

# Partial Information Decomposition for Continuous Variables based on Shared Exclusions: Analytical Formulation and Estimation

David A. Ehrlich,<sup>1,\*</sup> Kyle Schick-Poland,<sup>1,2,†</sup> Abdullah Makkeh,<sup>1</sup>  
Felix Lanfermann,<sup>2</sup> Patricia Wollstadt,<sup>2</sup> and Michael Wibral<sup>1</sup>

<sup>1</sup>*Göttingen Campus Institute for Dynamics of Biological Networks  
Universität Göttingen  
Göttingen*

<sup>2</sup>*Honda Research Institute Europe GmbH, Offenbach am Main  
(Dated: March 28, 2024)*

Describing statistical dependencies is foundational to empirical scientific research. For uncovering intricate and possibly non-linear dependencies between a single target variable and several source variables within a system, a principled and versatile framework can be found in the theory of Partial Information Decomposition (PID). Nevertheless, the majority of existing PID measures are restricted to categorical variables, while many systems of interest in science are continuous. In this paper, we present a novel analytic formulation for continuous redundancy—a generalization of mutual information—drawing inspiration from the concept of shared exclusions in probability space as in the discrete PID definition of  $I_{\cap}^{\text{sx}}$ . Furthermore, we introduce a nearest-neighbor based estimator for continuous PID, and showcase its effectiveness by applying it to a simulated energy management system provided by the Honda Research Institute Europe GmbH. This work bridges the gap between the measure-theoretically postulated existence proofs for a continuous  $I_{\cap}^{\text{sx}}$  and its practical application to real-world scientific problems.

## I. INTRODUCTION

The pursuit of discovering and quantifying dependencies between different experimental variables lies at the heart of empirical and data-driven scientific research. However, conventional tools such as correlation analysis might fail to capture relevant associations if these associations are non-linear. This ascertains the need for a comprehensive framework capable of capturing both linear and non-linear dependencies. Such a framework can be found in information theory, which was originally introduced for the analysis of communication channels by Claude Shannon [1] and has since become a general-purpose approach to data analysis.

A basic quantity of information theory is the mutual information [1, 2]

$$I(T : S) = \sum_{t,s} p_{T,S}(t,s) \log_2 \frac{p_T(t) p_S(s)}{p_{T,S}(t,s)},$$

which captures the dependency between two variables  $T$  and  $S$  by quantifying how much the uncertainty about one variable can be reduced by observing the respective other, and which is used as a general measure of dependency between two variables in many applications [3].

In many situations, however, there are variables of interest  $T$  which depend not only on one but multiple variables  $\mathbf{S} = (S_1, \dots, S_n)$ , and where the distinction between

the  $S_i$  is conceptually important. Hence arises the need to unravel the specific contributions of these source variables  $S_i$  to the information contained in the target  $T$ . Note that this information may be distributed among the source variables in very different ways: While parts of the information might exclusively be available from one variable but not from others (“unique information”), other parts can be obtained from either one or another (“redundant information”) and finally some pieces of information might only be revealed when considering multiple sources simultaneously (“synergistic information”). Identifying and quantifying these information *atoms* is the subject of Partial Information Decomposition (PID) [4, 5], which has been gaining popularity as a tool for the detailed information-theoretic analysis of variable dependencies, for example in neuroscience [6–12], machine learning [13–20], engineering [21], sociology [22], linguistics [23] and climatology [24].

The PID framework was originally envisioned by Williams and Beer [4], who showed that the information atoms of unique, redundant, and synergistic information could not be defined using classical information-theoretic terms, but required the introduction of novel axioms. Based on their proposed set of axioms, the authors introduced a measure of redundant information as a generalization of mutual information, that allowed to quantify the PID atoms. Following the original work of Williams and Beer, a series of different quantification schemes for PIDs has been proposed, which mostly adopt the proposed overall structure and either adopt the original set of axioms or modify it [25–29]. In particular, most works propose alternative measures of redundancy, which are grounded in a multitude of different and partly non-compatible desiderata [25, 26].

However, while most variables of interest encountered

\* Equally contributing first author; [davidalexander.ehrlich@uni-goettingen.de](mailto:davidalexander.ehrlich@uni-goettingen.de)

† Equally contributing first author; [kyle.schick-poland@uni-goettingen.de](mailto:kyle.schick-poland@uni-goettingen.de)

in science and engineering are continuous-valued, most of the PID measures suggested to this date are only valid for categorical random variables, i.e., variables with a discrete alphabet. Only a few continuous PID measures have been introduced in recent time, which each have distinct operational interpretations: Barrett [30] were among the first to introduce a fully continuous PID for Gaussian random variables. Drawing on concepts of game theory, the measure by Ince [28] based on common changes in surprisal can be straightforwardly applied to continuous variables. Similarly, the measure by Kolchinsky [29], which is build on the notion of the Blackwell order, and the redundant information neural estimation by Kleinman et al. [31] also transfer canonically to the continuous domain. Finally, the measure introduced by Pakman et al. [32] provides a continuous PID definition based on the decision-theoretic discrete concepts introduced by Bertschinger et al. [25], while Milzman and Lyzinski [33] introduce continuous generalizations of  $I_{\cap}^{\min}$  [4] and  $I_{\cap}^{\text{PM}}$  [27].

Overall, these measures all draw on different desiderata which make the measures applicable to different situations that fit their operational interpretation. If the system is best described by a memoryless agent acting on an ensemble of states, the shared-exclusion PID measure  $I_{\cap}^{\text{sx}}$  is most suitable, which has been introduced by Makkeh et al. [26] for categorical variables. Recently, Schick-Poland et al. [34] showed that an extension of this measure to continuous variables is possible, by proving that a continuous version of the measure is measure-theoretically well-defined. However, since the provided existence proofs are not fully constructive, neither an analytical definition nor a practical estimation procedure for such a continuous  $I_{\cap}^{\text{sx}}$  has been suggested to this date. With this paper, we fill the gap for a continuous PID measure based on the shared-exclusion principle which draws only on concepts of probability and information theory.

The main contributions of this paper are (1) the introduction of a tractable analytical PID definition inspired by the measure-theoretic definition of continuous  $I_{\cap}^{\text{sx}}$  from Schick-Poland et al. [34] and its application to simple theoretical examples in Section II B, (2) the development of an estimator for the associated redundancy measure, which draws on concepts of the k-nearest-neighbours based estimator for mutual information by Kraskov et al. [35] in Section II C and (3) the demonstration of the efficacy of our continuous  $I_{\cap}^{\text{sx}}$  measure in uncovering variable dependencies in data from an energy management system in Section IV.

## II. THE PID PROBLEM FOR TWO SOURCE VARIABLES

To provide a succinct account of the PID problem and its intuitive meaning, we start with a discussion of PID using only two source variables—referred to as *bivariate* PID—in this section. A comprehensive treatment of the

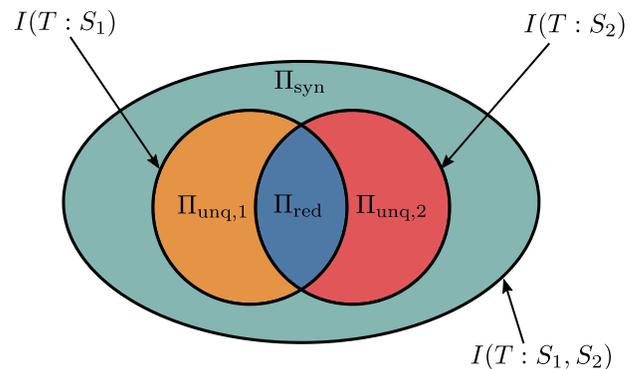


FIG. 1. **Partial Information Decomposition reveals the intricate interdependencies between classical information-theoretic quantities** [4]. Using PID, the mutual information  $I(T : S_1, S_2)$  (large oval) and the marginal information terms  $I(T : S_1)$  and  $I(T : S_2)$  (circles) can simultaneously be dissected into a total of four information atoms (colored areas): The unique information that  $S_1$  carries about  $T$ ,  $\Pi_{\text{unq},1}$  (orange), and likewise the unique information of  $S_2$ ,  $\Pi_{\text{unq},2}$  (red); the redundant information of  $S_1$  and  $S_2$  about  $T$ ,  $\Pi_{\text{red}}$  (blue), and the synergistic information exclusively contained in  $S_1$ , and  $S_2$ ,  $\Pi_{\text{syn}}$  (teal). The marginal mutual information terms  $I(T : S_i)$  can be constructed from the *redundant* and the corresponding *unique* atoms, while the joint mutual information  $I(T : S_1, S_2)$  contains of all four atoms.

general multivariate PID is deferred to Section III.

The mutual information  $I(T : \mathbf{S})$  that two source variables  $\mathbf{S} = (S_1, S_2)$  hold about a target variable  $T$  can be conveyed by the source variables in four distinct ways (see Figure 1): Some parts of the information might be *unique* to either  $S_1$  or  $S_2$ , meaning they are inaccessible from the other variable. Other parts may be *redundant* between both, meaning that they can be obtained from either source, while finally some parts only become accessible *synergistically* when both sources are observed together—like a cipher text and its corresponding password are jointly necessary to recover the plain text. We denote these PID atoms by  $\Pi_{\text{unq},1}$  and  $\Pi_{\text{unq},2}$  for the unique information, by  $\Pi_{\text{red}}$  for the redundant information, and by  $\Pi_{\text{syn}}$  for the synergistic information.

To understand how the joint mutual information  $I(T : \mathbf{S})$  relates to the PID atoms, a first step is to look at the *marginal* mutual information terms  $I(T : S_1)$  and  $I(T : S_2)$ , which quantify the statistical dependence between the target  $T$  and a single source variable. These two information quantities, however, need not be disjoint, because each mutual information, while containing the respective unique atoms  $\Pi_{\text{unq},1}$  or  $\Pi_{\text{unq},2}$ , also contain the *redundant* atom  $\Pi_{\text{red}}$ . Finally, besides the two unique and the redundant term, the joint mutual information  $I(T : \mathbf{S})$  also contains the synergistic atom  $\Pi_{\text{syn}}$ , which is part of neither of the marginal mutual information terms because it is inaccessible from any single source alone. These requirements lead to the so-called

consistency equations for two source variables [4, 5]:

$$\begin{aligned} I(T : S_1) &= \Pi_{\text{red}} + \Pi_{\text{unq},1} \\ I(T : S_2) &= \Pi_{\text{red}} + \Pi_{\text{unq},2} \\ I(T : \mathbf{S}) &= \Pi_{\text{red}} + \Pi_{\text{unq},1} + \Pi_{\text{unq},2} + \Pi_{\text{syn}}. \end{aligned} \quad (1)$$

Note that, in the bivariate case, there are four unknown PID atoms, yet only three consistency equations providing constraints. Hence, the system is underdetermined and has to be resolved by defining an additional quantity. Typically, this is done by explicitly quantifying the redundancy  $\Pi_{\text{red}}$  (e.g., [4, 26, 27]). For more than two source variables, a generalized notion of redundancy  $I_{\cap}$  needs to be introduced which captures the redundancy between *sets* of source variables (see Section III A).

In recent years, a variety of different redundancy measures has been suggested which fulfill a number of partly incompatible assumptions and have different operationalizations, e.g. from decision [25] or game theory [28]. In this work, we focus on the  $I_{\cap}^{\text{sx}}$  redundancy measure [26], which is rooted in the probability- and information-theoretic considerations of Fano [2].

#### A. The $I_{\cap}^{\text{sx}}$ redundancy measure and the measure-theoretic definition of continuous $I_{\cap}^{\text{sx}}$

The  $I_{\cap}^{\text{sx}}$  redundancy measure has originally been defined for discrete random variables [26]. To understand the intuition behind the definition, consider a discrete target random variable  $T$  and two discrete source random variables  $S_1$  and  $S_2$ , taking on realizations  $t$  and  $(s_1, s_2)$ , respectively. Notice that the joint mutual information  $I(T : S_1, S_2)$  between the target and both sources can be written as an expectation value of *local* mutual information terms  $i(t : s_1, s_2)$  as [3]

$$I(T : S_1, S_2) = \sum_{t, s_1, s_2} p_{T,S}(t, s_1, s_2) i(t : s_1, s_2), \quad (2)$$

where

$$i(t : s_1, s_2) = \log_2 \left[ \frac{p_{T|S}(t|s_1, s_2)}{p_T(t)} \right] \quad (3a)$$

$$= \log_2 \left[ \frac{p_{T|S}(t|S_1 = s_1 \wedge S_2 = s_2)}{p_T(t)} \right]. \quad (3b)$$

The local mutual information reflects how we need to update our beliefs about the occurrence of the target event  $T = t$  when we know both source events  $S_1 = s_1$  and  $S_2 = s_2$  have occurred [2]: If knowledge of the event  $S_1 = s_1$  and  $S_2 = s_2$  happening increases the likelihood of guessing the correct target event  $T = t$  (i.e.,  $p_{T|S_1, S_2}(t|s_1, s_2) > p_T(t)$ ), the event is called *informative* and  $i(t : s_1, s_2)$  is positive. Vice versa, if one is less likely to guess the target event  $T = t$  after knowing the source events (i.e.,  $p_{T|S_1, S_2}(t|s_1, s_2) < p_T(t)$ ), the event is called *misinformative* and  $i(t : s_1, s_2) < 0$ .

In the second transformation step, i.e., in Equation (3b), we explicitly state the logical statement whose observation leads us to update our beliefs about  $T = t$ . Makkeh et al. [26] suggest that building on this basic idea, a measure for redundancy can be constructed by instead measuring how our beliefs about  $T = t$  are updated if we only know that *either*  $S_1 = s_1$  or  $S_2 = s_2$  have occurred, which leads to the local expression

$$i_{\cap}^{\text{sx}}(t : s_1; s_2) = \log_2 \left[ \frac{p_{T|S}(t|S_1 = s_1 \vee S_2 = s_2)}{p_T(t)} \right]. \quad (4)$$

Averaging this quantity with weights  $p(t, s_1, s_2)$  as in Equation (2) then yields a *global* measure of redundant information between sources  $S_1$  and  $S_2$  about  $T$ , denoted by  $I_{\cap}^{\text{sx}}$ , i.e., [26]

$$I_{\cap}^{\text{sx}} = \sum_{s_1, s_2, t} p_{T,S}(t, s_1, s_2) i_{\cap}^{\text{sx}}(t : s_1; s_2). \quad (5)$$

This redundancy measure is symmetric with respect to swapping the sources  $S_1$  and  $S_2$  and invariant towards any arbitrary bijective relabelling of the discrete realizations of any of the random variables. Furthermore, it fulfills a target chain rule and is differentiable with respect to the underlying probability distribution [26].

Equivalently, Equation (4) can be derived from the regions of probability space which  $S_1 = s_1$  and  $S_2 = s_2$  jointly render impossible (see Appendix A for a more visual explanation), which is why the PID definition is referred to as “shared exclusion PID” or “SxPID” for short. For more than two source variables, the logical expression in the numerator has to be extended to account for disjunctions between sets of source events occurring simultaneously (see Section III B).

In previous work, Schick-Poland et al. [34] showed that the  $I_{\cap}^{\text{sx}}$  redundancy measure can be generalized to continuous source and target variables. In particular, the authors proved the existence of a measure-theoretic generalization of  $I_{\cap}^{\text{sx}}$ , which can handle arbitrarily many source variables while being invariant under isomorphic transformations and differentiable with respect to the joint distribution of sources and target. The proof of the existence of this generalization builds on established theorems from probability and measure theory. It is shown that the derived measure inherits all the above mentioned desirable properties of the discrete  $I_{\cap}^{\text{sx}}$  measure and is applicable to any possible discrete or continuous variable settings.

However, due to the provided proofs not being entirely constructive, a concrete analytic formulation of this continuous  $I_{\cap}^{\text{sx}}$  measure remains elusive. While proven to exist from measure-theoretic principles, the way to construct a transformation-invariant shared-exclusion redundancy measure is still not explicit. Therefore, in the next section we provide an analytical definition for a purely continuous  $I_{\cap}^{\text{sx}}$  measure which, while not adhering to all properties described by Schick-Poland et al. [34], provides a practical and interpretable shared-exclusion

based continuous PID definition. In the following, we show how such a formulation can be obtained, present an estimation procedure and discuss its relation to the measure-theoretic derivation.

## B. An analytical formulation of a continuous $I_{\cap}^{\text{sx}}$ measure

### 1. Definition

To make a continuous  $I_{\cap}^{\text{sx}}$  measure available to practitioners, an important first step is to find an analytic formulation. In this section, we present such a formulation and show how it can be computed numerically for a given probability density function. Since in most practical applications one has only access to a finite sample of data from an unknown distribution, we introduce an estimator for this redundancy measure based on k-nearest-neighbour distances in Section II C.

The main difficulty for finding an analytical formulation for the continuous  $I_{\cap}^{\text{sx}}$  measure is due to the disjunction (logical “or”) in Equation (4). While conjunctions (logical “and”) are ubiquitous in classical continuous information theory and correspond to simple joint probability densities, disjunctive statements have no established counterpart in the continuous regime. In the following we will introduce an intuitive notion of how such a quantity can be defined, while we defer a more rigorous derivation to Appendix B.

For discrete random variables, the probability of the event  $S_1 = s_1 \vee S_2 = s_2$  is given by the inclusion-exclusion-rule as

$$\begin{aligned} & \mathbb{P}(S_1 = s_1 \vee S_2 = s_2) \\ &= p_{S_1}(s_1) + p_{S_2}(s_2) - p_{S_1, S_2}(s_1, s_2), \end{aligned} \quad (6)$$

where the last term  $p_{S_1, S_2}(s_1, s_2)$  is the probability of both  $S_1 = s_1$  and  $S_2 = s_2$  occurring, which is double-counted in the sum of marginals and thus needs to be subtracted once.

To make the leap from discrete to continuous variables, we follow the steps of Jaynes’ derivation of the differential entropy [36] by reinterpreting the probability mass functions  $p_{S_i}$  of Equation (6) as binned distributions of an underlying continuous probability density  $f_{S_i}$ . That is, we assume the density  $f_{S_i}$  to be given, such that the probability mass of a specific binned realization can be calculated as

$$p(s_i) = \int_{U_{s_i}} ds'_i f_{S_i}(s'_i) \approx |U_{s_i}| f_{S_i}(s_i),$$

where  $\{U_{s_i}\}_{s_i}$  is a partition of the support of  $f_{S_i}$  and  $|U_{s_i}|$  is the area of one of the sets of the partition. The probability of the binned disjunction, i.e., the probability of either  $S_1$  falling into the bin  $U_{s_1}$  or  $S_2$  falling into  $U_{s_2}$

then becomes

$$\begin{aligned} & \mathbb{P}(S_1 \in U_{s_1} \vee S_2 \in U_{s_2}) \\ & \approx |U_{s_1}| f_{S_1}(s_1) + |U_{s_2}| f_{S_2}(s_2) \\ & \quad - |U_{s_1}| |U_{s_2}| f_{S_1, S_2}(s_1, s_2). \end{aligned} \quad (7)$$

To get to a fully continuous description, one now has to take the limit of the partitionings  $\{U_{s_1}\}_{s_1}$  and  $\{U_{s_2}\}_{s_2}$  becoming increasingly fine. If this limit is sufficiently well-behaved, the local density of the evaluation points  $s_i$  can in the limit be represented by a function (referred to by Jaynes as an “invariant measure” [37])  $m(s_i)$  such that

$$\lim_{n \rightarrow \infty} (n|U_{s_i}|) = \frac{1}{m(s_i)},$$

with

$$\lim_{n \rightarrow \infty} \frac{1}{n} (\text{number of points in } A) = \int_A ds_i m(s_i)$$

for arbitrary open sets  $A$  in the sample space of  $S_i$ . Specifically, the measure  $m_{S_i}$  represents the density of evaluation points, such that a more dense collection of evaluation points causes the corresponding binning intervals to shrink more rapidly. From this, we arrive at the formulation

$$\begin{aligned} & \mathbb{P}(S_1 \in U_{s_1} \vee S_2 \in U_{s_2}) \\ & \approx \frac{f_{S_1}(s_1)}{n m_{S_1}(s_1)} + \frac{f_{S_2}(s_2)}{n m_{S_2}(s_2)} - \frac{f_{S_1, S_2}(s_1, s_2)}{n^2 m_{S_2}(s_2) m_{S_1}(s_1)} \end{aligned} \quad (8)$$

for the probability mass of a sufficiently small disjunction in the continuous regime. Note that, while the whole expression naturally tends to zero in the limit  $n \rightarrow \infty$ , the subtracted term scales with the number of discrete points squared,  $n^2$ . Thus, it will become negligible compared to the two terms before under the assumption that  $f_{S_1, S_2}(s_1, s_2)$  is finite everywhere.

Nonetheless, neglecting the third term still leaves a choice of how the invariant measures  $m_{S_1}$  and  $m_{S_2}$  are chosen with respect to each other. Intuitively, this freedom is reflective of the *relative scale* with which the variables are considered: Since ultimately, probability densities are defined by what finite probability masses they integrate to in small neighbourhoods, the relative scale between two variables defines how to compare the neighborhoods for two variables. If the two marginal distributions of the sources are the same, a canonical choice is to take the same partitioning for both variables. On the other hand, if the variables come from very different distributions, the invariant measures  $m_{S_i}$  need to be selected to account for that.

