
Partial Information Decomposition for Data Interpretability and Feature Selection

Charles Westphal
University College London
charles.westphal.21@ucl.ac.uk

Stephen Hailes
University College London
s.hailes@ucl.ac.uk

Mirco Musolesi
University College London
University of Bologna
m.musolesi@ucl.ac.uk

Abstract

In this paper, we introduce Partial Information Decomposition of Features (PIDF), a new paradigm for simultaneous data interpretability and feature selection. Contrary to traditional methods that assign a single importance value, our approach is based on three metrics per feature: the mutual information shared with the target variable, the feature’s contribution to synergistic information, and the amount of this information that is redundant. In particular, we develop a novel procedure based on these three metrics, which reveals not only how features are correlated with the target but also the additional and overlapping information provided by considering them in combination with other features. We extensively evaluate PIDF using both synthetic and real-world data, demonstrating its potential applications and effectiveness, by considering case studies from genetics and neuroscience.

1 INTRODUCTION

Data interpretability and feature selection are key areas of research in machine learning (ML). One way in which researchers quantify the relative significance of a model’s input attributes is by assigning each of them a feature importance. Multiple paradigms for deriving feature importance have been developed over the past decades. Some of these, inspired by techniques from economics, are based on *Shapley values* (Keinan et al., 2004; Cohen et al., 2007; Apley and Zhu, 2020; Debeer and Strobl, 2020; Catav et al., 2021; Kwon and Zou,

2022; Janssen et al., 2023); others instead rely on the quantification of the correlation between the model’s target and its features (Peng et al., 2005; Brown et al., 2012; Gao et al., 2016; Chen et al., 2018; Schnapp and Sabato, 2021; Covert et al., 2023). Feature importance can be used to explain the relationships within the underlying data and select features (Covert et al., 2020; Weinberger et al., 2023). However, the authors Catav et al. (2021) and Janssen et al. (2023) highlighted inconsistencies in achieving these dual objectives simultaneously. To emphasize this, they considered two perfectly correlated features. If selecting from them, an optimal method would assign one of the two as important, discarding the other. On the other hand, for an exact understanding of how they reduce the uncertainty of the target, they should be assigned equal importance.

In this paper, we develop *partial information decomposition of features (PIDF)*, a novel method that simultaneously explains data and selects features, even in the presence of complex interactions such as those highlighted by Catav et al. (2021) and Janssen et al. (2023). PIDF relies on novel definitions of synergy, redundancy, and mutual information *per feature*¹. PIDF builds upon partial information decomposition (PID) (Chechik et al., 2001; Williams and Beer, 2010), the classic framework for studying variable interactions. However, unfortunately, PID is based on the calculation of a large number of intractable quantities (Makkeh et al., 2019), making it essentially inapplicable to real-world datasets. Instead, we elect to use the interaction information (II) to detect synergy and redundancy. However, II is designed to be applied to generic sets of variables and does not possess the ability to differentiate between synergy and redundancy on a per-feature basis (Williams and Beer, 2010). In response to this, we introduce the concepts of *feature-wise synergy (FWS)* and *feature-wise redundancy (FWR)*, which are

Proceedings of the 28th International Conference on Artificial Intelligence and Statistics (AISTATS) 2025, Mai Khao, Thailand. PMLR: Volume 258. Copyright 2025 by the author(s).

¹Code available at: <https://github.com/c-s-westphal/PIDF>.

not only able to capture these relationships but are also computationally tractable even when describing large datasets. PIDF isolates the FWR, FWS and MI per feature and allows them to be presented in an interpretable manner as illustrated in Figure 1. We show that PIDF effectively explains the data and selects optimal features. We evaluate PIDF’s versatility and performance on various synthetic and real-world datasets, demonstrating its robustness and efficacy.

2 RELATED WORK

Dataset interpretability. Feature importance has been used to enhance interpretability and select features (Fisher et al., 2019; Covert et al., 2020; Aas et al., 2021). However, Catav et al. (2021) highlighted that these dual perspectives do not appear universally consistent. This dilemma prompted Catav et al. (2021) to propose axioms for assigning feature importance when the goal was solely to explain the underlying data. Building on these principles, Janssen et al. (2023) further refined Catav’s axioms for interpretability, while also enhancing speed and ensuring unrelated variables do not receive a non-zero importance rating. That said, these axioms fail to handle complex redundancies, consider the following two cases. First, suppose there exist three features, two of which are indistinguishable, one is unique, and all three reduce the uncertainty of the target equally. According to both the ‘Invariance under Redundant Information and Symmetry under Duplication’ and ‘Elimination’ axioms presented in Janssen et al. (2023) and Catav et al. (2021) all three variables will be assigned the same importance. We contend that to explain the underlying relationships within data, it is essential to identify when information is redundant and consequently more disposable. Secondly, while these axioms of feature importance do account for higher-order interactions, they cannot distinguish when the interactions are synergistic and which features are involved.

Partial information decomposition of features for feature selection. Within the context of feature selection, the most relevant work to ours is that produced by Wollstadt et al. (2023) as they truly develop a framework for the PID of features. Here, they are not using their PID adaptation to describe the relationships within data and select features. Rather, it is used to explain why certain features were chosen over others. Their selection method is a derivative of the conditional likelihood maximization framework unified by Brown et al. (2012); in which a feature is selected if augmenting it to the set of those already chosen maximizes the correlation with the target. At each round of selection, they decompose the contributions of each potential feature according to their adaptation

of the PID paradigm. Unfortunately, this method is unable to identify fully synergistic variables (i.e, variables that considered alone are non-informative, but combine to describe the target). Such features will not be selected using this method, and their synergy will remain undiscovered. Another feature selection method built upon the theoretical foundations of PID is developed in Westphal et al. (2024). The authors study how entropy is transferred from the features to the target, sequentially eliminating those with a negligible contribution. This method effectively identifies synergistic variables; however, it does not consistently identify the smaller of two redundant feature subsets. Consider a scenario where a subset of features conveys the same information as a single variable. The method proposed in Westphal et al. (2024) falls short in consistently choosing the more parsimonious option, which in this case, would be a single variable.

3 RESEARCH QUESTIONS

Following the discussion of the related work in the previous section, we can identify three open research questions for which we will demonstrate that PIDF represents an effective solution (in particular, the first two relates to interpretability and the final one to feature selection):

- *The redundant variables question (RVQ):* Is it possible to distinguish two fully redundant features from one of their non-redundant counterparts if all three reduce the uncertainty of the target equally?
- *The synergistic variables question (SVQ):* Let us consider the case where two features combine to be fully informative regarding the target, whereas, if taken singularly, they are completely uninformative. Is it possible to reveal these variables are only useful in combination?
- *The minimal subsets question (MSQ):* Suppose there exist multiple perfectly redundant feature subsets, is it possible to consistently select the smallest one?

4 PIDF AT A GLANCE

4.1 Overview

Let the input to a ML model be denoted $x = (f_1, f_2, \dots, f_N)$, where f_i is a feature instance. By random sampling from the elements of our input list f_i , and from the realizations of the ground truth y , we represent not only all possible features as a set of random variables (written as $\mathcal{F} = \{F_1, F_2 \dots F_N\}$), but also the ground truth (written as Y). Using these distributions,

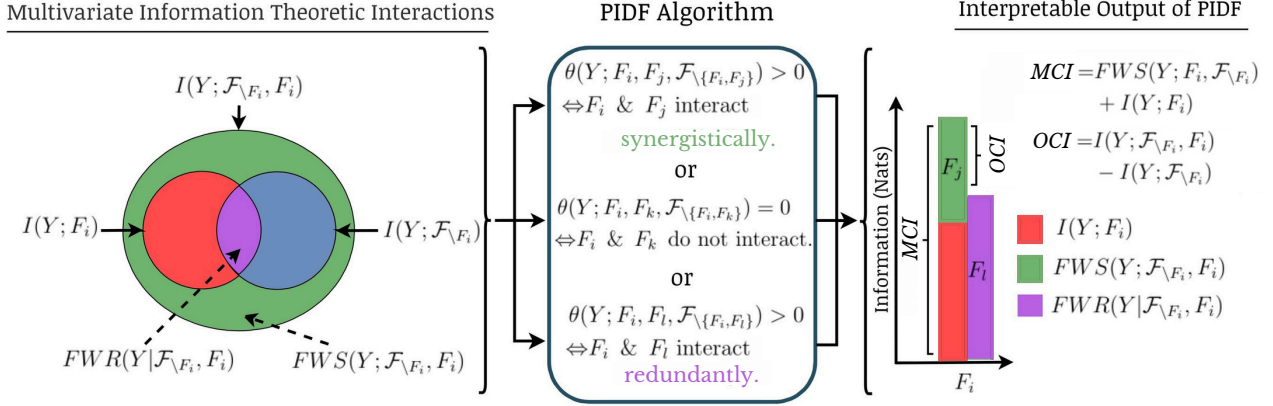


Figure 1: PIDF at a glance. The diagram on the left shows the interactions that characterize how feature F_i interacts with the remaining features (F_j, F_k) to describe the target. The bar graph on the right shows how PIDF can be used to represent these quantities in an interpretable manner.

PIDF aims to elucidate *not only the importance of a feature, but also the inter-feature interactions*. PIDF derives three information-theoretic quantities per feature, namely FWS, FWR, and MI. These quantities are then used to understand how the feature of interest (labeled $F_i \in \mathcal{F}$) interacts with the remaining features (labeled $F_j, F_k, F_l \in \mathcal{F}_{\setminus F_i}$) to describe the target. By considering this fine-grained decomposition and analysis, we are also able to deal with limit cases, such as the perfectly correlated features problem highlighted by Catav et al. (2021) and Janssen et al. (2023).

To enhance the interpretability of PIDF’s results, we present these three quantities as shown in Figure 1. Per feature we will have three bars, the first of which (red), will describe the MI shared with the target. The second (green), will quantify the FWS this feature provides in combination with the remaining features. These red and green bars will stack to give the *maximal conditional information* (MCI). To the right-hand side of this, we will present a purple bar that reveals the extent to which the MCI provided by this feature is redundant. FWS and FWR arise due to interactions; consequently, we label these bars with the other features with which these interactions occur. For instance, in Figure 1, we observe that F_i interacts synergistically with F_j and redundantly with F_l .

4.2 A Practical Example

To provide a better intuition of the proposed approach, let us consider the following example. Suppose we are trying to predict housing prices in Northern California, using the following per-property data: longitude, latitude, and distance to the coast. Individually, given its geography, longitude, and latitude offer limited insight

into housing prices. However, they combine synergistically to provide an exact location — a highly informative feature. Furthermore, longitude is redundant when compared to coastal proximity. The application of PID to this example would require the estimation of 20 different terms, which is not practically possible with the existing techniques (Williams and Beer, 2010; Makkeh et al., 2018, 2019). On the other hand, tractable methods, such as Chechik et al. (2001); Williams and Beer (2010); Griffith and Koch (2014); Bertschinger et al. (2013); Griffith and Ho (2015), are unable to capture the interactions of individual features, instead suggesting that the combination of longitude, latitude, and distance to the coast provides overall redundant information. For this reason, we re-formulate both synergy and redundancy as FWS and FWR, respectively. These quantities can reveal detailed interactions on a per-feature basis, while remaining tractable. Leveraging these quantities, we develop PIDF, an algorithm whose interpretable representation of FWS, FWR, and MI, allows for simultaneous data interpretation and feature selection.

5 FEATURE-WISE PID

In this section, we introduce FWS and FWR more formally. First of all, we use the definition of MI as introduced in Shannon (1948). PID exhaustively describes all possible feature interactions in a set of variables (Williams and Beer, 2010). Consequently, one may expect that FWS and FWR are derived from the PID terms. However, as discussed in the introduction, these values are far too numerous and their calculation too computationally-expensive to be considered appli-

cable to real-world datasets². Consequently, the definitions FWS and FWR extend the concept of interaction information (II) (Williams and Beer, 2010; Chechik et al., 2001). The II for two features is defined as follows: $I(Y; F_1; F_2) = I(Y; F_1, F_2) - I(Y; F_1) - I(Y; F_2)$, if these two features interact synergistically then $I(Y; F_1, F_2) > 0$, if the opposite is true, this indicates redundancy. The simplest approach would be to use II in its current form to calculate both the redundancy and synergy components of our output (purple and green bars in Figure 1). However, this approach has considerable shortcomings as detailed in the following section.

Measuring feature-wise synergy. When assessing the synergistic contribution of an individual feature, the aim is to determine how much extra information a feature can reveal about the target by combining with the remaining features. One method, adopted in Wollstadt et al. (2023), rewrites the MI considering the feature set as a two-body system, where the feature of interest interacts with those remaining, such that: $\mathcal{F} = \{F_i, \mathcal{F}_{\setminus F_i}\}$. We re-apply this reasoning to the II and we obtain: $I(Y; F_i; \mathcal{F}_{\setminus F_i}) = I(Y; F_i, \mathcal{F}_{\setminus F_i}) - I(Y; \mathcal{F}_{\setminus F_i}) - I(Y; F_i)$. However, this definition cannot describe features that interact both redundantly and synergistically. To explain this, we work through the following simple example: suppose that F_i is completely uninformative in itself, but combines with $F_j \in \mathcal{F}_{\setminus F_i}$ via an XOR function to fully describe Y (satisfying $I(Y; F_i, F_j) = H(Y)$). Moreover, let us suppose there also exists $F_k \in \mathcal{F}$ which is identical to F_i ($F_i \equiv F_k$). In this case, the perfect redundancy between F_i and F_k ensures that $I(Y; F_i; \mathcal{F}_{\setminus F_i}) = 0$. However, for a full understanding of the underlying data we would like to know that the variable F_i did provide synergistic information about Y in combination with F_j , despite it being redundant. Motivated by this example, we define the FWS as the maximum extra information gained by adding a feature (F_i) to any subset of $\mathcal{P}(\mathcal{F})$. More formally, this becomes: $FWS(Y; F_i; \mathcal{F}_{\setminus F_i}) = \max_{\mathcal{P}^{ms} \in \mathcal{P}(\mathcal{F}_{\setminus F_i})} I(Y; F_i; \mathcal{P}^{ms})$. This describes the extra reduction in uncertainty of the target caused by considering F_i collectively with \mathcal{P}^{ms} , the subset of maximum synergy. Combining $FWS(Y; F_i; \mathcal{F}_{\setminus F_i})$ with $I(Y; F_i)$ results in the MCI ($MCI(Y; F_i, \mathcal{F}_{\setminus F_i}) = I(Y; F_i) + FWS(Y; F_i; \mathcal{F}_{\setminus F_i})$), i.e., the maximum possible reduction in uncertainty of the target caused by the feature F_i .

