A Perspective on Unique Information: Directionality, Intuitions, and Secret Key Agreement

Ryan G. James,[*] Jeffrey Emenheiser,[†] and James P. Crutchfield[‡]
Complexity Sciences Center and Physics Department, University of California at Davis, One Shields Avenue, Davis, CA 95616

(Dated: August 27, 2018)

Recently, the partial information decomposition emerged as a promising framework for identifying the meaningful components of the information contained in a joint distribution. Its adoption and practical application, however, have been stymied by the lack of a generally-accepted method of quantifying its components. Here, we briefly discuss the bivariate (two-source) partial information decomposition and two implicitly directional interpretations used to intuitively motivate alternative component definitions. Drawing parallels with secret key agreement rates from information-theoretic cryptography, we demonstrate that these intuitions are mutually incompatible and suggest that this underlies the persistence of competing definitions and interpretations. Having highlighted this hitherto unacknowledged issue, we outline several possible solutions.

PACS numbers: 05.45.-a 89.75.Kd 89.70.+c 02.50.-r
Keywords: information theory, partial information decomposition, secret key agreement, cryptography

I. INTRODUCTION

Consider a joint distribution over "source" variables X_0 and X_1 and "target" Y. Such distributions arise in many settings: sensory integration, logical computing, neural coding, functional network inference, and many others. One promising approach to understanding how the information shared between X_0, X_1, and Y is organized is the *partial information decomposition* (PID) [1]. This decomposition seeks to quantify how much of the information shared between X_0, X_1, and Y is done so *redundantly*, how much is *uniquely* attributable to X_0, how much is *uniquely* attributable to X_1, and finally how much arises *synergistically* by considering both X_0 and X_1 together.

Unfortunately, the lack of a commonly accepted method of quantifying these components has hindered PID's adoption. In point of fact, several proposed axioms are not mutually consistent. And, to date, there is little agreement as to which should hold. Here, we take a step toward rectifying these issues by bringing to light a potentially fundamental inconsistency in the intuitions commonly and often implicitly brought to bear upon information decomposition. We make the intuitions quantitative by appealing to information-theoretic cryptography. Taken together, our observations suggest that the context in which PID is applied should determine how its components are quantified.

Our development proceeds as follows. Section II briefly describes the two-source PID. Section III calls out the two distinct intuitions often used in interpreting PID. Section IV introduces a prototype distribution that highlights the issues and we interpret it through the lenses of the two intuitions. Section V defines secret key agreement rates and computes them for the prototype distribution. Section VI then discusses how the two intuitions relate to secret key agreement rates and identifies when the latter result in viable decompositions. Finally, Section VII summarizes our findings and speculates as to how future developments can bring consistency to PID.

II. PARTIAL INFORMATION DECOMPOSITION

Two-source PID seeks to decompose the mutual information $\mathrm{I}[X_0 X_1 : Y]$ between "sources" X_0 and X_1 and a "target" Y into four nonnegative components. The components identify information that is redundant, uniquely associated with X_0, uniquely associated with X_1, and synergistic:

$$\begin{aligned} \mathrm{I}[X_0 X_1 : Y] = {} & \mathrm{I}_\partial\left[X_0 \cdot X_1 \to Y\right] && \textit{redundant} \\ & + \mathrm{I}_\partial\left[X_0 \to Y \setminus X_1\right] && \textit{unique from } X_0 \\ & + \mathrm{I}_\partial\left[X_1 \to Y \setminus X_0\right] && \textit{unique from } X_1 \\ & + \mathrm{I}_\partial\left[X_0 X_1 \to Y\right] ~. && \textit{synergistic} \end{aligned}$$

Furthermore, the mutual information between X_0 and Y is decomposed into two components:

$$\begin{aligned} \mathrm{I}[X_0 : Y] = {} & \mathrm{I}_\partial\left[X_0 \cdot X_1 \to Y\right] && \textit{redundant} \\ & + \mathrm{I}_\partial\left[X_0 \to Y \setminus X_1\right] ~. && \textit{unique with } X_0 \end{aligned}$$

[*] rgjames@ucdavis.edu
[†] jemenheiser@ucdavis.edu
[‡] chaos@ucdavis.edu

And, similarly:

$$\begin{aligned}\mathrm{I}[X_1 : Y] = &\mathrm{I}_\partial\left[X_0 \cdot X_1 \to Y\right] && \textit{redundant} \\ &+ \mathrm{I}_\partial\left[X_1 \to Y \setminus X_0\right] \ . && \textit{unique with } X_1\end{aligned}$$

In this way, PID relates the four component informations. However, it does not uniquely determine how to quantify them. To do this, a definition must be supplied for one of them and then the others follow.