On a practical note, the division by the invariant measures  $m_{S_i}$  can be absorbed into the definition of the probability densities  $f_{S_i}$  by a suitable variable pretransformation. How the variables should be preprocessed, if at all, depends on the nature of the variables in question. Firstly, if the variables describe equivalent source

processes, no pretransformation is necessary, and linear pretransformations such as standardization do not make a difference (see Appendix C for a more detailed comparison). On the other hand, if two variables measure the same physical quantity but with different scales, e.g., body height measured in centimeters versus in feet, a reasonable preprocessing step would be the rescaling of one of the variables or—if the variables come from a similar distribution—a standardization of both. Alternatively, if the measurement precision, i.e., the relative pointwise deviation, can be assumed to be proportional to the probability density at that point, the original probability densities  $f$  may be replaced with the copula densities  $c$  [38], which furthermore happens to make the measure invariant under invertible mappings of individual variables before this step.

Finally, these considerations lead to the following definition for the probability of a continuous disjunction for two random variables  $S_1$  and  $S_2$ , which have possibly been preprocessed beforehand.

**Definition 1.** *Let  $S_1, S_2$  be continuous random variables with associated densities  $f_{S_1}, f_{S_2}$ . Then we define the quasi-density of the disjunctive logic statement  $S_1 = s_1 \vee S_2 = s_2$  to be*

$$f_{S_1 \vee S_2}(s_1 \vee s_2) := f_{S_1}(s_1) + f_{S_2}(s_2).$$

Note that in the above definition, it is assumed that the precision of one variable does not depend on the value of the other variable and that the relative precision between the two variables is constant. Furthermore, while the local density of a disjunction as defined above does not fulfill all properties of a density function in the mathematical sense (particularly, it does not integrate to one), it nevertheless serves as a local approximate proportionality of probability in relation to magnitude of area in state space.

Given this definition for a continuous disjunction, we can write down an analytical formulation for the continuous local shared-exclusion redundancy.

**Definition 2.** *Let  $S_1, S_2$  be continuous random variables with associated densities  $f_{S_1}, f_{S_2}$ .*

*Then we define the continuous local shared-exclusion redundancy for two sources as*

$$\begin{aligned} i_{\cap}^{\text{sx}}(t, \{s_1\}\{s_2\}) &:= \log_2 \frac{f_{T|S_1 \vee S_2}(t|s_1 \vee s_2)}{f_T(t)} \\ &= \log_2 \frac{f_{T,S_1}(t, s_1) + f_{T,S_2}(t, s_2)}{f_T(t)(f_{S_1}(s_1) + f_{S_2}(s_2))}. \end{aligned}$$

This measure for continuous redundancy shares many of the favourable properties of its discrete counterpart. In particular, it is localizable (i.e., the global measure can be expressed as an expected value of local values) and can be generalized to arbitrary numbers of source variables (see Section III). Furthermore, the continuous  $i_{\cap}^{\text{sx}}$  measure is differentiable with respect to the underlying probability density function (see Appendix D).

Despite being initially inspired by the ansatz from [34], Definition 3 deviates from the measure-theoretic derivations. The main difference is the absence of the auxiliary random variable designed in [34] to carry the logical structure of the setting. This variable, however, stems from a perspective of abstractness that disregards all differences in the measurement scales between the individual density functions and relies on an auxiliary random variable carrying all the irregularity in this existence theorem. In actual applications this perspective needs to be concretized and made evaluable, which introduces the non-uniqueness of the definition due to the measurement scales and preprocessing. Thus Definition 1 is a single choice of concretization that can be shown to vary smoothly with the overall density as shown in Appendix B.

Comparing this definition for continuous local shared-exclusion redundancy to the discrete local shared-exclusion redundancy of Equation (4), which can be rewritten as

$$\begin{aligned} &i_{\cap}^{\text{sx,discrete}}(t, \{s_1\}\{s_2\}) \\ &= \log_2 \frac{p_{T,S_1}(t, s_1) + p_{T,S_2}(t, s_2) - p_{T,S_1,S_2}(t, s_1, s_2)}{p_T(t)(p_{S_1}(s_1) + p_{S_2}(s_2) - p_{S_1,S_2}(s_1, s_2))}, \end{aligned}$$

reveals that the continuous definition can be obtained from the discrete one by replacing probability masses  $p$  with probability densities  $f$  and ignoring the exclusion terms in the disjunctions. Because of these similarities in the analytic form, the derivation from shared exclusions of Makkeh et al. [26] carries over to the continuous domain, which is why the continuous redundancy measure inherits the same operational interpretation of a memoryless agent trying to guess the target from only the redundant local information on an ensemble of realizations.

As in the case of the mutual information, the average redundancy is defined as the expectation value of the local values, i.e.,

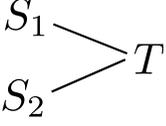
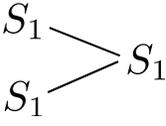
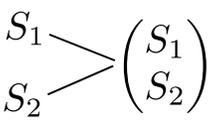
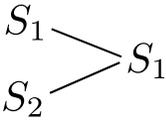
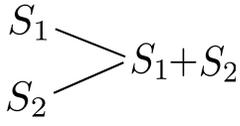
$$\begin{aligned} &I_{\cap}^{\text{sx}}(T, \{S_1\}\{S_2\}) \\ &= \int dt ds_1 ds_2 f_{T,S_1,S_2}(t, s_1, s_2) i_{\cap}^{\text{sx}}(t, \{s_1\}\{s_2\}). \end{aligned}$$

Finally, given this quantification of the redundant information in conjunction with the three classical mutual information terms between  $T$  and the sources, the four information atoms can be computed by solving Equations 1, which in the general case is known as a Moebius inversion [4] (see Section III A).

## 2. Results for toy examples

To gain an intuition for the continuous  $I_{\cap}^{\text{sx}}$  measure, we apply the analytical formulation established in the previous section to simple and intuitive toy examples. Such simple example cases can be found in logic gates, which

TABLE I. Numerical evaluations for four example distributions show qualitative similarities between discrete and the continuous  $I_{\cap}^{\text{sx}}$  on the three selected logic gates, while qualitative differences can be observed for the sum example. In the continuous case, the sources are distributed normally,  $S_i \sim \text{Normal}(\mu = 0, \sigma = 1)$ , except for  $S_2$  in the redundant gate which is set equal to  $S_1$ , and the target is constructed according to the pictograms with added Gaussian noise ( $\sigma = 0.01$ ) to make the mutual information  $I(T : S_1, S_2)$  finite. In the discrete setting, the variables are drawn from uniform binary distributions and no noise is added to the target.

|   | Redundant Gate  | Copy Gate   | Unique Gate  | Sum   |
|---|---|---|--|---|
|  |  |  |  |  |
| Continuous Gates  |   |   |  |   |
| $\Pi_{\text{syn}}$  | 0.000 bits  | 6.644 bits  | 5.525 bits   | 6.497 bits  |
| $\Pi_{\text{unq},1}$  | 0.000 bits  | 0.000 bits  | 1.119 bits   | 0.147 bits  |
| $\Pi_{\text{unq},2}$  | 0.000 bits  | 0.000 bits  | -5.525 bits  | 0.147 bits  |
| $\Pi_{\text{red}}$  | 6.644 bits  | 6.644 bits  | 5.525 bits   | 0.353 bits  |
| $I(T : S_1, S_2)$   | 6.644 bits  | 13.288 bits   | 6.644 bits   | 7.144 bits  |
| Discrete Gates  |   |   |  |   |
| $\Pi_{\text{syn}}$  | 0.000 bits  | 0.415 bits  | 0.415 bits   | 0.415 bits  |
| $\Pi_{\text{unq},1}$  | 0.000 bits  | 0.585 bits  | 0.585 bits   | 0.585 bits  |
| $\Pi_{\text{unq},2}$  | 0.000 bits  | 0.585 bits  | -0.415 bits  | 0.585 bits  |
| $\Pi_{\text{red}}$  | 1.000 bits  | 0.415 bits  | 0.415 bits   | -0.085 bits   |
| $I(T : S_1, S_2)$   | 1.000 bits  | 2.000 bits  | 1.000 bits   | 1.500 bits  |

have been employed as examples for discrete PID measures throughout the literature (e.g., [25, 39]). We here use continuous versions of the well-established “redundant”, “unique” and “copy” gate as well as a sum between source variables and discuss the results compared to their discrete counterparts. In the discrete domain, the two source variables are binary with two equiprobable realizations, whereas in the continuous domain, we draw the variables from a standard normal distribution. To ensure a finite mutual information, which is a prerequisite for a well-defined continuous PID, Gaussian noise with standard deviation  $\sigma = 0.01$  is added to the target variable  $T$  in the continuous case. While the choice of  $\sigma$  significantly affects the numerical results, the qualitative analysis of the results remains unaffected by the precise value chosen (see Appendix E).

While some logic gates (like the “and” gate) have no canonical continuous counterparts, we selected three gates whose basic concept can be straightforwardly transferred to the continuous domain: The redundant gate, in which the two sources are the same and equal to the target, the copy gate, in which the target is a two-dimensional variable whose components consist of the two source variables, and the unique gate, in which the target is a copy of the first source variable while the second source variable is drawn independently from the target. In addition to the three logic gates, we added the “sum” example, in which the source realizations are interpreted as numerical values and the target variable is the sum of the two sources.

While the computation of the discrete PID atoms is a straightforward exercise, the computation of the con-

tinuous atoms is more difficult due to the integral expressions. Note, first, that all continuous examples can be expressed as multivariate Gaussian distributions, for which an analytical closed form for the joint and marginal mutual information terms exists (see Appendix E). The integrals involved in computing the redundancy, however, cannot be easily evaluated analytically. Therefore, we resorted to simple Monte Carlo integration techniques to numerically estimate the involved integrals. To compute the expectation value, we drew randomly from the given joint probability distributions, and computed the local  $i_{\cap}^{\text{sx}}$  redundancy for  $2 \cdot 10^8$  samples. The PID atoms can then be computed from the analytical mutual information terms (see for instance Appendix E) and the numerically computed redundancy via the consistency equations (see Equations 1).

Table I shows the results for all four examples. For the continuous redundant gate,  $I_{\cap}^{\text{sx}}$  attributes all mutual information between the sources and the target to redundancy. This result, which can equivalently be observed in its discrete counterpart, is expected from an information measure constructed for quantifying redundant information, and stems from the fact that the probability of a disjunction between two equal events, e.g.  $S_1 = s_1 \vee S_1 = s_1$  is trivially equal to just the probability the individual event.

In the copy gate, the two marginal mutual information terms each contribute half of the total mutual information, i.e.,  $I(T : \mathbf{S}) = 2I(T : S_1) = 2I(T : S_2)$ . Since the two source variables are independent, learning about source event  $S_1 = s_1$  provides no information about the state of the second source, and vice versa, resulting in no

*source redundancy* [40]. Nevertheless, the redundancy measured by  $I_{\cap}^{\text{sx}}$  is non-zero, since the two variables inform redundantly about the target event because of the mechanism with which the target is constructed from the sources (i.e., *mechanistic redundancy* [40]): When an agent is informed that either  $S_1 = s_1$  or  $S_2 = s_2$  occurred, the agent can redundantly exclude all target events  $T = t = (s'_1, s'_2)$  for which both  $s'_1 \neq s_1$  (which is excluded by observing  $S_1 = s_1$ ) and  $s'_2 \neq s_2$  (which is excluded by observing  $S_2 = s_2$ ). In contrast, for the marginal mutual information  $I(T : S_1)$ , the hypothetical agent is informed which of the two source events actually took place, e.g.,  $S_1 = s_1$ , which enables the agent to additionally rule out all events for which  $s'_1 \neq s_1$  but  $s'_2 = s_2$ . In the discrete case, these additional events for which  $s_1$  does not match the first component of the target event  $t$  but  $s_2$  matches the second, carry a significant probability mass, which makes the redundancy smaller than the mutual information, which thus produces non-zero unique information. For the continuous copy gate, however, the additionally excluded regions carry only negligible probability, which results in negligible unique information.

The continuous unique gate produces similar PID results as its discrete counterpart. Despite  $S_2$  being independent from  $S_1$  and  $T$ , a significant amount of the joint mutual information is assigned to redundancy, which is attributable to incidental overlaps between exclusions in probability space. Note that since only the first variable provides the full information about the target, i.e.,  $I(T : \mathbf{S}) = I(T : S_1)$ , while  $I(T : S_2) = 0$ , the redundancy  $\Pi_{\text{red}}$  and the unique information  $\Pi_{\text{unq},1}$  add up to the total mutual information  $I(T : \mathbf{S})$ , resulting in a positive unique contribution of  $S_1$ . Conversely, the redundancy  $\Pi_{\text{red}}$  and the unique information  $\Pi_{\text{unq},2}$  need to cancel exactly, resulting in negative unique information for the uncorrelated source  $S_2$ . Since the unique information of  $S_2$  also cancels the redundant information in the consistency equation for the full mutual information, the synergy finally needs to be equal to the redundancy to compensate. While the quantitative results differ, these results are qualitatively similar to the PID of the discrete unique gate.

Finally, for the sum gate, most of the information about the target is carried synergistically between the two source variable, reflecting the fact that knowledge of one variable hardly restricts the possible target realizations that can be taken on. Some information is unique to either of the two sources, since knowledge of just one value already makes some values for the sum less plausible, while the small redundancy in this example can likely again be attributed to mechanistic effects. For this gate, the discrete and continuous examples differ qualitatively: While in the discrete case, the unique atoms are positive and the redundancy is negative, all atoms are positive in the continuous case. These discrepancies are likely due to the fact that the discrete sum gate is not a good analogue of the continuous case, since the values of the realizations the discrete random variable can take on has no natural

total ordering, while neighborhoods between values are crucial to the continuous definition.

Depending on the setting that the variables arise in, different preprocessing schemes might need to be applied to the variables before computing the PID atoms. Using different preprocessing schemes, one finds that the values agree qualitatively on most gates, which is mostly due to them being drawn from the same distributions. For a more detailed discussion, refer to Appendix C.

### C. KSG nearest-neighbour based estimator for mutual information

While we established in Section II B how continuous  $I_{\cap}^{\text{sx}}$  can be computed from a given probability distribution, in most research settings this distribution is unknown. More commonly, one has access to only a finite sample drawn from this distribution and needs to *estimate* the information-theoretic quantities of interest from this sample. For mutual information, several different approaches to this estimation problem have been discussed in literature, for instance, binning or kernel density estimation (see, e.g., [41] for a review). However the nearest-neighbour-based approach introduced by Kraskov, Stögbauer and Grassberger [35] (KSG) has been widely adopted due to its favourable variance-bias characteristics. In the following, we explain how local probabilities can be estimated using nearest-neighbor techniques [42] and introduce the key ideas behind the KSG estimator. Then, we show how these concepts can be built upon to derive an estimator for continuous  $I_{\cap}^{\text{sx}}$ .

Many information-theoretic quantities like the mutual information (see Equation (2)) can be expressed as expectation values of *local* quantities. For discrete data, a simple (but not optimal [43]) way to approximate these local values is to replace the probability masses by their frequencies in the sample, which is referred to as *plug-in* estimation. In the continuous case, however, all samples are almost always unique, which is why no continuous analogue to the plug-in principle exists. Thus, one has to revert to more elaborate methods to estimate the local densities  $f(x_i)$  necessary for computing the local information-theoretic values.

*a. Nearest-neighbor-based estimation of local probability densities* To estimate the local probability densities at the sample points, we exploit the intuitive idea that sample points are more likely to be closer together in regions of higher probability density. Note that this intuition crucially rests on the assumption that the sample is sufficiently large such that the true probability density varies only little between neighbouring data points. If this prerequisite is met, a natural approach to estimate the local probability density is to take a ball  $B_{\epsilon}(x_i)$  of fixed radius  $\epsilon$  around the evaluation point  $x_i$ , over which the probability density is assumed to be approximately constant, and count the number of neighbours  $k_i$  within the ball (see Figure 2.A). A point estimate for the prob-

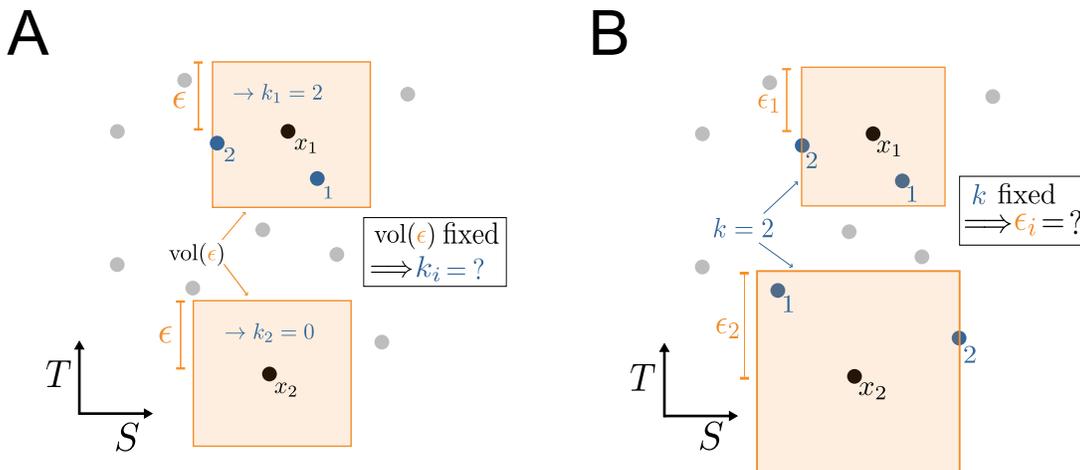


FIG. 2. **Local densities can be estimated from data in two opposite ways.** **A** Kernel density estimation using a fixed volume  $\text{vol}(\epsilon)$  around the reference point  $x_i$ . **B**  $k$ -nearest-neighbor-based estimation using an adaptive search volume, determined by the distance to the  $k$ -th neighbor.

ability density is then given by the fraction of sample points  $(k_i + 1)/N$  within that  $\epsilon$ -ball divided by its volume  $\text{vol}(B_\epsilon(x_i))$  as

$$\hat{f}(x_i) = \frac{k_i + 1}{N \text{vol}(B_\epsilon(x_i))}. \quad (9)$$

Kernel density estimators (KDE) use this principle, but typically weigh the neighbors by their respective distance using a fixed *kernel*, which in general does not have to be the characteristic function of an  $\epsilon$ -neighborhood.

