its ϕ -value and use this as our approximation for Φ .

Note that this approximation can be trivially generalized to asymmetric bipartitions (the subtle conceptual challenges of how to weight or otherwise handle asymmetric partitions [10, 11, 13, 22] are neither ameliorated nor exacerbated by our fast approximation).

In practice, we determine ϵ by using the interval halving method. A final technical point is that we have two separate definitions of graph connectivity to choose between: weak and strong. A graph is *strongly connected* if you can move between any pair of elements following the directional arrows on the edges. This means that every element can (at least through intermediaries) affect and be affected by every other element, precisely capturing the integration spirit of [10]. Strong connectivity is therefore the logical choice when using our approximation to compute $\Phi^{2.5}$, $\Phi^{2.5'}$, $\Phi^{2.5''}$, since it will reflect their property that integration vanishes for afferent and efferent pathways. A graph is *weakly connected* if you can move between any pair of elements ignoring edge arrows—in other words, if it simply looks connected when drawn. Using weak connectivity is arguably the better approximation for the Φ -measures that do not vanish for afferent/efferent pathways, and numerical experiments confirm this.

Fig 5 illustrates the accuracy of our approximation. For this example, we randomly generate 7,000 different 16×16 matrices \mathbf{A} and compute Φ^M both exactly (as the minimum of ϕ^M over

