

# Causal Entropy and Information Gain for Measuring Causal Control

Francisco Nunes Ferreira Quialheiro Simoes<sup>✉</sup>, Mehdi Dastani<sup>✉</sup>, and Thijs van Ommen<sup>✉</sup>

Department of Information and Computing Sciences, Utrecht University  
`{f.simoes,m.m.dastani,t.vanommen}@uu.nl`

**Abstract.** Artificial intelligence models and methods commonly lack causal interpretability. Despite the advancements in interpretable machine learning (IML) methods, they frequently assign importance to features which lack causal influence on the outcome variable. Selecting causally relevant features among those identified as relevant by these methods, or even before model training, would offer a solution. Feature selection methods utilizing information theoretical quantities have been successful in identifying statistically relevant features. However, the information theoretical quantities they are based on do not incorporate causality, rendering them unsuitable for such scenarios. To address this challenge, this article proposes information theoretical quantities that incorporate the causal structure of the system, which can be used to evaluate causal importance of features for some given outcome variable. Specifically, we introduce causal versions of entropy and mutual information, termed causal entropy and causal information gain, which are designed to assess how much control a feature provides over the outcome variable. These newly defined quantities capture changes in the entropy of a variable resulting from interventions on other variables. Fundamental results connecting these quantities to the existence of causal effects are derived. The use of causal information gain in feature selection is demonstrated, highlighting its superiority over standard mutual information in revealing which features provide control over a chosen outcome variable. Our investigation paves the way for the development of methods with improved interpretability in domains involving causation.

**Keywords:** Causal Inference · Information Theory · Interpretable Machine Learning · Explainable Artificial Intelligence

## 1 Introduction

Causality plays an important role in enhancing not only the prediction power of a model [19] but also its interpretability [4]. Causal explanations are more appropriate for human understanding than purely statistical explanations [12]. Accordingly, comprehending the causal connections between the variables of a system can enhance the interpretability of interpretable machine learning (IML) methods themselves.

Interpretable models such as linear regression or decision trees do not, despite their name, always lend themselves to *causal* interpretations. To illustrate this point, consider running multilinear regression on the predictors  $X_1, X_2$  and outcome  $Y$  within a system whose variables are causally related as depicted in the graph of Figure 1. The regression coefficients  $\beta_1$  and  $\beta_2$  of  $X_1$  and  $X_2$  might yield large values, which may be (and are often in practice) interpreted as suggesting a causal relationship. However, a causal interpretation of  $\beta_1$  would not be appropriate. Although  $X_1$  might provide predictive power over  $Y$ , this does not imply a causal relationship, since this predictive power is due to the confounder  $W$ . Consequently, intervening on  $X_1$  would not impact the outcome  $Y$ .

In current model-agnostic methods, a causal interpretation is often desirable but rarely possible. In partial dependence plots (PDPs) [6], the partial dependence of a model outcome  $\hat{Y}$  on a variable  $X_i$  coincides with the backdoor criterion formula [15] when the conditioning set encompasses all the other covariates  $X_{j \neq i}$  [24]. Consequently, there is a risk of disregarding statistical dependence or, conversely, finding spurious dependence, by conditioning on causal descendants of  $X_i$  [24]. Therefore, PDPs (along with the closely related individual conditional expectation (ICE) lines [7]) generally lack a causal interpretation. Similarly, when utilizing (Local Interpretable Model-Agnostic Explanations) LIME [18] to evaluate the importance of a feature for an individual, a causal interpretation cannot be guaranteed. LIME fits a local model around the point of interest and assesses which features, when perturbed, would cause the point to cross the decision boundary of the model. However, intervening on a feature in such a way as to cross the model’s decision boundary does not guarantee an actual change in the outcome in reality. This is because the model was trained on observational data, and that feature may merely be correlated with the outcome through a confounding factor, for example, rather than having a causal effect on the outcome.

In both cases just described, it is the presence of confounders, selection bias, or an incorrect direction of causality seemingly implied by the model that can lead to misleading predictions and interpretations. We need a way to select which features are causally relevant — *i.e.* give us control over the chosen outcome variable. Information theoretical quantities such as mutual information are often used to assess the relevance of a feature with respect to a given outcome variable [20,2,25], but this relevance is still purely statistical. This is a common issue when using standard information theoretical quantities in situations that require consideration of the underlying causal relationships. A version of mutual information which takes into account the causal structure of the system would solve this problem. This is what we set out to develop in this work.

In our research, we extend traditional conditional entropy and mutual information to the realm of *interventions*, as opposed to simple conditioning. This extension drew inspiration from the conceptual and philosophical work presented in<sup>1</sup>[8]. We dub these constructs “causal entropy” and “causal information gain”. They are designed to capture changes in the entropy of a given variable in re-

---

<sup>1</sup>The reader is referred to Section 6 for a detailed discussion about this.

sponse to manipulations affecting other variables. We derive fundamental results connecting these quantities to the presence of causal effect. We end by illustrating the use of causal information gain in selecting a variable which allows us to control an outcome variable, and contrast it with standard mutual information.

The novelty of our work consists of providing rigorous definitions for causal entropy and causal information gain, as well as deriving some of their key properties for the first time. These contributions set the foundations for the development of methods which correctly identify features which provide causal control over an outcome variable.

This paper is organized as follows. In Section 2, we introduce the definitions of quantities from the fields of causal inference and information theory that will be used throughout the rest of the paper. Section 3 includes a simple example of a structural causal model where standard entropy and mutual information are inadequate for obtaining the desired causal insights. In Section 4, we define causal entropy and explore its relation to total effect. Section 5 discusses the definition of causal information gain and investigates its connection with causal effect. Furthermore, it revisits the example from Section 3, showing that causal entropy and causal information gain allow us to arrive at the correct conclusions about causal control. In Section 6, we compare the definitions and results presented in this paper with those of previous work. Finally, in Section 7, we discuss the obtained results and propose future research directions.

