

Multivariate Dependence Beyond Shannon Information

Ryan G. James* and James P. Crutchfield†

*Complexity Sciences Center and
Physics Department
University of California at Davis,
One Shields Avenue, Davis, CA 95616*

(Dated: September 13, 2016)

Accurately determining dependency structure is critical to discovering a system’s causal organization. We recently showed that the transfer entropy fails in a key aspect of this—measuring information flow—due to its conflation of dyadic and polyadic relationships. We extend this observation to demonstrate that this is true of all such Shannon information measures when used to analyze multivariate dependencies. This has broad implications, particularly when employing information to express the organization and mechanisms embedded in complex systems, including the burgeoning efforts to combine complex network theory with information theory. Here, we do not suggest that any aspect of information theory is *wrong*. Rather, the vast majority of its informational measures are simply inadequate for determining the meaningful dependency structure within joint probability distributions. Therefore, such information measures are inadequate for discovering intrinsic causal relations. We close by demonstrating that such distributions exist across an arbitrary set of variables.

PACS numbers: 89.70.+c 05.45.Tp 02.50.Ey 02.50.-r

Keywords: stochastic process, transfer entropy, causation entropy, partial information decomposition, network science, causality

INTRODUCTION

Information theory is a general, broadly applicable framework for understanding a system’s statistical properties [1]. Due to its abstraction from details of the underlying system coordinates and its focus on probability distributions, it has found many successes outside of its original domain of communication in the physical, biological, and social sciences [2–19]. Often, the issue on which it is brought to bear is discovering and quantifying dependencies or causal relations [20–25].

The past two decades, however, produced a small but important body of results detailing how standard Shannon information measures are unsatisfactory for determining some aspects of dependency and shared information. Within information-theoretic cryptography, the conditional mutual information has proven to be a poor bound on secret key agreement [26, 27]. The conditional mutual information has also been shown to be unable to accurately measure information flow [28, and references therein]. Finally, the inability of standard methods of decomposing the joint entropy to provide any semantic understanding of how information is shared has motivated entirely new methods of decomposing information [29, 30]. Common to all these is the fact that conditional mutual information conflates intrinsic dependence with conditional dependence.

Here, we demonstrate a related, but deeper issue: Shan-

non information measures—entropy, mutual information, and their conditional and multivariate versions—can fail to distinguish joint distributions with vastly differing internal dependencies.

Concretely, we start by constructing two joint distributions, one with dyadic subdependencies and the other with strictly triadic subdependencies. From there, we demonstrate that no standard Shannon-like information measure, and exceedingly few nonstandard methods, can distinguish the two. Stately plainly: when viewed through Shannon’s lens, these two distributions are erroneously equivalent. While distinguishing these two (and their internal causality) may not be relevant to a mathematical theory of communication, it is absolutely critical to a mathematical theory of information storage, transfer, and modification [31]. We then demonstrate two ways in which these failures generalize to the multivariate case. The first generalizes our two distributions to the multivariate and polyadic case via “dyadic camouflage”. And, the second details a method of embedding an arbitrary distribution into a larger variable space using hierarchical dependencies, a technique we term “dependency diffusion”. In this way, one sees that the initial concerns about information measures can arise in virtually any statistical multivariate analysis.

In this short development, we assume a working knowledge of information theory, such as found in standard textbooks [32–35].

X	Y	Z	Pr	X	Y	Z	Pr
0	0	0	1/8	0	0	0	1/8
0	2	1	1/8	1	1	1	1/8
1	0	2	1/8	0	2	2	1/8
1	2	3	1/8	1	3	3	1/8
2	1	0	1/8	2	0	2	1/8
2	3	1	1/8	3	1	3	1/8
3	1	2	1/8	2	2	0	1/8
3	3	3	1/8	3	3	1	1/8

(a) Dyadic

(b) Triadic

TABLE I. The (a) dyadic and (b) triadic probability distributions over the three random variables X , Y , and Z that take values in the four-letter alphabet $\{0, 1, 2, 3\}$. Though not directly apparent from their tables of joint probabilities, the dyadic distribution is built from dyadic (pairwise) subdependencies while the triadic from triadic (three-way) subdependencies.

DEVELOPMENT

We begin by considering the two joint distributions shown in Table I. The first represents *dyadic* relationships among three random variables X , Y , and Z . And, the second *triadic*¹ among them. These appellations are used for reasons that will soon be apparent. How are these distributions structured? Are they structured identically, or are they qualitatively distinct?

We can develop a direct picture of underlying dependency structure by casting the random variables' four-symbol alphabet used in Table I into composite binary random variables, as displayed in Table II. It can be readily verified that the dyadic distribution follows three simple rules: $X_0 = Y_1$, $Y_0 = Z_1$, and $Z_0 = X_1$. Three dyadic rules. The triadic distribution similarly follows simple rules: $X_0 + Y_0 + Z_0 = 0 \pmod 2$ (the XOR relation [37], or, equivalently, any one of them is the XOR of the other two), and $X_1 = Y_1 = Z_1$. Two triadic rules. These dependency structures are represented pictorially in Fig. 1. Our development from this point on will not use any knowledge of this structure, but rather it will attempt to determine the structure using only information measures. Bayesian networks [38] underlie common methods of dependency determination. And so, naturally, we applied Bayesian network inference to the two distributions; specifically, the GS, IAMB, Fast-IAMB, Inter-IAMB, MMPC, SI-HITON-PC, HC, Tabu, MMHC, and RSMAX2 methods of the

