
Quantifying intrinsic causal contributions via structure preserving interventions

Dominik Janzing¹

Patrick Blöbaum¹

Atalanti A. Mastakouri¹

Philipp M. Faller^{1,2}

Lenon Minorics¹

Kailash Budhathoki¹

1 Amazon Research Tübingen, Germany

2 Karlsruhe Institute of Technology, Germany

Abstract

We propose a notion of causal influence that describes the ‘intrinsic’ part of the contribution of a node on a target node in a DAG. By recursively writing each node as a function of the upstream noise terms, we separate the intrinsic information added by each node from the one obtained from its ancestors. To interpret the intrinsic information as a *causal* contribution, we consider ‘structure-preserving interventions’ that randomize each node in a way that mimics the usual dependence on the parents and does not perturb the observed joint distribution. To get a measure that is invariant with respect to relabelling nodes we use Shapley based symmetrization and show that it reduces in the linear case to simple ANOVA after resolving the target node into noise variables. We describe our contribution analysis for variance and entropy, but contributions for other target metrics can be defined analogously. The code is available in the package `gcm` of the open source library `DoWhy`.

1 INTRODUCTION

Quantification of causal influence not only plays a role in expert’s research on scientific problems, but also in highly controversial public discussions. For instance, the question to what extent environmental factors versus genetic disposition influence human intelligence, is an ongoing debate (Krapohl et al., 2014). Given the relevance of these questions, there is surprisingly little clarity about how to define strength of influence in the first place, see e.g. Rose (2006). More recent discussions on feature relevance quantification

in explainable artificial intelligence have raised the problem of quantification of influence from a different perspective (Datta et al., 2016; Lundberg and Lee, 2017; Frye et al., 2020; Janzing et al., 2020; Wang et al., 2021; Jung et al., 2022). The quantification of causal influence proposed here is based on the intuition of measuring to what extent several factors influencing a target variable of interest ‘explain’ the *variation* (aka uncertainty) of the latter. Attributing *uncertainty* of the target to the influencing factors thus follows the spirit of *Analysis of Variance* (ANOVA) (Northcott, 2008).

Quantifying the causal influence has already resulted in a broad variety of proposals. A tentative taxonomy could classify measures according to the following three aspects: **1/** What type of intervention is used: Ay and Polani (2008); Datta et al. (2016); Lundberg and Lee (2017); Frye et al. (2020); Janzing et al. (2020); Heskes et al. (2020); Jung et al. (2022) use do-interventions (also called point-interventions) on *nodes* of a causal directed acyclic graph (DAG) in the sense of Pearl (2000), while Janzing et al. (2013); Schamberger et al. (2020) use interventions on the *edges*, and our paper uses *conditional* interventions.

2/ Which ones of the context variables are adjusted (and to which values) when the intervention is performed: Here and in most recent approaches one averages over all possible choices of adjustment sets based on Shapley values (but most approaches adjust *values* of nodes, while we adjust their *mechanisms*).

3/ Which target metric is used to assess the impact of those interventions: Here, we use *uncertainty* in the sense of entropy or variance, while Datta et al. (2016); Lundberg and Lee (2017); Frye et al. (2020); Janzing et al. (2020); Wang et al. (2021) use shift of the *target value* itself or its *expectation* for an attribution referring to individual instances at hand (as opposed to average contributions over populations). As a further example, Janzing et al. (2013); Schamberger et al. (2020) measure distribution change in terms of *relative entropy*.

Variations of these three different aspects can mostly be constructed independently in a modular manner. It is there-

fore pointless to blindly compare measures of influence that differ by more than one aspect. To explain, for instance, why our paper uses *conditional* interventions instead of *point* interventions, we will compare it to hypothetical modifications of existing measures that would also be based on variance or entropy of a target node, rather than the target metric used in the respective literature.¹

The crucial aspect of our proposal is what we consider the so-called ‘intrinsic’ contribution. We explain this intuition via the following toy example which is paradigmatic for a large class of scheduling processes. Consider the schedule of three trains A, B, C , where a delay in train A causes a delay of train B , which, in turn, causes a delay in train C . If we ask for the ‘intrinsic contribution’ of B to the delay of C , we are not asking for the hypothetical reduction of delay of C if B had arrived on time. Instead, we compare the delay of C to the scenario where B does not add any delay in addition to the one that it inherited from A . In other words, the intrinsic contribution of B on the delay of C is obtained via a comparison with a scenario in which the dependences to the causal ancestors remain intact. Note that this distinction between ‘intrinsic’ versus ‘inherited’ part makes sense regardless of whether we are interested in the delay for a specific instance, the mean, or the variation of the delay of C in the statistical population (although this paper focuses on the latter).

Probabilistic versus structural causal models While Causal Bayesian Networks (Spirtes et al., 1993; Pearl, 2000) model causal relations via a DAG with random variables as nodes, a more ‘fine-grained’ causal model is given by structural causal models:

Definition 1 (Structural Causal Model (SCM)). *An SCM corresponding to a DAG G with observed variables X_1, \dots, X_n as nodes is given by*

- (i) *exogenous noise variables N_1, \dots, N_n being jointly statistically independent*
- (ii) *functions f_j that express each X_j deterministically in terms of its parents and the noise term, that is,*

$$X_j = f_j(\mathbf{PA}_j, N_j). \quad (1)$$

Moreover, (1) entails the counterfactual statement that, for any particular observation (x_1, \dots, x_n) , setting \mathbf{PA}_j to \mathbf{pa}'_j instead of \mathbf{pa}_j would have changed x_j to $x'_j = f_j(\mathbf{pa}'_j, n_j)$ (where n_j denotes the value attained by N_j for that particular statistical instance).

Throughout the paper we leave it open whether or not the noise variables N_j are observed or not. The idea is that X_j are variables whose values are more immediately accessible, while the N_j or only its distribution $P(N_j)$ may

¹We have seen that verbal discussions on causal attribution often blur the importance of *all three* aspects, e.g. by claims like ‘we measure causal impact via Shapley values’.

be inferred from X_1, \dots, X_n . Here and henceforth, we denote probability distributions by the capital letter P and the corresponding densities or probability mass functions by p (e.g. $P(X_1)$ versus $p(x_1)$).

The existence of an SCM implies that $P(X_1, \dots, X_n)$ satisfies the Markov condition with respect to G (Pearl, 2000). On the other hand, every joint distribution that is Markovian relative to G can be generated by an SCM, but this construction is not unique. This is because knowing the causal DAG and the joint distribution alone does not determine all counterfactual causal statements (‘rung 3 in the ladder of causation’ according to Pearl and Mackenzie (2018)).²

Section 2 defines Intrinsic Causal Contribution (ICC) using the resolution into noise terms and quantifies the part of the uncertainty contributed by each noise variable. Section 3 explains that ICC has a causal interpretation in terms of interventions on the variables X_j although it requires adjustments of the noise at first glance. Then we discuss some properties of ICC in Section 4. Section 5 compares ICC to related information based approaches in the literature. Finally, Section 6 shows experiments with real data.

We will see that the quantification of the influence of a variable on a target crucially depends on the type of interventions considered. Therefore, we argue that different notions of causal influence coexist for good reasons, since they formalize different ideas on what causal influence is about. We emphasize that this paper is mostly theoretical and its main contribution is conceptual, rather than containing new mathematical insights. Further, also questions on statistical robustness and scalability for complex causal DAGs go beyond its scope. Constructing good proxies and adding assumptions that enable good estimates is left to future work, here we want to focus on providing an understanding of ICC to let readers judge themselves for which scenarios it is conceptually the right measure to ask for.

2 DEFINING ICC

2.1 Resolving variables into noises

To discuss the idea of ‘intrinsic’ within the language of SCMs, consider first the following metaphoric example. An empty donation box which starts at person A , is handed over to person B and finally arrives at C . Each person adds the donation N_A, N_B, N_C , respectively. If X_j denotes the amount of money after the box passed person j , we obtain the simple SCM:

$$X_A = N_A \quad (2)$$

$$X_B = X_A + N_B \quad (3)$$

$$X_C = X_B + N_C, \quad (4)$$

²See also Peters et al. (2017), Section 3.4, for an explicit description of the ambiguity.

where we consider N_j the exogenous ‘noise’ variables (which we know in this case), and the causal DAG reads

$$X_A \rightarrow X_B \rightarrow X_C. \quad (5)$$

Obviously, we would consider N_B the *contribution* of B to the final amount X_C . Hence, the *noise terms* carry the information about what has been added at the respective node X_j , not the values of the nodes X_j directly. Clearly, N_j captures what we would describe as one’s *contribution* in our everyday language, not X_j . The example is metaphoric in the sense that it deals with contribution to a *fixed amount*, while we will later focus on contributions to the *uncertainty*. We only want to motivate why our notion of contribution focuses on what is added by the *noise terms*. For this example one may argue that one could equally well define an augmented DAG containing N_X explicitly as nodes in order to avoid focusing on *exogenous* variables. We argue, however, that the values T_X are more directly observable, while the ‘contributions’ N_X are reconstructed from the former, using the above SCM. Likewise, the total delay of a train is *more directly observable* than the amount contributed by the respective train itself, and we argue that similar statements apply to a large class of problems of causal contribution analysis (see also Section 6), which justifies our view of looking at contributions as something to be *reconstructed* from the SCM.