## 2 Formal Setting

In this section we present the definitions from causal inference and information theory which are necessary for the rest of this paper. All random variables are henceforth assumed to be discrete and have finite range.

### 2.1 Structural Causal Models

One can model the causal structure of a system by means of a “structural causal model”, which can be seen as a Bayesian network [10] whose graph  $G$  has a causal interpretation and each conditional probability distribution (CPD)  $P(X_i | \text{PA}_{X_i})$  of the Bayesian network stems from a deterministic function  $f_{X_i}$  (called “structural assignment”) of the parents of  $X_i$ . In this context, it is common to separate the parent-less random variables (which are called “exogenous” or “noise” variables) from the rest (called “endogenous” variables). Only the endogenous variables are represented in the structural causal model graph. As is commonly done [16], we assume that the noise variables are jointly independent and that exactly one noise variable  $N_{X_i}$  appears as an argument in the structural assignment  $f_{X_i}$  of  $X_i$ . In full rigor<sup>2</sup>[16]:

**Definition 1 (Structural Causal Model).** *Let  $X$  be a random variable with range  $R_X$  and  $\mathbf{W}$  a random vector with range  $R_{\mathbf{W}}$ . A structural assignment for*

---

<sup>2</sup>We slightly rephrase the definition provided in [16] to enhance its clarity.

$X$  from  $\mathbf{W}$  is a function  $f_X: R_{\mathbf{W}} \rightarrow R_X$ . A structural causal model (SCM)  $\mathcal{C} = (\mathbf{X}, \mathbf{N}, S, p_{\mathbf{N}})$  consists of:

1. A random vector  $\mathbf{X} = (X_1, \dots, X_n)$  whose variables we call endogenous.
2. A random vector  $\mathbf{N} = (N_{X_1}, \dots, N_{X_n})$  whose variables we call exogenous or noise.
3. A set  $S$  of  $n$  structural assignments  $f_{X_i}$  for  $X_i$  from  $(PA_{X_i}, N_{X_i})$ , where  $PA_{X_i} \subseteq \mathbf{X}$  are called parents of  $X_i$ . The causal graph  $G^{\mathcal{C}} := (\mathbf{X}, E)$  of  $\mathcal{C}$  has as its edge set  $E = \{(P, X_i) : X_i \in \mathbf{X}, P \in PA_{X_i}\}$ . The  $PA_{X_i}$  must be such that the  $G^{\mathcal{C}}$  is a directed acyclic graph (DAG).
4. A jointly independent probability distribution  $p_{\mathbf{N}}$  over the noise variables. We call it simply the noise distribution.

We denote by  $\mathcal{C}(\mathbf{X})$  the set of SCMs with vector of endogenous variables  $\mathbf{X}$ . Furthermore, we write  $X := f_X(X, N_X)$  to mean that  $f_X(X, N_X)$  is a structural assignment for  $X$ .

Notice that for a given SCM the noise variables have a known distribution  $p_{\mathbf{N}}$  and the endogenous variables can be written as functions of the noise variables. Therefore the distributions of the endogenous variables are themselves determined if one fixes the SCM. This brings us to the notion of the entailed distribution<sup>2</sup> [16]:

**Definition 2 (Entailed distribution).** Let  $\mathcal{C} = (\mathbf{X}, \mathbf{N}, S, p_{\mathbf{N}})$  be an SCM. Its entailed distribution  $p_{\mathbf{X}}^{\mathcal{C}}$  is the unique joint distribution over  $\mathbf{X}$  such that  $\forall X_i \in \mathbf{X}, X_i = f_{X_i}(PA_{X_i}, N_{X_i})$ . It is often simply denoted by  $p^{\mathcal{C}}$ . Let  $\mathbf{x}_{-i} := (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ . For a given  $X_i \in \mathbf{X}$ , the marginalized distribution  $p_{X_i}^{\mathcal{C}}$  given by  $p_{X_i}^{\mathcal{C}}(x_i) = \sum_{\mathbf{x}_{-i}} p_{\mathbf{X}}^{\mathcal{C}}(\mathbf{x})$  is also referred to as entailed distribution (of  $X_i$ ).

An SCM allows us to model interventions on the system. The idea is that an SCM represents how the values of the random variables are generated, and by intervening on a variable we are effectively changing its generating process. Thus intervening on a variable can be modeled by modifying the structural assignment of said variable, resulting in a new SCM differing from the original only in the structural assignment of the intervened variable, and possibly introducing a new noise variable for it, in place of the old one. Naturally, the new SCM will have an entailed distribution which is in general different from the distribution entailed by the original SCM.

