

Article

Partial Information Decomposition: Redundancy as Information Bottleneck

Artemy Kolchinsky ^{1,2} 

¹ ICREA-Complex Systems Lab, Universitat Pompeu Fabra, 08003 Barcelona, Spain; artemyk@gmail.com

² Universal Biology Institute, The University of Tokyo, Tokyo 113-0033, Japan

Abstract: The partial information decomposition (PID) aims to quantify the amount of redundant information that a set of sources provides about a target. Here, we show that this goal can be formulated as a type of information bottleneck (IB) problem, termed the “redundancy bottleneck” (RB). The RB formalizes a tradeoff between prediction and compression: it extracts information from the sources that best predict the target, without revealing which source provided the information. It can be understood as a generalization of “Blackwell redundancy”, which we previously proposed as a principled measure of PID redundancy. The “RB curve” quantifies the prediction–compression tradeoff at multiple scales. This curve can also be quantified for individual sources, allowing subsets of redundant sources to be identified without combinatorial optimization. We provide an efficient iterative algorithm for computing the RB curve.

Keywords: partial information decomposition; information bottleneck; rate distortion; redundancy

1. Introduction

Many research fields that study complex systems are faced with multivariate probabilistic models and high-dimensional datasets. Prototypical examples include brain imaging data in neuroscience, gene expression data in biology, and neural networks in machine learning. In response, various information-theoretic frameworks have been developed in order to study multivariate systems in a universal manner. Here, we focus on two such frameworks, *partial information decomposition* and the *information bottleneck*.

The *partial information decomposition* (PID) considers how information about a target random variable Y is distributed among a set of source random variables X_1, \dots, X_n [1–4]. For example, in neuroscience, the sources X_1, \dots, X_n might represent the activity of n different brain regions and Y might represent a stimulus, and one may wish to understand how information about the stimulus is encoded in different brain regions. A central idea of the PID is that the information provided by the sources can exhibit *redundancy*, when the same information about Y is present in each source, and *synergy*, when information about Y is found only in the collective outcome of all sources. Moreover, it has been shown that standard information-theoretic quantities, such as entropy and mutual information, are not sufficient to quantify redundancy and synergy [1,5]. However, finding the right measures of redundancy and synergy has proven difficult. In recent work [4], we showed that such measures can be naturally defined by formalizing the analogy between set theory and information theory that lies at the heart of the PID [5]. We then proposed a measure of redundant information (*Blackwell redundancy*) that is motivated by algebraic, axiomatic, and operational considerations. We argued that Blackwell redundancy overcomes many limitations of previous proposals [4].

The *information bottleneck* (IB) [6,7] is a method for extracting compressed information from one random variable X that optimally predicts another target random variable Y . For instance, in the neuroscience example with stimulus Y and brain activity X , the IB method could be used to quantify how well the stimulus can be predicted using only one bit of



Citation: Kolchinsky, A. Partial Information Decomposition: Redundancy as Information Bottleneck. *Entropy* **2024**, *26*, 546. <https://doi.org/10.3390/e26070546>

Academic Editor: Daniel Chicharro

Received: 13 May 2024

Revised: 20 June 2024

Accepted: 21 June 2024

Published: 26 June 2024



Copyright: © 2024 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

information about brain activity. The overall tradeoff between the prediction of Y and compression of X is captured by the so-called *IB curve*. The IB method has been employed in various domains, including neuroscience [8], biology [9], and cognitive science [10]. In recent times, it has become particularly popular in machine learning applications [7,11–14].

In this paper, we demonstrate a formal connection between PID and IB. We focus in particular on the relationship between the IB and PID redundancy, leaving the connection to other PID measures (such as synergy) for future work. To begin, we show that Blackwell redundancy can be formulated as an information-theoretic constrained optimization problem. This optimization problem extracts information from the sources that best predict the target, under the constraint that the solution does not reveal which source provided the information. We then define a generalized measure of Blackwell redundancy by relaxing the constraint. Specifically, we ask how much predictive information can be extracted from the sources without revealing more than a certain number of bits about the identity of the source. Our generalization leads to an IB-type tradeoff between the prediction of the target (generalized redundancy) and compression (leakage of information about the identity of the source). We refer to the resulting optimization problem as the *redundancy bottleneck* (RB) and to the manifold of optimal solutions at different points on the prediction/compression tradeoff as the *RB curve*. We also show that the RB prediction and compression terms can be decomposed into contributions from individual sources, giving rise to an individual RB curve for each source.

Besides the intrinsic theoretical interest of unifying PID and the IB, our approach brings about several practical advantages. In particular, the RB curve offers a fine-grained analysis on PID redundancy, showing how redundant information emerges at various scales and across different sources. This fine-grained analysis can be used to uncover sets of redundant sources without performing intractable combinatorial optimization. Our approach also has numerical advantages. The original formulation of Blackwell redundancy was based on a difficult optimization problem that becomes infeasible for larger systems. By reformulating Blackwell redundancy as an IB-type problem, we are able to solve it efficiently using an iterative algorithm, even for larger systems (code available at <https://github.com/artemyk/pid-as-ib>, accessed on 12 May 2024). Finally, the RB has some attractive formal properties. For instance, unlike the original Blackwell redundancy, the RB curve is continuous in the underlying probability distributions.

This paper is organized as follows. In the next section, we provide the background on the IB, PID, and Blackwell redundancy. In Section 3, we introduce the RB, illustrate it with several examples, and discuss its formal properties. In Section 4, we introduce an iterative algorithm to solve the RB optimization problem. We discuss the implications and possible future directions in Section 5. All proofs are found in the Appendix A.

2. Background

We begin by providing relevant background on the information bottleneck, partial information decomposition, and Blackwell redundancy.

2.1. Information Bottleneck (IB)

The information bottleneck (IB) method provides a way to extract information that is present in one random variable X that is relevant for predicting another target random variable Y [6,15,16]. To do so, the IB posits a “bottleneck variable” Q that obeys the Markov condition $Q - X - Y$. This Markov condition guarantees that Q does not contain any information about Y that is not found in X . The quality of any particular choice of bottleneck variable Q is quantified via two mutual information terms: $I(X; Q)$, which decreases when Q provides a more compressed representation of X , and $I(Y; Q)$, which increases when Q allows a better prediction of Y . The IB method selects Q to maximize prediction given a constraint on compression [15–17]:

$$I_{\text{IB}}(R) = \max_{Q: Q-X-Y} I(Y; Q) \quad \text{where} \quad I(X; Q) \leq R. \quad (1)$$

The values of $I_{IB}(R)$ for different R specify the *IB curve*, which encodes the overall tradeoff between prediction and compression.

In practice, the IB curve is usually explored by considering the Lagrangian relaxation of the constrained optimization problem (1):

$$F_{IB}(\beta) := \max_Q I(Y; Q) - \frac{1}{\beta} I(X; Q) \tag{2}$$

Here, $\beta \geq 0$ is a parameter that controls the tradeoff between compression cost (favored for $\beta \rightarrow 0$) and prediction benefit (favored for $\beta \rightarrow \infty$). The advantage of the Lagrangian formulation is that it avoids the non-linear constraint in Equation (1). If the IB curve is strictly concave, then the two Equations (1) and (2) are equivalent, meaning that there is a one-to-one map between the solutions of both problems [18]. When the IB curve is not strictly concave, a modified objective such as the “squared Lagrangian” or “exponential Lagrangian” should be used instead; see Refs. [18–20] for more details.

Since the original proposal, many reformulations, generalizations, and variants of the IB have been developed [7]. Notable examples include the “conditional entropy bottleneck” (CEB) [13,21], the “multi-view IB” [22], the “distributed IB” [23], as well as a large family of objectives called the “multivariate IB” [24]. All of these approaches consider some tradeoff between two information-theoretic terms: one that quantifies the prediction of target information that should be maximized and one that quantifies the compression of unwanted information that should be minimized. We refer to an optimization that involves a tradeoff between information-theoretic prediction and compression terms as an *IB-type problem*.

2.2. Partial Information Decomposition

The PID considers how information about a target random variable Y is distributed across a set of source random variables X_1, \dots, X_n . One of the main goals of the PID is to quantify redundancy, the amount of shared information that is found in each of the individual sources. The notion of redundancy in PID was inspired by an analogy between sets and information that has re-appeared in various forms throughout the history of information theory [25–31]. Specifically, if the amount of information provided by each source is conceptualized as the size of a set, then the redundancy is conceptualized as the size of the intersection of those sets [1,4,5]. Until recently, however, this analogy was treated mostly as an informal source of intuition, rather than a formal methodology.

In a recent paper [4], we demonstrated that the terms of PID can be defined by formalizing this analogy to set theory. Recall that, in set theory, the intersection of sets A_1, \dots, A_n is defined as the largest set B that is contained in each set A_s for $s \in \{1 \dots n\}$. Thus, the size of the intersection of finite sets A_1, \dots, A_n is

$$\left| \bigcap_{s=1}^n A_s \right| = \max_B |B| \quad \text{where } B \subseteq A_s \quad \forall s \in \{1 \dots n\}.$$

We showed that PID redundancy can be defined in a similar way: the redundancy between sources X_1, \dots, X_n is the maximum mutual information in any random variable Q that is less informative about the target Y than each individual source [4]:

$$I_{\sqcap}^E := \max_Q I(Q; Y) \quad \text{where } Q \sqsubseteq X_s \quad \forall s \in \{1 \dots n\}. \tag{3}$$

The notation $Q \sqsubseteq X_s$ indicates that Q is “less informative” about the target than X_s , given some pre-specified ordering relation \sqsubseteq . The choice of the ordering relation completely determines the resulting redundancy measure I_{\sqcap}^E . We discuss possible choices in the following subsection.

We used a similar approach to define “union information”, which in turn leads to a principled measure of synergy [4]. Note that union information and redundancy are

related algebraically but not numerically; in particular, unlike in set theory, the principle of inclusion–exclusion does not always hold.

As mentioned above, here, we focus entirely on redundancy and leave the exploration of connections between IB and union information/synergy for future work.

2.3. Blackwell Redundancy

Our definition of PID redundancy (3) depends on the definition of the “less informative” relation \sqsubseteq . Although there are many relations that can be considered [25,32–35], arguably the most natural choice is the *Blackwell order*.