To quantify the contribution of each ancestor to some target node X_n of interest (which is assumed to have no descendants without loss of generality), we recursively insert structural equations (1) into each other and write X_n entirely in terms of the unobserved noise variables:

$$X_n = F_n(N_1, \dots, N_n). \quad (6)$$

Now we can think of X_n as being the effect of the *independent causes*³ N_1, \dots, N_n . Note that we do not introduce further assumptions regarding F_n , i.e. we allow arbitrary SCMs and do not limit it to a specific form such as additive noise models. While this resolution coincides with the root cause analysis of outliers by Budhathoki et al. (2022) and also with backtracking counterfactuals by Von Kügelgen et al. (2023), ICC does not require the assumption of *invertible* SCMs, that is, we do not need to be able to reconstruct the *values* of the noise from observed variables for single observations, we only need to know their distribution. This is because ICC is not about contributions to single events.

2.2 Quantifying conditional reduction of uncertainty

Next we quantify the reduction of uncertainty in X_n caused by a hypothetical adjustment of N_j (which coincides with usual conditioning due to exogeneity).

³Writing each node in terms of independent noise variables is also possible for some cyclic causal models, to which we could easily generalize ICC (Bongers et al., 2021).

Definition 2 (intrinsic causal contribution (ICC)). *Let F_n as in (6) express X_n in terms of all noise variables $(N_1, \dots, N_n) =: \mathbf{N}$. Then the ICC of node X_j (with $j = 1, \dots, n$), given some additional adjustment set $T \subset \{1, \dots, n\}$, is defined by*

$$ICC_\psi(X_j \rightarrow X_n|T) := \psi(X_n|\mathbf{N}_T) - \psi(X_n|N_j, \mathbf{N}_T), \quad (7)$$

where ψ can be any kind of conditional uncertainty measure satisfying **monotonicity** $\psi(X_n|\mathbf{N}_T) - \psi(X_n|N_j, \mathbf{N}_T) \geq 0$ and **calibration** $\psi(X_n|\mathbf{N}) = 0$. Here, $\psi(\cdot|\mathbf{N}_T)$ denotes conditioning on all noise variables N_j with $j \in T$.

Monotonicity is not strictly needed. But in real-world applications positive contributions are easier to interpret and visualize (e.g., in pie charts), therefore we prefer a contribution that is non-negative. Although the does not refer to interventions, Section 3 explains that it is nevertheless causal because it can be recast in terms of (a non-standard type of) interventions, but we keep Definition 2 because of its simplicity.

Note that we decided to write $ICC_\psi(\cdot|T)$ instead of $ICC_\psi(\cdot|\mathbf{X}_T)$ to emphasize that we do not condition on the random variables \mathbf{X}_T , but the corresponding noise terms. Possible choices of ψ are:

Example 1 (conditional Shannon entropy). *With $\psi(X_n|\mathbf{N}_T) := H(X_n|\mathbf{N}_T)$ we obtain conditional mutual information (Cover and Thomas, 1991)*

$$\begin{aligned} ICC_H(X_j \rightarrow X_n|T) &= H(X_n|\mathbf{N}_T) - H(X_n|N_j, \mathbf{N}_T) \\ &= I(N_j : X_n|\mathbf{N}_T). \end{aligned}$$

Non-negativity of the r.h.s. implies monotonicity of ψ . Calibration is satisfied for discrete X_n , while continuous entropy tends to minus infinity for a distribution approaching point measure.

We only use entropy when the target X_n is discrete. Note that it does not matter whether the N_j are continuous or discrete, in both cases we can measure the conditional uncertainty of X_n in terms of Shannon entropy.

Although information theoretic quantification of influence comes with the advantage of being applicable to variables with arbitrary finite range, e.g., categorical variables, quantification in terms of variance is more intuitive and often easier to estimate from finite data:

Example 2 (expected conditional variance). *For $\psi(X_n|\mathbf{N}_T) := \mathbb{E}[\text{Var}(X_n|\mathbf{N}_T)]$ we obtain*

$$\begin{aligned} ICC_{\text{Var}}(X_j \rightarrow X_n|T) &= \\ &= \mathbb{E}[\text{Var}(X_n|\mathbf{N}_T)] - \mathbb{E}[\text{Var}(X_n|N_j, \mathbf{N}_T)]. \end{aligned}$$

The monotonicity of expected conditional variance follows from the law of total variance, $\text{Var}(Y) = \mathbb{E}[\text{Var}(Y |$

X) + $\text{Var}(\mathbb{E}[Y | X])$. Thus, $\mathbb{E}[\text{Var}(Y)] \geq \mathbb{E}[\text{Var}(Y | X)]$ (where $\mathbb{E}[\text{Var}(Y)] = \text{Var}(Y)$). Further, $\mathbb{E}[\text{Var}(Y)] \geq \mathbb{E}[\text{Var}(Y | X)]$, which entails also the conditional version $\mathbb{E}[\text{Var}(Y | \mathbf{W})] \geq \mathbb{E}[\text{Var}(Y | \mathbf{W}, X)]$ for any random vector \mathbf{W} . Note that Analysis of Variance (ANOVA) relies on the same idea for the special case of a linear model where variances are just additive for independent factors (Lewontin, 1974; Northcott, 2008). Variance based sensitivity analysis (Sobol, 2001) allows for non-linear models, but does not necessarily measure *causal* influence since it considers reduction of variance by conditioning on observed nodes regardless of whether the statistical relation to the target is causal or confounded. In Section 3 we explain a rephrasing of ICC in terms of interventions on observed nodes, which also admits a generalization for semi-Markovian models.

2.3 Symmetrization via Shapley values

Unfortunately, the contribution of each node X_j in (7) depends on the index subset $T \subset \{1, \dots, n\}$ given as context. For any ordering π in the symmetric group S_n we can certainly define ‘Plain ICC’ by decomposing $\psi(X_n)$ into contributions $ICC_\psi(X_{\pi(j)} \rightarrow X_n | T_\pi^j)$, where T_π^j denotes the set of indices that occur before j in the ordering π . Unfortunately, the dependence on π introduces an undesired arbitrariness. Feature relevance quantification in explainable AI (Datta et al., 2016; Lundberg and Lee, 2017; Frye et al., 2020; Janzing et al., 2020) addresses similar problems via Shapley values from cooperative game theory (Shapley, 1953), which implicitly amounts to symmetrizing over all orderings π . For the definition of Shapley values see Section 1 in the appendix.

Definition 3 (Shapley ICC). *Let the ‘worth of a coalition of noise terms \mathbf{N}_T ’ for $T \subset \{1, \dots, n\}$ be given by $\nu(T) := -\psi(X_n | \mathbf{N}_T)$. Then the (Shapley based) ICC of each node X_j to the uncertainty of X_n reads:*

$$\begin{aligned} ICC_\psi^{Sh}(X_j \rightarrow X_n) & \\ := \sum_{T \subseteq U \setminus \{j\}} \frac{1}{n \binom{n-1}{|T|}} [\nu(T \cup \{j\}) - \nu(T)] & \quad (8) \\ = \sum_{T \subseteq U \setminus \{j\}} \frac{1}{n \binom{n-1}{|T|}} ICC_\psi(X_j \rightarrow X_n | T), & \end{aligned}$$

where $U := \{1, \dots, n\}$

Due to general properties of Shapley values (see Section A), Shapley based ICC values sum up to the uncertainty of X_n :

$$\begin{aligned} \sum_{j=1}^n ICC_\psi^{Sh}(X_j \rightarrow X_n) &= \nu(\{1, \dots, n\}) - \nu(\emptyset) \\ &= \psi(X_n) - \psi(X_n | \mathbf{N}) = \psi(X_n). \end{aligned} \quad (9)$$

Using an alternative equivalent definition of Shapley values, which enables simple Monte Carlo estimators

(Mitchell et al., 2022), we can also average over S_n instead:

$$ICC_\psi^{Sh} = \frac{1}{n!} \sum_{\pi \in S_n} \psi(X_n | \mathbf{N}_{T_\pi^j}) - \psi(X_n | \mathbf{N}_{\{j\} \cup T_\pi^j}), \quad (10)$$

which we will also use later.

Example 3 (linear SCMs). *Given the linear SCM $\mathbf{X} = A\mathbf{X} + \mathbf{N}$ where $\mathbf{X} = (X_1, \dots, X_n)$ and A is a lower triangular matrix. Then, $\mathbf{X} = (I - A)^{-1}\mathbf{N}$ and thus X_n is the linear combination of n independent noise terms $X_n = \sum_{j=1}^n \beta_{nj} N_j$, where β_{nj} is the last row of $(I - A)^{-1}$. The contribution of each N_j to the variance of X_n thus reads $\beta_{nj}^2 \cdot \text{Var}(N_j)$, independently of the conditioning set S . Hence, we simply obtain*

$$ICC_{\text{Var}}^{Sh}(X_j \rightarrow X_n) = \beta_{nj}^2 \cdot \text{Var}(N_j),$$

which amounts to usual ANOVA (without the computational load entailed by Shapley values) after resolving X_n into the noise terms.