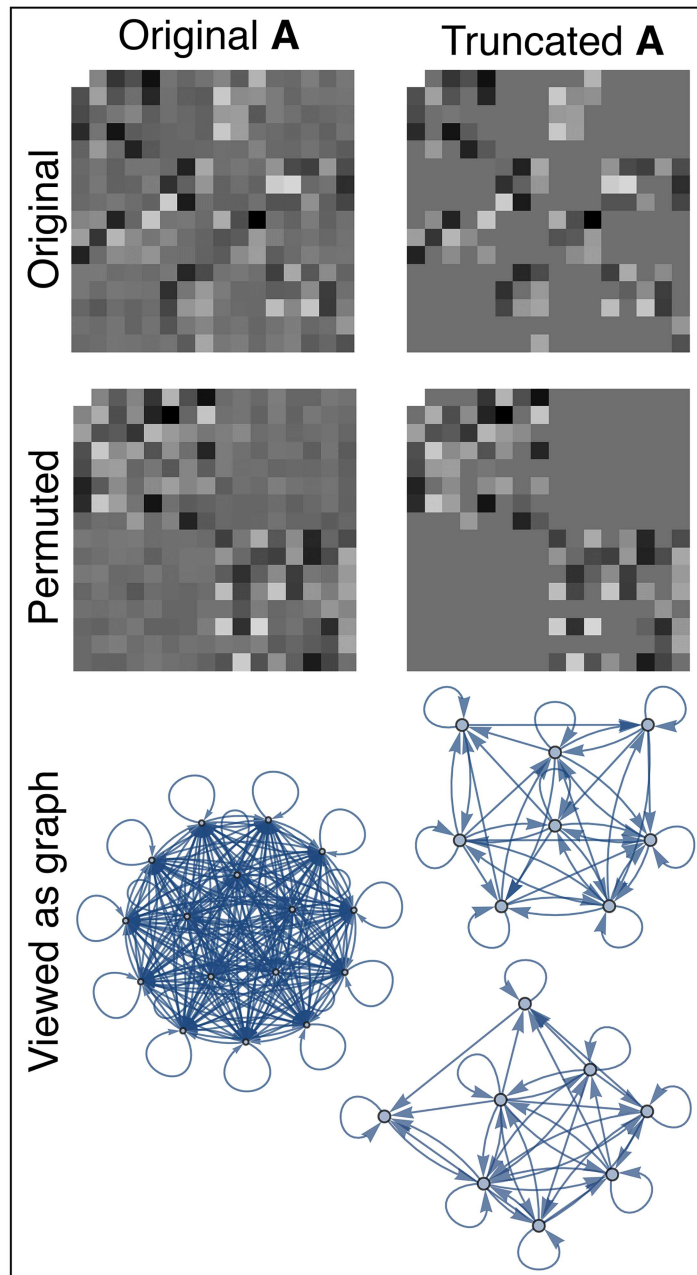


Fig 4. Illustration of our fast Φ -approximation for an $n = 16$ example. The structure of the \mathbf{A} -matrix can be visualized either as a grid (top four examples) where each pixel color shows the value of the corresponding element A_{ij} ranging from the smallest (black) to the largest (white), or as a graph (bottom examples) showing all non-zero matrix elements. Both of the matrices on the left correspond to the same graph below them, and both of the matrices on the right correspond to the same (disconnected) graph below them. Our method zeros all matrix elements $|A_{ij}| < \epsilon$ below the threshold ϵ that makes the largest connected graph component involve merely half of the elements, which in the matrix picture means that there is a permutation of the elements (rows and columns) rendering the matrix block-diagonal (middle right). Whereas it would take exponentially long to try all matrix permutations, graph connectivity can be determined in polynomial time, thus enabling us to rapidly find a good approximation for the “cruellest cut” bipartition.

doi:10.1371/journal.pcbi.1005123.g004

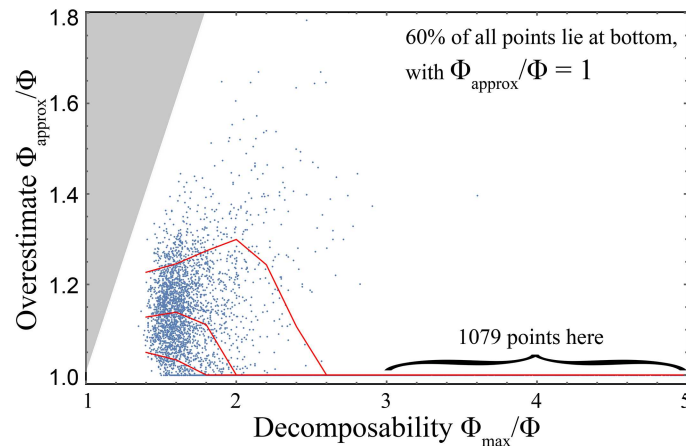


Fig 5. How well our fast Φ -approximation works for 7,000 simulations of the $n = 16 \Phi^M$ -example described in the text. Whereas it is seen to be excellent at finding the best bipartition when not all are comparably good, (*i.e.*, when $\Phi_{\max}/\Phi_{\min} \gg 1$), the approximation is seen to overestimate Φ by up to 15% (the median) when there is no clear winner (left side). From top to bottom, the three curves show the 95th, 50th and 5th percentiles of the overestimation factor. The shaded region delimits the largest overestimation possible, when $\Phi_{\text{approx}} = \Phi_{\max}$.

doi:10.1371/journal.pcbi.1005123.g005

all $\binom{16}{8} = 12,870$ symmetric bipartitions) and using our approximation. Here we generate \mathbf{A} -matrices by first computing

$$\mathbf{A} = \eta \mathbf{A}_0 + \begin{pmatrix} \mathbf{A}_1 & 0 \\ 0 & \mathbf{A}_2 \end{pmatrix}, \quad (78)$$

where \mathbf{A}_0 , \mathbf{A}_1 and \mathbf{A}_2 are random matrices (whose elements are independent Gaussian random variables with zero mean), each normalized to have their largest eigenvalue equal to unity. We then renormalize \mathbf{A} so that its largest eigenvalue equals 0.99. The parameter η controls the typical level of integration: $\eta = 0$ gives $\Phi = 0$ since \mathbf{A} is block-diagonal, whereas $\eta \rightarrow \infty$ gives maximal integration, with no special cut put in by hand; η is randomly chosen to be 0.1, 0.3, 0.5, 0.7, 1, 2 or 10 with equal probability. Once we have generated \mathbf{A} , we compute \mathbf{C} as the solution to the Lyapunov equation $\mathbf{C} = \mathbf{A}\mathbf{C}\mathbf{A}^t + \Sigma$ with $\Sigma = \mathbf{I}$.

