A Compositional Approach to Higher-Order Structure in Complex Systems: Carving Nature at its Joints

Abel Jansma^{*}

Max Planck Institute for Mathematics in the Sciences, Germany and School of Informatics, University of Edinburgh, United Kingdom (Dated: April 24, 2024)

Relating microscopic interactions to macroscopic observables is a central challenge in the study of complex systems. Addressing this question requires understanding both pairwise and higher-order interactions, but the latter are less well understood. Here, we show that the Möbius inversion theorem provides a general mathematical formalism for deriving higher-order interactions from macroscopic observables, relative to a chosen decomposition of the system into parts. Applying this framework to a diverse range of systems, we demonstrate that many existing notions of higherorder interactions, from epistasis in genetics and many-body couplings in physics, to synergy in game theory and artificial intelligence, naturally arise from an appropriate mereological decomposition. By revealing the common mathematical structure underlying seemingly disparate phenomena, our work highlights the fundamental role of decomposition choice in the definition and estimation of higher-order interactions. We discuss how this unifying perspective can facilitate the transfer of insights between domains, guide the selection of appropriate system decompositions, and motivate the search for novel interaction types through creative decomposition strategies. More broadly, our results suggest that the Möbius inversion theorem provides a powerful lens for understanding the emergence of complex behaviour from the interplay of microscopic parts, with applications across a wide range of disciplines.

^{*} abel.jansma@mis.mpg.de

I.	Introduction A. Higher-order interactions B. Aim C. Related work	2 2 3 3
II.	Möbius Inversion as a Mereological Framework A. Definitions and notation B. Decomposing Systems from macro to micro	4
	C. Inverting decompositions with Möbius functions	5
III.	Möbius inversions in Complex Systems	7
	A. Statistics	7
	B. Information theory	8
	1. Entropy and Mutual Information	8
	2. The Partial Information Decomposition	9
	C. Biology	10
	1. Epistasis	10
	2. Epidemiology	11
	3. Neuroscience	12
	D. Physics	12
	1. Equilibrium dynamics	12
	2. Statistical mechanics	13
	3. Quantum & Statistical Field Theory	14
	E. Chemistry	15
	F. Game Theory	15
	G. Artificial Intelligence	16
IV.	Discussion	17
	Acknowledgements	18
	References	18

I. INTRODUCTION

Much of the study of complex systems is focused on characterizing general principles of complexity. Famously, certain network or graph theoretical quantities have proven useful to describe the structure of complex systems in fields ranging from physics to biology and the social sciences [12, 25, 31, 55, 62, 66, 83]. In addition, concepts from a specific discipline, like phase transitions in physics or evolution in biology, have led to both qualitative and quantitative understanding of complex systems in other fields [19, 32, 70, 74, 91]. This paper makes another such attempt at revealing a general principle of complex systems, and aims to show that there is a unified theory of higher-order interactions that is intimately related to the decomposition of a system into its parts. We show that various notions of higher-order interactions across scientific fields are manifestations of a single principle—the Möbius inversion theorem—that describes how higher-order interactions emerge from decompositions of complex systems. In doing so, the framework presented here provides both a unified theory, as well as a precise definition of the term *higher-order* structure, which is often used in a somewhat vague way to refer to the non-additive interactions between parts of a system. In particular, we aim to show that when a system is decomposed into parts that allow for a partial ordering, then the notion of higher-order structure is uniquely determined, and certain quantities are higher-order with respect to this precisely this partial order. We give a range of examples that show that this approach reproduces famous examples of higher-order quantities in various scientific disciplines.

A. Higher-order interactions

Throughout the sciences, it is often said that a set of variables interact if their joint configuration affects an outcome differently than the sum of their individual effects. In this sense, interactions are a kind of nonadditivity, and they form the basis for all properties of complex systems that are not simply the sum of the properties of their parts. It is hard to imagine a scientific question that does not reduce to the problem of understanding interactions and their effects. Still, in the majority of cases, interactions are only studied among pairs of variables. For this reason, interactions that involve more than two variables are often collectively referred to as *higher-order* interactions. The reason higher-order interactions have been largely ignored is manifold: they require more data to measure, they are harder to interpret, and pairwise models have been surprisingly successful. In addition, pairwise interactions can be concisely represented as a graph, allowing them to be analysed with the powerful tools of graph theory. Higher-order interactions, in contrast, can only be represented as a *hypergraph*, a generalisation of a graph that allows for edges to connect more than two nodes. Hypergraphs are not as well-understood as graphs, and the tools to analyse them less developed, though significant progress is being made [23, 28, 49, 52].

While it has been shown that there are certain situations in which pairwise models are generally good approximations [57, 80], higher-order interactions have proven to be of crucial importance to the rich dynamics and multistability of complex systems [6, 69, 73, 79]. To motivate the attempt at a unified theory of higher-order interactions made in this paper, we first highlight the importance of higher-order interactions to particular problems in biology and physics.

In biology, higher-order interactions among genetic variants and mutations have been shown to play a key role in the emergence of phenotype from genotype [22, 47]. Similarly, at the level of transcription, non-additive higher-order effects are responsible for the complexity of BMP signalling [2, 45], embryonic development of *Drosophila* [3] and the emergence of cell type from gene expression [38, 40]. On an ecological scale, certain species of lichen crucially depend on a symbiosis that involves more than two species [77], and higher-order interactions among species in the *Drosophila* gut microbiome affect the longevity of the host [29]. Also in the brain, which is generally thought of as a network of pairwise connected neurons, higher-order and synergistic functional interactions among brain regions are associated with more complex and integrative cognitive processes than additive ones [53]. Each of these examples will be shown to be expressible in the framework of Möbius inversions in Section III C.

However, higher-order interactions are perhaps most ubiquitous in physics, which has a long history of introducing non-pairwise terms into models. For example, while the Ising model on a lattice is commonly defined with only pairwise nearest-neighbour interactions, a coarse-graining or renormalisation of the lattice necessarily introduces higher-order interactions among the spins [54]. Furthermore, spin models with varying and arbitraryorder interactions have been extensively studied in their own right, and are generally referred to as spin glasses due to their importance to the study of glassy materials [27, 41, 43]. In quantum field theory, scattering amplitudes are generally approximated by perturbative methods that sum increasing orders of particle interactions [63, 84]. Both these examples will be shown to be expressible in the framework of Möbius inversions in Section III D.

B. Aim

The goal of this study is to show how the higher-order structure of a complex system uniquely emerges from a decomposition of the system into parts. The goal is not to introduce new notions of higher-order interactions. Rather, we offer a unified perspective through the lens of Möbius inversions. While the Möbius inversion theorem has been previously applied in some of the individual examples covered here, viewing them together through this unifying lens provides several key insights. First, it reveals a common mathematical structure underlying seemingly disparate notions of higher-order interactions across scientific domains. Second, it provides a general framework for deriving microscopic interactions from macroscopic observables, relative to a chosen decomposition of the system into parts. We hope that this perspective inspires novel system decompositions that yield new notions of higher-order interactions, and allows insights in one scientific discipline to be more straightforwardly transferred to others. Finally, it reveals the choice of decomposition itself to be a key factor in determining the nature of the higher-order interactions. It emphasis that the chosen decomposition should be natural and useful for the given system and property, echoing Plato's claim that the best way to understand Nature is to carve it at its joints [64].

C. Related work

That Möbius inversions can be useful in the study of complex systems is itself not a novel observation. For example, Section III will show that many instances of its use are based on the relationship between moments and cumulants of a probability distribution, which were already realised as Möbius inversions on the appropriate lattices in [68, 75]. Furthermore, within information theory, deriving general principles of complexity based on system decompositions has been explored before in [4], and the authors of [37, 48] use Möbius inversion to connect different concepts from information theory. However, each of these examples only considers decompositions with the structure of a Boolean algebra, and as such essentially reduces to the inclusion-exclusion principle. While solving the system of equations associated to the partial information decomposition (Section III B) has historically been referred to as a Möbius inversion on a lattice of antichains, the corresponding Möbius function has only recently been identified [39]. The framework presented here holds for arbitrary partial orders, and it will be demonstrated that decompositions into lattices different from Boolean algebras are ubiquitous.

II. MÖBIUS INVERSION AS A MEREOLOGICAL FRAMEWORK

In this section, we introduce various ways to decompose a system and relate the parts to the whole. Bringing together parts to form a whole is the original meaning of algebra (*al-jabr* being Arabic for reunion of parts), and we indeed find that system decomposition can be fruitfully described with an algebraic technique known as a Möbius inversion. But before we present the central framework, we first define the terms and notations used in this paper.

A. Definitions and notation

Definition 1. Let P be a set. A partial order on P is a binary relation \leq with the following properties for all $a, b, c \in P$:

Reflexivity:
$$a \le a$$
 (1)

Transitivity:
$$a \le b$$
 and $b \le c \implies a \le c$ (2)

Antisymmetry:
$$a \le b$$
 and $b \le a \iff a = b$ (3)

Whenever $a \leq b$ we say that a is less than or equal to b. Two elements $a, b \in P$ are comparable when either $a \leq b$ or $b \leq a$, and incomparable otherwise. When $a \neq b$ and $a \leq b$, then we write b > a, and say that a is greater than b.

The tuple (P, \leq) is referred to as a partially ordered set, or poset, but we often just write P when the ordering is clear from context. Given a subset $S \subseteq P$, an element $b \in P$ is a lower bound for S if $\forall s \in S : b \leq s$ (upper bounds are defined similarly). An interval [a, b] on P is a set $\{x : a \leq x \leq b\}$, and a poset is called locally finite if all such intervals are finite.

One can impose extra structure on the poset to form special cases. For example, a poset in which every two elements are comparable is called a totally ordered set. An obvious example of this is the poset (\mathbb{N}, \leq) of the natural numbers with their usual ordering. A poset P in which every two elements $a, b \in P$ have a unique greatest lower bound $a \wedge b \in P$ and least upper bound $a \vee b$ is called a *lattice*.