Measuring feature-wise redundancy. When assessing feature-wise redundancy, our goal is to quantify to what extent the MCI ($FWS(Y; F_i; \mathcal{F}_{\setminus F_i}) + I(Y; F_i)$) is redundant. Consequently, we require a value with

the following properties: Firstly, it should be positive. Secondly, it should be smaller than the MCI. Consequently, we define the FWR as: $FWR(Y; F_i; \mathcal{F}_{\setminus F_i}) = FWS(Y; F_i; \mathcal{F}_{\setminus F_i}) - I(Y; F_i; \mathcal{F}_{\setminus F_i})$. This expression quantifies the discrepancy between the maximum and overall reduction in uncertainty of the target variable Y caused by feature F_i when combined with any subset of features and with all remaining features, respectively. Intuitively, such a definition aligns with standard definitions of redundancy, that measure the difference between the maximum and true entropy of a variable ensemble (Paluš, 1996). Moreover, FWR satisfies the desirable properties outlined to begin this section and can be related to the MCI via the following theorem:

Theorem 1: *The difference between the MCI (i.e., $FWS(Y; F_i; \mathcal{F}_{\setminus F_i}) + I(Y; F_i)$) and the $FWR(Y; F_i; \mathcal{F}_{\setminus F_i})$ is the overall conditional information (OCI) contributed by F_i . More formally: $FWS(Y; F_i; \mathcal{F}_{\setminus F_i}) + I(Y; F_i) - FWR(Y; F_i; \mathcal{F}_{\setminus F_i}) = I(Y; \mathcal{F}_{\setminus F_i}, F_i) - I(Y; \mathcal{F}_{\setminus F_i})$.*

Proof. See Appendix B.1.

In Theorem 1, we have shown that the difference between the MCI, and our definition of feature-wise redundancy $FWR(Y; F_i; \mathcal{F}_{\setminus F_i})$, is the non-redundant information gained by adding F_i to the set $\mathcal{F}_{\setminus F_i}$ (labeled OCI in Figure 1). It follows that we have developed a consistent theory of feature-wise PID.

6 PIDF

Thus far, we have introduced FWS and FWR, explaining that we aim to isolate these quantities alongside MI for data interpretability. The simplest method to do this would be to maximize $I(Y; F_i; \mathcal{F}_{\setminus \{F_i\}})$ across all possible subsets to find \mathcal{P}^{ms} for each feature F_i . However, this is computationally inefficient. To overcome expensive subset searches, we perform the following steps. We first introduce a measure with the ability to evaluate pairwise synergy or redundancy. This allows us to systematically assess whether feature F_j should be added to F_i 's set of maximum synergy, avoiding costly subset searches. We then discuss potential applications of this measure and highlight some flaws. Before using these examples to motivate an assumption about the data that bounds the values our measure can take. Finally, we leverage these bounds for a practical implementation of PIDF.

6.1 Interaction Information Based Measure

To begin, we introduce the following measure:

$$\theta(Y; F_i; F_j; \mathcal{P}_{\{F_i, F_j\}}) = I(Y; F_i; \mathcal{P}_{\{F_i\}}) - I(Y; F_i; \mathcal{P}_{\{F_i, F_j\}}), \quad (1)$$

²For a more detailed comparison of PIDF to PID please refer to Appendix A.

where $\mathcal{P}_{\{F_i, F_j\}} \in \mathcal{P}(\mathcal{F}_{\{F_i, F_j\}})$ and $\mathcal{P}_{\{F_i\}} \in \mathcal{P}(\mathcal{F}_{\{F_i\}})$. By comparing two IIs, $\theta(Y; F_i; F_j; \mathcal{P}_{\{F_i, F_j\}})$ quantifies how feature F_j affects F_i 's ability to reduce the uncertainty of the target. Positive values of θ indicate that feature F_j interacts synergistically with F_i , and should be added to the F_i 's set of maximum synergy. Meanwhile, negative values imply redundant information is shared between the two features. Using θ we aim to illuminate all synergistic interactions, enabling the calculation of FWS and FWR.

6.2 Handling Redundant Information

We now describe how, if features in the set $\mathcal{F}_{\setminus F_i}$ provide redundant information regarding feature F_i , it impedes the ability to calculate $FWS(Y; F_i; \mathcal{F}_{\setminus F_i})$ and $FWR(Y; F_i; \mathcal{F}_{\setminus F_i})$. To do this, we first point out that $\theta(Y; F_i; F_j; \mathcal{P}_{\{F_i, F_j\}})$, if positive, indicates that feature F_j can be added to F_i 's set of maximum synergy. For a simple and intuitive example of the application of this measure, suppose we calculate $\theta(Y; F_i; F_j; \mathcal{F}_{\setminus \{F_i, F_j\}}) \forall j$, adding those with positive values to F_i 's set of maximum synergy. However, this is potentially problematic. Suppose there exists another feature F_k that is identical to feature F_i . Even if F_i combines synergistically with other features, all values of $\theta(Y; F_i; F_j; \mathcal{F}_{\setminus \{F_i, F_j\}})$ will be equal to zero. To uncover synergistic interactions, we first need to remove redundant information from the set of all features.

A potential solution would be to remove all features ($F_j, F_k \in \mathcal{F}_{\setminus F_i}$) with a non-negligible MI with F_i . However, this approach suffers from the following drawback. The MI between two features does not account for redundant interactions beyond pairwise relationships. For instance, multiple variables can interact synergistically to provide redundant information about F_i . In practice, this is a limit case. As a result, we assume that other features ($F_j \in \mathcal{F}_{\setminus F_i}$) do not combine synergistically ($\mathcal{P} \in \mathcal{P}(\mathcal{F}_{\setminus F_i})$) to provide redundant information about the feature of interest (F_i).

Assumption 1: *Features cannot combine synergistically to provide redundant information about other features. More formally: $I(F_i; \mathcal{P}_{\setminus F_j}) + I(F_i; F_j) \geq I(F_i; \mathcal{P})$.*

Under this assumption, redundancy arises solely from pairwise MI between features, greatly simplifying the search for redundant features. We now show that, given this assumption, $I(F_i; F_j)$ serves as an indicator of whether F_j is redundant or synergistic with respect to F_i .

Theorem 2: *The upper and lower bounds of $\theta(Y; F_i; F_j; \mathcal{P}_{\{F_i, F_j\}})$ are a function of the pairwise MI between F_i and F_j ($I(F_i; F_j)$). More for-*

mally: $2H(F_i) - I(F_i; F_j) \geq \theta(Y; F_i; F_j; \mathcal{P}_{\{F_i, F_j\}}) \geq -I(F_i; F_j)$.

Proof. See Appendix B.2.

The above states that the lower bound of our measure is $-I(F_i; F_j)$. Given that if $\theta(Y; F_i; F_j; \mathcal{P}_{\{F_i, F_j\}}) < 0$ then F_j is redundant with respect to F_i , it is reasonable to use $I(F_i; F_j)$ as an indicator of redundancy. Meanwhile, the upper bound is also a function of $-I(F_i; F_j)$. This indicates the opposite, i.e., low values of $I(F_i; F_j)$ imply that F_j is more likely to interact synergistically with F_i .

As introduced in Section 5, deriving FWS and FWR requires identifying the set of maximum synergy, \mathcal{P}^{ms} , which includes all features, F_j , that have a positive value of $\theta(Y; F_i; F_j; \mathcal{P}_{\setminus \{F_i, F_j\}})$. To achieve this, it is necessary to identify and eliminate all redundant elements. Conveniently, in this section, we have shown that $I(F_i; F_j)$ is an indicator of redundancy. Combining these observations forms the basis of our straightforward implementation of PIDF, where we first order the features in descending order of $I(F_i; F_j)$ (from most to least redundant) and then remove those deemed redundant according to $\theta(Y; F_i; F_j; \mathcal{P}_{\setminus \{F_i, F_j\}})$ ³. Via this assumption and simple implementation, we can calculate FWS and FWR without costly subset searches. We now describe PIDF in detail.

6.3 PIDF in Practice

Description of PIDF. For all possible F_i , repeat the following steps: firstly, calculate $I(Y; F_i)$ to appear as the red bar in Figure 1. Then, rank the remaining features $F_j \in \mathcal{F}_{\setminus F_i}$ in descending order of $I(F_i, F_j)$ as, according to our assumption, this is an indicator of redundancy. Thirdly, iterate through the ranked features checking if they are truly redundant according to $\theta(Y; F_i; F_j; \mathcal{F}_{\setminus \{F_i, F_j\}})$ (the box titled ‘PIDF Algorithm’ in Figure 1 illustrates this process). If so, they can be removed from the set $\mathcal{F}_{\setminus F_i}$. We continue to remove such features until the remainder are no longer redundant, indicating the identification of the set of maximum synergy for the feature F_i . By completing these steps for all $F_i \in \mathcal{F}$ we derive the FWS, FWR, and MI for each feature as desired. For an algorithmic definition of this process, refer to Algorithm C.1 in Appendix C, and, for an example calculation, refer to Appendix D.

Adapting PIDF for varying estimates. II, FWS, and FWR can be directly derived from MI. That said,

³The reason that we still evaluate θ for all F_j , even for those features with positive values of $I(F_i; F_j)$, is because a non-zero MI between two features does not completely rule out the possibility that they interact synergistically in the presence of other variables.

unless restricting oneself to discreet data, exact calculations of MI are not possible. Consequently, we estimate it using methods outlined by Belghazi et al. (2018) (we motivate this choice in Appendix K.1). These estimates can randomly fluctuate, which can lead to the erroneous identification of redundant features during the evaluation of line 7 in Algorithm C.1. To overcome this, when estimating any MI, the process is repeated five times. For feature F_j to then be considered redundant with respect to the chosen feature F_i there has to be a 95% certainty that $\theta(Y; F_i; F_j; \mathcal{F}_{\setminus\{F_i, F_j\}}) < 0$. Otherwise, the feature is classed as non-redundant.

Interpretability of FWR. We assume that features can provide redundant information regarding one another, provided that it is not synergistic. This implies that the total redundant information in the features must be a linear combination of individual contributions. Consequently, we need not present redundancy as just one bar, but rather as a stack of all the individual contributions from all the redundant features. Please refer to Appendix B.3 for a formal proof of this statement. As a concrete example, consider the diagram labeled as MSQ in Figure 2. For feature F_0 , the FWR on the right-hand side of the MCI is not represented by a single bar, but rather as one composed by multiple stacked redundant contributions from individual features. The possibility of discriminating between the different redundant contributions through PIDF allows for better overall interpretability, by highlighting which features are the most redundant.

6.4 PIDF for Feature Selection

We now provide a procedure for feature selection based on the output of PIDF. First of all, we select all features with a non-redundant contribution, i.e., any features whose MCI exceeds their FWR. We then rank the remaining features by their MCI and check in descending order if any features with which they had a non-negligible MI have already been selected. If so, we do not add this feature; otherwise, we proceed to add it. As was the case for PIDF, exceeding thresholds must be done with a minimum of 95% certainty, where the MIs at the basis of these calculations have been estimated using multiple random seeds. The steps of the procedure are detailed in Algorithm 2 (Appendix C). For a fully worked example of these steps, please refer to Appendix D.

7 EXPERIMENTAL EVALUATION

7.1 PIDF for Data Interpretability

We first evaluate PIDF’s ability to explain data, comparing it to three baseline methods: UMFI (Janssen

et al., 2023), MCI (Catav et al., 2021) and PI (Breiman, 2001). For a full discussion of the hyperparameters used for all methods please refer to Appendix E. First, we provide a detailed explanation on how to interpret the results of PIDF. Following this, we present the remaining results, initially on synthetic data and subsequently on real-world data.

7.1.1 Example of How to Interpret the Results of PIDF

In this section, we discuss an example of how to interpret the results of PIDF, we refer to Figure 2. In this illustration, the top row represents the results of PIDF. We guide the reader through the Wollstadt toy example dataset (Wollstadt et al., 2023), as it is known to be characterized by both redundant and synergistic relationships (as shown in the graph titled WT). Firstly, we investigate the bars associated with F_0 . We observe that F_0 shares a large MI with the target (red bar). However, it is also clear that the majority of this information is redundant with respect to F_1 . We know the feature is redundant with respect to F_1 looking at the number (1) reported in F_0 ’s purple bar (i.e., the numbers in the bars represent the subscripts of the remaining features involved). If we now investigate the relationships associated with the feature F_1 , we observe that this feature also provides significant information; however, it is entirely redundant with respect to F_0 . Finally, investigating F_2 , we see this final feature combines with F_0 to provide some non-redundant synergistic information (see the green 0 above the bar). Consequently, using Figure 2, we have identified the following relationships in the data:

1. F_0 provides significant information; although, some of it is redundant with respect to F_1 .
2. F_1 provides significant information; however, *all* of it is redundant with respect to F_0 (this indicates that F_0 contains all the information in F_1 and some).
3. F_2 combines synergistically with F_0 to provide some non-redundant information (Wollstadt et al., 2023).

7.1.2 Synthetic Datasets

We now discuss the results obtained when applying PIDF to synthetic datasets. The first three of these datasets were designed to highlight how current methods fail to properly address the questions introduced in Section 2. The following five are taken from pre-existing works. A description of these datasets can be found in Appendix F.