X		Y		Z		Pr	X		Y		Z		Pr
X_0	X_1	Y_0	Y_1	Z_0	Z_1		X_0	X_1	Y_0	Y_1	Z_0	Z_1	
0	0	0	0	0	0	1/8	0	0	0	0	0	0	1/8
0	0	1	0	0	1	1/8	0	1	0	1	0	1	1/8
0	1	0	0	1	0	1/8	0	0	1	0	1	0	1/8
0	1	1	0	1	1	1/8	0	1	1	1	1	1	1/8
1	0	0	1	0	0	1/8	1	0	0	0	1	0	1/8
1	0	1	1	0	1	1/8	1	1	0	1	1	1	1/8
1	1	0	1	1	0	1/8	1	0	1	0	0	0	1/8
1	1	1	1	1	1	1/8	1	1	1	1	0	1	1/8

(a) Dyadic

(b) Triadic

TABLE II. Expansion of the (a) dyadic and (b) triadic distributions. In both cases, the variables from Table I were interpreted as two binary random variables, translating *e.g.* $X = 3$ into $(X_0, X_1) = (1, 1)$. In this light, it becomes apparent that the dyadic distribution consists of the subdependencies $X_0 = Y_1$, $Y_0 = Z_1$, and $Z_0 = X_1$ while the triadic distribution consists of $X_0 + Y_0 + Z_0 = 0 \pmod 2$ and $X_1 = Y_1 = Z_1$. These relationships are pictorially represented in Fig. 1.

bnlearn [39] package, version 4.0. Each of these methods, when run with its default parameters, resulted in the Bayesian network depicted in Fig. 2: three lone nodes. Despite the popularity of Bayesian networks for modeling dependencies, the failure of these methods is not surprising: the dyadic and triadic distributions violate a basic premise of Bayesian networks. Namely, that their dependency structure cannot be represented by a directed acyclic graph. This also implies that the methods of Pearl [40] and its generalizations [41] give no insight on structure in these distributions. And so, we leave Bayesian network inference aside.

What does an informational analysis say? Both the dyadic and triadic distributions describe events over three variables, each with an alphabet size of four. Each consists of eight joint events, each with a probability of $1/8$. As such, each has a joint entropy of $H[X, Y, Z] = 3 \text{ bit}$.² Our first observation is that *any* entropy—conditional or not—and any mutual information—conditional or not—will be identical for the two distributions. Specifically, the entropy of any variable conditioned on the other two vanishes: $H[X | Y, Z] = H[Y | X, Z] = H[Z | X, Y] = 0 \text{ bit}$; the mutual information between any two variables conditioned on the third is unity: $I[X : Y | Z] =$

¹ This distribution was first considered in Ref. [36], though for other reasons.

² The SI standard *symbol* for information is bit, analogous to s being the symbol for time. As such it is inappropriate to write 3 bits, just as it would be inappropriate to write 3 ss.




FIG. 1. Dependency structure for the (a) dyadic and (b) triadic distributions. Here, \sim denotes that two or more variables are distributed identically and \oplus denotes the enclosed variables form the XOR relation. Note that although these dependency structures are fundamentally distinct, their Bayesian network (Fig. 1) and their information diagrams (Fig. 3) are identical.




FIG. 2. Graphical depiction of the result from applying several Bayesian network inference algorithms to both the dyadic and triadic distributions. The algorithms do not find any dependence among the variables X , Y , and Z , resulting in three disconnected nodes. The algorithms' failure is not surprising: the dyadic and triadic distributions cannot be represented by a directed acyclic graph, a basic assumption of Bayesian network inference.




FIG. 3. Information diagrams for the (a) dyadic and (b) triadic distributions. For the three variable distributions depicted here, the diagram consists of seven atoms: three conditional entropies (each with value 0 bit), three conditional mutual informations (each with value 1 bit), and one co-information (0 bit). Note that the two diagrams are *identical*, meaning that although the two distributions are fundamentally distinct, no standard information-theoretic measure can differentiate the two.

$I[X : Z | Y] = I[Y : Z | X] = 1$ bit; and the three-way co-information also vanishes: $I[X : Y : Z] = 0$ bit. These conclusions are compactly summarized in the form of the *information diagrams* (I-diagrams) [42, 43] shown in Fig. 3. This diagrammatically represents all of the possible Shannon information measures (*I-measures*) [43] of the distribution: effectively, all the multivariate extensions of the standard Shannon measures, called *atoms*. The values of the information atoms are identical.

As a brief aside, it is of interest to note that it has been suggested (*e.g.*, in Refs. [44, 45], among others) that zero co-information implies that at least one variable is independent of the others—that is, in this case, a lack of three-way interactions. Krippendorff [46] early on demonstrated that this is not the case, though these examples more clearly exemplify this fact.

We now turn to the implications of the two information

diagrams, Figs. 3a and 3b, being identical. There are measures [20, 22, 44, 47–53] and expansions [54–56] purporting to measure or otherwise extract the complexity, magnitude, or structure of dependencies within a multivariate distribution. Many of these techniques, including those just cited, are sums and differences of atoms in these information diagrams. As such, they are unable to differentiate these distributions.

