

 Open access • Book Chapter • DOI:10.1007/978-3-642-53734-9_6

Quantifying Synergistic Mutual Information — [Source link](#)

Virgil Griffith, Christof Koch

Institutions: California Institute of Technology, Allen Institute for Brain Science

Published on: 01 Jan 2014 - arXiv: Information Theory (Springer Berlin Heidelberg)

Topics: Mutual information

Related papers:

- [Nonnegative Decomposition of Multivariate Information](#)
- [Quantifying unique information](#)
- [Bivariate measure of redundant information.](#)
- [A mathematical theory of communication](#)
- [Elements of information theory](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/quantifying-synergistic-mutual-information-2xd4gdpwo9>

Quantifying synergistic mutual information

Virgil Griffith^{1,*} and Christof Koch^{1,2}

¹Computation and Neural Systems, Caltech, Pasadena, CA 91125

²Allen Institute for Brain Science, Seattle, WA 98103

Abstract

Quantifying cooperation or synergy among random variables in predicting a single target random variable is an important problem in many complex systems. We review three prior information-theoretic measures of synergy and introduce a novel synergy measure defined as the difference between the whole and the union of its parts. We apply all four measures against a suite of binary circuits to demonstrate that our measure alone quantifies the intuitive concept of synergy across all examples. We show that for our measure of synergy that independent predictors can have positive redundant information.

1 Introduction

Synergy is a fundamental concept in complex systems that has received much attention in computational biology [1, 2]. Several papers [3–6] have proposed measures for quantifying synergy, but there remains no consensus which measure is most valid.

The concept of synergy spans many fields and theoretically could be applied to any non-subadditive function. But within the confines of Shannon information theory, synergy—or more formally, *synergistic information*—is a property of a set of n random variables $\mathbf{X} = \{X_1, X_2, \dots, X_n\}$ cooperating to predict (reduce the uncertainty of) a single target random variable Y .

One clear application of synergistic information is in computational genetics. It is well understood that most phenotypic traits are influenced not only by single genes but by interactions among genes—for example, human eye-color is cooperatively specified by more than a dozen genes [7]. The magnitude of this “cooperative specification” is the synergistic information between the set of genes \mathbf{X} and a phenotypic trait Y . Another application is neuronal firings where potentially thousands of presynaptic neurons influence the firing rate of a single post-synaptic (target) neuron. Yet another application is discovering the “informationally synergistic modules” within a complex system.

The prior literature [8, 9] has termed several distinct concepts as “synergy”. This paper defines synergy as how much the whole is greater than (the union of) its atomic elements.¹

The prior works on Partial Information Decomposition [6, 12–14] start with properties that a measure of redundant information, I_{\cap} satisfies and builds a measure of synergy from I_{\cap} . Although this paper deals directly with measures of synergy on “easy” examples, we are immensely sympathetic to this approach. Our proposed measure of synergy does give rise to an I_{\cap} measure.

*To whom correspondence should be addressed. Email: virgil@caltech.edu

¹The techniques here are unrelated to the information geometry perspective provided by [10]. The well-known “total correlation” measure [11], does not satisfy the desired properties for a measure of synergy.

The properties our I_{\cup} satisfies are discussed in Appendix C. For pedagogical purposes all examples are *deterministic*, however, these methods equally apply to non-deterministic systems.

1.1 Notation

We use the following notation throughout. Let

n : The number of predictors X_1, X_2, \dots, X_n . $n \geq 2$.

$X_{1\dots n}$: The *joint* random variable (coalition) of all n predictors $X_1 X_2 \dots X_n$.

X_i : The i 'th predictor random variable (r.v.). $1 \leq i \leq n$.

\mathbf{X} : The *set* of all n predictors $\{X_1, X_2, \dots, X_n\}$.

Y : The *target r.v.* to be predicted.

y : A particular state of the target r.v. Y .

All random variables are discrete, all logarithms are \log_2 , and all calculations are in *bits*. Entropy and mutual information are as defined by [15], $H(X) \equiv \sum_{x \in X} \Pr(x) \log \frac{1}{\Pr(x)}$, as well as $I(X:Y) \equiv \sum_{x,y} \Pr(x,y) \log \frac{\Pr(x,y)}{\Pr(x)\Pr(y)}$.

1.2 Understanding PI-diagrams

Partial information diagrams (PI-diagrams), introduced by [6], extend Venn diagrams to properly represent synergy. Their framework has been invaluable to the evolution of our thinking on synergy.

A PI-diagram is composed of nonnegative *partial information regions* (PI-regions). Unlike the standard Venn entropy diagram in which the sum of all regions is the joint entropy $H(X_{1\dots n}, Y)$, in PI-diagrams the sum of all regions (i.e. the space of the PI-diagram) is the mutual information $I(X_{1\dots n}:Y)$. PI-diagrams are immensely helpful in understanding how the mutual information $I(X_{1\dots n}:Y)$ is distributed across the coalitions and singletons of \mathbf{X} .²

How to read PI-diagrams. Each PI-region is uniquely identified by its “set notation” where each element is denoted solely by the predictors’ indices. For example, in the PI-diagram for $n = 2$ (Figure 1a): $\{1\}$ is the information about Y only X_1 carries (likewise $\{2\}$ is the information only X_2 carries); $\{1, 2\}$ is the information about Y that X_1 as well as X_2 carries, while $\{12\}$ is the information about Y that is specified only by the coalition (joint random variable) $X_1 X_2$. For getting used to this way of thinking, common informational quantities are represented by colored regions in Figure 2.

The general structure of a PI-diagram becomes clearer after examining the PI-diagram for $n = 3$ (Figure 1b). All PI-regions from $n = 2$ are again present. Each predictor (X_1, X_2, X_3) can carry unique information (regions labeled $\{1\}, \{2\}, \{3\}$), carry information redundantly with another predictor ($\{1,2\}, \{1,3\}, \{2,3\}$), or specify information through a coalition with another predictor ($\{12\}, \{13\}, \{23\}$). New in $n = 3$ is information carried by all three predictors ($\{1,2,3\}$) as well as information specified through a three-way coalition ($\{123\}$). Intriguingly, for three predictors, information can be provided by a coalition as well as a singleton ($\{1,23\}, \{2,13\}, \{3,12\}$) or specified by multiple coalitions ($\{12,13\}, \{12,23\}, \{13,23\}, \{12,13,23\}$).

2 Information can be redundant, unique, or synergistic

Each PI-region represents an irreducible nonnegative slice of the mutual information $I(X_{1\dots n}:Y)$ that is either:

²Formally, how the mutual information is distributed across the set of all nonempty antichains on the powerset of \mathbf{X} [16, 17].




Figure 1: PI-diagrams for two and three predictors. Each PI-region represents nonnegative information about Y . A PI-region's color represents whether its information is redundant (yellow), unique (magenta), or synergistic (cyan). To preserve symmetry, the PI-region “{12, 13, 23}” is displayed as three separate regions each marked with a “*”. All three *-regions should be treated as though they are a single region.




Figure 2: PI-diagrams for $n = 2$ representing standard informational quantities.

1. **Redundant.** Information carried by a singleton predictor as well as available somewhere else. For $n = 2$: {1,2}. For $n = 3$: {1,2}, {1,3}, {2,3}, {1,2,3}, {1,23}, {2,13}, {3,12}.
2. **Unique.** Information carried by exactly one singleton predictor and is available nowhere else. For $n = 2$: {1}, {2}. For $n = 3$: {1}, {2}, {3}.
3. **Synergistic.** Any and all information in $I(X_{1\dots n}:Y)$ that is not carried by a singleton predictor. $n = 2$: {12}. For $n = 3$: {12}, {13}, {23}, {123}, {12,13}, {12,23}, {13,23}, {12,13,23}.

Although a single PI-region is either redundant, unique, or synergistic, a single state of the target can have any combination of positive PI-regions, i.e. a single state of the target can convey redundant, unique, and synergistic information. This surprising fact is demonstrated in Figure 9.

2.1 Example Rdn: Redundant information

If X_1 and X_2 carry some identical³ information (reduce the same uncertainty) about Y , then we say the set $\mathbf{X} = \{X_1, X_2\}$ has some *redundant information* about Y . Figure 3 illustrates a simple case of redundant information. Y has two equiprobable states: \mathbf{r} and \mathbf{R} (\mathbf{r}/\mathbf{R} for “redundant bit”). Examining X_1 or X_2 identically specifies one bit of Y , thus we say set $\mathbf{X} = \{X_1, X_2\}$ has one bit of redundant information about Y .




Figure 3: Example RDN. Figure 3a shows the joint distribution of r.v.’s X_1 , X_2 , and Y , the joint probability $\Pr(x_1, x_2, y)$ is along the right-hand side of (a), revealing that all three terms are fully correlated. Figure 3b represents the joint distribution as an electrical circuit. Figure 3c is the PI-diagram indicating that set $\{X_1, X_2\}$ has 1 bit of redundant information about Y . $I(X_1 X_2 : Y) = I(X_1 : Y) = I(X_2 : Y) = H(Y) = 1$ bit.