3 CAUSAL INTERPRETATION

3.1 Structure-preserving interventions

We now argue that ICC can be rephrased using interventions on the nodes X_1, \dots, X_n : Going through the nodes in any topological ordering of the DAG, *replace each X_j with $f_j(\mathbf{pa}'_j, n'_j)$ where n'_j is a fixed value of an independent copy N'_j of N_j , and \mathbf{pa}'_j denotes the values of the parents obtained from the upstream interventions. In contrast, standard interventions $do(X_j = x_j)$ replace structural assignment (1) with the trivial structural assignment $X_j := x_j$. Although our intervention may be interpreted as setting the noise variable N_j to n'_j (which is unfeasible if we think of the exogenous noise of something that is not under our control, or even worse, not even observable), we suggest to interpret it as an intervention on X_j instead that depends on the values of the parents. Given assignment (1), we can intervene on X_j without perturbing the joint distribution of all observable nodes X_1, \dots, X_n as follows: after observing that \mathbf{PA}_j attained the values \mathbf{pa}_j , we set X_j to the value $x'_j := f_j(\mathbf{pa}_j, n'_j)$ where n'_j is randomly drawn from $P(N_j)$. Any statistical dependence between X_i and our i.i.d. copy N'_j of N_j indicates causal impact of X_j on X_i . This is because N'_j is randomized and thus interventional and observational probabilities coincide, i.e. $P(\cdot | do(N'_j = n'_j)) = P(\cdot | N'_j = n'_j)$. Generalized interventions that replace the structural assignment (1) with a different one, have been studied in the literature earlier (Eberhardt and Scheines, 2007; Korb et al., 2004; Tian and Pearl, 2001; Daniel and Murphy, 2007; Markowitz et al., 2005; Correa and Bareinboim, 2020), but we keep (1) and replace only N_j with an *observable and accessible* copy N'_j .*

3.2 Adjusting mechanisms at a node, not their values

The fact that we do not measure the reduction of uncertainty entailed by adjusting the *value* x_j of the observed node X_j , but instead the reduction caused by adjusting the *noise*, can be nicely interpreted in terms of the response function formulation of SCMs (Greenland and Robins, 1986). Then each value n_j of N_j corresponds to a function $f_j^{n_j} : \mathbf{pa}_j \mapsto x_j$ and the SCM is merely a probability distribution over the set of functions (‘deterministic mechanisms’) $f_j^{n_j}$. In this sense, ICC measures the extent to which the variability of the respective mechanisms contribute to the variability of the target variable. While phrasing ICC in terms of adjusting some of the N_j resulted in a simple formal definition (to avoid distributions over functions), the more principled way of thinking of ICC is given by adjustments of the mechanism $f_j^{n_j}$. The drawing of values of the iid copy N_j' above now translates into iterating over randomized adjustments of the mechanisms.

Note that the decision of whether we are seeking for a contribution measure that adjusts *values* or one that adjusts *mechanisms* is related to philosophical questions of blaming and the normative aspect of ‘actual causation’ (Hitchcock and Knobe, 2009; Halpern and Hitchcock, 2013). If we, for instance, blame a train for its *entire* delay, we implicitly consider the option of departing in time regardless of delays of ‘ancestor’ trains (corresponding to adjusting the value), while intrinsic contribution accepts the dependence as mandatory by adjusting the mechanism.

3.3 Generalization of ICC to the confounded case

Let us now assume a causal model with common causes of nodes X_i, X_j which are not blocked by any set of nodes in $\{X_1, \dots, X_n\}$. The easiest way to generalize ICC to this case is given by allowing for *dependent* noise terms. We then define the worth of the coalition for any $T \subset \{1, \dots, n\}$ by

$$\begin{aligned} \nu(T) &:= -\psi(X_n | do(\mathbf{N}_T)) \\ &:= \begin{cases} \sum_{\mathbf{n}_T} H(X_n | do(\mathbf{N}_T = \mathbf{n}_T)) p(\mathbf{n}_T) \\ \sum_{\mathbf{n}_T} \text{Var}(X_n | do(\mathbf{N}_T = \mathbf{n}_T)) p(\mathbf{n}_T) \end{cases}, \end{aligned}$$

where the meaning of $do(N_T)$ is defined by the respective right hand sides. Note that in the unconfounded scenario with independent noise variables N_j we did not explicitly introduce interventional probabilities since they coincided with conditional probabilities anyway. Assuming that the dependence of noise variables relies on common causes, the interventional probabilities are now given by the back-door formula (Pearl, 2000)

$$P(X_n | do(\mathbf{N}_T = \mathbf{n}_T)) = \sum_{\mathbf{n}_{\bar{T}}} P(X_n | \mathbf{n}_{\bar{T}}, \mathbf{n}_T) p(\mathbf{n}_{\bar{T}}), \quad (11)$$

where \bar{T} denotes the complement of T . Assuming that the noises are only connected by a common cause does

not entail any loss of generality: assume we had a path from one noise to another ($N_i \rightarrow N_j$), then we can introduce a mediator N_i' between N_i and X_i , and N_i becomes a common cause of N_i' and N_j . Since the only relevance of the noise consists in its effect on the response function $\mathbf{pa}_j \mapsto f_j(\mathbf{pa}_j, n_i)$, the modified model responds in the same way to structure preserving interventions.

For ease of notation, we have phrased $\nu(T)$ in terms of interventions on the noise variables N_j although we have emphasized in Section 3.2 in the main paper that ICC only requires interventions on the variables X_j , provided that we consider structure preserving interventions instead of usual do-interventions. We can do the same here, but the interventions are now controlled by a simulated noise vector \mathbf{N}'_T which is a copy of \mathbf{N}_T , that is, has the same *joint* distribution (rather than having separate copies N_j' of each N_j). Note that we have $\nu(\{1, \dots, n\}) = -\psi(X_n | do(\mathbf{N} = \mathbf{n})) = 0$ and $\nu(\emptyset) = -\psi(X_n)$. Therefore, Eq. (9) in the main paper holds for the confounded case as well, which is crucial for the interpretation of attributing *the total uncertainty* of X_n to its ancestors. Note, however, that ICC can become now negative because dropping a *correlated* noise variable can increase the variance of the target.

4 SOME PROPERTIES OF ICC

Here we describe properties that help getting an intuition on the behavior of ICC, particularly with respect to extending or marginalizing models. These properties will later help to understand the difference to other measures of causal influence in the literature and help the reader decide whether ICC is appropriate for the problem at hand.

4.1 Inserting nodes being perfect wires

Assume we are given the causal DAG $X \rightarrow Y$ with the structural equations $X := N_X$ and $Y := f_Y(X, N_Y)$. Then, application of Definition 2, eq. (7) and (8) yields

$$\begin{aligned} ICC_{\psi}^{Sh}(X \rightarrow Y) &= \frac{1}{2} [\psi(Y) - \psi(Y | N_X)] \\ &+ \psi(Y | N_Y) - \psi(Y | N_X, N_Y)]. \end{aligned} \quad (12)$$

Let us now insert an intermediate node \tilde{X} that is just an exact copy of X , that is, we define the modified SCM

$$X = N_X \quad (13)$$

$$\tilde{X} = X \quad (14)$$

$$Y = f_Y(\tilde{X}, N_Y). \quad (15)$$

The corresponding DAG reads $X \rightarrow \tilde{X} \rightarrow Y$. From a physicists perspective, such a refinement of the description should always be possible because any causal influence propagates via a signal that can be inspected right after it leaves the source. Lemma 1 in Section A shows that (13)

to (15) entail the same value for $ICC_{\psi}^{Sh}(X \rightarrow Y)$ as (12) because the constant ‘dummy’ noise $N_{\tilde{X}}$ corresponding to \tilde{X} is irrelevant for the contribution of the other nodes.

4.2 Marginalizing over intermediate nodes or grandparents

While we inserted a deterministic node in Subsection 4.1 we now marginalize over an intermediate node that depends non-deterministically on its cause.⁴ Let us again consider the chain

$$X \rightarrow Y \rightarrow Z, \quad (16)$$

with the structural equations

$$X := N_X \quad (17)$$

$$Y := f_Y(X, N_Y) \quad (18)$$

$$Z := f_Z(Y, N_Z). \quad (19)$$

Recall that, in case of entropy, $ICC_H^{Sh}(X \rightarrow Z)$ contains the terms $I(N_X : Z)$, $I(N_X : Z | N_Y)$, $I(N_X : Z | N_Z)$, $I(N_X : Z | N_Y, N_Z)$.

Marginalizing over Y yields the causal DAG $X \rightarrow Z$ with the structural equations

$$X := N_X \quad (20)$$

$$Z := \tilde{f}_Z(X, \tilde{N}_Z), \quad (21)$$

where $\tilde{N}_Z := (N_Y, N_Z)$ and

$$\tilde{f}_Z(X, \tilde{N}_Z) = f_Z(f_Y(X, N_Y), N_Z).$$

For the reduced structure, $ICC_H^{Sh}(X \rightarrow Z)$ contains only terms of the form $I(N_X : Z)$ and

$$I(N_X : Z | \tilde{N}_Z) = I(N_X : Z | N_Y, N_Z),$$

while the terms $I(N_X : Z | N_Y)$, $I(N_X : Z | N_Z)$ do not occur any longer. Hence, the Shapley based ICC is not invariant with respect to the marginalization. The reason is that Shapley symmetrization averages the relevance of N_X over all possible combinations of background conditions. Reducing the possible combinations by ignoring nodes can result in different values for ICC_{ψ}^{Sh} . One may consider this as a caveat: we would not like the contribution of train A for the delay of train C depend on whether B is explicitly taken into account or not. However, the reason is a caveat that other Shapley based quantification of feature relevance show as well: Assume we are given $Y = f(X_1, X_2, X_3)$ and compute relevance of feature X_1 for Y according to Lundberg and Lee (2017). One can then verify that the relevance of X_1 changes when we merge X_2, X_3 to a vector valued feature X'_2 and write $Y = f'(X_1, X'_2)$. This

⁴Note that consistency of causal structures under various coarse-grainings is an interesting topic in a more general context too (Rubenstein et al., 2017).

is exactly what happens in our case: marginalizing over Y merges the noise variables N_Y and N_Z to one noise variable for Z .