The most common type of interventions are the so-called “atomic interventions”, where one sets a variable to a chosen value, effectively replacing the distribution of the intervened variable with a point mass distribution. In particular, this means that the intervened variable has no parents after the intervention. This is the only type of intervention that we will need to consider in this work. Formally<sup>2</sup> [16]:

**Definition 3 (Atomic intervention).** Let  $\mathcal{C} = (\mathbf{X}, \mathbf{N}, S, p_{\mathbf{N}})$  be an SCM,  $X_i \in \mathbf{X}$  and  $x \in R_{X_i}$ . The atomic intervention  $do(X_i = x)$  is the function

$\mathcal{C}(\mathbf{X}) \rightarrow \mathcal{C}(\mathbf{X})$  given by  $\mathcal{C} \mapsto \mathcal{C}^{do(X_i=x)}$ , where  $\mathcal{C}^{do(X_i=x)}$  is the SCM that differs from  $\mathcal{C}$  only in that the structural assignment  $f_{X_i}(PA_{X_i}, N_{X_i})$  is replaced by the structural assignment  $\tilde{f}_{X_i}(\tilde{N}_{X_i}) = \tilde{N}_{X_i}$ , where  $\tilde{N}_{X_i}$  is a random variable with range  $R_{X_i}$  and  $p_{\tilde{N}_{X_i}}(x_i) = \mathbf{1}_x(x_i)$  for all  $x_i \in R_{X_i}$ . Such SCM is called the post-atomic-intervention SCM. One says that the variable  $X_i$  was (atomically) intervened on. The distribution  $p^{do(X_i=x)} := p^{\mathcal{C}^{do(X_i=x)}}$  entailed by  $\mathcal{C}^{do(X_i=x)}$  is called the post-intervention distribution (w.r.t. the atomic intervention  $do(X_i = x)$  on  $\mathcal{C}$ ).

We can also define what we mean by “ $X$  having a total causal effect on  $Y$ ”. Following [16,14], there is such a total causal effect if there is an atomic intervention on  $X$  which modifies the initial distribution of  $Y$ <sup>2</sup> [16]:

**Definition 4 (Total Causal Effect).** *Let  $X, Y$  be random variables of an SCM  $\mathcal{C}$ .  $X$  has a total causal effect on  $Y$ , denoted by  $X \rightarrow Y$ , if there is  $x \in R_X$  such that  $p_Y^{do(X=x)} \neq p_Y$ .*

In this work, all variables of the form  $X_i, Y_i$  or  $Z_i$  are taken to be endogenous variables of some SCM  $\mathcal{C}$ .

## 2.2 Entropy and Mutual Information

Since the quantities defined and studied in this article build upon the standard entropy and mutual information, it is important for the reader to be familiar with these. In this subsection we will state the definitions of entropy, conditional entropy and mutual entropy. In the interest of space, we will not try to motivate these definitions. For a pedagogical introduction the reader is referred to [5,11]. We will also clarify what we precisely mean by causal control.

**Definition 5 (Entropy and Conditional Entropy [5]).** *Let  $X$  be a discrete random variable with range  $R_X$  and  $p$  be a probability distribution for  $X$ . The entropy of  $X$  w.r.t. the distribution  $p$  is<sup>4</sup>*

$$H_{X \sim p}(X) := - \sum_{x \in R_X} p(x) \log p(x). \quad (1)$$

*Entropy is measured in bit. If the context suggests a canonical probability distribution for  $X$ , one can write  $H(X)$  and refers to it simply as the entropy of  $X$ . The conditional entropy  $H(Y | X)$  of  $Y$  conditioned on  $X$  is the expected value w.r.t.  $p_X$  of the entropy  $H(Y | X = x) := H_{Y \sim p_{Y|X=x}}(Y)$ :*

$$H(Y | X) := \mathbb{E}_{x \sim p_X} [H(Y | X = x)]. \quad (2)$$

---

<sup>3</sup>We denote by  $\mathbf{1}_x$  the indicator function of  $x$ , so that  $\mathbf{1}_x(x_i) = \begin{cases} 1, & x_i = x \\ 0, & \text{otherwise} \end{cases}$ .

<sup>4</sup>In this article,  $\log$  denotes the logarithm to the base 2.

This means that the conditional entropy  $H(Y | X)$  is the entropy of  $H(Y)$  that remains on average if one conditions on  $X$ .

An essential concept closely associated with entropy is that of “uncertainty.” This qualitative concept is often present when interpreting information-theoretical quantities. The entropy of a variable  $X$  purports to measure the uncertainty regarding  $X$ . In this paper, we use another qualitative concept called “causal control” (or simply “control”). The (causal) control that variable  $X$  has over variable  $Y$  is the level of uncertainty remaining about  $Y$  after intervening on  $X$ . It indicates how close we are to fully specifying  $Y$  by intervening on  $X$ . This understanding of the term “control” has been implicitly utilized in the philosophy of science literature [17,3].

*Remark 1.* Notice that  $H(Y | X = x)$  is seen as a function of  $x$  and the expected value in Equation (2) is taken over the random variable  $x$  with distribution  $p_X$ . This disrespects the convention that random variables are represented by capital letters, but preserves the convention that the specific value conditioned upon (even if that value can be randomly realized — *i.e.* is a random variable) is represented by a lower case letter. Since we cannot respect both, we will follow the common practice and opt to use lower case letters for random variables in these cases.

There are two common equivalent ways to define mutual information (often called information gain).

**Definition 6 (Mutual Information [5]).** *Let  $X$  and  $Y$  be discrete random variables with ranges  $R_X$  and  $R_Y$  and distributions  $p_X$  and  $p_Y$ , respectively. The mutual information between  $X$  and  $Y$  is the KL divergence between the joint distribution  $p_{X,Y}$  and the product distribution  $p_X p_Y$ , i.e.:*

$$I(X; Y) := \sum_{x,y \in R_X \times R_Y} p_{X,Y}(x,y) \log \frac{p_{X,Y}(x,y)}{p_X(x)p_Y(y)}. \quad (3)$$

Or equivalently:

$$\begin{aligned} I(X; Y) &:= H(Y) - H(Y | X) \\ &= H(X) - H(X | Y). \end{aligned} \quad (4)$$


The view of mutual information as entropy reduction from Equation (4) will be the starting point for our definition of causal information gain.