For comparison, we also compute the maximum Φ_{\max}^M over the bipartitions. The ratio $\Phi_{\max}/\Phi \geq 1$ (where $\Phi \equiv \Phi_{\min}$) quantifies how relatively decomposable a system is, whereas the ratio $\Phi_{\text{approx}}/\Phi \geq 1$ quantifies how well our approximation works, with a value of unity signifying that it is perfect and found the optimal bipartition. Fig 5 plots these two quantities against each other, and reveals that they are strongly related. For fairly separable systems, the approximation tends to be excellent: it gives exactly the correct answer 95% of the time when $\Phi_{\max}/\Phi > 2$ and 99.96% of the time when $\Phi_{\max}/\Phi > 3$. When $\Phi_{\max}/\Phi \lesssim 2$, on the other hand, so that there is less of a clear winner among the bipartitions, our approximation is seen to overestimate the true Φ -value by up to 15% on average (this is the median).

An alternative implementation, which we find works even better for some examples, is to apply the above-mentioned graph-based bipartition-finding scheme not to the evolution matrix \mathbf{A} but to the covariance matrix \mathbf{C} . We therefore recommend computing two approximate bipartitions, one based on \mathbf{A} and one based on \mathbf{C} , and selecting the one producing the smaller ϕ -value.

Discussion

Motivated by the growing interest in measuring integrated information Φ in computational and cognitive systems, we have presented a simple taxonomy of Φ -measures where they are each characterized by their choice of factorization method (5 options), choice of probability distributions to compare (3×4 options) and choice of measure for comparing probability distributions (5 options). We classify all the integration measures revealed in this taxonomy by various desirable properties, as summarized in [Table 1](#). When requiring the Φ -measures to satisfy a minimum of attractive properties, the hundreds of options reduce to a mere handful, some of which turn out to be identical. All leading contenders are summarized in [Table 2](#).

Unfortunately, these most general integration measures are unfeasible to evaluate in practice, with the computational cost growing doubly exponentially with b , the number of bits in the system: they involve a Markov matrix of size $n = 2^b$, and they also involve minimizing over approximately $N = 2^n = 2^{2^b}$ bipartitions. Generalizing the pioneering work of [13], we derive formulas for the Gaussian case that are exponentially faster, involving manipulations of a matrix whose size grows as $2b$ rather than 2^b with the number of variables b . Moreover, we show how the second exponential can also be avoided using an approximation using graph theory, thus reducing the computational cost from doubly exponential to merely polynomial in the system size b .

Which Φ -measures are best?

As described in detail in Results, six Φ -measures stand out from the taxonomy of hundreds of measures as particularly attractive: Φ^M , $\Phi_{kk'}^M$, $\Phi^{3.0}$, $\Phi^{2.5}$, $\Phi^{2.5'}$ and $\Phi^{2.5''}$. Φ^M retains all the attractive features of the Barrett/Seth measure Φ^B and adds further improvements: it is guaranteed to vanish for separable systems and to never be negative. If state-dependence is viewed as desirable, then its cousin $\Phi_{kk'}^M$ adds that feature too.

$\Phi^{3.0}$ is the measure advocated by IIT3.0 and has the many attractive features described in [11]. It has the drawback of being the slowest of all the measures to evaluate numerically: its definition involves a linear programming problem which needs to be solved numerically, and even with the fastest algorithms currently available, the computation for a given bipartition grows faster than quadratically with the number of system states—which in turn grows exponentially with the number of bits, and is infinite for continuous variables.