One also checks easily that the contribution of Y changes when marginalizing over X . In the limiting case where Y is just a copy of X we even obtain $ICC_{\psi}^{Sh}(Y \rightarrow Z) = 0$ for DAG (16) while the DAG $Y \rightarrow Z$ is blind for the fact that Y has ‘inherited’ all its information from its grandparent. Following our intuition of contribution, this change of ICC is required because we marginalization blurs that part of the uncertainty of Y is just propagated from its parent.

4.3 Continuity w.r.t. removing weak edges

To define a family of models for which an edge gets arbitrarily weak we consider variables attaining finitely many values and again use the response function formulation of SCMs, where (1) turns into a probability distribution on the set \mathcal{F}_j of functions from the set of possible states of \mathbf{PA}_j to the possible states of X_j . For each parent PA_j^i , there is a subset $\mathcal{F}_j^i \subset \mathcal{F}_j$ of functions that are constant in PA_j^i . ‘Fading away’ of an edge from PA_j^i can then be described by continuously turning the probability distribution into one with support \mathcal{F}_j^i . For discrete variables, entropy and variance are continuous on finite sets, thus ICC will not show any discontinuity when the probability of the complement of \mathcal{F}_j^i reaches zero.

4.4 Dependence on the SCM

ICC may differ for different SCMs (with the same DAG) describing the same interventional probabilities for interventions $do(X_i = x_i)$ (rung 3 versus rung 2 causal statements). As an example, let us take a look at $X \rightarrow Y$ with binary variables X, Y . First consider the structural equations

$$X := N_X, \quad Y := X \oplus N_Y \quad (22)$$

where \oplus denotes the XOR operator. With $\psi = H$ and assuming ‘unbiased coins’ $P(N_X = 1) = P(N_Y = 1) = 1/2$ we then obtain

$$\begin{aligned} ICC_H^{Sh}(X \rightarrow Y) &= \frac{1}{2}(I(N_X : Y) + I(N_X : Y | N_Y)) \\ &= \frac{1}{2}(0 + 1) = 1/2. \end{aligned}$$

The same joint distribution $P(X, Y)$ can also be generated by

$$X := N_X, \quad Y := N_Y \quad (23)$$

for which we obtain⁵ $ICC_H^{Sh}(X \rightarrow Y) = 0$. Hence, ICC also captures counterfactual influence which is invisible by usual interventions on the node itself. This can be desirable, for instance, in the following scenario: Let X be a random variable denoting a text in binary encoding. Let Y be its encrypted version generated by bitwise XOR with a randomly generated secret key N . If we have no access to N , we cannot detect any statistical dependence between X and Y . However, we would not argue that the binary encoding of X *did not contribute* to the encryption Y just because the statistical dependence gets only visible after knowing N . Even if we are not able to decrypt Y because N is unknown to us, the mere knowledge that Y *could* be decrypted after knowing N suffices to acknowledge the contribution of X to Y .⁶ As far as we can see, all known attempts based on usual point interventions fail to formalize our intuitive notion of ‘intrinsic’, despite being reasonable concepts in their own right. This will be explained in the following section.

5 PREVIOUS WORK ON QUANTIFYING CAUSAL INFLUENCE

Since do-interventions are accepted as a crucial concept of causality, we want to check how they could potentially capture ‘intrinsic’ contribution without relying on rung 3 causal models. To this end we consider the contribution of X_2 on X_3 in the causal chain

$$X_1 \rightarrow X_2 \rightarrow X_3, \quad (24)$$

and emphasize that we want it to be zero when X_2 is a copy of X_1 . We will see that most approaches fail in this regard, except for one case which we reject because it violated continuity.

Information Flow: This measure, introduced by Ay and Polani (2008) and denoted by $I(X_2 \rightarrow X_3)$, measures the *mutual information* of X_3 and X_2 with respect to the joint distribution obtained when X_2 is subjected to randomized adjustments according to $P(X_2)$ (see Section C for the formal definition). This concept is certainly causal, but does not separate the information generated at X_2 from the one inherited from X_1 . One option to achieve this separation would be to randomize X_2 according to the conditional distribution $P(X_2|do(X_1 = x_1))$ instead of $P(X_2)$, which yields the *conditional* Information Flow $I(X_2 \rightarrow X_3|do(x_1))$ (Ay and Polani, 2008). Its average over $P(X_1)$

⁵Readers may find it disturbing that (22) generates an *unfaithful* distribution and (23) actually corresponds to a DAG without arrows. But we have chosen this example only to keep the solution simple, dependence on SCM also holds for more generic models.

⁶The idea that noise terms N_j in SCMs may be unobservable in one context, but not *in principle*, is also emphasized by Pearl (2000) to justify the scientific content of counterfactuals.

is denoted by $I(X_2 \rightarrow X_3|do(X_1))$. For the example where X_2 is just a deterministic copy of X_1 , we obtain $I(X_2 \rightarrow X_3|do(X_1)) = 0$, as desired, since adjusting X_1 sets X_2 to a constant. However, if $I(X_2 \rightarrow X_3|do(X_1))$ is a better candidate for ‘intrinsic’ contribution, generalization to arbitrary DAGs is unclear. The example suggests to measure the intrinsic contribution of any node X_i in a DAG with n nodes to the information on X_j by $I(X_i \rightarrow X_j|do(PA_i))$. Accordingly, in the DAG in Figure 1a, we would consider $I(X_2 \rightarrow X_3|do(X_1))$ the intrinsic contribution of X_2 on X_3 . In Figure 1b, X_1 is no longer a parent




Figure 1: Left: Causal DAG for which it is already non-trivial to define the strength of the influence of X_1 on X_3 – if one demands that this definition should also apply to the limiting case on the right (where the edge $X_1 \rightarrow X_2$ disappeared).

of X_2 and then we would consider choosing $I(X_2 \rightarrow X_3)$ instead. This raises a conceptual problem raised by switching from conditional to unconditional Information Flow: let the edge $X_1 \rightarrow X_2$ be arbitrarily weak, such that it disappears entirely, which entails a *discontinuous* change from $I(X_2 \rightarrow X_3|do(X_1))$ to $I(X_2 \rightarrow X_3)$. This is because the difference between $I(X_2 \rightarrow X_3|do(X_1))$ and $I(X_2 \rightarrow X_3)$ remains even when the edge disappears, which is easily seen for binary variables linked by a logical XOR gate: let $X_3 = X_1 \oplus X_2$ and X_1 be an unbiased coin. Without adjusting X_1 , adjusting X_2 has no average influence on X_3 , hence $I(X_2 \rightarrow X_3) = 0$. When adjusting X_1 , however, X_2 controls X_3 entirely and when X_2 is unbiased too, we thus obtain $I(X_2 \rightarrow X_3|do(X_1)) = 1$ Bit.

Causal Shapley Values and do-Shapley Values: Using the idea of these concepts (Heskes et al., 2020; Jung et al., 2022), but also modifying it to *uncertainty* as our target metric, we define the worth of a coalition T as $\psi(X_n|do(\mathbf{X}_T)) := \sum_{\mathbf{x}_T} \psi(X_n|do(\mathbf{X}_T = \mathbf{x}_T))p(\mathbf{x}_T)$. Then the contribution of X_2 on X_3 in the DAG (24) is given by a sum of two terms of the form $\psi(X_n|do(\mathbf{X}_T)) - \psi(X_n|do(X_2), do(\mathbf{X}_T))$, which is non-zero for $T = \emptyset$ in contrast to what we demand.

Asymmetric Shapley Values: Let us now modify (10) by averaging only over all possible *topological orderings* of the DAG. For those π we have

$$\psi(X_n|\mathbf{N}_{T_\pi^j}) = \psi(X_n|\mathbf{X}_{T_\pi^j}) = \psi(X_n|do(\mathbf{X}_{T_\pi^j})).$$

The first equality follows from $X_n \perp\!\!\!\perp \mathbf{X}_{T_\pi^j}|\mathbf{N}_{T_\pi^j}$ and because $\mathbf{X}_{T_\pi^j}$ is a function of $\mathbf{N}_{T_\pi^j}$. The second one follows

because conditioning on all ancestors blocks all backdoor paths. This way, we end up with a causal contribution measure that contains only observational conditionals and, more importantly, avoids rung 3 causal models. It turns out that this measure would be a straightforward modification of Asymmetric Shapley Values (Frye et al., 2020), with the only difference that we measure contribution to *uncertainty* in the population, while the former measures contribution of a specific value to the target value.