Given a poset (P, \leq) , one can study how functions on the underlying set P interact with the ordering. For example, a function $f: P \to \mathbb{R}$ is called *monotone* if $a \leq b$ implies $f(a) \leq f(b)$. However, there is a particular rich theory of functions on *intervals* on locally finite posets. Functions on intervals exploit the full structure of the poset, and form an algebraic structure known as the incidence algebra, where the algebra's multiplication operation * on two functions $f, g: P \times P \to \mathbb{R}$ is defined as the convolution of their values over the interval:

$$(f*g)(a,b) = \sum_{a \le x \le b} f(a,x)g(x,b) \tag{4}$$

Of particular interest are three elements of the incidence algebra: the delta function δ_P , zeta function ζ_P and the Möbius function μ .

Definition 2 (Delta function). Let (P, \leq) be a locally finite poset. Then the delta function $\delta_P : P \times P \to \mathbb{R}$ is defined as

$$\delta_P(x,y) = \begin{cases} 1 & \text{if } x = y \\ 0 & \text{otherwise} \end{cases}$$
(5)

Definition 3 (Zeta function). Let (P, \leq) be a locally finite poset. Then the zeta function $\zeta_P : P \times P \to \mathbb{R}$ is defined as

$$\zeta_P(x,y) = \begin{cases} 1 & \text{if } x \le y \\ 0 & \text{otherwise} \end{cases}$$
(6)

Definition 4 (Möbius function). Let (P, \leq) be a locally finite poset. Then the Möbius function $\mu_P : P \times P \to \mathbb{R}$ is defined as

$$\mu_P(x,y) = \begin{cases} 1 & \text{if } x = y \\ -\sum_{z:x \le z < y} \mu_P(x,z) & \text{if } x < y \\ 0 & \text{otherwise} \end{cases}$$
(7)

When the underlying poset is clear from context, or irrelevant, we sometimes omit the subscript. Note that δ corresponds to the multiplicative unit ($\delta * f = f * \delta = f$ for all f in the incidence algebra), and that multiplying a function $f : P \to \mathbb{R}$ with the zeta function amounts an integral over the poset: $(f * \zeta)(x, y) = \sum_{x \leq z \leq y} f(z)$ (to properly define the multiplication, interpret f as a function on $P \times P$ that is constant in its first argument). One of the most important results in the theory of incidence algebra then states that the Möbius function is the multiplicative inverse of the zeta function ($\mu * \zeta = \zeta * \mu = \delta$) and that therefore the following theorem holds:

Theorem 1 (Möbius inversion theorem, Rota [67]). Let $P = (S, \leq)$ be a finite poset and $\tau, \eta \in S$. Let $f : P \to \mathbb{R}$ be a function on P, and let μ_P be the Möbius function on P. Then the following two statements are equivalent:

$$f(\tau) = \sum_{\eta \le \tau} g(\eta) \tag{8}$$

$$g(\tau) = \sum_{\eta \le \tau} \mu_P(\eta, \tau) f(\eta) \tag{9}$$

The Möbius inversion theorem states that decompositions over a poset can be inverted by looking up the Möbius function of the poset. This powerful result forms the basis of this study. In the following sections, we will show how this theorem can be applied to a wide range of systems, and how it can be used to define and estimate many well-established notions of higher-order structure in complex systems.

B. Decomposing Systems from macro to micro

Given a system S with parts s_i , consider an arbitrary property Q(S) of the system. A purely additive property would be one where

$$Q(S) = \sum_{i} q(s_i) \tag{10}$$

where $q(s_i)$ is the contribution of the part s_i to the property Q(S). For example, if a person's height is determined purely additively in terms of the presence of genetic variants $g_i \in G$, then one could write $H(G) = h(\emptyset) + \sum_i h(g_i)$, where $h(\emptyset)$ is the height of a person without any of the variants, and $h(g_i)$ is the contribution of the genetic variant g_i to the person's height H(G). Note the difference in interpretation between H and hhere: H is the length of a person and easy to measure, whereas h is the length added by a genetic variant which is generally impossible to directly observe.

However, as discussed in the introduction, many interesting properties of complex systems are non-additive and emerge from the potentially complex interactions among the components. Still, it's reasonable to assume that Q could be decomposed into properties of the components in a nonadditive or interacting way, as anything else would entail Q being *strongly* emergent (i.e. *fundamentally* not derivable from a description of its parts). One option to allow for nonadditive effects is to introduce nonadditive terms (e.g. multiplicative) by hand. Another way is to still only include additive contributions, but from elements of a more complex decomposition of the system. Note that this is fully analogous to introducing more complex terms in the design matrix of a linear regression problem. This invites one to instead write

$$Q(S) = \sum_{t \in \mathcal{D}(S)} q(t) \tag{11}$$

where $\mathcal{D}(S)$ is some kind of decomposition of S that relates the whole to its parts. An obvious choice for $\mathcal{D}(S)$ would be the powerset of S, but there are other choices. For example, while the powerset is essentially the set of bipartitions, one could assume Q to be determined by the set of all partitions of S. Both powersets and partitions are very simple and structured decompositions, and can be partially ordered by inclusion and refinement, respectively, as shown in Figures 1 and 2. Choosing a particular $\mathcal{D}(S)$ amounts to a mereological claim about the system S and property Q, so should be informed by prior knowledge of the system or the method of observation. In the example of a person's height, it is natural to decompose the set of genetic variants into subsets of variants, based on the assumption that the only property of a variant that affects height is its presence or absence. If we for simplicity assume that G contains just two variants, then a person's height $H(S) = h(\emptyset) + h(s_1) + h(s_2) + h(s_1, s_2)$, indicating that the height is determined by the individual variants, but also by an interaction term $h(s_1, s_2)$ among them.

C. Inverting decompositions with Möbius functions

If the decomposition $\mathcal{D}(S)$ is indeed a partial order with a largest element S, then we can write

$$Q(S) = \sum_{t \le S} q(t) \tag{12}$$



FIG. 1. The powerset of a set of n variables, ordered by set inclusion, forms a lattice known as a Boolean algebra. Shown here are the transitive reductions (Hasse diagrams) of the Boolean algebras on 2 (left), 3 (middle) and 4 (right) variables. For arbitrary n, the Hasse diagrams of a Boolean algebra describes an n-cube.



FIG. 2. The partitions of a set of n variables, ordered by refinement, form a lattice. Shown here are the transitive reductions (Hasse diagrams) of the partition lattices on 2 (left), 3 (middle) and 4 (right) variables.

where the sum is now over elements from the partial order. If the microscopic contributions q(t) are known, then predicting the macroscopic quantity Q(S) is known as the *forward problem*. However, in many cases we might only be able to observe Q(S), and never the individual contributions q(t). To reverse-engineer q(t) from Q(S), Equation (12) should be inverted to express q(t) in terms of observations of Q on different parts of the decomposition. To invert sums like equation (12) over posets, one can use the Möbius inversion theorem:

$$q(\tau) = \sum_{\eta \le \tau} \mu_P(\eta, \tau) Q(\eta) \tag{13}$$

This is a powerful result: the problem of inverting the decomposition (which amounts to solving a large system of equations) has been reduced to looking up the Möbius function of the partial order implied by the decomposition.

The result can be represented diagrammatically as follows. Let Σ be the set of full systems, and Δ the set of decompositions. Let $\mathcal{F}_{\mathcal{D}} = \{f | f : \mathcal{D}(S) \to \mathbb{R}\}$ be the set of functions on a decomposition $\mathcal{D}(S)$. Given a decomposition map \mathcal{D} , consider a map $\mathcal{O}_q : \Delta \to \mathcal{F}_{\mathcal{D}}$ that assigns a definition q of microscopic interactions to a decomposition, and a map $\mathcal{O}_Q : \Delta \to \mathcal{F}_{\mathcal{D}}$ that assigns a definition Q of macroscopic observables. The two definitions \mathcal{O}_q and \mathcal{O}_Q are called \mathcal{D} -compatible if the following diagram commutes:



If the definitions of Q and q are \mathcal{D} -compatible, then the forward problem is solved by multiplying q by $\zeta_{\mathcal{D}}$, and the inverse problem is solved by multiplying Q by $\mu_{\mathcal{D}}$. In fact, the Möbius inversion theorem states that for any macroscopic observable Q there is a *unique* notion of microscopic higher-order interaction q that is \mathcal{D} -compatible with Q, and *vice versa*. Since many decompositions use the same partial order, the Möbius function can be precomputed, and the inversion can be done in a single step. In fact, by far the most common decompositions of a system into parts are the powerset and partition lattices, ordered by inclusion and refinement respectively, and the Möbius function of these lattices is well-known. Denoting the powerset lattice by $\mathcal{P}(S)$, and the partition lattice by $\Pi(S)$, the Möbius functions are given by

$$\mu_{\mathcal{P}(S)}(x,y) = \begin{cases} (-1)^{|y|-|x|} & \text{if } x \le y \\ 0 & \text{otherwise} \end{cases}$$
(14)

$$\mu_{\Pi(S)}(x,y) = \begin{cases} (-1)^{|x|-1}(|x|-1)! & \text{if } x \le y \\ 0 & \text{otherwise} \end{cases}$$
(15)

where |x| is the cardinality of the set x. These functions are well-known and have been used to invert decompositions across scientific disciplines, as will be discussed in more detail in Section III.

In the example of genetic variants determining a person's height, inserting Equation (14) into (13) allows us to calculate the effect of individual genetic variants, as well as the interaction between them, from observations of a population's height. To see this, order the powerset of genetic variants $\mathcal{P}(S) = \{\emptyset, \{g_1\}, \{g_2\}, \{g_1, g_2\}\}$ by inclusion. Then Eq. (13) yields

$$h(g_1) = Q(g_1) - Q(\emptyset) \tag{16}$$

$$h(g_2) = Q(g_2) - Q(\emptyset) \tag{17}$$

$$h(g_1, g_2) = Q(g_1, g_2) - Q(g_1) - Q(g_2) + Q(\emptyset)$$
(18)

$$= (Q(g_1, g_2) - Q(g_2)) - (Q(g_1) - Q(\emptyset))$$
(19)

This is easily interpreted: the effect of a single genetic variant is the difference between the height of a person with only that variant and the height of a person without any of the variants. The effect of the interaction among two variants g_1 and g_2 is the difference between the effect of g_1 in people with g_2 , and the effect of g_1 in people without g_2 . Given a population sample containing people with all combinations of these variants, both of these quantities can be directly estimated. This argument straightforwardly extends to higher-order effects: the third-order interaction among three genetic variants g_1, g_2, g_3 is given by

$$h(g_1, g_2, g_3) = H(g_1, g_2, g_3) - H(g_1, g_2) - H(g_1, g_3) - H(g_2, g_3) + H(g_1) + H(g_2) + H(g_3) - H(\emptyset)$$
(20)

In genetics, interaction terms like $h(g_1, g_2)$ are called epistatic effects and commonly defined and estimated using exactly this estimator [9, 29] (see Section III C), though not generally linked to Möbius inversions. Section III will show in more detail how this construction can be applied to any system with a well-defined partial order of parts, and how it reproduces many well-established notions of higher-order structure in complex systems.