To drive home the point that the concerns raised here are very broad, Table III enumerates the results of applying a great many information measures. It is organized in to four sections: entropies, mutual informations, common informations, and other measures.

None of the entropies, dependent only on the probability mass function of the distribution, can distinguish the two distributions. Nor can any of the mutual informations, as they are functions of the information atoms in the I-diagrams of Fig. 3.

The common informations, defined via auxiliary variables satisfying particular properties, can potentially isolate differences in the dependencies. Though only one of them—the Gács-Körner common information $K[\bullet]$ [57, 58], involving the construction of the largest “subrandom variable” common to the variables and highlighted in the table—discerns that the two distributions are not equivalent because the triadic distribution contains the subrandom variable $X_1 = Y_1 = Z_1$ common to all three variables.

Finally, only two of the other measures (also highlighted) identify any difference between the two. Some fail because they are functions of the probability mass function. Others, like the TSE complexity [59] and erasure entropy [60], fail since they are functions of the I-diagram atoms. Only the intrinsic mutual information $I[\bullet \downarrow \bullet]$ [26] and the reduced mutual information $I[\bullet \Downarrow \bullet]$ [27] distinguish the two since the dyadic distribution contains three dyadic subvariables each of which is independent of the third variable, whereas in the triadic distribution the conditional dependence of the XOR relation can be destroyed.

Figure 4 demonstrates three different information expansions—that, roughly speaking, group variables into subsets of difference sizes or “scales”—applied to our distributions of interest. The first is the complexity profile [55]. At scale k , the complexity profile is the sum of all I-diagram atoms consisting of at least k variables conditioned on the others. Here, since the I-diagrams are identical so are the complexity profiles. The second profile is the marginal utility of information [56], which is a derivative of a linear programming problem whose constraints are given by the I-diagram so here, again,

TABLE III. Suite of information measures applied to the dyadic and triadic distributions, where: $H[\bullet]$ is the Shannon entropy [32], $H_\alpha[\bullet]$ is the Rényi entropy [61], $S_q[\bullet]$ is the Tsallis entropy [62], $I[\bullet]$ is the co-information [44], $T[\bullet]$ is the total correlation [47], $B[\bullet]$ is the dual total correlation [48, 63], $J[\bullet]$ is the CAEKL mutual information [49], $II[\bullet]$ is the interaction information [64], $K[\bullet]$ is the Gács-Körner common information [57], $C[\bullet]$ is the Wyner common information [65, 66], $G[\bullet]$ is the exact common information [67], $F[\bullet]$ is the functional common information^a, $M[\bullet]$ is the MSS common information^b, $I[\bullet \downarrow \bullet]$ is the intrinsic mutual information [26]^c, $I[\bullet \Downarrow \bullet]$ is the reduced intrinsic mutual information [27]^{cd}, $X[\bullet]$ is the extropy [68], $R[\bullet]$ is the residual entropy or erasure entropy [60, 63], $P[\bullet]$ is the perplexity [69], $D[\bullet]$ is the disequilibrium [51], $C_{\text{LMRP}}[\bullet]$ is the LMRP complexity [51], and TSE $[\bullet]$ is the TSE complexity [59]. Only the Gács-Körner common information and the intrinsic mutual informations, highlighted, are able to distinguish the two distributions; the Gács-Körner common information via the construction of a subvariable ($X_1 = Y_1 = Z_1$) common to X , Y , and Z , and the intrinsic mutual informations via the relationship $X_0 = Y_1$ being independent of Z .

Measures	Dyadic	Triadic
$H[X, Y, Z]$	3 bit	3 bit
$H_2[X, Y, Z]$	3 bit	3 bit
$S_2[X, Y, Z]$	0.875 bit	0.875 bit
$I[X : Y : Z]$	0 bit	0 bit
$T[X : Y : Z]$	3 bit	3 bit
$B[X : Y : Z]$	3 bit	3 bit
$J[X : Y : Z]$	1.5 bit	1.5 bit
$II[X : Y : Z]$	0 bit	0 bit
$K[X : Y : Z]$	0 bit	1 bit
$C[X : Y : Z]$	3 bit	3 bit
$G[X : Y : Z]$	3 bit	3 bit
$F[X : Y : Z]^a$	3 bit	3 bit
$M[X : Y : Z]^b$	3 bit	3 bit
$I[X : Y \downarrow Z]^c$	1 bit	0 bit
$I[X : Y \Downarrow Z]^cd$	1 bit	0 bit
$X[X, Y, Z]$	1.349 bit	1.349 bit
$R[X : Y : Z]$	0 bit	0 bit
$P[X, Y, Z]$	8	8
$D[X, Y, Z]$	0.761 bit	0.761 bit
$C_{\text{LMRP}}[X, Y, Z]$	0.381 bit	0.381 bit
$\text{TSE}[X : Y : Z]$	2 bit	2 bit

^a $F[\{X_i\}] = \min_{\substack{\perp X_i|V \\ V=f(\{X_i\})}} H[V]$, where $\perp X_i|V$ means that the X_i are conditionally independent given V .

^b $M[\{X_i\}] = H[\bigvee (X_i \searrow X_j)]$, where $X \searrow Y$ is the minimal sufficient statistic [32] of X about Y and \bigvee denotes the informational union of variables.