The Blackwell order is a preorder relation over “channels”, that is conditional distributions with full support. A channel $\kappa_{B|Y}$ is said to be less informative than $\kappa_{C|Y}$ in the sense of the Blackwell order if there exists some other channel $\kappa_{B|C}$ such that

$$\kappa_{B|Y} = \kappa_{B|C} \circ \kappa_{C|Y}. \tag{4}$$

Throughout, we use the notation \circ to indicate the composition of channels, as defined via matrix multiplication. For instance, $\kappa_{B|Y} = \kappa_{B|C} \circ \kappa_{C|Y}$ is equivalent to the statement $\kappa_{B|Y}(b|y) = \sum_c \kappa_{B|C}(b|c)\kappa_{C|Y}(c|y)$ for all b and y . Equation (4) implies that $\kappa_{B|Y}$ is less informative than $\kappa_{C|Y}$ if $\kappa_{B|Y}$ can be produced by downstream stochastic processing of the output of channel $\kappa_{C|Y}$. We use the notation

$$\kappa_{B|Y} \preceq \kappa_{C|Y}, \tag{5}$$

to indicate that $\kappa_{B|Y}$ is less Blackwell-informative than $\kappa_{C|Y}$. The Blackwell order can also be defined over random variables rather than channels. Given a target random variable Y with full support, random variable B is said to be less Blackwell-informative than random variable C , written as

$$B \preceq_Y C, \tag{6}$$

when their corresponding conditional distributions obey the Blackwell relation, $p_{B|Y} \preceq p_{C|Y}$ [36]. It is not hard to verify that any random variable B that is independent of Y is lowest under the Blackwell order, obeying $B \preceq_Y C$ for all C .

The Blackwell order plays a fundamental role in statistics, and it has an important operational characterization in decision theory [36–38]. Specifically, $p_{B|Y} \preceq p_{C|Y}$ if and only if access to channel $p_{C|Y}$ is better for every decision problem than access to channel $p_{B|Y}$. See Refs. [4,39] for details of this operational characterization and Refs. [4,36,39–42] for more discussion of the relation between the Blackwell order and the PID.

Combining the Blackwell order (6) with Equation (3) gives rise to *Blackwell redundancy* [4]. Blackwell redundancy, indicated here as I_{\sqcap} , is the maximal mutual information in any random variable that is less Blackwell-informative than each of the sources:

$$I_{\sqcap} := \max_Q I(Q; Y) \quad \text{where} \quad Q \preceq_Y X_s \quad \forall s. \tag{7}$$

The optimization is always well defined because the feasible set is not empty, given that any random variable Q that is independent of Y satisfies the constraints. (Note also that, for continuous-valued or countably infinite sources, max may need to be replaced by a sup; see also Appendix A.)

I_{\sqcap} has many attractive features as a measure of PID redundancy, and it overcomes several problems with previous approaches [4]. For instance, it can be defined for any number of sources, it uniquely satisfies a natural set of PID axioms, and it has fundamental statistical and operational interpretations. Statistically, it is the maximum information transmitted across any channel that can be produced by downstream processing of any one of the sources. Operationally, it is the maximum information that any random variable can have about Y without being able to perform better on any decision problem than any one of the sources.

As we showed [4], the optimization problem (7) can be formulated as the maximization of a convex function subject to a set of linear constraints. For a finite-dimensional system, the feasible set is a finite-dimensional polytope, and the maximum will lie on one of its extreme points; therefore, the optimization can be solved exactly by enumerating the vertices of the feasible set and choosing the best one [4]. However, this approach is limited to small systems, because the number of vertices of the feasible set can grow exponentially.

Finally, it may be argued that Blackwell redundancy is actually a measure of redundancy in the channels $p_{X_1|Y}, \dots, p_{X_n|Y}$, rather than in the random variables X_1, \dots, X_n . This is because the joint distribution over (Y, X_1, \dots, X_n) is never explicitly invoked in the definition of I_\cap ; in fact, any joint distribution is permitted as long as it is compatible with the correct marginals. (The same property holds for several other redundancy measures ([4], Table 1), and Ref. [39] even suggested this property as a requirement for any valid measure of PID redundancy.) In some cases, the joint distribution may not even exist, for instance when different sources represent mutually exclusive conditions. To use a neuroscience example, imagine that $p_{X_1|Y}$ and $p_{X_2|Y}$ represent the activity of some brain region X in response to stimulus Y , measured either in younger ($p_{X_1|Y}$) or older ($p_{X_2|Y}$) subjects. Even though there is no joint distribution over (Y, X_1, X_2) in this case, redundancy is still meaningful as the information about the stimulus that can be extracted from the brain activity of either age group. In the rest of this paper, we generally work within the channel-based interpretation of Blackwell redundancy.

3. Redundancy Bottleneck

In this section, we introduce the redundancy bottleneck (RB) and illustrate it with examples. Generally, we assume that we are provided with the marginal distribution p_Y of the target random variable Y , as well as n source channels $p_{X_1|Y}, \dots, p_{X_n|Y}$. Without loss of generality, we assume that p_Y has full support. We use calligraphic letters (like \mathcal{Y} and \mathcal{X}_s) to indicate the set of outcomes of random variables (like Y and X_s). For simplicity, we use notation appropriate for discrete-valued variables, such as in Equation (4), though most of our results also apply to continuous-valued variables.

3.1. Reformulation of Blackwell Redundancy

We first reformulate Blackwell redundancy (7) in terms of a different optimization problem. Our reformulation will make use of the random variable Y , along with two additional random variables, S and Z . The outcomes of S are the indexes of the different sources, $\mathcal{S} = \{1, \dots, n\}$. The set of outcomes of Z is the union of the outcomes of the individual sources, $\mathcal{Z} = \bigcup_{s=1}^n \mathcal{X}_s$. For example, if there are two sources with outcomes $\mathcal{X}_1 = \{0, 1\}$ and $\mathcal{X}_2 = \{0, 1, 2\}$, then $\mathcal{S} = \{1, 2\}$ and $\mathcal{Z} = \{0, 1\} \cup \{0, 1, 2\} = \{0, 1, 2\}$. The joint probability distribution over (Y, S, Z) is defined as

$$p_{YSZ}(y, s, z) = \begin{cases} p_Y(y) \nu_S(s) p_{X_s|Y}(z|y) & \text{if } z \in \mathcal{X}_s \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

In other words, y is drawn from the marginal p_Y , the source s is then drawn independently from the distribution ν_S , and finally z is drawn from the channel $p_{X_s|Y}(z|y)$ corresponding to source s . In this way, the channels corresponding to the n sources ($p_{X_1|Y}, \dots, p_{X_n|Y}$) are combined into a single conditional distribution $p_{Z|SY}$.

We treat the distribution ν_S as an arbitrary fixed parameter, and except where otherwise noted, we make no assumptions about this distribution except that it has full support. As we will see, different choices of ν_S cause the different sources to be weighed differently in the computation of the RB. We return to the question of how to determine this distribution below.

Note that, under the distribution defined in Equation (8), Y and S are independent, so

$$I(Y; S) = 0. \tag{9}$$

Actually, many of our results can be generalized to the case where there are correlations between S and Y . We leave exploration of this generalization for future work.

In addition to Y, Z , and S , we introduce another random variable Q . This random variable obeys the Markov condition $Q - (Z, S) - Y$, which ensures that Q does not contain any information about Y that is not contained in the joint outcome of Z and S . The full joint distribution over (Y, S, Z, Q) is

$$p_{YSZQ}(y, s, z, q) = p_{YSZ}(y, s, z)p_{Q|SZ}(q|s, z). \tag{10}$$

We sometimes refer to Q as the “bottleneck” random variable.

The set of joint outcomes of (S, Z) with non-zero probability is the disjoint union of the outcomes of the individual sources. For instance, in the example above with $\mathcal{X}_1 = \{0, 1\}$ and $\mathcal{X}_2 = \{0, 1, 2\}$, the set of joint outcomes of (S, Z) with non-zero probability is $\{(1, 0), (1, 1), (2, 0), (2, 1), (2, 2)\}$. Because Q depends jointly on S and Z , our results do not depend on the precise labeling of the source outcomes, e.g., they are the same if $\mathcal{X}_2 = \{0, 1, 2\}$ is relabeled as $\mathcal{X}_2 = \{2, 3, 4\}$.

Our first result shows that Blackwell redundancy can be equivalently expressed as a constrained optimization problem. Here, the optimization is over bottleneck random variables Q , i.e., over conditional distributions $p_{Q|SZ}$ in Equation (10).

Theorem 1. *Blackwell redundancy (7) can be expressed as*

$$I_{\cap} = \max_{Q:Q-(Z,S)-Y} I(Q;Y|S) \quad \text{where} \quad I(Q;S|Y) = 0. \tag{11}$$

Importantly, Theorem 1 does not depend on the choice of the distribution ν_S , as long as it has full support.

In Theorem 1, the Blackwell order constraint in Equation (7) has been replaced by an information-theoretic constraint $I(Q;S|Y) = 0$, which states that Q does not provide any information about the identity of source S , additional to that already provided by the target Y . The objective $I(Q;Y)$ has been replaced by the conditional mutual information $I(Q;Y|S)$. Actually, the objective can be equivalently written in either form, since $I(Q;Y|S) = I(Q;Y)$ given our assumptions (see the proof of Theorem 1 in the Appendix A). However, the conditional mutual information form will be useful for further generalization and decomposition, as discussed in the next sections.

3.2. Redundancy Bottleneck

To relate Blackwell redundancy to the IB, we relax the constraint in Theorem 1 by allowing the leakage of R bits of conditional information about the source S . This defines the *redundancy bottleneck* (RB) at compression rate R :

$$I_{RB}(R) := \max_{Q:Q-(Z,S)-Y} I(Q;Y|S) \quad \text{where} \quad I(Q;S|Y) \leq R. \tag{12}$$

We note that, for $R > 0$, the value of $I_{RB}(R)$ does depend on the choice of the source distribution ν_S .

Equation (12) is an IB-type problem that involves a tradeoff between prediction $I(Q;Y|S)$ and compression $I(Q;S|Y)$. The prediction term $I(Q;Y|S)$ quantifies the generalized Blackwell redundancy encoded in the bottleneck variable Q . The compression term $I(Q;S|Y)$ quantifies the amount of conditional information that the bottleneck variable leaks about the identity of the source. The set of optimal values of $(I(Q;S|Y), I(Q;Y|S))$ defines the *redundancy bottleneck curve* (RB curve) that encodes the overall tradeoff between prediction and compression.

We prove a few useful facts about the RB, starting from monotonicity and concavity.

Theorem 2. *$I_{RB}(R)$ is non-decreasing and concave as a function of R .*

Since $I_{RB}(R)$ is non-decreasing in R , the lowest RB value is achieved in the $R = 0$ regime, when it equals the Blackwell redundancy (Theorem 1):

$$I_{RB}(R) \geq I_{RB}(0) = I_{\cap}. \tag{13}$$

The largest value is achieved as $R \rightarrow \infty$, when the compression constraint vanishes. It can be shown that $I(Q; Y|S) \leq I(Z; Y|S) = I(Y; Z, S)$ using the Markov condition $Q - (Z, S) - Y$ and the data-processing inequality (see the next subsection). This upper bound is achieved by the bottleneck variable $Q = Z$. Combining implies

$$I_{RB}(R) \leq I(Z; Y|S) = \sum_s v_S(s) I(X_s; Y), \tag{14}$$

where we used the form of the distribution p_{YSZ} in Equation (8) to arrive at the last expression. The range of necessary compression rates can be restricted as $0 \leq R \leq I(Z; S|Y)$.