The disadvantage of approaches like this is that the same ball or kernel is used irrespective of the local density: In areas of low density there may be insufficiently many points within the neighborhood to arrive at a good estimation, while conversely in areas of high density the high numbers  $k_i$  would have allowed for a smaller search region, partly relaxing the smoothness assumption on  $f$ . A way to get around these problems is to invert the procedure: Instead of fixing the search radius  $\epsilon$  and counting the number of neighbouring points  $k_i$ , one can fix a natural number  $k$  and compute the distance  $\epsilon_i$  from point  $x_i$  to its  $k$ -th nearest neighbor (see Figure 2.B). Analogously to Equation (9), a point estimate of the local probability density can now be expressed as

$$\hat{f}(x_i) = \frac{k + 1}{N \text{vol}(B_{\epsilon_i}(x_i))}. \quad (10)$$

*b. The Kozachenko-Leonenko Estimator for Shannon Differential Entropy* Kozachenko and Leonenko [42] used the nearest-neighbour-based approach to build an estimator for the Shannon differential entropy [3]

$$H(X) = \mathbb{E}[-\log_2 p(X)] = - \int dx p(x) \log_2 p(x). \quad (11)$$

However, instead of using the point estimate of Equation (10), the authors derive the expectation value for

$\log p(x_i)$  for the probability mass of the  $\epsilon$ -ball

$$p(x_i) = \int_{B_{\epsilon_i}(x_i)} dx' f(x') \approx \text{vol}(B_{\epsilon_i}(x_i)) f(x_i)$$

given the number of neighbours  $k$  for each sample point, which results in the expression

$$\log \widehat{p_X}(x_i) = \psi(k) - \psi(N),$$

where  $\Psi$  represents the Digamma function. Plugging this result into Equation (11), the authors obtain an estimator for the Shannon differential entropy as (for a more formal derivation, refer to Appendix F)

$$\begin{aligned} \hat{H}(X) &= -\frac{1}{N} \sum_{i=1}^N \log [\widehat{p_X}(x_i)] \\ &= -\psi(k) + \psi(N) + \frac{1}{N} \sum_{i=1}^N \log [\text{vol}(B_{\epsilon_i})]. \end{aligned} \quad (12)$$

*c. The Kraskov-Stögbauer-Grassberger (KSG) Estimator for Mutual Information* The most straightforward way to construct an estimator for continuous mutual information from the results of Kozachenko and Leonenko [42] is by using the identity  $I(T : S) = H(T) + H(S) - H(T, S)$ . However, in doing so, the biases of the individual entropy estimates likely do not cancel, which led Kraskov et al. [35] to take a different approach: First, like for the Kozachenko-Leonenko estimator, the distance  $\epsilon_i$  of point  $x_i = (t_i, s_i)$  to its  $k$ th nearest neighbour is determined in the joint space (see Figure 3.A). Subsequently, the number of neighbours  $n_T$  which have a distance of less than  $\epsilon_i$  in only the marginal space of the first variable  $T$  is determined (see Figure 3.B), and the procedure is repeated analogously to determine the number of neighbors  $n_S$  in radius  $\epsilon_i$  in the marginal space of the second random variable  $S$  (see Figure 3.C). The advantage

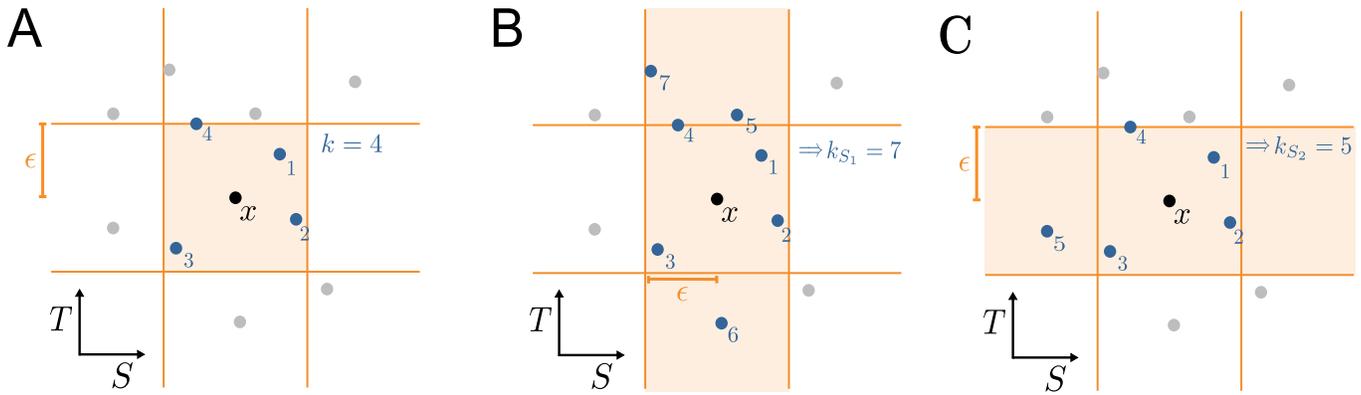


FIG. 3. The KSG estimator for mutual information [35] works by considering  $k$ th nearest neighbors in the joint and marginal spaces in three steps. **A** Determining search radius in joint space **B** counting neighbors in marginal space of  $S$  and **C** counting neighbors in marginal space of  $T$ .

of obtaining the marginal probability densities using the distance determined in the joint space is that—when using the maximum norm—the volume terms which appear in Equation (12) cancel exactly, leading to the succinct expression

$$\hat{I}(T : S) = \Psi(k) + \Psi(N) - \langle \Psi(n_T(i)) + \Psi(n_S(i)) \rangle_i, \quad (13)$$

for the estimated mutual information. Here, angled brackets  $\langle \cdot \rangle_i$  denote the average over the sample points.

Clearly, the estimated density as well as the information quantity assessed via the density estimation will depend on  $k$ . However, Kraskov et al. [35] have empirically studied the influence of  $k$ , or more exactly a fixed  $\frac{k}{N}$  on the estimation result and concluded that in the limit of infinitely many sample points the exact choice of  $k$  does not matter as all choices will converge.

#### D. Generalizing the KSG estimator for $I^{\text{sx}}$

*a. Adapting local probability density estimation for redundant information* To understand how to estimate the continuous redundancy  $I_{\cap}^{\text{sx}}$  using an approach similar to Kraskov et al. [35]’s, first consider the case of estimating the mutual information  $I(T : \mathbf{S})$  between a single target random variable  $T$  and two joint source variables  $\mathbf{S} = (S_1, S_2)$  using the KSG estimator. In the first step, the radius  $\epsilon_i$  is determined as the smallest radius in the joint space for which the  $k$ -th neighbor of the point  $x_i = (t_i, s_{1,i}, s_{2,i})$  lies within the Ball  $B_{\epsilon_i}(x_i)$ . Using the maximum norm, this ball in the joint space is equal to the product of the  $\epsilon$ -balls in the marginal spaces, i.e.,  $B_{\epsilon_i}(t_i) \times B_{\epsilon_i}(s_{1,i}) \times B_{\epsilon_i}(s_{2,i})$  (square region in Figure 4.A). In the second step, neighbours are counted within the marginal space of  $\mathbf{S}$ , where only those points are counted which lie within less than the distance  $\epsilon_i$  to  $x_i$  in both dimensions  $S_1$  and  $S_2$  simultaneously.

As laid out in detail in Section II B, the difference in the mathematical formulation between mutual information between a single target  $T$  and two (or more) source

variables  $\mathbf{S} = (S_1, S_2)$  and  $I_{\cap}^{\text{sx}}$  redundancy is that while the former is expressed with conjunctions, which amount to *intersections* of source events in sample space, the latter deals with disjunctions, which result in *unions*. The KSG estimator can be readily adapted to this concept by changing the search regions accordingly: To compute the redundancy in the example case of two source variables introduced above, the radius  $\epsilon_i$  now has to be determined not as the smallest radius laying on the intersection of  $B_{\epsilon_i}(t_i)$ ,  $B_{\epsilon_i}(s_{1,i})$  and  $B_{\epsilon_i}(s_{2,i})$  (“square” region in Figure 4.A) but on the union of the source variable neighbourhoods intersected with the target ball  $B_{\epsilon_i}(t_i) \times B_{\epsilon_i}(s_{1,i}) \times S_2 \cup B_{\epsilon_i}(t_i) \times S_1 \times B_{\epsilon_i}(s_{2,i})$  (“cross” region in Figure 4.B). Similarly,  $n_S$  is determined by counting all points which lie within the union of marginal balls with radius  $\epsilon$ .

Note that, analogously to Kraskov’s [35] estimator for mutual information, this estimation procedure assumes the maximum norm when combining the unions between the source with the target variables in order for the volume terms of the KL estimator to cancel, which simplifies the analytical formulation considerably. Although the choice of any particular distance norm will matter for small sample sizes (e.g. Gao et al. [44] suggest using the Euclidean norm), the assumed continuity of the density function and all its marginals ensures that regardless of the choice of distance norm the estimation of the local densities *asymptotically* converges to the same values. Intuitively, the average probability over decreasing neighbourhoods of a point converges to the density at that point in the limit of infinitely many samples, regardless of the exact choice of the shape of the neighbourhood, the shape being determined by the distance norm. Technically, however, these neighbourhoods must not be chosen entirely arbitrarily—the shrinking procedure obey the *nicely shrinking* [45] property, guaranteeing that the shrinking neighbourhoods indeed cover neighbourhoods in each dimension of the point. For the nicely shrinking property to hold, one can choose any distance norm as all imply the nicely shrinking property. This asymptotic

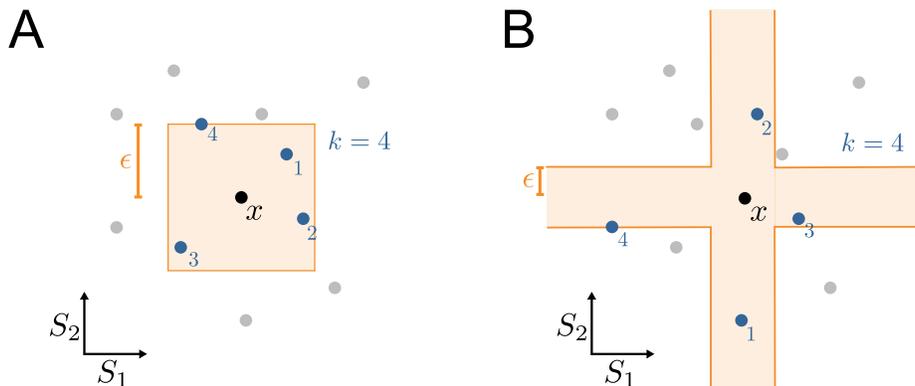


FIG. 4. To estimate local redundant information for a two-variable example, the search regions of the original mutual information KSG estimator [35] need to be adapted. **A** For mutual information estimation, the radius  $\epsilon$  at a point  $x$  is assessed in the joint space by searching for the minimal distance such that *all* of the variables’ marginal distances to  $x$  is smaller or equal to  $\epsilon$ . **B** For redundancy estimation, on the other hand,  $\epsilon$  is inferred by searching for the minimal distance such that *at least one* of the variables’ marginal distances to  $x$  is smaller or equal to  $\epsilon$ .

convergence behavior of the estimated densities then directly implies the asymptotic convergence of the estimate of the redundant information to its analytical value.

*b. Nearest-neighbor-based estimator for continuous redundant information* Adapting the steps by Kraskov et al. [35] using the described estimation procedure for the probability density of a logical disjunction we find

$$\begin{aligned} \hat{I}_{\cap}^{sx}(T : X, Y) &= \langle \log[\widehat{f_{T, X \vee Y}^i}(t, x, y)] \rangle_i \\ &\quad - \langle \log[\widehat{f_{X \vee Y}^i}(x, y)] \rangle_i - \langle \log[\widehat{f_T^i}(t)] \rangle_i \\ &= \psi(k) + \psi(N) - \langle \psi(n_{X \vee Y}(i)) \rangle_i - \langle \psi(n_T(i)) \rangle_i. \end{aligned} \quad (14)$$

More precisely, we have retaken the approach of choosing the  $\epsilon$  to be determined by a nearest neighbor search in the joint space. We then use the same  $\epsilon$  in the marginal spaces themselves, such that the volume terms cancel exactly. The search with the predetermined  $\epsilon$  will then cause an adapted number of nearest neighbors in the marginal spaces, denoted by  $n_{X \vee Y}(i)$  and  $n_T(i)$ , respectively. The exact steps leading to Equation (14) are outlined in Appendix G. Furthermore, details of the implementation and code availability are described in Appendix H.

Note, however, one subtle difference in the meaning of the parameter  $\epsilon$  between the mutual information and redundancy case: While scaling  $\epsilon$  by a constant factor for one but not the other source variable changes the estimated mutual information for any finite sample, it does not change what the mutual information estimator converges to in the limit of infinite samples. This is different in the case of redundancy: Because of the addition of marginal densities in Definition 1, scaling one but not the other search radius  $\epsilon$  makes a difference also in the limit of infinite samples. Using the same  $\epsilon$  for both source variables here thus reflects the choice of treating the source variables on “equal footing”. This train of thought aligns with the dependence of the result on the chosen distance function, as in many cases a transformation of one of the

variables can be reflected with a change in the utilized norm.

For a fixed small  $k$ , the specific value of  $k$  does not influence the estimation results if sufficient samples  $N$  are provided. Analogously to the KSG estimator for mutual information, the estimation results only depend on the fraction  $k/N$ , as a fixed fraction  $k/N$  hinders the estimator in resolving high-frequency fluctuations in the probability density. Fluctuations on all length scales are resolved in the limit  $k/N \rightarrow 0$ , for which the estimator converges to the analytical solutions (see Figure 14 and Figure 15 in Appendix I). Following the suggestions from Kraskov et al. [35], we typically recommend  $k$  to be set to 2, 3 or 4, where estimates with higher  $k$  have smaller variance due to the stochasticity of the distances neighbours averaging out, but also higher bias due to their inability to resolve small-scale features.

### E. Convergence of the estimator for simple toy examples

In order to assess the efficacy of our estimation procedure, we applied it to the toy examples outlined in Section II B. To this end, different numbers of samples have been drawn from the known probability distributions and have been plugged into the estimator to check for convergence.

Figure 5 shows how the PID atoms converge with more samples for the four example cases. The copy gate and sum example converge slower than the other values. This behaviour is inherited from the KSG estimator, which also converges more slowly for higher-dimensional distributions and distributions with higher mutual information. All in all, we conclude that the estimator succeeds in reproducing the numerically evaluated analytical results from Section II B given sufficiently many samples.

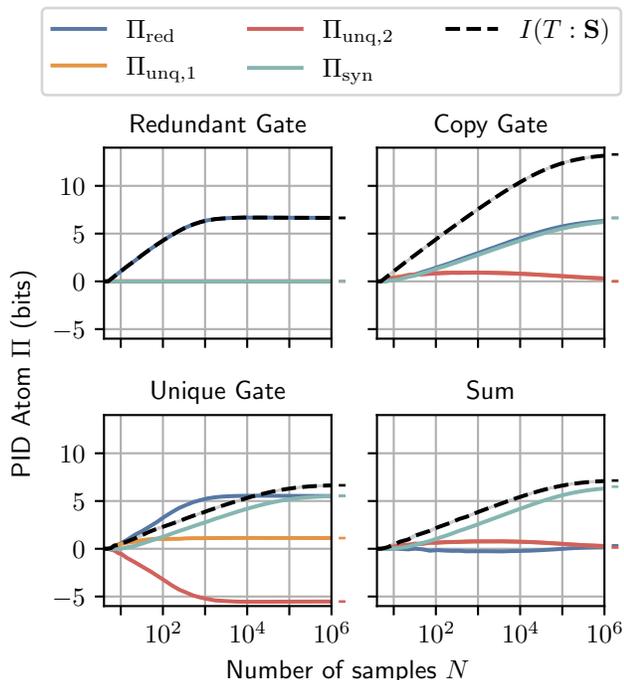


FIG. 5. **With increasing sample size  $N$ , the estimated PID atoms converge to the numerically evaluated analytical results.** The subfigures show the four PID atoms  $\Pi$  and the mutual information  $I(T : \mathbf{S})$  for the continuous redundant, unique and copy gates as well as the sum example estimated using the nearest-neighbors approach outlined above with  $k = 4$ . The colored bars to the right of the graphs correspond to the respective analytical value for each atom and gate from Table I.

### III. DEFINITION OF MULTIVARIATE CONTINUOUS $I_{\cap}^{\text{sx}}$

#### A. The multivariate PID Problem

In the previous sections we introduced the continuous  $I_{\cap}^{\text{sx}}$  measure for two source variables. However, as was already mentioned, the PID framework as envisioned by Williams and Beer [4] is more general and in principle applicable to an arbitrary number of sources. In the following, we will explain the PID problem for the general, multivariate case and generalize the continuous  $I_{\cap}^{\text{sx}}$  measure accordingly. This section may be skipped if the reader has a particular application in mind for which two source variables suffice and this section is found too technical.

Before we generalize to an arbitrary number of source variables, note that already in the bivariate case, each atom can be uniquely identified by which mutual information terms it contributes to via the consistency equations (see Equations 1). For instance, the unique atom  $\Pi_{\text{unq},1}$  is the only atom that is part of both  $I(T : S_1)$  and  $I(T : \mathbf{S})$  but not  $I(T : S_2)$  while the synergistic informa-

tion atom  $\Pi_{\text{syn}}$  is precisely the one atom whose information is contained in neither  $I(T : S_1)$  nor in  $I(T : S_2)$ , but only in  $I(T : S_1, S_2)$ . In the following, we show how this notion of parthood, also called *mereology*, can serve as a general framework for identifying and ordering PID atoms for arbitrarily many source variables following the work by Gutknecht et al. [5].

Mathematically, the mereological relations between an atom and the marginal mutual information terms can be captured by a *parthood distribution*  $\Phi : \mathcal{P}(\mathbf{S}) \rightarrow \{0, 1\}$ : A boolean function defined on the power set of source variables  $\mathcal{P}(\mathbf{S})$  which is equal to 1 for exactly those sets of sources  $\mathbf{S}_a = \{S_i \mid i \in \mathbf{a}\}$  for which the atom  $\Pi_{\Phi}$  is part of the marginal mutual information  $I(T : \mathbf{S}_a)$  [5]. Note that not all boolean functions constitute valid parthood distributions: Since all information about  $T$  that is contained in the mutual information  $I(T : \mathbf{S}_a)$  with a set of sources  $\mathbf{S}_a$  must naturally also be contained in the mutual information  $I(T : \mathbf{S}_b)$  with any superset  $\mathbf{S}_b \supset \mathbf{S}_a$ ,  $\Phi$  has to be *monotone*, in the sense that  $\mathbf{S}_b \supset \mathbf{S}_a \Rightarrow (\Phi(\mathbf{S}_a) = 1 \rightarrow \Phi(\mathbf{S}_b) = 1)$ . For instance, any information provided by source  $S_1$  alone must naturally also be present when taking sources  $S_1$  and  $S_2$  together. Furthermore, we must always require that  $\Phi(\emptyset) = 0$  and  $\Phi(\mathbf{S}) = 1$ , indicating that no atom is part of the mutual information with the empty set and, conversely, all information atoms are part of the full joint mutual information  $I(T : \mathbf{S})$ .

The marginal mutual information terms can be constructed additively from these PID atoms according to the generalized consistency equations

$$I(T : \mathbf{S}_a) = \sum_{\Phi: \Phi(\mathbf{S}_a)=1} \Pi_{\Phi}. \quad (15)$$

The number of such parthood distributions, and thus of the PID atoms, grows superexponentially like the Dedekind numbers [5]. On the other hand, there are only  $2^n - 1$  classical mutual information quantities providing constraints through Equation (15), leaving these consistency equations underdetermined and making it impossible to determine the sizes of the PID atoms. This underdeterminedness is typically resolved by introducing a measure for generalized redundancy  $I_{\cap}(T : \mathbf{S}_{\Phi})$  for each parthood distribution  $\Phi$ . This redundancy generalizes the mutual information and fulfills

$$I_{\cap}(T : \mathbf{S}_{\Phi}) = \sum_{\Psi: \Psi \prec \Phi} \Pi_{\Psi}, \quad (16)$$

where the sum is over all atoms  $\Pi_{\Psi}$  with a parthood distribution  $\Psi$  which are part of the generalized redundancy  $I_{\cap}(T : \mathbf{S}_{\Phi})$ .

This condition is similar to the monotonicity requirement of the mutual information mentioned before and is mathematically captured in the ordering relation  $\prec$  given by

$$\Psi \prec \Phi \Leftrightarrow \Phi(\mathbf{S}_a) = 1 \rightarrow \Psi(\mathbf{S}_a) = 1 \text{ for all } \mathbf{S}_a \subseteq \mathbf{S}.$$

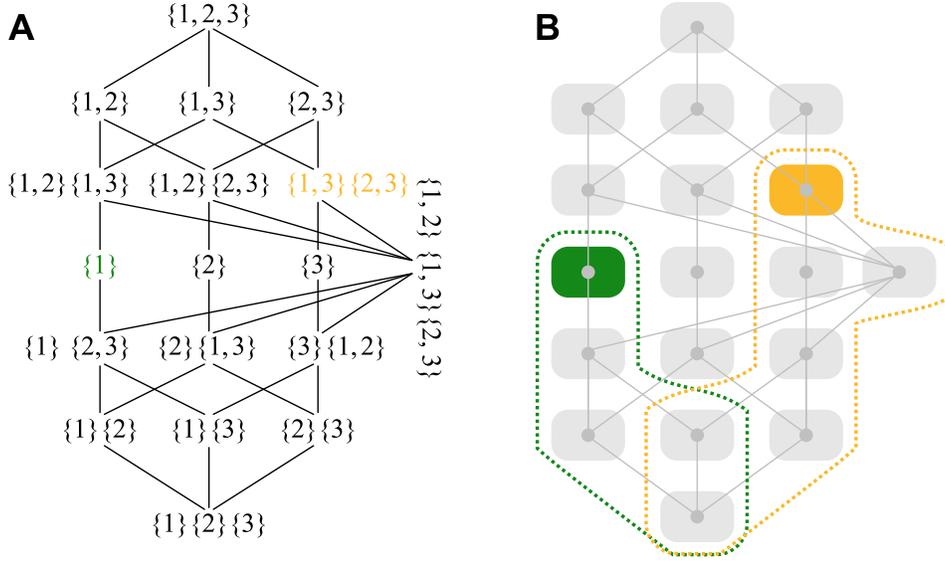


FIG. 6. **The PID lattice structure for three source variables reveals how generalized redundancies can be dissected into PID atoms [4, 5].** **A** Every node in the redundancy lattice represents a generalized redundancy  $I_\alpha$  for a given antichain, while edges indicate parthood of generalized redundant information in information terms higher up the lattice. For instance, the mutual information  $I(T : S_1)$ , which corresponds to the self-redundancy with the antichain  $\{1\}$ , contains all redundancies in its downset, i.e., those with antichains  $\{1\}\{2,3\}$ ,  $\{1\}\{2\}$ ,  $\{1\}\{3\}$  and  $\{1\}\{2\}\{3\}$ , which, however, themselves overlap and contain parts of each other. **B** The PID atoms (shaded boxes), which are referenced by the same antichains, only comprise the information that the redundancy of the same antichain contains *in excess of what the all the atoms below contain together*. Atoms thus reflect the incremental increases of the redundancies which conversely can be used to construct the generalized redundancies. For instance, the redundancy with antichain  $\{1\}$  is the sum of the four atoms in its downset plus the atom  $\{1\}$  itself (green shaded box and green dashed outline, respectively). Analogously, the figure shows how the redundancy with the antichain  $\{1,3\}\{2,3\}$  (orange) is constructed from the atoms in its downset.

This partial ordering bestows a structure onto the set of PID atoms which is known as the “redundancy lattice” [4, 5] (see Figure 6).