^c Though this measure is generically dependent on which variable(s) are chosen to be conditioned on, due to the symmetry of the dyadic and triadic distributions the values reported here are insensitive to permutations of the variables.

^d The original work [27] used the slightly more verbose notation $I[\bullet \Downarrow \bullet]$.

they are identical. Finally, we have Schneidman *et al.*'s connected informations [70], which are the differences in entropies of the maximum entropy distributions whose k - and $k - 1$ -way marginals are fixed to match those of the distribution of interest. Here, all dependencies are detected once pairwise marginals are fixed in the dyadic distribution, but it takes the full joint distribution to realize the XOR subdependency in the triadic distribution. Neither the transfer entropy [20], the transinformation [53], the directed information [52], the causation entropy [22], nor any of their generalizations based on conditional mutual information are capable of determining that the dependency structure, and therefore the causal structure, of the two distributions are qualitatively different. This defect fundamentally precludes them from determining causal structure within a system of unknown dependencies. It underlies our prior criticism of these functions as measures of information flow [28].³

A promising approach to understanding informational dependencies is the *partial information decomposition* (PID) [29]. This framework seeks to decompose a mutual information of the form $I[(I_0, I_1) : O]$ into four nonnegative components: the information R that both inputs I_0 and I_1 *redundantly* provide the output O , the information U_0 that I_0 *uniquely* provides O , the information U_1 that I_1 *uniquely* provides O , and finally the information S that both I_0 and I_1 *synergistically* or *collectively* provide O .

Under this decomposition, our two distributions take on very different characteristics.⁴ For both, the decomposition is invariant as far as which variables are selected as I_0 , I_1 , and O . For the dyadic distribution PID identifies both bits in $I[(I_0, I_1) : O]$ as unique, one from each input I_i , corresponding to the dyadic subdependency shared by I_i and O . Orthogonally, for the triadic distribution PID identifies one of the bits as redundant, stemming from $X_1 = Y_1 = Z_1$, and the other as synergistic, resulting from the XOR relation among X_0 , Y_0 and Z_0 . These decompositions are displayed pictorially in Fig. 5.




FIG. 4. Suite of information expansions applied to the dyadic and triadic distributions: the complexity profile [55], the marginal utility of information [56], and the connected information [70]. The complexity profile and marginal utility of information profiles are identical for the two distributions as a consequence of the information diagrams (Fig. 3) being identical. The connected informations, quantifying the amount of dependence realized by fixing k -way marginals, is able to distinguish the two distributions. Note that although all the x-axes are each scale, exactly what that means depends on the measure.

³ As discussed there, the failure of these measures stems from the possibility of conditional dependence, whereas the aim for these directed measures is to quantify the information flow from one time series to another *excluding* influences of the second. In this light, $T'_{X \rightarrow Y} = I[X_0^t : Y_t \downarrow Y_0^t]$ [26] is certainly an incremental improvement over the transfer entropy.

⁴ Here, we quantified the partial information lattice using the now-commonly accepted technique of Ref. [71], though calculations using two other techniques [72, 73] match. The original PID measure I_{\min} , however, assigns both distributions 1 bit of redundant information and 1 bit of synergistic information. We have yet to compute two other recent proposals [74, 75], though we suspect they will match Ref. [71]'s values.

Another somewhat similar approach is that of *integrated information theory* [76]. However, this approach requires a known dynamic over the variables and is, in addition, highly sensitive to the dynamic. Here, in contrast, we considered only simple probability distributions without any assumptions as to how they might arise from the dynamics of interacting agents. That said, one might associate an integrated information measure with a distribution via the maximal information integration over all possible dynamics that give rise to the distribution. We leave this task for a later study.




FIG. 5. Partial information decomposition diagrams for the (a) dyadic and (b) triadic distributions. Here, X and Y are treated as inputs and Z as output, but in both cases the decomposition is invariant to permutations of the variables. In the dyadic case, the relationship is realized as 1 bit of unique information from X to Z and 1 bit of unique information from Y to Z . In the triadic case, the relationship is quantified as X and Y providing 1 bit of redundant information about Z while also supplying 1 bit of information synergistically about Z .

DISCUSSION

The broad failure of Shannon information measures to differentiate the dyadic and polyadic distributions has far-reaching consequences. Consider, for example, an experiment where a practitioner places three probes into a cluster of neurons, each probe touching two neurons and reporting 0 when they are both quiescent, 1 when the first is excited but second quiescent, 2 when the second is excited but the first quiescent, and 3 when both are excited. Shannon-like measures—including the transfer entropy—would be unable to differentiate between the dyadic situation consisting of three pairs of synchronized neurons, the triadic situation consisting of a trio of synchronized neurons, and a trio exhibiting the XOR relation—a relation requiring nontrivial sensory integration. Such a situation might arise when probing the circuitry of the

drosophila melanogaster connectome [77], for instance.

Furthermore, while partitioning each variable into sub-variables made the dependency structure clear, we do not believe that such a refinement should be a necessary step in discovering such structure. Consider that we demonstrated that refinement is not strictly needed, since the partial information decomposition was able to discover the distribution’s internal structure without it.

These results, observations, and the broad survey clearly highlight the need to extend Shannon’s theory. In particular, the extension must introduce a fundamentally new measure, not merely sums and differences of the standard Shannon information measures. While the partial information decomposition was initially proposed to overcome the interpretational difficulty of the (potentially negative valued) co-information, we see here that it actually overcomes a vastly more fundamental weakness with Shannon information measures. While negative information atoms can *subjectively* be seen as a flaw, the inability to distinguish dyadic from polyadic relations is a much deeper and objective issue.