The remaining three top measures, $\Phi^{2.5}$, $\Phi^{2.5'}$ and $\Phi^{2.5''}$, share with $\Phi^{3.0}$ the arguably desirable feature of vanishing for afferent and efferent systems, but are much quicker to compute. $\Phi^{2.5}$ combines core ideas from IIT3.0 with the computational speed of IIT2.0 [10, 11] and elegantly depends only on the system's dynamics and present state, not on any assumptions about which states are more probable. Its drawback of being infinite for continuous variables is overcome by its cousin $\Phi^{2.5'}$.

A potential philosophical objection to both $\Phi^{2.5}$ and $\Phi^{2.5'}$ is that they are arguably not measures of integration, but measures of how suboptimal the factorizations “ n ” and “ m ” are, since they would both vanish if an optimal factorization were used—the measure $\Phi^{2.5''}$ eliminates this concern.

Outlook

Although the results in this paper will hopefully prove useful, there is ample worthwhile work left to do on integration measures.

One major open question is how to best handle asymmetric partitions. We deliberately sidestepped this challenge in the present paper, since it is independent of our results, which is why the subtle normalization issue raised by [10, 11, 13, 22] never entered. The crux is that if we apply any of the measures in our taxonomy with an asymmetric bipartition, the resulting ϕ -value will tend to get small when any of the two subsystems is very small, so simply defining Φ as the minimum of ϕ over all bipartitions (symmetric or not) makes no sense. IIT3.0 makes an interesting proposal [11] for how to handle asymmetric partitions, and it is worthwhile exploring whether there are other attractive options as well.

Another foundational question is whether our taxonomy can be placed on a firmer logical footing. Although our classification based on factorization, comparison, conditioning and measure may seem sensible and exhaustive, it is interesting to consider whether one or several Φ -measures can be rigorously derived from a small set of attractive axioms alone, in the same spirit as Claude Shannon derived his famous entropy formula, eq (30).

Yet another foundational question is whether integration maximization can be placed on a firmer physical footing, as advocated by [33, 34] in the context of continuous physical fields and by [12] in the context of quantum systems. The formulas in our taxonomy take information, measured in bits, as a starting point. But when I view a brain or computer through my physicist eyes, as myriad moving particles, then what physical properties of the system should be interpreted as logical bits of information? I interpret as a “bit” both the position of certain electrons in my computer’s RAM memory (determining whether the micro-capacitor is charged) and the position of certain sodium ions in your brain (determining whether a neuron is firing), but on the basis of what principle? Surely there should be some way of identifying consciousness from the particle motions alone, or from the quantum state evolution, even without this information interpretation? If so, what aspects of the behavior of particles corresponds to conscious integrated information? In other words, how can we generalize the quest for neural correlates of consciousness to *physical correlates of consciousness*? IIT argues that the consciousness occurs at precisely the level of course-graining in space and time that maximizes Φ [10], which is a prediction that should be tested.

A more practical question involves exploring ways of generalizing and further improving our graph-theory-based approximation for exponential speedup. One obvious generalization would involve taking advantage of the structure of Σ (which our method ignored) and the effect of \mathbf{x} (for those Φ -measures that are state-dependent). Another interesting opportunity is to generalize from continuous Gaussian systems to arbitrary discrete systems. For example, if the system consists of b different bits coupled by a nonlinear network of gates, one can apply a similar graph-theory approach by defining a $b \times b$ coupling matrix A_{ij} that in some way quantifies how strongly flipping the j^{th} bit would affect the i^{th} bit at the next timestep. As an example, consider defining A_{ij} as the probability that flipping the j^{th} bit will flip the i^{th} bit at the next timestep. If we have six bits evolving according to