Since the causal chain (24) only admits one unique ordering, we obtain $\psi(X_3|do(X_2)) - \psi(X_3|do(X_2), do(X_1))$, which is zero when X_2 is a copy of X_1 , as desired. Unfortunately, Asymmetric Shapley Values can change discontinuously when edges ‘fade away’, because a DAG suddenly admits more topological orderings after removing an edge, which we wanted to avoid. For this reason, we believe that ‘fair’ intrinsic causal contribution should account for all orderings, and thus also account for terms that result from adjusting mechanisms (aka noise variables) *downstream* while randomizing those of the *ancestors* (which requires causal models of rung 3, see Subsection 4.4).

Nevertheless, ICC with averaging over topological orderings only may be a reasonable approach, for instance in causal models of time series without instantaneous effects. There, non-topological orderings would result in counterintuitive terms anyway which randomize past noise variables, while adjusting later ones.

Strength of edges and paths: We sketch proposals for quantifying the strength of *edges* or *paths* (Janzing et al., 2013; Wang et al., 2021) in Section B.

6 EXPERIMENTS

6.1 River flows

We have used a data set with 3,953 daily records of the river flows (in m^3/s) at 5 different stations in England at Henthorn, Hodder Place, Whalley Weir, New Jumbles Rock, and Samlesbury⁷ between 1979/05/01 and 2021/12/11. Samlesbury is considered the target since the other 4 stations are upstream, as seen on the map on Figure 2. New Jumbles Rock lies at a confluence point of the 3 rivers passing Henthorn, Hodder Place, and Whalley Weir. The water passing a certain station is certainly a mixture of some fraction of the amount observed at the next stations further upstream plus some amount contributed by streams and little rivers entering the river in between from the countryside around. Again, the intrinsic contribution is the *unrecorded* amount required to explain the water flows observed at the respective station, given the flows at the adjacent stations upstream. Therefore, choosing the DAG (Figure 2, mid-

dle) containing only the observed flows seems more natural than an *augmented* DAG containing the unrecorded influx as additional nodes, given that the latter values are reconstructed. Due to strong confounding effects (we have seen strong correlations of root nodes), estimating causal regression coefficients and noise by OLS would be heavily biased. Therefore we have inferred the SCM from common sense knowledge and set all regression coefficients to 1, assuming, for simplicity, that most of the water flow at each node will also reach downstream nodes. Hence, the noise (i.e., the hidden influx) is simply the difference to the sum over all parents. Afterwards, we have applied ICC for the confounded scenario, see Subsection 3.3, and Section D for more details.

Figure 2, right, shows the ICC values. The low values for New Jumbles Rock and Samlesbury show that the flow at the root nodes already explains most of the variation of the former. We have also applied variance based do-Shapley, as explained in Section 5, and obtained high contributions for New Jumbles Rock, see Section D for details.

Adjusting values versus adjusting mechanisms: The low value for New Jumbles Rock distinguishes ICC from causal influence measures that quantify contribution via the impact of a *do*-intervention on the respective node: adjusting the flow at New Jumbles Rock reduces the variance of Samlesbury significantly. Assume politicians of Samlesbury want to reduce the fluctuations of water flows observed there (to enable safe shipping or reduce flooding risk) by building *retention dams* spread over the *countryside around* the rivers. Then the intrinsic contribution of each node is the maximal reduction achievable by dams in the respective region. A different strategy for getting more constant flow through Samlesbury would be to build dams *inside the river*, directly controlling the flow. Assume all the 5 location were candidates for possible places. Trivially, having a dam in Samlesbury has the strongest impact, but there may be other aspects for preferring one of the other 4 locations. In this case, a measure of causal influence that quantifies the variance reduction by usual *do*-interventions would be the right one to assess the impact of these dams. This example shows, again, that different measures of causal influence are appropriate for different goals.

6.2 Fuel consumption of engines

Here we consider a data set with non-linearities, as verified by scatter plots of bivariate dependences between some of the variables. Although we have still fitted an additive noise model as a convenient approximation, non-linearity of structural equations entail that then noise variables influence the target quantity in a non-additive way.

The dataset AUTO MPG in the UCI Machine Learning Repository (Bache and Lichman, 2013), contributed by R.

⁷<https://environment.data.gov.uk/hydrology/explore>




Figure 2: Left: location of the 5 stations at which water flows are recorded. Middle: Causal model for the flows with Weather as latent confounder. Right: ICC with bootstrap confidence bounds for each of the 4 upstream stations and the target station (Samlesbury) itself.




Figure 3: Causal DAG of the AUTO MPG dataset for factors influencing fuel consumption with cyl: Cylinders, dis: displacement, hp: horsepower, wgt: weight, mpg: miles per gallon

Quinlan contain 7 features that are relevant for the prediction of miles per gallon (mpg) of a car engine. The causal DAG given in Wang and Mueller (2017), after removing variables that have no influence on mpg, is shown in figure 3. We have obtained the following values for ICC_{Var}^{Sh} after normalizing with the total variance: cyl 64%, dis 17%, hp < 1%, wgt 3%, mpg 15%. It turns out that the number of cylinders already explains a large fraction of the fuel consumption in the sense that the number commonly entails design decisions that involve parameters like the recorded ones. The intermediate nodes like dis, hp, and wgt mostly inherit uncertainty from their parents, but roughly 1/7 of the variance of mpg remains unexplained by all of the above factors.

7 DISCUSSION

Computational complexity of ICC. The exact computation of the Shapley values is an expensive job which thus requires approximations. The package SHAP for explainable AI (Lundberg and Lee, 2017) uses a probabilistic approach with random subsets (for a different kind of attribution though). We meanwhile work with a method that samples random permutations. This way, we simply replace the population average with an empirical mean, then the Hoeffding inequality yields error bounds that, a priori, do not depend on the number n of nodes (but on bounds on information differences of subsets). The choice of f_j in (1) and ψ also influences the complexity of ICC. Example 3 shows a choice that simplifies computation. In particular, in linear models, variance based ICC does not involve costly averaging over subsets S . If we were to compute

the reduction of variance entailed by *do-interventions* on the X_j instead, even the linear case would involve costly computation of Shapley values. To see this, consider the chain $X_1 \rightarrow X_2 \rightarrow \dots \rightarrow X_{n-1} \rightarrow X_n$, with the structural equation $X_j = \alpha_{j-1}X_{j-1} + N_j$. The contribution of X_{n-1} on X_n is computed via terms $\text{Var}(X_n|do(\mathbf{X}_T)) - \text{Var}(X_n|do(X_{n-1}), do(\mathbf{X}_T))$, and certainly the variance reduction by adjusting X_{n-1} is smaller when upstream nodes that are close by are already adjusted (such that the variance of X_{n-1} is also low without adjusting it). Thus, ICC raises less computational challenges than other Shapley value based measures for the linear case, and similar ones for non-linear models.

Can we learn the SCM? Although the observed conditionals $P_{X_j|\mathbf{PA}_j}$ do not in general determine the SCM, the hardness of inferring $P_{X_j|\mathbf{PA}_j}$ from finite data often requires strong restriction of model classes such that the SCM follows ‘for free’ (Kano and Shimizu, 2003; Zhang and Hyvärinen, 2009; Peters et al., 2014; Mooij et al., 2016). For instance, additive noise model based inference (Mooij et al., 2016) infers $P_{X_j|\mathbf{PA}_j}$ by fitting the structural equation $X_j = \tilde{f}_j(\mathbf{PA}_j) + N_j$ with N_j independent of \mathbf{PA}_j and $\tilde{f}_j(\mathbf{pa}_j) := \mathbb{E}[X_j|\mathbf{pa}_j]$. Hence, the SCM and also N_j can then be entirely derived from observable entities due to $N_j = X_j - \mathbb{E}[X_j|\mathbf{PA}_j]$. For more SCM-learning methods see e.g. Storkey et al. (2006); Pawlowski et al. (2020).

Does ICC reduce to some known measure after augmenting the DAG? One may wonder whether ICC reduces to some known measure of causal influence in a DAG that contains all X_j and N_j . However, in this augmented DAG usually interventions on X_j and N_j both can have an impact on the target. This way, we then obtain separate contributions for X_j and for N_j ⁸, while ICC yields one contribution per X_j . In other words, no obvious modification of existing proposals seems to capture our notion of intrinsic by simply augmenting the DAG.

⁸Here we assume that we apply Causal Shapley Values or Shapley to all nodes in the augmented DAG, as usual. If we consider all variables N_j as independent causes of X_n , then they coincide with ICC on this DAG.