Next, we show that, for finite-dimensional sources, it suffices to consider finite-dimensional Q . Thus, for finite-dimensional sources, the RB problem (12) involves the maximization of a continuous objective over a compact domain, so the maximum is always achieved by some Q . (Conversely, in the more general case of infinite-dimensional sources, it may be necessary to replace max with sup in Equation (12); see Appendix A.)

Theorem 3. *For the optimization problem (12), it suffices to consider Q of cardinality $|Q| \leq \sum_s |\mathcal{X}_s| + 1$.*

Interestingly, the cardinality bound for the RB is the same as for the IB if we take $X = (Z, S)$ in Equation (1) [16,20]. It is larger than the cardinality required for Blackwell redundancy (7), where $|Q| \leq (\sum_s |\mathcal{X}_s|) - n + 1$ suffices [4].

The Lagrangian relaxation of the constrained RB problem (12) is given by

$$F_{RB}(\beta) = \max_{Q: Q-(Z,S)-Y} I(Q; Y|S) - \frac{1}{\beta} I(Q; S|Y). \tag{15}$$

The parameter β controls the tradeoff between prediction and compression. The $\beta \rightarrow 0$ limit corresponds to the $R = 0$ regime, in which case, Blackwell redundancy is recovered, while the $\beta \rightarrow \infty$ limit corresponds to the $R = \infty$ regime, when the compression constraint is removed. The RB Lagrangian (15) is often simpler to optimize than the constrained optimization (12). Moreover, when the RB curve $I_{RB}(R)$ is strictly concave, there is a one-to-one relationship between the solutions to the two optimization problems (12) and (15). However, when the RB curve is not strictly concave, there is no one-to-one relationship and the usual Lagrangian formulation is insufficient. This can be addressed by optimizing a modified objective that combines prediction and compression in a nonlinear fashion, such as the “exponential Lagrangian” [19]:

$$F_{RB}^{exp}(\beta) = \max_{Q: Q-(Z,S)-Y} I(Q; Y|S) - \frac{1}{\beta} e^{I(Q; S|Y)}. \tag{16}$$

(See an analogous analysis for IB in Refs. [18,19].)

3.3. Contributions from Different Sources

Both the RB prediction and compression terms can be decomposed into contributions from different sources, leading to an individual RB curve for each source. As we show in the examples below, this decomposition can be used to identify groups of redundant sources without having to perform intractable combinatorial optimization.

Let Q be an optimal bottleneck variable at rate R , so that $I_{\text{RB}}(R) = I(Q; Y|S)$ and $I(Q; S|Y) \leq R$. Then, the RB prediction term can be expressed as the weighted average of the prediction contributions from individual sources:

$$I_{\text{RB}}(R) = I(Q; Y|S) = \sum_s v_S(s) I(Q; Y|S = s). \tag{17}$$

Here, we introduce the specific conditional mutual information:

$$I(Q; Y|S = s) := D(p_{Q|Y, S=s} \| p_{Q|S=s}), \tag{18}$$

where $D(\cdot \| \cdot)$ is the Kullback–Leibler (KL) divergence. To build intuitions about this decomposition, we may use the Markov condition $Q - (Z, S) - Y$ to express the conditional distributions in Equation (18) as compositions of channels:

$$\begin{aligned} p_{Q|Y, S=s} &= p_{Q|Z, S=s} \circ p_{Z|Y, S=s} \\ p_{Q|S=s} &= p_{Q|Z, S=s} \circ p_{Z|S=s} \end{aligned}$$

Using the data-processing inequality for the KL divergence and Equation (8), we can then write

$$I(Q; Y|S = s) \leq D(p_{Z|Y, S=s} \| p_{Z|S=s}) = D(p_{X_s|Y} \| p_{X_s}).$$

The last term is simply the mutual information $I(Y; X_s)$ between the target and source s . Thus, the prediction contribution from source s is bounded between 0 and the mutual information provided by that source:

$$0 \leq I(Q; Y|S = s) \leq I(Y; X_s). \tag{19}$$

The difference between the mutual information and the actual prediction contribution:

$$I(Y; X_s) - I(Q; Y|S = s) \geq 0,$$

quantifies the unique information in source s . The upper bound in Equation (19) is achieved in the $R \rightarrow \infty$ limit by $Q = Z$, leading to Equation (14). Conversely, for $R = 0$, $p_{Q|Y, S=s} = p_{Q|Y}$ (from $I(Q; S|Y) = 0$) and $p_{Q|S=s} = p_Q$ (from Equation (9)), so

$$I(Q; Y|S = s) = I(Q; Y) = I_{\text{RB}}(0) = I_{\cap}.$$

Thus, when $R = 0$, the prediction contribution from each source is the same, and it is equal to the Blackwell redundancy.

The RB compression cost can also be decomposed into contributions from individual sources:

$$I(Q; S|Y) = \sum_s v_S(s) I(Q; S = s|Y). \tag{20}$$

Here, we introduce the specific conditional mutual information:

$$I(Q; S = s|Y) := D(p_{Q|Y, S=s} \| p_{Q|Y}). \tag{21}$$

The source compression terms can be related to so-called *deficiency*, a quantitative generalization of the Blackwell order. Although various versions of deficiency can be defined [43–45], here we consider the “weighted deficiency” induced by the KL divergence. For any two channels $p_{B|Y}$ and $p_{C|Y}$, it is defined as

$$\delta_D(p_{C|Y}, p_{B|Y}) := \min_{\kappa_{B|C}} D(\kappa_{B|C} \circ p_{C|Y} \| p_{B|Y}). \tag{22}$$

This measure quantifies the degree to which two channels violate the Blackwell order, vanishing when $\kappa_{B|Y} \preceq \kappa_{C|Y}$. To relate the source compression terms (20) to deficiency, observe that $p_{Q|Y,S=s} = p_{Q|Z,S=s} \circ p_{Z|Y,S=s}$ and that $p_{Z|Y,S=s} = p_{X_s|Y}$. Given Equation (21), we then have

$$I(Q; S = s|Y) \geq \delta_D(p_{X_s|Y}, p_{Q|Y}). \quad (23)$$

Thus, each source compression term is lower bounded by the deficiency between the source channel $p_{X_s|Y}$ and the bottleneck channel $p_{Q|Y}$. Furthermore, the compression constraint in the RB optimization problem (12) sets an upper bound on the deficiency of $p_{Q|Y}$ averaged across all sources.

Interestingly, several recent papers have studied the relationship between deficiency and PID redundancy in the restricted case of two sources [38,41,45–47]. To our knowledge, we provide the first link between deficiency and redundancy for the general case of multiple sources. Note also that previous work considered a slightly different definition of deficiency where the arguments of the KL divergence are reversed. Our definition of deficiency is arguably more natural, since it is more natural to minimize the KL divergence over a convex set with respect to the first argument [48].

Finally, observe that, in both decompositions (17) and (20), the source contributions are weighted by the distribution $\nu_S(s)$. Thus, the distribution ν_S determines how different sources play into the tradeoff between prediction and compression. In many cases, ν_S can be chosen as the uniform distribution. However, other choices of ν_S may be more natural in other situations. For example, in a neuroscience context where different sources correspond to different brain regions, $\nu_S(s)$ could represent the proportion of metabolic cost or neural volume assigned to region s . Alternatively, when different sources represent mutually exclusive conditions, as in the age group example mentioned at the end of Section 2, $\nu_S(s)$ might represent the frequency of condition s found in the data. Finally, it may be possible to set ν_S in an “adversarial” manner so as to maximize the resulting value of $I_{RB}(R)$ in Equation (12). We leave the exploration of this adversarial approach for future work.

3.4. Examples

We illustrate our approach using a few examples. For simplicity, in all examples, we use a uniform distribution over the sources, $\nu_S(s) = 1/n$. The numerical results are calculated using the iterative algorithm described in the next section.

Example 1. We begin by considering a very simple system, called the “UNIQUE gate” in the PID literature. Here, the target Y is binary and uniformly-distributed, $p_Y(y) = 1/2$ for $y \in \{0, 1\}$. There are two binary-valued sources, X_1 and X_2 , where the first source is a copy of the target, $p_{X_1|Y}(x_1|y) = \delta_{x_1,y}$, while the second source is an independent and uniformly-distributed bit, $p_{X_2|Y}(x_2|y) = 1/2$. Thus, source X_1 provides 1 bit of information about the target, while X_2 provides none. The Blackwell redundancy is $I_\cap = 0$ [4], because it is impossible to extract any information from the sources without revealing that this information came from X_1 .

We performed RB analysis by optimizing the RB Lagrangian $F_{RB}(\beta)$ (15) at different β . Figure 1a,b show the prediction $I(Q; Y|S)$ and compression $I(Q; S|Y)$ values for the optimal bottleneck variables Q . At small β , the prediction converges to the Blackwell redundancy, $I(Q; Y|S) = I_\cap = 0$, and there is complete loss of information about source identity, $I(Q; S|Y) = 0$. At larger β , the prediction approaches the maximum $I(Q; Y|S) = 0.5 \times I(X_1; Y) = 0.5$ bit, and compression approaches $I(Q; S|Y) = I(Z; S|Y) \approx 0.311$ bit. Figure 1c shows the RB curve, illustrating the overall tradeoff between prediction and compression.

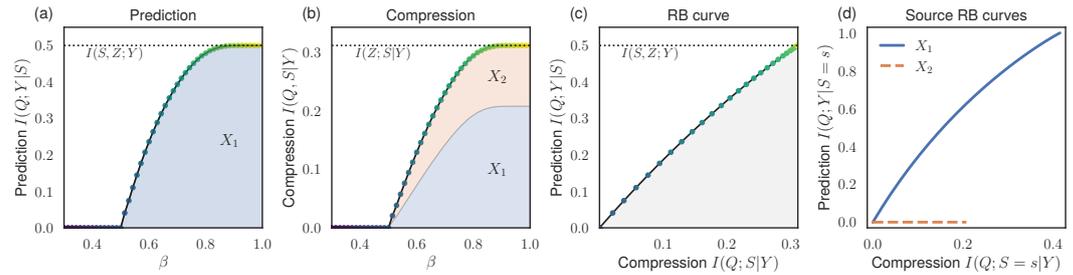


Figure 1. RB analysis for the UNIQUE gate (Example 1). (a) Prediction values found by optimizing the RB Lagrangian (15) at different β . Colored regions indicate contributions from different sources, $\nu_S(s)I(Q; Y|S = s)$ from Equation (17). For this system, only source X_1 contributes to the prediction. (b) Compression costs found by optimizing the RB Lagrangian at different β . Colored regions indicate contributions from different sources, $\nu_S(s)I(Q; S = s|Y)$ from Equation (20). (c) The RB curve shows the tradeoff between optimal compression and the prediction values; the marker colors correspond to the β values as in (a,b). All bottleneck variables Q must fall within the accessible grey region. (d) RB curves for individual sources.