### 3 Running Example - Comparing Control Over an Outcome

We provide a simple example showcasing how the standard entropy and mutual information can fail to assess which variable gives us more control over a chosen outcome variable. We will later (Section 5) check that using causal entropy and causal information gain enable us to correctly make this assessment.

*Example 1.* Let us consider an ice-cream shop where the sales volume  $Y$  on a given day can be categorized as low ( $Y = 0$ ), medium ( $Y = 1$ ), or high ( $Y = 2$ ). We would like to find a way to control  $Y$ . Assume that the sales volume is influenced by two factors: the temperature  $W$ , characterized as warm ( $W = 1$ ) or cold ( $W = 0$ ), and whether the ice-cream shop is being advertised, represented by the binary variable  $X_2$ . Additionally, we introduce a discrete variable  $X_1$  to represent the number of individuals wearing shorts, which can be categorized as few ( $X_1 = 0$ ), some ( $X_1 = 1$ ), or many ( $X_1 = 2$ ). Naturally, higher temperatures have a positive influence on the variable  $X_1$ . We do not consider any other variables.

One can crudely model this situation using an SCM with endogenous variables  $X_1, X_2, W$  and  $Y$ , as specified in Figure 1. The chosen structural assignments and noise distributions reflect the specific scenario where: the temperature  $W$  is warm about half of the time; the number  $X_1$  of people wearing shorts is highly determined by the weather conditions; and the ice-cream shop is advertised occasionally.  $W, X_2$  and all noise variables of the SCM are binary variables, while  $X_1, Y \in \{0, 1, 2\}$ . Assume we cannot intervene on  $W$ . We would like to decide which of the variables  $X_1$  or  $X_2$  provide us with the most control over  $Y$ .



**Fig. 1.** An SCM<sup>6</sup>. It models the real-world scenario described in Example 1, where  $Y$  is the sales volume of an ice-cream shop,  $W$  is the temperature,  $X_1$  is the amount of people wearing shorts, and  $X_2$  stands for the advertisement efforts of the ice-cream shop. The notation  $N_Z \sim \text{Bern}(q)$  signifies that the random variable  $Z$  follows the Bernoulli probability distribution with parameter  $q$ . Grayed out variables cannot be intervened on.

It is clear that being able to intervene on  $X_1$  gives us no control whatsoever over  $Y$ . Any observed statistical dependence between  $X_1$  and  $Y$  comes purely from the confounder  $W$ . Consequently, interpreting a non-zero correlation or mutual information between  $X_1$  and  $Y$  as indicative of a causal connection between these variables would be a mistake, and an instance of conflation between correlation and causation.

<sup>6</sup>The careful reader may notice that there is no noise variable  $N_Y$  for  $Y$ , which seems to conflict with Definition 1. Such apparent conflicts are resolved by seeing a deterministic assignment function such as  $Y := X_2 + W$  as having a trivial additive dependence on a noise variable  $N_Y$  with a point mass distribution at 0.

If we naively use the mutual information to assess whether one should intervene on  $X_1$  or  $X_2$  for controlling  $Y$ , one wrongly concludes that one should use  $X_1$ . Intuitively, this happens because knowing  $X_2$  provides us with less information about  $Y$  than  $W$ , and  $X_1$  is very close to  $W$ . The (approximate) values can be consulted<sup>6</sup> in Table 1.

**Table 1.** Information theoretical values for Figure 1.

$H(Y) \approx 1.41$	$H(Y   X_1) \approx 0.85$	$H(Y   X_2) = 1$
$I(Y; W) \approx 0.60$	$I(Y; X_1) \approx 0.56$	$I(Y; X_2) \approx 0.41$

Notice that  $I(Y; W) > I(Y; X_1)$ , as it should be:  $W$  has more information about  $Y$  than  $X_1$  has. We also see that  $I(Y; X_2) < I(Y; X_1)$ . If mutual information were a suitable criterion for selecting the variable to intervene on, the contrary would be expected. In the context of our real-world scenario, intervening on the number  $X_1$  of people wearing shorts would not be a logical approach for controlling ice cream sales. Instead, allocating more resources to advertising efforts (represented by  $X_2$ ) would be more appropriate.

The issue is that the mutual information  $I(Y; X_1)$  includes the information that one has about  $Y$  by *observing*  $X_1$  which flows through the confounder  $W$ . But what we want is a metric quantifying how much control we can have over  $Y$  by *intervening* on  $X_1$ . We will see that the generalization of mutual information studied in this paper (“causal information gain”) satisfies these requirements.

## 4 Causal Entropy

The causal entropy of  $Y$  for  $X$  will be the entropy of  $Y$  that is left, on average, after one atomically intervenes on  $X$ . In this section we give a rigorous definition of causal entropy and study its connection to causal effect.

We define causal entropy in a manner analogous to conditional entropy (see Definition 5). It will be the average uncertainty one has about  $Y$  if one sets  $X$  to  $x$  with probability  $p_{X'}(x)$ , where  $X'$  is a new auxiliary variable with the same range as  $X$  but independent of all other variables, including  $X$ . In contrast with the non-causal case, here one needs to make a choice of distribution over  $X'$  corresponding to the distribution over the atomic interventions that one is intending to perform.

**Definition 7 (Causal entropy,  $H_c$ ).** *Let  $Y$ ,  $X$  and  $X'$  be random variables such that  $X$  and  $X'$  have the same range and  $X'$  is independent of all variables in  $\mathcal{C}$ . We say that  $X'$  is an intervention protocol for  $X$ .*

<sup>6</sup>The details of the computations can be found in Appendix A.