References

- N. Ay and D. Polani. Information flows in causal networks. *Advances in Complex Systems*, 11(1):17–41, 2008.
- K. Bache and M. Lichman. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- P. Blöbaum, P. Götz, K. Budhathoki, A. A. Mastakouri, and D. Janzing. Dowhy-gcm: An extension of dowhy for causal inference in graphical causal models. [arXiv:2206.06821](https://arxiv.org/abs/2206.06821), 2022.
- S. Bongers, P. Forré, J. Peters, and J. M. Mooij. Foundations of structural causal models with cycles and latent variables. *Annals of Statistics*, 49(5):2885–2915, 2021.
- K. Budhathoki, L. Minorics, P. Bloebaum, and D. Janzing. Causal structure-based root cause analysis of outliers. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 2357–2369. PMLR, 17–23 Jul 2022.
- J. Correa and E. Bareinboim. A calculus for stochastic interventions: causal effect identification and surrogate experiments. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(06):10093–10100, Apr. 2020.
- T. Cover and J. Thomas. *Elements of Information Theory*. Wiley Series in Telecommunications, New York, 1991.
- E. Daniel and K. Murphy. Exact bayesian structure learning from uncertain interventions. In M. M. and X. Shen, editors, *Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics*, volume 2 of *Proceedings of Machine Learning Research*, pages 107–114, San Juan, Puerto Rico, 2007. PMLR.
- A. Datta, S. Sen, and Y. Zick. Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 598–617, 2016.
- F. Eberhardt and R. Scheines. Interventions and causal inference. *Philosophy of Science*, 74(5):981–995, 2007.
- C. Frye, C. Rowat, and I. Feige. Asymmetric shapley values: incorporating causal knowledge into model-agnostic explainability. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- S. Greenland and J. Robins. Identifiability, exchangeability, and epidemiological confounding. *Int. Journal for Epidemiology*, 15(3):413–9, 1986.
- J. Halpern and C. Hitchcock. Graded causation and defaults. *The British Journal for the Philosophy of Science*, 66:413–457, 2013.
- T. Heskes, E. Sijben, I. G. Bucur, and T. Claassen. Causal shapley values: Exploiting causal knowledge to explain individual predictions of complex models. *preprint arXiv:2011.01625*, 2020.
- C. Hitchcock and J. Knobe. Cause and norm. *The Journal of Philosophy*, 106:587–612, 2009.
- D. Janzing, D. Balduzzi, M. Grosse-Wentrup, and B. Schölkopf. Quantifying causal influences. *Annals of Statistics*, 41(5):2324–2358, 2013.
- D. Janzing, L. Minorics, and P. Bloebaum. Feature relevance quantification in explainable ai: A causal problem. In S. Chiappa and R. Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 2907–2916, Online, 26–28 Aug 2020. PMLR.
- Y. Jung, S. Kasiviswanathan, D. Janzing, P. Blöbaum, and B. Elias. *do*-shapley: Towards causal interpretation of model prediction. In *Proceedings of 39. International Conference on Machine Learning (ICML)*. 2022.
- Y. Kano and S. Shimizu. Causal inference using nonnormality. In *Proceedings of the International Symposium on Science of Modeling, the 30th Anniversary of the Information Criterion*, pages 261–270, Tokyo, Japan, 2003.
- K. Korb, L. Hope, A. Nicholson, and K. Axnick. Varieties of causal intervention. In C. Zhang, G. H., and W. Yeap, editors, *Trends in Artificial Intelligence*, volume 3157 of *Lecture Notes in Computer Science*. Springer, 2004.
- E. Krapohl, K. Rimfeld, N. Shakeshaft, M. Trzaskowski, A. McMillan, J.-B. Pingault, K. Asbury, N. Harlaar, Y. Kovas, P. Dale, and R. Plomin. The high heritability of educational achievement reflects many genetically influenced traits, not just intelligence. *Proceedings of the National Academy of Sciences*, 111(42):15273–15278, 2014.
- R. Lewontin. Annotation: the analysis of variance and the analysis of causes. *American Journal Human Genetics*, 26(3):400–411, 1974.
- S. Lundberg and S. Lee. A unified approach to interpreting model predictions. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4765–4774. Curran Associates, Inc., 2017.
- F. Markowetz, S. Grossmann, and R. Spang. Probabilistic soft interventions in conditional gaussian networks. In *Proceedings of the conference Artificial Intelligence and Statistics (AISTATS)*, PMLR R5:214–221, 2005.

- R. Mitchell, J. Cooper, E. Frank, and G. Holmes. Sampling permutations for shapley value estimation. *Journal of Machine Learning Research*, 23(43):1–46, 2022.
- J. Mooij, J. Peters, D. Janzing, J. Zscheischler, and B. Schölkopf. Distinguishing cause from effect using observational data: methods and benchmarks. *Journal of Machine Learning Research*, 17(32):1–102, 2016.
- R. Northcott. Can ANOVA measure causal strength? *The Quarterly Review of Biology*, 83(1):47–55, 2008.
- N. Pawlowski, D. Coelho de Castro, and B. Glocker. Deep structural causal models for tractable counterfactual inference. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2020.
- J. Pearl. *Causality*. Cambridge University Press, 2000.
- J. Pearl. Direct and indirect effects. In *Proceedings of the Seventh Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 411–420, San Francisco, CA, 2001. Morgan Kaufmann.
- J. Pearl. Interpretation and identification of causal mediation. *Psychological methods*, 19(4):459–481, 2014.
- J. Pearl and J. Mackenzie. *The book of why*. Basic Books, USA, 2018.
- J. Peters, J. Mooij, D. Janzing, and B. Schölkopf. Causal discovery with continuous additive noise models. *Journal of Machine Learning Research*, 15:2009–2053, 2014.
- J. Peters, D. Janzing, and B. Schölkopf. *Elements of Causal Inference – Foundations and Learning Algorithms*. MIT Press, 2017.
- S. P. R. Rose. Commentary: heritability estimates—long past their sell-by date. *International journal of epidemiology*, 35 3:525–7, 2006.
- P. K. Rubenstein, S. Weichwald, S. Bongers, J. M. Mooij, D. Janzing, M. Grosse-Wentrup, and B. Schölkopf. Causal consistency of structural equation models. In *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence (UAI 2017)*, 2017.
- G. Schamberg, W. Chapman, S.-P. Xie, and T. P. Coleman. Direct and indirect effects—an information theoretic perspective. *Entropy*, 22(8), 2020.
- L. Shapley. A value for n-person games. *Contributions to the Theory of Games (AM-28)*, 2, 1953.
- I. Sobol. Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates. *Mathematics and Computers in Simulation*, 55(1):271 – 280, 2001. The Second IMACS Seminar on Monte Carlo Methods.
- P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. Springer-Verlag, New York, NY, 1993.
- A. Storkey, E. Simonotto, H. Whalley, S. Lawrie, L. Murray, and G. McGonigle. Learning structural equation models for fmri. In *NeurIPS 2006*, 2006.
- J. Tian and J. Pearl. Causal discovery from changes. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence, UAI’01*, page 512–521. Morgan Kaufmann Publishers Inc., 2001.
- J. Von Kügelgen, A. Mohamed, and S. Beckers. Backtracking counterfactuals. In M. van der Schaar, C. Zhang, and D. Janzing, editors, *Proceedings of the Second Conference on Causal Learning and Reasoning*, volume 213 of *Proceedings of Machine Learning Research*, pages 177–196. PMLR, 11–14 Apr 2023.
- J. Wang and K. Mueller. Visual causality analysis made practical. In *IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 151–161, 2017.
- J. Wang, J. Wiens, and S. Lundberg. Shapley flow: A graph-based approach to interpreting model predictions. In A. Banerjee and K. Fukumizu, editors, *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15, 2021, Virtual Event*, volume 130 of *Proceedings of Machine Learning Research*, pages 721–729. PMLR, 2021.
- K. Zhang and A. Hyvärinen. On the identifiability of the post-nonlinear causal model. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, Montreal, Canada, 2009.

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: No
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): Not Applicable
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider): Not Applicable
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: Not Applicable
 - (b) The license information of the assets, if applicable: Yes
 - (c) New assets either in the supplemental material or as a URL, if applicable: Not Applicable
 - (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 SHAPLEY VALUES

Definition 4 (Shapley values). *Let N be a set with n elements (called ‘players’ in the context of game theory) and $\nu : 2^N \rightarrow \mathbb{R}$ be a set function with $\nu(\emptyset) = 0$ (assigning a ‘worth’ to each ‘coalition’). Then the Shapley value of $i \in N$ is given by*

$$\varphi_i(\nu) := \sum_{S \subset N \setminus \{i\}} \frac{|S|!(n - |S| - 1)!}{n!} (\nu(S \cup \{i\}) - \nu(S)) = \sum_{S \subset N \setminus \{i\}} \frac{1}{n \binom{n-1}{|S|}} (\nu(S \cup \{i\}) - \nu(S)). \quad (25)$$

$\varphi_i(\nu)$ is thought of measuring the contribution of each player in a fair way and satisfies

$$\sum_{i=1}^n \varphi_i(\nu) = \nu(N). \quad (26)$$

Lemma 1 (dummy noise variables). *Let N_1, \dots, N_n be noise variables of an SCM \mathcal{M} with observed nodes X_1, \dots, X_n . Let $\tilde{\mathcal{M}}$ be a modified SCM with observed variables X_1, \dots, X_{n+k} and noise variables N_1, \dots, N_{n+k} modeling the same joint distribution on $X_1, \dots, X_n, N_1, \dots, N_n$. Assume that the additional noise variables N_{n+1}, \dots, N_{n+k} are irrelevant for X_j , that is*

$$N_{n+1}, \dots, N_{n+k} \perp\!\!\!\perp X_j \mid \mathbf{N}_T, \quad (27)$$

for all $T \subset \{1, \dots, n\}$. Then \mathcal{M} and $\tilde{\mathcal{M}}$ yield the same values for $ICC_{\psi}^{Sh}(X_i \rightarrow X_j)$ for all $i = 1, \dots, n$.