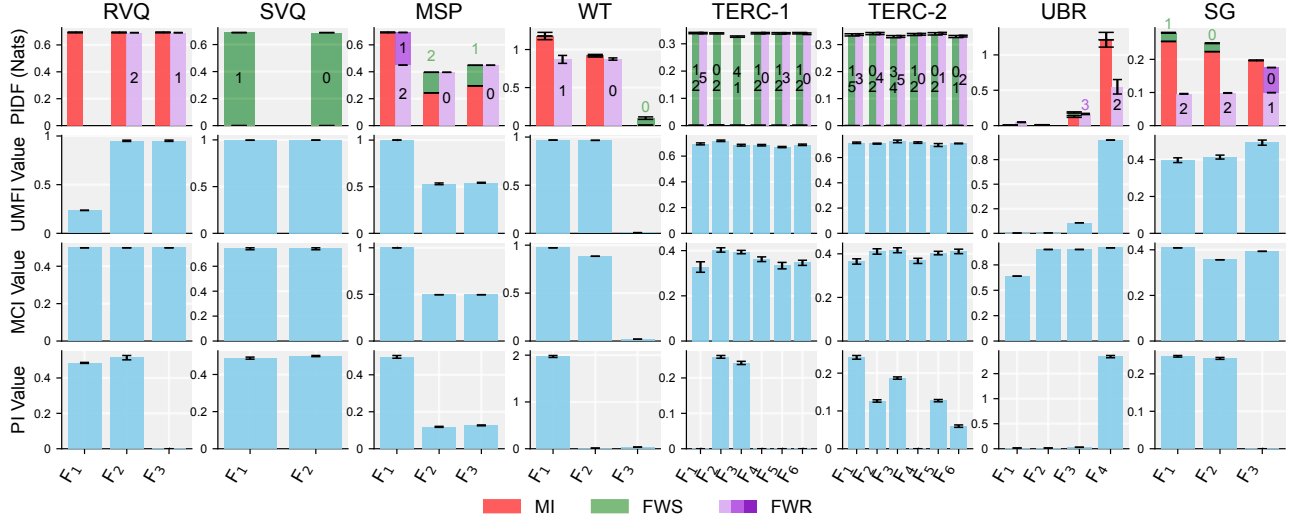


Figure 2: Comparison of feature importance indicators using synthetic datasets.




Figure 3: Comparison of feature importance indicators applied to the California housing dataset.

Redundant variables question (RVQ) dataset.

We now design an extremely simple dataset to demonstrate our method’s ability to answer the RVQ, as introduced in Section 2. The dataset is designed such that $F_1 \equiv F_2$, meaning these two features are fully redundant with respect to one another, as revealed by PIDF in Figure 2. Meanwhile, using MCI one would assume that we have three features of equal importance, without knowledge of the redundancy.

Synergistic variables question (SVQ) dataset.

We now present the evaluation conducted to answer the SVQ introduced in Section 2. In this dataset, the target is formed by the synergistic combination of the two features. This is indicated by the subscripts in the synergy bar for features F_0 and F_1 , as shown in Figure 2. We observe that only PIDF reveals that F_0 combines with F_1 synergistically.

Multiple subsets question (MSQ) dataset. This dataset is characterized by redundancy between the subset $\{F_1, F_2\}$ and the variable F_0 . In Figure 2, we observe that our method again resolves the complex relationships correctly.

TERC-1 & TERC-2. These datasets, introduced by Westphal et al. (2024), are characterized by features that are simultaneously both redundant and synergistic, as illustrated in Figure 2. Moreover, through the labels of the green and purple bars, one can understand with which other features these interactions occur.

UMFI blood relation dataset (UBR). Janssen et al. (2023) used the UBR dataset to demonstrate that their method only assigns non-negligible importance to features that are related to the target via a directed path or causal ancestor (Pearl, 1986). Despite neither F_1 nor F_2 satisfying these conditions, the methods introduced in Catav et al. (2021) still assigned them as important. PIDF avoids these issues and, additionally, unlike UMFI, it reveals that the information provided in F_2 is redundant with respect to F_3 .

Synthetic genes dataset (SG). Our final synthetic dataset was developed in Anastassiou (2007) to demonstrate that two genes could interact synergistically, while others are redundant. We observe that our method, similarly to the PID analysis in Anastassiou (2007), has the capability to reveal such relationships.




Figure 4: Gene importance in the BRCA dataset.

However, we will show that our method is also scalable.

7.1.3 Real-World Datasets

California housing dataset. In this section, we use our method to explain the well-adopted California housing dataset. In Figure 3, we observe that the longitude and latitude combine synergistically. Alone, neither the longitude nor the latitude have great predictive power of housing prices. On the contrary, when considered together, these variables reveal an exact location and become a very useful indicator. Another interesting feature of these results is that the ocean proximity is only redundant with respect to the longitude, as the coast is to the east. Other methods are unable to reveal such relationships. For extra experiments on classic ML datasets, please refer to Appendix H.

BRCA dataset. We now demonstrate PIDF’s ability to describe the relationships between genes in the BRCA dataset (available here: <https://portal.gdc.cancer.gov/projects/TCGA-BRCA>) (Tomczak et al., 2015; Ashton et al., 2018). This dataset details the RNA expression levels of many genes in patients with and without breast cancer. Of these genes, 10 are known chemically to cause cancer. Previously, those using this dataset to study feature importance (Catav et al., 2021; Janssen et al., 2023) have aimed to develop methods that identify these 10 as the most important. In this study, we instead focus on how these 10 genes interact to cause breast cancer. Figure 4 reveals that CDK6 interacts synergistically with multiple genes. This is because CDK6 regulates gene expression levels in the cell, which promotes or inhibits division (Nebenfuhr et al., 2020; Goel et al., 2022). Consequently, in a healthy cell CDK6 expression levels are correlated with those of the genes it regulates. On the contrary, this is not the case in cancerous cells (Malumbres and Barbacid, 2009), where the CDK6 present becomes overactive, leading to expression levels that exceed what is expected. PIDF, unlike the baselines, indicates that we require the expression levels of both CDK6 and the genes it regulates to predict cancer, which is consistent with the overactive CDK6 hypothesis.




Figure 5: Average FWR, FWS, and MI per neuron as synaptic connections are re-established.

Neuron growth experiment. In this experiment, we study the information-theoretic quantities that characterize the spiking interactions that form within a dissociated neural culture over 33 days. Full details of the dataset are available in Appendix G. We now discuss the results presented in Figure 5, where we observe the average FWR, FWS and MI for each neuron at each day over different neuron subsets. An initial increase in redundancy is followed by a decrease. This phenomenon aligns with the principles of Hebbian theory (Hebb, 2005), where the initial surge in redundancy is attributed to the exploratory phase of network formation, during which connections are established randomly. Subsequently, synaptic links that do not exhibit synchronized firing patterns are pruned. This adaptive reconfiguration results in not only a reduction in the observed redundancy, but also a gradual increase in synergy and mutual information.

7.2 PIDF for Feature Selection

In Section 6.4, we developed a simple algorithm based on PIDF for feature selection. We now evaluate this technique. Each synthetic dataset in Section 7.1.2, characterized by certain relationships, has an associated set of optimal features. In this section, we verify our method selects this optimal subset. We adopt the methods described in Wollstadt et al. (2023), West-

Table 1: Number of true positive (TP), false positive (FP), true negative (TN), and false negative (FN) features selected using each method. Red values represent optimal performance, and confidence intervals were omitted since negligible variations were observed.

Dataset	PIDF				Wollstadt				TERC				PI			
	TP	FP	TN	FN	TP	FP	TN	FN	TP	FP	TN	FN	TP	FP	TN	FN
RVQ	2	0	1	0	2	0	1	0	2	0	1	0	2	0	1	0
SVQ	2	0	0	0	0	0	0	2	2	0	0	0	2	0	0	0
MSQ	1	0	2	0	1	0	2	0	0	2	0	1	1	2	0	0
WT	2	0	1	0	2	0	1	0	2	0	1	0	2	0	1	0
TERC-1	3	0	3	0	0	0	3	3	3	0	3	0	2	0	3	1
TERC-2	3	0	3	0	0	0	3	3	3	0	3	0	3	3	0	0
UMFIBR	1	0	3	0	1	0	3	0	1	0	3	0	1	3	0	0
SG	3	0	0	0	3	0	0	0	3	0	0	0	2	0	0	1

phal et al. (2024) (TERC), and Breiman (2001) (PI) as baselines. In Table 1 we observe that PIDF selects the optimal number of features for all the datasets investigated; however, this is not the case for the baselines. The method developed in Wollstadt et al. (2023) fails to recognize fully synergistic variables as important (hence, the poor results achieved on the SVQ, TERC-1, and TERC-2 datasets). Meanwhile, TERC is not able to address the MSQ due to its inability to rank features based on their MCI.

8 CONCLUDING REMARKS

Summary of the Contributions. In this paper, we have presented PIDF, a novel paradigm for simultaneous data interpretability and feature selection based on per-feature synergy, redundancy and mutual information. Through an extensive evaluation using both synthetic and real-world data, we have demonstrated its effectiveness in analyzing and interpreting systems characterized by complex interactions. As part of our future agenda, we plan to explore the potential of applying PIDF to open problems in a variety of scientific fields.

Limitations. The main limitation of PIDF is its temporal complexity, which scales as a function $\mathcal{O}(k^2)$, where k is the number of features. It is worth noting that it is not directly comparable with the other methods since it captures feature-wise synergistic interactions, an intrinsically computationally expensive task. Consequently, in order to scale-up to very-large feature spaces, it is necessary to adopt hierarchical solutions for the comparison of sets of variables in PIDF.

References

Aas, K., Jullum, M., and Løland, A. (2021). Explaining individual predictions when features are dependent: More accurate approximations to Shapley values. *Artificial Intelligence*, 298(4):103502.

Anastassiou, D. (2007). Computational analysis of the

synergy among multiple interacting genes. *Molecular Systems Biology*, 3(1):83.

- Apley, D. W. and Zhu, J. (2020). Visualizing the effects of predictor variables in black box supervised learning models. *Journal of the Royal Statistical Society Series B*, 82(4):1059–1086.
- Ashton, B. et al. (2018). A comprehensive pan-cancer molecular study of gynecologic and breast cancers. *Cancer Cell*, 33(4):690–705.e9.
- Belghazi, I., Rajeswar, S., Baratin, A., Hjelm, R. D., and Courville, A. (2018). MINE: Mutual information neural estimation. In *ICML’18*.
- Bertschinger, N., Rauh, J., Olbrich, E., and Jost, J. (2013). Shared information—new insights and problems in decomposing information in complex systems. In Gilbert, T., Kirkilionis, M., and Nicolis, G., editors, *ECS’12*, pages 251–269. Springer International Publishing.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1):5–32.
- Brown, G., Pocock, A., Zhao, M.-J., and Luján, M. (2012). Conditional likelihood maximisation: A unifying framework for information theoretic feature selection. *Journal of Machine Learning Research*, 13(2):27–66.
- Cang, Z. and Nie, Q. (2020). Inferring spatial and signaling relationships between cells from single cell transcriptomic data. *Nature communications*, 11(1):2084.
- Catav, A., Fu, B., Zoabi, Y., Meilik, A. L. W., Shomron, N., Ernst, J., Sankararaman, S., and Gilad-Bachrach, R. (2021). Marginal contribution feature importance - an axiomatic approach for explaining data. In *ICML’21*.
- Chechik, G., Globerson, A., Anderson, M., Young, E., Nelken, I., and Tishby, N. (2001). Group redundancy measures reveal redundancy reduction in the auditory pathway. In *NeurIPS’01*.
- Chen, J., Song, L., Wainwright, M., and Jordan, M. (2018). Learning to explain: An information-theoretic perspective on model interpretation. In *ICML’18*.
- Cohen, S., Dror, G., and Ruppín, E. (2007). Feature selection via coalitional game theory. *Neural Computation*, 19(8):1939–61.
- Covert, I. C., Lundberg, S., and Lee, S.-I. (2020). Understanding global feature contributions with additive importance measures. In *NeurIPS’20*.
- Covert, I. C., Qiu, W., Lu, M., Kim, N. Y., White, N. J., and Lee, S.-I. (2023). Learning to maximize mutual information for dynamic feature selection. In *ICML’23*.

- Debeer, D. and Strobl, C. (2020). Conditional permutation importance revisited. *BMC Bioinformatics*, 21(1):1–30.
- Fisher, A., Rudin, C., and Dominici, F. (2019). All models are wrong, but many are useful: Learning a variable’s importance by studying an entire class of prediction models simultaneously. *Journal of Machine Learning Research*, 20(177):1–81.
- Gao, S., Ver Steeg, G., and Galstyan, A. (2016). Variational information maximization for feature selection. In *NeurIPS’16*.
- Goel, S., Bergholz, J. S., and Zhao, J. J. (2022). Targeting CDK4 and CDK6 in cancer. *Nature Reviews Cancer*, 22(6):356–372.
- Griffith, V. and Ho, T. (2015). Quantifying redundant information in predicting a target random variable. *Entropy*, 17(7):4644–4653.
- Griffith, V. and Koch, C. (2014). Quantifying synergistic mutual information. In *Guided Self-Organization: Inception*, pages 159–190. Springer.
- Hebb, D. O. (2005). *The organization of behavior: A neuropsychological theory*. Psychology Press.
- Janssen, J., Guan, V., and Robeva, E. (2023). Ultramarginal feature importance: Learning from data with causal guarantees. In *ICAIS’23*.
- Keinan, A., Hilgetag, C. C., Meilijson, I., and Ruppin, E. (2004). Causal localization of neural function: the Shapley value method. *Neurocomputing*, 58-60(3):215–222.
- Kwon, Y. and Zou, J. Y. (2022). WeightedSHAP: analyzing and improving Shapley based feature attributions. In *NeurIPS’22*.
- Lau, C. W., Nair, C., and Ng, D. (2023). A mutual information inequality and some applications. *IEEE Transactions on Information Theory*, 69(10).
- Liang, P. P., Cheng, Y., Fan, X., Ling, C. K., Nie, S., Chen, R. J., Deng, Z., Allen, N., Auerbach, R., Mahmood, F., et al. (2023). Quantifying & modeling multimodal interactions: An information decomposition framework. In *NeurIPS’23*.
- Luppi, A. I., Rosas, F. E., Mediano, P. A., Menon, D. K., and Stamatakis, E. A. (2024). Information decomposition and the informational architecture of the brain. *Trends in Cognitive Sciences*, 28(4):352–368.
- Makkeh, A., Chicharro, D., Theis, D. O., and Vicente, R. (2019). Maxent3d_pid: An estimator for the maximum-entropy trivariate partial information decomposition. *Entropy*, 21(9):862.
- Makkeh, A., Theis, D. O., and Vicente, R. (2018). Broja-2pid: A robust estimator for bivariate partial information decomposition. *Entropy*, 20(4):271.
- Malumbres, M. and Barbacid, M. (2009). Cell cycle, cdks and cancer: a changing paradigm. *Nature Reviews Cancer*, 9(3):153–166.
- Nebenfuehr, S., Kollmann, K., and Sexl, V. (2020). The role of CDK6 in cancer. *International Journal of Cancer*, 147(11):2988–2995.
- Nirenberg, S., Carcieri, S. M., Jacobs, A. L., and Latham, P. E. (2001). Retinal ganglion cells act largely as independent encoders. *Nature*, 411(6838):698–701.
- Pakman, A., Nejatbakhsh, A., Gilboa, D., Makkeh, A., Mazzucato, L., Wibrall, M., and Schneidman, E. (2021). Estimating the unique information of continuous variables. In *NeurIPS’21*.
- Paluš, M. (1996). Detecting nonlinearity in multivariate time series. *Physics Letters A*, 213(3):138–147.
- Pearl, J. (1986). Fusion, propagation, and structuring in belief networks. *Artificial Intelligence*, 29(3):241–288.
- Peng, H., Long, F., and Ding, C. (2005). Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1226–1238.
- Rosas, F. E., Mediano, P. A., Gastpar, M., and Jensen, H. J. (2019). Quantifying high-order interdependencies via multivariate extensions of the mutual information. *Physical Review E*, 100(3):032305.
- Rosas, F. E., Mediano, P. A., Luppi, A. I., Varley, T. F., Lizier, J. T., Stramaglia, S., Jensen, H. J., and Marinazzo, D. (2022). Disentangling high-order mechanisms and high-order behaviours in complex systems. *Nature Physics*, 18(5):476–477.
- Schnapp, S. and Sabato, S. (2021). Active feature selection for the mutual information criterion. In *AAAI’21*.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423.
- Steudel, B. and Ay, N. (2015). Information-theoretic inference of common ancestors. *Entropy*, 17(4):2304–2327.
- Timme, N. M. and Lapish, C. (2018). A tutorial for information theory in neuroscience. *eNeuro*, 5(3).
- Tomczak, K., Czerwińska, P., and Wiznerowicz, M. (2015). Review The Cancer Genome Atlas (TCGA): an immeasurable source of knowledge. *Contemporary Oncology/Współczesna Onkologia*, 2015(1):68–77.
- Varley, T. F. and Hoel, E. (2022). Emergence as the conversion of information: a unifying theory. *Philosophical Transactions of the Royal Society A*, 380(2227):20210150.