2.2 Example Unq: Unique information

Predictor X_i carries *unique information* about Y if and only if X_i specifies information about Y that is not specified by anything else (a singleton or coalition of the other $n - 1$ predictors). Figure 4 illustrates a simple case of unique information. Y has four equiprobable states: \mathbf{ab} , \mathbf{aB} , \mathbf{Ab} , and \mathbf{AB} . X_1 uniquely specifies bit \mathbf{a}/\mathbf{A} , and X_2 uniquely specifies bit \mathbf{b}/\mathbf{B} . If we had instead labeled the Y -states: 0 , 1 , 2 , and 3 , X_1 and X_2 would still have strictly unique information about Y . The state of X_1 would specify between $\{0, 1\}$ and $\{2, 3\}$, and the state of X_2 would specify between $\{0, 2\}$ and $\{1, 3\}$ —together fully specifying the state of Y . Accepting the property **(Id)** from [12] is sufficient but not necessary for the desired decomposition of example UNQ.




Figure 4: Example UNQ. X_1 and X_2 each uniquely specify a single bit of Y . $I(X_1 X_2 : Y) = H(Y) = 2$ bits. The joint probability $\Pr(x_1, x_2, y)$ is along the right-hand side of (a).

2.3 Example Xor: Synergistic information

A set of predictors $\mathbf{X} = \{X_1, \dots, X_n\}$ has synergistic information about Y if and only if the whole $(X_{1\dots n})$ specifies information about Y that is not specified by any singleton predictor.

³ X_1 and X_2 providing identical information about Y is different from providing the same *magnitude* of information about Y , i.e. $I(X_1 : Y) = I(X_2 : Y)$. Example UNQ (Figure 4) is an example where $I(X_1 : Y) = I(X_2 : Y) = 1$ bit yet X_1 and X_2 specify “different bits” of Y . Providing the same magnitude of information about Y is neither necessary or sufficient for providing some identical information about Y .

The canonical example of synergistic information is the XOR-gate (Figure 5). In this example, the whole X_1X_2 fully specifies Y ,

$$I(X_1X_2:Y) = H(Y) = 1 \text{ bit}, \quad (1)$$

but the singletons X_1 and X_2 specify *nothing* about Y ,

$$I(X_1:Y) = I(X_2:Y) = 0 \text{ bits}. \quad (2)$$

With both X_1 and X_2 themselves having zero information about Y , we know that there can not be any redundant or unique information about Y —that the three PI-regions $\{1\} = \{2\} = \{1,2\} = 0$ bits. As the information between X_1X_2 and Y must come from somewhere, by elimination we conclude that X_1 and X_2 synergistically specify Y .




Figure 5: Example XOR. X_1 and X_2 synergistically specify Y . $I(X_1X_2:Y) = H(Y) = 1$ bit. The joint probability $\Pr(x_1, x_2, y)$ is along the right-hand side of (a).

3 Two examples elucidating properties of synergy

To help the reader develop intuition for a proper measure of synergy we illustrate two desired properties of synergistic information with pedagogical examples derived from XOR. Readers solely interested in the contrast with prior measures can skip to Section 4.

3.1 Duplicating a predictor does not change synergistic information

Example XORDUPLICATE (Figure 6) adds a third predictor, X_3 , a copy of predictor X_1 , to XOR. Whereas in XOR the target Y is specified only by coalition X_1X_2 , duplicating predictor X_1 as X_3 makes the target equally specifiable by coalition X_3X_2 .

Although now two different coalitions identically specify Y , mutual information is invariant to duplicates, e.g. $I(X_1X_2X_3:Y) = I(X_1X_2:Y)$ bit. Likewise for synergistic information to be likewise bounded between zero and the total mutual information $I(X_{1..n}:Y)$, synergistic information must similarly be invariant to duplicates, e.g. the synergistic information between set $\{X_1, X_2\}$ and Y must be the same as the synergistic information between $\{X_1, X_2, X_3\}$ and Y . This makes sense because if synergistic information is defined as the information in the whole beyond its parts, duplicating a part does not increase the net information provided by the parts. Altogether, we assert that *duplicating a predictor does not change the synergistic information*. Synergistic information being invariant to duplicated predictors follows from the equality condition of the monotonicity property (**M**) from [13].⁴

3.2 Adding a new predictor can decrease synergy

Example XORLOSES (Figure 7) adds a third predictor, X_3 , to XOR and concretizes the distinction between synergy and “redundant synergy”. In XORLOSES the target Y has one bit of uncertainty and just as in example XOR the coalition X_1X_2 fully specifies the target, $I(X_1X_2:Y) = H(Y) = 1$ bit. However, XORLOSES has *zero* intuitive synergy because the newly added singleton predictor, X_3 , fully specifies Y by itself. This makes the synergy between X_1 and X_2 *completely redundant*—everything the coalition X_1X_2 specifies is now already specified by the singleton X_3 .

⁴For a proof see Appendix E.




Figure 6: Example XORDUPLICATE shows that duplicating predictor X_1 as X_3 turns the single-coalition synergy $\{12\}$ into the multi-coalition synergy $\{12, 23\}$. After duplicating X_1 , the coalition X_3X_2 as well as coalition X_1X_2 specifies Y . Synergistic information is unchanged from XOR, $I(X_3X_2:Y) = I(X_1X_2:Y) = H(Y) = 1$ bit.

4 Prior measures of synergy

4.1 I_{\max} synergy: $\mathcal{S}_{\max}(\mathbf{X} : Y)$

I_{\max} synergy, denoted \mathcal{S}_{\max} , derives from [6]. \mathcal{S}_{\max} defines synergy as the whole beyond the state-dependent *maximum* of its parts,

$$\mathcal{S}_{\max}(\mathbf{X} : Y) \equiv I(X_{1\dots n} : Y) - I_{\max}(\{X_1, \dots, X_n\} : Y) \quad (3)$$

$$= I(X_{1\dots n} : Y) - \sum_{y \in Y} \Pr(Y = y) \max_i I(X_i : Y = y), \quad (4)$$

where $I(X_i : Y = y)$ is [18]’s “specific-surprise”,

$$I(X_i : Y = y) \equiv D_{\text{KL}}\left[\Pr(X_i|y) \parallel \Pr(X_i)\right] \quad (5)$$

$$= \sum_{x_i \in X_i} \Pr(x_i|y) \log \frac{\Pr(x_i, y)}{\Pr(x_i) \Pr(y)}. \quad (6)$$

There are two major advantages of \mathcal{S}_{\max} synergy. First, \mathcal{S}_{\max} obeys the bounds of $0 \leq \mathcal{S}_{\max}(X_{1\dots n} : Y) \leq I(X_{1\dots n} : Y)$. Second, \mathcal{S}_{\max} is invariant to duplicate predictors. Despite these desired properties, \mathcal{S}_{\max} sometimes miscategorizes merely unique information as synergistic. This can be seen in example UNQ (Figure 4). In example UNQ the wires




Figure 7: Example XORLOSES. Target Y is fully specified by the coalition X_1X_2 as well as by the singleton X_3 . $I(X_1X_2:Y) = I(X_3:Y) = H(Y) = 1$ bit. Therefore the information synergistically specified by coalition X_1X_2 is a redundant synergy.

in Figure 4b don't even touch, yet \mathcal{S}_{\max} asserts there is one bit of synergy and one bit of redundancy—this is palpably strange.

A more abstract way to understand why \mathcal{S}_{\max} overestimates synergy is to imagine a hypothetical example where there are exactly two bits of unique information for every state $y \in Y$ and no synergy or redundancy. \mathcal{S}_{\max} would be the whole (both unique bits) minus the *maximum* over both predictors—which would be the $\max[1, 1] = 1$ bit. The \mathcal{S}_{\max} synergy would then be $2 - 1 = 1$ bit of synergy—even though by definition there was no synergy, but merely two bits of unique information.

Altogether, we conclude that \mathcal{S}_{\max} *overestimates* the intuitive synergy by miscategorizing merely unique information as synergistic whenever two or more predictors have unique information about the target.

4.2 WholeMinusSum synergy: $WMS(X : Y)$

The earliest known sightings of bivariate WholeMinusSum synergy (WMS) is [19, 20] with the general case in [21]. WholeMinusSum synergy is a signed measure where a positive value signifies synergy and a negative value signifies redundancy. WholeMinusSum synergy is defined by eq. (7) and interestingly reduces to eq. (9)—the difference of two *total correlations*.⁵

⁵ $TC(X_1; \dots; X_n) = -H(X_{1\dots n}) + \sum_{i=1}^n H(X_i)$ per [11].

$$\text{WMS}(\mathbf{X} : Y) \equiv I(X_{1\dots n} : Y) - \sum_{i=1}^n I(X_i : Y) \quad (7)$$

$$= \sum_{i=1}^n H(X_i|Y) - H(X_{1\dots n}|Y) - \left[\sum_{i=1}^n H(X_i) - H(X_{1\dots n}) \right] \quad (8)$$

$$= \text{TC}(X_1; \dots; X_n|Y) - \text{TC}(X_1; \dots; X_n) \quad (9)$$

Representing eq. (7) for $n = 2$ as a PI-diagram (Figure 8a) reveals that WMS is the synergy between X_1 and X_2 *minus* their redundancy. Thus, when there is an equal magnitude of synergy and redundancy between X_1 and X_2 (as in RDNXOR, Figure 9), WholeMinusSum synergy is *zero*—leading one to *erroneously* conclude there is no synergy or redundancy present.⁶

The PI-diagram for $n = 3$ (Figure 8b) reveals that WholeMinusSum double-subtracts PI-regions $\{1,2\}$, $\{1,3\}$, $\{2,3\}$ and triple-subtracts PI-region $\{1,2,3\}$, revealing that for $n > 2$ $\text{WMS}(\mathbf{X} : Y)$ becomes synergy minus the redundancy *counted multiple times*.