In the shaded regions of Figure 1a,b, we show the additive contributions to the prediction and compression terms from the individual sources, $\nu_S(s)I(Q; Y|S = s)$ from Equation (17) and $\nu_S(s)I(Q; S = s|Y)$ from Equation (20), respectively. We also show the resulting RB curves for individual sources in Figure 1d. As expected, only source X_1 contributes to the prediction at any level of compression.

To summarize, if some information about the identity of the source can be leaked (non-zero compression cost), then improved prediction of the target is possible. At the maximum needed compression cost of 0.311, it is possible to extract 1 bit of predictive information from X_1 and 0 bits from X_2 , leading to an average of 0.5 bits of prediction.

Example 2. We now consider the “AND gate”, another well-known system from the PID literature. There are two independent and uniformly distributed binary sources, X_1 and X_2 . The target Y is also binary-valued and determined via $Y = X_1 \text{ AND } X_2$. Then, $p_Y(0) = 3/4$ and $p_Y(1) = 1/4$, and both sources have the same channel:

$$p_{X_s|Y}(x|y) = \begin{cases} 2/3 & \text{if } y = 0, x = 0 \\ 1/3 & \text{if } y = 0, x = 1 \\ 0 & \text{if } y = 1, x = 0 \\ 1 & \text{if } y = 1, x = 1 \end{cases}$$

Because the two source channels are the same, the Blackwell redundancy obeys $I_\cap = I(Y; X_1) = I(Y; X_2) = 0.311$ bits [4]. From Equations (13) and (14), we see that $I_{RB}(R) = I_\cap$ across all compression rates. In this system, all information provided by the sources is redundant, so there is no strict tradeoff between prediction and compression. The RB curve (not shown) consists of a single point, $(I(Q; Y|S), I(Q; S|Y)) = (0.311, 0)$.

Example 3. We now consider a more sophisticated example with four sources. The target is binary-valued and uniformly distributed, $p_Y(y) = 1/2$ for $y \in \{0, 1\}$. There are four binary-valued sources, where the conditional distribution of each source $s \in \{1, 2, 3, 4\}$ is a binary symmetric channel with error probability ϵ_s :

$$p_{X_s|Y}(x|y) = \begin{cases} 1 - \epsilon_s & \text{if } y = x \\ \epsilon_s & \text{if } y \neq x \end{cases} \tag{24}$$

We take $\epsilon_1 = \epsilon_2 = 0.1$, $\epsilon_3 = 0.2$, and $\epsilon_4 = 0.5$. Thus, sources X_1 and X_2 provide most information about the target; X_3 provides less information; X_4 is completely independent of the target.

We performed our RB analysis and plot the RB prediction values in Figure 2a and the compression values in Figure 2b, as found by optimizing the RB Lagrangian at different β . At small β , the prediction converges to the Blackwell redundancy, $I(Q; Y|S) = I_\cap = 0$, and there is complete loss of information about source identity, $I(Q; S|Y) = 0$. At large β , the prediction is equal to the maximum $I(Z; Y|S) \approx 0.335$ bit, and compression is equal to $I(Q; S|Y) \approx 0.104$ bit. Figure 2c shows the RB curve.

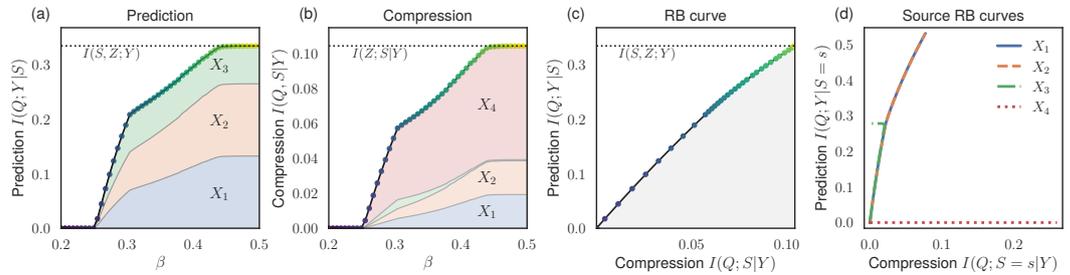


Figure 2. RB analysis for the system with 4 binary symmetric channels (Example 3). (a,b) Prediction and compression values found by optimizing the RB Lagrangian (15) at different β . Contributions from individual sources are shown as shaded regions. (c) The RB curve shows the tradeoff between optimal compression and prediction values; marker colors correspond to the β values as in (a,b). (d) RB curves for individual sources.

In Figure 2a,b, we show the additive contributions to the prediction and compression terms from the individual sources, $\nu_S(s)I(Q; Y|S = s)$ and $\nu_S(s)I(Q; S = s|Y)$, respectively, as shaded regions. We also show the resulting RB curves for individual sources in Figure 2d.

As expected, source X_4 does not contribute to the prediction at any level of compression, in accord with the fact that $I(Q; Y|S = s) \leq I(X_4; Y) = 0$. Sources X_1 and X_2 provide the same amount of prediction and compression at all points, up to the maximum $I(X_1; Y) = I(X_2; Y) \approx 0.531$. Source X_3 provides the same amount of prediction and compression as sources X_1 and X_2 , until it hits its maximum prediction $I(X_3; Y) \approx 0.278$. As shown in Figure 2d, at this point, X_3 splits off from sources X_1 and X_2 and its compression contribution decreases to 0; this is compensated by increasing the compression cost of sources X_1 and X_2 . The same behavior can also be seen in Figure 2a,b, where we see that the solutions undergo phase transitions as different optimal strategies are uncovered at increasing β . Importantly, by considering the prediction/compression contributions from the the individual sources, we can identify that sources X_1 and X_2 provide the most redundant information.

Let us comment on the somewhat surprising fact that, at larger β , the compression cost of X_3 decreases—even while its prediction contribution remains constant and the prediction contribution from X_1 and X_2 increases. At first glance, this appears counter-intuitive if one assumes that, in order to increase prediction from X_1 and X_2 , the bottleneck channel $p_{Q|Y}$ should approach $p_{X_1|Y} = p_{X_2|Y}$, thereby increasing the deficiency $\delta_D(p_{X_3|Y}, p_{Q|Y})$ and the compression cost of X_3 via the bound (23). In fact, this is not the case, because the prediction is quantified via the conditional mutual information $I(Q; Y|S)$, not the mutual information $I(Q; Y)$. Thus, it is possible that the prediction contributions from X_1 and X_2 are large, even when the bottleneck channel $p_{Q|Y}$ does not closely resemble $p_{X_1|Y} = p_{X_2|Y}$.

More generally, this example shows that it is possible for the prediction contribution from a given source to stay the same, or even increase, while its compression cost decreases. In other words, as can be seen from Figure 2d, it is possible for the RB curves of the individual sources to be non-concave and non-monotonic. It is only the overall RB curve, Figure 2c, representing the optimal prediction–compression tradeoff on average, that must be concave and monotonic.

Example 4. In our final example, the target consists of three binary spins with a uniform distribution, so $Y = (Y_1, Y_2, Y_3)$ and $p_Y(y) = 1/8$ for all y . There are three sources, each of which contains two binary spins. Sources X_1 and X_2 are both equal to the first two spins of the target Y , $X_1 = X_2 = (Y_1, Y_2)$. Source X_3 is equal to the first and last spin of the target, $X_3 = (Y_1, Y_3)$.

Each source provides $I(Y; X_s) = 2$ bits of mutual information about the target. The Blackwell redundancy I_\cap is 1 bit, reflecting the fact that there is a single binary spin that is included in all sources (Y_1).

We performed our RB analysis and plot the RB prediction values in Figure 3a and the compression values in Figure 3b, as found by optimizing the RB Lagrangian at different β . At small β , the prediction converges to the Blackwell redundancy, $I(Q; Y|S) = I_\cap = 1$, and $I(Q; S|Y) = 0$. At large β , the prediction is equal to the maximum $I(Z; Y|S) = 2$ bit, and compression is equal to $I(Z; S|Y) \approx 0.459$. Figure 3c shows the RB curve. As in the previous example, the RB curve undergoes phase transitions as different optimal strategies are uncovered at different β .

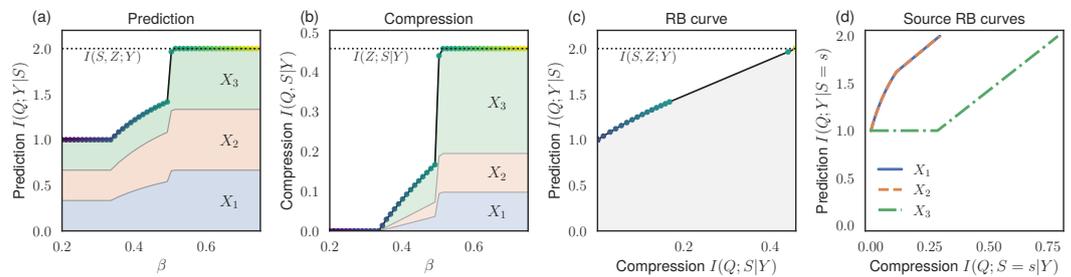


Figure 3. RB analysis for the system with a 3-spin target (Example 4). (a,b) Prediction and compression values found by optimizing the RB Lagrangian (15) at different β . Contributions from individual sources are shown as shaded regions. (c) The RB curve shows the tradeoff between optimal compression and prediction values; marker colors correspond to the β values as in (a,b). (d) RB curves for individual sources.

In Figure 3a,b, we show the additive contributions to the prediction and compression terms from the individual sources, $v_S(s)I(Q; Y|S = s)$ and $v_S(s)I(Q; S = s|Y)$, as shaded regions. We also show the resulting RB curves for individual sources in Figure 3d.

Observe that sources X_1 and X_2 provide more redundant information at a given level of compression. For instance, as shown in Figure 3d, at source compression $I(Q; S = s|Y) \approx .25$, X_1 and X_2 provide 2 bits of prediction, while X_3 provides only a single bit. This again shows how the RB source decomposition can be used for identifying sources with high levels of redundancy.

3.5. Continuity

It is known that the Blackwell redundancy I_\cap can be discontinuous as a function of the probability distribution of the target and source channels [4]. In Ref. [4], we explain the origin of this discontinuity in geometric terms and provide sufficient conditions for Blackwell redundancy to be continuous. Nonetheless, the discontinuity of I_\cap is sometimes seen as an undesired property.

On the other hand, as we show in this section, the value of RB is continuous in the probability distribution for all $R > 0$.

Theorem 4. For finite-dimensional systems and $R > 0$, $I_{RB}(R)$ is a continuous function of the probability values $p_{X_s|Y}(x|y)$, $p_Y(y)$, and $v_S(s)$.

Thus, by relaxing the compression constraint in Theorem 1, we “smooth out” the behavior of Blackwell redundancy and arrive at a continuous measure. We illustrate this using a simple example.