We first need the following property of Shapley values:

Lemma 2 (adding zero value players). *For an extended set $\tilde{N} \supset N$, define a coalition function $\tilde{\nu} : 2^{\tilde{N}} \rightarrow \mathbb{R}$ by*

$$\tilde{\nu}(\tilde{S}) := \nu(\tilde{S} \cap N),$$

that is, $\tilde{\nu}$ is an extension of ν to irrelevant elements. Then

$$\varphi_i(\tilde{\nu}) = \varphi_i(\nu) \quad \forall i \in N.$$

Proof. It is sufficient to show the claim for the case where \tilde{N} contains just one additional element, say $n + 1$, since the remaining part follows by induction.

When computing $\varphi_i(\tilde{\nu})$ via a sum over all $\tilde{S} \subset \tilde{N}$ we can always merge two corresponding terms: one set S not containing $n + 1$ and one corresponding set $S' := S \cup \{n + 1\}$. Due to the irrelevance of $n + 1$ we have

$$\tilde{\nu}(S' \cup \{i\}) - \tilde{\nu}(S') = \tilde{\nu}(S \cup \{i\}) - \tilde{\nu}(S) = \nu(S \cup \{i\}) - \nu(S),$$

that is, both terms are the same as for the set S in (9), up to the combinatorial factors. For the computation of $\tilde{\varphi}_i$, the term with S comes with the factor $|S|!(n - |S|)!/(n + 1)!$, while S' comes with $(|S| + 1)!(n - |S| - 1)!/(n + 1)!$. The sum of these terms reads

$$\frac{|S|!(n - |S| - 1)!}{(n + 1)!} ((n - |S|) + (|S| + 1)) = \frac{|S|!(n - |S| - 1)!}{n!},$$

which coincides with the factor in (8). \square

To prove Lemma 1, we note that ICC_{ψ}^{Sh} is given by first defining the coalition function $\nu(S) := -\psi(X_j|\mathbf{N}_S)$, for each $S \subset \{1, \dots, n\}$. Then $ICC_{\psi}^{Sh}(X_i \rightarrow X_j) = \varphi_i(\nu)$. Further, define

$$\tilde{\nu}(\tilde{S}) := -\psi(X_j|\mathbf{N}_{\tilde{S}}). \quad (28)$$

for each $\tilde{S} \subset \{1, \dots, n+k\}$. Then, $\tilde{\nu}(\tilde{S}) = \nu(\tilde{S} \cap \{1, \dots, n\})$. To see this, set $S := \tilde{S} \cap \{1, \dots, n\}$. Then (28) implies

$$\psi(X_j|\mathbf{N}_{\tilde{S}}) = \psi(X_j|\mathbf{N}_S),$$

since conditioning on the dummy nodes does not change the distribution of X_j . Hence, $\tilde{\nu}$ defines an extension of ν to irrelevant elements in the sense of Lemma 2. Since $ICC_{\psi}^{Sh}(X_i \rightarrow X_j)$ with respect to the extended SCM is given by $\tilde{\varphi}_i(\tilde{\nu})$, the statement follows from Lemma 2.

B STRENGTH OF EDGES AND PATHS

B.1 Strength of causal arrows and indirect causal influence

Janzing et al. (2013) defined ‘strength of an edge’ in the sense of an information theoretic quantity. It is based on an operation they called ‘cutting of edges’. To quantify the information transferred along an arrow, one thinks of arrows as ‘channels’ that propagate information through space – for instance ‘wires’ that connect electrical devices. To measure the impact of an arrow $X_j \rightarrow X_i$, they ‘cut’ it and feed it with a random input that is an i.i.d. copy of X_j . This results in the following ‘post-cutting’ distribution:

Definition 5 (Single arrow post-cutting distribution). *Let G be a causal DAG with nodes X_1, \dots, X_n and P_{X_1, \dots, X_n} be Markovian with respect to G . Further, let \mathbf{PA}_i^j denote the parents of X_i without X_j . Define the ‘post-cutting conditional’ by*

$$p_{X_j \rightarrow X_i}(x|\mathbf{pa}_i^j) := \sum_{x_j} p(x_i|\mathbf{pa}_i^j, x_j)p(x_j). \quad (29)$$

Then, the post-cutting distribution $P_{X_j \rightarrow X_i}(x_1, \dots, x_n)$ is defined by replacing $p(x_i|\mathbf{pa}_i)$ in the causal factorization $p(x_1, \dots, x_n) = \prod_j p(x_j|\mathbf{pa}_j)$ with (29).

The relative entropy between the observed joint distribution $P(x_1, \dots, x_n)$ and the post-cutting distribution $P_{X_i \rightarrow X_j}$ now measures the strength of the arrow:

$$\mathfrak{C}_{X_i \rightarrow X_j} := D(P||P_{X_i \rightarrow X_j}).$$

This measure is one of the concepts for which it is most apparent that its intention is different from ICC, and even complementary in a sense. To see this, note that Postulate 2 in Janzing et al. (2013) states that a measure of strength of an edge should be independent from how its tail node depends on its parents. This implies, in particular, that the values $\mathfrak{C}_{X_2 \rightarrow X_3}$ for the DAGs $X_1 \rightarrow X_2 \rightarrow X_3$ and $X_2 \rightarrow X_3$ coincide. Thus, the postulate explicitly requires $\mathfrak{C}_{X_i \rightarrow X_j}$ to ignore whether the information of X_i has been inherited or not.

This fundamental conceptual difference to ICC carries over to many other quantifications of causal influence, in particular the information theoretic *indirect* and *path-specific* causal influence in Schamberg et al. (2020), which generalizes (Janzing et al., 2013). For the quantification of the influence of X_2 on X_3 in $X_1 \rightarrow X_2 \rightarrow X_3$, the indirect influence (which is the same as the direct one here) by Schamberg et al. (2020) coincides with the strength of the arrow $X_2 \rightarrow X_3$ from Janzing et al. (2013). Further, note that also more classical approaches to mediation analysis and quantifying indirect effect, e.g., Pearl (2001, 2014) have a different intention and do not distinguish whether the information node X_2 propagated to the target X_3 has been inherited from X_2 ’s parents or generated at the node itself.

B.2 Shapley Flow

Wang et al. (2021) aim to quantify the contribution of *edges* to the value attained by a specific target node. They argue that these attributions can also be used to measure the contribution of a *node* to the target. First, we are going to explain the concept of Shapley Flow. In section B.3 we want to argue, that ICC can be seen as a special case of a modified version of Shapley Flow.

Boundary consistency Wang et al. (2021) state that a quantification of the influence of edges must be invariant to *refinement* of the causal model. This is formalized in an axiom, which they call *boundary consistency*. They partition the nodes of a DAG G into a data side D and a model side F (note that the terminology is motivated by applying their concept to model explainability in the context of explainable AI), where all root nodes (i.e. nodes without parents) of the DAG must lie in D and the target node in F . Further, there must not be edges from F to D . The idea is that F represents a new model with D as input data. Let $\phi_{\mathcal{B}}(e)$ be the contribution of edge e to the target node with respect to the boundary $\mathcal{B} = (D, F)$. Let further $\text{cut}(D, F)$ be the set of all edges with one end in D and one end in F . Then ϕ is boundary consistent, when

$$\phi_{\mathcal{B}_1}(e) = \phi_{\mathcal{B}_2}(e) \text{ iff } e \in (\text{cut}(\mathcal{B}_1) \cap \text{cut}(\mathcal{B}_2)). \quad (30)$$

Update of edges The goal of Shapley Flow is to attribute the change in the target node between a foreground setting \mathbf{x} and a background setting \mathbf{x}^* of the variables. To derive an attribution that fulfils the boundary consistency axiom, they traverse the DAG in a depth-first search (DFS) order and *update* edges along the way. We will discuss later, how the attribution is symmetrized over all possible DFS orders. When updated, an edge propagates an *updated* value (or the given foreground value for root nodes) to one child at a time. The other children will be updated later, when they are explored by the DFS. The child node uses this new value to update its own value according to the SCM (and stays at this value until updated again). For its other parents the child resorts back to the last value it received in an update or their background values if the respective edge has not been updated yet. This is motivated by the analogy to a message passing system that propagates the updated values step by step through the graph. Formally, the *history* of updates is represented as a list of edges. The value of a variable $X_i \in D$ after updating the edges in history h is calculated according to the SCM and the values of the parents \mathbf{PA}_i that X_i has been sent and stored, namely

$$\tilde{x}_i(h) = f_i(m_{i,1}(h), \dots, m_{i,l}(h)), \quad (31)$$

where $l = |\mathbf{PA}_i|$ and $m_{i,j}(h)$ is the value in the memory slot where X_i stores the value of its j -th parent node, after history h . Let $\mathbf{PA}_{i,j}$ denote the j -th parent of X_i . Initially (for the empty history), all these memory slots contain the respective baseline values from \mathbf{x}^* , i.e. with $X_k = \mathbf{PA}_{i,j}$ we have

$$m_{i,j}([\]) = x_k^*. \quad (32)$$

Let h be a history, $(X_k \rightarrow X_i)$ an edge and $h' = h + (X_k \rightarrow X_i)$ be the history that we get when we append $(X_k \rightarrow X_i)$ to the end of h . When we already know the values $m_{i,j}(h)$ and $\tilde{x}_i(h)$ for $i = 1, \dots, n$ and $j = 1, \dots, |\mathbf{PA}_i|$ we get the updated memory for h' by changing the memory slot, where X_i stores the value of X_k . This means, if $X_k = \mathbf{PA}_{i,j}$ we get

$$m_{i,j}(h') = \begin{cases} x_k & \text{if } X_k \text{ is a root node} \\ \tilde{x}_k(h) & \text{else} \end{cases} \quad (33)$$

where x_k is the foreground value of X_k from \mathbf{x} and $\tilde{x}_k(h)$ is the updated value of the node X_k after history h (as defined in equation 31). The other memory values stay the same, i.e.