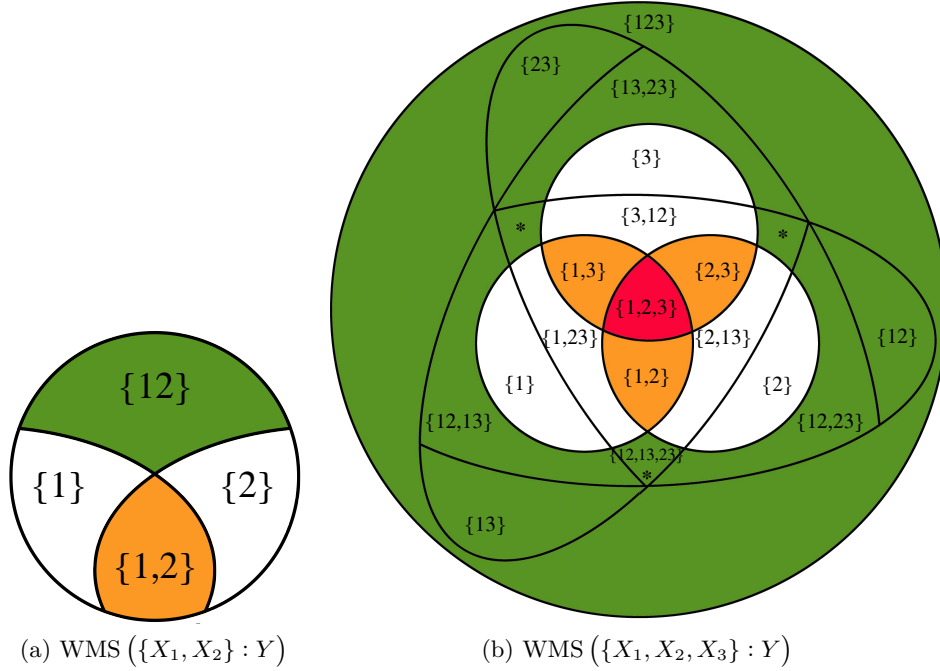


Figure 8: PI-diagrams illustrating WholeMinusSum synergy for $n = 2$ (left) and $n = 3$ (right). For this diagram the colors denote the added and subtracted PI-regions. $\text{WMS}(\mathbf{X} : Y)$ is the green PI-region(s), minus the orange PI-region(s), minus two times any red PI-region.

A concrete example demonstrating WholeMinusSum’s “synergy minus redundancy” behavior is RDNXOR (Figure 9) which overlays examples RDN and XOR to form a single system. The target Y has two bits of uncertainty, i.e. $H(Y) = 2$. Like RDN, either X_1 or X_2 identically specifies the letter of Y (r/R), making one bit of redundant information. Like XOR, only the coalition X_1X_2 specifies the digit of Y (0/1), making one bit of synergistic information. Together this makes one bit of redundancy and one bit of synergy.

⁶This is deeper than [3]’s point that a mish-mash of synergy and redundancy across different states of $y \in Y$ can average to zero. Figure 9 evaluates to zero for *every state* $y \in Y$.

Note that in RDNXOR every state $y \in Y$ conveys one bit of redundant information and one bit of synergistic information, e.g. for the state $y = r0$ the letter “r” is specified redundantly and the digit “0” is specified synergistically. Example RDNUNQXOR (Appendix A) extends RDNXOR to demonstrate redundant, unique, and synergistic information for every state $y \in Y$.

In summary, WholeMinusSum *underestimates* synergy for all n with the potential gap increasing with n . Equivalently, we say that WholeMinusSum synergy is a *lowerbound* on the intuitive synergy with the bound becoming looser with n .

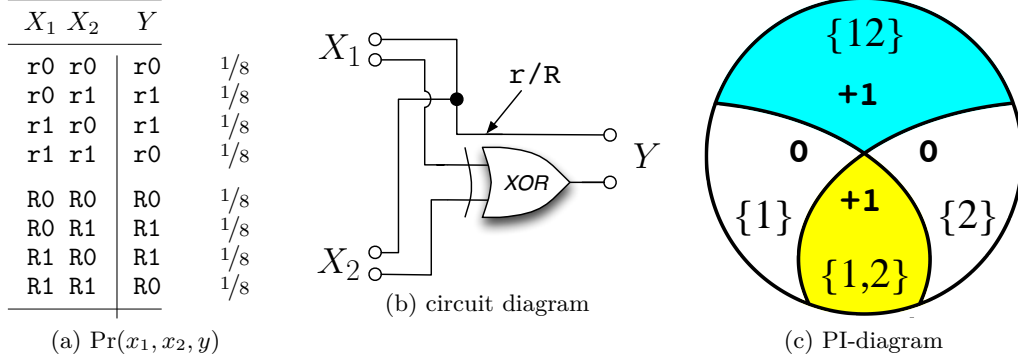


Figure 9: Example RDNXOR has one bit of redundancy and one bit of synergy. Yet for this example, $\text{WMS}(\mathbf{X} : Y) = 0$ bits.

4.3 Correlational importance: $\Delta I(\mathbf{X}; Y)$

Correlational importance, denoted ΔI , comes from [5, 22–25]. Correlational importance quantifies the “informational importance of conditional dependence” or the “information lost when ignoring conditional dependence” among the predictors decoding target Y . As conditional dependence is necessary for synergy, ΔI seems related to our intuitive conception of synergy. ΔI is defined as,

$$\Delta I(\mathbf{X}; Y) \equiv D_{\text{KL}}\left[\Pr(Y|X_{1\dots n}) \parallel \Pr_{\text{ind}}(Y|\mathbf{X})\right] \quad (10)$$

$$= \sum_{y, \mathbf{x} \in Y, \mathbf{X}} \Pr(y, x_{1\dots n}) \log \frac{\Pr(y|x_{1\dots n})}{\Pr_{\text{ind}}(y|\mathbf{x})}, \quad (11)$$

where $\Pr_{\text{ind}}(y|\mathbf{x}) \equiv \frac{\Pr(y) \prod_{i=1}^n \Pr(x_i|y)}{\sum_{y'} \Pr(y') \prod_{i=1}^n \Pr(x_i|y')}$. After some algebra⁷ eq. (11) becomes,

$$\Delta I(\mathbf{X}; Y) = \text{TC}(X_1; \dots; X_n|Y) - D_{\text{KL}}\left[\Pr(X_{1\dots n}) \parallel \sum_y \Pr(y) \prod_{i=1}^n \Pr(X_i|y)\right]. \quad (12)$$

ΔI is conceptually innovative and moreover agrees with our intuition for all of our examples thus far. Yet further examples reveal that ΔI measures something ever-so-subtly different from intuitive synergistic information.

The first example is [3]’s Figure 4 where ΔI exceeds the mutual information $I(X_{1\dots n}; Y)$ with $\Delta I(\mathbf{X}; Y) = 0.0145$ and $I(X_{1\dots n}; Y) = 0.0140$. This fact alone prevents interpreting ΔI as a loss of mutual information from $I(X_{1\dots n}; Y)$.⁸

⁷See Appendix F for the steps between eqs. (11) and (12).

⁸Although ΔI can not be a loss of mutual information, it could still be a loss of some alternative information such as Wyner’s common information [26].

Could ΔI upperbound synergy instead? We turn to example AND (Figure 10) with $n = 2$ independent binary predictors and target Y is the AND of X_1 and X_2 . Although AND’s PI-region exact decomposition remains uncertain, we can still bound the synergy. For example AND, the $\text{WMS}(\{X_1, X_2\} : Y) \approx 0.189$ and $\mathcal{S}_{\max}(\{X_1, X_2\} : Y) = 0.5$ bits. So we know the synergy must be between $(0.189, 0.5]$ bits. Despite this, $\Delta I(\mathbf{X}; Y) = 0.104$ bits, thus ΔI does not upperbound synergy.

Finally, in the face of duplicate predictors ΔI often *decreases*. From example AND to ANDDUPLICATE (Appendix A.0.1, Figure 13) ΔI drops 63% to 0.038 bits.

Taking all three examples together, we conclude ΔI measures something fundamentally different from synergistic information.




Figure 10: Example AND. The exact PI-decomposition of an AND-gate remains uncertain. But we can bound a , b , and c using WMS and \mathcal{S}_{\max} . In section 5 these bounds will be tightened. Most intriguingly, we’ll show that $a > 0$ despite $I(X_1 : X_2) = 0$.

5 Synergistic mutual information

We are all familiar with the English expression describing synergy as when the whole exceeds the “sum of its parts”. Although this informal adage captures the intuition underlying synergy, the formalization of this adage, WholeMinusSum synergy, “double-counts” whenever there is duplication (redundancy) among the parts. A mathematically correct adage should change “sum” to “union”—meaning synergy occurs when the whole exceeds the *union* of its parts. The sum adds duplicate information multiple times, whereas the union adds duplicate information only once. The union of parts never exceeds the sum.

The guiding intuition of “whole minus union” leads us to a novel measure denoted $\mathcal{S}_{\text{VK}}(\{X_1, \dots, X_n\} : Y)$, or $\mathcal{S}_{\text{VK}}(\mathbf{X} : Y)$, as the mutual information in the whole beyond the union of elements $\{X_1, \dots, X_n\}$.