Example 5. We consider the COPY gate, a standard example in the PID literature. Here, there are two binary-valued sources jointly distributed according to

$$p_{X_1 X_2}(x_1, x_2) = \begin{cases} 1/2 - \epsilon/4 & \text{if } x_1 = x_2 \\ \epsilon/4 & \text{if } x_1 \neq x_2 \end{cases}$$

The parameter ϵ controls the correlation between the two sources, with perfect correlation at $\epsilon = 0$ and complete independence at $\epsilon = 1$. The target Y is a copy of the joint outcome of the two sources, $Y = (X_1, X_2)$.

It is known that Blackwell redundancy I_\cap is discontinuous for this system, jumping from $I_\cap = 1$ at $\epsilon = 0$ to $I_\cap = 0$ for $\epsilon > 0$ [4]. On the other hand, the RB function $I_{RB}(R)$ is continuous for $R > 0$. Figure 4 compares the behavior of Blackwell redundancy and RB as a function of ϵ , at $R = 0.01$ bits. In particular, it can be seen that $I_{RB}(R) = 1$ at $\epsilon = 0$ and then decays continuously as ϵ increases.

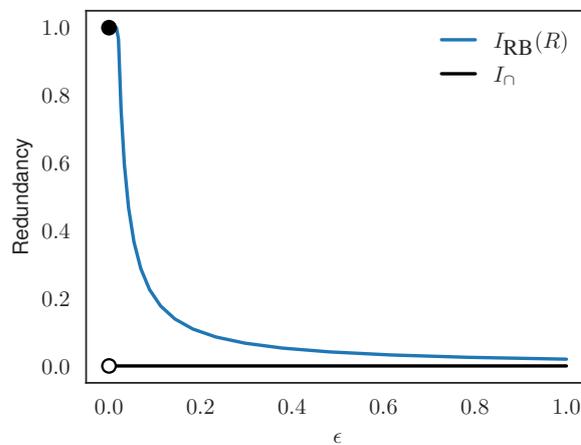


Figure 4. The RB function $I_{RB}(R)$ is continuous in the underlying probability distribution for $R > 0$, while Blackwell redundancy can be discontinuous. Here illustrated on the COPY gate, $Y = (X_1, X_2)$, as a function of correlation strength ϵ between X_1 and X_2 (perfect correlation at $\epsilon = 0$, independence at $\epsilon = 1$). Blackwell redundancy jumps from $I_\cap = 1$ at $\epsilon = 0$ to $I_\cap = 0$ at $\epsilon > 0$, while $I_{RB}(R)$ (at $R = 0.01$) decays continuously.

4. Iterative Algorithm

We provide an iterative algorithm to solve the RB optimization problem. This algorithm is conceptually similar to the Blahut–Arimoto algorithm, originally employed for rate distortion problems and later adapted to solve the original IB problem [6]. A Python implementation of our algorithm is available at <https://github.com/artemyk/pid-as-ib>; there, we also provide updated code to exactly compute Blackwell redundancy (applicable to small systems).

To begin, we consider the RB Lagrangian optimization problem, Equation (15). We rewrite this optimization problem using the KL divergence:

$$F_{RB}(\beta) = \max_{r_{Q|SZ}} D(r_{Y|QS} \| p_{Y|S}) - \frac{1}{\beta} D(r_{Q|SY} \| r_{Q|Y}). \tag{25}$$

Here, notation like $r_{Q|Y}$, $r_{Y|QS}$, etc., refers to distributions that include Q and therefore depend on the optimization variable $r_{Q|SZ}$, while notation like $p_{Y|S}$ refers to distributions that do not depend on Q and are not varied under the optimization. Every choice of conditional distribution $r_{Q|SZ}$ induces a joint distribution $r_{YSZQ} = p_{YSZ} r_{Q|SZ}$ via Equation (10).

We can rewrite the first KL term in Equation (25) as

$$D(r_{Y|QS} \| p_{Y|S}) = D(r_{Y|QS} \| p_{Y|S}) - \min_{\omega_{YSZQ}} D(r_{Y|QS} \| \omega_{Y|QS})$$

$$= \max_{\omega_{YSZQ}} \mathbb{E}_{p_{YSZ} r_{Q|SZ}} \left[\ln \frac{\omega(y|q, s)}{p(y|s)} \right].$$

where \mathbb{E} indicates the expectation, and we introduced the variational distribution ω_{YSZQ} . The maximum is achieved by $\omega_{YSZQ} = r_{YSZQ}$, which gives $\omega_{Y|QS} = r_{Y|QS}$. We rewrite the second KL term in Equation (25) as

$$D(r_{Q|SY} \| r_{Q|Y}) = D(r_{Q|SY} r_{Z|SYQ} \| r_{Q|Y} r_{Z|SYQ})$$

$$= \min_{\omega_{YSZQ}} D(r_{Q|SY} r_{Z|SYQ} \| \omega_{Q|Y} \omega_{Z|SYQ}).$$

Here, we introduce the variational distribution ω_{YSZQ} , where the minimum is achieved by $\omega_{YSZQ} = r_{YSZQ}$. The term $r_{Q|SY} r_{Z|SYQ}$ can be rewritten as

$$r(q|s, y) r(z|s, y, q) = \frac{r(z, s, y, q)}{p(s, y)} = \frac{r(q|s, z) p(z, s, y)}{p(s, y)} = r(q|s, z) p(z|s, y)$$

where we used the Markov condition $Q - (S, Z) - Y$. In this way, we separate the contribution from the conditional distribution $r_{Q|SZ}$ being optimized.

Combining the above allows us to rewrite Equation (25) as

$$F_{RB}(\beta) = \max_{r_{Q|SZ}, \omega_{YSZQ}} \mathbb{E}_{p_{YSZ} r_{Q|SZ}} \left[\ln \frac{\omega(y|q, s)}{p(y|s)} \right] - \frac{1}{\beta} D(r_{Q|SZ} p_{Z|SY} \| \omega_{Q|Y} \omega_{Z|SYQ}). \quad (26)$$

We now optimize this objective in an iterative and alternating manner with respect to $r_{Q|SZ}$ and ω_{YSZQ} . Formally, let $\mathcal{L}(r_{Q|SZ}, \omega_{YSZQ})$ refer to the objective in Equation (26). Then, starting from some initial guess $r_{Q|SZ}^{(0)}$, we generate a sequence of solutions

$$\omega_{YSZQ}^{(t+1)} = \arg \max_{\omega_{YSZQ}} \mathcal{L}(r_{Q|SZ}^{(t)}, \omega_{YSZQ}) \quad (27)$$

$$r_{Q|SZ}^{(t+1)} = \arg \max_{r_{Q|SZ}} \mathcal{L}(r_{Q|SZ}, \omega_{YSZQ}^{(t+1)}) \quad (28)$$

Each optimization problem can be solved in closed form. As already mentioned, the optimizer in Equation (27) is

$$\omega_{YSZQ}^{(t+1)} = r_{YSZQ}^{(t)} = r_{Q|SZ}^{(t)} p_{SZY}.$$

The optimization (28) can be solved by taking derivatives, giving

$$r^{(t+1)}(q|s, z) \propto e^{\sum_y p(y|s, z) \left[\beta \ln \omega^{(t)}(y|q, s) - \ln \frac{p(z|s, y)}{\omega^{(t)}(q|y) \omega^{(t)}(z|s, y, q)} \right]},$$

where the proportionality constant in \propto is fixed by normalization $\sum_q r^{(t+1)}(q|s, z) = 1$.

Each iteration increases the value of objective \mathcal{L} . Since the objective is upper bounded by $I(Z; Y|S)$, the algorithm is guaranteed to converge. However, as in the case of the original IB problem, the objective is not jointly convex in both arguments, so the algorithm may converge to a local maximum or a saddle point, rather than a global maximum. This can be partially alleviated by running the algorithm several times starting from different initial guesses $r_{Q|SZ}^{(0)}$.

When the RB is not strictly concave, it is more appropriate to optimize the exponential RB Lagrangian (16) or another objective that combines the prediction and compression

terms in a nonlinear manner [18,19]. The algorithm described above can be used with such objectives after a slight modification. For instance, for the exponential RB Lagrangian, we modify (26) as

$$F_{RB}^{exp}(\beta) = \max_{r_{Q|SZ}, \omega_{YSZQ}} \mathbb{E}_{p_{YSZ} r_{Q|SZ}} \left[\ln \frac{\omega(y|q,s)}{p(y|s)} \right] - \frac{1}{\beta} e^{D(r_{Q|SZ} p_{Z|SY} \| \omega_{Q|Y} \omega_{Z|SYQ})}. \quad (29)$$

A similar analysis as above leads to the following iterative optimization scheme:

$$\begin{aligned} \omega_{YSZQ}^{(t+1)} &= r_{Q|SZ}^{(t)} p_{SZY} \\ r^{(t+1)}(q|s,z) &\propto e^{\sum_y p(y|s,z) \left[\beta^{(t)} \ln \omega^{(t)}(y|q,s) - \ln \frac{p(z|s,y)}{\omega^{(t)}(q|y) \omega^{(t)}(z|s,y,q)} \right]}, \end{aligned}$$

where $\beta^{(t)} = \beta e^{-I_{r^{(t)}}(Q;S|Y)}$ is an effective inverse temperature. (Observe that, unlike the squared Lagrangian [18], the exponential Lagrangian leads to an effective inverse temperature $\beta^{(t)}$ that is always finite and converges to β as $I_{r^{(t)}}(Q;S|Y) \rightarrow 0$.)

When computing an entire RB curve, as in Figure 1a–c, we found good results by annealing, that is by re-using the optimal $r_{Q|SZ}$ found for one β as the initial guess at higher β . For quantifying the value of the RB function $I_{RB}(R)$ at a fixed R , as in Figure 4, we approximated $I_{RB}(R)$ via a linear interpolation of the RB prediction and compression values recovered from the RB Lagrangian at varying β .

5. Discussion

In this paper, we propose a generalization of Blackwell redundancy, termed the redundancy bottleneck (RB), formulated as an information-bottleneck-type tradeoff between prediction and compression. We studied some implications of this formulation and proposed an efficient numerical algorithm to solve the RB optimization problem.

We briefly mention some directions for future work.

The first direction concerns our iterative algorithm. The algorithm is only applicable to systems where it is possible to enumerate the outcomes of the joint distribution p_{QYSZ} . This is impractical for discrete-valued variables with very many outcomes, as well as continuous-valued variables as commonly found in statistical and machine learning settings. In future work, it would be useful to develop RB algorithms suitable for such datasets, possibly by exploiting the kinds of variational techniques that have recently gained traction in machine learning applications of IB [11–13].