This allows for a range of choices. In the case that one wishes to directly quantify the unique informations $\mathrm{I}_\partial\left[X_0 \to Y \setminus X_1\right]$ and $\mathrm{I}_\partial\left[X_1 \to Y \setminus X_0\right]$, a consistency relation must hold when they are computed independently:

$$\begin{aligned}&\mathrm{I}_\partial\left[X_0 \to Y \setminus X_1\right] + \mathrm{I}[X_1 : Y] \\ &\qquad = \mathrm{I}_\partial\left[X_1 \to Y \setminus X_0\right] + \mathrm{I}[X_0 : Y] \ . \quad (1)\end{aligned}$$

III. THE CAMEL AND THE ELEPHANT

There are two common ways of thinking about PID. These approaches differ only in the (implied) directionality of cause and effect—a property unspecified by PID.

In the first approach, one thinks of X_0 and X_1 as "inputs" that, when combined, produce Y, a "output". While seemingly helpful labels, their use already imports an unwarranted semantics to the relationship between the three random variables. In this, it inadvertently begs the main issue we wish to raise here, while at the same time illustrating the issue.

When taking this view of PID, one generally asks questions such as "How much information in X_0 is uniquely conveyed to Y?". From this vantage, considering the role of the individual channels $X_0 \to Y$ and $X_1 \to Y$ might or might not help develop intuition. Recalling the aphorism "a camel is a horse designed by committee", we call this the *camel intuition* as particular input events X_0 and X_1 come together to describe an output Y.

In the second approach, one considers X_0 and X_1 as "noisy observations" or "representations" of a single underlying object Y. When taking this view, one might ask a question such as "How much information in Y is uniquely captured by X_0?". Under this, the individual channels $Y \to X_0$ and $Y \to X_1$ take on primary importance. After the parable of the blind men describing an elephant, we call this the *elephant intuition* since particular objects Y may be described by various, possibly partial, representations, X_0 and X_1.

Pnt. Unq.			
X_0	X_1	Y	Pr
0	1	1	1/4
1	0	1	1/4
0	2	2	1/4
2	0	2	1/4

TABLE I. The *pointwise unique* distribution.

IV. THE POINTWISE UNIQUE DISTRIBUTION

The *pointwise unique* distribution [2] is given by the events and probabilities displayed in Table I: at any time exactly one of X_0 or X_1 is a '1' or '2' and matches Y, while the other is '0'. Let's now interpret this distribution by adopting the camel and elephant intuitions in turn. We will see that they provide contradictory interpretations of the relationships between the variables.

Adopting the camel intuition, we consider the ways in which X_0 influences Y. It is easy to see that half of the time (Table I's 1st and 3rd rows) X_0 is unable to say anything about the state of Y. The other half of the time (the 2nd and 4th rows) X_0 and Y are perfectly correlated, while X_1 is ignorant as to their state. Analogously, this is true when considering how X_1 influences Y. In this way, we interpret the distribution's PID as consisting entirely of unique informations. The camel intuition is summarized in Table II.

When adopting the elephant intuition, however, a strikingly different picture emerges. Taking the viewpoint of Y, both single channel distributions $p(X_0|Y)$ and $p(X_1|Y)$ are identical. So, any information shared with one must be redundantly shared with the other. These channels do not allow one to determine the states of either X_0 or X_1. What is learned, however, is that exactly one of them matches Y, while the other is '0'. Furthermore, removing the remaining uncertainty in the values of X_0 and X_1 requires observing one of them—a synergistic effect. The resulting elephant analysis is also summarized in Table II.