Comparing Equations (15) and (16), one finds that the mutual information can be interpreted as a “self-redundancy”—a special case of a redundancy which describes how the mutual information provided by a set of sources about the target is trivially redundant with itself. Finally, to quantify the atoms  $\Pi_\Phi$  given the redundancies  $I_\alpha(T : \mathbf{S}_\Phi)$ , Equation (16) can be inverted, which, due to the partial ordering, is known as a *Moebius inversion* [4].

Equivalently, each parthood distribution can be uniquely identified by the set of sets of source variables for which the parthood distribution is equal to 1, i.e.,  $\Phi^{-1}[\{1\}]$ . Because of the monotonicity requirement mentioned before, we can make this representation even more concise by removing all sets which are supersets of others, which gives the so-called *antichains*

$$\alpha = \{\mathbf{a} \in \Phi^{-1}[\{1\}] \mid \nexists \mathbf{b} \in \Phi^{-1}[\{1\}] \text{ s.t. } \mathbf{b} \subseteq \mathbf{a}\}, \quad (17)$$

which is the most prevalent way how PID atoms are referred to throughout literature and also the formalism that PID was originally conceived in [4]. Throughout this paper, the outer braces of the set of sets are neglected for the sake of brevity (e.g.,  $\{1\}$ ,  $\{2\}$  instead of  $\{\{1\}, \{2\}\}$ ). The term antichain stems from the fact that the sets contained within  $\alpha$  are incomparable with respect to the

partial order of set inclusion. These antichains make the intuitive meaning of the redundancy terms explicit, e.g.,  $I_{\alpha, \{1\}\{2,3\}}$  refers to the information that can either be obtained from source  $S_1$  alone, or—redundantly to that—from sources  $S_2$  and  $S_3$  taken together.

### B. Definition of multivariate continuous $I_\alpha^{\text{sx}}$

Since in the multivariate case the generalized redundancy  $I_{\alpha, \mathbf{a}}$  refers, for a particular antichain  $\alpha$ , to information that can equivalently be obtained from all sets  $\mathbf{S}_\mathbf{a}$  for  $\mathbf{a} \in \alpha$  of source variables taken together, the logical statement in the conditional probability of Definition 2 needs to be generalized to accommodate for such general disjunctions as

$$i_{\alpha, \mathbf{a}}^{\text{sx}}(t : s_1; s_2) = \log_2 \left[ \frac{p(t \mid \bigvee_{\mathbf{a} \in \alpha} \mathbf{S}_\mathbf{a} = \mathbf{s}_\mathbf{a})}{p(t)} \right] \quad (18)$$

$$= \log_2 \left[ \frac{p(t \mid \bigvee_{\mathbf{a} \in \alpha} \bigwedge_{i \in \mathbf{a}} S_i = s_i)}{p(t)} \right], \quad (19)$$

where  $\mathbf{S}_\mathbf{a} = \{S_i \mid i \in \mathbf{a}\}$  and likewise  $\mathbf{s}_\mathbf{a} = \{s_i \mid i \in \mathbf{a}\}$ . This generalization is analogous to how the multivariate redundancy is introduced for the discrete case by Makkeh et al. [26].

### C. Estimation of multivariate continuous $I_{\cap}^{sx}$

To estimate multivariate continuous  $I_{\cap}^{sx}$ , the search regions used by the nearest neighbor estimator need to be expanded to higher dimensions. For  $d > 2$  source variables  $\mathbf{S} = (S_1, \dots, S_d)$  and a specific antichain  $\alpha$ , the search space for determining the radius  $\epsilon_i$  is given by the

Using this notion of neighborhood, the estimator for continuous  $I_{\cap}^{sx}$  can be generalized to arbitrary antichains  $\alpha$  as in Definition 3, yielding

$$\begin{aligned} & \hat{I}_{\cap}^{sx}(T : \alpha) \\ &= \langle \log[\widehat{f_{\alpha, T}^i}(t, s_k)] \rangle_i - \langle \log[\widehat{f_{\alpha}^i}(s_k)] \rangle_i - \langle \log[\widehat{f_T^i}(t)] \rangle_i \\ &= \psi(k) + \psi(N) - \langle \psi(n_{\alpha}(i)) \rangle_i - \langle \psi(n_T(i)) \rangle_i. \end{aligned}$$

## IV. APPLICATION OF CONTINUOUS $I_{\cap}^{sx}$ ESTIMATOR TO SIMULATED DATA FROM AN ENERGY MANAGEMENT SYSTEM

### A. Simulation of an energy management system

To demonstrate the applicability of the proposed definition for continuous  $I_{\cap}^{sx}$  and its corresponding estimator, we applied it to simulated data for an energy management system and show how the simulated relationships between system components can be recovered from the estimated PID. In a hybrid simulation approach, we integrated real-world sensor readings from a Honda R&D facility into simulation components, which were in turn modelled via fundamental physical differential equations [46]. We informed our simulation by recordings of the actual facility energy consumption profile and local weather patterns, obtained by smart meter measurement. Specifically, the measured time series consist of all environmental inputs such as temperature, wind speeds, and humidity, as well as the power consumption of subsystems of the Honda R&D facility, while all other variables were simulated. The simulations have been done in the Modelica simulation language, using the commercial tool SimulationX [47].

The simulation comprises multiple different agents representing energy sources, energy storage, consumers and grid interconnections. The energy originates from four sources: An on-premise wind turbine, a photovoltaic (PV) system, a combined heat and power natural gas plant (CHP), as well as energy transfers from the electricity grid. This energy is utilized as building power consumption, while excess power is stored in the battery storage system or fed back into the external power grid (see Figure 8).

In the simulation, the different agents of the system are given set rules for how they respond to each other or to external changes. Firstly, as the CHP not only provides electrical energy to the facility but also heating, it is set

union of intersections

$$\bigcup_{\mathbf{a} \in \alpha} B_{\epsilon_i}(t_i) \times \prod_{j \in \mathbf{a}} B_{\epsilon_i}(s_{i,j}) \times \prod_{j \notin \mathbf{a}} S_j. \quad (20)$$

For illustration, the search regions for three source variables (trivariate PID) are visualized in Figure 7 for all antichains in the redundancy lattice.

to operate only during times in which the ambient temperature is equal to or lower than 16 °C. During those times, the plant produces a continuous electric power output depending on the overall energy consumption of the facility (see Algorithm 1). The stationary battery is either charged or discharged depending on the overall power consumption level of the facility (see Algorithm 2). Specifically, the storage system is set to **Discharge** according to a binary hysteresis with respective threshold values of 330 kW and 350 kW for the overall power consumption (see Figure 9). The model was simulated for a full year to capture seasonal variations in both renewable energy availability and power demand.

---

#### Algorithm 1: CHP controller

---

```

if Temperature  $\leq$  16 °C then
  | PowerOutput  $\leftarrow$  0.4 PowerReference
else
  | PowerOutput  $\leftarrow$  0
end

```

---



---

#### Algorithm 2: Battery controller

---

```

if Discharge then
  | PowerOutput  $\leftarrow$  -200 000
else if PowerReference  $\leq$  200 000 then
  | PowerOutput  $\leftarrow$  100 000
else
  | PowerOutput  $\leftarrow$  100 000
end

```

---

The goal of the PID analysis is to find the environmental variables which best predict the net energy transfer to the power grid  $P_{\text{Grid}}$ . Specifically, we performed analyses between the net energy transfer as the target variable and used the ambient temperature,  $T_A$ , wind speed,  $v_{\text{Wind}}$ , and time of day,  $t_{\text{Day}}$ , as sources. A full PID between these three source variables reveals how the predictive capabilities of the environmental variables is distributed among the individual factors.

## B. Results of experiments

### 1. Bivariate PID results

As a first step, we analyzed how well the net energy transfer  $P_{\text{Grid}}$  is predictable from the ambient temperature  $T_A$  and the wind speed  $v_{\text{Wind}}$ . The results of

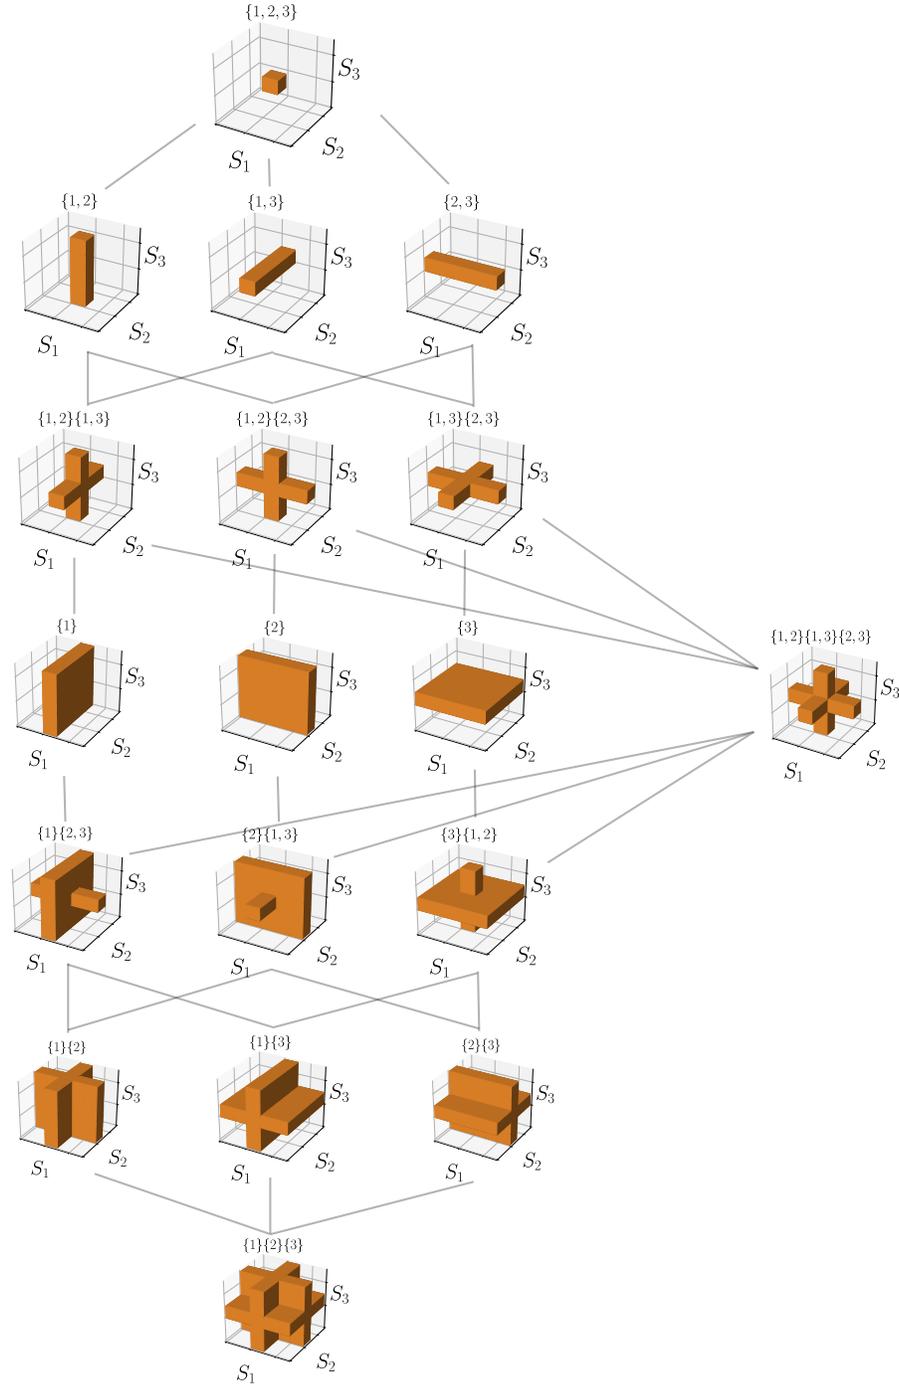


FIG. 7. **The nearest neighbor search regions for three source variables are cumulative when going from the full synergy (top) towards the full redundancy (bottom) in the redundancy lattice.** The orange regions depict where to search for nearest neighbors, and its border signifies the surfaces of equal distance from the logical statement corresponding to each antichain. Here the finite heights, widths and depths  $\epsilon$  are chosen equally large. Also note that the search regions for all classical mutual information terms (i.e., antichains  $\{1, 2, 3\}$ ,  $\{1, 2\}$ ,  $\{1, 3\}$ ,  $\{2, 3\}$ ,  $\{1\}$ ,  $\{2\}$  and  $\{3\}$ ) correspond to marginalized hypercubes, just as in the regular definition of the KSG estimator [35].

this bivariate PID are summarized in Table II. In total, 2.323 bits of information about  $P_{\text{Grid}}$  can be obtained by observing  $T_A$  and  $v_{\text{Wind}}$ . Intuitively, if all values for  $P_{\text{Grid}}$  were equally likely, one bit of information would be

equivalent to restricting the range of possible values for the energy transfer by half after observation of the environmental variables. Accordingly, the observed 2.323 bits would then equate to a reduction of range of the possible

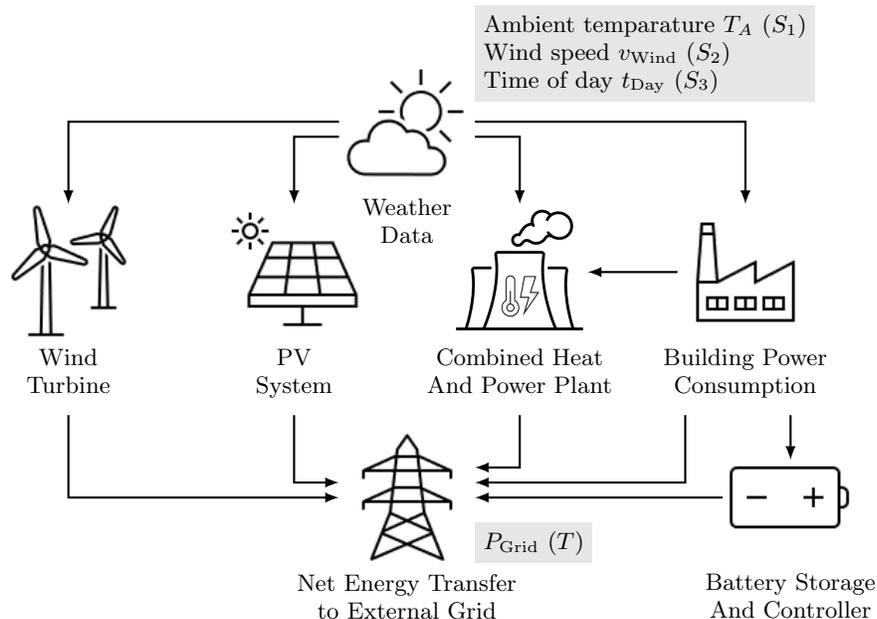


FIG. 8. **Causal structure of the energy management system model.** An arrow from A to B reads as “A influences B”. The model simulates the building power and heat demand based on time and weather conditions. Energy is provided by the grid connection, a PV-system, a combined heat and power plant (CHP), a wind turbine, and a stationary battery. The battery’s charging and discharging behavior is controlled depending on the overall power consumption level. The CHP produces energy depending on the overall power consumption and the ambient temperature. The gray boxes highlight the choices for source ( $S_i$ ) and target ( $T$ ) variables for the PID analysis.

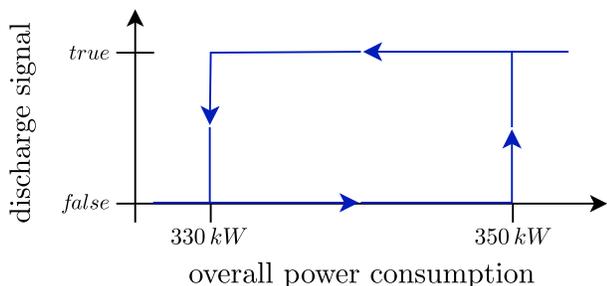


FIG. 9. **Hysteresis function of the battery controller.** Depending on the overall power consumption level  $u$ , a discharge signal  $y$  is sent to the controller.

TABLE II. **Most information that the ambient temperature ( $S_1 = T_A$ ) and the wind speed ( $S_2 = v_{Wind}$ ) hold about the net energy transfer ( $P_{Grid}$ ) is held synergistically.** The PID is calculated considering the target variable to be the net energy transfer to the electricity grid. The estimator was applied to  $N = 10^5$  data points.

|                   | Size of information atom |
|-------------------|--------------------------|
| $\Pi_{syn}$       | 1.363 bits               |
| $\Pi_{unq,1}$     | 0.296 bits               |
| $\Pi_{unq,2}$     | 0.034 bits               |
| $\Pi_{red}$       | 0.630 bits               |
| $I(T : S_1, S_2)$ | 2.323 bits               |

outcomes by a factor of about 1/5.

Most of this information, namely 1.363 bits are carried synergistically between the two environmental variables, meaning that this information is not accessible from any individual source alone. This implies that the net energy transfer is, to a large extent, predictable only from the interaction between the temperature and wind speed. Some of the information can also be determined redundantly from both sources ( $\Pi_{red} = 0.630$  bits) and finally, some information can be gained from the ambient temperature alone ( $\Pi_{unq,1} = 0.296$  bits), while the wind speed conveys hardly any information exclusively by itself ( $\Pi_{unq,2} = 0.034$  bits).

Note the strong similarity in values to the sum gate Table I, which constructs the target as a net sum of its two source inputs. One can assume that this results from the dependency on the ambient temperature in the building power consumption and the CHP controller Algorithm 1 and the dependency of the energy generation of the wind turbine on the wind speed leading to these two source variables’ effects would *add up* together to inform about the net energy transfer to the external grid. However, the inequality in effects on other parts of the system such as potential temperature and wind speed dependency of for instance the photovoltaic system lead to a less balanced summation than was done in the equiweighted sum gate, and becomes conceptually more similar to a skewed sum gate as in Table VI in Appendix J.

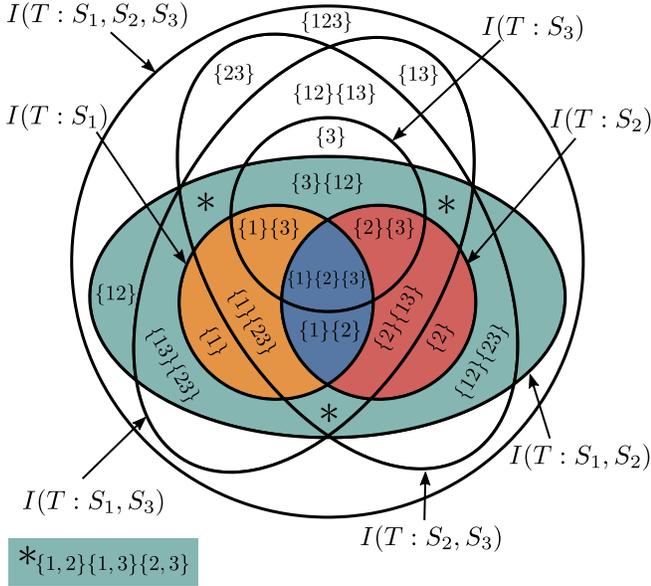


FIG. 10. By adding a third source variable  $S_3$ , the four PID atoms between the target  $T$  and the sources  $S_1$  and  $S_2$  of Figure 1 are further subdivided and new ones are added, resulting in a total of 16 atoms [4]. The individual atoms are identified by their antichains, e.g.,  $\{1\}\{2,3\}$  is the redundancy between source  $S_1$  taken alone and sources  $S_2$  and  $S_3$  synergistically taken together.

## 2. Trivariate PID results

A possible common cause for the correlation between the environmental variables and the net energy transfer from or to the power grid might be the time of day. To investigate this hypothesis, we augmented our analysis by incorporating the time of day as a third source variable in our PID analysis. By adding a third variable, the bivariate PID atoms themselves are dissected into smaller parts depending on their relation with the third variable (see Figure 10). For instance, the bivariate synergy  $\Pi_{\text{syn}}$  splits into five parts upon the introduction of the third variable, which are each characterized by a different antichain  $\alpha$ :

- The part of the synergy between only the sources  $S_1$  and  $S_2$  which is **not** part of  $S_3$  ( $\alpha = \{1, 2\}$ ),
- the part of the synergy between  $S_1$  and  $S_2$  which is redundant with  $S_3$  ( $\alpha = \{1, 2\}\{3\}$ ),
- the part of the synergy between  $S_1$  and  $S_2$  which is redundant with the synergy between  $S_1$  and  $S_3$  ( $\alpha = \{1, 2\}\{1, 3\}$ ),
- the part of the synergy between  $S_1$  and  $S_2$  which is redundant with the synergy between  $S_2$  and  $S_3$  ( $\alpha = \{1, 2\}\{2, 3\}$ ) and finally
- the part of the synergy between  $S_1$  and  $S_2$  which is redundant with both the synergy between  $S_1$

and  $S_3$  and the synergy between  $S_2$  and  $S_3$  ( $\alpha = \{1, 2\}\{1, 3\}\{2, 3\}$ ).