- Wagenaar, D. A., Pine, J., and Potter, S. M. (2006). An extremely rich repertoire of bursting patterns during the development of cortical cultures. *BMC Neuroscience*, 7(1):1–18.
- Weinberger, E., Covert, I., and Lee, S.-I. (2023). Feature selection in the contrastive analysis setting. In *NeurIPS’23*.
- Westphal, C., Hailes, S., and Musolesi, M. (2024). Information-theoretic state variable selection for reinforcement learning. *arXiv preprint arXiv:2401.11512*.
- Williams, P. L. and Beer, R. D. (2010). Nonnegative decomposition of multivariate information. *arXiv preprint arXiv:1004.2515*.
- Wollstadt, P., Schmitt, S., and Wibral, M. (2023). A rigorous information-theoretic definition of redundancy and relevancy in feature selection based on (partial) information decomposition. *Journal of Machine Learning Research*, 24(131):1–44.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Not Applicable]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Yes]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Yes]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

A In Depth Comparison with PID

In Section 2, we motivate PIDF by comparing it to existing feature importance and selection methods. However, we were unable to do an in-depth comparison of PIDF to pre-existing methods of calculating synergy and redundancy. Consequently, in this section, we provide a more thorough comparison.

In Williams and Beer (2010), the authors established a framework that identifies all possible interactions among features predicting a target, which has since been applied across multiple fields, including machine learning (Liang et al., 2023), neuroscience (Luppi et al., 2024), complexity studies (Varley and Hoel, 2022), cellular biology (Cang and Nie, 2020), and physics (Rosas et al., 2022). In this paper, we aim to develop a method that allows us to understand with what other features and to what extent does each feature interact synergistically and redundantly. One would expect PID may be the solution. However, in its current form it comes with the following associated drawbacks: first, calculating the terms is intractable, while estimating them is subject to convergence issues and size limitations (Makkeh et al., 2018, 2019; Pakman et al., 2021). Second, the number of terms needed to describe the system equals the $n - 1$ 'th Dedekind number, where n is the number of features. This number is impractically large, consider that a system with nine variables would require approximately 5×10^{22} terms, while for ten variables, the Dedekind number remains unknown. Varley and Hoel (2022) reduced the number of investigable quantities by averaging the contribution of 'layers' in the PID lattice. The different layers can be thought of as representing different levels of redundancy or synergy. Consequently, through this spectral per-layer view one comes to understand how overall synergistic or redundant the system is. However, despite these simplifications, the calculations remained too complex for applications involving more than a few features. O-information (Rosas et al., 2019), II (Chechik et al., 2001), correlational importance (Nirenberg et al., 2001) and synergistic mutual information (Griffith and Ho, 2015) can all be used to estimate the synergy or redundancy of large sets of variables. However, they are unable to consider individual features. The calculation of such quantities does not assist us in answering the question: does a feature interact synergistically, or redundantly, and with what? Instead, they reveal emergent properties of sets. In this work we overcome such issues, by adopting the feature-wise perspective introduced in Wollstadt et al. (2023). Specifically, we re-write the II as a two-body system, containing the feature of interest, as a single variable, and the remaining features as a second variable. Re-formulating the II in this way will allow us to develop measures of synergy and redundancy on a per-feature basis (labeled FWS and FWR), which interact in a consistent way. Unlike its predecessors, this re-formulation will remain computable at scales beyond what is currently possible in PID. To emphasize the differences in practicality of these methods, we provide the following example.

Let us suppose we have four features, F_0 and F_1 are random binary strings, while $F_0 = F_2$, $F_3 = F_1$, and $Y = F_0 + F_1$. To begin, let us outline the interactions in this system that should be present after we have decomposed the information:

- F_0 is redundant with respect to F_2 and vice versa;
- F_1 is redundant with respect to F_3 and vice versa;
- F_0 or F_2 combines synergistically with F_1 or F_3 and vice versa.

This is because the knowledge of one binary string alone only reduces the possible realizations of Y by 1/3. For instance, if $f_0 = 1$ then $y = 1$ or $y = 2$, leading to only a 0.58 bit reduction in uncertainty. On the other hand, having both F_0 and F_1 fully describes Y , leading to a 1.58 bit reduction in uncertainty. Consequently, $I(Y; F_0, F_1) > I(Y; F_0) + I(Y; F_1)$ and our assumption is satisfied.

If we were to now attempt to decompose this information using the system developed in Williams and Beer (2010), we would have to calculate 120 terms. Moreover, no method currently exists with the ability to estimate these quantities. One can estimate the redundancy/synergy of the system by calculating its II. However, such methods fail to describe any properties of individual features. The only viable option is the PIDF algorithm. As a result, we guide the reader through the calculations for the feature F_0 .

As stated, $I(Y; F_0) = 0.58$ bits. We now move onto the calculation of the FWS. The aim here is to isolate the subset of the existing features that maximize the synergistic contribution of F_0 , in accordance with the definition of FWS. For our four features, we obtain the following:

$$FWS(Y; F_0; \mathcal{F} \setminus F_0) = \max_{\mathcal{P}^{ms} \in \{\{F_1\}, \{F_2\}, \{F_3\}, \{F_1, F_2\}, \dots\}} (I(Y; F_0, \mathcal{P}^{ms}) - I(Y; \mathcal{P}^{ms}) - I(Y; F_0)).$$

The result is that we have three subsets for which we always obtain the maximum: $\{F_1\}$, $\{F_3\}$ and $\{F_1, F_3\}$. The members of these subsets are therefore the features with which F_0 combines synergistically, where the final result is $FWS(Y; F_0, \mathcal{F} \setminus F_0) = 0.42$. Due to redundancy among the synergistic features, there are three sets of maximum synergy, with the least cardinal sets being free of redundancy. Notably, our algorithm would only derive one of these subsets as the set of maximum synergy, which does lead to a slight loss in information. However, this comes with the associated speedups. Now, we calculate:

$$\begin{aligned} FWR(Y; F_0; \mathcal{F} \setminus F_0) &= FWS(Y; F_0; \mathcal{F} \setminus F_0) - I(Y; F_0; \mathcal{F} \setminus F_0) \\ &= \max_{\mathcal{P}^{ms} \in \{\{F_1\}, \{F_2\}, \{F_3\}, \{F_1, F_2\}, \dots\}} (I(Y; F_0, \mathcal{P}^{ms}) - I(Y; \mathcal{P}^{ms}) - I(Y; F_0)) \\ &\quad - (I(Y; F_0, F_1, F_2, F_3) - I(Y; F_0) - I(Y; F_1, F_2, F_3)). \end{aligned}$$

From our synergy calculations, we know this leads to:

$$(I(Y; F_0, F_1) - I(Y; F_0) - I(Y; F_1)) - (I(Y; F_0, F_1, F_2, F_3) - I(Y; F_0) - I(Y; F_1, F_2, F_3)) = 0.58 + 0.42,$$

implying that this binary feature is fully redundant. Overall, we have again demonstrated that our method can be used to illuminate synergistic and redundant contributions of individual features.

B Proofs

B.1 Proof of Theorem 1

Proof. We now prove Theorem 1.

$$\begin{aligned} MCI(Y; F_i; \mathcal{F}_{\setminus F_i}) - FWR(Y; F_i; \mathcal{F}_{\setminus F_i}) &= \\ &= I(Y; F_i) + FWS(Y; F_i; \mathcal{F}_{\setminus F_i}) - FWR(Y; F_i; \mathcal{F}_{\setminus F_i}) \\ &\quad \text{(Via our definition of MCI)} \\ &= I(Y; F_i) + I(Y; F_i; \mathcal{F}_{\setminus F_i}) \\ &\quad \text{(Via our definition of FWR and cancelling FWS)} \\ &= I(Y; F_i) + I(Y; F_i, \mathcal{F}_{\setminus F_i}) - I(Y; F_i) - I(Y; \mathcal{F}_{\setminus F_i}) \\ &= I(Y; F_i, \mathcal{F}_{\setminus F_i}) - I(Y; \mathcal{F}_{\setminus F_i}) \end{aligned} \tag{2}$$

B.2 Proof of Theorem 2

Proof. We begin by proving the lower bound. This proof shares the first few steps with those described in Steudel and Ay (2015); Lau et al. (2023); Janssen et al. (2023), before we then use our assumption to complete the proof.

$$\begin{aligned}
 & I(Y; F_i, \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) - I(Y; \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) \\
 &= I(Y; F_i | \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) + I(Y; \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) - I(Y; \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) \\
 & \quad (\text{via the chain rule}) \\
 &= I(F_i; Y | \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) \\
 & \quad (\text{via symmetry}) \\
 &= I(F_i; Y, \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) - I(F_i; \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) \\
 & \quad (\text{via the chain rule}) \\
 &\geq I(F_i; Y, \mathcal{P}_{\setminus\{F_i, F_j\}}) - I(F_i, \mathcal{P}_{\setminus\{F_i, F_j\}}) - I(F_i, F_j) \\
 & \quad (\text{via the monotonicity of MI and our assumption}) \\
 &\geq I(F_i; Y | \mathcal{P}_{\setminus\{F_i, F_j\}}) + I(F_i; \mathcal{P}_{\setminus\{F_i, F_j\}}) - I(F_i, \mathcal{P}_{\setminus\{F_i, F_j\}}) - I(F_i, F_j) \\
 & \quad (\text{via the chain rule}) \\
 &\geq I(Y; F_i | \mathcal{P}_{\setminus\{F_i, F_j\}}) - I(F_i, F_j) \\
 & \quad (\text{via symmetry}) \\
 &\geq I(Y; F_i, \mathcal{P}_{\setminus\{F_i, F_j\}}) - I(Y; \mathcal{P}_{\setminus\{F_i, F_j\}}) - I(F_i, F_j) \\
 & \quad (\text{via chain rule}) \\
 &\theta(Y; F_i; F_j; \mathcal{P}_{\setminus\{F_i, F_j\}}) \geq -I(F_i, F_j)
 \end{aligned} \tag{3}$$

We now provide the upper bound to complete the proof.

$$\begin{aligned}
 & I(Y; F_i, \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) - I(Y; \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) \\
 &= I(Y; F_i | \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) + I(Y; \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) - I(Y; \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) \\
 & \quad (\text{via the chain rule}) \\
 &= I(F_i; Y | \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) \quad (\text{via symmetry}) \\
 &= I(F_i; Y, \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) - I(F_i; \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) \quad (\text{via the chain rule}) \\
 &\leq I(F_i; Y, \mathcal{P}_{\setminus\{F_i, F_j\}}, F_j) - I(F_i; F_j) \quad (\text{via the monotonicity of MI}) \\
 &\leq I(F_i; Y, F_j | \mathcal{P}_{\setminus\{F_i, F_j\}}) + I(F_i; Y, \mathcal{P}_{\setminus\{F_i, F_j\}}) - I(F_i; F_j) \\
 & \quad (\text{via the chain rule}) \\
 &\leq I(F_i; Y, F_j | \mathcal{P}_{\setminus\{F_i, F_j\}}) + I(F_i; Y | \mathcal{P}_{\setminus\{F_i, F_j\}}) + I(F_i; \mathcal{P}_{\setminus\{F_i, F_j\}}) - I(F_i; F_j) \\
 & \quad (\text{via the chain rule}) \\
 &\leq I(F_i; Y, F_j | \mathcal{P}_{\setminus\{F_i, F_j\}}) + I(Y; F_i | \mathcal{P}_{\setminus\{F_i, F_j\}}) + I(F_i; \mathcal{P}_{\setminus\{F_i, F_j\}}) - I(F_i; F_j) \\
 & \quad (\text{via symmetry}) \\
 &\leq I(F_i; Y, F_j | \mathcal{P}_{\setminus\{F_i, F_j\}}) + I(Y; F_i, \mathcal{P}_{\setminus\{F_i, F_j\}}) \\
 & \quad - I(Y; \mathcal{P}_{\setminus\{F_i, F_j\}}) + I(F_i; \mathcal{P}_{\setminus\{F_i, F_j\}}) - I(F_i; F_j) \\
 & \quad (\text{via the chain rule}) \\
 &\leq I(F_i; Y, F_j | \mathcal{P}_{\setminus\{F_i, F_j\}}) + \theta(Y; F_i, \mathcal{P}_{\setminus\{F_i, F_j\}}) + I(F_i; \mathcal{P}_{\setminus\{F_i, F_j\}}) - I(F_i; F_j) \\
 & \quad (\text{via the chain rule}) \\
 &\theta(Y; F_i; F_j; \mathcal{P}_{\setminus\{F_i, F_j\}}) \leq I(F_i; Y, F_j | \mathcal{P}_{\setminus\{F_i, F_j\}}) + I(F_i; \mathcal{P}_{\setminus\{F_i, F_j\}}) - I(F_i; F_j) \\
 &\leq 2H(F_i) - I(F_i; F_j) \quad (\text{via the monotonicity of MI})
 \end{aligned} \tag{4}$$

B.3 Proof that under Assumption 1 FWR is a Linear Combination of Redundant Contributions.