In short, the two directional PID interpretations lead to contradictory quantifications. From the viewpoint of camels, elephant approaches create redundancy where there is none. From the vantage of elephants, camels draw distinctions where none exist. This has been discussed by Ref. [3] regarding whether or not unique information should depend on $\mathrm{I}[X_0 : X_1]$. From the camel's point of view, ignoring this as a constraint may "artificially correlate" X_0 and X_1 and thereby inflate redundancy. This viewpoint can be more directly illustrated by considering the intermediate distribution from which $\mathrm{I}_{\mathrm{BROJA}}$ [4]—an elephant—computes unique information

Decompositions by Intuition		
	camel	elephant
$\mathrm{I}_\partial\left[X_0 \cdot X_1 \to Y\right]$	0 bit	1/2 bit
$\mathrm{I}_\partial\left[X_0 \to Y \setminus X_1\right]$	1/2 bit	0 bit
$\mathrm{I}_\partial\left[X_1 \to Y \setminus X_0\right]$	1/2 bit	0 bit
$\mathrm{I}_\partial\left[X_0 X_1 \to Y\right]$	0 bit	1/2 bit

TABLE II. Camel and elephant intuitions applied to Table I's pointwise unique distribution. The camel intuition takes the view that X_0 and X_1 supply Y with unique informations, though only one of them at a time. The elephant intuition takes the view that Y provides both X_0 and X_1 with the same information, but it gets erased on the way to exactly one of them.

for the pointwise unique distribution:

X_0	X_1	Y	Pr
0	0	1	1/4
0	0	2	1/4
1	1	1	1/4
2	2	2	1/4

From the elephant's view, $\mathrm{I}[X_0 : X_1]$ is irrelevant.

V. SECRET KEY AGREEMENT

Secret key agreement is a fundamental concept within information-theoretic cryptography [5]. The central idea is that if three parties, Alice, Bob, and Eve, observe some joint probability distribution $ABE \sim p(a, b, e)$ where Alice has access only to a, Bob b, and Eve e, is it possible for Alice and Bob to agree upon a secret key of which Eve has no knowledge. The degree to which they may generate such a secret key immediately depends upon the structure of the joint distribution ABE. It also depends upon whether Alice and Bob are allowed to publicly communicate.

Concretely, consider Alice, Bob, and Eve each receiving n independent, identically distributed samples from ABE—Alice receiving A^n, Bob B^n, and Eve E^n. A *secret key agreement scheme* consists of functions f and g, as well as a protocol for public communication (h) allowing either Alice, Bob, neither, or both to communicate. In the case of a single party being permitted to communicate—say, Alice—she constructs $C = h(A^n)$ and then broadcasts it to all parties. In the case that both parties are permitted communication, they take turns constructing and broadcasting messages of the form $C_i = h_i(A^n, C_{[0\ldots i-1]})$ (Alice) and $C_i = h_i(B^n, C_{[0\ldots i-1]})$ (Bob) [6].

Formally, a secret key agreement scheme is considered R-achievable if for all $\epsilon > 0$:

$$
\begin{aligned}
K_A &\overset{(1)}{=} f(A^n, C) \\
K_B &\overset{(2)}{=} g(B^n, C) \\
p(K_A = K_B = K) &\overset{(3)}{\geq} 1 - \epsilon \\
\mathrm{I}[K : CE^n] &\overset{(4)}{\leq} \epsilon \\
\frac{1}{n}\,\mathrm{H}[K] &\overset{(5)}{\geq} R - \epsilon
\end{aligned}
$$

where (1) and (2) denote the method by which Alice and Bob construct their keys K_A and K_B, respectively, (3) states that their keys must agree with arbitrarily high probability, (4) states that the information about the key which Eve—armed with both her private information E^n as well as the public communication C—be arbitrarily small, and (5) states that the key consists of approximately R bits per sample.