Examining the estimation results of the trivariate PID, as listed in Table III and visualized in Figure 11, we can observe that the synergy between ambient temperature and wind speed primarily resides within the trivariate atom characterized by the antichain  $\{1, 2\}\{3\}$ . This implies that the bivariate synergy between  $T_A$  and  $v_{\text{Wind}}$  largely duplicates the information conveyed by the time of day. A similar trend emerges when considering bivariate redundancy  $\Pi_{\text{red}}$ , which also predominantly overlaps with the time of day, i.e., resides mainly in the trivariate atom  $\{1\}\{2\}\{3\}$ , and also the bivariate unique information of the ambient temperature  $\Pi_{\text{unq},1}$ , which mostly resides in the redundant atom  $\{1\}\{3\}$ . These findings substantiate our hypothesis that weather variables are primarily redundant with the time of day. One possible explanation for this is that the weather variables primarily act as intermediaries, conveying information that the time of day carries about the net energy transfer.

Nevertheless, there are also other PID atoms which contribute non-negligibly to the total mutual information. In particular, some information is carried synergistically between the ambient temperature and the time of day ( $\{1, 3\}$ ) or, to a lesser extent, between the synergy between all variables ( $\{1, 2, 3\}$ ) and the synergy between wind speed and the time of day ( $\{2, 3\}$ ). Yet other relevant contributions come from the redundancy between the ambient temperature and the time of day ( $\{1\}\{3\}$ ) and the redundancy between all three two-variable synergies ( $\{1, 2\}\{1, 3\}\{2, 3\}$ ).

Further, note that the information that the time of day provides uniquely and independent of ambient temperature and wind speed is on average negative, i.e., misinformative about the total power draw. Average misinformation can arise in  $I_{\bar{\alpha}}^{\text{sx}}$  since the operational interpretation, which has been put forward in Makkeh et al. [26] and carries over to the continuous case, assumes a memoryless agent and interprets the average information-theoretic quantities as ensemble averages. For a hypothetical agent incapable of learning from past events, this means that gaining only the purely unique information of the time of day diminishes the agent's accuracy in predicting the net energy transfer.

## V. DISCUSSION

This work presents a tractable analytical definition of continuous PID which is based on the intuitive ideas of the discrete  $I_{\bar{\alpha}}^{\text{sx}}$  measure [26] and the measure-theoretical existence proofs by Schick-Poland et al. [34]. It is in a local form by definition, is applicable to arbitrary many source variables and differentiable. We exemplify the intuition of our definition on four toy examples and observe that the measure gives mostly qualitatively comparable results to the the discrete case when computed on

TABLE III. The three source PID of the Honda energy management system between the ambient temperature ( $S_1 = T_A$ ), the wind speed ( $S_2 = v_{\text{Wind}}$ ), and the time of day ( $S_3 = t_{\text{Day}}$ ) with the net energy transfer to an external grid ( $T = P_{\text{grid}}$ ) as target variable reveals that most information that the sources carry about the target is redundant with the time of day. The bivariate PID atoms from Table II are subdivided into trivariate atoms according to Figure 10, and the estimated information carried by each atom  $\widehat{\Pi}(T : \alpha)$  as well as the cumulative atoms, i.e., the corresponding estimated redundant information  $\widehat{I}_{\cap}^{\text{sx}}(T : \alpha)$  are shown. The estimator was applied to the same  $N = 10^5$  data points as in Table II. The total mutual information is  $I(T : S_1, S_2, S_3) = 3.092$  bits, agreeing with the top-most redundancy in the lattice. The bold rows each represent the most informative trivariate atom from the decomposition of the corresponding bivariate atom. Note that all bold rows contain an atom which is in some way a redundancy with source variable three.

| Antichain $\alpha$ of trivariate atom | Contained in bivariate atom | Estimation results          |  |
|---------------------------------------|-----------------------------|-----------------------------|--|
|                                       |                             | $\widehat{\Pi}(T : \alpha)$ | $\widehat{I}_{\cap}^{\text{sx}}(T : \alpha)$ |
| $\{1, 2, 3\}$                         | -                           | 0.429 bits                  | 3.092 bits                                   |
| $\{1, 3\}$                            | -                           | 0.688 bits                  | 2.440 bits                                   |
| $\{2, 3\}$                            | -                           | 0.223 bits                  | 1.971 bits                                   |
| $\{1, 3\}\{2, 3\}$                    | -                           | 0.177 bits                  | 1.742 bits                                   |
| $\{3\}$                               | -                           | -0.749 bits                 | 1.239 bits                                   |
| $\{1, 2\}$                            | $\Pi_{\text{syn}}$          | 0.000 bits                  | 2.330 bits                                   |
| $\{1, 2\}\{1, 3\}$                    | $\Pi_{\text{syn}}$          | 0.013 bits                  | 2.324 bits                                   |
| $\{1, 2\}\{2, 3\}$                    | $\Pi_{\text{syn}}$          | 0.006 bits                  | 2.320 bits                                   |
| $\{3\}\{1, 2\}$                       | $\Pi_{\text{syn}}$          | <b>1.053 bits</b>           | 1.988 bits                                   |
| $\{1, 2\}\{1, 3\}\{2, 3\}$            | $\Pi_{\text{syn}}$          | 0.29 bits                   | 2.314 bits                                   |
| $\{1\}$                               | $\Pi_{\text{unq},1}$        | -0.003 bits                 | 0.927 bits                                   |
| $\{1\}\{2, 3\}$                       | $\Pi_{\text{unq},1}$        | 0.022 bits                  | 0.930 bits                                   |
| $\{1\}\{3\}$                          | $\Pi_{\text{unq},1}$        | <b>0.278 bits</b>           | 0.905 bits                                   |
| $\{2\}$                               | $\Pi_{\text{unq},2}$        | 0.000 bits                  | 0.664 bits                                   |
| $\{2\}\{1, 3\}$                       | $\Pi_{\text{unq},2}$        | 0.004 bits                  | 0.664 bits                                   |
| $\{2\}\{3\}$                          | $\Pi_{\text{unq},2}$        | <b>0.03 bits</b>            | 0.657 bits                                   |
| $\{1\}\{2\}$                          | $\Pi_{\text{red}}$          | 0.003 bits                  | 0.630 bits                                   |
| $\{1\}\{2\}\{3\}$                     | $\Pi_{\text{red}}$          | <b>0.627 bits</b>           | 0.627 bits                                   |

continuous versions of certain logic gates. Furthermore, we provide a nearest-neighbor based estimator and show that it converges to the numerically obtained values on the before-mentioned examples. Lastly, we demonstrate how this estimator can be used in a practical example by using it to uncover variable dependencies in a simulated energy management system.

The proposed measure fills a gap not addressed by existing continuous PID measures. The space of PID measures can be subdivided into mainly two groups: Those based on optimization of an auxiliary variable with respect to some ordering criterion [29, 48, 49] as explained in the framework of Kolchinsky [29], and those which have a *pointwise* definition, i.e., which can be defined for an individual realization in a self-contained functional way dependent solely on the probability densities at that realization. Kolchinsky [29] highlights that for ordering-based PID measures, the inclusion-exclusion-rule for defining PID atoms is inconsistent and thus cannot be applied. However, being defined for local realizations, our proposed measure is fully compatible to the mereological framework of Gutknecht et al. [5] and is therefore not affected by this criticism.

Furthermore, it is crucial to remark that mutual information as well as PID are measures of statistical dependence only, and therefore do not imply any causation in the interventional sense. Note that a causality decomposition into causal atoms must have a different struc-

ture from the partial information decomposition lattice as causality does not follow the same monotonicity criterion as redundancy [50]. This is because multiple causal influences can hinder and prevent each other: Consider, for instance, two stones being thrown at a bottle. While each stone individually would hit and destroy the bottle, their trajectories may intersect in such a way that the stones themselves collide when thrown simultaneously, diverting each other enough to both miss the bottle. In this case, two causes which would have individually been sufficient to cause an event have cancelled each other, which is an interaction that is impossible when considering redundancy instead.

As outlined in Section II B, the definition of continuous  $I_{\cap}^{\text{sx}}$  leaves open a choice of how to set the relative scale between variables that needs to be made in accordance with the setup in question. While such a choice marks a departure from the purely model-free nature of classical information theory, note that this is likely unavoidable in a PID context. For PID itself, choosing one of the multiple competing PID measures according to their operational interpretation similarly introduces some notion of “proto-semantics” to the analysis. Overall, we see the choice of relative scale not as a weakness but as a necessity for a continuous PID measure, since what it means to be in a neighborhood of a value might differ drastically between different variables.

This work builds on the definition of the discrete  $I_{\cap}^{\text{sx}}$

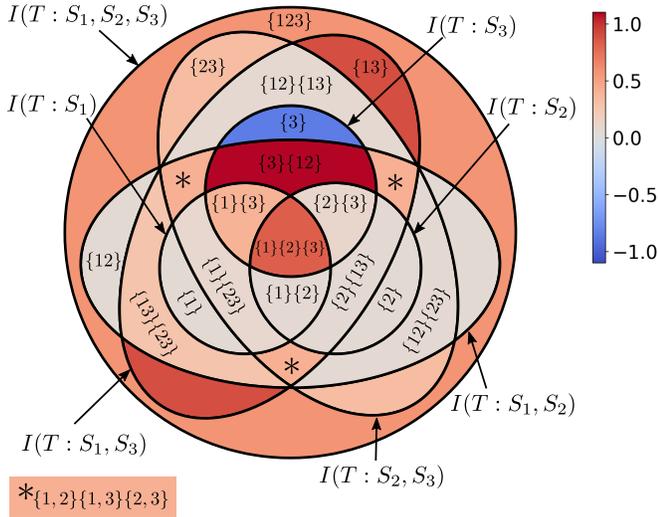


FIG. 11. **Most of the information in  $S_1$  (ambient temperature) and  $S_2$  (wind speed) is redundant with  $S_3$  (time of day).** Visualization of the results of Table III in the mereological diagram from Figure 10. Here, red encodes positive, i.e., informative, information contributions carried by the associated information atom, while blue encodes negative, i.e., misinformative contributions in bits.

measure [26] and its measure-theoretic extension to continuous variables [34] to introduce a tractable analytical definition and practical nearest-neighbor based estimator for a continuous  $I_{\bar{n}}^{\text{sx}}$  measure. However, this definition turns out not to be invariant with respect to bijective transformations of the individual variables—at least not without prescribing a preprocessing scheme (see Appendix C). Yet, this invariance is a property central to the measure-theoretic derivation in [34]. For this reason, the analytic definition presented here cannot be directly interpreted as a concretization of the latter but should rather be seen as a practical implementation rooted in the intuitive meaning of the  $I_{\bar{n}}^{\text{sx}}$  measure.

Like in the discrete case, this definition of redundancy might yield negative information atoms. In the context originally introduced by Makkeh et al. [26], namely for memory-less agents and interpreting the average in Equation (5) as an ensemble average, these are easily interpreted as misinforming the agents of the ensemble on average. For time-series data, for which the averages are taken over time, however, negative atoms can be a distracting technicality hindering a straightforward interpretation of the results. In these cases, practitioners may opt to clamp the atoms to non-negative values and adjust the other atoms accordingly so that the consistency equations (Equations (16) for the general case) continue to hold.

In this paper, we propose a novel shared-exclusion based PID measure designed specifically for purely continuous variables. However, it is important to acknowledge that real-world systems often involve a combination of continuous and discrete source variables. Addi-

tionally, even certain individual variables may comprise both discrete and continuous parts. To address these scenarios, we introduce an analytical formulation in Appendix K, outlining an ansatz for handling mixed systems. While we defer the presentation of specific examples demonstrating its application and the development of a mixed continuous-discrete estimator to future research, this work lays the groundwork for a comprehensive understanding of information decomposition in systems with diverse variable types.

Overall, with this paper we introduce a new analytical formulation for a shared-exclusion based continuous PID measure and provide an estimator that makes it applicable to practitioners from all sciences and technology.

## ACKNOWLEDGEMENTS

We would like to thank Anja Sturm for fruitful discussions and valuable feedback on this paper. We further want to thank everyone in the Wibrall lab and the Honda Research Institute for their feedback and indirect contributions.

DE, KSP, AM and MW are employed at the Göttingen Campus Institute for Dynamics of Biological Networks (CIDBN). DE and MW were supported by funding from the Ministry for Science and Education of Lower Saxony and the Volkswagen Foundation through the “Niedersächsisches Vorab” under the program “Big Data in den Lebenswissenschaften” – project “Deep learning techniques for association studies of transcriptome and systems dynamics in tissue morphogenesis”. KSP, FL and PW are funded by the Honda Research Institute Europe GmbH.

## Appendix A: On the shared-exclusions measure of redundancy

The definition of bivariate  $I_{\bar{n}}^{\text{sx}}$  in Equation (4) can also be derived from shared exclusions of parts of the underlying abstract probability space  $(\Omega, \mathbb{P})$ , i.e., the parts of  $(\Omega, \mathbb{P})$  that are rendered impossible by observation of either variable. To determine the precise subsets of the probability space to be excluded, we assume that a realization  $\mathbf{S} = (S_1, S_2) = \mathbf{s} = (s_1, s_2)$  has been taken by the system and we have gained the information that either  $S_1 = s_1$  or  $S_2 = s_2$  have materialized. Through this observation, the part of the probability space corresponding to any other realization  $\mathbf{s}' = (s'_1, s'_2)$  is rendered impossible unless either  $s'_1 = s_1$  or  $s'_2 = s_2$ . If neither of those equations holds, then  $\mathbf{s}'$  would not fulfill the logical or-statement in Equation (4) and would hence be impossible under the assumption (see Figure 12).

When parts of the probability space are rendered impossible by observing a certain event, the resulting posterior probability measure  $\mathbb{P}$  needs to be normalized on  $\Omega$ , yielding an adapted probability measure  $\mathbb{P}'$ , known

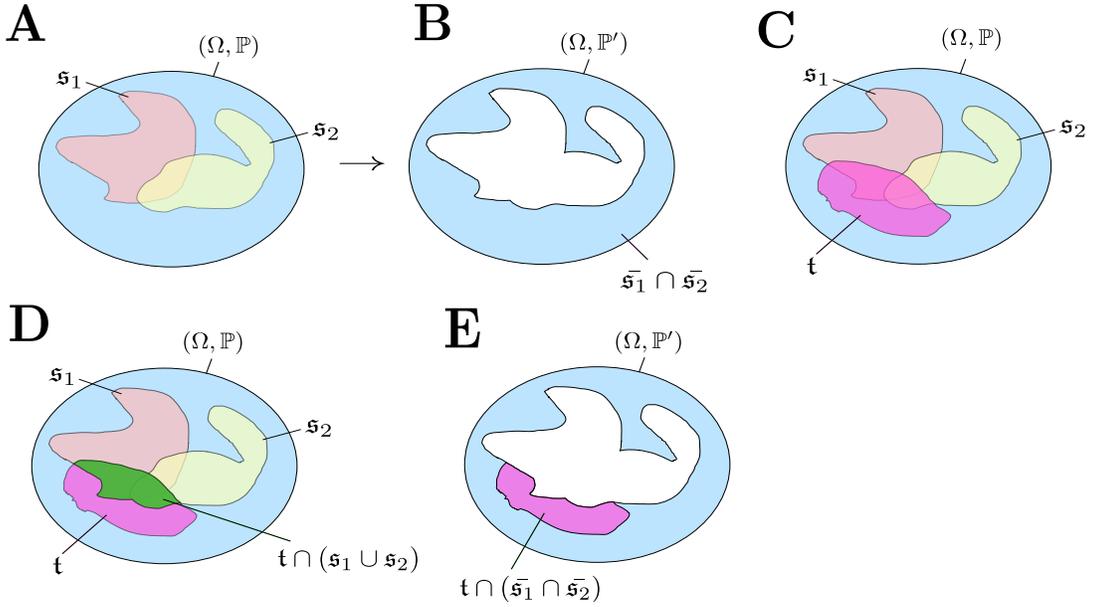


FIG. 12. **The pointwise redundancy  $I_{\cap}^{\text{sx}}$  can be derived from principles of probability exclusion.** Local redundant information for two sources  $S_1$  and  $S_2$  with respect to target  $T$  at a joint realization  $(s_1, s_2, t)$ . Each realization corresponds to one of the events in the probability space denoted by  $s_1, s_2$  and  $t$  in the probability space  $(\Omega, \mathbb{P})$ . Here we aim at assessing the benefit of knowing that either  $S_1 = s_1$  or  $S_2 = s_2$  in predicting  $t$ . In panel **A** the events  $s_1$  and  $s_2$  are displayed. In panel **B** there is the intersection of both complements from the events in panel **A**, which is the area of probability mass which either event excludes. In panel **C** we consider a target variable and its event set  $t$  in  $\Omega$ , while panel **D** shows the area of interest, the green set represents the part of event  $t$  that is included in also either  $s_1$  or in  $s_2$  which corresponds to the part of the probability of realization  $t$ , which is accessible when assuming that the logical statement  $S_1 = s_1 \vee S_2 = s_2$  is true, or, correspondingly,  $W_{\{1\},\{2\}} = 1$ . While the green area can be measured via the probability measure directly, i.e.,  $\mathbb{P}(t \cap (s_1 \cup s_2))$ , we can similarly represent it using the shared exclusion overlapping with  $t$  as in panel **E**, by  $\mathbb{P}'(t) - \mathbb{P}'(t \cap (\bar{s}_1 \cap \bar{s}_2)) = \frac{\mathbb{P}(t) - \mathbb{P}(t \cap (\bar{s}_1 \cap \bar{s}_2))}{1 - \mathbb{P}(s_1 \cup s_2)}$ . This green area then has to be compared to the regular probability of obtaining the realization  $t$  without knowledge of any of the sources,  $\mathbb{P}(t)$ .

as the conditional probability measure. The visualization Figure 12 shows the exclusion of specific parts of the probability space leading to descriptions of the information that can be assessed from either  $S_1$  or  $S_2$  locally about  $T$ .

Similarly, this concept can then be generalized to collections  $\alpha = \{\mathbf{a}_j\}_j = \{\{s_{ij}\}_i\}_j$  of three or more source variables by excluding the part of the probability space that is redundantly excluded by all joint collections of random variables  $a_j$ . For instance, in a quadrivariate (i.e., four source variable) system, the redundancy between the collections  $\{S_1, S_3\}$  and  $\{S_2, S_4\}$  is defined by the joint exclusions between the collections, which is given by all realizations  $\mathbf{s}' = (s'_1, s'_2, s'_3, s'_4)$  for which neither  $(s'_1, s'_3) = (s_1, s_3)$  nor  $(s'_2, s'_4) = (s_2, s_4)$  holds. This generalization to collections of subsets of the source variables turns out to be of the form of local mutual information with respect to an auxiliary random variable

$$W_{\Phi}(\mathbf{s}) = \bigvee_{a \in \mathcal{P}(\mathbf{S}) : \Phi(a)=1} \bigwedge_{i \in a} (S_i = s_i)$$

and reverts to a classical mutual information quantity for *self-redundancies*, i.e.,  $\alpha = \{\mathbf{a}\}$ . Further, its dependence on elements of the power set of the index set of

sources (set of all collections of sources) implies that the  $I_{\cap}^{\text{sx}}$  measure of redundant information gives rise to a full lattice of PID atoms as described in Section III A. Note also that this measure, as a composition of differentiable functions, is differentiable itself with respect to minor perturbations of the underlying probability mass function. Additionally, it fulfills a target chain rule [26].

### Appendix B: Construction of densities of logical statements

In this appendix section we argue for the explicit form of the quasi-densities appearing in  $I_{\cap}^{\text{sx}}$ , which stem from a logical statement allowing also for disjunctions, and not only conjunctions as is usual in traditional probability theory.

Let us recall that when measuring a union of two sets, by the law of inclusion-exclusion, we can express the result as the following sum:

$$\mathbb{P}^{S_1, S_2}(B \cup C) = \mathbb{P}^{S_1, S_2}(B) + \mathbb{P}^{S_1, S_2}(C) - \mathbb{P}^{S_1, S_2}(B \cap C).$$

Similarly to Appendix A, a union can be utilized to represent the logical disjunction  $S_1 \in B \vee S_2 \in C$  via mea-

asuring the set  $B \times E_{S_2} \cup E_{S_1} \times C$  where  $E_{S_i}$  denotes the respective space of possible realizations for  $S_i$ . Measuring  $B \times E_{S_2} \cup E_{S_1} \times C$  then yields

$$\begin{aligned} & \mathbb{P}^{S_1, S_2}(B \times E_{S_2} \cup E_{S_1} \times C) \\ &= \mathbb{P}^{S_1}(B) + \mathbb{P}^{S_2}(C) - \mathbb{P}^{S_1, S_2}(B \cap C), \end{aligned}$$

i.e., by using the marginal measures, just as it is done in discrete settings. Suppose now that we know that the statement  $S_1 \in B \vee S_2 \in C$  is true. How would we measure the probability of additional events under this presumption? Assuming the event  $S_1 \in B \vee S_2 \in C$  is true means that we need to use a conditional probability measure to measure additional events  $A \in E_{S_1} \times E_{S_2}$ . Thus, to obtain a probability measure under the assumption that  $S_1 \in B \vee S_2 \in C$  is true, we can write

$$\begin{aligned} \mathbb{P}_{S_1 \in B \vee S_2 \in C}^{(S_1, S_2)}(A) &:= \mathbb{P}^{S_1}(A \cap B) + \mathbb{P}^{S_2}(A \cap C) \\ &\quad - \mathbb{P}^{(S_1, S_2)}(A \cap (B \times C)). \end{aligned}$$

with  $C_\epsilon \xrightarrow{\epsilon \searrow 0} C$ , where  $C$  is a constant. This can easily be seen when using the triangle inequality and recalling

$$\left\| \mathbb{P}_{(S_1 \in B_{s_1}(\epsilon)) \vee (S_2 \in B_{s_2}(\epsilon))}^{(S_1, S_2)}(A) - \epsilon \cdot \mathbb{P}_{(S_1 \in \{s_1\}) \vee (S_2 \in \{s_2\})}^{(S_1, S_2)}(A) \right\| \leq C_\epsilon \cdot \epsilon \xrightarrow{\epsilon \searrow 0} 0 \quad \forall A \in \mathcal{E},$$

Applying this to infinitesimally small sets  $B = \{s_1\}$  and  $C = \{s_2\}$ , we can achieve the same form by adding slices of conditional measures with respect to the constraints set by the disjunction. We define

$$\begin{aligned} & \mathbb{P}_{(S_1=s_1) \vee (S_2=s_2)}^{(S_1, S_2)}(A) := \mathbb{P}^{S_1}(\{x | (x, y) \in A, y = s_2\}) \\ & + \mathbb{P}^{S_2}(\{y | (x, y) \in A, x = s_1\}) - \mathbb{P}^{(S_1, S_2)}(\{(s_1, s_2)\} \cap A) \end{aligned}$$

for an event  $A$ .