Unfortunately, there’s no established measure of “union-information” in contemporary information theory. We introduce a novel technique, inspired by [27], for defining the union information among n predictors. We numerically compute the union information by noisifying the joint distribution $\Pr(X_{1\dots n} | Y)$ such that only the correlations with singleton predictors are preserved. This is achieved like so,

$$\begin{aligned}
 I_{\text{VK}}(\{X_1, \dots, X_n\} : Y) &\equiv \min_{\Pr^*(X_1, \dots, X_n, Y)} I^*(X_{1\dots n} : Y) & (13) \\
 &\text{subject to: } \Pr^*(X_i, Y) = \Pr(X_i, Y) \quad \forall i,
 \end{aligned}$$

where $I^*(X_{1\dots n} : Y) \equiv D_{\text{KL}}[\text{Pr}^*(X_{1\dots n}, Y) \parallel \text{Pr}^*(X_{1\dots n}) \text{Pr}^*(Y)]$.

Without any constraint on the distribution $\text{Pr}^*(X_1, \dots, X_n, Y)$, the minimum of eq. (13) is trivially found to be zero bits because simply setting $\text{Pr}^*(X_{1\dots n})$ to a constant makes $I^*(X_{1\dots n} : Y) = 0$ bits. Therefore we must put some constraint on $\text{Pr}^*(X_1, \dots, X_n, Y)$. As all bits a singleton X_i knows about Y are determined by the joint distribution $\text{Pr}(X_i, Y)$, we simply prevent the minimization from altering these distributions, and presto we arrive at the constraint $\text{Pr}^*(X_i, Y) = \text{Pr}(X_i, Y) \forall i$.⁹ Finally, we prove that a minimum of eq. (13) always exists because setting $\text{Pr}^*(x_1, \dots, x_n, y) = \text{Pr}(y) \prod_{i=1}^n \text{Pr}(x_i|y)$ always satisfies the constraints.

Unfortunately, we currently have no analytic way to calculate eq. (13), however, we do have an analytic upperbound on it. Applying this to AND's PI-decomposition allows us to tighten the bounds in Figure 10 to those in Figure 11.




Figure 11: Revisiting example AND. Using the analytic upperbound on I_{VK} in Appendix D, we arrive at the Pr^* distribution in (a). Using this distribution, we tighten the bounds on a , b , and c . Intriguingly, we see that despite $I(X_1 : X_2) = 0$, that $a > 0$. **Note:** Previous versions (preprints) of this paper *erroneously asserted* independent predictors could not convey redundant information, i.e. that $I(X_1 : X_2) = 0$ entailed $I_{\cap}(\{X_1, X_2\} : Y) = 0$.

Our union-information measure I_{VK} satisfies several desired properties for a union-information measure.¹⁰ Once the union information is computed, the \mathcal{S}_{VK} synergy is simply,

$$\mathcal{S}_{\text{VK}}(\{X_1, \dots, X_n\} : Y) \equiv I(X_{1\dots n} : Y) - I_{\text{VK}}(\{X_1, \dots, X_n\} : Y) . \quad (14)$$

\mathcal{S}_{VK} synergy quantifies the total “informational work” strictly the coalitions within $X_{1\dots n}$ perform in reducing the uncertainty of Y . Pleasingly, \mathcal{S}_{VK} is bounded¹¹ by the WholeMinus-Sum synergy (which underestimates the intuitive synergy) and \mathcal{S}_{max} (which overestimates intuitive synergy),

$$\max [0, \text{WMS}(\mathbf{X} : Y)] \leq \mathcal{S}_{\text{VK}}(\mathbf{X} : Y) \leq \mathcal{S}_{\text{max}}(\mathbf{X} : Y) \leq I(X_{1\dots n} : Y) . \quad (15)$$

6 Properties of I_{VK}

Our measure of the union information I_{VK} satisfies several desirable properties for the union-information¹²:

(GP) Global Positivity. $I_{\text{VK}}(\mathbf{X} : Y) \geq 0$

(SR) Self-Redundancy. The union information a single predictor X_1 has about the target Y is equal to the Shannon mutual information between the predictor and the target, i.e. $I_{\text{VK}}(X_1 : Y) = I(X_1 : Y)$.

⁹We could have instead chosen the *looser* constraint $I^*(X_i : Y) = I(X_i : Y) \forall i$, but $\text{Pr}^*(X_i, Y) = \text{Pr}(X_i, Y) \forall i$ ensures we preserve the “same bits”, not just the same magnitude of bits.

¹⁰For details see Section 6 and Appendix C.

¹¹Proven in Appendix E.2.

¹²For proofs see Appendix C.

- (**S**₀) Weak Symmetry. $I_{\text{VK}}(X_1, \dots, X_n : Y)$ is invariant under reordering X_1, \dots, X_n .
- (**M**) Monotonicity. $I_{\text{VK}}(X_1, \dots, X_n : Y) \leq I_{\text{VK}}(X_1, \dots, X_n, W : Y)$ with equality if W is “informationally poorer” than some $X_i \in \{X_1, \dots, X_n\}$, i.e. $\exists H(W|X_i) = 0$ for some $i \in \{1, \dots, n\}$.
- (**TM**) Target Monotonicity. For all random variables Y and Z , $I_{\text{VK}}(\mathbf{X} : Y) \leq I_{\text{VK}}(\mathbf{X} : YZ)$.
- (**LP**₀) Weak Local Positivity. For $n = 2$ predictors, the derived “partial informations” [6] are nonnegative. This is equivalent to,

$$\max [I(X_1 : Y), I(X_2 : Y)] \leq I_{\text{VK}}(X_1, X_2 : Y) \leq I(X_1 X_2 : Y) .$$

- (**Id**₁) Strong Identity. $I_{\text{VK}}(X_1, \dots, X_n : X_{1\dots n}) = H(X_{1\dots n})$.

7 Applying the measures to our examples

Table 1 summarizes the results of all four measures applied to our examples.

RDN (Figure 3). There is exactly one bit of redundant information and all measures reach their intended answer. For the axiomatically minded, the equality condition of (**M**) is sufficient for the desired answer.

UNQ (Figure 4). \mathcal{S}_{max} ’s miscategorization of unique information as synergistic reveals itself. Intuitively, there are two bits of unique information and no synergy. However, \mathcal{S}_{max} reports one bit of synergistic information. For the axiomatically minded, property (**Id**) is sufficient (but not necessary) for the desired answer.

XOR (Figure 5). There is exactly one bit of synergistic information. All measures reach the desired answer of 1 bit.

XORDUPLICATE (Figure 6). Target Y is specified by the coalition $X_1 X_2$ as well as by the coalition $X_3 X_2$, thus $I(X_1 X_2 : Y) = I(X_3 X_2 : Y) = H(Y) = 1$ bit. All measures reach the expected answer of 1 bit.

XORLOSES (Figure 7). Target Y is specified by the coalition $X_1 X_2$ as well as by the singleton X_3 , thus $I(X_1 X_2 : Y) = I(X_3 : Y) = H(Y) = 1$ bit. Together this means there is one bit of redundancy between the coalition $X_1 X_2$ and the singleton X_3 as illustrated by the +1 in PI-region $\{3, 12\}$. All measures account for this redundancy and reach the desired answer of 0 bits.

RDNXOR (Figure 9). This example has one bit of synergy as well as one bit of redundancy. In accordance with Figure 8a, WholeMinusSum measures *synergy minus redundancy* to calculate $1 - 1 = 0$ bits. On the other hand, \mathcal{S}_{max} , ΔI , and \mathcal{S}_{VK} are not misled by the co-existence of synergy and redundancy and correctly report 1 bit of synergistic information.

AND (Figure 10). This example is a simple case where correlational importance, $\Delta I(\mathbf{X}; Y)$, disagrees with the intuitive value for synergy. The WholeMinusSum synergy—an unambiguous *lowerbound* on the intuitive synergy—is 0.189 bits, yet $\Delta I(\mathbf{X}; Y) = 0.104$ bits. We can’t perfectly determine \mathcal{S}_{VK} , but we can lowerbound \mathcal{S}_{VK} using our analytic bound, as well as upperbound it using \mathcal{S}_{max} . This gives $0.270 \leq \mathcal{S}_{\text{VK}} \leq 1/2$.

The three supplementary examples in Appendix A: RDNUNQXOR, ANDDUPLICATE, and XORMULTICOAL aren’t essential for understanding this paper and are for the intellectual pleasure of advanced readers.

Table 1 shows that no prior measure of synergy consistently matches intuition even for $n = 2$. To summarize,

1. I_{max} synergy, \mathcal{S}_{max} , overestimates the intuitive synergy when two or more predictors convey unique information about the target (e.g. UNQ).
2. WholeMinusSum synergy, WMS, inadvertently double-subtracts redundancies and thus underestimates the intuitive synergy (e.g. RDNXOR). Duplicating predictors often decreases WholeMinusSum synergy (e.g. ANDDUPLICATE).

Example	\mathcal{S}_{\max}	WMS	ΔI	\mathcal{S}_{VK}
RDN	0	-1	0	0
UNQ	1	0	0	0
XOR	1	1	1	1
XORDUPLICATE	1	1	1	1
XORLOSES	0	0	0	0
RDNXOR	1	0	1	1
AND	1/2	0.189	0.104	[0.270, 1/2]
RDNUNQXOR	2	0	1	1
ANDDUPLICATE	1/2	-0.123	0.038	[0.270, 1/2]
XORMULTICOAL	1	1	1	1

Table 1: Synergy measures for our examples. Answers conflicting with intuitive synergistic information are in **red**. The \mathcal{S}_{VK} value for AND and ANDDUPLICATE is not conclusively known, but can be bounded.