The second direction would explore connections between the RB and other information-theoretic objectives for representation learning. To our knowledge, the RB problem is novel to the literature. However, it has some similarities to existing objectives, including among others the conditional entropy bottleneck [13], multi-view IB [22], and the privacy funnel and its variants [49]. Showing formal connections between these objectives would be of theoretical and practical interest, and could lead to new interpretations of the concept of PID redundancy.

Another direction would explore the relationship between RB and information-theoretic measures of causality [50,51]. In particular, if the different sources represent some mutually exclusive conditions—such as the age group example provided at the end of Section 2—then redundancy could serve as a measure of causal information flow that is invariant to the conditioning variable.

Finally, one of the central ideas of this paper is to treat the identity of the source as a random variable in its own right, which allows us to consider what information different bottleneck variables reveal about the source. In this way, we convert the search for topological or combinatorial structure in multivariate systems into an interpretable and differential information-theoretic objective. This technique may be useful in other problems that consider how information is distributed among variables in complex systems,

including other PID measures such as synergy [4], information-theoretic measures of modularity [52,53], and measures of higher-order dependency [54,55].

Funding: This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 101068029.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No research data was used in this study.

Acknowledgments: I thank Nihat Ay, Daniel Polani, Fernando Rosas, and especially, André Gomes for useful feedback. I also thank the organizers of the “Decomposing Multivariate Information in Complex Systems” (DeMICS 23) workshop at the Max Planck Institute for the Physics of Complex Systems (Dresden, Germany), which inspired this work.

Conflicts of Interest: The author declares no conflicts of interest.

Appendix A

We provide proofs of the theorems in the main text. Throughout, we use D for the Kullback–Leibler (KL) divergence, H for the Shannon entropy, and I for the mutual information.

Appendix A.1. Proof of Theorem 1

We begin by proving a slightly generalized version of Theorem 1, that is we show the equivalence between the two optimization problems:

$$I_{\cap} = \sup_Q I(Q; Y) \quad \text{where } Q \preceq_Y X_s \forall s \tag{A1}$$

$$I_{\cap} = \sup_{Q:Q-(Z,S)-Y} I(Q; Y|S) \quad \text{where } I(Q; S|Y) = 0. \tag{A2}$$

The slight generalization comes from replacing max by sup, so that the result also holds for systems with infinite-dimensional sources, where the supremum is not guaranteed to be achieved. For finite-dimensional systems, the supremum is always achieved, and we reduce to the simpler case of Equations (7) and (11).

Proof. Let V_1 indicate the supremum in Equation (A1) and V_2 the supremum in Equation (A2), given some $\nu_S(s)$ with full support. We prove that $V_1 = V_2$.

We will use that, for any distribution that has the form of Equation (10) and obeys $I(Q; S|Y) = 0$, the following holds:

$$\begin{aligned} I(Q; Y|S) &= H(Q|S) - H(Q|S, Y) \\ &= H(Q) - H(Q|Y) = I(Q; Y) \end{aligned} \tag{A3}$$

Here, we used the Markov condition $Q - Y - S$, as well as $I(Q; S) = H(Q) - H(Q|S) = 0$, as follows from Equation (9) and the data-processing inequality.

Let Q be a feasible random variable that comes within $\epsilon \geq 0$ of the objective in (A1), $I(Q; Y) \geq V_1 - \epsilon$. Define the joint distribution:

$$p_{QYSZ}(q, y, s, z) = \kappa_{Q|X_s}(q|z)p_Y(y)\nu_S(s)p_{X_s|Y}(z|y)$$

whenever $z \in \mathcal{X}_s$, otherwise $p_{QYSZ}(q, y, s, z) = 0$. Here, we used the channels $\kappa_{Q|X_s}$ associated with the Blackwell relation $Q \preceq_Y X_s$, so that $p_{Q|Y} = \kappa_{Q|X_s} \circ p_{X_s|Y}$. Under the distribution p_{QYSZ} , the Markov conditions $Q - (S, Z) - Y$ and $Q - Y - S$ hold, the latter since

$$p_{QS|Y}(q, s|y) = p_{Q|Y}(q|y)\nu_S(s). \tag{A4}$$

Therefore, this distribution has the form of Equations (8) and (10) and satisfies the constraints in Equation (A2). Using Equation (A3), we then have

$$V_1 - \epsilon \leq I(Q; Y) = I(Q; Y|S) \leq V_2.$$

Conversely, let p_{YSZQ} be a feasible joint distribution for the optimization of Equation (A2) that comes within $\epsilon \geq 0$ of the supremum, $I(Q; Y|S) \geq V_2 - \epsilon$. Using the form of this joint distribution from Equation (10), we can write

$$\begin{aligned} p_{Q|Y}(q|y) &\stackrel{(a)}{=} p_{Q|YS}(q|y, s) \\ &= \sum_z p_{Q|YSZ}(q|y, s, z) p_{Z|YS}(z|y, s) \\ &\stackrel{(b)}{=} \sum_z p_{Q|SZ}(q|z, s) p_{Z|YS}(z|y, s) \\ &\stackrel{(c)}{=} \sum_z p_{Q|SZ}(q|z, s) p_{X_s|Y}(z|y) \end{aligned}$$

In (a), we used $I(Q; S|Y) = 0$; in (b), we used $Q - (S, Z) - Y$; in (c), we used that $p_{Z|YS=s} = p_{X_s|Y}$. This implies that $p_{Q|Y} \preceq p_{X_s|Y}$ for all s . Therefore, $p_{Q|Y}$ satisfies the constraints in Equation (A1), so $I(Q; Y) \leq V_1$. Combining with Equation (A3) implies

$$V_2 - \epsilon \leq I(Q; Y|S) = I(Q; Y) \leq V_1.$$

Taking the limit $\epsilon \rightarrow 0$ gives the desired result. \square

Appendix A.2. Proof of Theorem 2

We now prove a slightly generalized version of Theorem 2. We show that the solution to the following optimization problem is non-decreasing and concave in R :

$$I_{RB}(R) := \sup_{Q:Q-(Z,S)-Y} I(Q; Y|S) \quad \text{where} \quad I(Q; S|Y) \leq R. \tag{A5}$$

The slight generalization comes from replacing max in Equation (12) by sup, so that the result also holds for systems with infinite-dimensional sources where the supremum is not guaranteed to be achieved.

Proof. $I_{RB}(R)$ is non-decreasing in R because larger R give weaker constraints (larger feasible set) in the maximization problem (A5).

To show concavity, consider any two points on the RB curve as defined by Equation (A5): $(R, I_{RB}(R))$ and $(R', I_{RB}(R'))$. For any $\epsilon > 0$, there exist Q and Q' such that

$$\begin{aligned} I(Q; S|Y) &\leq R & I(Q; Y|S) &\geq I_{RB}(R) - \epsilon \\ I(Q'; S|Y) &\leq R' & I(Q'; Y|S) &\geq I_{RB}(R') - \epsilon \end{aligned}$$

Without loss of generality, suppose that both variables have the same set of outcomes \mathcal{Q} . Then, we define a new random variable Q_λ with outcomes $\mathcal{Q}_\lambda = \{1, 2\} \times \mathcal{Q}$, as well as a family of conditional distributions parameterized by $\lambda \in [0, 1]$:

$$\begin{aligned} p_{Q_\lambda|ZS}(1, q|z, s) &= \lambda p_{Q|ZS}(q|z, s) \\ p_{Q_\lambda|ZS}(2, q|z, s) &= (1 - \lambda) p_{Q'|ZS}(q|z, s) \end{aligned}$$

In this way, we define Q_λ via a disjoint convex mixture of Q and Q' onto non-overlapping subspaces, with λ being the mixing parameter. With a bit of algebra, it can be verified that, for every λ ,

$$H(Q_\lambda|Y) = \lambda H(Q|Y) + (1 - \lambda) H(Q'|Y),$$

and similarly for $H(Q_\lambda|Y, S)$ and $H(Q_\lambda|S)$. Therefore,

$$\begin{aligned} I(Q_\lambda; S|Y) &= \lambda I(Q; S|Y) + (1 - \lambda)I(Q'; S|Y) \\ &\leq \lambda R + (1 - \lambda)R' \\ I(Q_\lambda; Y|S) &= \lambda I(Q; Y|S) + (1 - \lambda)I(Q'; Y|S) \\ &\geq \lambda I_{RB}(R) + (1 - \lambda)I_{RB}(R') - \epsilon \end{aligned}$$

Since I_{RB} is defined via a maximization, we have

$$I_{RB}(\lambda R + (1 - \lambda)R') \geq I(Q_\lambda; Y|S) \geq \lambda I_{RB}(R) + (1 - \lambda)I_{RB}(R') - \epsilon.$$

Taking the limit $\epsilon \rightarrow 0$ proves the concavity. \square

Appendix A.3. Proof of Theorem 3

Proof. We show that, for any Q that achieves $I(Q; S|Y) \leq R$, there is another Q' with cardinality $|Q'| \leq \sum_s |\mathcal{X}_s| + 1$ that satisfies $I(Q'; S|Y) \leq R$ and $I(Q'; Y|S) \geq I(Q; Y|S)$.

Consider any joint distribution p_{QSZY} from Equation (10) that achieves $I(Q; S|Y) \leq R$, and let Q be the corresponding set of outcomes of Q . Fix the corresponding conditional distribution $p_{SZ|Q}$, and note that it also determines the conditional distributions:

$$p_{YSZ|Q}(y, s, z|q) = p_{Y|SZ}(y|s, z)p_{SZ|Q}(s, z|q) \tag{A6}$$

$$= \frac{v_S(s)p_{X_s|Y}(z|y)p_Y(y)}{p_{SZ}(s, z)}p_{SZ|Q}(s, z|q) \tag{A7}$$

$$p_{Y|SQ}(y|s, q) = \frac{\sum_z p_{YSZ|Q}(y, s, z|q)}{\sum_{z, y'} p_{YSZ|Q}(y', s, z|q)} \tag{A8}$$

$$p_{S|YQ}(s|y, q) = \frac{\sum_z p_{YSZ|Q}(y, s, z|q)}{\sum_{z, s'} p_{YSZ|Q}(y, s', z|q)} \tag{A9}$$

Next, consider the following linear program:

$$V = \max_{\omega_{Q'} \in \Delta} \sum_q \omega_{Q'}(q) D(P_{Y|SQ=q} \| P_{Y|S}) \tag{A10}$$

$$\text{where } \sum_q \omega_{Q'}(q) p_{SZ|Q}(s, z|q) = p_{SZ}(s, z) \quad \forall s, z \tag{A11}$$

$$\sum_q \omega_{Q'}(q) H(S|Y, Q = q) = H(S|Y, Q) \tag{A12}$$

where Δ is the $|Q|$ -dimensional unit simplex, and we use the notation $H(S|Y, Q = q) = -\sum_{y, s} p_{YS|Q}(y, s|q) \ln p_{S|YQ}(s|y, q)$. The first set of constraints (A11) guarantees that $\omega_{Q'} p_{YSZ|Q}$ belongs to the family (10) and, in particular, that the marginal over (S, Z, Y) is $v_S(s)p_{X_s|Y}(z|y)p_Y(y)$ (see Equation (A7)). There are $\sum_s |\mathcal{X}_s|$ possible outcomes of (s, z) , but $\sum_{s, z} p_{SZ}(s, z) = 1$ by the conservation of probability. Therefore, Equation (A11) effectively imposes $\sum_s |\mathcal{X}_s| - 1$ constraints. The last constraint (A12) guarantees that $H(S|Y, Q') = H(S|Y, Q)$; hence,

$$\begin{aligned} I(Q'; S|Y) &= H(S|Y) - H(S|Y, Q') \\ &= H(S|Y) - H(S|Y, Q) = I(Q; S|Y) \leq R. \end{aligned}$$