This may lead one to consider the partial information decomposition as the needed extension to Shannon theory. We do not. The partial information decomposition depends on interpreting some random variables as “inputs” and others as “outputs”. While this may be perfectly natural in some contexts, it is not satisfactory in general. Consider: how should the triadic distribution’s information be allotted? Certainly one of its three bits is redundant and one is synergistic, but what about the third? The XOR (dyadic) distribution contains two bits—are both to be considered synergy [78]? In any case, the partial information framework does not address this question.

DYADIC CAMOUFLAGE & DEPENDENCY DIFFUSION

The dyadic and triadic distributions we analyzed thus far were deliberately chosen to have small dimensionality in an effort to make them and the failure of Shannon information measures as comprehensible and intuitive as possible. Since a given data set may have exponentially many different three-variable subsets, even just this pair of trivariate distributions will stymie most any assessment of variable dependency. However, this is simply a starting point. We will now demonstrate that there exist distributions of arbitrary size which mask their k -way dependencies in such a way that they match, from




FIG. 6. Dyadic camouflage distribution: This distribution, when uniformly and independently mixed with the 4-variable parity distribution (in which each variable is the parity of the other three), results in a distribution whose I-diagram incorrectly implies that the distribution contains only dyadic dependencies.

a Shannon information theory perspective, a distribution of the same size containing only dyadic relationships. Furthermore, we show how any such distribution may be obfuscated over any larger set of variables. This likely mandates a search over all partitions of all subsets of a system, making the problem of finding such distributions in the EXPTIME computational complexity class [79].

Specifically, consider the 4-variable parity distribution consisting of four binary variables such that each variable’s value is equal to the parity of the remaining three. This is a straightforward generalization of the XOR distribution used in constructing the triadic distribution. We next need a generalization of the “giant bit” [63]—which we call *dyadic camouflage*—to mix with the parity, informationally “canceling out” the higher-order mutual informations even though dependencies of such orders exist in the distribution. An example dyadic camouflage distribution for four variables is given in Fig. 6.

Generically, an n -variable dyadic camouflage distribution has an alphabet size of 2^{n-2} and consists of $2^{\frac{(n-2)(n-1)}{2}}$ equally likely outcomes, both numbers determined due to entropy considerations. The distribution is constrained such that any two variables are completely determined by the remaining $n - 2$. Moreover, each m -variable sub-distribution has equal entropy and, otherwise, is of maximal joint entropy. One method of constructing such a distribution is to begin by writing down one variable in increasing lexicographic order such that it has the correct number of outcomes; *e.g.*, column W in Fig. 6a. For each remaining variable, uniformly apply permutations of increasing size to its set of values.

Finally, one can obfuscate any distribution by embedding it in a larger collection of random variables. Given

a distribution D over n variables, associate each random variable i of D with a k -variable subset of a distribution D' in such a way that there is a mapping from the k outcomes in the subset of D' to the outcome of the variable i in D . For example, one can embed the XOR distribution over X, Y, Z into six variables $X_0, X_1, Y_0, Y_1, Z_0, Z_1$ via $X_0 \oplus X_1 = X$, $Y_0 \oplus Y_1 = Y$, and $Z_0 \oplus Z_1 = Z$. In other words, the parity of (Z_0, Z_1) is equal to the XOR of the parities of (X_0, X_1) and (Y_0, Y_1) . In this way one must potentially search over all partitions of all subsets of D' in order to discover the distribution D hiding within. We refer to this method of obfuscation as *dependency diffusion*.

The first conclusion is that the challenges of conditional dependence can be found in joint distributions over arbitrarily large sets of random variables. The second conclusion, one that heightens the challenge to discovery, is that even finding which variables are implicated in polyadic dependencies can be exponentially difficult. Together the camouflage and diffusion constructions demonstrate how challenging it is to discover, let alone work with, multivariate dependencies. This difficulty strongly implies that the current state of information-theoretic tools is vastly underpowered for the types of analyses required of our modern, data-rich sciences.

It is unlikely that the parity plus dyadic camouflage distribution discussed here is the only example of Shannon measures conflating the arity of dependencies and thus producing an information diagram identical to that of a qualitatively distinct distribution. This suggests an important challenge: find additional, perhaps simpler, joint distributions exhibiting this phenomenon.

CONCLUSION

To conclude, we constructed two distributions that cannot be distinguished using conventional (and many non-conventional) Shannon-like information measures. In fact, of the more than two dozen measures we surveyed only five were able to separate the distributions: the Gács-Körner common information, the intrinsic mutual information, the reduced intrinsic mutual information, the connected informations, and the partial information decomposition. We also noted in an aside that causality detection approaches that assume an underlying directed acyclic graph structure are structurally impotent.

The failure of the Shannon-type measures is perhaps not surprising: nothing in the standard mathematical theories of information and communication suggests that

such measures *should* be able to distinguish these distributions [80]. However, distinguishing dyadic from triadic relationships and the related causal structure is of the utmost importance to the sciences. Critically, since interpreting dependencies in random distributions is traditionally the domain of information theory, we propose that new extensions to information theory are needed.