In Section 6.3, we explained that, for enhanced interpretability, FWR can be decomposed into individual redundant contributions from features. In this section, we provide formal evidence for this statement. First, we note that:

$$\begin{aligned}
FWR(Y; F_i; \mathcal{F}_{\setminus F_i}) &= I(Y; F_i; \mathcal{P}^{ms}) - I(Y; F_i; \mathcal{F}_{\setminus F_i}) \\
&= I(Y; F_i; \mathcal{P}^{ms}) - I(Y; F_i) - I(Y; \mathcal{P}^{ms}) \\
&\quad - (I(Y; F_i; \mathcal{F}_{\setminus F_i}) - I(Y; F_i) - I(Y; \mathcal{F}_{\setminus F_i})) \\
&= I(Y; F_i; \mathcal{P}^{ms}) - I(Y; \mathcal{P}^{ms}) - I(Y; F_i; \mathcal{F}_{\setminus F_i}) + I(Y; \mathcal{F}_{\setminus F_i}).
\end{aligned}$$

Next, consider the upper bounds of FWR. Suppose $I(Y; \mathcal{P}^{ms}) = 0$, while all remaining terms in the expression above are equal to $H(Y)$. We obtain:

$$\begin{aligned}
FWR(Y; F_i; \mathcal{F}_{\setminus F_i}) &\geq I(Y; F_i; \mathcal{P}^{ms}) - I(Y; F_i; \mathcal{F}_{\setminus F_i}) + I(Y; \mathcal{F}_{\setminus F_i}) \\
&\geq H(Y) - H(Y) + I(Y; \mathcal{F}_{\setminus F_i}) \\
&\geq I(Y; \mathcal{F}_{\setminus F_i}).
\end{aligned}$$

Because F_i is fully redundant and maximally reduces the uncertainty of Y , we can further write:

$$FWR(Y; F_i; \mathcal{F}_{\setminus F_i}) \geq I(F_i; \mathcal{F}_{\setminus F_i}).$$

Under our assumption, we can then “un-roll” the FWR into a sum of pairwise mutual information, completing the proof.

C Algorithms

C.1 Partial Information Decomposition of Features

The full implementation of PIDF is presented in Algorithm 1.

C.2 PIDF for Feature Selection

The algorithm for feature selection based on PIDF is presented in Algorithm 2.

D Example Calculations

In this Appendix, we perform example calculations for the RVQ dataset. We first re-introduce the form of the data, before then explicitly calculating from lines 2-17 in Algorithm C.1 for feature F_0 . We only consider F_0 for space and clarity. We then calculate the results of Algorithm 2 in full. The RVQ dataset is of the following form:

$$\begin{aligned}
F_0 &\sim \text{Bernoulli}(\mu = 0.5, \sigma = 0.5) \\
F_1 &\sim \text{Bernoulli}(\mu = 0.5, \sigma = 0.5) \\
F_2 &\equiv F_1 \\
Y &= F_0 + 2F_1.
\end{aligned} \tag{5}$$

D.1 Partial Information Decomposition of Features Calculation

Line 2: We begin by calculating $I(F_0; Y)$. Given that $y = f_0 + 2f_1$, where f_1 and f_2 can be either 0 or 1 (according to Equation 5), y can take the following values with equal chance $y \in \{0, 1, 2, 3\}$. It follows that $H(Y) = -\sum_{y \in A(Y)} p_Y(y) \log_e(p_Y(y)) = \sum_{y \in A(Y)} 0.25 \cdot 1.386 = 1.386$. Where $A(\cdot)$ is an operator that produces all possible realizations of a random variable. Upon obtaining the value of f_0 , we reduce y ’s

Algorithm 1 Partial Information Decomposition of Features

Input: Sampled features \mathcal{F} , and target Y .

Output: $(I(F_i; Y), FWS(Y; F_i; \mathcal{F}_{\setminus F_i}), CR(Y; F_i; \mathcal{F}_{\setminus F_i}), \mathcal{R}_{\setminus F_i}) \forall F_i \in \mathcal{F}$.

```

1: for  $F_i$  in  $\mathcal{F}$  do
2:   Calculate  $I(F_i; Y)$ 
3:    $\mathcal{F}_{\setminus F_i}^{sorted} = \text{argsort}_{F_j \in \mathcal{F}_{\setminus F_i}} I(F_i; F_j)$ 
4:   Initialize  $FWR(Y; F_i; \mathcal{F}_{\setminus F_i}) = 0$  and  $\mathcal{R}_{\setminus F_i} = \emptyset$ 
5:   for  $F_j$  in  $\mathcal{F}_{\setminus F_i}^{sorted}$  do
6:     if  $I(F_i, F_j) > 0$  then
7:        $\mathcal{R}_{\setminus F_i} \cup F_j$ 
8:     end if
9:     Calculate  $\theta(Y; F_i; F_j; \mathcal{F}_{\setminus \{F_i, F_j\}}^{sorted})$ 
10:    if  $\theta(Y; F_i; F_j; \mathcal{F}_{\setminus \{F_i, F_j\}}^{sorted}) \leq 0$  then
11:       $FWR(Y; F_i; \mathcal{F}_{\setminus F_j}^{sorted}) = FWR(Y; F_i; \mathcal{F}_{\setminus F_j}^{sorted}) - \theta(Y; F_i; F_j; \mathcal{F}_{\setminus \{F_i, F_j\}}^{sorted})$ 
12:       $\mathcal{F}_{\setminus F_i}^{sorted} = \mathcal{F}_{\setminus \{F_i, F_j\}}^{sorted}$ 
13:    end if
14:  end for
15:   $FWS(Y; F_i; \mathcal{F}_{\setminus F_i}) = I(Y; F_i, \mathcal{F}_{\setminus F_i}) - I(\mathcal{F}_{\setminus F_i}; Y) - I(F_i; Y)$ 
16:   $\mathcal{P}^{ms} = \mathcal{F}_{\setminus F_i}^{sorted}$ 
17:  return  $I(F_i; Y), FWS(Y; F_i; \mathcal{F}_{\setminus F_i}), FWR(Y; F_i; \mathcal{F}_{\setminus F_i}), \mathcal{P}^{ms}, \mathcal{R}_{\setminus F_i}$  ( $F_i$ 's redundant features)
18: end for

```

potential realizations by half. To explain this, let us suppose $f_0 = 0$. In this case, it must be true that $y = 0$ or 2 , whereas if $f_0 = 1$ it then must be true that $y = 1$ or 3 . Therefore, the following holds $H(Y|F_0) = -\sum_{y \in A(Y), f_0 \in A(F_0)} p_{Y, F_0}(y|f_0) \log_e(p_{Y, F_0}(y|f_0)) = \sum_{y \in A(Y), f_0 \in A(F_0)} 0.5 \cdot 0.693 = 0.693$. Given that $I(F_0; Y) = H(Y) - H(Y|F_0)$ it must be true that $I(F_0; Y) = 1.386 - 0.693 = 0.693$. This result agrees with the MI presented in Figure 2.

Lines 3-4: These lines are simple for variable F_0 as F_1 and F_2 have been defined such that they are independent random processes; therefore, $I(F_0; F_1) = I(F_0; F_2) = 0$ and the precise ordering of set $\mathcal{F}_{\setminus F_i}^{sorted}$ is unimportant. In this case, we let $\mathcal{F}_{\setminus F_i}^{sorted} = \{F_1, F_2\}$.

Lines 5-14: To begin we do not add F_1 to the set $\mathcal{R}_{\setminus F_i}$ as $I(F_0; F_1) = 0$. We now calculate $\theta(Y; F_0, F_1, \mathcal{F}_{\{F_0, F_1\}}) = (I(Y; F_0, \mathcal{F}_{\setminus F_0}) - I(Y; \mathcal{F}_{\setminus F_0})) - (I(Y; F_0, \mathcal{F}_{\setminus \{F_0, F_1\}}) - I(Y; \mathcal{F}_{\setminus \{F_0, F_1\}}))$. To begin, we calculate $I(Y; F_0, \mathcal{F}_{\setminus F_0}) = H(Y) = 1.386$, because the knowledge of variables $F_0, \mathcal{F}_{\setminus F_0}$ completely describes the target Y . Meanwhile, $I(Y; \mathcal{F}_{\setminus F_0}) = 0.693$, because without F_0 (via a calculation identical to that completed for $I(F_0; Y)$ above), we can only reduce the uncertainty of Y by half. Therefore, $(I(Y; F_0, \mathcal{F}_{\setminus F_0}) - I(Y; \mathcal{F}_{\setminus F_0})) = 0.693$. We now investigate the values of $I(Y; F_0, \mathcal{F}_{\setminus \{F_0, F_1\}})$ and $I(Y; \mathcal{F}_{\setminus \{F_0, F_1\}})$. Let us first note that $F_1 \equiv F_2$; consequently, by removing F_1 from \mathcal{F} , we do not diminish the set's ability to reduce the uncertainty of Y . It follows that, $I(Y; \mathcal{F}_{\setminus F_0}) = I(Y; \mathcal{F}_{\setminus \{F_0, F_1\}})$, and therefore $I(Y; F_0, \mathcal{F}_{\setminus \{F_0, F_1\}}) - I(Y; \mathcal{F}_{\setminus \{F_0, F_1\}}) = 0.693$, where $\theta(Y; F_0, F_1, \mathcal{F}_{\{F_0, F_1\}}) = 0$. Consequently, F_1 remains in $\mathcal{F}_{\setminus F_0}^{sorted}$ according to line 12.

We now repeat lines 4-14, but iterating from F_1 to F_2 . Again, we do not add F_2 to the set $\mathcal{R}_{\setminus F_i}$ as $I(F_0; F_2) = 0$. We now move onto lines 9-14 by calculating $\theta(Y; F_0, F_2, \mathcal{F}_{\setminus \{F_0, F_2\}}^{sorted})$. We note that $F_2 \equiv F_1$, as a result it must be true that $\theta(Y; F_0, F_2, \mathcal{F}_{\{F_0, F_2\}}) = \theta(Y; F_0, F_1, \mathcal{F}_{\{F_0, F_1\}}) = 0$. Therefore, for feature F_0 , we return $I(F_0; Y) = 0.693, FWS(Y; F_0; \mathcal{F}_{\setminus F_0}) = 0, FWR(Y; F_0; \mathcal{F}_{\setminus F_0}) = 0, \mathcal{R}_{\setminus F_i} = \{\}$.

Above we have calculated from lines 2-17 in Algorithm C.1 for variable F_0 in the RVQ dataset. By applying the same procedure, one can obtain the following values for F_1 : $I(F_1; Y) = 0.693, FWS(Y; F_1; \mathcal{F}_{\setminus F_1}) = 0, FWR(Y; F_1; \mathcal{F}_{\setminus F_1}) = 0.693$, with $\mathcal{R}_{\setminus F_i} = \{F_2\}$. Similarly, for F_2 : $I(F_2; Y) = 0.693, FWS(Y; F_2; \mathcal{F}_{\setminus F_2}) = 0, FWR(Y; F_2; \mathcal{F}_{\setminus F_2}) = 0.693$, and again $\mathcal{R}_{\setminus F_i} = \{F_1\}$.

Algorithm 2 PIDF for Feature Selection**Input:** $(FWS(Y; F_i; \mathcal{F}_{\setminus F_i}), FWR(Y; F_i; \mathcal{F}_{\setminus F_i}), I(Y; F_i), \mathcal{R}_{\setminus F_i}) \forall F_i \in \mathcal{F}$ **Output:** \mathcal{F}_* (a desirable feature set).

```

1: Initialize  $\mathcal{F}_* = \{\}$ .
2: for  $F_i$  in  $\mathcal{F}$  do
3:   if  $FWS(Y; F_i; \mathcal{F}_{\setminus F_i}) + I(Y; F_i) > FWR(Y; F_i; \mathcal{F}_{\setminus F_i})$  then
4:      $\mathcal{F}_* \cup F_i$ 
5:   end if
6: end for
7:  $\mathcal{F}^{sorted} = \text{argsort}_{F_j \in \mathcal{F} \setminus \mathcal{F}_*} FWS(Y; F_j; \mathcal{F}_{\setminus F_j}) + I(Y; F_j)$ 
8: for  $F_i$  in  $\mathcal{F}^{sorted}$  do
9:   if  $\mathcal{F}_* - \mathcal{R}_{\setminus F_i} = \mathcal{F}_*$  then
10:     $\mathcal{F}_* \cup F_i$ 
11:   end if
12: end for
13: return  $\mathcal{F}_*$ 

```

D.2 Feature Selection Calculation

We will now utilize the results obtained from the PIDF calculations in the previous section to select features. This will be done by explicitly carrying out the calculations as per Algorithm 2.

Lines 1-6: From the results generated in the last section, it is clear that $I(F_0; Y) + FWS(Y; F_0; \mathcal{F}_{\setminus F_0}) - FWR(Y; F_0; \mathcal{F}_{\setminus F_0}) = 0.693$, and therefore F_0 is included in the set of selected features ($\mathcal{F}_* = \{F_0\}$). However, this is not the case for F_1 and F_2 . For F_1 we calculate $I(F_1; Y) + FWS(Y; F_1; \mathcal{F}_{\setminus F_1}) - FWR(Y; F_1; \mathcal{F}_{\setminus F_1}) = 0.693 + 0 - 0.693 = 0$. Similarly, for F_2 we get $I(F_2; Y) + FWS(Y; F_2; \mathcal{F}_{\setminus F_2}) - FWR(Y; F_2; \mathcal{F}_{\setminus F_2}) = 0.693 + 0 - 0.693 = 0$. Consequently, neither F_1 or F_2 are added to our set of selected features at this stage.