III. MÖBIUS INVERSIONS IN COMPLEX SYSTEMS

We aim to show that the framework presented in Section II reproduces many notions of higher-order interactions throughout the sciences. First, the role of Möbius inversions in statistics is discussed in Section III A. We next recover various quantities from information theory in Section III B. Within the natural sciences, multiple kinds of interactions in biology (Section III C), physics (Section III D), and chemistry (Section III E) are discussed. Finally, we discuss the role that Möbius inversions play in game theory and artificial intelligence in Sections III F and III G, respectively. An overview of all dualities between macroscopic and microscopic quantities is given in Table I of the Discussion. The wide range of examples covered here serve mainly to illustrate the broad applicability of the framework, but it is by no means necessary to understand the details of all of them to appreciate the general idea.

A. Statistics

Of central importance in statistics are the moments of a distribution. Given a joint distribution p over N variables $(X_1, \ldots, X_N) \in \mathcal{X}$, the mixed moment of a set of variables S is given by

$$\mathbb{E}(\prod_{i} S_{i}) = \sum_{x_{1},\dots,x_{N} \in \mathcal{X}} p(X_{1} = x_{1},\dots,X_{N} = x_{N}) \prod_{i} S_{i}$$
(21)

$$\coloneqq \langle S \rangle \tag{22}$$

Since the variables might be dimensionful quantities, any decomposition of this moment into a sum over more elemental contributions e(X) must comprise only terms with the same dimension. If we decompose τ into its powerset, it is therefore natural to write:

$$\langle S \rangle = \sum_{T \in \mathcal{P}(S)} e(T) \prod_{S_i \in S \setminus T} e(S_i)$$
(23)

where $\mathcal{P}(S)$ is the powerset of S. To find out what the elemental contributions $e(S_i)$ are, we can invert this sum using a Möbius inversion:

$$e(S) = \sum_{T \le S} \mu_{\mathcal{P}}(T, S) \langle T \rangle \prod_{S_i \in S \setminus T} \langle S_i \rangle$$
(24)

$$=\sum_{T\leq S} (-1)^{|T|-|S|} \langle T \rangle \prod_{S_i \in S \setminus T} \langle S_i \rangle$$
(25)

which for the case of three variables X_1, X_2, X_3 yields

$$e(X_i) = \langle X_i \rangle \tag{26}$$

$$e(X_i, X_j) = \langle X_i, X_j \rangle - \langle X_i \rangle \langle X_j \rangle$$

$$(27)$$

$$V = V = \langle Y = V \rangle \langle$$

$$e(X_1, X_2, X_3) = \langle X_1, X_2, X_3 \rangle - \langle X_1, X_2 \rangle \langle X_3 \rangle - \langle X_1, X_3 \rangle \langle X_2 \rangle - \langle X_2, X_3 \rangle \langle X_1 \rangle + 2 \langle X_1 \rangle \langle X_2 \rangle \langle X_3 \rangle$$
(28)

These are exactly the mixed central moments (it can be readily verified that this construction generalises to all higher-order moments), so that central moments are the Möbius inverse of mixed moments with respect to the powerset lattice.

What would happen if we chose a different decomposition of the moments? For example, one might decompose a moment into contributions from all possible partitions π of the system. Again, making sure that dimensions match, we write this as

$$\langle S \rangle = \sum_{\pi \in \Pi(S)} \prod_{\pi_i \in \pi} \kappa(\pi_i) \tag{29}$$

where now $\kappa(\pi_i)$ is the contribution of the partition π_i to the mixed moment, and $\Pi(S)$ is the set of all partitions of S. Inverting this sum over the lattice of partitions ordered by refinement yields

$$\kappa(S) = \sum_{\pi \in \Pi(S)} \mu_{\Pi}(\pi, S) \prod_{\pi_i \in \pi} \langle \pi_i \rangle$$
(30)

$$= \sum_{\pi \in \Pi(S)} (-1)^{|\pi|-1} (|\pi|-1)! \prod_{\pi_i \in \pi} \langle \pi_i \rangle$$
(31)

where we have written S for the partition $\{S, \emptyset\}$. This happens to be the same as the central moments for up to three variables, but is different from them afterwards. The κ are called the mixed cumulants, and famously offer an equivalent way to characterise the distribution p.

Note that a decomposing a set of variables S into partitions only makes sense if the variables commute. In the noncommutative case, it is more natural to decompose S into so-called *non-crossing* partitions. A non-crossing partition is a partition of S such that no two elements of the partition cross each other. For example, the partition $\{\{1,2\},\{3,4\}\}$ is non-crossing, but $\{\{1,3\},\{2,4\}\}$ is not. In general, a partition is non-crossing if there is no chain of elements A > B > C > D such that A and C are in the same block, B and D are in the same block, but A and D are in different blocks. The Möbius function of the lattice of non-crossing partitions is known to be given by signed Catalan numbers, and can be used to define the noncommutative version of cumulants, called *free cumulants* [76]. When sampling from stochastic processes, one might encounter path-valued variables. The statistics of such path-values samples can be summarised in a sequence of *path signature moments*. The authors of [11] argue that the sequential nature of these path-valued variables makes it most natural to decompose them into *ordered partitions*, which is a different subset of partitions that can again be ordered by refinement. Inverting the path signature moments over this lattice of ordered partitions yields the path signature cumulants, which are the natural generalisation of the cumulants to path-valued variables.

B. Information theory

1. Entropy and Mutual Information

Given a set X of random variables X_i , the joint entropy H(X) is the total amount of uncertainty about the state of X before an observation, or equivalently, the amount of information gained by observing X. It is defined as

$$H(X) = -\sum_{x \in \mathcal{X}} p(X = x) \log p(X = x)$$
(32)

where \mathcal{X} is the set of all possible states of X, and p(X = x) is the probability of observing state x. One might assume that the total information in the system can be decomposed into information contained in different parts of the system. To do this, we decompose the system into the powerset of X. The joint entropy can then be written as

$$H(X) = \sum_{S \in \mathcal{P}(X)} I(S) \tag{33}$$

where I(S) is the information contributed by the part S. These individual contributions to the joint entropy are given by a Möbius inversion on the powerset of X, ordered by inclusion:

$$I(X) = \sum_{S < X} \mu_{\mathcal{P}}(S, X) H(S) \tag{34}$$

$$=\sum_{S\leq X} (-1)^{|S|-|X|} H(S)$$
(35)

For the example of two variables X_1 and X_2 this yields

$$I(X_1) = H(X_1) \tag{36}$$

$$I(X_2) = H(X_2) \tag{37}$$

$$I(X_1, X_2) = H(X_1, X_2) - H(X_1) - H(X_2)$$
(38)

This is, up to a minus sign, exactly the definition of mutual information, which is the amount of information shared between two variables (or equivalently: the Kullback-Leibler divergence between the joint distribution and the product of the marginals):

$$I(X_1, X_2) = \sum_{x_1, x_2} p(x_1, x_2) \log \frac{p(x_1, x_2)}{p(x_1)p(x_2)}$$
(39)

In general, the mutual information among a set of variables, also referred to as their *interaction information*, is indeed given by the Möbius inversion of the entropy of the powerset of the variables, with some authors adding a minus sign to even orders to make sure that the entropies of single variables always get a + sign. This is a well-established result, and is based on the analogy between Shannon information theory and set theory [89].

In fact, the same argument holds for the pointwise information, or surprisal h(x), which is the amount of information gained by observing a particular realisation X = x:

$$h(x) = -\log p(X = x) \tag{40}$$

A Möbius inversion on $\mathcal{P}(X)$ then leads to the definition of pointwise mutual information:

$$i(X_1, X_2) = \log \frac{p(x_1, x_2)}{p(x_1)p(x_2)}$$
(41)

This construction is covered in more detail in [37]. The powerset decomposition of entropy and surprisal is the simplest option, by far the most common, and the basis of all Shannon information theory. However, in recent years, other decompositions have been proposed, motivated at least in part by the fact that higher-order mutual information can become negative, which has hindered the operational interpretation of higher-order information quantities.

2. The Partial Information Decomposition

The partial information decomposition (PID) framework, introduced by Williams and Beer [87], proposes that the information a set of predictor variables X contains about a target Y can be decomposed into various terms representing synergistic information (available only in the joint state of the variables), unique information (exclusively contained in a single predictor variable), and redundant information (shared among multiple predictor variables). In their seminal work, Williams and Beer [87] utilized the PID to demonstrate that threeway mutual information is the difference between synergistic and redundant information, with negative values indicating the prevalence of synergistic effects.