$$\mathbf{x}_1 = \begin{pmatrix} a_1 \\ b_1 \\ c_1 \\ d_1 \\ e_1 \\ f_1 \end{pmatrix} = f(\mathbf{x}_0) = \begin{pmatrix} a_0 \\ \text{NOT } a_0 \\ \text{RANDOM} \\ c_0 \text{ XOR } d_0 \\ c_0 \text{ AND } d_0 \\ c_0 \text{ AND } d_0 \text{ AND } e_0 \end{pmatrix},$$

then the coupling matrix is

$$\mathbf{A} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & p_d & p_c & 0 & 0 \\ 0 & 0 & p_{de} & p_{ce} & p_{cd} & 0 \end{pmatrix},$$

where p_c denotes the probability that $c_0 = 1$, p_{de} denotes the probability that $d_0 = 1$ and $e_0 = 1$, etc. This coupling matrix is block-diagonal, showing that the bits a , b are completely independent of the others. For a state-independent Φ -measure, these probabilities can be computed as time-averages, otherwise they are each zero or one depending on the state. In either case, some elements of the \mathbf{A} -matrix can be small but non-zero (making the graph-theory approximation useful) if the system involves noisy gates or other randomness.

As regards practical challenges, it is important to note that there are many other issues besides speed that deserve further work because they have hindered the practical computation of integration Φ -measures from real brain data, including non-stationarity, statistical issues with estimating large numbers of parameters from short data windows without overfitting, possibilities of statistical bias, numerical instabilities, etc.

Last but not least, a veritable goldmine of data is becoming available in neuroscience and other fields, and it will be fascinating to measure Φ for these emerging data sets. In particular, the exponentially faster Φ -measures we have proposed will hopefully facilitate quantitative tests of theories of consciousness.

Acknowledgments

The author would like to thank Meia Chita-Tegmark, Henry Lin, Adam Barrett, Christof Koch, Masafumi Oizumi and Giulio Tononi for stimulating conversations, useful suggestions and proofreading help and Dan Fitch for catching typographical errors.

Author Contributions

Analyzed the data: MT.

Wrote the paper: MT.

References

1. Rees G, Kreiman G, Koch C. Neural correlates of consciousness in humans. *Nature Reviews Neuroscience*. 2002; 3(4):261–270. doi: [10.1038/nrn783](https://doi.org/10.1038/nrn783) PMID: [11967556](https://pubmed.ncbi.nlm.nih.gov/11967556/)
2. Metzinger T. *Neural correlates of consciousness: Empirical and conceptual questions*. MIT press; 2000. doi: [10.1006/ccog.1998.0335](https://doi.org/10.1006/ccog.1998.0335) PMID: [9521836](https://pubmed.ncbi.nlm.nih.gov/9521836/)
3. Chalmers DJ. Facing up to the problem of consciousness. *Journal of consciousness studies*. 1995; 2(3):200–219. doi: [10.1093/acprof:oso/9780195311105.003.0001](https://doi.org/10.1093/acprof:oso/9780195311105.003.0001)
4. Dehaene S, Charles L, King JR, Marti S. Toward a computational theory of conscious processing. *Current opinion in neurobiology*. 2014; 25:76–84. doi: [10.1016/j.conb.2013.12.005](https://doi.org/10.1016/j.conb.2013.12.005) PMID: [24709604](https://pubmed.ncbi.nlm.nih.gov/24709604/)
5. Dehaene S. Conscious and nonconscious processes: distinct forms of evidence accumulation? In: *Biological Physics*. Springer; 2011. p. 141–168. doi: [10.7551/mitpress/9780262195805.003.0002](https://doi.org/10.7551/mitpress/9780262195805.003.0002)
6. Shadlen MN, Kiani R. Consciousness as a decision to engage. In: *Characterizing consciousness: from cognition to the clinic?* Springer; 2011. p. 27–46. doi: [10.1007/978-3-642-18015-6_2](https://doi.org/10.1007/978-3-642-18015-6_2)