The causal entropy  $H_c(Y \mid do(X \sim X'))$  of  $Y$  given the intervention protocol  $X'$  for  $X$  is the expected value w.r.t.  $p_{X'}$  of the entropy  $H(Y \mid do(X = x)) := H_{Y \sim p_Y^{do(X=x)}}(Y)$  of the interventional distribution  $p_Y^{do(X=x)}$ . That is:

$$H_c(Y \mid do(X \sim X')) := \mathbb{E}_{x \sim p_{X'}} [H(Y \mid do(X = x))] \quad (5)$$

We will now see that, unsurprisingly, if there is no total effect of  $X$  on  $Y$ , then the causal entropy is just the initial entropy  $H(Y)$ . Perhaps more unexpectedly, the converse is not true: it is possible to have  $H_c(Y \mid X \sim X') = H(Y)$  while  $X \rightarrow Y$ . One way this can happen is due to the non-injectivity of entropy when seen as a mapping from the set of distributions over  $Y$ , i.e. it may happen that  $p_Y^{do(X=x)} \neq p_Y$  but  $H_{Y \sim p_Y^{do(X=x)}}(Y) = H_{Y \sim p_Y}(Y)$ .

**Proposition 1.** *If there is no total effect of  $X$  on  $Y$ , then  $H_c(Y \mid do(X \sim X')) = H(Y)$  for any intervention protocol  $X'$  for  $X$ . The converse does not hold.*

*Proof.* The proof can be found in Appendix B.  $\square$

If there is a total causal effect of  $X$  on  $Y$ , there cannot be a total causal effect of  $Y$  on  $X$  (if  $X$  is a cause of  $Y$ ,  $Y$  cannot be a cause of  $X$ ) [16]. This immediately yields the following corollary.

**Corollary 1.** *If  $H_c(Y \mid do(X \sim X')) \neq H(Y)$  for some intervention protocol  $X'$  for  $X$ , then  $H_c(X \mid do(Y \sim Y')) = H(X)$  for any intervention protocol  $Y'$  for  $Y$ .*

*Proof.* Suppose that  $H_c(Y \mid X \sim X') \neq H(Y)$ . By the contrapositive of Proposition 1, this means that there is a total effect of  $X$  on  $Y$ . Hence there is no total effect of  $Y$  on  $X$ , which again by Proposition 1 yields the desired result.  $\square$

## 5 Causal Information Gain

Causal information gain extends mutual information to the causal context. The causal information gain of  $Y$  for  $X$  will be the average decrease in the entropy of  $Y$  after one atomically intervenes on  $X$ . We start this section by giving a rigorous definition of causal information gain, and proceed to study its connection with causal effect. We end this section by revisiting Example 1 armed with this new information theoretical quantity. We will confirm in this example that causal information is the correct tool for assessing which variable has the most causal control over the outcome, as opposed to standard mutual information.

Recall the entropy-based definition of mutual information in Equation (4). The mutual information between two variables  $X$  and  $Y$  is the average reduction in uncertainty about  $Y$  if one observes the value of  $X$  (and vice-versa, by symmetry of the mutual information). This view of mutual information allows for a straightforward analogous definition in the causal case, so that one can take causal information gain  $I_c(Y \mid do(X \sim X'))$  to signify the average reduction in uncertainty about  $Y$  if one sets  $X$  to  $x$  with probability  $p_{X'}(x)$ .

**Definition 8 (Causal Information Gain,  $I_c$ ).** *Let  $Y$ ,  $X$  and  $X'$  be random variables such that  $X'$  is an intervention protocol for  $X$ . The causal information gain  $I_c(Y \mid do(X \sim X'))$  of  $Y$  for  $X$  given the intervention protocol  $X'$  is the difference between the entropy of  $Y$  w.r.t. its prior and the causal entropy of  $Y$  given the intervention protocol  $X'$ . That is:*

$$I_c(Y \mid do(X \sim X')) := H(Y) - H_c(Y \mid do(X \sim X')). \quad (6)$$

A few properties of causal information gain can be immediately gleaned from its definition. First, in contrast with mutual information, causal information gain is *not* symmetric. Also, similarly to causal entropy, one needs to specify an intervention protocol with a distribution to be followed by interventions on  $X$ .

We can make use of the relation between causal entropy and causal effect to straightforwardly deduce the relation between causal information gain and causal effect.

**Proposition 2.** *If  $I_c(Y \mid do(X \sim X')) \neq 0$  for some protocol  $X'$  for  $X$ , then  $X \rightarrow Y$ . The converse does not hold.*

*Proof.* The implication in this proposition follows directly from Definition 8 and the contrapositive of the implication in Proposition 1. The converse does not hold simply because it is equivalent to the converse of the contrapositive of the implication in Proposition 1, which also does not hold.  $\square$

**Corollary 2.** *Let  $X'$  and  $Y'$  be intervention protocols for  $X$  and  $Y$ , respectively. At least one of  $I_c(Y \mid do(X \sim X'))$  or  $I_c(X \mid do(Y \sim Y'))$  is zero.*

*Proof.* Suppose both  $I_c(Y \mid do(X \sim X'))$  and  $I_c(X \mid do(Y \sim Y'))$  are non-zero. Then by Proposition 2 we have both  $X \rightarrow Y$  and  $Y \rightarrow X$ , which is not possible in the context of an SCM.  $\square$

It is worth noting that the last part of Proposition 2 contradicts [17]. In that work, it is stated without proof that “causation is equivalent to non-zero specificity”, wherein the term “specificity” coincides with what we refer to as causal information gain given a uniformly distributed intervention protocol.

### 5.1 Comparison of Causal Information Gain and Mutual Information in Running Example

Consider again Example 1. Compare the causal entropy and causal information gain values<sup>7</sup> in Table 2 with the conditional entropy and mutual information values from Table 1.