10

The PID is constructed by defining the types of information sources that can be built from predictor variables. For instance, given two predictor variables X_1 and X_2 , the information they carry about Y, denoted $I(\{X_1, X_2\}; Y)$, can be decomposed as follows:

$$I({X_1, X_2}; Y) = \Pi({X_1}; Y) + \Pi({X_2}; Y) + \Pi({X_1}, {X_2}; Y) + \Pi({X_1, X_2}; Y)$$
(42)

where $\Pi({X_i}; Y)$ represents the unique information carried by X_i about Y, $\Pi({X_1, X_2}; Y)$ denotes the synergistic information available only from the joint state of the variables, and $\Pi({X_1}, {X_2}; Y)$ represents the redundant information about Y shared by X_1 and X_2 . To construct information sources from an arbitrary set S of predictor variables, one should consider redundancies among all possible combinations of subsets of S. However, note that the redundancy among a set a and a set b reduces to the unique information of a if $a \subseteq b$. This restricts the set of information sources to combinations of predictor subsets that are mutually incomparable by the ordering \subseteq . Such incomparable sets are called *antichains* of the poset ($\mathcal{P}(S), \subseteq$). Two antichains A and B can be ordered by redundancy by setting $A \leq B$ if for every $b \in B$ there is an $a \in A$ such that $a \subseteq b$. That is, $A \leq B$ if B contains at least the redundancy with respect to the outcome Y that A does. The Hasse diagram of this lattice of antichains is shown in Figure 3 for up to |S| = 4. Note that the redundancy lattice is different from both the powerset and partition lattices, but is completely isomorphic up to N = 2 variables (*cf*. Figures 1 and 2).

Let \mathcal{R}_n be the redundancy lattice of *n* variables. The PID decomposition of the information I(S;Y) that a set of variables *S* carries about *Y* can then be written as

$$I(S;Y) = \sum_{R \le S} \Pi(R;Y) \tag{43}$$

This can be inverted to get expressions of the individual contributions $\Pi(R; Y)$ in terms of the mutual information I(T; Y):

$$\Pi(R;Y) = \sum_{T \le R} \mu_{\mathcal{R}_n}(T,R)I(T;Y)$$
(44)

This inversion, however, is far from trivial. One problem is that $|\mathcal{R}_n| = \mathcal{D}(n)$, where $\mathcal{D}(n)$ is the *n*th Dedekind number. The series of Dedekind numbers grows so quickly that it is only known up to D(9) [36, 81]. This makes the Möbius inversion of the PID computationally infeasible for large n, although a closed-form expression for the Möbius function on the redundancy lattice has recently been derived [39].

However, even when the Möbius function can be calculated, this inversion leads to an ambiguity, as the definition of I(T; Y) is unclear when the antichain T contains more than a single set of variables. Consequently, solving the PID requires a well-defined notion of information on arbitrary antichains. A significant portion of the literature on the PID has focused on constructing such definitions, but no consensus has been reached thus far.

Recently, the PID framework has been extended to accommodate multiple target variables and information dynamics in time [56]. The associated Φ ID decomposition lattice is the product of the standard redundancy lattices associated to the individual targets. Since the Möbius function of a product of lattices is the product of Möbius functions of the lattices [78], the Φ ID decomposition reduces to that of the normal PID when the Möbius function of the redundancy lattice is known.

C. Biology

1. Epistasis

In the last decade, improved genetic sequencing of genomes and transcriptomes has revealed that traits often depend on many, if not most genes. This change in perspective has been described as a transition from monogenic, to polygenic and omnigenic models [13]. Furthermore, the effect of a single genetic variant can depend on the presence of other genetic variants, a type of genetic interaction called epistasis. As in the example of a person's height in the introduction, a general phenotype F(G) that depends on the presence of a set of genetic variants $g_i \in G$ can be decomposed into the sum of the effects of subsets $g_i \in \mathcal{P}(G)$ of genetic variants. A Möbius inversion on the lattice of subsets of genetic variants gives the following definitions for genetic effects and epistatic interactions I(g):

$$F(g = \vec{1}, G \setminus g = \vec{0}) = \sum_{s \in \mathcal{P}(g)} I(s)$$

$$\tag{45}$$

$$\iff I(g) = \sum_{s \subseteq g} (-1)^{|s| - |g|} F(s = \vec{1}, G \setminus s = \vec{0})$$

$$\tag{46}$$



FIG. 3. The redundancy terms of a powerset of n variables can be partially ordered: if A and B are redundancy terms, then $A \leq B$ if for every $b \in B$ there is an $a \in A$ such that $a \leq b$. Shown here are the transitive reductions (Hasse diagrams) of the redundancy lattices on 2 (left), 3 (middle) and 4 (right, labels not shown) variables. Note that redundancies among n variables are the antichains in the powerset of n-variables.

That is, the interaction among a set of genes g can be expressed in terms of observed phenotypes of individuals with varying genotypes. For example, the interaction among three genes g_1, g_2, g_3 is given by

$$I_{g_1g_2g_3} = I(\{g_1, g_2, g_3\}) = F(1, 1, 1, \vec{0}) - F(1, 1, 0, \vec{0}) - F(1, 0, 1, \vec{0}) - F(0, 1, 1, \vec{0}) + F(1, 0, 0, \vec{0}) + F(0, 1, 0, \vec{0}) + F(0, 0, 1, \vec{0}) - F(0, 0, 0, \vec{0})$$

$$(47)$$

This expression is indeed the one used to estimate epistasis in much of the literature. The authors of [9] arrive at a similar expression based on volumes of polytopes, which has been used to describe epistasis among genetic variants [22] as well as in the context of microbiomes [29].

Even though interactions are necessarily defined with respect to some outcome, this need not be a phenotype in the traditional sense of the word. One possible outcome that requires no macroscopic measurements, for example, is the population probability of a particular gene expression pattern. If we represent gene expression as a binary vector G, where $G_i = 1$ if gene i is expressed above some threshold, then inverting the log-probability of observing transcriptome G gives the following definition of the interaction among genes g_i :

$$I(g) = \sum_{s \subseteq g} (-1)^{|s| - |g|} \log p(s = 1, g \setminus s = 0)$$
(48)

which can be directly estimated from a population sample. A method to estimate such interactions from singlecell RNA-seq data was recently developed, and used to identify novel cell states and types in various organisms [40].

Instead of inferring the log-probability, i.e. the surprisal, one could also invert its expected value, i.e. the entropy. As explained in Section IIIB, this yields (higher-order) mutual information, which is used by the algorithm that won the DREAM2 challenge of inferring gene regulatory networks from expression data [82].

2. Epidemiology

The central question in epidemiology is to understand the influence of various factors on an individual's health outcomes. One obvious example is estimating the effect a vaccine has on disease risk. However, disease risk after vaccination is not only determined by the vaccine itself, but also by innate factors like the individual's immune system and external factors like the presence of other pathogens, as well as interactions among these. Typically, the effect on outcome Y of a factor X is isolated from all other factors by randomly dividing the population into groups, and then intervening such that X = 1 in group 1, and X = 0 in group 2. The Average Treatment Effect (ATE) is then defined as

$$ATE(Y;X) = E(Y|X=1) - E(Y|X=0)$$
(49)

However, in observational studies, where randomisation followed by intervention is not possible, the effect of X generally cannot be isolated, and the outcome Y depends on a potentially large set of factors S such that:

$$E(Y|S=1) = \sum_{s \in \mathcal{P}(S)} I(s)$$
(50)

where I(s) now represents the effect of the collection of factors s on the outcome Y. The effect of a single factor s = X is then seen to be given by an ATE in the absence of any of the other factors:

$$I(X) = E(Y|X = 1, S \setminus X = 0) - E(Y|X = 0, S \setminus X = 0)$$
(51)

whereas the second- and higher-order terms correspond to interactions among the factors, equivalent to the epidemiological interactions defined in [8]. For example, the interaction among two factors X_1, X_2 is given by

$$I(X_1, X_2) = E(Y|X_1 = 1, X_2 = 1, S \setminus \{X_1, X_2\} = 0) - E(Y|X_1 = 1, X_2 = 0, S \setminus \{X_1, X_2\} = 0)$$
(52)

$$-E(Y|X_1 = 0, X_2 = 1, S \setminus \{X_1, X_2\} = 0) + E(Y|X_1 = 0, X_2 = 0, S \setminus \{X_1, X_2\} = 0)$$

$$ATTP(Y, Y, Y_2, Y_3) = 0 - ATTP(Y, Y_1|Y_2, 0, C) \setminus \{X_1, X_2\} = 0)$$

$$(52)$$

$$=ATE(Y;X_1|X_2=1,S \setminus \{X_1,X_2\}=0) - ATE(Y;X_1|X_2=0,S \setminus \{X_1,X_2\}=0)$$
(53)

In theory, this allows one to estimate drug interactions from population statistics.

3. Neuroscience

As biology tends to deal with very complex systems, decompositions are ubiquitous. The history of neuroscience can accordingly be described by evolving views on what the appropriate decomposition of the brain is: from neurons, to neural circuits, to hierarchical structures, to the currently popular approach that focuses on different spatial and functional regions [7]. To describe the correlations and higher-order relationships between parts of the brain, both the classical approach to higher-order information theory, as well as the PID are commonly used, reflecting concurrent use of both the powerset and redundancy decomposition [24, 33, 59, 72, 86]

D. Physics

1. Equilibrium dynamics

Statistical physics is largely focused on relating the large-scale behaviour of a system to the microscopic interactions. Commonly, this is done through an energy function $E : S \to \mathbb{R}$ that maps a state $s \in S$ of the system to its energy. Writing down a form of E amounts to choosing a decomposition of the system into parts. For example, the Ising model assumes that the energy of a system $S = \{s_1, \ldots, s_n\}, s_i \in \{0, 1\}$, decomposes into singleton and pairwise contributions and is given by

$$E(S=s) = -\sum_{i,j=1}^{n} J_{ij} s_i s_j - \sum_{i=1}^{n} h_i s_i$$
(54)

or, for a subsystem $\hat{S} \subseteq S$

$$E(\hat{S} = \vec{1}, S \setminus \hat{S} = \vec{0}) = -\sum_{i,j: s_i, s_j \in \hat{S}} J_{ij} - \sum_{i: s_i \in \hat{S}}^n h_i$$
(55)

where h_i is the external field acting on spin *i* and J_{ij} is the interaction strength between spins *i* and *j*, typically only nonzero for nearest neighbours on a lattice. Limiting the interactions to linear and pairwise quantities only is an assumption imposed by the Ising model. A more general form of the energy function would be

$$E(S = s) = -\sum_{t \in \mathcal{P}(S)} J_t \prod_{i=1}^{|t|} t_i$$
(56)

where J_t is the interaction strength of the part t and $\mathcal{P}(S)$ denotes the powerset of S. If one is given the parameters J_t , the *forward problem* is to calculate the behaviour of the system S under the influence of the energy function E. The *inverse problem* is to infer the parameters J_t from observations of the system S. The inverse problem is in general intractable, as direct observations of the energy are not possible. However, at equilibrium, the probability of observing a state s is given by the Boltzmann distribution

$$p(s) = \frac{1}{Z}e^{-E(s)}$$
(57)

where $Z = \sum_{s \in S} e^{-E(s)}$ is the partition function. This means that one can observe the energy (up to an unimportant global shift) of a state indirectly as $-\log p(s)$ by simply estimating the probability of a state from

a collection of samples. From this, a Möbius inversion quickly yields precise values for the parameters J_{η} [37]:

$$E(S=\vec{1}) = \sum_{t \in \mathcal{P}(S)} J_t \tag{58}$$

$$\iff J_S = \sum_{t \le S} \mu_P(t, S) \log p(t = 1, S \setminus t = 0)$$
(59)

$$=\sum_{t\leq S} (-1)^{|t|-|S|} \log p(t=1, S \setminus t=0)$$
(60)