Lines 7-12: We begin by ordering the non-selected features by their MCI $I(F_i; Y) + FWS(Y; F_i; \mathcal{F})$ (line 7). However, in this case, we have $I(F_1; Y) + FWS(Y; F_1; \mathcal{F}) = 0.693$ and $I(F_2; Y) + FWS(Y; F_2; \mathcal{F}) = 0.693$; therefore, the order is not meaningful. To complete the calculation, we let $\mathcal{F}^{sorted} = \{F_1, F_2\}$, and apply line 9. Currently, we have $\mathcal{F}_* = \{F_0\}$ and $\mathcal{R}_{\setminus F_1} = \{F_2\}$, therefore $\mathcal{F}_* - \mathcal{R}_{\setminus F_1} = \mathcal{F}_*$ and F_1 is added to the set \mathcal{F}_* , such that $\mathcal{F}_* = \{F_1, F_2\}$. Now repeating this step for F_2 , we have $\mathcal{F}_* = \{F_0, F_1\}$ and $\mathcal{R}_{\setminus F_2} = \{F_1\}$, it is trivial to see that $\mathcal{F}_* - \mathcal{R}_{\setminus F_2} = \{F_0\} \neq \mathcal{F}_*$. As a result F_2 is not added to our final set of selected features, and our final result is $\mathcal{F}_* = \{F_0, F_1\}$.

E Hyperparameter Selection

For both the feature selection and feature interpretation experiments we used a simple grid search when deciding hyperparameters. PIDF, TERC and the method introduced by Wollstadt et al. (2023) all rely on the calculation of MI. As discussed in Section 6.3, we estimate MI across all methods using the 5 repetitions of the techniques described by Belghazi et al. (2018). For clarity regarding hyperparameters, we present the details of this technique in Algorithm 3.

E.1 Explaining Data Hyperparameters

For PIDF we selected $b = 1000$, $N = 20,000$ as defined in Algorithm 4. Meanwhile, the learning rate of the network in use was $\alpha = 0.0001$ with an adam optimizer. Each network had one hidden layer consisting of 50 neurons. For UMFI and MCI, 100 trees were selected as optimal. Otherwise, hyperparameters were as described in Janssen et al. (2023) and Catav et al. (2021) respectively. For PI, we also selected 100 random trees.

Algorithm 4 MINE Belghazi et al. (2018).

Input: Training dataset $(f^1, y^1, f^2, y^2 \dots f^T, y^T)$
Output: $I(Y; F)$

- 1: Initialize weights θ .
 - 2: **for** 1 to N **do**
 - 3: Draw mini batch samples of length b from the joint distribution of the actions and the state with all possible variables included $p_{Y,F} \sim ((y^1, f^1), \dots, (y^b, f^b))$, and repeat for the marginal distribution $p_{Y \otimes F} \sim ((y^1, f^r), \dots, (y^b, f^r))$ (where the superscript r indicates random sampling).
 - 4: $I(Y; F) = \frac{1}{b} \sum^b F_\theta((y^1, f^1), \dots, (y^b, f^b)) - \frac{1}{b} \sum^b \log e^{F_\theta((y^1, f^r), \dots, (y^b, f^r))}$
 - 5: $\theta \leftarrow \tilde{\nabla}_\theta I(Y; F)$
 - 6: **end for**
 - 7: **return** $I(Y; F)$
-

E.2 Feature Selection Hyperparameters

For PIDF, TERC and the Wollstadt method we again selected $b = 1000$, $N = 20,000$ as defined in Algorithm 4. Meanwhile, the learning rate of the network in use was $\alpha = 0.0001$. For PI, we again selected 100 random trees.

F Statistical Characterization of Synthetic Datasets

Redundant variables question dataset (RVQ). This dataset is comprised of three stochastic binary arrays, F_0, F_1, F_2 of length 1000, two of which are fully redundant with respect to one another. More formally, we have:

$$\begin{aligned}
 F_0 &\sim \text{Bernoulli}(\mu = 0.5, \sigma = 0.5) \\
 F_1 &\sim \text{Bernoulli}(\mu = 0.5, \sigma = 0.5) \\
 F_2 &\equiv F_1 \\
 Y &= F_0 + 2 * F_1.
 \end{aligned} \tag{6}$$

Synergistic variables question dataset (SVQ). The features are two stochastic binary arrays, F_0, F_1 of length 1000, where $Y = \text{XOR}(F_0, F_1)$ and $F_0 \neq F_1$. F_0 and F_1 have been defined such that they combine to provide synergistic information about the target. More formally:

$$\begin{aligned}
 F_0 &\sim \text{Bernoulli}(\mu = 0.5, \sigma = 0.5) \\
 F_1 &\sim \text{Bernoulli}(\mu = 0.5, \sigma = 0.5) \\
 Y &= \text{XOR}(F_0, F_1).
 \end{aligned} \tag{7}$$

Multiple subsets question dataset (MSQ). The features are three stochastic binary arrays, F_0, F_1, F_2 of length 1000; where, $Y = F_0 = F_1 + F_2$. This can be re-written:

$$\begin{aligned}
 F_1 &\sim \text{Bernoulli}(\mu = 0.5, \sigma = 0.5) \\
 F_2 &\sim \text{Bernoulli}(\mu = 0.5, \sigma = 0.5) \\
 Y &= F_0 = F_1 + F_2.
 \end{aligned} \tag{8}$$

Wollstadt toy dataset (WT). The Wollstadt toy example is characterized by three features, formed by combining normal distributions:

$$\begin{aligned}
 \epsilon_1 &\sim \mathcal{N}(\mu = 0, \sigma = 1) \\
 \epsilon_2 &\sim \mathcal{N}(\mu = 0, \sigma = 1) \\
 \epsilon_3 &\sim \mathcal{N}(\mu = 0, \sigma = 1) \\
 F_0 &= \epsilon_1 + 0.1F_2 \\
 F_1 &= 0.8\epsilon_1 + 0.2\epsilon_2 + 0.01 * F_2 \\
 F_2 &\sim \mathcal{N}(\mu = 0, \sigma = 1) \\
 Y &= \sin(\epsilon_1) + 0.1\epsilon_3.
 \end{aligned} \tag{9}$$

TERC-1. This dataset, introduced by Westphal et al. (2024), contains six features, with the following characteristics:

$$\begin{aligned}
F_0 &\sim \text{Bernoulli}(\mu = 0.5, \sigma = 0.5) \\
F_1 &\sim \text{Bernoulli}(\mu = 0.5, \sigma = 0.5) \\
F_2 &\sim \text{Bernoulli}(\mu = 0.5, \sigma = 0.5) \\
F_3 &\equiv F_0 \\
F_4 &\equiv F_0 \\
F_5 &\equiv F_0 \\
Y &= \begin{cases} 0, & \text{if } f_1 = f_2 = f_2 \\ 1, & \text{otherwise} \end{cases}
\end{aligned} \tag{10}$$

TERC-2. This dataset, introduced by Westphal et al. (2024), contains six features, with the following characteristics:

$$\begin{aligned}
F_0 &\sim \text{Bernoulli}(\mu = 0.5, \sigma = 0.5) \\
F_1 &\sim \text{Bernoulli}(\mu = 0.5, \sigma = 0.5) \\
F_2 &\sim \text{Bernoulli}(\mu = 0.5, \sigma = 0.5) \\
F_3 &\equiv F_0 \\
F_4 &\equiv F_1 \\
F_5 &\equiv F_2 \\
Y &= \begin{cases} 0, & \text{if } f_1 = f_2 = f_2 \\ 1, & \text{otherwise} \end{cases}
\end{aligned} \tag{11}$$

UMFI blood relation dataset (UBR). Janssen et al. (2023) used the UBR dataset to demonstrate that their method had favorable properties in detecting blood relationships in causal graphs. This dataset is characterized by four features with the following characteristics.

$$\begin{aligned}
\epsilon_1 &\sim \mathcal{U}(\mu = -1, \sigma = 1) \\
\epsilon_2 &\sim \mathcal{U}(\mu = -0.5, \sigma = 0.5) \\
\epsilon_3 &\sim \text{Exp}(1) \\
\epsilon_4 &\sim \mathcal{N}(\mu = 0, \sigma = 1) \\
F_0 &\sim \mathcal{N}(\mu = 0, \sigma = 1) \\
F_1 &= 3F_1 + \epsilon_1 \\
F_2 &= \epsilon_4 + F_0 \\
F_3 &= Y + \epsilon_3 \\
Y &= \epsilon_4 + \epsilon_2.
\end{aligned} \tag{12}$$

Synthetic genes dataset (SG). Our final synthetic dataset was developed in Anastassiou (2007) to demonstrate the types of relationships that could be present in genetic interactions. Formally Y is:

$$Y \sim \text{Bernoulli}(\mu = 0.5, \sigma = 0.5), \tag{13}$$

where $y = 1$ indicates cancer. The features meanwhile can be explained via the following bullet points:

- In healthy samples, the probability that both $f_1 = f_2 = 1$ is 95%, whereas the remaining three joint states (00, 01 and 10) are equally likely. Meanwhile, in cancerous samples, the probability that both $f_1 = f_2 = 1$ is 5%, whereas the remaining three joint states (00, 01 and 10) remain equally likely.
- In healthy samples, if $y = 0$, then $F_3 \sim \text{Bernoulli}(\mu = 0.2, \sigma = 0.4)$. Meanwhile, in cancerous samples, we have $F_3 \sim \text{Bernoulli}(\mu = 0.8, \sigma = 0.4)$.

G Sampling Method for Neural Culture Data

We now provide more details regarding the firing neurons dataset, including the methods we used to sample from it. This dataset was introduced in Wagenaar et al. (2006) and is publicly accessible online⁴. These datasets comprise multiunit spiking activities recorded from each of the 60 electrodes in the multielectrode array used to study the activations of dissociated neural cells. Specifically, the data from neural culture 2-2 were utilized. Comprehensive details regarding the cultivation and maintenance of these cultures are documented in Wagenaar et al. (2006). The analysis focused on recordings from eight developmental stages of the culture, measured in days in vitro (DIV): 4, 7, 12, 16, 20, 25, 31, and 33. The recording from DIV 16 lasted for 60 minutes, while the recordings from the other days each lasted for 45 minutes. For the purposes of this analysis, the data were segmented into bins of 16 μs . The probability distributions required for calculating the three information-theoretic quantities of interest were generated by analyzing the spike trains from groups of ten distinct electrodes. Within each group, one electrode was randomly designated as the Y variable, and the remainder as the features. For every time step in the spike trains, the states (spiking or not) of the feature electrodes at time t and the Y electrode at time $t + 1$ were noted. This tabular data could then act as input to Algorithms 4 and C.1. This procedure was replicated for each set of non-identical electrodes, ensuring that groups with interchanged features were only considered once to prevent duplication. A full discussion of this dataset can be found in Timme and Lapish (2018).

H Further Experiments on Classic Machine Learning Datasets

H.1 Iris Dataset

In this section, we apply PIDF to the Iris dataset. One of the earliest and best-adopted datasets in ML, Iris was used to predict flower type given the petal length and width and sepal length and width.

In Figure 6, we observe that, unsurprisingly, there is a large amount of redundancy between the features in the Iris dataset. UMFI’s pre-processing step, which removes redundancies between pairwise features, is unable to resolve the more complex redundancies we see Figure 6; hence, the assignment of feature importance scores of 0. Meanwhile, PIDF is the only method able to reveal the complex set of redundancies, showing that sepal width and sepal length are both well correlated with the target. But the sepal width also provides redundant information with respect to the petal length and width, making this the most informative feature, as revealed using PI.

H.2 Abalone Dataset

In this section, we apply PIDF to the Abalone dataset, which comprised information about the size and weight of an Abalone, and is used to predict its age.

In Figure 7, it is possible to observe that in this everyday dataset we see synergistic interactions occurring. The whole weight and shucked weight combine synergistically with the shell weight. Intuitively, this makes sense as finer grained detail about the weights of individual components of the animal may lead to greater knowledge of the animals age. Meanwhile, we observe that the viscera weight and shell weight are redundant with respect to one another, while this is not the case with whole and shucked weights.

H.3 Whitewine Dataset

In this section, we apply PIDF to a classic ML dataset, which uses the chemical properties of wine to predict its quality.

In Figure 8, we observe that the alcohol levels and density provide redundant information with respect to each other. This redundancy can be attributed to the fact that, apart from alcohol, wine primarily consists of water-based solubles. Consequently, the density of the wine is a function of the alcohol-water ratios. Interestingly, we see that the free and total sulfur dioxide features are not highly redundant with respect to one another. This is because, during the wine-making process, free sulfur dioxide is added for anti-microbial effects. Meanwhile, the bound sulfur dioxide is more dependent on the grape used.

⁴<https://potterlab.gatech.edu/potter-lab-data-code-and-designs/>

H.4 Cleveland Heart Disease Dataset

In this section, we apply PIDF to a classic ML dataset, which uses medical data to predict cases of heart disease in Cleveland, Ohio.

In Figure 9, we observe that thalassemia intermedia and the number of major articles with blockages as detected by fluoroscopy combine synergistically to predict heart disease.

I Feature Duplication Experiments

In the main body of the paper, we demonstrated that our method could be used to successfully detect redundancies in synthetic data. Extending this validation, inspired by experiments in Catav et al. (2021), we now demonstrate our method’s ability to illuminate known redundancies in real-world datasets. This is achieved by duplicating features within the, widely recognized, California housing, Abalone, and Whitewine ML datasets. Subsequently, we provide evidence demonstrating the successful detection of these introduced redundancies by our methodology. These duplicated features will simply be added to the original datasets.

In Figure 10, we showcase the outcomes from our experiments on duplicated features. The findings demonstrate that FWS and MI of the duplicated features mirror that of their non-duplicated counterparts. However, the information provided by these features is now redundant.

J Computational Requirements

In its current form, the execution of PIDF requires more computational resources than its other counterparts for feature interpretability (Catav et al., 2021; Janssen et al., 2023). For example, executing the results for the TERC-1 and TERC-2 datasets, which contained six features each with 1000 instances took 45 minutes on a P100 NVIDIA GPU cluster. However, this can easily be reduced via the following methods:

1. As discussed in the conclusion, a hierarchical approach can be adopted. Rather than removing individual redundant features, as per line 12 of Algorithm C.1, many can be checked and removed at once.
2. Another way in which it is possible to scale the PIDF implementation is by enhancing the efficiency of the MI estimation. Currently, we use MINE (Belghazi et al., 2018), as the results are highly accurate and in Nats, which are easily interpretable. However, if one wishes just to compare the relative values of PIDF, without the need for values with units, more efficient methods can be used. In particular, we will make available a version of PIDF implemented using the method developed in Covert et al. (2020). By using this simple method of MI estimation, the time taken to execute PIDF is reduced by a factor of ten. However, the resulting values only represent relative importance and do not exactly correspond to the ground truth. We will discuss this method in detail in the following section.