We see that using causal information gain allows us to correctly conclude that using  $X_1$  to control  $Y$  would be fruitless: intervening on  $X_1$  does not change the

---

<sup>7</sup>In this particular case it does not matter what intervention protocol  $X'$  we choose, since  $H_c(Y \mid do(X_1 = x_1)) = H(Y) \approx 1.41$  for all  $x_1$  and  $H_c(Y \mid do(X_2 = x_2)) = 1$  for all  $x_2$ .

**Table 2.** Causal information theoretical values for Figure 1.

$H_c(Y \mid do(X_1 \sim X'_1)) \approx 1.41$	$H_c(Y \mid do(X_2 \sim X'_2)) = 1$
$I_c(Y \mid do(X_1 \sim X'_1)) = 0$	$I_c(Y \mid do(X_2 \sim X'_2)) \approx 0.41$

entropy of  $Y$ . This is reflected by the fact that the causal information gain of  $Y$  for  $X_1$  is zero. Since  $X_1$  has no causal effect on  $Y$ , this result was to be expected by the contrapositive of Proposition 2. On the other hand,  $X_2$  does provide us with some control over  $Y$ : intervening on  $X_2$  decreases the entropy of  $Y$  by 0.4 bit on average. In the real-world scenario described in Example 1, utilizing causal information gain to determine which variable to intervene on for controlling the sales volume  $Y$  would lead us to make the correct decision of intensifying advertising efforts ( $X_2$ ). Furthermore, it would enable us to conclude that manipulating the number of people wearing shorts ( $X_1$ ) provides no control whatsoever over  $Y$ . Thus, causal information gain could be used in this case to assess whether statistical dependence between  $Y$  and another variable in this causal system can be interpreted to have causal significance.

## 6 Related Work

Previous work has aimed to provide causal explanations of machine learning models through “counterfactual explanations” [21,13]. These explanations reveal what the model would have predicted under different feature values. However, they do not offer insights into the causal significance of a feature in influencing the outcome variable. Instead, they merely inform us about the behavior of the model itself. In other words, counterfactual explanations inform us about the changes required for the model to produce a different prediction, but not the changes necessary for the outcome to differ in reality. While counterfactual explanations can be useful, for instance, in advising loan applicants on improving their chances of approval [13], they fall short in providing causal interpretations for tasks such as scientific exploration [23], where it is crucial to understand the actual causal relationships between features and the chosen outcome. As discussed in Section 1, the quantities investigated in this paper can precisely address this need.

Information theoretical quantities aimed at capturing aspects of causality have been previously proposed. An important example is the work in [9]. In that paper, the authors suggest a list of postulates that a measure of causal strength should satisfy, and subsequently demonstrate that commonly used measures fall short of meeting them. They then propose their own measure (called “causal influence”), which does satisfy the postulates. Causal influence is the KL divergence of the original joint distribution and the joint distribution resulting from removing the arrows whose strength we would like to measure, and feeding noise to the orphaned nodes. Thus although it utilizes information theory, it does not purport to generalize entropy or mutual information to the causal context. One

information-theoretical measure mentioned in [9] is closer to ours. It is called “information flow” [1]. Similarly to causal information gain, this quantity is a causal generalization of mutual information. Their goal was to come up with a generalization of mutual information that would be a measure of “causal independence” in much the same way as standard mutual information is a measure of statistical independence. They take the route of starting from the definition of mutual information as the KL divergence between the joint distribution and the product of the marginal distributions (Equation (3)), and proceed to “make it causal” by effectively replacing conditioning with intervening everywhere. In contrast, we treat entropy as the main quantity of interest, and start from the definition of mutual entropy as the change in entropy due to conditioning (Equation (4)), and proceed to define its causal counterpart as the change in entropy due to intervening. This then results in a quantity that is the appropriate tool for evaluating the control that a variable has over another.

The basic idea of extending the concept of mutual information to the causal context as the average reduction of entropy after intervening was introduced in the philosophy of science literature, as part of an attempt to capture a property of causal relations which they refer to as “specificity” [8]. This property can be thought of as a measure of the degree to which interventions on the cause variable result in a deterministic one-to-one mapping [22]. This means that maximal specificity of a causal relationship is attained when: (a) performing an atomic intervention on the cause variable results in complete certainty about the effect variable’s value; and (b) no two distinct atomic interventions on the cause variable result in the same value for the effect variable [8]. Notice that (a) means precisely that the cause variable provides maximal causal control over the effect variable. The causal extension of mutual information proposed in [8] was named “causal mutual information”. They call “causal entropy” the average entropy of the effect variable after performing an atomic intervention on the cause variable. Their “causal mutual information” is then the difference between the initial entropy of the effect variable and the causal entropy. Although they do not say so explicitly, their definition of causal entropy assumes that one only cares about the entropy that results from interventions that are equally likely: the average of post-intervention entropies is taken w.r.t. a uniform distribution — hence their “causal entropy” is the same as the causal entropy defined in this paper, but restricted to uniform intervention protocols. This was also noted in [17], where the authors propose that other choices of distribution over the interventions would result in quantities capturing causal aspects that are distinct from the standard specificity. In this paper we both generalized and formalized the information theoretical notions introduced in [8]. We provided rigorous definitions of causal entropy and causal information gain which allow for the use of non-uniform distributions over the interventions. Our causal entropy can thus be seen as a generalized version of their causal entropy, while our causal information gain can be seen as a generalization of their causal mutual information<sup>8</sup>. Armed

---

<sup>8</sup>The term causal mutual information may be misleading given the directional nature of the relationship between cause and effect. We thus prefer the term causal infor-

with concrete, mathematical definitions, we are able to study key mathematical aspects of these quantities.