For example, the inverse Ising problem for pairwise nearest-neighbour interactions at equilibrium is solved by

$$J_{ij} = \log \frac{p(s_i = 1, s_j = 1, s_{-(i,j)} = 0)p(s_i = 0, s_j = 0, s_{-(i,j)} = 0)}{p(s_i = 1, s_j = 0, s_{-(i,j)} = 0)p(s_i = 0, s_j = 1, s_{-(i,j)} = 0)}$$
(61)

where $s_{-(i,j)}$ denotes all spins except *i* and *j*. Simplifying notation further by writing $p_{abc} = p(s_i = a, s_j = b, s_k = c, s_{-(i,j,k)} = 0)$, the 3-point coupling is given by

$$J_{ijk} = \log \frac{p_{111}p_{100}p_{010}p_{001}}{p_{000}p_{011}p_{101}p_{110}}$$
(62)

This solution to the inverse problem was already noted in [8], but the argument presented here shows that the forward and inverse problem are exactly related through a Möbius inversion.

The chosen powerset decomposition of the energy function restricts this approach to binary variables. However, categorical variables can obey different partial orders but be treated in the same way, leading to a new notion of interaction that is no longer related to spin models. This construction is briefly discussed in [8, 37].

2. Statistical mechanics

In the approach above, one still needs access to observations of all microscopic variables to estimate the energy of observed states and solve the inverse problems. However, one could also start the line of reasoning from more macroscopic quantities, like averages over an ensemble. For example, one might only be able to measure the expected value of variables and their products, called correlation functions. For example, the two-point correlation function is given by

$$\langle X_1 X_2 \rangle = \sum_{x_1, x_2} p(X_1 = x_1, X_2 = x_2) x_1 x_2$$
 (63)

where the summation is over the full state space of the joint system (X_1, X_2) . The observed correlations will be the result of microscopic interactions, but the exact form of the interactions might be unknown. To decompose the correlation function into contributions from individual interactions, it should be noted that to respect the multiplication rule of probability, contributions of multiple interactions happening in a single process should be multiplied. For this reason, a decomposition of a system of variables X into products over partitions is most appropriate:

$$\langle X \rangle = \sum_{\pi \in \Pi(X)} \prod_{\pi_i \in \pi} u(\pi_i) \tag{64}$$

where $\Pi(X)$ is the set of all partitions of X, and $u(\pi_i)$ is the contribution by the members of partition element π_i . For example, the 4-point correlation function can be decomposed into the following contributions, where variables that appear together in a given partition are connected by a line:

$$\langle X_1 X_2 X_3 X_4 \rangle = \frac{x_3 x_4}{x_1 x_2} + \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \left[+ \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \left[+ \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \left[+ \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \right] + \cdots + \left[+ \cdots + \left[$$

For example, a diagram like corresponds to a term $u(X_1, X_3)u(X_2)u(X_4)$, and might be interpreted as the contribution of the situation in which X_1 and X_3 interact, but X_1 and X_4 do not. The correlation functions can then be inverted over the lattice of partitions to give the contributions of individual diagrams, as well as those of terms involving a single u. For example, the first three orders of u are easily seen to be given by

$$u(X_1) = \langle X_1 \rangle \tag{66}$$

$$u(X_1, X_2) = \langle X_1 X_2 \rangle - \langle X_1 \rangle \langle X_2 \rangle \tag{67}$$

$$u(X_1, X_2, X_3) = \langle X_1 X_2 X_3 \rangle - \langle X_1 X_2 \rangle \langle X_3 \rangle - \langle X_1 X_3 \rangle \langle X_2 \rangle - \langle X_2 X_3 \rangle \langle X_1 \rangle + 2 \langle X_1 \rangle \langle X_2 \rangle \langle X_3 \rangle$$
(68)

These are exactly the famous Ursell functions of statistical mechanics. Ursell functions are also called connected correlation functions, because a single u term involves only the connected part of a diagram in the expansion of Equation (65). While the first three Ursell functions coincide with the mixed central moments, the higher-order Ursell functions are different, and used throughout statistical mechanics. The Ursell functions are related to the higher-order interactions of the theory and are essentially the cumulants of the dynamics, which is why they rely on exactly the same decomposition of moments as was described in Section III A. Note, however, that Ursell functions are the partition-inverse of the moments, while the higher-order interactions of equilibrium dynamics are the powerset-inverse of the energy.

3. Quantum & Statistical Field Theory

In quantum field theory, the random variables are replaced by field operators $\phi(x)$. The correlation functions or *Green's functions*—are then defined as expectation values of time-ordered products of field operators in the vacuum state $|\Omega\rangle$:

$$G^{(4)}(x_1, x_2, x_3, x_4) = \langle \Omega | T\phi(x_1)\phi(x_2)\phi(x_3)\phi(x_4) | \Omega \rangle$$
(69)

$$= \mathcal{Z}^{-1} \int \mathcal{D}\phi e^{iS[\phi]} \phi(x_1) \phi(x_2) \phi(x_3) \phi(x_4)$$
(70)

where T is the time-ordering operator, $S[\phi]$ is the action of the field theory, $\mathcal{Z} = \int \mathcal{D}\phi e^{iS[\phi]}$, and the integral is over all possible field configurations. The reason that such time-ordered correlation functions are of great interest, is that they determine the particle scattering amplitudes in a given quantum field theory through the LSZ reduction formula [50, 63]. To calculate the correlation functions, one typically introduces a functional $\mathcal{Z}[J]$, defined as

$$\mathcal{Z}[J] = \int \mathcal{D}\phi e^{iS[\phi] + i\int d^4x J(x)\phi(x)}$$
(71)

such that

$$G^{(4)}(x_1, x_2, x_3, x_4) = \frac{1}{i^4} \frac{\delta^4}{\delta J(x_1) \delta J(x_2) \delta J(x_3) \delta J(x_4)} Z[J] \bigg|_{J=0}$$
(72)

where $\frac{\delta}{\delta J(x)}$ is the functional derivative with respect to J(x). The functional $\mathcal{Z}[J]$ can be expanded in a power series in J, and the coefficients of this expansion are the correlation functions, which is why $\mathcal{Z}[J]$ is called the generating functional, in analogy with the moment generating function of probability theory. Note, however, that the analogy is not perfect: correlation functions are complex-valued, and are not defined with respect to a joint distribution over states so cannot be interpreted as the moments of a probability distribution.

Now, as in the case of statistical classical mechanics, one might decompose the quantum correlation functions over the lattice of partitions:

$$G^{(4)}(X) = \sum_{\pi \in \Pi(X)} \prod_{\pi_i \in \pi} u(\pi_i)$$
(73)

where the $u(\pi_i)$ are the contributions of the members of the partition element π_i . This can of course be represented by the same diagrams as in the case of statistical mechanics in Equation 65, but there is an important difference. By expanding $e^{iS[\phi]}$ as a power series around the non-interacting theory up to certain order and applying Wick's theorem to the resulting products of operators, one can show that the connected correlation functions are given by the sum of all Feynman diagrams with the same external vertices, but with an arbitrary number of internal vertices and loops. To emphasise that the diagrammatic representations of the decomposition may contain arbitrary internal processes, we draw the quantum diagrams with a shaded internal circle:



A single connected component of one of the diagrams now represents an infinite sum over Feynman diagrams. For example, in a theory with quartic interactions:



A Möbius inversion on the lattice of partitions then gives an expression for the connected components of these graphs in terms of the correlation functions. For example:

$$= \langle \Omega | T\phi(x_1)\phi(x_2) | \Omega \rangle - \langle \Omega | \phi(x_1) | \Omega \rangle \langle \Omega | \phi(x_2) | \Omega \rangle - 1$$
(75)

$$= \langle \Omega | T\phi(x_1)\phi(x_2)\phi(x_3) | \Omega \rangle - \langle \Omega | \phi(x_1) | \Omega \rangle \langle \Omega | \phi(x_2)\phi(x_3) | \Omega \rangle$$
(76)

$$-\langle \Omega | \phi(x_2) | \Omega \rangle \langle \Omega | \phi(x_1) \phi(x_3) | \Omega \rangle - \langle \Omega | \phi(x_3) | \Omega \rangle \langle \Omega | \phi(x_1) \phi(x_2) | \Omega \rangle$$
(77)

$$+ 2 \langle \Omega | \phi(x_1) | \Omega \rangle \langle \Omega | \phi(x_2) | \Omega \rangle \langle \Omega | \phi(x_3) | \Omega \rangle$$
(78)

The expansion of a quantum amplitude associated to a scattering process only involves such diagrams where all external lines are connected, so for any given quantum field theory, one can identify scattering amplitudes as the Möbius inverse of correlation functions on the lattice of partitions.