K Scalability Experiments

In our experiments, we estimated MI using MINE. In this section, we discuss how MINE, despite its high accuracy, is computationally expensive. Therefore, it may be preferable to use a faster method, even if it sacrifices some accuracy. We now present an implementation of PIDF using an alternative method of MI estimation. We will demonstrate that, using this approach, PIDF is highly scalable, making it applicable to real-world datasets, such as MNIST. We will discuss the inherent trade-offs between computation efficiency and accuracy of this alternative method.

K.1 Comparison of MI Estimation Techniques

First, we empirically represent the trade-off between efficiency and accuracy for three different MI estimation techniques (Covert et al., 2020). In Figure 11, we observe that MINE takes the longest to converge, but is the most accurate. Meanwhile, the method presented by Covert et al. (2020) with an MSE loss converges rapidly, but it is not accurate.

K.2 Alternative Method for the Calculation of MI and Trade-offs

Based on the results in the previous section, replacing MINE with the method presented by Covert et al. (2020) with an MSE loss would likely result in a more efficient but less accurate approach. We briefly discuss the validity of such a method before presenting results.

PIDF ideally would derive the values of FWR and FWS in bits; however, significant insights into the interactions can still be gained from the relative values of FWS, FWR, and MI. To clarify, Figure 11 illustrates the PIDF results obtained from two simple datasets using both MINE and our more efficient estimator. We observe that while the relative proportions of FWR, FWS, and MI remain consistent, the exact values in the two rightmost graphs are no longer accurate. Consequently, even when using the efficient estimator, it is still possible to discern the interactions and their respective elements.

Consequently, we apply this more scalable method to the MNIST dataset and present the results in Figure 13. The outer pixels in MNIST are often just black; consequently, only the central pixels contribute to strongly synergistic or redundant interactions, as they are the only pixels providing information. Furthermore, we observe a tendency for these pixels to interact synergistically. This occurs because, when considered together, their spatial relationships and patterns (such as edges, corners, and textures) emerge. These features are crucial for recognizing handwritten digits in MNIST. We might expect that pixels near one another would have high levels of redundancy. In fact, pixels close to one another did have high pairwise MIs. However, according to line 3 of Algorithm C.1, this caused them to be checked and removed first when identifying the set of maximum synergy. Consequently, other pixels were then identified as providing redundant information.

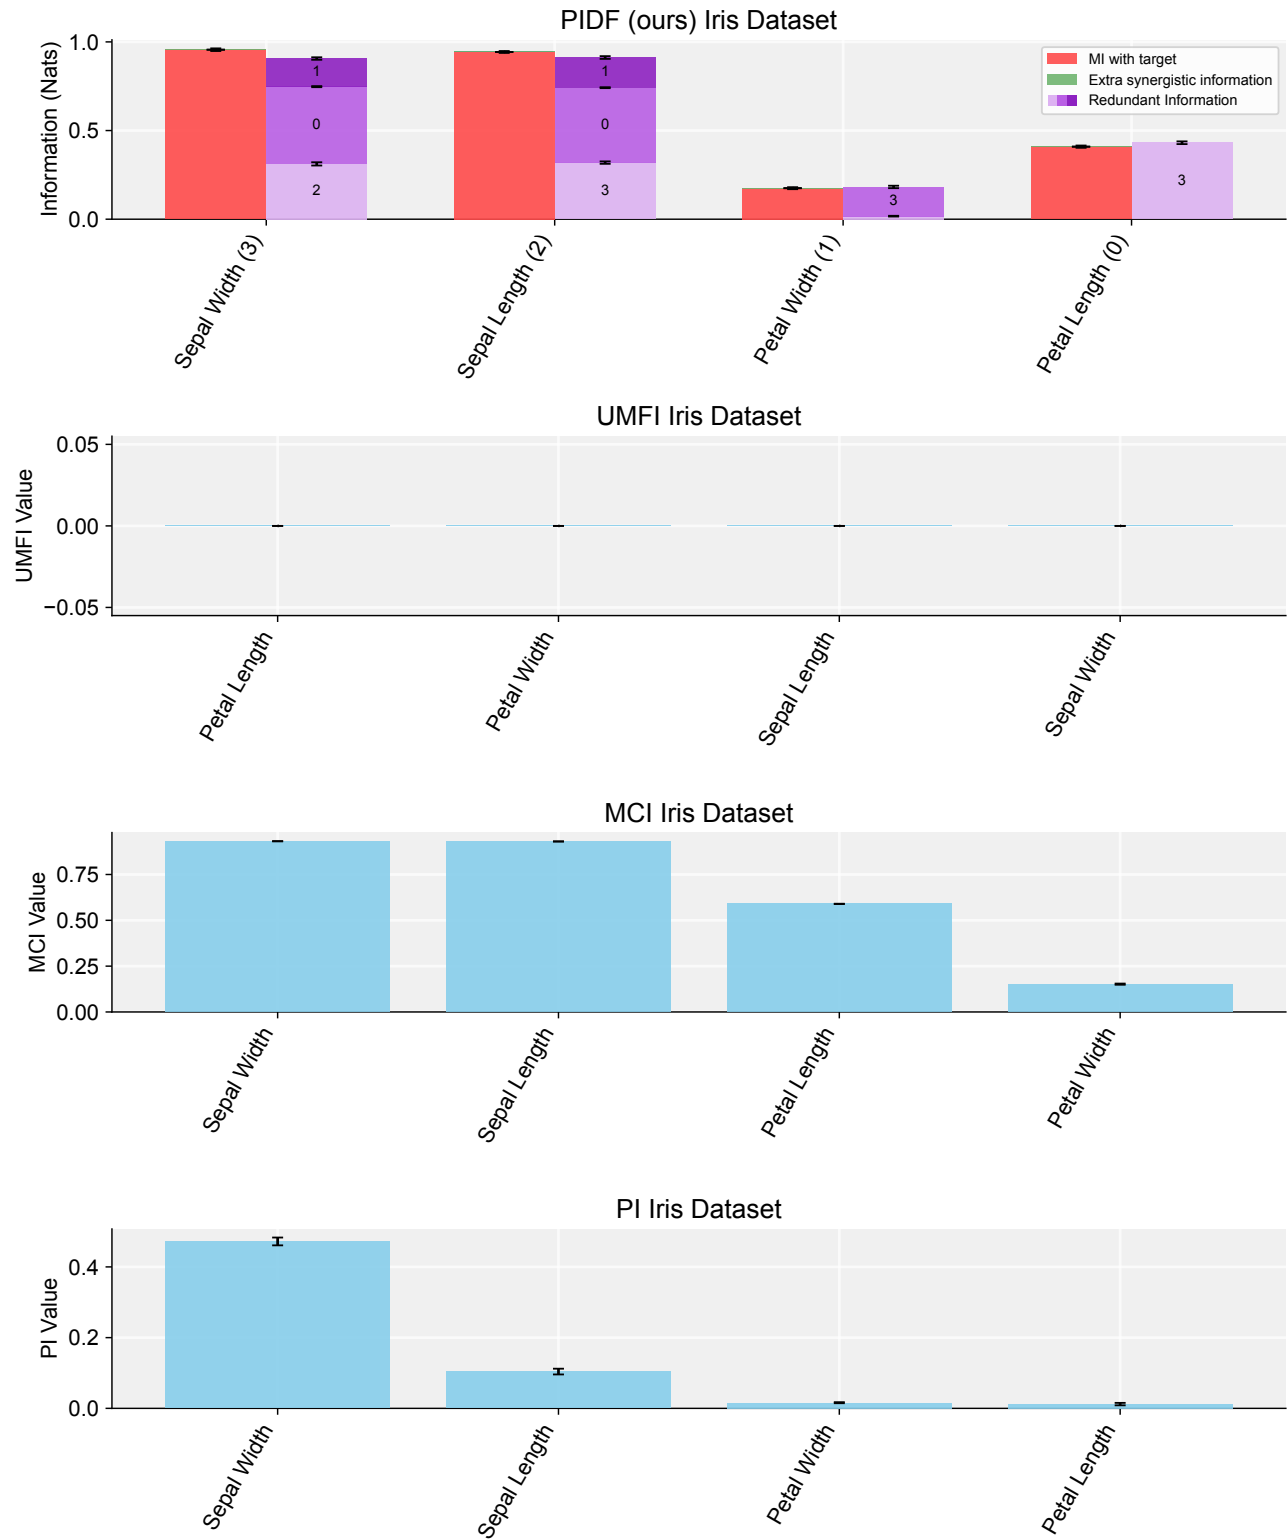


Figure 6: Comparison of feature importance methods using the classic Iris dataset.

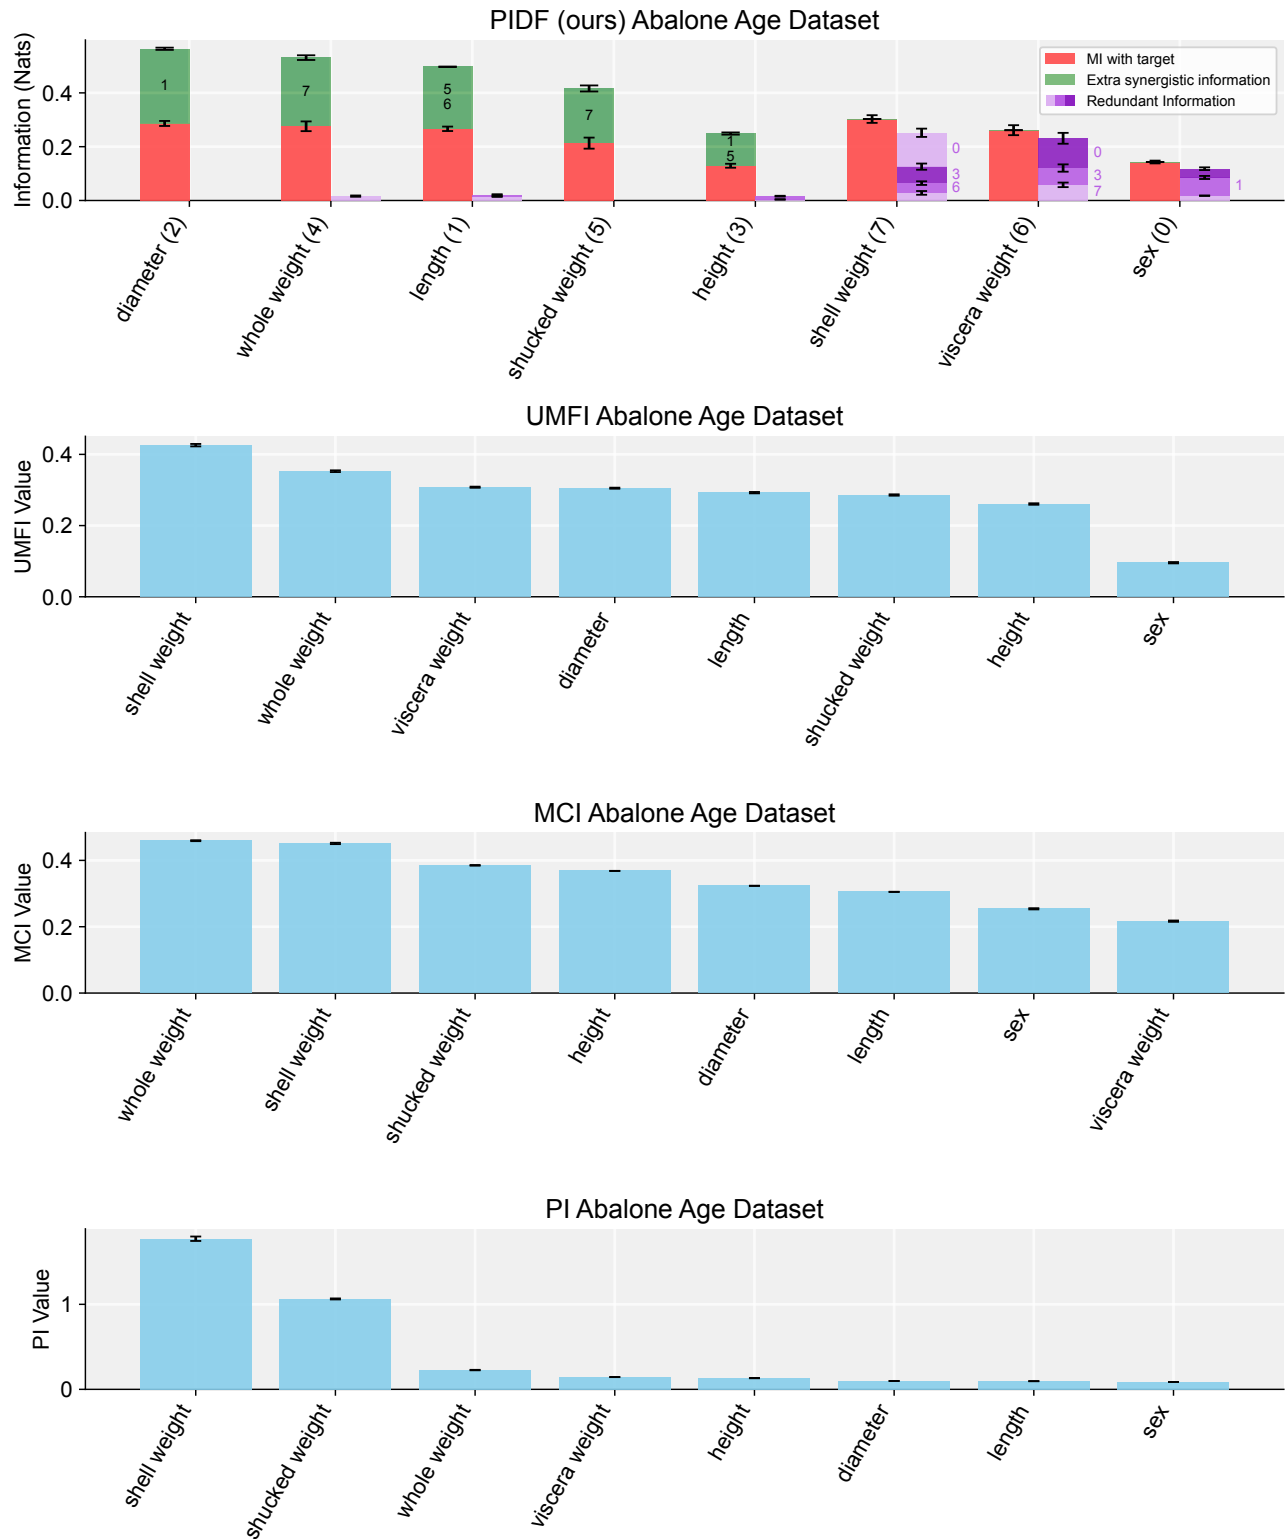


Figure 7: Comparison of feature importance methods using the classic Abalone dataset.