3. Correlational importance, ΔI , is not bounded by the Shannon mutual information, underestimates the known lowerbound on synergy (e.g. AND), and duplicating predictors often decreases correlational importance (e.g. ANDDUPLICATE). Altogether, ΔI does not quantify the intuitive synergistic information (nor was it intended to).

8 Conclusion

Fundamentally, we assert that synergy quantifies how much the whole exceeds the *union* of its parts. Considering synergy as the whole minus the *sum* of its parts inadvertently “double-subtracts” redundancies, thus *underestimating* synergy. Within information theory, PI-diagrams, a generalization of Venn diagrams, are immensely helpful in improving one’s intuition for synergy.

We demonstrated with RDNXOR and RDNUNQXOR that a single state can simultaneously carry redundant, unique, and synergistic information. This fact is underappreciated, and prior work often implicitly assumed these three types of information could not coexist in a single state.

We introduced a novel measure of synergy, \mathcal{S}_{VK} , (eq. (14)). Unfortunately our expression is not easily computable, and until we have an explicit analytic solution to the minimization in I_{VK} the best one can do is numerical optimization using our analytic upperbound (Appendix D) as a starting point.

Along with our examples, we consider our introduction of a candidate for the union information, I_{VK} (eq. (13)) and its upperbound our primary contributions to the literature.

Finally, by means of our analytic upperbound on I_{VK} we’ve shown that, at least for our measure, *independent predictors can convey redundant information about a target*, e.g. Figure 11.

Acknowledgments

We thank Suzannah Fraker, Tracey Ho, Artemy Kolchinsky, Chris Adami, Giulio Tononi, Jim Beck, Nihat Ay, and Paul Williams for extensive discussions. This research was funded by the Paul G. Allen Family Foundation and a DOE CSGF fellowship to VG.

References

- [1] Narayanan NS, Kimchi EY, Laubach M (2005) Redundancy and synergy of neuronal ensembles in motor cortex. The Journal of Neuroscience 25: 4207-4216.

- [2] Balduzzi D, Tononi G (2008) Integrated information in discrete dynamical systems: motivation and theoretical framework. *PLoS Computational Biology* 4: e1000091.
- [3] Schneidman E, Bialek W, II MB (2003) Synergy, redundancy, and independence in population codes. *Journal of Neuroscience* 23: 11539–53.
- [4] Bell AJ (2003) The co-information lattice. In: Amari S, Cichocki A, Makino S, Murata N, editors, *Fifth International Workshop on Independent Component Analysis and Blind Signal Separation*. Springer.
- [5] Nirenberg S, Carcieri SM, Jacobs AL, Latham PE (2001) Retinal ganglion cells act largely as independent encoders. *Nature* 411: 698–701.
- [6] Williams PL, Beer RD (2010) Nonnegative decomposition of multivariate information. *CoRR* abs/1004.2515.
- [7] White D, Rabago-Smith M (2011) Genotype-phenotype associations and human eye color. *Journal of Human Genetics* 56: 5–7.
- [8] Schneidman E, Still S, Berry MJ, Bialek W (2003) Network information and connected correlations. *Phys Rev Lett* 91: 238701-238705.
- [9] Anastassiou D (2007) Computational analysis of the synergy among multiple interacting genes. *Molecular Systems Biology* 3: 83.
- [10] Amari S (1999) Information geometry on hierarchical decomposition of stochastic interactions. *IEEE Transaction on Information Theory* 47: 1701–1711.
- [11] Han TS (1978) Nonnegative entropy measures of multivariate symmetric correlations. *Information and Control* 36: 133–156.
- [12] Harder M, Salge C, Polani D (2012) A bivariate measure of redundant information. *CoRR* abs/1207.2080.
- [13] Bertschinger N, Rauh J, Olbrich E, Jost J (2012) Shared information – new insights and problems in decomposing information in complex systems. *CoRR* abs/1210.5902.
- [14] Lizier JT, Flecker B, Williams PL (2013) Towards a synergy-based approach to measuring information modification. *CoRR* abs/1303.3440.
- [15] Cover TM, Thomas JA (1991) *Elements of Information Theory*. New York, NY: John Wiley.
- [16] Weisstein EW (2011). Antichain. <http://mathworld.wolfram.com/Antichain.html>.
- [17] Comtet L (1998) *Advanced Combinatorics: The Art of Finite and Infinite Expansions*. Dordrecht, Netherlands: Reidel, 271–273 pp.
- [18] DeWeese MR, Meister M (1999) How to measure the information gained from one symbol. *Network* 10: 325-340.
- [19] Gawne TJ, Richmond BJ (1993) How independent are the messages carried by adjacent inferior temporal cortical neurons? *Journal of Neuroscience* 13: 2758-71.
- [20] Gat I, Tishby N (1999) Synergy and redundancy among brain cells of behaving monkeys. In: *Advances in Neural Information Proceedings systems*. MIT Press, pp. 465–471.
- [21] Chechik G, Globerson A, Anderson MJ, Young ED, Nelken I, et al. (2002) Group redundancy measures reveal redundancy reduction in the auditory pathway. In: Dietterich TG, Becker S, Ghahramani Z, editors, *NIPS 2002*. Cambridge, MA: MIT Press, pp. 173–180.
- [22] Panzeri S, Treves A, Schultz S, Rolls ET (1999) On decoding the responses of a population of neurons from short time windows. *Neural Comput* 11: 1553–1577.
- [23] Nirenberg S, Latham PE (2003) Decoding neuronal spike trains: How important are correlations? *Proceedings of the National Academy of Sciences* 100: 7348–7353.
- [24] Pola G, Thiele A, Hoffmann KP, Panzeri S (2003) An exact method to quantify the information transmitted by different mechanisms of correlational coding. *Network* 14: 35–60.
- [25] Latham PE, Nirenberg S (2005) Synergy, redundancy, and independence in population codes, revisited. *Journal of Neuroscience* 25: 5195-5206.

- [26] Lei W, Xu G, Chen B (2010) The common information of n dependent random variables. Forty-Eighth Annual Allerton Conference on Communication, Control, and Computing abs/1010.3613: 836-843.
- [27] Maurer UM, Wolf S (1999) Unconditionally secure key agreement and the intrinsic conditional information. IEEE Transactions on Information Theory 45: 499-514.

A Three extra examples

For the reader's intellectual pleasure, we include three more sophisticated examples: RDNUNQXOR, ANDDUPLICATE, and XORMULTICOAL.

X_1	X_2	Y		X_1	X_2	Y	
ra0	rb0	rab0	$1/32$	Ra0	Rb0	Rab0	$1/32$
ra0	rb1	rab1	$1/32$	Ra0	Rb1	Rab1	$1/32$
ra1	rb0	rab1	$1/32$	Ra1	Rb0	Rab1	$1/32$
ra1	rb1	rab0	$1/32$	Ra1	Rb1	Rab0	$1/32$
ra0	rB0	raB0	$1/32$	Ra0	RB0	RaB0	$1/32$
ra0	rB1	raB1	$1/32$	Ra0	RB1	RaB1	$1/32$
ra1	rB0	raB1	$1/32$	Ra1	RB0	RaB1	$1/32$
ra1	rB1	raB0	$1/32$	Ra1	RB1	RaB0	$1/32$
rA0	rb0	rAb0	$1/32$	RA0	Rb0	RAb0	$1/32$
rA0	rb1	rAb1	$1/32$	RA0	Rb1	RAb1	$1/32$
rA1	rb0	rAb1	$1/32$	RA1	Rb0	RAb1	$1/32$
rA1	rb1	rAb0	$1/32$	RA1	Rb1	RAb0	$1/32$
rA0	rB0	rAB0	$1/32$	RA0	RB0	RAB0	$1/32$
rA0	rB1	rAB1	$1/32$	RA0	RB1	RAB1	$1/32$
rA1	rB0	rAB1	$1/32$	RA1	RB0	RAB1	$1/32$
rA1	rB1	rAB0	$1/32$	RA1	RB1	RAB0	$1/32$

(a) $\Pr(x_1, x_2, y)$




Figure 12: Example RDNUNQXOR weaves examples RDN, UNQ, and XOR into one. $I(X_1X_2:Y) = H(Y) = 4$ bits. This example is pleasing because it puts exactly one bit in each PI-region.

A.0.1 Example AndDuplicate

ANDDUPLICATE adds a duplicate predictor to example AND to show how ΔI responds to a duplicate predictor in a less pristine example than XOR. Unlike XOR, in example AND there's also unique and redundant information. Will this cause the loss of synergy in the spirit of XORLOSES? Taking each one at a time:

- Predictor X_2 is unaltered from example AND. Thus X_2 's unique information stays the same. AND's $\{2\} \rightarrow$ ANDDUPLICATE's $\{2\}$.
- Predictor X_3 is identical to X_1 . Thus all of X_1 's unique information in AND becomes redundant information between predictors X_1 and X_3 . AND's $\{1\} \rightarrow$ ANDDUPLICATE's $\{1, 3\}$.

- In AND there is synergy between X_1 and X_2 , and this synergy is still present in ANDDUPLICATE. Just as in XORDUPLICATE, the only difference is that now an identical synergy also exists between X_3 and X_2 . Thus AND's $\{12\} \rightarrow$ ANDDUPLICATE's $\{12, 23\}$.
- Predictor X_3 is identical to X_1 . Therefore any information in AND that is specified by both X_1 and X_2 is now specified by X_1 , X_2 , and X_3 . Thus AND's $\{1, 2\} \rightarrow$ ANDDUPLICATE's $\{1, 2, 3\}$.