Note that when taking the limit to infinitesimal sets, a choice needs to be made regarding the relative rate of convergence of the two variables. Choosing to interpret the two variables as being on equal scales, i.e., using the same  $\epsilon$  to search in the respective marginal neighborhoods, this uniquely defines the measure as the only one fulfilling

the form of a conditional probability measure, i.e.,

$$\begin{aligned} \mathbb{P}^{S_1, S_2}(\{(x, y) \in A | y \in B_{s_2}(\epsilon)\}) &= \underbrace{\frac{\mathbb{P}^{S_1, S_2}(\{(x, y) \in A | y \in B_{s_2}(\epsilon)\})}{\mathbb{P}^{S_2}(B_{s_2}(\epsilon))}}_{\xrightarrow{\epsilon \searrow 0} \mathbb{P}^{S_1}(\{(x, y) \in A | y = s_2\})} \underbrace{\frac{\mathbb{P}^{S_2}(B_{s_2}(\epsilon))}{\lambda_{S_2}(B_{s_2}(\epsilon))}}_{\xrightarrow{\epsilon \searrow 0} f_{S_2}(s_2)} \underbrace{\lambda_{S_2}(B_{s_2}(\epsilon))}_{2\epsilon} \\ &\stackrel{\epsilon \ll 1}{\approx} 2\epsilon \mathbb{P}^{S_1}(\{(x, y) \in A | y = s_2\}) f_{S_2}(s_2). \end{aligned}$$

This idea is further visualized in Figure 13.

Evidently, this ansatz fails if any of the marginal densities  $f_{S_i}$  is unbounded at  $s_i$ . The choice of this measure could also be argued for as a consequence of the simple function approximation, found for instance in [51], thrm. 7.10.

To arrive at a probability for an infinitesimal event, compare the above to the construction of the usual Lebesgue measure  $\lambda_S$  measuring the maximal size of nicely shrinking balls  $B_i = B_{(s_1, s_2)}(r_i)$  [51]. Note that the maximal size of a  $B_i$  is the diameter of the ball,  $2r_i$ , and in the continuous case, the last term in the numerator vanishes.

Then

$$\begin{aligned} & f_{S_1 \vee S_2}(s_1, s_2) \\ &= \lim_{B_i \searrow \{(s_1, s_2)\}} \frac{\mathbb{P}_{(S_1=s_1) \vee (S_2=s_2)}^{(S_1, S_2)}(B_{(s_1, s_2)}(r_i))}{\lambda_S(B_{(s_1, s_2)}(r_i))} \\ &= \lim_{r_i \searrow 0} \frac{\mathbb{P}^{S_1}(\{(x, y) \in B_i | y = s_2\})}{2r_i} \\ &\quad + \frac{\mathbb{P}^{S_2}(\{(x, y) \in B_i | x = s_1\})}{2r_i} \\ &= \lim_{r_i \searrow 0} \frac{\mathbb{P}^{S_1}(\{s_1 \pm r_i\}) + \mathbb{P}^{S_2}(\{s_2 \pm r_i\})}{2r_i} \\ &= f_{S_1}(s_1) + f_{S_2}(s_2), \end{aligned}$$

where  $\lambda_{S_1}(\{(x, y) \in B_i | y = s_2\}) + \lambda_{S_2}(\{(x, y) \in B_i | x =$

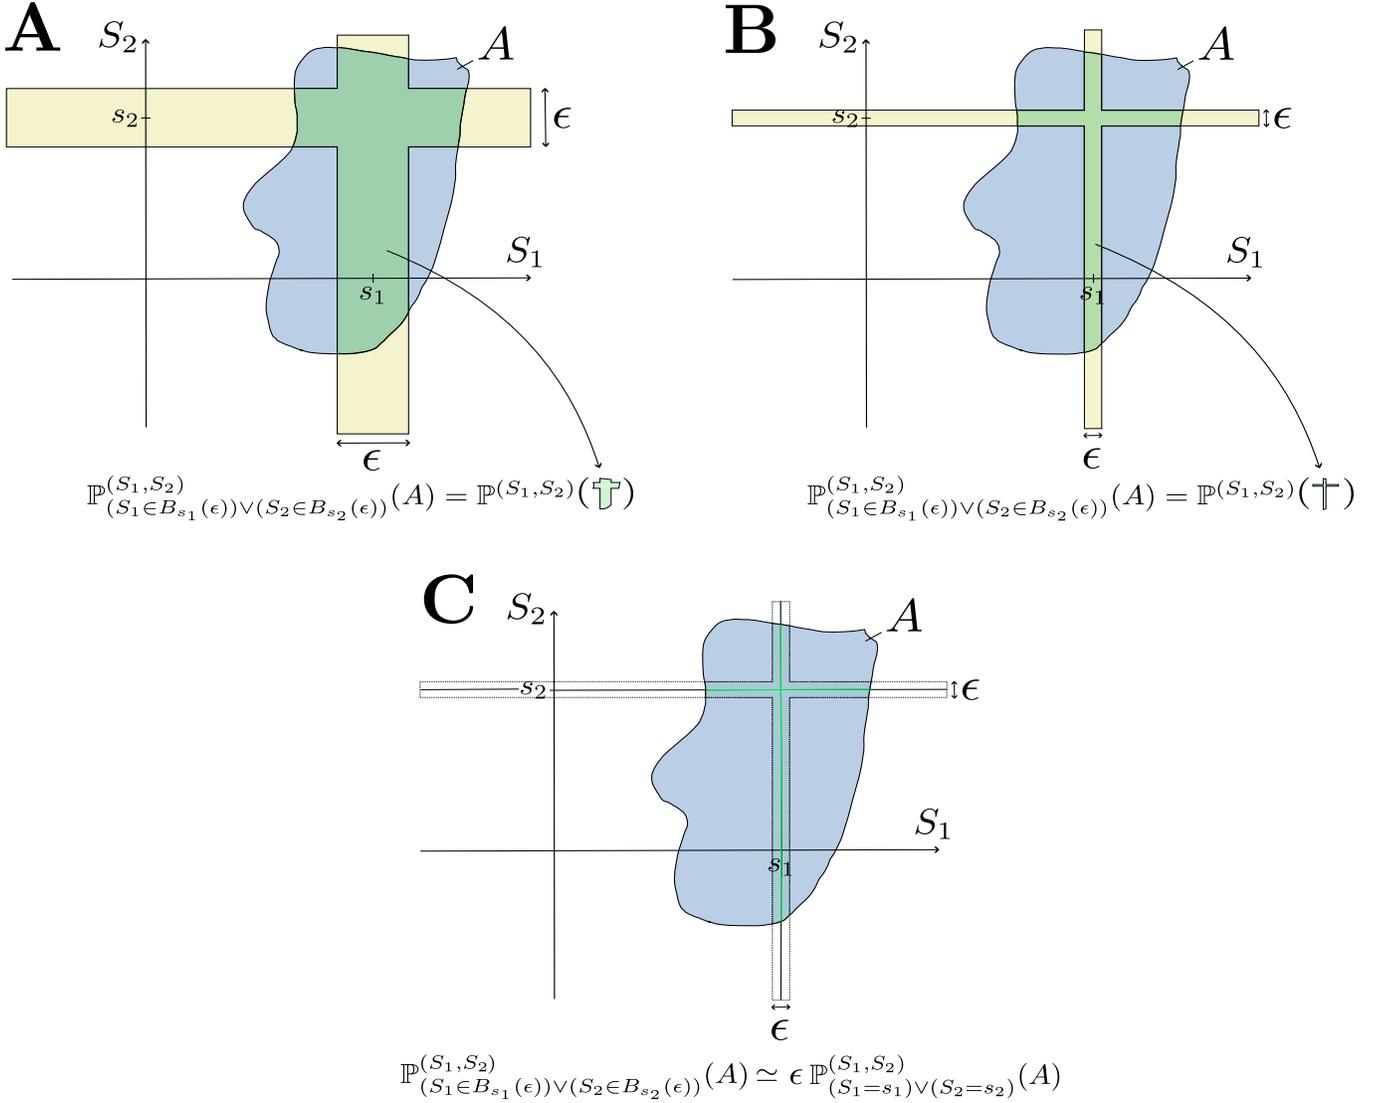


FIG. 13. We choose our localized conditional probability measure such that its extension for small neighbourhoods naturally results in the global probability measure. In panel **A** the intuitive meaning of the measure  $\mathbb{P}_{(S_1 \in B_{s_1}(\epsilon)) \vee (S_2 \in B_{s_2}(\epsilon))}^{(S_1, S_2)}(A)$  for a set  $A$  is illustrated. In panel **B** one sees the consequences of  $\epsilon$  decreasing, leading to  $\mathbb{P}_{(S_1 \in B_{s_1}(\epsilon)) \vee (S_2 \in B_{s_2}(\epsilon))}^{(S_1, S_2)}(A)$  approaching the conditional measure of our choice,  $\mathbb{P}_{(S_1 \in \{s_1\}) \vee (S_2 \in \{s_2\})}^{(S_1, S_2)}(A)$  multiplied with the width of the neighborhood,  $\epsilon$ , as shown in panel **C**.

$s_1\}$  is equal to  $2r_i$  since the measures  $\lambda_{S_i}$  measure the diameter of the uniform ball centered around  $(s_1, s_2)$ .

This method is readily generalized to  $n$  sources in the exact same way it worked for two, with the one exception being that conjunctions are treated as one variable. For example, a logical statement  $\alpha(s_1, s_2, s_3) = \{S_1 =$

$s_1\} \wedge \{S_2 = s_2\} \vee \{S_3 = s_3\}$  is treated as a disjunction of a single two dimensional statement and a single one dimensional statement.

**Definition 3.** Let for a generic logical statement  $\alpha$  of length  $l(\alpha)$  in its disjunctive normal form  $\alpha = \bigvee_{j \in J} a_j$ ,  $I_d \subset J$  be the subset of purely discrete statements. Then we define the density of  $\alpha$  as

$$\frac{\delta \mathbb{P}_\vee}{\delta \lambda_\vee}(s_1, \dots, s_n) = f_\alpha(x_1, \dots, x_n) := \sum_{k=1}^{|I_d|+1} (-1)^{k+1} \left( \sum_{\substack{m_1 < m_2 < \dots < m_{k-1} \in I_d \\ m_1 \neq m_2 \neq \dots \neq m_{k-1} \neq j \in J}} f_{\bigcup_{q=1}^{k-1} a_{m_q} \cup a_j} \left( \bigcup_{q=1}^{k-1} x_{a_{m_q}} \cup x_{a_j} \right) \right).$$

**Example 1.** 1. Take the example  $\alpha = (X_1 \wedge X_2) \vee (X_3) \vee (X_4 \wedge X_5)$ , where  $X_1, X_2$  and  $X_4$  are discrete, and  $X_3$  and  $X_5$  are continuous, such that  $J = \{1, 2, 3\}$  and  $I_d = \{1\}$ . The resulting quasi-density is

$$\begin{aligned} & f_\alpha(x_1, x_2, x_3, x_4, x_5) \\ &= \underbrace{p_{X_1 X_2}(x_1, x_2) + f_{x_3}(x_3) + f_{X_4 X_5}(x_4, x_5)}_{k=1} \\ - & \left( \underbrace{f_{X_1 X_2 X_3}(x_1, x_2, x_3) + f_{X_1 X_2 X_4 X_5}(x_1, x_2, x_4, x_5)}_{k=2} \right). \end{aligned}$$

2. Choosing  $\alpha = (X_1 \wedge X_2) \vee (X_1 \wedge X_3)$  with  $X_1, X_3$  discrete and  $X_2$  continuous on the other hand, leads to

$$\begin{aligned} & f_\alpha(x_1, x_2, x_3) \\ &= f_{X_1 X_2}(x_1, x_2) + p_{X_1 X_3}(x_1, x_3) - f_{X_1 X_2 X_3}(x_1, x_2, x_3). \end{aligned}$$

The above derived form of the density of an OR statement readily only works for purely continuous measures, as for purely discrete ones;  $\epsilon^i \rightarrow 1$  and thus the last term does not vanish, yielding the expression commonly derived via the inclusion-exclusion rule, that is,

$$\frac{\mathbb{P}_\vee(A_i, B_i)}{\epsilon^i} \rightarrow p_X(x_A) + p_Y(y_B) - p_{XY}(x_A, y_B).$$

Thus those corner cases of purely discrete and purely continuous systems can be understood as they have comparable  $\epsilon_A^i = \epsilon_B^i$ , but if one considers the case of a system consisting of both types of variables, one cannot fulfill this requirement. To circumvent this nuisance while still measuring the maximal possible amount of probability, we define the density of an OR-statement in the measure-theoretic case to be

$$\begin{aligned} & \frac{\delta \mathbb{P}_\vee}{\delta \lambda_\vee}(x_A, y_B) \\ := & \lim_{i \rightarrow \infty} \frac{\mathbb{P}_X(\Delta A^i)}{\epsilon_A^i} + \frac{\mathbb{P}_Y(\Delta B^i)}{\epsilon_B^i} - \frac{\mathbb{P}_{XY}(\Delta A^i \times B_{i-1}) + \mathbb{P}_{XY}(A_{i-1} \times \Delta B^i) - \mathbb{P}_{XY}(\Delta A^i \times \Delta B^i)}{\epsilon_A^i \epsilon_B^i} \max(\epsilon_A^i, \epsilon_B^i) \end{aligned}$$

with  $\Delta A^i = A_i \setminus A_{i+1}$ , and the same for  $\Delta B^i = B_i \setminus B_{i+1}$ . With this definition, the special cases of purely discrete and purely continuous arise again, and the mixed case yields sums of probability masses and densities of the same order of magnitude with respect to the underlying space vanishing, i.e., in the case where  $X$  is continuous and  $Y$  is discrete, we find

$$\frac{\delta \mathbb{P}_\vee}{\delta \lambda_\vee}(x_A, y_B) = f_X(x_A) + p_Y(y_B) - f_{XY}(x_A, y_B).$$

Requiring the same to hold for higher order statements  $\alpha$ , this can easily be generalized utilizing the inclusion-exclusion law, leading to Definition 3.

### Appendix C: Variable preprocessing

This appendix chapter discusses the effect that different preprocessing schemes have on the results of the introduced PID definition. Different to the mutual information and the measure-theoretic existence proofs by [34], the analytic definition for redundancy suggested in this paper is not invariant under isomorphic mappings of individual source variables. At the core, the reason for this fact is that in the continuous case, the meaning of a logical disjunction turns out to be contingent on the relative scale of the two variables, since densities are defined as the limits of neighbourhoods, and there is no canonical way to compare neighborhoods between two different variables to make them shrink at the same pace.

Definition 3 has been made under the assumption that the two variables are on a comparable scale. While this

TABLE IV. Different preprocessing schemes for continuous PID gates produce qualitatively similar results on the redundant and unique gate, while a copula pretransformation results in different results for the copy and sum gates. The results have been computed using a Monte-Carlo integration analogous to Table I.

|                       | Redundant Gate | Copy Gate   | Unique Gate | Sum         |
|-----------------------|----------------|-------------|-------------|-------------|
| No transformation     |                |             |             |             |
| $\Pi_{\text{syn}}$    | 0.000 bits     | 6.644 bits  | 5.525 bits  | 6.497 bits  |
| $\Pi_{\text{unq},1}$  | 0.000 bits     | 0.000 bits  | 1.119 bits  | 0.147 bits  |
| $\Pi_{\text{unq},2}$  | 0.000 bits     | 0.000 bits  | -5.525 bits | 0.147 bits  |
| $\Pi_{\text{red}}$    | 6.644 bits     | 6.644 bits  | 5.525 bits  | 0.353 bits  |
| $I(T : S_1, S_2)$     | 6.644 bits     | 13.288 bits | 6.644 bits  | 7.144 bits  |
| Standardization       |                |             |             |             |
| $\Pi_{\text{syn}}$    | 0.000 bits     | 6.644 bits  | 5.524 bits  | 6.497 bits  |
| $\Pi_{\text{unq},1}$  | 0.000 bits     | 0.000 bits  | 1.119 bits  | 0.147 bits  |
| $\Pi_{\text{unq},2}$  | 0.000 bits     | 0.000 bits  | -5.524 bits | 0.147 bits  |
| $\Pi_{\text{red}}$    | 6.644 bits     | 6.644 bits  | 5.524 bits  | 0.353 bits  |
| $I(T : S_1, S_2)$     | 6.644 bits     | 13.288 bits | 6.644 bits  | 7.144 bits  |
| Copula transformation |                |             |             |             |
| $\Pi_{\text{syn}}$    | 0.000 bits     | 6.922 bits  | 5.671 bits  | 6.791 bits  |
| $\Pi_{\text{unq},1}$  | 0.000 bits     | -0.279 bits | 0.973 bits  | -0.147 bits |
| $\Pi_{\text{unq},2}$  | 0.000 bits     | -0.279 bits | -5.671 bits | -0.147 bits |
| $\Pi_{\text{red}}$    | 6.644 bits     | 6.922 bits  | 5.671 bits  | 0.647 bits  |
| $I(T : S_1, S_2)$     | 6.644 bits     | 13.288 bits | 6.644 bits  | 7.144 bits  |

holds true for many scientific questions in which the variables in question are distributed equivalently, in other scenarios a preprocessing step might become necessary to compare the variables in a way adequate for the research question. Table IV shows how the PID results of the logic gate examples in Table I change when different preprocessing schemes are applied. First, note that preprocessing does not impact the mutual information quantities. While there are no qualitative differences between the preprocessing schemes for the redundant and the unique gate, the copula transformation, i.e. changing the random variables to their own cumulative density making them uniform on the unit interval, introduces some unique information in the copy gate and negative unique information in the sum gate. These results show that while different preprocessing schemes might often lead to comparable results, it is nevertheless important to choose a preprocessing scheme that matches the scientific question to yield reliable and interpretable results.

Note, also, that specifically the preprocessing based on standardization makes no difference compared to no transformation for the given examples. This, however, is not true in general but only if the two source random variables already have the same standard deviation, as is the case for all examples shown here.

#### Appendix D: Proof of the differentiability of $I_{\cap}^{\text{sx}}$

The analytical definition of the redundancy (see Definition 2) allows us to prove its differentiability and even smoothness for both the local and global measure.

**Theorem 1.** The local continuous measure of shared information

$$\begin{aligned} i_{\cap}^{\text{sx}}[f](t, \{s_1\}\{s_2\}) &= \log_2 \left[ \frac{f_{T|S_1 \vee S_2}(t|s_1 \vee s_2)}{f_T(t)} \right] \\ &= \log_2 \left[ \frac{f_{TS_1}(t, s_1) + f_{TS_2}(t, s_2)}{f_T(t)(f_{S_1}(s_1) + f_{S_2}(s_2))} \right] \end{aligned}$$

varies smoothly with respect to changes of the underlying joint probability density  $f_{TS_1S_2}$ . Moreover, for more than two source variables,  $i_{\cap}^{\text{sx}}[f_{TS_1S_2}](t : \alpha)$  is smooth for arbitrary antichains  $\alpha$ .

*Proof.* Let  $f_{TS_1S_2}$  be a density function and  $(t, s_1, s_2)$  a point in state space such that  $f_{TS_1S_2}(t, s_1, s_2) > 0$ . Now, let  $g_{TS_1S_2}$  be a function such that  $g_{TS_1S_2}(t, s_1, s_2) \neq 0$ .