It should be emphasised that while we used the language of quantum field theory throughout this section, similar methods lead to definitions of cumulants in classical- and mean-field theories as well. In fact, classical statistical field theory is directly related to quantum field theory through a Wick rotation (by replacing imaginary time by inverse temperature). While there is a strong analogy to the decomposition of the moments in Section III A, it should be emphasised that the correlation functions treated here are not the moments of some probability distribution over field configurations. Rather, the analogy holds only at the level of the expectation values, and the fact that one can define a moment generating function $\mathcal{Z}[J]$, as well as a cumulant generating function $\log(\mathcal{Z}[J])$. This fact has been used for decades to define cumulants in any setting where a suitable notion of *average* can be defined, and is known as the generalised cumulant expansion method [46]. Such approaches have been used to define connected correlations in networks of neurons [61], on belief propagation graphs [14, 85], and in chemical dynamics [21, 44].

E. Chemistry

A significant challenge in chemistry is predicting the properties of molecules from their configuration, commonly referred to as the quantitative structure-activity relationship (QSAR). To describe a molecular property X in terms of the configuration of a molecule M, one might associate a graph G_M with the molecule, where the nodes are atoms and the edges are bonds. The property X can then be expressed as a function of the subgraphs of G:

$$X(M) = \sum_{G \le G_M} x(G) \tag{79}$$

where x(G) is the contribution of the subgraph G to the property X and the sum is over all subgraphs of G. A straightforward Möbius inversion is of course possible, but some chemical properties are more naturally decomposed over other posets. For example, a molecule's resonance energy is most naturally decomposed into contributions from acyclic graphs only. Motivated such chemical questions, the authors of [5] derived a closed-form expression for the Möbius function on this poset of acyclic graphs.

Also studied in chemistry is how properties of a family of molecules are related. The authors of [34], for instance, study how toxic molecules from the family of Chlorobenzenes are to guppies. A chlorobenzene is a benzene molecule where one or more hydrogen atoms have been replaced by chlorine atoms. That means that there is a certain partial order one can impose on the set of 13 possible chlorobenzenes (taking into account the 6-fold rotational symmetry of the benzene ring). Namely, two chlorobenzenes c_1 and c_2 are related by $c_1 \leq c_2$ if c_2 can be created from c_1 by adding a single chlorine atom. The toxicity T can then be expressed as a sum over the toxic contributions t of the chlorobenzenes that come before it in the reaction chain:

$$T(c) = \sum_{c' \le c} t(c') \tag{80}$$

Similarly, the authors of [35] construct the poset of adding methyl groups to cyclobutanes. As the molecules higher-up in the partial order are constructed from the lower ones, one could expand the property of a molecule in terms of the contributions of the molecules that come before it. This means that every molecule defines its own poset and thus Möbius function, but these have been calculated and used to verify that higher-order contributions decrease so that truncating the expansion at a certain level leads to a good approximation of the property [35].

F. Game Theory

In coalitional game theory, players can form coalitions and cooperate, potentially increasing their expected payoff. The core idea is that value might add synergistically, or superadditively. For example, a pair of shoes might be worth more than twice the value of an individual shoe. Therefore, one could decompose the total value v(S) of a coalition S into the sum of the synergistic contributions w(R) of all subsets of S:

$$v(S) = \sum_{R \subseteq S} w(R) \tag{81}$$

If the value of the coalitions is known, then this can be inverted on the Boolean algebra of subsets of S to yield a definition of the synergistic contributions w(R):

$$w(R) = \sum_{R \subseteq S} (-1)^{|S| - |R|} v(S)$$
(82)

One reason that exact values for the synergy of coalitions might be relevant, is that they are a natural answer to the question of how to distribute the value of a coalition among its members. Since the synergy of a given coalition cannot be attributed to any single member, it should be distributed evenly among all members. The payout any individual player i should then expect in a grand coalition involving all of N players is then the average synergy that player i adds to each possible coalition that includes them:

$$\phi_i = \sum_{S \subseteq N: i \in S} \frac{w(S)}{|S|} \tag{83}$$

This is a very well-known quantity, known as the Shapley value for player i. It is more commonly written as

$$\phi_i = \frac{1}{|N|} \sum_{S \subseteq N \setminus \{i\}} {|N| - 1 \choose |S|} (v(S \cup \{i\}) - v(S))$$
(84)

and is the unique payout function that satisfies a number of favourable properties. In other words: the Shapley value of a player i is the Möbius inverse of the normalised synergy of a coalition on the inclusion lattice of subsets of N that include i.

Choosing a different normalisation in the decomposition of ϕ_i , one that for example depends on the identity of player *i*, the Möbius inversion recovers a family of distribution rules [10].

G. Artificial Intelligence

Modern artificial intelligence systems are complex and involve many interacting parts. While the practical and commercial success of machine learning models has been undeniable, a good understanding of the relationship between the microscopic structure of the model and its emergent macroscopic behaviour is still lacking in most applications. In many situations, like medical applications, a way to link the microscopic structure of a model to its macroscopic behaviour is a crucial step towards widespread adoption.

Predictive machine learning models are generally trained on many features, but once the final model has been constructed it can be difficult to determine how much each feature contributes to the prediction. One way to address this issue is to decompose the prediction of the model into contributions of individual features, and using the Shapley value to assign a total contribution to each feature, replacing the value function v(S) with the model's prediction, and marginalising over the features not in S [15, 26, 88]. This allows one to determine which features are most important for a particular model's prediction, as well as which groups of features show synergistic effects.

In generative machine learning, the features do not contribute to a prediction, but to a probability distribution. Since energy-based models have a closed-form expression for the generated probability distribution, they allow for a precise understanding of the relationship between their internal structure and the encoded distribution. One architecture that has led to particularly fruitful insights has been the Restricted Boltzmann machine, in which the probability of a sample is essentially the Boltzmann distribution of a statistical physics model. An argument similar as above reveals that the cumulants of the model are the Möbius inverse of the mixed moments of the model and describe the interactions among the features, with higher-order interactions corresponding to synergistic effects among the features [17, 37, 60].

Even before the modern success of machine learning techniques, artificial intelligence has created formal methods to treat reasoning agents. Dempster-Shafer theory is a generalisation of Bayesian probability theory [18, 71] (though its precise relationship to Bayesian inference is controversial [20]) that formalises how agents combine evidence to update beliefs. It decomposes the set of all properties of the universe X into the powerset of propositions about X. Each element (A) from the powerset $\mathcal{P}(X)$ then gets assigned a certain *belief mass* through the basic belief assignment function $m : \mathcal{P}(X) \to [0, 1]$, where $\sum_{A \in \mathcal{P}(X)} m(A) = 1$. The total belief an agent has about a proposition A is then given by the mass of A and the mass of all subsets of A

$$Bel(A) = \sum_{B \subseteq A} m(B) \tag{85}$$

such that the belief assignment function can be expressed in terms of the belief of agents after observing varying evidence

$$m(A) = \sum_{B \subseteq A} (-1)^{|A| - |B|} Bel(B)$$
(86)

which is used to define the generalised Bayesian update rule. While classical Dempster-Shafer theory always used the powerset decomposition, more modern approaches have generalised this to other lattices where properties of the Möbius function are known [30, 90].

Field of Study	Macro Quantity	Decomposition	Micro Quantity/Interactions
Statistics	Moments	Powerset	Central moments
	Moments	Partitions	Cumulants
	Free moments	Non-crossing partitions	Free cumulants
	Path signature moments	Ordered partitions	Path signature cumulants
Information Theory	Entropy	Powerset	Mutual information
	Surprisal	Powerset	Pointwise mutual information
	Joint Surprisal	Powerset	Conditional interactions
	Mutual Information	Antichains	Synergy/redundancy atoms
Biology	Pheno- & Genotype	Powerset	Epistasis
	Gene expression profile	Powerset	Genetic interactions
	Population statistics	Powerset	Synergistic treatment effects
Physics	Ensemble energies	Powerset	Ising interactions
	Correlation functions	Partitions	Ursell functions
	Quantum corr. functions	Partitions	Scattering amplitudes
Chemistry	Molecular property	Subgraphs	Fragment contributions
	Molecular property	Reaction poset	Cluster contributions
Game Theory	Coalition value	Powerset	Coalition synergy
	Shapley value	Powerset	Normalised coalition synergy
Artificial intelligence	Generative model probabilities	Powerset	Feature interactions
	Predictive model predictions	Powerset	Feature contributions
	Dempster-Shafer Belief	Distributive	Evidence weight

IV. DISCUSSION

TABLE I. An overview of the various ways in which macroscopic quantities can be linked to microscopic interactions by the Möbius inversion associated with a certain decomposition.

The aim of this study has been to provide a unified perspective on the various notions of higher-order interactions in complex systems. We have shown that decomposing a system into parts leads to a unique definition of the interactions among the parts, through a Möbius inversion of the outcome of interest. We found that this approach reproduces well-known notions of higher-order structure in a variety of scientific fields, an overview of which is provided in Table I. While some of the relationships in Table I have previously been described as Möbius inversions, to our knowledge this is the first time that the shared structure underlying these definitions has been made explicit.

The presented framework shows how the choice of system decomposition determines the nature of the derived interactions. This highlights decomposition selection as a key step in studying complex system behaviour. While there are no universal criteria for a "good" decomposition, in each case the decomposition should be motivated by knowledge of the system's structure and the macroscopic property of interest. For example, a partition-based decomposition may be natural when the property depends on how different parts come together, while a powerset-based decomposition is more suited to properties that depend binary configurations of the parts.

Fundamentally, this framework mathematises the notion that complex system properties depend on the way the system is carved up into parts. Different decompositions reveal different kinds of interactions; there is no single, privileged decomposition. The "correct" decomposition is the one that yields meaningful interactions for the question of interest, as seen in the various examples.

One way to interpret the Möbius inversion theorem is as a relationship between the macro- and microscopic degrees of freedom of a system. If $f(S) = \sum_{t \leq S} g(t)$, then since f is the sum of contributions of all parts, it can be considered a macroscopic feature of the system. The Möbius inversion theorem then states that the microscopic contributions g(t) can be recovered from the macroscopic feature f observed on the parts. In this

sense, it provides a way to study the emergent properties of complex systems, but it should be emphasised that the inversion is only possible if the macroscopic quantity is also defined for the smaller parts. For many classic examples of emergence, like bird flocks, temperature, *etc*, this is not possible.