7. Dehaene S, Naccache L. Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*. 2001; 79(1):1–37. doi: [10.1016/S0010-0277\(00\)00123-2](https://doi.org/10.1016/S0010-0277(00)00123-2) PMID: [11164022](https://pubmed.ncbi.nlm.nih.gov/11164022/)
8. Shanahan M, Baars B. Applying global workspace theory to the frame problem. *Cognition*. 2005; 98(2):157–176. doi: [10.1016/j.cognition.2004.11.007](https://doi.org/10.1016/j.cognition.2004.11.007) PMID: [16307957](https://pubmed.ncbi.nlm.nih.gov/16307957/)
9. Dehaene S, Kerszberg M, Changeux JP. A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences*. 1998; 95(24):14529–14534. doi: [10.1073/pnas.95.24.14529](https://doi.org/10.1073/pnas.95.24.14529) PMID: [9826734](https://pubmed.ncbi.nlm.nih.gov/9826734/)
10. Tononi G. Consciousness as integrated information: a provisional manifesto. *The Biological Bulletin*. 2008; 215(3):216–242. doi: [10.2307/25470707](https://doi.org/10.2307/25470707) PMID: [19098144](https://pubmed.ncbi.nlm.nih.gov/19098144/)
11. Oizumi M, Albantakis L, Tononi G. From the Phenomenology to the Mechanisms of Consciousness: Integrated Information Theory 3.0. *PLoS computational biology*. 2014; 10(5):e1003588. doi: [10.1371/journal.pcbi.1003588](https://doi.org/10.1371/journal.pcbi.1003588) PMID: [24811198](https://pubmed.ncbi.nlm.nih.gov/24811198/)
12. Tegmark, M. Consciousness as a State of Matter. arXiv preprint arXiv:14011219. 2014;. <http://www.sciencedirect.com/science/article/pii/S0960077915000958>.
13. Barrett AB, Seth AK. Practical measures of integrated information for time-series data. *PLoS computational biology*. 2011; 7(1):e1001052. doi: [10.1371/journal.pcbi.1001052](https://doi.org/10.1371/journal.pcbi.1001052) PMID: [21283779](https://pubmed.ncbi.nlm.nih.gov/21283779/)
14. Seth AK, Barrett AB, Barnett L. Causal density and integrated information as measures of conscious level. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*. 2011; 369(1952):3748–3767. doi: [10.1098/rsta.2011.0079](https://doi.org/10.1098/rsta.2011.0079) PMID: [21893526](https://pubmed.ncbi.nlm.nih.gov/21893526/)
15. Oizumi M, Amari Si, Yanagawa T, Fujii N, Tsuchiya N. Measuring integrated information from the decoding perspective. *PLoS Comput Biol*. 2016; 12(1):e1004654. doi: [10.1371/journal.pcbi.1004654](https://doi.org/10.1371/journal.pcbi.1004654) PMID: [26796119](https://pubmed.ncbi.nlm.nih.gov/26796119/)
16. Casali AG, Gosseries O, Rosanova M, Boly M, Sarasso S, Casali KR, et al. A theoretically based index of consciousness independent of sensory processing and behavior. *Science translational medicine*. 2013; 5(198):198ra105–198ra105. doi: [10.1126/scitranslmed.3006294](https://doi.org/10.1126/scitranslmed.3006294) PMID: [23946194](https://pubmed.ncbi.nlm.nih.gov/23946194/)
17. Sitt JD, King JR, El Karoui I, Rohaut B, Faugeras F, Gramfort A, et al. Large scale screening of neural signatures of consciousness in patients in a vegetative or minimally conscious state. *Brain*. 2014; 137(8):2258–2270. doi: [10.1093/brain/awu141](https://doi.org/10.1093/brain/awu141) PMID: [24919971](https://pubmed.ncbi.nlm.nih.gov/24919971/)