Furthermore, the dyadic camouflage distribution (Fig. 6) presents an acute challenge for traditional methods of dependency and causality inference. Let's close with an example. Consider the widely used Granger causality [81] applied to the camouflage distribution. Fixing any two variables, say X and Y , determines the remaining two, Z and W . What is one to conclude from this, other than that X cannot influence Y ? And yet, in conjunction with either Z or W , X completely determines Y . This makes clear the deep assumption of dyadic relationships that permeates and biases our ways of thinking about complex systems.

These results may seem like a deal-breaking criticism of employing information theory to determine dependencies. Indeed, these results seem to indicate that much existing empirical work and many interpretations have simply been wrong and, worse even, that the associated methods are misleading while appearing quantitatively consistent. We think not, though. With the constructive and detailed problem diagnosis given here, at least we can see the true problem. It is now a necessary step to address it. This leads us to close with a cautionary quote:

The tools we use have a profound (and devious!) influence on our thinking habits, and, therefore, on our thinking abilities.

Edsger W. Dijkstra [82]

ACKNOWLEDGMENTS

We thank N. Barnett and D. P. Varn for helpful feedback. JPC thanks the Santa Fe Institute for its hospitality during visits as an External Faculty member. This material is based upon work supported by, or in part by, the U. S. Army Research Laboratory and the U. S. Army Research Office under contracts W911NF-13-1-0390 and W911NF-13-1-0340.

A Python Discrete Information Package

Hand calculating the information quantities used in the main text, while profitably done for a few basic examples, soon becomes tedious and error prone. We provide a Jupyter notebook [83] making use of `dit` (“Discrete Information Theory”) [84], an open source Python package that readily calculates these quantities.

* rgjames@ucdavis.edu

† chaos@ucdavis.edu

- [1] S. Kullback. *Information Theory and Statistics*. Dover, New York, 1968. 1
- [2] H. Quastler. *Information Theory in Biology*. University of Illinois Press, Urbana-Champaign, Illinois, 1953. 1
- [3] H. Quastler. The status of information theory in biology—A roundtable discussion. In H. P. Yockey, editor, *Symposium on Information Theory in Biology*, page 399. Pergamon Press, New York, 1958.
- [4] J. Kelly. A new interpretation of information rate. *IRE Trans. Info. Th.*, 2(3):185–189, 1956.
- [5] L. Brillouin. *Science and Information Theory*. Academic Press, New York, second edition, 1962.
- [6] W. Bialek, F. Rieke, R. R. De Ruyter Van Steveninck, and D. Warland. Reading a neural code. *Science*, 252(5014):1854–1857, 1991.
- [7] S. P. Strong, R. Koberle, R. R. de Ruyter van Steveninck, and W. Bialek. Entropy and information in neural spike trains. *Phys. Rev. Lett.*, 80(1):197, 1998.
- [8] R. E. Ulanowicz. The central role of information theory in ecology. In M. Dehmer, A. Mehler, and F. Emmert-Streib, editors, *Towards an Information Theory of Complex Networks*, pages 153–167. Springer, 2011.
- [9] W. T. Grandy Jr. *Entropy and the Time Evolution of Macroscopic Systems*, volume 141. Oxford University Press, Oxford, United Kingdom, 2008.
- [10] J. Harte. *Maximum Entropy and Ecology: a Theory of Abundance, Distribution, and Energetics*. Oxford University Press, 2011.
- [11] R. F. Nalewajski. *Information Theory of Molecular Systems*. Elsevier, 2006.
- [12] J. Garland, R. G. James, and E. Bradley. Model-free quantification of time-series predictability. *Phys. Rev. E*, 90(5):052910, 2014.
- [13] O. Kafri. Information theoretic approach to social networks. *arXiv:1407.7964*, 2014.
- [14] D. P. Varn and J. P. Crutchfield. Chaotic crystallography: How the physics of information reveals structural order in materials. *Current Opinion in Chemical Engineering*, 7:47–56, 2015.
- [15] D. P. Varn and J. P. Crutchfield. What did Erwin mean? The physics of information from the materials genomics of aperiodic crystals and water to molecular information catalysts and life. *Phil. Trans. R. Soc. A*,