## 7 Discussion and Conclusion

The motivation behind extending traditional entropy and mutual information to interventional settings in the context of interpretable machine learning (IML) arises from the necessity to determine whether the high importance assigned to specific features by machine learning models and IML methods can be causally interpreted or is purely of a statistical nature.

Information theoretical quantities are commonly used to assess statistical feature importance. We extended these quantities to handle interventions, allowing them to capture the control one has over a variable by manipulating another. The proposed measures, namely causal entropy and causal information gain, hold promise for the development of new algorithms in domains where knowledge of causal relationships is available or obtainable. It is worth noting that the utility of these measures extends well beyond the field of IML, as both information-theoretical quantities and the need for causal control are pervasive in machine learning.

Moving forward, a crucial theoretical endeavor involves establishing a fundamental set of properties for the proposed causal information-theoretical measures. This can include investigating a data processing inequality and a chain rule for causal information gain, drawing inspiration from analogous properties associated with mutual information. Other important research directions involve the extension of these definitions to continuous variables, as well as investigating the implications of employing different intervention protocols. Furthermore, the design and study of appropriate estimators for these measures constitute important avenues for future research, as well as their practical implementation. Ideally, these estimators should be efficient to compute even when dealing with high-dimensional data and complex, real-world datasets. Additionally, they ought to be applicable to observational data. In cases where the structural causal model is known, this could be accomplished by utilizing a framework such as *do*-calculus [14] when devising the estimators. This could allow for their application in extracting causal insights from observational data.

## A Computations for the running example

We have

$$\begin{aligned} H(Y) &= p_Y(0) \log\left(\frac{1}{p_Y(0)}\right) + p_Y(1) \log\left(\frac{1}{p_Y(1)}\right) + p_Y(2) \log\left(\frac{1}{p_Y(2)}\right) \\ &= \frac{3}{8} \log\left(\frac{8}{3}\right) + \frac{1}{2} \log(2) + \frac{1}{8} \log(8) = 2 - \frac{3}{8} \log(3) \approx 1.41 \text{ (bit)}, \end{aligned}$$

---

mation gain, drawing inspiration from the alternate name “information gain”, which is frequently employed in discussions about decision trees when referring to mutual information.

and

$$H(Y | W) = H(Y | W = 0) = \frac{3}{4} \log\left(\frac{4}{3}\right) + \frac{1}{4} \log(4) \approx 0.81 \text{ (bit)},$$

where we used that  $H(Y | W = 0) = H(Y | W = 1)$ , so that taking the average is unnecessary.

Notice that  $X_1 = 0$  implies  $W = 0$ , in which case  $Y = X_2$ . Hence  $H(Y | X_1 = 0) = H(Y | W = 0) \approx 0.81$  (bit). By a similar argument,  $H(Y | X_1 = 2) = H(Y | W = 1) \approx 0.81$  (bit). Now, denote  $q = \frac{1}{64}$ . It is easy to check that  $p_{Y|X_1=1}(0) = \frac{3q}{4}$ ,  $p_{Y|X_1=1}(1) = \frac{3}{4} - \frac{q}{2}$  and  $p_{Y|X_1=1}(2) = \frac{1}{4}(1 - q)$ . Then

$$H(Y | X_1 = 1) = -\frac{3q}{4} \log\left(\frac{3q}{4}\right) - \left(\frac{3}{4} - \frac{q}{2}\right) \log\left(\frac{3}{4} - \frac{q}{2}\right) - \frac{1}{4}(1 - q) \log\left(\frac{1}{4}(1 - q)\right) \approx 0.89 \text{ (bit)}.$$

We can then compute:

$$\begin{aligned} H(Y | X_1) &= p_{X_1}(0) \overbrace{H(Y | X_1 = 0)}^{0.81} + p_{X_1}(1) \overbrace{H(Y | X_1 = 1)}^{0.89} + p_{X_1}(2) \overbrace{H(Y | X_1 = 2)}^{0.81} \\ &= \frac{1}{2} \times (1 - q) \times 0.81 + \frac{1}{2} \times 0.89 + \frac{q}{2} \times 0.81 \approx 0.85 \text{ (bit)}. \end{aligned}$$

We also have:

$$H(Y | X_2) = p_{X_2}(0) \overbrace{H(Y | X_2 = 0)}^1 + p_{X_2}(1) \overbrace{H(Y | X_2 = 1)}^1 = 1 \text{ (bit)},$$

It immediately follows that  $I(Y; W) \approx 0.60$ ,  $I(Y; X_1) \approx 0.56$  (bit) and  $I(Y; X_2) \approx 0.41$  (bit).

Moving on to the causal information theoretical quantities, we have  $H(Y | do(X_1 = x_1)) = H(Y) \approx 1.41$  (bit) for every  $x_1 \in R_{X_1}$  and  $H(Y | do(X_2 = x_2)) = H(W) = 1$  (bit) for every  $x_2 \in R_{X_2}$ . Hence  $H_c(Y | do(X_1 \sim X'_1)) \approx 1.41$  (bit) and  $H_c(Y | do(X_2 \sim X'_2)) = 1$  (bit) for any intervention protocols  $X'_1, X'_2$ . It follows that  $I_c(Y | do(X_1 \sim X'_1)) = 0$  (bit) and  $I(Y | do(X_2 \sim X'_2)) \approx 0.41$  (bit).

## B Proof of Proposition 1

*Proof.* Suppose  $X$  has no causal effect on  $Y$ . Then  $\forall x \in R_X$ ,  $p_Y^{do(X=x)} = p_Y$ . The expression for the causal entropy then reduces to  $\mathbb{E}_{x \sim X'} H(Y) = H(Y)$ . This shows the implication in the proposition.