An interesting example of this problem arose in the partial information decomposition. With the decomposition defined on the redundancy lattice of antichains, estimating the contribution of the information atoms is only possible given a suitable notion of redundant information on each of the antichains. Much of the PID literature has focused on resolving this ambiguity in different ways, but no concensus has been reached so far. However, this also highlights how versatile the Möbius inversion framework is: even when the decomposition results in ambiguity, different resolutions of this ambiguity can give rise to a rich set of higher-order structure, which in the case of the PID have been used to characterise different properties of neural information processing [53, 72, 86].

As a general rule, observations of a system provide access to the macroscopic features (moments, phenotypes, energies, predictions, *etc*), which can then be used to infer the microscopic interactions (cumulants, genetic interactions, Ising interactions, feature contributions, *etc*), so the Möbius inversion theorem is essentially a solution to the *inverse* problem of inferring microscopic structure from macroscopic observables. Note, however, that it solves the inverse problem relative to a chosen decomposition. If the chosen decomposition is not appropriate, then the Möbius inversion theorem will not provide a meaningful interpretation of the microscopic interactions either. One example of this is the Möbius inversion on the lattice of partitions that related scattering amplitudes of quantum field theory to correlation functions. The Möbius inversion still gave infinite sums over Feynman diagrams, not the contribution of individual diagrams, so the microscopic interactions were not directly accessible from the macroscopic observables. Furthermore, collections of particles can carry more structure, like fermion number and colour, which also implies that partitions are not the most appropriate decomposition for the problem.

One interesting phenomenon not explored in the present study is that of lattice dualities. It was observed in [37] that since the order-theoretic dual of a lattice is again a lattice, there are dual higher-order structures implied by a given decomposition lattice. In information theory, the dual to information theory was found to be conditional mutual information, with a similar dual quantity derived for Ising interactions. It is an unexplored but interesting question whether the quantities dual to the ones presented here also offer a meaningful interpretation, but this is left for future work.

Another exciting direction for future work is to use the current framework to transfer insights from one scientific discipline to another. For example, in the so-called *cluster variation method*, physical intuition has motivated the truncation of sums over lattice subsets to obtain approximations to thermodynamic quantities of crystals [42]. This summation can be seen as a truncated Möbius inversion over the lattice of physical crystal lattice sites [1, 58], and could be explored as an approximation scheme in other settings as well. In fact, precisely such a truncation has been suggested in decompositions of chemical properties of molecules [35]. In addition, the authors of [40] use causal discovery methods to improve the estimation of genetic interactions (the Möbius inverse of gene-expression profiles). Since the estimation of microscopic interactions from macroscopic observables can require many observations, similar methods could be used to improve the estimability of higher-order interactions in other fields.

Other future work could explore different kinds of decompositions and the higher-order structure they imply. For example, a recent study has started to explore the combinatorics of nucleotide sequences in polyploid genomes by decomposing sequencing data over the lattice of integer partitions ordered by refinement [65]. More speculatively, decompositions beyond locally finite partial orders could be explored by suitably generalising the Möbius inversion theorem. It is well-known that the Möbius inversion theorem can be generalised to a more general class of skeletal categories that includes not just posets but also monoids and groupoids [16, 51]. Whether these more general decompositions faithfully and fruitfully describe higher-order structure in complex systems is, to our best knowledge, a mostly unexplored and open question. We hope that this work inspires others to explore novel decompositions and discover new types of higher-order structures in complex systems.

ACKNOWLEDGEMENTS

The connections to game theory, chemistry, and path statistics were the result of fruitful discussions with Jürgen Jost, Guillermo Restrepo, and Darrick Lee, respectively. The author is also grateful to Eckehard Olbrich and Fernando Rosas for insightful discussions on the partial information decomposition, as well as to Hadleigh Frost for his comments on diagrammatic sums in quantum field theory.

^[1] G. An. A note on the cluster variation method. Journal of Statistical Physics, 52:727-734, 1988.

^[2] Y. E. Antebi, J. M. Linton, H. Klumpe, B. Bintu, M. Gong, C. Su, R. McCardell, and M. B. Elowitz. Combinatorial signal perception in the bmp pathway. *Cell*, 170(6):1184–1196, 2017.

- [3] D. N. Arnosti, S. Barolo, M. Levine, and S. Small. The eve stripe 2 enhancer employs multiple modes of transcriptional synergy. *Development*, 122(1):205–214, 1996.
- [4] N. Ay, E. Olbrich, N. Bertschinger, and J. Jost. A geometric approach to complexity. Chaos: An Interdisciplinary Journal of Nonlinear Science, 21(3), 2011.
- [5] D. Babić and N. Trinajstić. Möbius inversion on a poset of a graph and its acyclic subgraphs. Discrete applied mathematics, 67(1-3):5-11, 1996.
- [6] F. Battiston, E. Amico, A. Barrat, G. Bianconi, G. Ferraz de Arruda, B. Franceschiello, I. Iacopini, S. Kéfi, V. Latora, Y. Moreno, et al. The physics of higher-order interactions in complex systems. *Nature Physics*, 17(10):1093–1098, 2021.
- [7] W. Bechtel. Decomposing the mind-brain: A long-term pursuit. Brain and Mind, 3:229–242, 2002.
- [8] S. V. Beentjes and A. Khamseh. Higher-order interactions in statistical physics and machine learning: A modelindependent solution to the inverse problem at equilibrium. *Physical Review E*, 102(5):053314, 2020.
- [9] N. Beerenwinkel, L. Pachter, and B. Sturmfels. Epistasis and shapes of fitness landscapes. Statistica Sinica, pages 1317–1342, 2007.
- [10] A. Billot and J.-F. Thisse. How to share when context matters: The möbius value as a generalized solution for cooperative games. *Journal of Mathematical Economics*, 41(8):1007–1029, 2005.
- [11] P. Bonnier and H. Oberhauser. Signature cumulants, ordered partitions, and independence of stochastic processes. 2020.
- [12] S. P. Borgatti, A. Mehra, D. J. Brass, and G. Labianca. Network analysis in the social sciences. science, 323(5916):892–895, 2009.
- [13] E. A. Boyle, Y. I. Li, and J. K. Pritchard. An expanded view of complex traits: from polygenic to omnigenic. Cell, 169(7):1177–1186, 2017.
- [14] M. Chertkov and V. Y. Chernyak. Loop series for discrete statistical models on graphs. Journal of Statistical Mechanics: Theory and Experiment, 2006(06):P06009, 2006.
- [15] S. Cohen, G. Dror, and E. Ruppin. Feature selection via coalitional game theory. Neural computation, 19(7):1939– 1961, 2007.
- [16] M. Content, F. Lemay, and P. Leroux. Catégories de möbius et fonctorialités: un cadre général pour l'inversion de möbius. Journal of Combinatorial Theory, Series A, 28(2):169–190, 1980.
- [17] G. Cossu, L. Del Debbio, T. Giani, A. Khamseh, and M. Wilson. Machine learning determination of dynamical parameters: The ising model case. *Physical Review B*, 100(6):064304, 2019.
- [18] A. Dempster. Upper and lower probabilities induced by a multi-valued mapping. Annals of Mathematical Statistics, 38, 1967.
- [19] D. C. Dennett. Darwin's dangerous idea. The Sciences, 35(3):34-40, 1995.
- [20] J. Dezert, A. Tchamova, D. Han, and J.-M. Tacnet. Why dempster's fusion rule is not a generalization of bayes fusion rule. In *Proceedings of the 16th International Conference on Information Fusion*, pages 1127–1134. IEEE, 2013.
- [21] C. Domb and G. Joyce. Cluster expansion for a polymer chain. *Journal of Physics C: Solid State Physics*, 5(9):956, 1972.
- [22] H. Eble, M. Joswig, L. Lamberti, and W. B. Ludington. Master regulators of biological systems in higher dimensions. Proceedings of the National Academy of Sciences, 120(51):e2300634120, 2023.
- [23] M. Eidi and J. Jost. Ollivier ricci curvature of directed hypergraphs. Scientific Reports, 10(1):12466, 2020.
- [24] A. Erramuzpe, G. J. Ortega, J. Pastor, R. G. De Sola, D. Marinazzo, S. Stramaglia, and J. M. Cortes. Identification of redundant and synergetic circuits in triplets of electrophysiological data. *Journal of neural engineering*, 12(6):066007, 2015.
- [25] J. W. Essam. Graph theory and statistical physics. Discrete Mathematics, 1(1):83–112, 1971.
- [26] D. Fryer, I. Strümke, and H. Nguyen. Shapley values for feature selection: The good, the bad, and the axioms. *Ieee Access*, 9:144352–144360, 2021.
- [27] E. Gardner. Spin glasses with p-spin interactions. Nuclear Physics B, 257:747–765, 1985.
- [28] G. Ghoshal, V. Zlatić, G. Caldarelli, and M. E. Newman. Random hypergraphs and their applications. *Physical Review E*, 79(6):066118, 2009.
- [29] A. L. Gould, V. Zhang, L. Lamberti, E. W. Jones, B. Obadia, N. Korasidis, A. Gavryushkin, J. M. Carlson, N. Beerenwinkel, and W. B. Ludington. Microbiome interactions shape host fitness. *Proceedings of the National Academy of Sciences*, 115(51):E11951–E11960, 2018.
- [30] M. Grabisch. Belief functions on lattices. International Journal of Intelligent Systems, 24(1):76–95, 2009.
- [31] M. Grandjean. A social network analysis of twitter: Mapping the digital humanities community. Cogent Arts & Humanities, 3(1):1171458, 2016.
- [32] G. M. Hodgson. Generalizing darwinism to social evolution: Some early attempts. Journal of Economic Issues, 39(4):899–914, 2005.
- [33] R. A. Ince, B. L. Giordano, C. Kayser, G. A. Rousselet, J. Gross, and P. G. Schyns. A statistical framework for neuroimaging data analysis based on mutual information estimated via a gaussian copula. *Human brain mapping*, 38(3):1541–1573, 2017.
- [34] T. Ivanciuc, O. Ivanciuc, and D. J. Klein. Posetic quantitative superstructure/activity relationships (qssars) for chlorobenzenes. Journal of chemical information and modeling, 45(4):870–879, 2005.
- [35] T. Ivanciuc, D. J. Klein, and O. Ivanciuc. Posetic cluster expansion for substitution-reaction networks and application to methylated cyclobutanes. *Journal of Mathematical Chemistry*, 41:355–379, 2007.
- [36] C. Jäkel. A computation of the ninth dedekind number. arXiv preprint arXiv:2304.00895, 2023.
- [37] A. Jansma. Higher-order interactions and their duals reveal synergy and logical dependence beyond shannoninformation. *Entropy*, 25(4):648, 2023.
- [38] A. Jansma. Higher-order interactions in single-cell gene expression. PhD thesis, University of Edinburgh, 2023.