- 374(2063):20150067, 2016.
- [16] X.-Y. Zhou, C. Rong, T. Lu, P. Zhou, and S. Liu. Information functional theory: Electronic properties as functionals of information for atoms and molecules. *J. Phys. Chem. A*, 120(20):3634–3642, 2016.
- [17] C. Kirst, M. Timme, and D. Battaglia. Dynamic information routing in complex networks. *Nature Comm.*, 7, 2016.
- [18] E. J. Izquierdo, P. L. Williams, and R. D. Beer. Information flow through a model of the *C. elegans* klinotaxis circuit. *PLoS One*, 10(10):e0140397, 2015.
- [19] R. G. James, K. Burke, and J. P. Crutchfield. Chaos forgets and remembers: Measuring information creation, destruction, and storage. *Phys. Lett. A*, 378(30):2124–2127, 2014. 1
- [20] T. Schreiber. Measuring information transfer. *Phys. Rev. Lett.*, 85(2):461, 2000. 1, 4, 5
- [21] P. Fiedor. Partial mutual information analysis of financial networks. *arXiv:1403.2050*, 2014.
- [22] J. Sun and E. M. Bollt. Causation entropy identifies indirect influences, dominance of neighbors and anticipatory couplings. *Physica D: Nonlinear Phenomena*, 267:49–57, 2014. 4, 5
- [23] J. T. Lizier, M. Prokopenko, and A. Y. Zomaya. Local information transfer as a spatiotemporal filter for complex systems. *Phys. Rev. E*, 77(2):026110, 2008.
- [24] S. I. Walker, H. Kim, and P. C. W. Davies. The informational architecture of the cell. *arXiv:1507.03877*, 2015.
- [25] U. Lee, S. Blain-Moraes, and G. A. Mashour. Assessing levels of consciousness with symbolic analysis. *Phil. Trans. Roy. Soc. Lond. A*, 373(2034):20140117, 2015. 1
- [26] U. Maurer and S. Wolf. The intrinsic conditional mutual information and perfect secrecy. In *IEEE Intl. Symp. Info. Th.*, pages 88–88, 1997. 1, 4, 5
- [27] R. Renner, J. Skripsky, and S. Wolf. A new measure for conditional mutual information and its properties. In *IEEE Intl. Symp. Info. Th.*, pages 259–259, 2003. 1, 4
- [28] R. G. James, N. Barnett, and J. P. Crutchfield. Information flows? A critique of transfer entropies. *Phys. Rev. Lett.*, 116:238701, Jun 2016. 1, 5
- [29] P. L. Williams and R. D. Beer. Nonnegative decomposition of multivariate information. *arXiv:1004.2515*, 2010. 1, 5
- [30] N. Bertschinger, J. Rauh, E. Olbrich, and J. Jost. Shared information: New insights and problems in decomposing information in complex systems. In *Proceedings of the European Conference on Complex Systems 2012*, pages 251–269. Springer, 2013. 1
- [31] J. T. Lizier. *The Local Information Dynamics of Distributed Computation in Complex Systems*. PhD thesis, Citeseer, 2010. 1
- [32] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, New York, 2012. 1, 4
- [33] R. W. Yeung. *A First Course in Information Theory*. Springer Science & Business Media, 2012.
- [34] I. Csizsar and J. Körner. *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Cambridge University Press, 2011.
- [35] D. J. C. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003. 1
- [36] V. Griffith and C. Koch. Quantifying synergistic mutual information. In *Guided Self-Organization: Inception*, pages 159–190. Springer, 2014. 2
- [37] M. Cook. *Networks of Relations*. PhD thesis, California Institute of Technology, 2005. 2
- [38] F. V. Jensen. *An introduction to Bayesian networks*, volume 210. UCL Press London, 1996. 2
- [39] Marco S. Learning Bayesian networks with the bnlearn R package. *J. Stat. Software*, 35(3):1–22, 2010. 2
- [40] J. Pearl. *Causality*. Cambridge University Press, 2009. 2
- [41] N. Ay and D. Polani. Information flows in causal networks. *Adv. Complex Sys.*, 11(01):17–41, 2008. 2
- [42] F. M. Reza. *An Introduction to Information Theory*. Courier Corporation, 1961. 3
- [43] R. W. Yeung. A new outlook on Shannon’s information measures. *IEEE Trans. Info. Th.*, 37(3):466–474, 1991. 3
- [44] A. J. Bell. The co-information lattice. In S. Makino S. Amari, A. Cichocki and N. Murata, editors, *Proc. Fifth Intl. Workshop on Independent Component Analysis and Blind Signal Separation*, volume ICA 2003, pages 921–926, New York, 2003. Springer. 3, 4
- [45] L. M. A. Bettencourt, G. J. Stephens, M. I. Ham, and G. W. Gross. Functional structure of cortical neuronal networks grown in vitro. *Phys. Rev. E*, 75(2):021915, 2007. 3
- [46] K. Krippendorff. Information of interactions in complex systems. *Intl. J. General Systems*, 38(6):669–680, 2009. 3
- [47] S. Watanabe. Information theoretical analysis of multivariate correlation. *IBM J. Res. Dev.*, 4(1):66–82, 1960. 4
- [48] H. Te Sun. Multiple mutual informations and multiple interactions in frequency data. *Info. Control*, 46:26–45, 1980. 4
- [49] C. Chan, A. Al-Bashabsheh, J. B. Ebrahimi, T. Kaced, and T. Liu. Multivariate mutual information inspired by secret-key agreement. *Proc. IEEE*, 103(10):1883–1913, 2015. 4
- [50] R. G. James, C. J. Ellison, and J. P. Crutchfield. Anatomy of a bit: Information in a time series observation. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 21(3):037109, 2011.
- [51] P. W. Lamberti, M. T. Martin, A. Plastino, and O. A. Rosso. Intensive entropic non-triviality measure. *Physica A*, 334(1):119–131, 2004. 4
- [52] J. Massey. Causality, feedback and directed information. In *Proc. Int. Symp. Inf. Theory Applic.*, volume ISITA-90, pages 303–305, 1990. 5
- [53] H. Marko. The bidirectional communication theory: A generalization of information theory. *IEEE Trans. Commun.*, 21(12):1345–1351, 1973. 4, 5
- [54] L. M. A. Bettencourt, V. Gintautas, and M. I. Ham. Identification of functional information subgraphs in complex networks. *Phys. Rev. Lett.*, 100(23):238701, 2008. 4