18. Griffith V, Koch C. Quantifying synergistic mutual information. In: *Guided Self-Organization: Inception*. Springer; 2014. p. 159–190. doi: [10.1007/978-3-642-53734-9_6](https://doi.org/10.1007/978-3-642-53734-9_6)
19. Edlund JA, Chaumont N, Hintze A, Koch C, Tononi G, Adami C. Integrated information increases with fitness in the evolution of animats. *PLoS Comput Biol*. 2011; 7(10):e1002236. doi: [10.1371/journal.pcbi.1002236](https://doi.org/10.1371/journal.pcbi.1002236) PMID: [22028639](https://pubmed.ncbi.nlm.nih.gov/22028639/)
20. Engel D, Malone TW. Can groups be conscious? Integrated information as a metric for group interaction. preprint. 2015;.
21. Chicharro D, Ledberg A. When two become one: the limits of causality analysis of brain dynamics. *PLoS One*. 2012; 7(3):e32466. doi: [10.1371/journal.pone.0032466](https://doi.org/10.1371/journal.pone.0032466) PMID: [22438878](https://pubmed.ncbi.nlm.nih.gov/22438878/)
22. Balduzzi D, Tononi G. Integrated information in discrete dynamical systems: motivation and theoretical framework. *PLoS Comput Biol*. 2008; 4(6):e1000091. doi: [10.1371/journal.pcbi.1000091](https://doi.org/10.1371/journal.pcbi.1000091) PMID: [18551165](https://pubmed.ncbi.nlm.nih.gov/18551165/)
23. Kullback S, Leibler RA. On information and sufficiency. *The annals of mathematical statistics*. 1951; p. 79–86. doi: [10.1214/aoms/1177729694](https://doi.org/10.1214/aoms/1177729694)
24. Amari Si. *Information Geometry and Its Applications*. vol. 194. Springer; 2016. doi: [10.1007/978-4-431-55978-8](https://doi.org/10.1007/978-4-431-55978-8)
25. Griffith, V. A Principled Infotheoretic\ phi-like Measure. arXiv preprint arXiv:14010978. 2014;.
26. Endres DM, Schindelin JE. A new metric for probability distributions. *IEEE Transactions on Information theory*. 2003;. doi: [10.1109/TIT.2003.813506](https://doi.org/10.1109/TIT.2003.813506)
27. Rubner Y, Tomasi C, Guibas LJ. A metric for distributions with applications to image databases. In: *Computer Vision, 1998. Sixth International Conference on. IEEE*; 1998. p. 59–66.
28. Merhav N, Kaplan G, Lapidot A, Shitz SS. On information rates for mismatched decoders. *Information Theory, IEEE Transactions on*. 1994; 40(6):1953–1967. doi: [10.1109/18.340469](https://doi.org/10.1109/18.340469)
29. Shannon CE. A mathematical theory of communication. *Bell System Technical Journal*. 1948; 27:379–423. doi: [10.1145/584091.584093](https://doi.org/10.1145/584091.584093)
30. Ay N, et al. Information geometry on complexity and stochastic interaction. 2001;.
31. Ay N. Information geometry on complexity and stochastic interaction. *Entropy*. 2015; 17(4):2432–2458.
32. M. Oizumi, N. Tsuchiya, and S.-i. Amari, *arXiv preprint arXiv:1510.04455* (2015).

33. Barrett AB. An integration of integrated information theory with fundamental physics. *Frontiers in psychology*. 2014; 5. doi: [10.3389/fpsyg.2014.00063](https://doi.org/10.3389/fpsyg.2014.00063) PMID: [24550877](https://pubmed.ncbi.nlm.nih.gov/24550877/)
34. Barrett AB. A comment on Tononi & Koch (2015) "Consciousness: here, there and everywhere?". *Phil Trans R Soc B*. 2016; 371(1687):20140198. doi: [10.1098/rstb.2014.0198](https://doi.org/10.1098/rstb.2014.0198) PMID: [26729923](https://pubmed.ncbi.nlm.nih.gov/26729923/)