- [39] A. Jansma and F. E. Rosas. Decomposing information on the free distributive lattice (forthcoming). 2024.
- [40] A. Jansma, Y. Yao, J. Wolfe, L. Del Debbio, S. V. Beentjes, C. P. Ponting, and A. Khamseh. High order expression dependencies finely resolve cryptic states and subtypes in single cell data. *bioRxiv*, pages 2023–12, 2023.
- [41] E. T. Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957.
- [42] R. Kikuchi. A theory of cooperative phenomena. *Physical review*, 81(6):988, 1951.
- [43] T. R. Kirkpatrick and D. Thirumalai. p-spin-interaction spin-glass models: Connections with the structural glass problem. *Physical Review B*, 36(10):5388, 1987.
- [44] D. Klein, T. Schmalz, and L. Bytautas. Chemical sub-structural cluster expansions for molecular properties. SAR and QSAR in Environmental Research, 10(2-3):131–156, 1999.
- [45] H. E. Klumpe, M. A. Langley, J. M. Linton, C. J. Su, Y. E. Antebi, and M. B. Elowitz. The context-dependent, combinatorial logic of bmp signaling. *Cell systems*, 13(5):388–407, 2022.
- [46] R. Kubo. Generalized cumulant expansion method. Journal of the Physical Society of Japan, 17(7):1100–1120, 1962.
- [47] E. Kuzmin, B. VanderSluis, W. Wang, G. Tan, R. Deshpande, Y. Chen, M. Usaj, A. Balint, M. Mattiazzi Usaj, J. Van Leeuwen, et al. Systematic analysis of complex genetic interactions. *Science*, 360(6386):eaao1729, 2018.
- [48] L. Lang, P. Baudot, R. Quax, and P. Forré. Information decomposition diagrams applied beyond shannon entropy: A generalization of hu's theorem. arXiv preprint arXiv:2202.09393, 2022.
- [49] W. Leal, G. Restrepo, P. F. Stadler, and J. Jost. Forman-ricci curvature for hypergraphs. Advances in Complex Systems, 24(01):2150003, 2021.
- [50] H. Lehmann, K. Symanzik, and W. Zimmermann. Zur formulierung quantisierter feldtheorien. Il Nuovo Cimento (1955-1965), 1:205-225, 1955.
- [51] T. Leinster. Notions of möbius inversion. Bulletin of the Belgian Mathematical Society-Simon Stevin, 19(5):909–933, 2012.
- [52] Q. F. Lotito, F. Musciotto, A. Montresor, and F. Battiston. Higher-order motif analysis in hypergraphs. Communications Physics, 5(1):79, 2022.
- [53] A. I. Luppi, P. A. Mediano, F. E. Rosas, N. Holland, T. D. Fryer, J. T. O'Brien, J. B. Rowe, D. K. Menon, D. Bor, and E. A. Stamatakis. A synergistic core for human brain evolution and cognition. *Nature Neuroscience*, 25(6):771–782, 2022.
- [54] H. J. Maris and L. P. Kadanoff. Teaching the renormalization group. American journal of physics, 46(6):652–657, 1978.
- [55] O. Mason and M. Verwoerd. Graph theory and networks in biology. *IET systems biology*, 1(2):89–119, 2007.
- [56] P. A. Mediano, F. E. Rosas, A. I. Luppi, R. L. Carhart-Harris, D. Bor, A. K. Seth, and A. B. Barrett. Towards an extended taxonomy of information dynamics via integrated information decomposition. arXiv preprint arXiv:2109.13186, 2021.
- [57] L. Merchan and I. Nemenman. On the sufficiency of pairwise interactions in maximum entropy models of networks. *Journal of Statistical Physics*, 162:1294–1308, 2016.
- [58] T. Morita. Cluster variation method and möbius inversion formula. Journal of Statistical Physics, 59:819–825, 1990.
- [59] E. L. Newman, T. F. Varley, V. K. Parakkattu, S. P. Sherrill, and J. M. Beggs. Revealing the dynamics of neural information processing with multivariate information decomposition. *Entropy*, 24(7):930, 2022.
- [60] H. C. Nguyen, R. Zecchina, and J. Berg. Inverse statistical problems: from the inverse ising problem to data science. Advances in Physics, 66(3):197–261, 2017.
- [61] G. K. Ocker, K. Josić, E. Shea-Brown, and M. A. Buice. Linking structure and activity in nonlinear spiking networks. *PLoS computational biology*, 13(6):e1005583, 2017.
- [62] G. A. Pavlopoulos, M. Secrier, C. N. Moschopoulos, T. G. Soldatos, S. Kossida, J. Aerts, R. Schneider, and P. G. Bagos. Using graph theory to analyze biological networks. *BioData mining*, 4:1–27, 2011.
- [63] M. E. Peskin. An introduction to quantum field theory. CRC press, 2018.
- [64] Plato. Phaedrus 265e. http://data.perseus.org/citations/urn:cts:greekLit:tlg0059.tlg012.perseus-eng1: 265e, 370 BC.
- [65] T. R. Ranallo-Benavidez, K. S. Jaron, and M. C. Schatz. Genomescope 2.0 and smudgeplot for reference-free profiling of polyploid genomes. *Nature communications*, 11(1):1432, 2020.
- [66] R. Rosen. Some results in graph theory and their application to abstract relational biology. The bulletin of mathematical biophysics, 25(2):231–241, 1963.
- [67] G.-C. Rota. On the foundations of combinatorial theory i. theory of möbius functions. Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete, 2(4):340–368, 1964.
- [68] G.-C. Rota and J. Shen. On the combinatorics of cumulants. *Journal of Combinatorial Theory, Series A*, 91(1-2):283–304, 2000.
- [69] H. Schawe and L. Hernández. Higher order interactions destroy phase transitions in deffuant opinion dynamics model. Communications Physics, 5(1):32, 2022.
- [70] M. Scheffer, J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. Van Nes, M. Rietkerk, and G. Sugihara. Early-warning signals for critical transitions. *Nature*, 461(7260):53–59, 2009.
- [71] G. Shafer. A mathematical theory of evidence, volume 42. Princeton university press, 1976.
- [72] S. P. Sherrill, N. M. Timme, J. M. Beggs, and E. L. Newman. Partial information decomposition reveals that synergistic neural integration is greater downstream of recurrent information flow in organotypic cortical cultures. *PLoS computational biology*, 17(7):e1009196, 2021.
- [73] P. S. Skardal and A. Arenas. Higher order interactions in complex networks of phase oscillators promote abrupt synchronization switching. *Communications Physics*, 3(1):218, 2020.
- [74] R. Solé. Phase transitions, volume 3. Princeton University Press, 2011.
- [75] T. Speed and H. Silcock. Cumulants and partition lattices v. calculating generalized k-statistics. Journal of the Australian Mathematical Society, 44(2):171–196, 1988.

- [76] R. Speicher. Free probability theory and non-crossing partitions. Séminaire Lotharingien de Combinatoire [electronic only], 39:B39c-38, 1997.
- [77] T. Spribille, V. Tuovinen, P. Resl, D. Vanderpool, H. Wolinski, M. C. Aime, K. Schneider, E. Stabentheiner, M. Toome-Heller, G. Thor, et al. Basidiomycete yeasts in the cortex of ascomycete macrolichens. *Science*, 353(6298):488–492, 2016.
- [78] R. P. Stanley. Enumerative combinatorics volume 1 second edition. *Cambridge studies in advanced mathematics*, 2011.
- [79] T. Tanaka and T. Aoyagi. Multistable attractors in a network of phase oscillators with three-body interactions. *Physical Review Letters*, 106(22):224101, 2011.
- [80] G. Tkacik, E. Schneidman, M. J. Berry II, and W. Bialek. Ising models for networks of real neurons. arXiv preprint q-bio/0611072, 2006.
- [81] L. Van Hirtum, P. De Causmaecker, J. Goemaere, T. Kenter, H. Riebler, M. Lass, and C. Plessl. A computation of d (9) using fpga supercomputing. arXiv preprint arXiv:2304.03039, 2023.
- [82] J. Watkinson, K.-c. Liang, X. Wang, T. Zheng, and D. Anastassiou. Inference of regulatory gene interactions from expression data using three-way mutual information. *Annals of the New York Academy of Sciences*, 1158(1):302–313, 2009.
- [83] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. nature, 393(6684):440-442, 1998.
- [84] S. Weinberg. The quantum theory of fields, volume 2. Cambridge university press, 1995.
- [85] M. Welling, A. E. Gelfand, and A. T. Ihler. A cluster-cumulant expansion at the fixed points of belief propagation. arXiv preprint arXiv:1210.4916, 2012.
- [86] M. Wibral, C. Finn, P. Wollstadt, J. T. Lizier, and V. Priesemann. Quantifying information modification in developing neural networks via partial information decomposition. *Entropy*, 19(9):494, 2017.
- [87] P. L. Williams and R. D. Beer. Nonnegative decomposition of multivariate information. arXiv preprint arXiv:1004.2515, 2010.
- [88] B. Williamson and J. Feng. Efficient nonparametric statistical inference on population feature importance using shapley values. In *International conference on machine learning*, pages 10282–10291. PMLR, 2020.
- [89] R. W. Yeung. A new outlook on shannon's information measures. IEEE transactions on information theory, 37(3):466–474, 1991.
- [90] C. Zhou. Belief functions on distributive lattices. Artificial Intelligence, 201:1–31, 2013.
- [91] W. H. Zurek. Quantum darwinism. Nature physics, 5(3):181-188, 2009.