- [55] Y. Bar-Yam. Multiscale complexity/entropy. *Adv. Complex Sys.*, 7(01):47–63, 2004. 4, 5
- [56] B. Allen, B. C. Stacey, and Y. Bar-Yam. An information-theoretic formalism for multiscale structure in complex systems. *arXiv:1409.4708*, 2014. 4, 5
- [57] P. Gács and J. Körner. Common information is far less than mutual information. *Problems of Control and Information Theory*, 2(2):149–162, 1973. 4
- [58] H. Tyagi, P. Narayan, and P. Gupta. When is a function securely computable? *IEEE Trans. Info. Th.*, 57(10):6337–6350, 2011. 4
- [59] N. Ay, E. Olbrich, N. Bertschinger, and J. Jost. A unifying framework for complexity measures of finite systems. *Proc. European Conf. Complex Sys.*, 6, 2006. 4
- [60] S. Verdu and T. Weissman. The information lost in erasures. *IEEE Trans. Info. Th.*, 54(11):5030–5058, 2008. 4
- [61] A. Rényi. On measures of entropy and information. In *Proc. Fourth Berkeley Symp. Math. Stat. Prob.*, pages 547–561, 1960. 4
- [62] C. Tsallis. Possible generalization of Boltzmann-Gibbs statistics. *J. Stat. Phys.*, 52(1-2):479–487, 1988. 4
- [63] S. A. Abdallah and M. D. Plumbley. A measure of statistical complexity based on predictive information with application to finite spin systems. *Phys. Let. A*, 376(4):275–281, 2012. 4, 7
- [64] W. J. McGill. Multivariate information transmission. *Psychometrika*, 19(2):97–116, 1954. 4
- [65] A. D. Wyner. The common information of two dependent random variables. *IEEE Trans. Info. Th.*, 21(2):163–179, 1975. 4
- [66] W. Liu, G. Xu, and B. Chen. The common information of n dependent random variables. In *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, pages 836–843. IEEE, 2010. 4
- [67] G. R. Kumar, C. T. Li, and A. El Gamal. Exact common information. In *Proc. IEEE ISIT*, pages 161–165. IEEE, 2014. 4
- [68] F. Lad, G. Sanfilippo, G. Agrò, et al. Extropy: Complementary dual of entropy. *Statistical Science*, 30(1):40–58, 2015. 4
- [69] F. Jelinek, R. L. Mercer, L. R. Bahl, and J. K. Baker. Perplexity—a measure of the difficulty of speech recognition tasks. *J. Acoust. Soc. Am.*, 62(S1):S63–S63, 1977. 4
- [70] E. Schneidman, S. Still, M. J. Berry, and W. Bialek. Network information and connected correlations. *Phys. Rev. Lett.*, 91(23):238701, 2003. 5
- [71] N. Bertschinger, J. Rauh, E. Olbrich, J. Jost, and N. Ay. Quantifying unique information. *Entropy*, 16(4):2161–2183, 2014. 5
- [72] M. Harder, C. Salge, and D. Polani. Bivariate measure of redundant information. *Phys. Rev. E*, 87(1):012130, 2013. 5
- [73] V. Griffith, E. K.P. Chong, R. G. James, C. J. Ellison, and J. P. Crutchfield. Intersection information based on common randomness. *Entropy*, 16(4):1985–2000, 2014. 5
- [74] R. Quax, O. Har-Shemesh, S. Thurner, and P. Sloot. Stripping syntax from complexity: An information-theoretical perspective on complex systems. *arXiv:1603.03552*, 2016. 5
- [75] R. A. A. Ince. Measuring multivariate redundant information with pointwise common change in surprisal. *arXiv:1602.05063*, 2016. 5
- [76] L. Albantakis, M. Oizumi, and G. Tononi. From the phenomenology to the mechanisms of consciousness: Integrated information theory 3.0. *PLoS Comp. Biol.*, 10(5):e1003588, 05 2014. 5
- [77] S. Takemura, A. Bharioke, Z. Lu, A. Nern, S. Vitaladevuni, P. K. Rivlin, W. T. Katz, D. J. Olbris, S. M. Plaza, P. Winston, T. Zhao, J. A. Horne, R. D. Fetter, K. B. Takemura, L. Chang, O. Ogundeyi, M. A. Saunders, V. Shaprio, C. Sigmund, G. M. Rubin, L. K. Scheffer, I. A. Meinertzhagen, and D. B. Chklovskii. A visual motion detection circuit suggested by drosophila connectomics. *Nature*, 500(7461):175–181, 2013. 6
- [78] F. Rosas, V. Ntranos, C. J. Ellison, S. Pollin, and M. Verhelst. Understanding interdependency through complex information sharing. *Entropy*, 18(2):38, 2016. 6
- [79] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, New York, 1979. 7
- [80] C. E. Shannon. The bandwagon. *IEEE Tran. Info. Th.*, 2(3):3, 1956. 8
- [81] C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, pages 424–438, 1969. 8
- [82] E. W. Dijkstra. How do we tell truths that might hurt? In *Selected Writings on Computing: A Personal Perspective*, pages 129–131. Springer, 1982. 8
- [83] Jupyter. <https://github.com/jupyter/notebook>, 2016. 8
- [84] R. G. James, C. J. Ellison, and J. P. Crutchfield. dit: Discrete information theory in python. <https://github.com/dit/dit>, 2016. 8