Figure 13: Example ANDDUPLICATE. The total mutual information is the same as in AND, $I(X_1X_2:Y) = I(X_1X_2X_3:Y) = 0.811$ bits. Every PI-region in example AND maps to a PI-region in ANDDUPLICATE. The intuitive synergistic information is unchanged from AND. However, correlational importance, ΔI , arrives at 0.104 bits of synergy for AND, and 0.038 bits for ANDDUPLICATE. ΔI is not invariant to duplicate predictors.


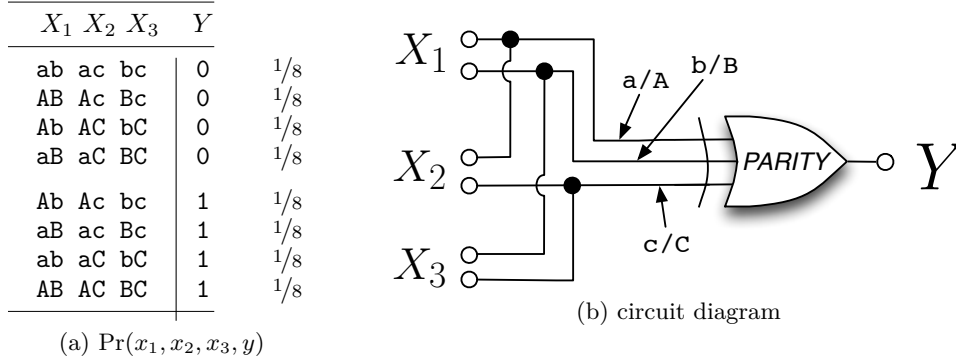


Figure 14: Example XORMULTICOAL demonstrates how the same information can be specified by multiple coalitions. In XORMULTICOAL the target Y has one bit of uncertainty, $H(Y) = 1$ bit, and Y is the *parity* of three incoming wires. Just as the output of XOR is specified only after knowing the state of both inputs, the output of XORMULTICOAL is specified only after knowing the state of all three wires. Each predictor is distinct and has access to two of the three incoming wires. For example, predictor X_1 has access to the a/A and b/B wires, X_2 has access to the a/A and c/C wires, and X_3 has access to the b/B and c/C wires. Although no single predictor specifies Y , any coalition of two predictors has access to all three wires and fully specifies Y , $I(X_1X_2:Y) = I(X_1X_3:Y) = I(X_2X_3:Y) = H(Y) = 1$ bit. In the PI-diagram this puts one bit in PI-region $\{12, 13, 23\}$ and zero everywhere else. All measures reach the expected answer of 1 bit of synergy.

B Connecting back to I_{\cap}

Our candidate measure of the union information, I_{VK} , gives rise to a measure of the intersection-information denoted I_{\cap}^{VK} . This is done by,

$$I_{\cap}^{VK}(\mathbf{X}:Y) = \sum_{\mathbf{S} \subseteq \mathbf{X}} (-1)^{|\mathbf{S}|+1} I_{VK}(\mathbf{S}:Y) . \quad (16)$$

C Desired properties of I_{\cup}

What properties does I_{\cap}^{VK} satisfy? We originally worked on proofs for which properties I_{\cap}^{VK} satisfies, but for $n > 2$ we were blocked by not having an analytic solution to I_{VK} . So we instead translated the I_{\cap} properties into the analogous I_{\cup} properties. Although one can't always prove the I_{\cap} version from the analogous I_{\cup} property, it is a start.

In addition to the properties in Section 6, we've proven that I_{VK} *does not satisfy* the property,

(S₁) Strong Symmetry. $I_{\cup}(\{X_1, \dots, X_n\}:Y)$ is invariant under reordering X_1, \dots, X_n, Y .

C.0.2 Proof of (GP)

Proven by the nonnegativity of mutual information.

C.0.3 Proof of (SR)

$$\begin{aligned} I_{VK}(X_1:Y) &\equiv \min_{\substack{p^*(x_1,y) \\ p^*(x_1,y)=p(x_1,y)}} I^*(X_1:Y) \\ &= I(X_1:Y) . \end{aligned}$$

C.0.4 Proof of (S₀)

There's only one instance of the terms in \mathbf{X} in the definition of I_{VK} , which is,

$$I_{VK}(\mathbf{X}:Y) \equiv \min_{\substack{p^*(X_1, \dots, X_n, Y) \\ p^*(X_i, Y)=p(X_i, Y) \quad \forall i}} I^*(X_1 \cdots X_n:Y) .$$

The term $I^*(X_1 \cdots X_n:Y)$ is invariant to the ordering of $X_1 \cdots X_n$. This is due to $\Pr^*(x_1, \dots, x_n) = \Pr^*(x_n, \dots, x_1)$. Thus I_{VK} is invariant to the ordering of $\{X_1, \dots, X_n\}$.

C.0.5 Proof of (LP₀)

$$I_{VK}(\mathbf{X}:Y) \leq I(X_{1\dots n}:Y) .$$

This is proven by the condition that $\Pr(X_1, \dots, X_n, Y)$ satisfies the constraints on the minimizing distribution in I_{VK} . Thus $I^*(X_{1\dots n}:Y) \leq I(X_{1\dots n}:Y)$.

C.0.6 Disproof of (S₁)

We show that, $I_{VK}(\{X, Y\}:Z) \neq I_{VK}(\{X, Z\}:Y)$ by setting $X = Y$ where $H(X) > 0$, and Z is a constant, $I_{VK}(\{X, Y\}:Z) = 0$ yet $I_{VK}(\{X, Z\}:Y) = H(X)$.

C.0.7 Proof of (Id₁)

$$\begin{aligned} I_{VK}(\mathbf{X} : X_{1\dots n}) &\equiv \min_{\substack{p^*(X_1, \dots, X_n, X_{1\dots n}) \\ p^*(X_i, X_{1\dots n})=p(X_i, X_{1\dots n}) \quad \forall i}} I^*(X_{1\dots n} : X_{1\dots n}) & (17) \\ &= \min_{\substack{p^*(X_1, \dots, X_n, X_{1\dots n}) \\ p^*(X_i, X_{1\dots n})=p(X_i, X_{1\dots n}) \quad \forall i}} H^*(X_{1\dots n}) , & (18) \end{aligned}$$

Then because $p^*(X_{1\dots n}) = p(X_{1\dots n})$,

$$I_{\text{VK}}(\mathbf{X} : X_{1\dots n}) = H(X_{1\dots n}) . \quad (19)$$

D Analytic upperbound on $I_{\text{VK}}(\mathbf{X} : Y)$

Our analytic upperbound on I_{VK} starts with the n joint distributions we wish to preserve: $\Pr(X_1, Y), \dots, \Pr(X_n, Y)$. From one these joint distributions, e.g. $\Pr(X_1, Y)$, we compute the marginal probability distribution $\Pr(Y)$ by summing over the index of $x_1 \in X_1$,

$$\Pr(Y) = \left\{ \sum_{x_1 \in X_1} \Pr(x_1, y) : \forall y \in Y \right\} . \quad (20)$$

Then, for every state $y \in Y$ we compute n conditional distributions $\Pr(X_1|y), \dots, \Pr(X_n|y)$ via,

$$\Pr(X_i|Y = y) = \left\{ \frac{\Pr(x_i, y)}{\Pr(y)} : \forall x_i \in X_i \right\} . \quad (21)$$

With the marginal distribution $\Pr(Y)$ and the $|Y| \cdot n$ conditional distributions, we construct a novel, artificial joint distribution $\Pr^*(X_1, \dots, X_n, Y)$ defined by,

$$\Pr^*(x_1, \dots, x_n, y) \equiv \Pr(y) \prod_{i=1}^n \Pr(x_i|y) . \quad (22)$$

This novel, artificial joint distribution $\Pr^*(X_1, \dots, X_n, Y)$ satisfies the constraints $\Pr^*(X_i, Y) = \Pr(X_i, Y) \forall i$. This is proven by,

$$\Pr^*(x_i, y) = \underbrace{\sum_{x_1 \in X_1} \dots \sum_{x_n \in X_n}}_{\text{All except } x_i \in X_i} \Pr^*(x_1, \dots, x_n, y) \quad (23)$$

$$= \underbrace{\sum_{x_1 \in X_1} \dots \sum_{x_n \in X_n}}_{\text{All except } x_i \in X_i} \Pr(y) \prod_{j=1}^n \Pr(x_j|y) \quad (24)$$

$$= \underbrace{\sum_{x_1 \in X_1} \dots \sum_{x_n \in X_n}}_{\text{All except } x_i \in X_i} \Pr(x_i, y) \prod_{\substack{j=1 \\ j \neq i}}^n \Pr(x_j|y) \quad (25)$$

$$= \Pr(x_i, y) \underbrace{\sum_{x_1 \in X_1} \dots \sum_{\substack{x_n \in X_n \\ j=1 \\ j \neq i}}}_{\text{All except } x_i \in X_i} \prod_{j=1}^n \Pr(x_j|y) \quad (26)$$

$$= \Pr(x_i, y) . \quad (27)$$

The upperbound on I_{VK} is then the mutual information using this artificial \Pr^* distribution,

$$I^*(X_1 \dots X_n : Y) = \sum_{x_1 \in X_1} \dots \sum_{x_n \in X_n} \sum_{y \in Y} \Pr^*(x_1, \dots, x_n, y) \log \frac{\Pr^*(x_1, \dots, x_n, y)}{\Pr^*(x_1, \dots, x_n) \Pr^*(y)} , \quad (28)$$