Then we construct the derivative  $D(i_{\cap}^{\text{sx}}[f])g$  in the direction of  $g$  by considering the limit of a shifted density  $f_{TS_1S_2} \mapsto f_{TS_1S_2} + \epsilon \cdot g_{TS_1S_2}$  with  $\epsilon > 0$ , yielding

$$\begin{aligned}
D(i_{\bar{\cap}}^{\text{sx}}[f])g &= \lim_{\epsilon \rightarrow 0} \frac{i_{\bar{\cap}}^{\text{sx}}[f_{TS_1S_2} + \epsilon \cdot g_{TS_1S_2}](t : s_1 \vee s_2) - i_{\bar{\cap}}^{\text{sx}}[f_{TS_1S_2}](t : s_1 \vee s_2)}{\epsilon} \\
&= \lim_{\epsilon \rightarrow 0} \frac{\log \left[ \frac{f_{TS_1}(t, s_1) + f_{TS_2}(t, s_2) + \epsilon(g_{TS_1}(t, s_1) + g_{TS_2}(t, s_2))}{(f_T(t) + \epsilon g_T(t))(f_{S_1}(s_1) + f_{S_2}(s_2) + \epsilon(g_{S_1}(s_1) + g_{S_2}(s_2)))} \right] - \log \left[ \frac{f_{TS_1}(t, s_1) + f_{TS_2}(t, s_2)}{f_T(t)(f_{S_1}(s_1) + f_{S_2}(s_2))} \right]}{\epsilon} \\
&= \lim_{\epsilon \rightarrow 0} \frac{\log_2 \left[ 1 + \epsilon \frac{g_{TS_1}(t, s_1) + g_{TS_2}(t, s_2)}{f_{TS_1}(t, s_1) + f_{TS_2}(t, s_2)} \right] - \log_2 \left[ 1 + \epsilon \frac{g_{S_1}(s_1) + g_{S_2}(s_2)}{f_{S_1}(s_1) + f_{S_2}(s_2)} \right] - \log_2 \left[ 1 + \epsilon \frac{g_T(t)}{f_T(t)} \right]}{\epsilon}.
\end{aligned}$$

At this point, we use the Taylor expansion  $\ln[1+x] = \sum_{n=1}^{\infty} \frac{(-1)^{n+1}x^n}{n} = x + \mathcal{O}(x^2)$  to determine the final form of

the derivative

$$D(i_{\bar{\cap}}^{\text{sx}}[f])g = \frac{g_{TS_1}(t, s_1) + g_{TS_2}(t, s_2)}{f_{TS_1}(t, s_1) + f_{TS_2}(t, s_2)} - \frac{g_{S_1}(s_1) + g_{S_2}(s_2)}{f_{S_1}(s_1) + f_{S_2}(s_2)} - \frac{g_T(t)}{f_T(t)},$$

which is finite by assumption.<sup>1</sup>

Considering the form that results from the derivative, it can be concluded that the  $i_{\bar{\cap}}^{\text{sx}}[f_{TS_1S_2}](t : s_1 \vee s_2)$  is not only differentiable at point  $(t, s_1, s_2)$ , but even smooth, since the following derivatives will be linear combinations of powers of terms of the form  $\frac{1}{f_{TS_1} + f_{TS_2}}$ ,  $\frac{1}{f_{S_1} + f_{S_2}}$  and  $\frac{1}{f_T}$ .

Note further that this proof naturally extends to any arbitrary antichain  $\alpha$  of interest by an analogous argument.  $\square$

**Corollary 1.** The global measure of redundant information  $I_{\bar{\cap}}^{\text{sx}}[f] = \int f i_{\bar{\cap}}^{\text{sx}}[f]$  is a smooth functional of the underlying density  $f$  if  $f$  and its marginals are bounded from below by some  $\xi > 0$ .

*Proof.* The change of  $I_{\bar{\cap}}^{\text{sx}}$  when shifting  $f$  infinitesimally in the direction of  $g$  is given by

$$\left[ \frac{dI_{\bar{\cap}}^{\text{sx}}[f + \epsilon g]}{d\epsilon} \right]_{\epsilon=0} = \int dt ds_1 ds_2 \frac{\delta I_{\bar{\cap}}^{\text{sx}}[f]}{\delta f(t, s_1, s_2)} g(t, s_1, s_2).$$

where  $\delta I_{\bar{\cap}}^{\text{sx}}[f]/\delta f(t, s_1, s_2)$  is the functional derivative of  $I_{\bar{\cap}}^{\text{sx}}$  [52]. By applying a product rule of differentiation to the functional

$$I_{\bar{\cap}}^{\text{sx}}[f] = \int dt ds_1 ds_2 f(t, s_1, s_2) i_{\bar{\cap}}^{\text{sx}}[f](t, s_1, s_2),$$

which depends on  $f$  in two places, we obtain

$$\begin{aligned}
&\left[ \frac{dI_{\bar{\cap}}^{\text{sx}}[f + \epsilon g]}{d\epsilon} \right]_{\epsilon=0} \\
&= \left[ \frac{d}{d\epsilon} \int dt ds_1 ds_2 (f + \epsilon g)(t, s_1, s_2) i_{\bar{\cap}}^{\text{sx}}[f + \epsilon g](t, s_1, s_2) \right]_{\epsilon=0} \\
&= \int dt ds_1 ds_2 \{ i_{\bar{\cap}}^{\text{sx}}[f](t, s_1, s_2) g(t, s_1, s_2) \\
&\quad + f(t, s_1, s_2) [D(i_{\bar{\cap}}^{\text{sx}}[f])g](t, s_1, s_2) \} \\
&= \int dt ds_1 ds_2 [\{ i_{\bar{\cap}}^{\text{sx}}[f] + f D(i_{\bar{\cap}}^{\text{sx}}[f]) \} g](t, s_1, s_2),
\end{aligned}$$

from which we can read off the functional derivative as the functional

$$\frac{\delta I_{\bar{\cap}}^{\text{sx}}[f]}{\delta f(t, s_1, s_2)} = i_{\bar{\cap}}^{\text{sx}}[f] + f D(i_{\bar{\cap}}^{\text{sx}}[f]).$$

If the marginals of the joint density  $f$  are bounded from below, the magnitude of the Fréchet derivative,  $|D_g(i_{\bar{\cap}}^{\text{sx}}[f])|$ , is bounded, too. Hence the functional derivative is well-defined for all points  $(t, s_1, s_2)$ . Thus the total change of the functional  $I_{\bar{\cap}}^{\text{sx}}[f]$  with respect to an infinitesimal change of  $f$  in the direction of  $g$  exists if and only if both  $i_{\bar{\cap}}^{\text{sx}}[f]g$  and  $f D(i_{\bar{\cap}}^{\text{sx}}[f])g$  are integrable functions over the space of realizations.  $\square$

<sup>1</sup> Interpreted as a linear functional  $D(i_{\bar{\cap}}^{\text{sx}}[f])$  acting on  $g$ , this expression can be seen to be the Fréchet derivative of the local redundancy fulfilling [52]

$$\lim_{\| \epsilon g \| \rightarrow 0} \frac{\| i_{\bar{\cap}}^{\text{sx}}[f + \epsilon g] - i_{\bar{\cap}}^{\text{sx}}[f] - D(i_{\bar{\cap}}^{\text{sx}}[f])g \|}{\| \epsilon g \|} = 0.$$

## Appendix E: Analytical probability distributions for the logic gates and sum examples

The probability densities for the three continuous logic gates as well as the sum example can be expressed as multivariate Gaussian distributions of the form

$$p_{T,S_1,S_2}(t, s_1, s_2) = \frac{1}{\sqrt{(2\pi)^l \det \Sigma}} \exp\left\{-\frac{1}{2} \mathbf{x}^T \Sigma^{-1} \mathbf{x}\right\},$$

where  $\mathbf{x} = (t \ s_1 \ s_2)^T$  and the covariance matrix

$$\Sigma = \begin{pmatrix} \Sigma_T & \Sigma_{TS} \\ \Sigma_{ST} & \Sigma_S \end{pmatrix}$$

depends on the gate. Note, further, that the integer  $l$  is equal to 3 for the redundant and unique gate as well as the sum example while it is equal to 4 for the copy gate because of the two-dimensional target.

An analytical formulation for the mutual information for such Gaussian models is given by [3]

$$I(T, \mathbf{S}) = \frac{1}{2} \log_2 \left( \frac{\det \Sigma_T \det \Sigma_S}{\det \Sigma} \right).$$

To make the mutual information finite, noise with standard deviation  $\sigma = 0.01$  is added to the target variable. While the choice of this parameter influences the absolute sizes of the PID atoms, it does not majorly affect the qualitative analysis of the logic gates (Table V). However, analogously to the KSG estimator for mutual information, a higher number of samples is necessary to correctly estimate the redundant information when the dependence between the variables is strong [35]. Furthermore, the constant  $\delta = 10^{-9}$  is introduced in the redundant gate for numerical reasons to avoid a singular matrix. It has been carefully chosen small enough to not affect the results up to the presented precision.

The covariance matrices are found as follows:

a. *Redundant gate*

$$\Sigma = \begin{pmatrix} 1 + \sigma^2 & 1 - \delta & 1 - \delta \\ 1 - \delta & 1 & 1 - \delta \\ 1 - \delta & 1 - \delta & 1 \end{pmatrix}$$

b. *Copy gate*

$$\Sigma = \begin{pmatrix} 1 + \sigma^2 & 0 & 1 & 0 \\ 0 & 1 + \sigma^2 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}$$

where  $l = 4$  and  $\mathbf{x} = (t_1, t_2, s_1, s_2)^T$ ,

c. *Unique gate*

$$\Sigma = \begin{pmatrix} 1 + \sigma^2 & 1 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}$$

d. *Continuous sum*

$$\Sigma = \begin{pmatrix} 2 + \sigma^2 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}.$$

## Appendix F: Formal derivation of the Kozachenko-Leonenko estimator for Shannon differential entropy

The proposed estimator for continuous  $I_{\cap}^{\text{sx}}$  is based on ideas of the KSG estimator, originally introduced by Kraskov et al. [35] in 2004, which is explained here for completeness. The main advancement of the KSG estimator was to apply the Kosachenko-Leonenko estimator [42] for (differential) entropy in a specific way, allowing for estimation of  $I(X; Y) = H(X) + H(Y) - H(X, Y)$ . Here we will continue in a very similar fashion.

Kosachenko and Leonenko estimate (differential) entropy on the grounds of the assumption that the density to consider for the entropy is locally constant around a measured data point  $x_i$ , i.e.,  $q_i(\epsilon) = \int_{B_\epsilon(x_i)} dx p_X(x) \approx \lambda_X(B_\epsilon(x_i)) p_X(x_i)$ , where  $B_\epsilon(x_i)$  is the  $\epsilon$ -ball around  $x_i$ . Then Kosachenko and Leonenko introduce a probability density function  $P_k(\epsilon)$ , such that  $P_k(\epsilon)d\epsilon$  determines the probability of  $k - 1$  points lying within radius  $\epsilon$  around  $x_i$ , a  $k$ -th lying within the shell of radii within  $[\epsilon, \epsilon + d\epsilon]$  around  $x_i$ , and the remaining  $N - k - 1$  points lying outside.

Using well-known integral identities, we find

$$\int_{\mathbb{R}^+} d\epsilon \log[q_i(\epsilon)] P_k(\epsilon) = \psi(k) - \psi(N).$$

On the other hand, since  $p_X(x_i) \lambda_X(B_{\epsilon_i}(x_i)) \approx q_i(\epsilon_i)$ , it follows that

$$\log[p_X(x_i) \lambda_X(B_{\epsilon_i}(x_i))] \approx \mathbb{E}[\log[q_i]] = \psi(k) - \psi(N)$$

such that a point estimate  $\hat{H}$  of the differential entropy can be determined as

$$\begin{aligned} \hat{H}(X) &= -\frac{1}{N} \sum_{i=1}^N \log[\widehat{p_X}(x_i)] \\ &= -\psi(k) + \psi(N) + \frac{1}{N} \sum_{i=1}^N \log[\lambda_X(B_{\epsilon_i}(x_i))]. \end{aligned}$$

The volume bias terms  $\lambda_X(B_{\epsilon_i}(x_i))$  are numerically obtained by performing searches for the  $k$ -th nearest neighbor in the data set, the distance to whom is  $\epsilon_i$ .

TABLE V. The PID of the four toy examples does not qualitatively differ when varying the amount of noise  $\sigma$  that is added to the source variable  $T$ . The results have been computed using a Monte-Carlo integration analogous to Table I.

|                      | Redundant Gate | Copy Gate   | Unique Gate | Sum         |
|----------------------|----------------|-------------|-------------|-------------|
| $\sigma = 0.01$      |                |             |             |             |
| $\Pi_{\text{syn}}$   | 0.000 bits     | 6.644 bits  | 5.525 bits  | 6.497 bits  |
| $\Pi_{\text{unq},1}$ | 0.000 bits     | 0.000 bits  | 1.119 bits  | 0.147 bits  |
| $\Pi_{\text{unq},2}$ | 0.000 bits     | 0.000 bits  | -5.525 bits | 0.147 bits  |
| $\Pi_{\text{red}}$   | 6.644 bits     | 6.644 bits  | 5.525 bits  | 0.353 bits  |
| $I(T : S_1, S_2)$    | 6.644 bits     | 13.288 bits | 6.644 bits  | 7.144 bits  |
| $\sigma = 0.1$       |                |             |             |             |
| $\Pi_{\text{syn}}$   | 0.000 bits     | 3.322 bits  | 2.388 bits  | 3.186 bits  |
| $\Pi_{\text{unq},1}$ | 0.000 bits     | 0.000 bits  | 0.942 bits  | 0.143 bits  |
| $\Pi_{\text{unq},2}$ | 0.000 bits     | 0.000 bits  | -2.388 bits | 0.143 bits  |
| $\Pi_{\text{red}}$   | 3.329 bits     | 3.329 bits  | 2.388 bits  | 0.353 bits  |
| $I(T : S_1, S_2)$    | 3.329 bits     | 6.658 bits  | 3.329 bits  | 3.826 bits  |
| $\sigma = 0.001$     |                |             |             |             |
| $\Pi_{\text{syn}}$   | 0.000 bits     | 9.966 bits  | 8.822 bits  | 9.819 bits  |
| $\Pi_{\text{unq},1}$ | 0.000 bits     | 0.000 bits  | 1.143 bits  | 0.147 bits  |
| $\Pi_{\text{unq},2}$ | 0.000 bits     | 0.000 bits  | -8.822 bits | 0.147 bits  |
| $\Pi_{\text{red}}$   | 9.965 bits     | 9.966 bits  | 8.822 bits  | 0.353 bits  |
| $I(T : S_1, S_2)$    | 9.966 bits     | 19.932 bits | 9.966 bits  | 10.466 bits |

### Appendix G: Estimation steps

In Section IID we estimate  $\log[f_\alpha(t, s_i)]$  for  $\alpha = \bigvee_j a_j$  containing an OR statement by

$$\begin{aligned} \log[\widehat{f_\alpha^i}(t, s_k)] &= \log \left[ \widehat{\sum_j f_{a_j}^i(t, s_k)} \right] \\ &\approx \log \left[ \sum_j e^{\widehat{\log[f_{a_j}^i]}} \right] \end{aligned}$$

$$\stackrel{\text{KSG}}{=} -\psi(N) + \psi(k) + \log \left[ \sum_j v_{a_j, \epsilon}^i \right]$$

where the  $v_{a_j, \epsilon}^i$  each represent the volume terms in the estimation at the  $i$ -th point. To the end of arriving at a form which is treatable by considerations akin to the KSG estimator, we have intermediately used an expansion of the sort  $z = \exp[\log[z]]$ ,  $z \geq 0$ . Then, since by usage of the maximum norm  $v_{a_j, T, \epsilon}^i = v_{a_j, \epsilon}^i \cdot v_{T, \epsilon}^i$ , we find

$$\begin{aligned} \hat{I}_\cap^{sx} &= \langle \log[\widehat{f_{\alpha, T}^i}(t, s_k)] \rangle_i - \langle \log[\widehat{f_\alpha^i}(t, s_k)] \rangle_i - \langle \log[\widehat{f_T^i}(t, s_k)] \rangle_i \\ &= \psi(k) + \psi(N) - \langle \psi(n_\alpha(i)) \rangle_i - \langle \psi(n_T(i)) \rangle_i \end{aligned}$$

### Appendix H: Implementation details

Having introduced an estimator for continuous  $I_\cap^{sx}$  in Sections IIC and IIIC, this section explains how this estimator can be realized in an efficient algorithm.

In the estimation procedure, the most computationally expensive step is the search for  $k$  nearest neighbors for each sample point in the joint space and counting of neighbours within a ball in the marginal spaces. For the maximum norm, efficient  $\mathcal{O}(n \log n)$  algorithms can be found for both low- ( $kd$ -trees) and high-dimensional (ball tree) data. Furthermore, approximate procedures exist for scaling beyond what the exact algorithms can compute.

In order to determine the distance  $\epsilon$  to the  $k$ th nearest neighbor from a given disjunction statement, we proceed as follows: First, we obtain the  $k$  nearest neighbors in each of the sets of variables of the corresponding antichain. Then we do a merging procedure: To get the  $k$  closest points to any of the source variables, we successively take points with the smallest distance in any one of the subspaces until we have a total of  $k$  points. While the kd-tree algorithm could be adapted to directly find neighbours according to this custom distance function, using the merging-procedure we achieve the same runtime complexity and can use highly optimized existing kd-tree implementations. The pseudocode for this function is given in Algorithm 3.

Following an analogous merging procedure, the number of points within a radius of  $\epsilon$  from the disjunction statement in the source space (see Algorithm 4) as well as the number of points within a radius  $\epsilon$  from the query point in the target space (see Algorithm 5) is determined. From these results, the redundancy can readily be computed according to Algorithm 6.

We implemented the described algorithm as a python package using SciPy [53] for efficient nearest-neighbor searches. The package is publicly available under

[gitlab.gwdg.de/wibr/continuouspidestimator](https://gitlab.gwdg.de/wibr/continuouspidestimator)

TABLE VI. **PID results for skewed sum example** The results have been computed using a Monte-Carlo integration analogous to Table I.

|                      | Skewed Sum  |
|----------------------|-------------|
| $\Pi_{\text{syn}}$   | 9.884 bits  |
| $\Pi_{\text{unq},1}$ | 0.214 bits  |
| $\Pi_{\text{unq},2}$ | 0.082 bits  |
| $\Pi_{\text{red}}$   | 0.355 bits  |
| $I(T : S_1, S_2)$    | 10.535 bits |

### Appendix I: Role of the parameter $k$

The characteristics of the estimates produced by the procedure laid out in Section II C are influenced by the choice of the parameter  $k$ , which determines which neighbors are used to compute the distance  $\epsilon$  in the joint space by requiring that the  $k - 1$  closest neighbors are skipped. For larger  $k$ , these distances  $\epsilon$  will thus be larger as well, which has two important ramifications: Firstly, since the considered  $\epsilon$ -neighborhoods are larger and contain more points, fluctuations in the positions of individual points have less impact on the estimation leading to less variance. Secondly, however, since the estimator relies on the assumption that the probability density is approximately constant over the  $\epsilon$ -neighborhood, larger  $k$  for a fixed  $N$  may make the estimator unable to resolve high-frequency fluctuations which occur in the data. This may lead to a higher bias in the estimation for a given number of sam-

ples  $N$ . The estimation results do not depend on the parameter  $k$  alone but only on the fraction  $k/N$  (as visible in Figure 14 and Figure 15), as this determines the average distance  $\epsilon$ .