The greatest rate R such that an achievable scheme exists is known as the *secret key agreement rate*. Notational variations indicate which parties are permitted to communicate. In the case that Alice and Bob are not allowed to communicate, their rate of secret key agreement is denoted $\mathrm{S}(A : B \mid\mid E)$. When only Alice is allowed to communicate their secret key agreement rate is $\mathrm{S}(A \to B \mid\mid E)$. And, similarly, if only Bob is permitted to communicate. When both Alice and Bob are allowed to communicate, their secret key agreement rate is denoted $\mathrm{S}(A \leftrightarrow B \mid\mid E)$. In this, we modified the standard notation for secret key agreement rates to emphasize which party or parties communicate.

In the case of no communication, $\mathrm{S}(A : B \mid\mid E)$ is given by [7]:

$$
\mathrm{S}(A : B \mid\mid E) = \mathrm{H}[A \curlywedge B | E] \tag{2}
$$

where $X \curlywedge Y$ denotes the Gács-Körner common random variable [8]. It is worth noting that this quantity does not vary continuously with the distribution and generically vanishes.

In the case of one-way communication, $\mathrm{S}(A \to B \mid\mid E)$ is given by [9]:

$$
\mathrm{S}(A \to B \mid\mid E) = \max \left\{ \mathrm{I}[B : K | C] - \mathrm{I}[E : K | C] \right\} \tag{3}
$$

where the maximum is taken over all variables C and K, such that the following Markov condition holds: $C - \circ - K - \circ - A - \circ - BE$. It suffices to consider K and C such that $|K| \leq |A|$ and $|C| \leq |A|^2$.

There is no such solution for $\mathrm{S}(A \leftrightarrow B \mid\mid E)$, however both upper- and lower-bounds are known [6].

Secret Key Agreement Rates	
$\mathrm{S}(X_0 : Y \mid\mid X_1)$	0 bit
$\mathrm{S}(X_1 : Y \mid\mid X_0)$	0 bit
$\mathrm{S}(Y \to X_0 \mid\mid X_1)$	0 bit
$\mathrm{S}(Y \to X_1 \mid\mid X_0)$	0 bit
$\mathrm{S}(X_0 \to Y \mid\mid X_1)$	1/2 bit
$\mathrm{S}(X_1 \to Y \mid\mid X_0)$	1/2 bit
$\mathrm{S}(X_0 \leftrightarrow Y \mid\mid X_1)$	1/2 bit
$\mathrm{S}(X_1 \leftrightarrow Y \mid\mid X_0)$	1/2 bit

TABLE III. The variety of secret sharing schemes and their rates for the pointwise unique distribution of Table I.

Let us now consider the pointwise unique distribution of Table I and the ability of X_0 and Y to agree upon a secret key while X_1 eavesdrops.[1] This can be interpreted four different ways. First, neither X_0 nor Y may be allowed to communicate. Second, only Y can communicate. Third, only X_0 is permitted to communicate. Finally, both X_0 and Y may be allowed to communicate. Note that the eavesdropper X_1 is not allowed to communicate in any secret sharing schemes here. Looking at this distribution, a general strategy becomes clear: both X_0 and Y need some scheme to determine when they agree (the 2$^{\text{nd}}$ and 4$^{\text{th}}$ rows).

Broadly, the only way in which both X_0 and Y can come to understand if they match or not is if X_0 is permitted to broadcast whether she observed a 0 or not. Therefore, in the instances where X_0 is not communicating there is no ability to agree upon a key: $\mathrm{S}(X_0 : Y \mid\mid X_1) = \mathrm{S}(Y \to X_0 \mid\mid X_1) = 0\,\text{bit}$. However, when X_0 is allowed communication a key can be agreed upon: $\mathrm{S}(X_0 \to Y \mid\mid X_1) = \mathrm{S}(X_0 \leftrightarrow Y \mid\mid X_1) = 1/2\,\text{bit}$.[2] These rates are summarized in Table III.