Figure 15: The Directed Acyclic Graph generating the joint distribution $\Pr^*(x_1, \dots, x_n, y)$. This is a graphical representation of eq. (22).

where the terms $\Pr^*(x_1, \dots, x_n)$ and $\Pr^*(y)$ are defined by summing over the relevant indices of joint distribution $\Pr^*(X_1, \dots, X_n, Y)$,

$$\Pr^*(x_1, \dots, x_n) = \sum_{y' \in Y} \Pr^*(x_1, \dots, x_n, y') \quad (29)$$

$$= \sum_{y' \in Y} \Pr(y') \prod_{i=1}^n \Pr(x_i | y') ; \quad (30)$$

$$\Pr^*(y) = \sum_{x_1 \in X_1} \dots \sum_{x_n \in X_n} \Pr^*(x_1, \dots, x_n, y) \quad (31)$$

$$= \sum_{x_1 \in X_1} \dots \sum_{x_n \in X_n} \Pr(y) \prod_{i=1}^n \Pr(x_i | y) \quad (32)$$

$$= \Pr(y) \underbrace{\sum_{x_1 \in X_1} \dots \sum_{x_n \in X_n} \prod_{i=1}^n \Pr(x_i | y)}_{\text{sums to 1}} \quad (33)$$

$$= \Pr(y) . \quad (34)$$

Putting everything together, our analytic upperbound on I_{VK} is,

$$I_{\text{VK}}(\{X_1, \dots, X_n\} : Y) \leq I^*(X_{1..n} : Y) \quad (35)$$

$$= \sum_{x_1} \dots \sum_{x_n} \sum_y \Pr^*(x_1, \dots, x_n, y) \log \frac{\Pr^*(x_1, \dots, x_n, y)}{\Pr^*(x_1, \dots, x_n) \Pr^*(y)} \quad (36)$$

$$= \sum_{x_1} \dots \sum_{x_n} \sum_y \Pr^*(x_1, \dots, x_n, y) \log \frac{\Pr(y) \prod_{i=1}^n \Pr(x_i | y)}{\Pr^*(x_1, \dots, x_n) \Pr(y)} \quad (37)$$

$$= \sum_{x_1} \dots \sum_{x_n} \sum_y \Pr^*(x_1, \dots, x_n, y) \log \frac{\prod_{i=1}^n \Pr(x_i | y)}{\Pr^*(x_1, \dots, x_n)} \quad (38)$$

$$= \sum_y \Pr(y) \sum_{x_1} \dots \sum_{x_n} \prod_{i=1}^n \Pr(x_i | y) \log \frac{\prod_{i=1}^n \Pr(x_i | y)}{\sum_{y' \in Y} \Pr(y') \prod_{i=1}^n \Pr(x_i | y')} .$$

E Essential proofs

These proofs underpin essential claims about our introduced measure, synergistic mutual information.

E.1 Proof duplicate predictors don't increase synergy

We show that synergy being invariant to duplicate predictors follows from the equality condition of **(M)** of the intersection (as well as union) information.

We show that,

$$\mathcal{S}_{\text{VK}}(\mathbf{X}:Y) = \mathcal{S}_{\text{VK}}(\mathbf{X}':Y) ,$$

where $\mathbf{X}' \equiv \{X_1, \dots, X_n, X_1\}$. We show that $\mathcal{S}_{\text{VK}}(\mathbf{X}:Y) - \mathcal{S}_{\text{VK}}(\mathbf{X}':Y) = 0$.

$$0 = \mathcal{S}_{\text{VK}}(\mathbf{X}:Y) - \mathcal{S}_{\text{VK}}(\mathbf{X}':Y) \tag{39}$$

$$= I(X_{1\dots n}:Y) - I_{\text{VK}}(\mathbf{X}:Y) - I(X_{1\dots n}X_1:Y) + I_{\text{VK}}(\mathbf{X}':Y) \tag{40}$$

$$= I_{\text{VK}}(\mathbf{X}':Y) - I_{\text{VK}}(\mathbf{X}:Y) \tag{41}$$

$$= \sum_{\mathbf{T} \subseteq \mathbf{X}'} (-1)^{|\mathbf{T}|+1} I_{\cap}^{\text{VK}}(\mathbf{T}:Y) - \sum_{\mathbf{S} \subseteq \mathbf{X}} (-1)^{|\mathbf{S}|+1} I_{\cap}^{\text{VK}}(\mathbf{S}:Y) . \tag{42}$$

The terms that \mathbf{S} enumerates over is a subset of the terms that \mathbf{T} enumerates. Therefore the $\sum_{\mathbf{S} \subseteq \mathbf{X}}$ completely cancels, leaving,

$$0 = \sum_{\mathbf{T} \subseteq \mathbf{X}} (-1)^{|\mathbf{T}|} I_{\cap}^{\text{VK}}(\{X_1, T_1, \dots, T_{|\mathbf{T}|}\}:Y) . \tag{43}$$

If I_{\cap}^{VK} obeys **(M)**, then each term of eq. (43) s.t. $X_1 \notin \mathbf{T}$ cancels with the same term but with $X_1 \in \mathbf{T}$. This makes eq. (43) sum to zero, and completes the proof. Note we don't explicitly prove that I_{\cap}^{VK} satisfies **(M)**, but if it does, then duplicate predictors do not increase synergy.

E.2 Proof of bounds of $\mathcal{S}_{\text{VK}}(\mathbf{X}:Y)$

We show that,

$$\text{WMS}(\mathbf{X}:Y) \leq \mathcal{S}_{\text{VK}}(\mathbf{X}:Y) \leq \mathcal{S}_{\text{max}}(\mathbf{X}:Y) . \tag{44}$$

E.2.1 Proof that $\mathcal{S}_{\text{VK}}(\mathbf{X}:Y) \leq \mathcal{S}_{\text{max}}(\mathbf{X}:Y)$

We invoke the standard definitions of \mathcal{S}_{VK} and \mathcal{S}_{max} ,

$$\begin{aligned} \mathcal{S}_{\text{VK}}(\mathbf{X}:Y) &\equiv I(X_{1\dots n}:Y) - I_{\text{VK}}(\mathbf{X}:Y) \\ \mathcal{S}_{\text{max}}(\mathbf{X}:Y) &\equiv I(X_{1\dots n}:Y) - I_{\text{max}}(\mathbf{X}:Y) , \end{aligned}$$

where I_{VK} and I_{max} are defined as,

$$\begin{aligned} I_{\text{VK}}(\mathbf{X}:Y) &= \mathbb{E}_Y I_{\text{VK}}(\mathbf{X}:Y=y) \\ &= \mathbb{E}_Y \min_{\substack{p^*(X_1, \dots, X_n, Y) \\ p^*(X_i, Y) = p(X_i, Y) \quad \forall i}} I^*(X_{1\dots n}:Y=y) \end{aligned} \tag{45}$$

$$I_{\text{max}}(\mathbf{X}:Y) \equiv \mathbb{E}_Y \max_i I(X_i:Y=y) . \tag{46}$$

Now we prove $\mathcal{S}_{\text{VK}}(\mathbf{X}:Y) \leq \mathcal{S}_{\text{max}}(\mathbf{X}:Y)$ by showing that $I_{\text{VK}}(\mathbf{X}:Y) \geq I_{\text{max}}(\mathbf{X}:Y)$.

Proof.

$$\mathbb{E}_Y \text{I}_{\text{VK}}(\mathbf{X} : Y = y) \geq \mathbb{E}_Y \text{I}_{\text{max}}(\mathbf{X} : Y = y) \quad (47)$$

$$\mathbb{E}_Y [\text{I}_{\text{VK}}(\mathbf{X} : Y = y) - \text{I}_{\text{max}}(\mathbf{X} : Y = y)] \geq 0. \quad (48)$$

Now expanding $\text{I}_{\text{VK}}(\mathbf{X} : Y = y)$ and $\text{I}_{\text{max}}(\mathbf{X} : Y = y)$,

$$\mathbb{E}_Y \left[\left(\min_{\substack{p^*(X_1, \dots, X_n, Y) \\ p^*(X_i, Y) = p(X_i, Y) \quad \forall i}} \text{I}^*(X_{1\dots n} : Y = y) \right) - \max_i \text{I}(X_i : Y = y) \right] \geq 0. \quad (49)$$

We define the index $m \in \{1, \dots, n\}$ such that $m = \operatorname{argmax}_i \text{I}(X_i : Y = y)$. The predictor with the most information about state $Y = y$ is thus X_m . This yields,

$$\mathbb{E}_Y \left[\left(\min_{\substack{p^*(X_1, \dots, X_n, Y) \\ p^*(X_i, Y) = p(X_i, Y) \quad \forall i}} \text{I}^*(X_{1\dots n} : Y = y) \right) - \text{I}(X_m : Y = y) \right] \geq 0. \quad (50)$$