We will check that the converse does not hold by giving an example where  $X$  has a causal effect on  $Y$  but  $H_c(Y | X \sim X') = H(Y)$ . Consider the SCM

with three binary endogenous variables  $X, Y$  and  $M$  specified by:

$$\begin{cases} f_M(N_M) = N_M \\ f_X(M, N_X) = \begin{cases} (N_X + 1) \mod 2, M = 1 \\ N_X, M = 0 \end{cases} \\ f_Y(M, N_X) = \begin{cases} X, M = 1 \\ (X + 1) \mod 2, M = 0 \end{cases} \\ N_X, N_M \sim \text{Bern}(q), \text{ for some } q \in (0, 1). \end{cases} \quad (7)$$

Then  $p_Y^{do(X=0)} \sim \text{Bern}(q)$  and  $p_Y^{do(X=1)} \sim \text{Bern}(q)$ . Also,

$$p_Y = p_{X|M=1}(1)p_M(1) + p_{X|M=0}(0)p_M(0) = 1 - q \Rightarrow Y \sim \text{Bern}(1 - q) \quad (8)$$

Hence  $p_Y \neq p_Y^{do(X=1)}$ , meaning that  $X \rightarrow Y$ . And since both post-intervention distributions have the same entropy  $H_{Y \sim \text{Bern}(q)}(Y) = H_{Y \sim \text{Bern}(1-q)}(Y)$ , then the causal entropy will also be  $H_c(Y \mid X \sim X') = H_{Y \sim \text{Bern}(1-q)}(Y) = H(Y)$  (for any chosen of  $X'$ ).  $\square$

## References

1. Ay, N., Polani, D.: Information flows in causal networks. *Advances in complex systems* **11**(01), 17–41 (2008)
2. Beraha, M., Metelli, A.M., Papini, M., Tirinzoni, A., Restelli, M.: Feature selection via mutual information: New theoretical insights. *CoRR* **abs/1907.07384** (2019), <http://arxiv.org/abs/1907.07384>
3. Bourrat, P.: Variation of information as a measure of one-to-one causal specificity. *European Journal for Philosophy of Science* **9**(1), 1–18 (2019)
4. Confalonieri, R., Coba, L., Wagner, B., Besold, T.R.: A historical perspective of explainable artificial intelligence. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* **11**(1), e1391 (2021)
5. Cover, T.M., Thomas, J.A.: *Elements of information theory*. Wiley-Interscience (2006)
6. Friedman, J.H.: Greedy function approximation: a gradient boosting machine. *Annals of statistics* pp. 1189–1232 (2001)
7. Goldstein, A., Kapelner, A., Bleich, J., Pitkin, E.: Peeking inside the black box: Visualizing statistical learning with plots of individual conditional expectation. *Journal of Computational and Graphical Statistics* **24**(1), 44–65 (2015)
8. Griffiths, P.E., Pocheville, A., Calcott, B., Stotz, K., Kim, H., Knight, R.: Measuring causal specificity. *Philosophy of science* **82**(4), 529–555 (2015)
9. Janzing, D., Balduzzi, D., Grosse-Wentrup, M., Schölkopf, B.: Quantifying causal influences. *The Annals of Statistics* **41**(5), 2324–2358 (2013)
10. Koller, D., Friedman, N.: *Probabilistic graphical models: principles and techniques*. MIT press (2009)
11. MacKay, D.J., Mac Kay, D.J.: *Information theory, inference and learning algorithms*. Cambridge university press (2003)
12. Miller, T.: Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence* **267**, 1–38 (2019)

13. Mothilal, R.K., Sharma, A., Tan, C.: Explaining machine learning classifiers through diverse counterfactual explanations. In: Proceedings of the 2020 conference on fairness, accountability, and transparency. pp. 607–617 (2020)
14. Pearl, J.: Causality. Cambridge university press (2009)
15. Pearl, J., Glymour, M., Jewell, N.P.: Causal inference in statistics: A primer. John Wiley & Sons (2016)
16. Peters, J., Janzing, D., Schölkopf, B.: Elements of causal inference: foundations and learning algorithms. The MIT Press (2017)
17. Pocheville, A., Griffiths, P., Stotz, K.: Comparing causes – an information-theoretic approach to specificity, proportionality and stability. 15th Congress of Logic, Methodology, and Philosophy of Science (08 2015)
18. Ribeiro, M.T., Singh, S., Guestrin, C.: ” why should i trust you?” explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining. pp. 1135–1144 (2016)
19. Schölkopf, B., Janzing, D., Peters, J., Sgouritsa, E., Zhang, K., Mooij, J.: On causal and anticausal learning. arXiv preprint arXiv:1206.6471 (2012)
20. Vergara, J.R., Estévez, P.A.: A review of feature selection methods based on mutual information. Neural computing and applications **24**, 175–186 (2014)
21. Wachter, S., Mittelstadt, B., Russell, C.: Counterfactual explanations without opening the black box: Automated decisions and the gdpr. Harv. JL & Tech. **31**, 841 (2017)
22. Woodward, J.: Causation in biology: stability, specificity, and the choice of levels of explanation. Biology & Philosophy **25**(3), 287–318 (2010)
23. Zednik, C., Boelsen, H.: Scientific exploration and explainable artificial intelligence. Minds and Machines **32**(1), 219–239 (2022)
24. Zhao, Q., Hastie, T.: Causal interpretations of black-box models. Journal of Business & Economic Statistics (2019)
25. Zhou, H., Wang, X., Zhu, R.: Feature selection based on mutual information with correlation coefficient. Applied Intelligence pp. 1–18 (2022)