Equation (A10) involves a maximization of a linear function over the simplex, subject to $\sum_s |\mathcal{X}_s|$ hyperplane constraints. The feasible set is compact, and the maximum is achieved at one of the extreme points of the feasible set. By Dubin's theorem [56], any extreme point of this feasible set can be expressed as a convex combination of at most $\sum_s |\mathcal{X}_s| + 1$ extreme

points of Δ . Thus, the maximum is achieved by a marginal distribution $\omega_{Q'}$ with support on at most $\sum_s |\mathcal{X}_s| + 1$ outcomes. This distribution satisfies:

$$\sum_q \omega_{Q'}(q) D(P_{Y|SQ=q} \| P_{Y|S}) \geq \sum_q p_Q(q) D(P_{Y|SQ=q} \| P_{Y|S})$$

since the actual marginal distribution p_Q is an element of the feasible set. Finally, note that

$$\begin{aligned} \sum_q \omega_{Q'}(q) D(P_{Y|SQ=q} \| P_{Y|S}) &= I(Q'; Y|S) \\ \sum_q p_Q(q) D(P_{Y|SQ=q} \| P_{Y|S}) &= I(Q; Y|S); \end{aligned}$$

therefore, $I(Q'; Y|S) \geq I(Q; Y|S)$. \square

Appendix A.4. Proof of Theorem 4

Proof. For a finite-dimensional system, we may restrict the optimization problem in Theorem 1 to Q with cardinality $|Q| \leq \sum_s |\mathcal{X}_s| + 1$ (Theorem 3). In this case, the feasible set can be restricted to a compact set, and the objective is continuous; therefore, the maximum will be achieved.

Now, consider a tuple of random variables (S, Z, Y, Q) that obey the Markov conditions $S - Y - Z$ and $Q - (S, Z) - Y$. Suppose that Q achieves the maximum in Theorem 1 for a given $R > 0$:

$$I(Q; Y|S) = I_{RB}(R), \quad I(Q; S|Y) \leq R. \tag{A13}$$

Consider also a sequence of random variables (S_k, Z_k, Y_k, Q_k) for $k = 1, 2, 3 \dots$ where each tuple has the same outcomes as (S, Z, Y, Q) and obeys the Markov conditions $S_k - Y_k - Z_k$ and $Q_k - (S_k, Z_k) - Y_k$. Let $I_{RB}^k(R)$ indicate the redundancy bottleneck defined in Theorem 1 for random variables Z_k, Y_k, S_k , and suppose that Q_k achieves the optimum for problem k :

$$I(Q_k; Y_k|S_k) = I_{RB}^k(R) \quad I(Q_k; S_k|Y_k) \leq R. \tag{A14}$$

To prove continuity, we assume that the joint distribution of (S_k, Z_k, Y_k) approaches the joint distribution of (S, Z, Y) ,

$$\lim_k \|p_{S_k Z_k Y_k} - p_{S Z Y}\|_1 = 0.$$

We first show that

$$I_{RB}(R) \geq \lim_k I_{RB}^k(R). \tag{A15}$$

First, observe that given our assumption that p_{SZ} has full support, we can always take k sufficiently large so that each $p_{S_k Z_k}$ has full support. Next, we define the random variable Q'_k that obeys the Markov condition $Q'_k - (S, Z) - Y$, with conditional distribution:

$$p_{Q'_k|SZ}(q|s, z) := p_{Q_k|S_k Z_k}(q|s, z).$$

This conditional distribution is always well-defined, given that $p_{S_k Z_k}$ has the same support as p_{SZ} . By assumption, $p_{S_k Z_k Y_k} \rightarrow p_{SZY}$; therefore,

$$p_{Q'_k SZY} - p_{Q_k S_k Z_k Y_k} \rightarrow 0. \tag{A16}$$

Conditional mutual information is (uniformly) continuous due to the (uniform) continuity of entropy (Theorem 17.3.3, [57]). Therefore,

$$0 = \lim_k [I(Q_k; S_k | Y_k) - I(Q'_k; S | Y)] \leq R - \lim_k I(Q'_k; S | Y), \tag{A17}$$

where we used Equation (A14). We also define another random variable Q_k^α , which also obeys the Markov condition $Q_k^\alpha - (S, Z) - Y$, whose conditional distribution is defined in terms of the convex mixture:

$$p_{Q_k^\alpha | SZ}(q | s, z) := \alpha_k p_{Q'_k | SX}(q | s, z) + (1 - \alpha_k) p_U(q),$$

$$\alpha_k := \min \left\{ 1, \frac{R}{I(Q'_k; S | Y)} \right\} \in [0, 1] \tag{A18}$$

Here, $p_U(q) = 1/|\mathcal{Q}|$ is a uniform distribution over an auxiliary independent random variable U with outcomes \mathcal{Q} . From the convexity of conditional mutual information [58],

$$I(Q_k^\alpha; S | Y) \leq \alpha_k I(Q'_k; S | Y) + (1 - \alpha_k) I(U; S | Y) \leq R.$$

In the last inequality, we used $I(U; S | Y) = 0$ and plugged in the definition of α_k . Observe that the random variable Q_k^α falls in the feasible set of the maximization problem that defines I_{RB} , so

$$I_{RB}(R) \geq I(Q_k^\alpha; Y | S). \tag{A19}$$

Combining Equations (A17) and (A18), and $R > 0$ implies that $\alpha_k \rightarrow 1$, so

$$p_{Q_k^\alpha; SZY} - p_{Q'_k; SZY} \rightarrow 0.$$

Combining with Equations (A14), (A16) and (A19), along with continuity of conditional mutual information, gives

$$I_{RB}(R) \geq \lim_k I(Q_k^\alpha; Y | S) = \lim_k I(Q'_k; Y | S) = \lim_k I(Q_k; Y_k | S_k) = \lim_k I_{RB}^k(R).$$

We now proceed in a similar way to prove

$$I_{RB}(R) \leq \lim_k I_{RB}^k(R). \tag{A20}$$

We define the random variable Q'' that obeys the Markov condition $Q'' - (S_k, Z_k) - Y_k$, with conditional distribution

$$p_{Q'' | S_k Z_k}(q | s, z) = p_{Q | SX}(q | s, z).$$

Since $p_{S_k Z_k Y_k} \rightarrow p_{SZY}$ by assumption,

$$p_{Q''; S_k Z_k Y_k} - p_{QSZY} \rightarrow 0. \tag{A21}$$

We then have

$$0 = \lim_k [I(Q; S | Y) - I(Q''; S_k | Y_k)] \leq R - \lim_k I(Q''; S_k | Y_k). \tag{A22}$$

where we used Equation (A13). We also define the random variable $Q^{\alpha'}$ that obeys the Markov condition $Q^{\alpha'} - (S_k, Z_k) - Y_k$, with conditional distribution

$$p_{Q^{\alpha'} | S_k Z_k}(q | s, z) = \alpha'_k p_{Q'' | S_k Z_k}(q | s, z) + (1 - \alpha'_k) p_U(q),$$

$$\alpha'_k := \min \left\{ 1, \frac{R}{I(Q''; S_k | Y_k)} \right\} \in [0, 1] \tag{A23}$$

Using the convexity of conditional mutual information, $I(U; S_k; Y_k) = 0$, and the definition of α'_k , we have

$$I(Q_k^{\alpha'}; S_k | Y_k) \leq \alpha'_k I(Q_k''; S_k | Y_k) + (1 - \alpha'_k) I(U; S_k | Y_k) \leq R.$$

Therefore, the random variable $Q_k^{\alpha'}$ falls in the feasible set of the maximization problem that defines I_{RB}^k , so

$$I_{RB}^k(R) \geq I(Q_k^{\alpha'}; Y_k | S_k). \quad (\text{A24})$$

Combining Equations (A22) and (A23), and $R > 0$ implies $\alpha'_k \rightarrow 1$; therefore,

$$p_{Q_k^{\alpha'} S_k Z_k Y_k} - p_{Q_k'' S_k Z_k Y_k} \rightarrow 0.$$

Combining this with Equations (A13), (A21) and (A24), along with the continuity of conditional mutual information, gives

$$\lim_k I_{RB}^k(R) \geq \lim_k I(Q_k^{\alpha'}; Y_k | S_k) = \lim_k I(Q_k''; Y_k | S_k) = I(Q; Y | S) = I_{RB}(R).$$