VI. DIRECTIONALITY, NATURALNESS, AND CONSISTENCY

We are now in a position to integrate the two intuitions with the results of secret key agreement rates. The camel intuition, with the channels $X_0 \to Y$ and $X_1 \to Y$ taking center stage, most closely aligns with the one-way secret key agreement rates $\mathrm{S}(X_0 \to Y \mid\mid X_1)$ and $\mathrm{S}(X_1 \to Y \mid\mid X_0)$. This also agrees with Section IV's quantification (compare Tables II and III):

$$\mathrm{I}_\partial\left[X_0 \to Y \setminus X_1\right] = \mathrm{S}(X_0 \to Y \mid\mid X_1) \text{ and}$$
$$\mathrm{I}_\partial\left[X_1 \to Y \setminus X_0\right] = \mathrm{S}(X_1 \to Y \mid\mid X_0) \ .$$

The elephant intuition, with its focus on the channels $Y \to X_0$ and $Y \to X_1$ is more naturally aligned with the one-way secret key agreement rates $\mathrm{S}(Y \to X_0 \mid\mid X_1)$ and $\mathrm{S}(Y \to X_0 \mid\mid X_1)$. This again accords with Section IV's quantification:

$$\mathrm{I}_\partial\left[X_0 \to Y \setminus X_1\right] = \mathrm{S}(Y \to X_0 \mid\mid X_1) \text{ and}$$
$$\mathrm{I}_\partial\left[X_1 \to Y \setminus X_0\right] = \mathrm{S}(Y \to X_1 \mid\mid X_0) \ .$$

There are, however, difficulties with these approaches.

The first difficulty concerns the camel intuition. If the one-way secret key agreement rates $\mathrm{S}(X_0 \to Y \mid\mid X_1)$ and $\mathrm{S}(X_1 \to Y \mid\mid X_0)$ are used to quantify the unique informations $\mathrm{I}_\partial\left[X_0 \to Y \setminus X_1\right]$ and $\mathrm{I}_\partial\left[X_1 \to Y \setminus X_0\right]$, respectively, the consistency relation given by Eq. (1) is not necessarily satisfied. Importantly, though, if $\mathrm{S}(Y \to X_0 \mid\mid X_1)$ and $\mathrm{S}(Y \to X_1 \mid\mid X_0)$ are used, the resulting PID is always consistent. One concludes that the elephant intuition is the more natural of the two when using one-way secret key agreement rates to quantify unique informations.

There is another difficulty. PID is defined to be agnostic to directionality. Furthermore, only one of the myriad proposed PID axioms is contingent on any inherent directionality—the Blackwell Property [13] and it is an elephant. In this sense, neither the camel nor the elephant intuitions are consistent with PID. Again relating to secret key agreement, this implies that unique informations should more closely align with either the pair $\mathrm{S}(X_0 : Y \mid\mid X_1)$ and $\mathrm{S}(X_1 : Y \mid\mid X_0)$ or with the pair $\mathrm{S}(X_0 \leftrightarrow Y \mid\mid X_1)$ and $\mathrm{S}(X_1 \leftrightarrow Y \mid\mid X_0)$; neither of which adopt any sort of directionality.

Both approaches bring their own further difficulties. On the one hand, the no-communication secret key agreement rate is not continuous in the space of distributions, whereas PID is generally considered to vary continuously. On the other hand, the two-way secret key agreement rate $\mathrm{S}(X_0 \leftrightarrow Y \mid\mid X_1)$ has no known closed-form solution, only upper and lower bounds, and so it cannot be practically computed. Furthermore and perhaps more fundamentally, whether or not the two-way secret key agreement rate results in a consistent decomposition is not known. That said, our extensive searches of examples for which the upper and lower bounds converge are encouraging—they have not resulted in any violations of Eq. (1).

[1] Secret key agreement rates have been associated with unique informations before. An upper bound on $\mathrm{S}(A \leftrightarrow B \mid\mid E)$—the intrinsic mutual information [10]—is known to not satisfy the consistency condition Eq. (1) [11]. More recently, the relationship between a particular method of quantifying unique information and one-way secret key agreement has been considered [12].

[2] It is known that $\mathrm{S}(X_0 \leftrightarrow Y \mid\mid X_1) = 1/2\,\text{bit}$ due to the convergence of upper and lower bounds in this instance.

VII. CONCLUSION

At present, a primary barrier for PID's general adoption as a useful and possibly a central tool in analyzing how complex systems store and process information is an agreement on a method to quantify its component informations. Here, we posited that one reason for disagreement stems from conflicting intuitions regarding the decomposition's operational behavior. This suggests several possibilities.