The constraint $p^*(X_i, Y) = p(X_i, Y)$ entails that $\text{I}(X_m : Y = y) = \text{I}^*(X_m : Y = y)$. Therefore we can pull $\text{I}(X_m : Y = y)$ inside the minimization as a constant,

$$\mathbb{E}_Y \left[\min_{\substack{p^*(X_1, \dots, X_n, Y) \\ p^*(X_i, Y) = p(X_i, Y) \quad \forall i}} \text{I}^*(X_{1\dots n} : Y = y) - \text{I}^*(X_m : Y = y) \right] \geq 0. \quad (51)$$

As X_m is a subset of predictors $X_{1\dots n}$, we can subtract it yielding,

$$\mathbb{E}_Y \left[\min_{\substack{p^*(X_1, \dots, X_n, Y) \\ p^*(X_i, Y) = p(X_i, Y) \quad \forall i}} \text{I}^*(X_{1\dots n \setminus m} : Y = y | X_m) \right] \geq 0. \quad (52)$$

The state-dependent conditional mutual information $\text{I}^*(X_{1\dots n \setminus m} : Y = y | X_m)$ is a Kullback-Liebler divergence. As such it is nonnegative. Likewise the minimum of a nonnegative quantity is also nonnegative.

$$\mathbb{E}_Y \left[\underbrace{\min_{\substack{p^*(X_1, \dots, X_n, Y) \\ p^*(X_i, Y) = p(X_i, Y) \quad \forall i}} \text{I}^*(X_{1\dots n \setminus m} : Y = y | X_m)}_{\geq 0} \right] \geq 0. \quad (53)$$

Finally, the expected value of a list of nonnegative quantities is nonnegative. And the proof that $\mathcal{S}_{\text{VK}}(\mathbf{X} : Y) \leq \mathcal{S}_{\text{max}}(\mathbf{X} : Y)$ is complete. \square

E.2.2 Proof that $\text{WMS}(\mathbf{X} : Y) \leq \mathcal{S}_{\text{VK}}(\mathbf{X} : Y)$

We invoke the standard definitions of WMS and \mathcal{S}_{VK} ,

$$\text{WMS}(\mathbf{X} : Y) \equiv \text{I}(X_{1\dots n} : Y) - \sum_{i=1}^n \text{I}(X_i : Y) \quad (54)$$

$$\mathcal{S}_{\text{VK}}(\mathbf{X} : Y) \equiv \text{I}(X_{1\dots n} : Y) - \text{I}_{\text{VK}}(X_{1\dots n} : Y) \quad (55)$$

$$= \text{I}(X_{1\dots n} : Y) - \min_{\substack{p^*(X_1, \dots, X_n, Y) \\ p^*(X_i, Y) = p(X_i, Y) \quad \forall i}} \text{I}^*(X_{1\dots n} : Y) . \quad (56)$$

We prove the conjecture $\text{WMS}(\mathbf{X} : Y) \leq \mathcal{S}_{\text{VK}}(\mathbf{X} : Y)$ by showing,

$$\min_{\substack{p^*(X_1, \dots, X_n, Y) \\ p^*(X_i, Y) = p(X_i, Y) \quad \forall i}} \text{I}^*(X_{1\dots n} : Y) \leq \sum_{i=1}^n \text{I}(X_i : Y) . \quad (57)$$

Given:

$$\begin{aligned} & \min_{\substack{p^*(X_1, \dots, X_n, Y) \\ p^*(X_1, Y) = p(X_1, Y) \\ \vdots \\ p^*(X_n, Y) = p(X_n, Y)}} \text{I}^*(X_{1\dots n} : Y) , \end{aligned} \quad (58)$$

the individual constraint $p^*(X_1, Y) = p(X_1, Y)$ can add at most $\text{I}(X_1 : Y)$ bits to $\text{I}^*(X_{1\dots n} : Y)$. Therefore we can upperbound eq. (58) by dropping the constraint $p^*(X_1, Y) = p(X_1, Y)$ and adding $\text{I}(X_1 : Y)$. This yields,

$$\begin{aligned} \text{I}_{\text{VK}}(\mathbf{X} : Y) \leq & \min_{\substack{p^*(X_1, \dots, X_n, Y) \\ p^*(X_2, Y) = p(X_2, Y) \\ \vdots \\ p^*(X_n, Y) = p(X_n, Y)}} \text{I}^*(X_{1\dots n} : Y) + \text{I}(X_1 : Y) . \end{aligned} \quad (59)$$

Likewise, the righthand-side of eq. (59) can be upperbounded by dropping the constraint $p^*(X_2, Y) = p(X_2, Y)$ and adding $\text{I}(X_2 : Y)$. This yields,

$$\begin{aligned} \min_{\substack{p^*(X_2, \dots, X_n, Y) \\ p^*(X_2, Y) = p(X_2, Y) \\ \vdots \\ p^*(X_n, Y) = p(X_n, Y)}} \text{I}^*(X_{1\dots n} : Y) \leq & \min_{\substack{p^*(X_3, \dots, X_n, Y) \\ p^*(X_3, Y) = p(X_3, Y) \\ \vdots \\ p^*(X_n, Y) = p(X_n, Y)}} \text{I}^*(X_{1\dots n} : Y) + \text{I}(X_1 : Y) + \text{I}(X_2 : Y) . \end{aligned} \quad (60)$$

Repeating this process n times yields,

$$\text{I}_{\text{VK}}(\mathbf{X} : Y) \leq \min_{p^*(X_1, \dots, X_n, Y)} \text{I}^*(X_{1\dots n} : Y) + \sum_{i=1}^n \text{I}(X_i : Y) \quad (61)$$

$$= \sum_{i=1}^n \text{I}(X_i : Y) . \quad (62)$$

F Algebraic simplification of ΔI

Prior literature [5, 23–25] defines $\Delta I(\mathbf{X}; Y)$ as,

$$\Delta I(\mathbf{X}; Y) \equiv D_{\text{KL}}\left[\Pr(Y|X_{1\dots n}) \parallel \Pr_{\text{ind}}(Y|\mathbf{X})\right] \quad (63)$$

$$= \sum_{\mathbf{x}, y \in \mathbf{X}, Y} \Pr(\mathbf{x}, y) \log \frac{\Pr(y|\mathbf{x})}{\Pr_{\text{ind}}(y|\mathbf{x})}. \quad (64)$$

Where,

$$\Pr_{\text{ind}}(Y = y|\mathbf{X} = \mathbf{x}) \equiv \frac{\Pr(y) \Pr_{\text{ind}}(\mathbf{X} = \mathbf{x}|Y = y)}{\Pr_{\text{ind}}(\mathbf{X} = \mathbf{x})} \quad (65)$$

$$= \frac{\Pr(y) \prod_{i=1}^n \Pr(x_i|y)}{\Pr_{\text{ind}}(\mathbf{x})} \quad (66)$$

$$\Pr_{\text{ind}}(\mathbf{X} = \mathbf{x}) \equiv \sum_{y \in Y} \Pr(Y = y) \prod_{i=1}^n \Pr(x_i|y) \quad (67)$$

The definition of ΔI , eq. (63), reduces to,

$$\Delta I(\mathbf{X}; Y) = \sum_{\mathbf{x}, y \in \mathbf{X}, Y} \Pr(\mathbf{x}, y) \log \frac{\Pr(y|\mathbf{x})}{\Pr_{\text{ind}}(y|\mathbf{x})} \quad (68)$$

$$= \sum_{\mathbf{x}, y \in \mathbf{X}, Y} \Pr(\mathbf{x}, y) \log \frac{\Pr(y|\mathbf{x}) \Pr_{\text{ind}}(\mathbf{x})}{\Pr(y) \prod_{i=1}^n \Pr(x_i|y)} \quad (69)$$

$$= \sum_{\mathbf{x}, y \in \mathbf{X}, Y} \Pr(\mathbf{x}, y) \log \frac{\Pr(\mathbf{x}|y)}{\prod_{i=1}^n \Pr(x_i|y)} \frac{\Pr_{\text{ind}}(\mathbf{x})}{\Pr(\mathbf{x})} \quad (70)$$

$$= \sum_{\mathbf{x}, y \in \mathbf{X}, Y} \Pr(\mathbf{x}, y) \log \frac{\Pr(\mathbf{x}|y)}{\prod_{i=1}^n \Pr(x_i|y)} + \sum_{\mathbf{x}, y \in \mathbf{X}, Y} \Pr(\mathbf{x}, y) \log \frac{\Pr_{\text{ind}}(\mathbf{x})}{\Pr(\mathbf{x})}$$

$$= \sum_{\mathbf{x}, y \in \mathbf{X}, Y} \Pr(\mathbf{x}, y) \log \frac{\Pr(\mathbf{x}|y)}{\prod_{i=1}^n \Pr(x_i|y)} - \sum_{\mathbf{x} \in \mathbf{X}} \Pr(\mathbf{x}) \log \frac{\Pr(\mathbf{x})}{\Pr_{\text{ind}}(\mathbf{x})} \quad (71)$$

$$= D_{\text{KL}}\left[\Pr(X_{1\dots n}|Y) \parallel \prod_{i=1}^n \Pr(X_i|Y)\right] - D_{\text{KL}}[\Pr(X_{1\dots n}) \parallel \Pr_{\text{ind}}(\mathbf{X})]$$

$$= \text{TC}(X_1; \dots; X_n|Y) - D_{\text{KL}}[\Pr(X_{1\dots n}) \parallel \Pr_{\text{ind}}(\mathbf{X})]. \quad (72)$$

where $\text{TC}(X_1; \dots; X_n|Y)$ is the conditional total correlation among the predictors given Y .