□

References

- Williams, P.L.; Beer, R.D. Nonnegative decomposition of multivariate information. *arXiv* **2010**, arXiv:1004.2515.
- Wibral, M.; Priesemann, V.; Kay, J.W.; Lizier, J.T.; Phillips, W.A. Partial information decomposition as a unified approach to the specification of neural goal functions. *Brain Cogn.* **2017**, *112*, 25–38. [[CrossRef](#)] [[PubMed](#)]
- Lizier, J.; Bertschinger, N.; Jost, J.; Wibral, M. Information decomposition of target effects from multi-source interactions: Perspectives on previous, current and future work. *Entropy* **2018**, *20*, 307. [[CrossRef](#)] [[PubMed](#)]
- Kolchinsky, A. A Novel Approach to the Partial Information Decomposition. *Entropy* **2022**, *24*, 403. [[CrossRef](#)]
- Williams, P.L. Information Dynamics: Its Theory and Application to Embodied Cognitive Systems. Ph.D. Thesis, Indiana University, Bloomington, IN, USA, 2011.
- Tishby, N.; Pereira, F.; Bialek, W. The information bottleneck method. In Proceedings of the 37th Allerton Conference on Communication, Monticello, IL, USA, 22–24 September 1999.
- Hu, S.; Lou, Z.; Yan, X.; Ye, Y. A Survey on Information Bottleneck. *IEEE Trans. Pattern Anal. Mach. Intell.* **2024**, 1–20. [[CrossRef](#)] [[PubMed](#)]
- Palmer, S.E.; Marre, O.; Berry, M.J.; Bialek, W. Predictive information in a sensory population. *Proc. Natl. Acad. Sci. USA* **2015**, *112*, 6908–6913. [[CrossRef](#)]
- Wang, Y.; Ribeiro, J.M.L.; Tiwary, P. Past–future information bottleneck for sampling molecular reaction coordinate simultaneously with thermodynamics and kinetics. *Nat. Commun.* **2019**, *10*, 3573. [[CrossRef](#)] [[PubMed](#)]
- Zaslavsky, N.; Kemp, C.; Regier, T.; Tishby, N. Efficient compression in color naming and its evolution. *Proc. Natl. Acad. Sci. USA* **2018**, *115*, 7937–7942. [[CrossRef](#)] [[PubMed](#)]
- Alemi, A.A.; Fischer, I.; Dillon, J.V.; Murphy, K. Deep variational information bottleneck. In Proceedings of the International Conference on Learning Representations (ICLR), Toulon, France, 24–26 April 2017. Available online: <https://openreview.net/forum?id=HyxQzBceg> (accessed on 12 May 2024).
- Kolchinsky, A.; Tracey, B.D.; Wolpert, D.H. Nonlinear information bottleneck. *Entropy* **2019**, *21*, 1181. [[CrossRef](#)]
- Fischer, I. The conditional entropy bottleneck. *Entropy* **2020**, *22*, 999. [[CrossRef](#)]
- Goldfeld, Z.; Polyanskiy, Y. The information bottleneck problem and its applications in machine learning. *IEEE J. Sel. Areas Inf. Theory* **2020**, *1*, 19–38. [[CrossRef](#)]
- Ahlsweide, R.; Körner, J. Source Coding with Side Information and a Converse for Degraded Broadcast Channels. *IEEE Trans. Inf. Theory* **1975**, *21*, 629–637. [[CrossRef](#)]
- Witsenhausen, H.; Wyner, A. A conditional entropy bound for a pair of discrete random variables. *IEEE Trans. Inf. Theory* **1975**, *21*, 493–501. [[CrossRef](#)]
- Gilad-Bachrach, R.; Navot, A.; Tishby, N. An Information Theoretic Tradeoff between Complexity and Accuracy. In *Learning Theory and Kernel Machines*; Goos, G., Hartmanis, J., van Leeuwen, J., Schölkopf, B., Warmuth, M.K., Eds.; Springer: Berlin/Heidelberg, Germany, 2003; Volume 2777, pp. 595–609. [[CrossRef](#)]
- Kolchinsky, A.; Tracey, B.D.; Van Kuyk, S. Caveats for information bottleneck in deterministic scenarios. In Proceedings of the International Conference on Learning Representations (ICLR), New Orleans, LA, USA, 6–9 May 2019. Available online: <https://openreview.net/forum?id=rke4HiAcY7> (accessed on 12 May 2024).
- Rodríguez Gálvez, B.; Thobaben, R.; Skoglund, M. The convex information bottleneck lagrangian. *Entropy* **2020**, *22*, 98. [[CrossRef](#)] [[PubMed](#)]

20. Bengier, E.; Asoodeh, S.; Chen, J. The cardinality bound on the information bottleneck representations is tight. In Proceedings of the 2023 IEEE International Symposium on Information Theory (ISIT), Taipei, Taiwan, 25–30 June 2023; pp. 1478–1483.
21. Geiger, B.C.; Fischer, I.S. A comparison of variational bounds for the information bottleneck functional. *Entropy* **2020**, *22*, 1229. [[CrossRef](#)] [[PubMed](#)]
22. Federici, M.; Dutta, A.; Forré, P.; Kushman, N.; Akata, Z. Learning robust representations via multi-view information bottleneck. *arXiv* **2020**, arXiv:2002.07017.
23. Murphy, K.A.; Bassett, D.S. Machine-Learning Optimized Measurements of Chaotic Dynamical Systems via the Information Bottleneck. *Phys. Rev. Lett.* **2024**, *132*, 197201. [[CrossRef](#)] [[PubMed](#)]
24. Slonim, N.; Friedman, N.; Tishby, N. Multivariate Information Bottleneck. *Neural Comput.* **2006**, *18*, 1739–1789. [[CrossRef](#)]
25. Shannon, C. The lattice theory of information. *Trans. IRE Prof. Group Inf. Theory* **1953**, *1*, 105–107. [[CrossRef](#)]
26. McGill, W. Multivariate information transmission. *Trans. IRE Prof. Group Inf. Theory* **1954**, *4*, 93–111. [[CrossRef](#)]
27. Reza, F.M. *An Introduction to Information Theory*; Dover Publications: Mineola, NY, USA, 1961.
28. Ting, H.K. On the amount of information. *Theory Probab. Its Appl.* **1962**, *7*, 439–447. [[CrossRef](#)]
29. Han, T. Linear dependence structure of the entropy space. *Inf. Control* **1975**, *29*, 337–368. [[CrossRef](#)]
30. Yeung, R.W. A new outlook on Shannon’s information measures. *IEEE Trans. Inf. Theory* **1991**, *37*, 466–474. [[CrossRef](#)]
31. Bell, A.J. The co-information lattice. In Proceedings of the Fifth International Workshop on Independent Component Analysis and Blind Signal Separation: ICA, Nara, Japan, 1–4 April 2003; Volume 2003.
32. Gomes, A.F.; Figueiredo, M.A. Orders between Channels and Implications for Partial Information Decomposition. *Entropy* **2023**, *25*, 975. [[CrossRef](#)] [[PubMed](#)]
33. Griffith, V.; Koch, C. Quantifying synergistic mutual information. In *Guided Self-Organization: Inception*; Springer: Berlin/Heidelberg, Germany, 2014; pp. 159–190.
34. Griffith, V.; Chong, E.K.; James, R.G.; Ellison, C.J.; Crutchfield, J.P. Intersection information based on common randomness. *Entropy* **2014**, *16*, 1985–2000. [[CrossRef](#)]
35. Griffith, V.; Ho, T. Quantifying redundant information in predicting a target random variable. *Entropy* **2015**, *17*, 4644–4653. [[CrossRef](#)]
36. Bertschinger, N.; Rauh, J. The Blackwell relation defines no lattice. In Proceedings of the 2014 IEEE International Symposium on Information Theory, Honolulu, HI, USA, 29 June–4 July 2014; pp. 2479–2483.
37. Blackwell, D. Equivalent comparisons of experiments. *Ann. Math. Stat.* **1953**, *24*, 265–272. [[CrossRef](#)]
38. Rauh, J.; Banerjee, P.K.; Olbrich, E.; Jost, J.; Bertschinger, N.; Wolpert, D. Coarse-Graining and the Blackwell Order. *Entropy* **2017**, *19*, 527. [[CrossRef](#)]
39. Bertschinger, N.; Rauh, J.; Olbrich, E.; Jost, J.; Ay, N. Quantifying unique information. *Entropy* **2014**, *16*, 2161–2183. [[CrossRef](#)]
40. Rauh, J.; Banerjee, P.K.; Olbrich, E.; Jost, J.; Bertschinger, N. On extractable shared information. *Entropy* **2017**, *19*, 328. [[CrossRef](#)]
41. Venkatesh, P.; Schamberg, G. Partial information decomposition via deficiency for multivariate gaussians. In Proceedings of the 2022 IEEE International Symposium on Information Theory (ISIT), Espoo, Finland, 26 June–1 July 2022; pp. 2892–2897.
42. Mages, T.; Anastasiadi, E.; Rohner, C. Non-Negative Decomposition of Multivariate Information: From Minimum to Blackwell Specific Information. *Entropy* **2024**, *26*, 424. [[CrossRef](#)]
43. Le Cam, L. Sufficiency and approximate sufficiency. *Ann. Math. Stat.* **1964**, *35*, 1419–1455. [[CrossRef](#)]
44. Raginsky, M. Shannon meets Blackwell and Le Cam: Channels, codes, and statistical experiments. In Proceedings of the 2011 IEEE International Symposium on Information Theory, St. Petersburg, Russia, 31 July–5 August 2011; pp. 1220–1224.
45. Banerjee, P.K.; Olbrich, E.; Jost, J.; Rauh, J. Unique informations and deficiencies. In Proceedings of the 2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton), Monticello, IL, USA, 2–5 October 2018; pp. 32–38.
46. Banerjee, P.K.; Montufar, G. The Variational Deficiency Bottleneck. In Proceedings of the 2020 International Joint Conference on Neural Networks (IJCNN), Glasgow, UK, 19–24 July 2020; pp. 1–8. [[CrossRef](#)]
47. Venkatesh, P.; Gurushankar, K.; Schamberg, G. Capturing and Interpreting Unique Information. In Proceedings of the 2023 IEEE International Symposium on Information Theory (ISIT), Taipei, Taiwan, 25–30 June 2023; pp. 2631–2636. [[CrossRef](#)]
48. Csiszár, I.; Matus, F. Information projections revisited. *IEEE Trans. Inf. Theory* **2003**, *49*, 1474–1490. [[CrossRef](#)]
49. Makhdoumi, A.; Salamati, S.; Fawaz, N.; Médard, M. From the information bottleneck to the privacy funnel. In Proceedings of the 2014 IEEE Information Theory Workshop (ITW 2014), Hobart, Australia, 2–5 November 2014; pp. 501–505.
50. Janzing, D.; Balduzzi, D.; Grosse-Wentrup, M.; Schölkopf, B. Quantifying causal influences. *Ann. Stat.* **2013**, *41*, 2324–2358. [[CrossRef](#)]
51. Ay, N. Confounding ghost channels and causality: A new approach to causal information flows. *Vietnam. J. Math.* **2021**, *49*, 547–576. [[CrossRef](#)]
52. Kolchinsky, A.; Rocha, L.M. Prediction and modularity in dynamical systems. In Proceedings of the European Conference on Artificial Life (ECAL), Paris, France, 8–12 August 2011. Available online: <https://direct.mit.edu/isal/proceedings/ecal2011/23/65/111139> (accessed on 12 May 2024).
53. Hidaka, S.; Oizumi, M. Fast and exact search for the partition with minimal information loss. *PLoS ONE* **2018**, *13*, e0201126. [[CrossRef](#)] [[PubMed](#)]
54. Rosas, F.; Ntranos, V.; Ellison, C.J.; Pollin, S.; Verhelst, M. Understanding interdependency through complex information sharing. *Entropy* **2016**, *18*, 38. [[CrossRef](#)]

55. Rosas, F.E.; Mediano, P.A.; Gastpar, M.; Jensen, H.J. Quantifying high-order interdependencies via multivariate extensions of the mutual information. *Phys. Rev. E* **2019**, *100*, 032305. [[CrossRef](#)]
56. Dubins, L.E. On extreme points of convex sets. *J. Math. Anal. Appl.* **1962**, *5*, 237–244. [[CrossRef](#)]
57. Cover, T.M.; Thomas, J.A. *Elements of Information Theory*; John Wiley & Sons: Hoboken, NJ, USA, 2006.
58. Timo, R.; Grant, A.; Kramer, G. Lossy broadcasting with complementary side information. *IEEE Trans. Inf. Theory* **2012**, *59*, 104–131. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.