The first is that PID is inherently context-dependent and quantification depends on a notion of directionality. In this case, the elephant intuition is apparently more natural, as adopting closely related notions from cryptography results in a consistent PID. If context demands the camel intuition, though, either a noncryptographic method of quantifying unique information is needed or consistency must be enforced by augmenting the secret key agreement rate.

The second possibility suggested by our observations is that intuitions which project a directionality on the decomposition are inherently flawed and that any correct quantification must be independent of direction. Interestingly, cryptographic notions may still play a role here. Though, since there is as yet no known way to compute the two-way secret key agreement rate, its application remains open.

A final possibility is that associating secret key agreement rates with unique information is fundamentally flawed and that, ultimately, PID quantifies unique information as something distinct from the ability to agree upon a secret key.

Given that one of the main factors driving PID's creation was the need for interpretability, ensuring that the intuitions brought to bear are consistent with the quantitative values is of the utmost importance. We described three quantitative regimes, each corresponding to a specific directionality or the lack thereof. While it is possible that each can play a distinct role in the understanding of complex systems, our hope is that a single method will emerge as the most useful and accepted approach to understanding the organization of information within a joint probability distribution.

ACKNOWLEDGMENTS

All calculations herein were performed using the `dit` Python package [14]. We thank P. Banerjee, E. Olbrich, and D. Feldspar for many helpful discussions. As a faculty member, JPC thanks the Santa Fe Institute and the Telluride Science Research Center for their hospitality during visits. This material is based upon work supported by, or in part by, Foundational Questions Institute grant FQXi-RFP-1609, the U.S. Army Research Laboratory and the U.S. Army Research Office under contracts W911NF-13-1-0390 and W911NF-13-1-0340 and grant W911NF-18-1-0028, and via Intel Corporation support of CSC as an Intel Parallel Computing Center.

---

[1] P. L. Williams and R. D. Beer. Nonnegative decomposition of multivariate information. *arXiv:1004.2515*. 1

[2] C. Finn and J. T. Lizier. Pointwise partial information decomposition using the specificity and ambiguity lattices. *Entropy*, 20(4):297, 2018. 2

[3] R. A.A. Ince. Measuring multivariate redundant information with pointwise common change in surprisal. *Entropy*, 19(7):318, 2017. 2

[4] N. Bertschinger, J. Rauh, E. Olbrich, J. Jost, and N. Ay. Quantifying unique information. *Entropy*, 16(4):2161–2183, 2014. 2

[5] U. M. Maurer. Secret key agreement by public discussion from common information. *IEEE Trans. Info. Th.*, 39(3):733–742, 1993. 3

[6] A. Gohari, O. Günlü, and G. Kramer. Coding for positive rate in the source model key agreement problem. *arXiv:1709.05174*. 3

[7] E. Chitambar, B. Fortescue, and M.-H. Hsieh. The conditional common information in classical and quantum secret key distillation. *IEEE Trans. Info. Th.*, 2018. 3

[8] P. Gács and J. Körner. Common information is far less than mutual information. *Problems of Control and Information Theory*, 2(2):149–162, 1973. 3

[9] R. Ahlswede and I. Csiszár. Common randomness in information theory and cryptography. i. secret sharing. *IEEE Trans. Info. Th.*, 39(4):1121–1132, 1993. 3

[10] U. M. Maurer and S. Wolf. Unconditionally secure key agreement and the intrinsic conditional information. *IEEE Trans. Info. Th.*, 45(2):499–514, 1999. 4

[11] 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. 4

[12] P. K. Banerjee, E. Olbrich, J. Jost, and J. Rauh. Unique informations and deficiencies. *arXiv:1807.05103*. 4

[13] J. Rauh, P. Banerjee, E. Olbrich, J. Jost, and N. Bertschinger. On extractable shared information. *Entropy*, 19(7):328, 2017. 4

[14] R. G. James, C. J. Ellison, and J. P. Crutchfield. dit: a Python package for discrete information theory. *The Journal of Open Source Software*, 3(25):738, 2018. 5