### Appendix J: Skewed sum

We define the covariance matrix for the skewed sum as

$$\Sigma = \begin{pmatrix} 2.2 + \sigma^2 & 1.2 & 1 \\ 1.2 & 1.2 & 0 \\ 1 & 0 & 1 \end{pmatrix}.$$

The results of the continuous PID analysis are summarized in table Table VI

### Appendix K: A Glimpse into Mixed systems of discrete and continuous variables

Although systems including both discrete and continuous variables do at this point in time exceed the capabilities of the estimator developed in this work, we still intend to share some theoretical aspects about potential future endeavours. Thus, we propose the following treatment of systems, which are composed of both discrete and continuous variables. Assuming a mixture of discrete and continuous sources  $S_j$  and an either discrete or continuous target  $T$ , we suggest utilizing the same exponential expansion as in the purely continuous case Definition 3. This leads to the approximation

$$\widehat{\log[f_\alpha^i]} \approx \log \left[ \sum_{k=1}^{|I_d|+1} (-1)^{k+1} \left( \sum_{\substack{m_1 < m_2 < \dots < m_{k-1} \in I_d \\ m_1 \neq m_2 \neq \dots \neq m_{k-1} \neq j \in J}} e^{\log \left[ f_{\bigcup_{q=1}^{k-1} a_{m_q} \cup a_j} \right]} \right) \right].$$

Here then we have a number of the possible combinations of statements  $\bigcup_{q=1}^{k-1} a_{m_q} \cup a_j$  containing at least one discrete variable (denoted by  $a_j$ ). The statements con-

sisting solely of continuous variables will be treated individually by KSG-like estimation while the others will be handled by isolating the discrete variables and conditioning on those, using the identity  $H(\{S_k\}_{k \in K}) - H(\{S_m\}_{m \in K \setminus L}) = H(\{S_l\}_{l \in L} | \{S_n\}_{n \in K \setminus L})$  for some index sets  $L \subset K$ . This leads to

**Algorithm 3:** compute\_εs

---

**Data:** Source samples  $(\mathcal{S}_j)_j$ ; Target samples  $(\mathcal{T}_j)_j$ ; Antichain  $\alpha \in \mathcal{P}(\mathcal{P}(\{1, \dots, n\}))$

**Result:** Distances from the disjunction to the  $k$ -th nearest neighbor in the joint source-target space.

```

// For each sample point indexed by j
for j ∈ (1, ..., n_samples) do
  // Find the k nearest neighbors in each
  // of the spaces S_a × T for all sets a
  // in the antichain α
  for a ∈ α do
    kNN_{a_j}^{ST} ←
    | find_k_nearest_neighbors(S_{a_j} × T_j)
  end
  // Merge points from all separate nearest
  // neighbor searches in the individual
  // spaces to get the k unique samples
  // with the smallest distances to their
  // respective query point j
  kNN_merged_j^{ST} ← merge({kNN_{a_j}^{ST} | a ∈ α})
  // Get search radius ε as the kth
  // smallest distance
  ε_j ← dist_to_kth(kNN_merged_j^{ST})
end
return (ε_j)_j

```

---

**Algorithm 4:** compute\_n\_α

---

**Data:** Source samples  $(\mathcal{S}_j)_j$ ; Antichain  $\alpha \in \mathcal{P}(\mathcal{P}(\{1, \dots, n\}))$ ; Distances from the disjunction to the  $k$ -th nearest neighbor in the joint source-target space  $(\epsilon_j)_j$

**Result:** Number of points in the source disjunction space with distance less than  $\epsilon_j$  from the disjunction

```

for j ∈ (1, ..., n_samples) do
  // Get nearest neighbors in ball of
  // radius ε in the individual conjunction
  // spaces
  for a ∈ α do
    NN_{a_j}^S ←
    | find_points_within_distance(S_{a_j}, ε_j)
  end
  // Count unique samples within any of the
  // epsilon balls for the individual
  // conjunction spaces
  n_{α_j} ← count_unique({NN_{a_j}^S | a ∈ α})
end
return (n_{α_j})_j

```

---

**Algorithm 5:** compute\_n\_T

---

**Data:** Target samples  $(\mathcal{T}_j)_j$ ; Antichain  $\alpha \in \mathcal{P}(\mathcal{P}(\{1, \dots, n\}))$ ; Distances from the disjunction to the  $k$ -th nearest neighbor in the joint source-target space  $(\epsilon_j)_j$

**Result:** Number of points in the target space with distance less than  $\epsilon_j$  from the disjunction

```

for j ∈ (1, ..., n_samples) do
  // Get nearest neighbors in ball of
  // radius ε in the target space
  NN_j^T ← find_points_within_distance(T_j, ε_j)
  // Count samples within the epsilon ball
  // for the target space
  n_{T_j} ← count(NN_j^T)
end
return (n_{T_j})_j

```

---

**Algorithm 6:** compute\_redundancy

---

**Data:** Source samples  $(\mathcal{S}_j)_j$ ; Target samples  $(\mathcal{T}_j)_j$ ; Antichain  $\alpha \in \mathcal{P}(\mathcal{P}(\{1, \dots, n\}))$

**Result:** Distances from the disjunction to the  $k$ -th nearest neighbor in the joint source-target space.

```

// Compute distances in joint space
ε ← compute_εs((S_j)_j, (T_j)_j)
// Compute number of neighbors in source
// space
n_α ← compute_n_α((S_j)_j, ε)
// Compute number of neighbors in target
// space
n_T ← compute_n_α((T_j)_j, ε)
// Compute estimate for redundancy
I_∩ ← Ψ(k) + Ψ(n_samples) - (Ψ(n_α[j]) + Ψ(n_T[j]))_j
return I_∩;

```

---

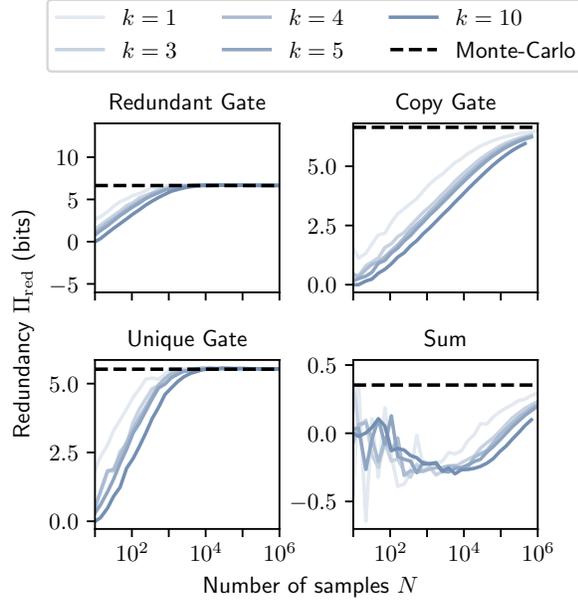


FIG. 14. On the demo gates, the estimation converges to the Monte-Carlo results with increasing number of samples  $N$  for all fixed parameters  $k$

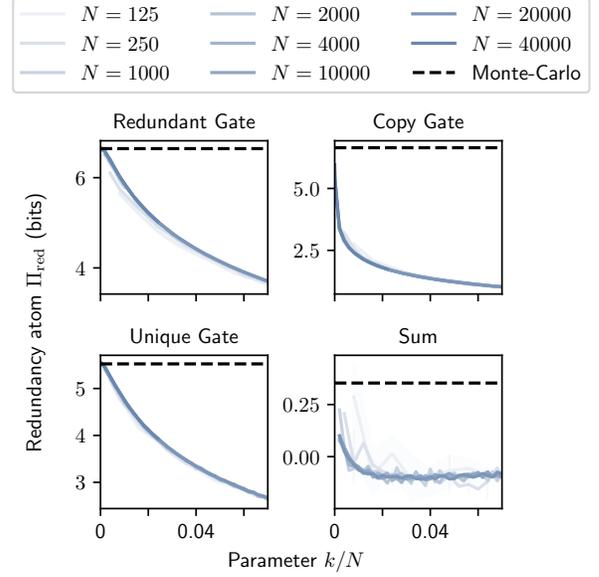


FIG. 15. The result of the estimation does not depend on the value of the parameter  $k$  alone but only on its fraction  $k/N$ .

$$\begin{aligned}
 & \log \left[ \widehat{f_{\bigcup_{q=1}^{k-1} a_{m_q} \cup a_j}^i} \right] = \log \left[ \widehat{f_{c|d(\bigcup_{q=1}^{k-1} a_{m_q} \cup a_j)}^i} \right] + \log \left[ \widehat{f_{\text{disc}(\bigcup_{q=1}^{k-1} a_{m_q} \cup a_j)}^i} \right] \\
 \Rightarrow & \log \left[ \widehat{f_{\bigcup_{q=1}^{k-1} a_{m_q} \cup a_j}^i} \right] = \log \left[ \widehat{f_{c|d(\bigcup_{q=1}^{k-1} a_{m_q} \cup a_j)}^i} \right] + \log \left[ \widehat{f_{\text{disc}(\bigcup_{q=1}^{k-1} a_{m_q} \cup a_j)}^i} \right] \\
 \stackrel{\text{KSG}}{\cong} & \psi(k) - \psi(N) - \log \left[ v_{\epsilon, c|d(\bigcup_{q=1}^{k-1} a_{m_q} \cup a_j)}^i \right] + \log \left[ \widehat{f_{\text{disc}(\bigcup_{q=1}^{k-1} a_{m_q} \cup a_j)}^i} \right],
 \end{aligned}$$

where we have omitted the explicit dependent variables  $x_{a_q} \cup x_{a_j}$  for simplicity. Further,  $c|d(a)$  denotes the continuous random variables in  $a$  conditioned on the discrete variables in  $a$ . Here we have treated the first es-

timated logarithm via KSG estimation, as here the only dependent variables are purely continuous, while the latter term can be evaluated via a standard plugin estimator as it is a purely discrete entropy. Thus the estimators for  $\alpha$  and  $(\alpha, T)$  become

$$\widehat{\log[f_\alpha^i]} = \psi(k) - \psi(N) - \underbrace{\left\langle \log \left[ \sum_{k=1}^{|I_d|+1} (-1)^{k+1} \left( \sum_{\substack{m_1 < m_2 < \dots < m_{k-1} \in I_d \\ m_1 \neq m_2 \neq \dots \neq m_{k-1} \neq j \in J}} v_{\epsilon, c|d(\bigcup_{q=1}^{k-1} a_{m_q} \cup a_j)}^i \widehat{f_{\text{disc}(\bigcup_{q=1}^{k-1} a_{m_q} \cup a_j)}^i} \right) \right] \right\rangle}_=: G_\alpha \quad \text{and}$$

$$\widehat{\log[f_{\alpha,T}^i]} = \psi(k) - \psi(N) - \left\langle \log \left[ \sum_{k=1}^{|I_d|+1} (-1)^{k+1} \left( \sum_{\substack{m_1 < m_2 < \dots < m_{k-1} \in I_d \\ m_1 \neq m_2 \neq \dots \neq m_{k-1} \neq j \in J}} v^i_{\epsilon, c|d(\bigcup_{q=1}^{k-1} a_{m_q} \cup a_j, T)} \hat{f}^i_{disc(\bigcup_{q=1}^{k-1} a_{m_q} \cup a_j, T)} \right) \right] \right\rangle_i .$$

$=: G_{\alpha,T}$

Using the joint space for initial determination of  $\epsilon^i$  as

the distance to  $(x_1, \dots, x_n)^i$ -th  $k$ -th nearest neighbor, the final estimator reads

$$\hat{I}_{\cap}^{\text{sx}} = \begin{cases} \psi(k_{c|d(\alpha,T)}) + \psi(N) - \langle \psi(n_{c|d(\alpha)}(i)) \rangle_i - \langle \psi(n_T(i)) \rangle_i - \Delta \text{vol}_{\alpha,T} & T \text{ cont.} \\ \psi(k_{c|d(\alpha,T)}) + \psi(N) - \langle \psi(n_{c|d(\alpha)}(i)) \rangle_i & T \text{ disc.} \end{cases} ,$$

with  $\Delta \text{vol}_{\alpha,T} = G_{\alpha,T} - G_{\alpha} - \langle \log[v_{\epsilon,T}^i] \rangle_i$  the resulting differential volume term,  $c|d(\alpha)$  being the continuous statement, each conditioned on the discrete variables.

$c|d(\alpha) = X_3 \vee X_5 | X_4$ . Note  $X_1 \vee X_2$  vanished from the continuous statement as they have already been treated inside  $G_{\alpha,T}$  and  $G_{\alpha}$ .

**Example 2.** Considering  $\alpha = X_1 \wedge X_2 \vee X_3 \vee X_4 \wedge X_5$ , with  $X_1, X_2$  and  $X_4$  discrete,  $X_3, X_5$  continuous, we find

- [1] Claude E Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423, 1948.
- [2] Robert M Fano and David Hawkins. Transmission of information: A statistical theory of communications. *American Journal of Physics*, 29(11):793–794, 1961.
- [3] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory, Second Edition*. John Wiley & Sons, Ltd, 2005.
- [4] Paul L. Williams and Randall D. Beer. Nonnegative decomposition of multivariate information. *arXiv Preprint arXiv:1004.2515 [cs.IT]*, 2010.
- [5] Aaron J. Gutknecht, Michael Wibral, and Abdullah Makkeh. Bits and pieces: Understanding information decomposition from part-whole relationships and formal logic. *Proceedings of the Royal Society A*, 477(2251):20210110, 2021.
- [6] Michael Wibral, Conor Finn, Patricia Wollstadt, Joseph T Lizier, and Viola Priesemann. Quantifying information modification in developing neural networks via partial information decomposition. *Entropy*, 19(9):494, 2017.
- [7] Patricia Wollstadt, Daniel L Rathbun, W Martin Usrey, André Moraes Bastos, Michael Lindner, Viola Priesemann, and Michael Wibral. Information-theoretic analyses of neural data to minimize the effect of researchers’ assumptions in predictive coding studies. *arXiv preprint arXiv:2203.10810*, 2022.
- [8] Samantha P Sherrill, Nicholas M Timme, John M Beggs, and Ehren L Newman. Partial information decomposition reveals that synergistic neural integration is greater downstream of recurrent information flow in organotypic cortical cultures. *PLoS computational biology*, 17(7):e1009196, 2021.
- [9] Giuseppe Pica, Eugenio Piasini, Houman Safaai, Caroline Runyan, Christopher Harvey, Mathew Diamond, Christoph Kayser, Tommaso Fellin, and Stefano Panzeri. Quantifying how much sensory information in a neural code is relevant for behavior. *Advances in Neural Information Processing Systems*, 30, 2017.
- [10] Luca Faes, Daniele Marinazzo, and Sebastiano Stramaglia. Multiscale information decomposition: Exact computation for multivariate gaussian processes. *Entropy*, 19(8):408, 2017.
- [11] Jan M Schulz, Jim W Kay, Josef Bischofberger, and Matthew E Larkum. Gaba b receptor-mediated regulation of dendro-somatic synergy in layer 5 pyramidal neurons. *Frontiers in cellular neuroscience*, 15:718413, 2021.
- [12] Andrea I Luppi, Pedro AM Mediano, Fernando E Rosas, Judith Allanson, John D Pickard, Robin L Carhart-Harris, Guy B Williams, Michael M Craig, Paola Finoia, Adrian M Owen, et al. A synergistic workspace for human consciousness revealed by integrated information decomposition. *BioRxiv*, pages 2020–11, 2020.
- [13] Tycho MS Tax, Pedro AM Mediano, and Murray Shanahan. The partial information decomposition of generative neural network models. *Entropy*, 19(9):474, 2017.
- [14] David Alexander Ehrlich, Andreas Christian Schneider, Viola Priesemann, Michael Wibral, and Abdullah Makkeh. A measure of the complexity of neural representations based on partial information decomposition. *Transactions on Machine Learning Research*, 2023.
- [15] Patricia Wollstadt, Sebastian Schmitt, and Michael Wibral. A rigorous information-theoretic definition of redundancy and relevancy in feature selection based on (partial) information decomposition. *J. Mach. Learn. Res.*, 24:131–1, 2023.
- [16] Marcel Graetz, Abdullah Makkeh, Andreas C Schneider, David A Ehrlich, Viola Priesemann, and Michael Wibral. Infomorphic networks: Locally learning neural networks derived from partial information decomposition. *arXiv preprint arXiv:2306.02149*, 2023.
- [17] Seiya Tokui and Issei Sato. Disentanglement analysis with partial information decomposition. *arXiv preprint arXiv:2108.13753*, 2021.
- [18] Paul Pu Liang, Chun Kai Ling, Yun Cheng, Alex Obolenskiy, Yudong Liu, Rohan Pandey, Alex Wilf, Louis-Philippe Morency, and Ruslan Salakhutdinov. Multimodal learning without labeled multimodal data: Guarantees and applications. *arXiv preprint arXiv:2306.04539*, 2023.
- [19] Paul Pu Liang, Zihao Deng, Martin Ma, James Zou, Louis-Philippe Morency, and Ruslan Salakhutdinov. Factorized contrastive learning: Going beyond multi-view redundancy. *arXiv preprint arXiv:2306.05268*, 2023.
- [20] Anti Ingel, Abdullah Makkeh, Oriol Corcoll, and Raul Vicente. Quantifying reinforcement-learning agent’s autonomy, reliance on memory and internalisation of the environment. *Entropy*, 24(3):401, 2022.
- [21] Patricia Wollstadt and Sebastian Schmitt. Interaction-aware sensitivity analysis for aerodynamic optimization results using information theory. In *2021 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 01–08. IEEE, 2021.
- [22] Thomas F Varley and Patrick Kaminski. Untangling synergistic effects of intersecting social identities with partial information decomposition. *Entropy*, 24(10):1387, 2022.
- [23] Michaela Socolof, Jacob Louis Hoover, Richard Futrell, Alessandro Sordani, and Timothy O’Donnell. Measuring morphological fusion using partial information decomposition. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 44–54, 2022.
- [24] Rita Glowienka-Hense, Andreas Hense, Sebastian Brune, and Johanna Baehr. Comparing forecast systems with multiple correlation decomposition based on partial correlation. *Advances in Statistical Climatology, Meteorology and Oceanography*, 6(2):103–113, 2020.
- [25] Nils Bertschinger, Johannes Rauh, Eckehard Olbrich, Jürgen Jost, and Nihat Ay. Quantifying unique information. *Entropy*, 16(4):2161–2183, 2014.
- [26] Abdullah Makkeh, Aaron J. Gutknecht, and Michael Wibral. Introducing a differentiable measure of pointwise shared information. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, 103(3):032149, 2021.
- [27] Conor Finn and Joseph T. Lizier. Pointwise partial information decomposition using the specificity and ambiguity lattices. *Entropy*, 20(4):297, 2018.
- [28] Robin AA Ince. Measuring multivariate redundant information with pointwise common change in surprisal. *Entropy*, 19(7):318, 2017.
- [29] Artemy Kolchinsky. A novel approach to the partial in-

- formation decomposition. *Entropy*, 24(3):403, 2022.
- [30] Adam B Barrett. Exploration of synergistic and redundant information sharing in static and dynamical gaussian systems. *Physical Review E*, 91(5):052802, 2015.
- [31] Michael Kleinman, Alessandro Achille, Stefano Soatto, and Jonathan C Kao. Redundant information neural estimation. *Entropy*, 23(7):922, 2021.
- [32] Ari Pakman, Amin Nejatbakhsh, Dar Gilboa, Abdullah Makkeh, Luca Mazzucato, Michael Wibral, and Elad Schneidman. Estimating the unique information of continuous variables. *Advances in Neural Information Processing Systems*, 34:20295–20307, 2021.
- [33] Jesse Milzman and Vince Lyzinski. Signed and unsigned partial information decompositions of continuous network interactions. *Journal of Complex Networks*, 10(5):cna026, 2022.
- [34] Kyle Schick-Poland, Abdullah Makkeh, Aaron J. Gutknecht, Patricia Wollstadt, Anja Sturm, and Michael Wibral. A partial information decomposition for discrete and continuous variables. *ArXiv preprint arXiv:2106.12393 [cs.IT]*, 2021.
- [35] Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. Estimating mutual information. *Phys. Rev. E*, 69:066138, Jun 2004. doi:10.1103/PhysRevE.69.066138. URL <https://link.aps.org/doi/10.1103/PhysRevE.69.066138>.
- [36] E. T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003. doi:10.1017/CBO9780511790423.
- [37] Edwin T Jaynes. Prior probabilities. *IEEE Transactions on systems science and cybernetics*, 4(3):227–241, 1968.
- [38] Paul Appell and Edouard Goursat. Théorie des fonctions algébriques et de leurs intégrales étude des fonctions analytiques sur une surface de Riemann, 1895.
- [39] Virgil Griffith and Christof Koch. Quantifying synergistic mutual information. In *Guided self-organization: inception*, pages 159–190. Springer, 2014.
- [40] Malte Harder, Christoph Salge, and Daniel Polani. Bivariate measure of redundant information. *Phys. Rev. E*, 87:012130, Jan 2013. doi:10.1103/PhysRevE.87.012130. URL <https://link.aps.org/doi/10.1103/PhysRevE.87.012130>.
- [41] Katerina Hlaváčková-Schindler, Milan Paluš, Martin Vejmelka, and Joydeep Bhattacharya. Causality detection based on information-theoretic approaches in time series analysis. *Physics Reports*, 441(1):1–46, 2007.
- [42] L.F. Kozachenko and N.N. Leonenko. A statistical estimate for the entropy of a random vector. *Problemy Peredachi Informatsii*, 2(23):9–16, 1987.
- [43] Ilya Nemenman, Fariel Shafee, and William Bialek. Entropy and inference, revisited. *Advances in neural information processing systems*, 14, 2001.
- [44] Weihao Gao, Sewoong Oh, and Pramod Viswanath. Demystifying fixed  $k$ -nearest neighbor information estimators. *IEEE Transactions on Information Theory*, 64(8):5629–5661, 2018. doi:10.1109/TIT.2018.2807481.
- [45] Walter Rudin. *Real and Complex Analysis*. McGraw-Hill Science/Engineering/Math, 1986.
- [46] René Unger, Beate Mikoleit, Torsten Schwan, Bernard Bäker, Christian Kehrer, and Tobias Rodemann. "Green Building"—Modelling renewable building energy systems and electric mobility concepts using modelica. In *Proceedings of the 9th International MODELICA Conference*, Munich, Germany, 2012. Modelica Association.
- [47] SimulationX, 2023. URL <https://www.esi-group.com/products/system-simulation>.
- [48] Virgil Griffith, Edwin KP Chong, Ryan G James, Christopher J Ellison, and James P Crutchfield. Intersection information based on common randomness. *Entropy*, 16(4):1985–2000, 2014.
- [49] Virgil Griffith and Tracey Ho. Quantifying redundant information in predicting a target random variable. *Entropy*, 17(7):4644–4653, 2015.
- [50] Aaron Julian Gutknecht. *Information, Logic, and Inference in the Analysis of Complex Networks*. PhD thesis, Georg-August-Universität Göttingen, 2023.
- [51] Walter Rudin. *Real and Complex Analysis*. McGraw-Hill Book Company, 1987.
- [52] Henri Paul Cartan. Differential calculus. (*No Title*), 1971.
- [53] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi:10.1038/s41592-019-0686-2.