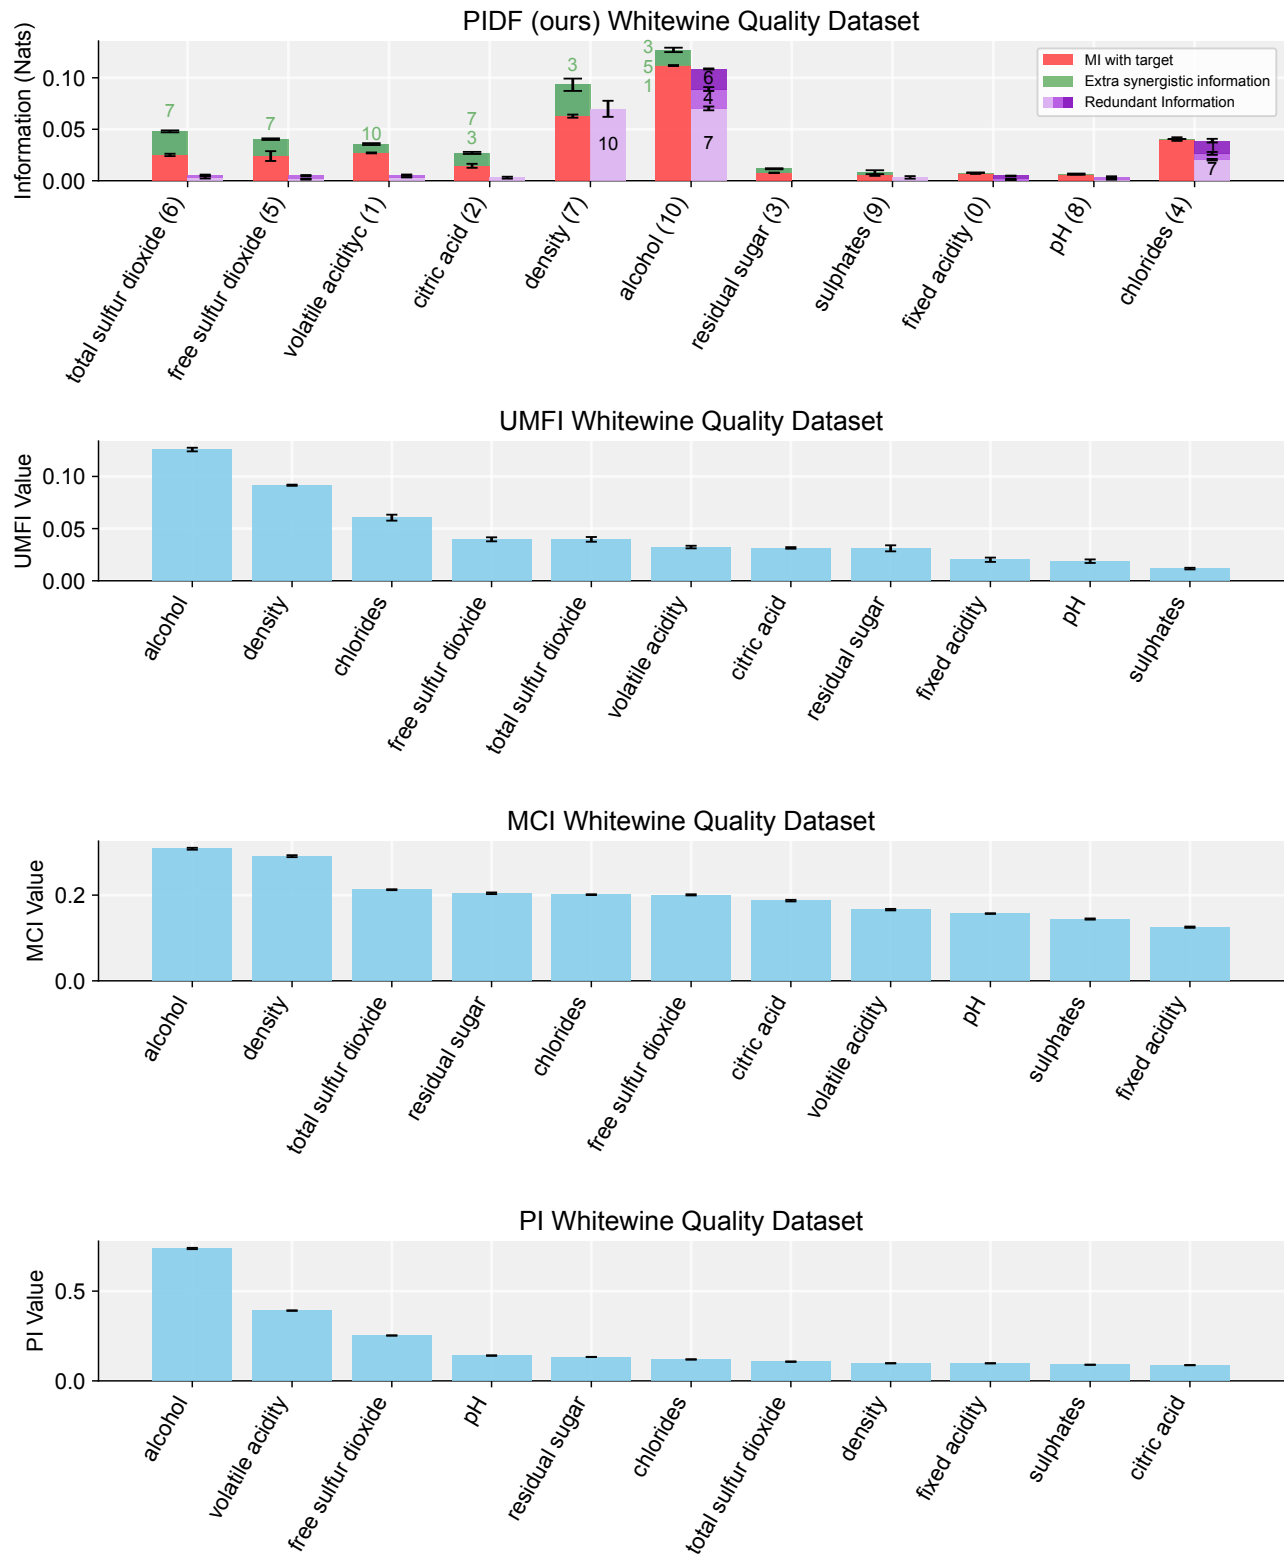


Figure 8: Comparison of feature importance methods using the classic whitewine quality dataset.

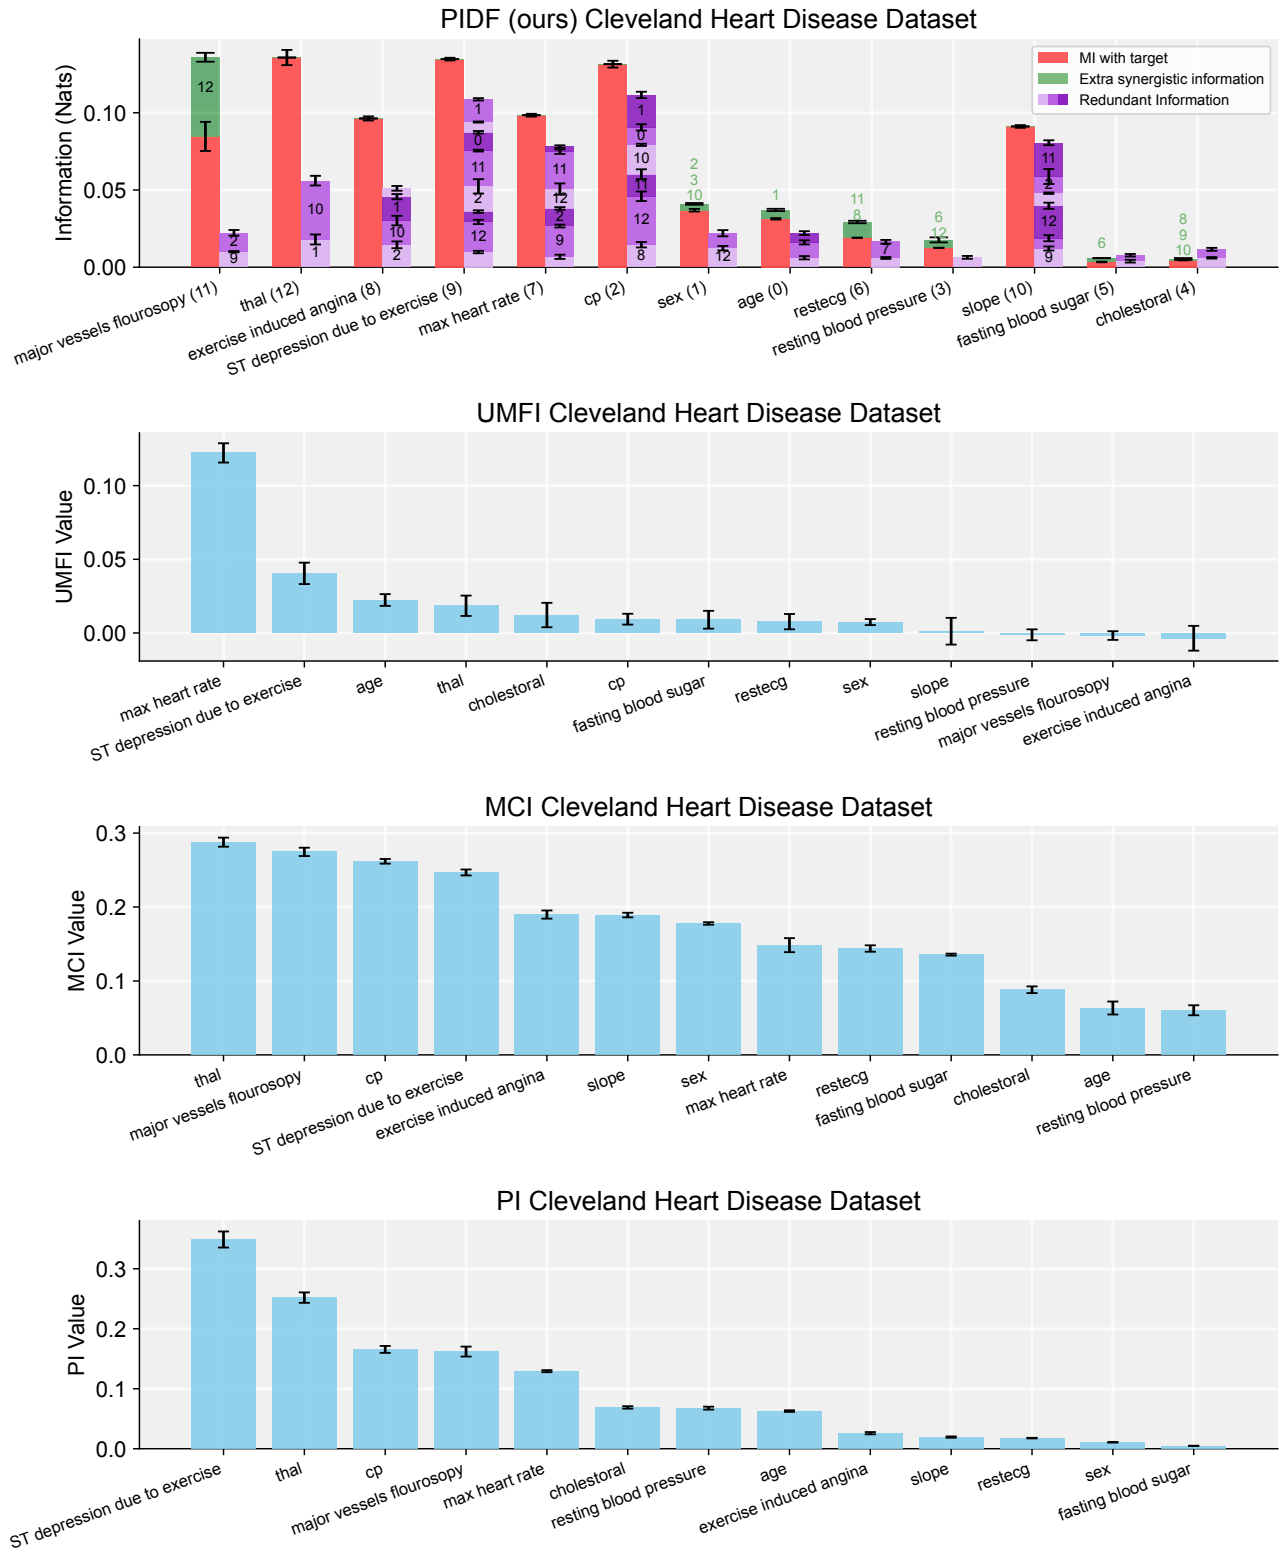


Figure 9: Cleveland heart disease prediction experiments.



Figure 10: Feature duplication experiments. We apply PIDF to a modified version of the California housing, Abalone, and Whitewine datasets. These modified datasets are obtained by adding a duplicate feature to the original ones. The duplicate features we used in our experiments are longitude, diameter and density, respectively.




Figure 11: In these plots, we compare MINE with other MI estimators by illustrating the number of iterations required to converge on the ground truth MI. MINE takes the longest to converge but yields the most accurate estimates with minimal variation. CEL quickly produces accurate estimates on discrete data but exhibits high variation. An MSE loss converges rapidly but results in incorrect estimates.




Figure 12: Comparison of the results of PIDF using MINE and PIDF using a different estimator.




Figure 13: The results achieved when applying scalable PIDF to MNIST.