$$m_{\hat{i},j}(h') = m_{\hat{i},j}(h) \quad (34)$$

for $\hat{i} = i$ and $\mathbf{PA}_{i,j} \neq X_k$ or $\hat{i} \neq i$ and $j = 1, \dots, |\mathbf{PA}_{\hat{i}}|$. Note, that a node X_k can have multiple children and can thus appear in the set \mathbf{PA}_i for several X_i . Yet, the respective value does not have to be the same, since these children might have been sent an updated value of X_k at different points along history h . I.e., if $\mathbf{PA}_{i,j} = X_k = \mathbf{PA}_{\hat{i},\hat{j}}$ for $i \neq \hat{i}$ that does not necessarily mean, that $m_{i,j}(h) = m_{\hat{i},\hat{j}}(h)$ for all h .

Once an update crosses the boundary $\mathcal{B} = (D, F)$, it immediately reaches the target Y over all paths, i.e. edges between two nodes in F are assumed to instantly propagate the values of their parents. This reflects the idea that F is thought of as a new model that operates normally, while Shapley Flow only alters the dependencies in the data. Let for $X_i \in F$ the set \mathbf{PA}_i^D be the parents of X_i in D and \mathbf{PA}_i^F the parents in F . Then replace the functions of the SCM according to

$$f_i^h(\mathbf{pa}_i^F) = f_i(\mathbf{pa}_i^F, \mathbf{pa}_i^D(h)). \quad (35)$$

Note, that for a fix h the value $\mathbf{pa}_i^D(h)$ is a constant. As a special case, the equation of an X_i does not change, if all its parents are in F . The equations can recursively be solved for the variables in D (like a normal SCM would be solved for the noise variables). For a variable $X_i \in F$ the notation $\tilde{x}_i(h)$ refers to the value of X_i after solving the functional equations for the nodes in D and inserting the respective values for history h .

As an intermediate step, we define a function $\nu_{\mathcal{B}}$ that assigns a value to a history, according to

$$\nu_{\mathcal{B}}(h) = \tilde{y}(h). \quad (36)$$

Depth-first search The histories that Shapley Flow actually considers, are the ones that emerge from a DFS. In a DFS, the graph is traversed in a recursive manner, i.e., each node recursively does a DFS on all its children. Once the target Y is reached, the algorithm returns. Through this, the DFS finds all paths from the root nodes to the target node. Note, that DFS does not dictate an order in which a node has to explore its children. This means, there are several orderings in which a DFS can find the paths. Further note, that the number of these possible orderings is smaller, than the number of all permutations of paths. Once DFS is called on a node X , it will be called recursively on all children of X before it backtracks. Therefore, the respective paths always appear next to each other in an ordering derived from a DFS.

The purpose of these DFS orderings is to ensure boundary consistency. To get some intuition, consider two boundaries $\mathcal{B}_1 = (D, F)$, $\mathcal{B}_2 = (D \cup \{X\}, F \setminus \{X\})$, that only differ in one node $X \notin D$. Once a message passes the boundary, it is immediately propagated further by all downstream edges. So, once an update reaches X w.r.t. \mathcal{B}_1 , all outgoing edges of X instantly propagate this update. Once an update reaches X w.r.t. \mathcal{B}_2 , all outgoing edges of X will be updated, before any update can cross the boundary elsewhere. This means for the other edges it *looks like* nothing changed from \mathcal{B}_1 to \mathcal{B}_2 .

Shapley contribution of edges The attribution to an edge e is the sum of all path attributions to paths, that contain e . This procedure can also be formulated as follows. Let Π_{DFS} be the set of orderings of paths, in which they can be found through depth-first search. For path i from root node to sink node define

$$\tilde{\phi}(i) := \frac{1}{|\Pi_{DFS}|} \sum_{\pi \in \Pi_{DFS}} (\tilde{\nu}([j \mid \pi(j) \leq \pi(i)]) - \tilde{\nu}([j \mid \pi(j) < \pi(i)])), \quad (37)$$

where $\tilde{\nu}(S)$ operates on an list of lists of edges and evaluates the function $\nu_{\mathcal{B}}$ on the concatenation of all paths in S and with respect to a boundary \mathcal{B} . Here we assume that a path is represented as a list of edges and that these lists of paths are sorted according to π . Then we have

$$\tilde{\nu}(S) := \nu_{\mathcal{B}}(s_1 + \dots + s_m), \quad m = |S|, s_i \in S. \quad (38)$$

Eventually, the Shapley Flow for an edge e is

$$\phi(e) := \sum_{p \in \text{paths in } G} \mathbb{1}[e \in p] \tilde{\phi}(p). \quad (39)$$

Wang et al. (2021) implicitly assume that all relationships in the DAG are deterministic, with only the root nodes being random. Put into the framework defined in Eq. (1) in the main paper, this means, that all nodes depend either on endogenous parents or on exogenous noise, i.e., all structural equations have one of the following forms

$$X_j = f_j(\mathbf{PA}_j) \quad \text{or} \quad X_j = f_j(N_j). \quad (40)$$

If one wishes to calculate the Shapley Flow on a general SCM, the *augmented* DAG is needed, containing all the noise variables N_j as additional nodes with the edges $N_j \rightarrow X_j$.

B.3 Describing ICC using modified Shapley Flow

Adding the noise variables Now consider the augmented DAG G' with $2n$ nodes and the boundary $\mathcal{B}' := (\mathbf{N}, \mathbf{X})$. Since the SCM assumes causal sufficiency (i.e. all noise terms are only causing one variable) the cut of our boundary \mathcal{B}' contains the n edges $N_j \rightarrow X_j$. Further, the set of DFS paths from the nodes N_j to the boundary is also n (namely the paths $N_j \rightarrow X_j$) and the number of DFS orderings thereof is $n!$. Wang et al. (2021) emphasize that in general attributions to nodes can be made by summing up the attribution to all outgoing edges. Since the nodes N_j only have $N_j \rightarrow X_j$ as outgoing edges, we can interpret their contribution as the contribution of X_j . We will therefore identify the edges with their respective noise node in the following equation. Rephrasing Equation (37) with respect to boundary \mathcal{B}' results in

$$\tilde{\phi}(i) = \frac{1}{n!} \sum_{\pi \in \Pi} [\tilde{\nu}(\{j \mid \pi(j) \leq \pi(i)\}) - \tilde{\nu}(\{j \mid \pi(j) < \pi(i)\})] = \sum_{S \subset N \setminus \{i\}} \frac{1}{n \binom{n-1}{|S|}} (\tilde{\nu}(S \cup \{i\}) - \tilde{\nu}(S)),$$

where Π is the set of all permutations over N_1, \dots, N_n . Hence, the Shapley Flow amounts to measuring the impact of updating the edges $N_j \rightarrow X_j$ in a symmetrized way. This way, it reduces to the standard way of measuring the impact of n independent causes N_1, \dots, N_n on the value of the target node.

Replacing the target metric Since Shapley Flow quantifies contribution of edges to the *value* of a target node, while we want to measure the uncertainty of the target node, we need a different choice for the set function \tilde{v} : instead of measuring the value of the target node, we need a metric that measures the uncertainty of the target after activating an edge by feeding with an independent random input. This amounts to replacing the set function above with an uncertainty measure ψ to finally obtain *ICC*.

Remark With this reinterpretation, the flow conservation property of Shapley Flow nicely captures the intuition that the strengths of the outgoing edges sum up to the flow coming from the parents plus the intrinsic contribution. This is seen as follows. The flow property implies

$$\sum_{X_j \in \mathbf{PA}'_i} \phi(X_j \rightarrow X_i) = \sum_{X_k \in \mathbf{CH}'_i} \phi(X_i \rightarrow X_k), \quad (41)$$

with \mathbf{PA}'_i and \mathbf{CH}'_i being the parents and children, respectively, of X_i in the augmented DAG G' . This amounts to

$$\sum_{X_j \in \mathbf{PA}_i} \phi(X_j \rightarrow X_i) + \phi(N_i \rightarrow X_i) = \sum_{X_k \in \mathbf{CH}_i} \phi(X_i \rightarrow X_k),$$

where \mathbf{PA}_i and \mathbf{CH}_i are the parents and children, respectively, of X_i in the original DAG G , and $\phi(N_i \rightarrow X_i)$ measuring the intrinsic contribution.

C DEFINITION OF INFORMATION FLOW

As pointed out by Ay and Polani (2008), quantifying causal influence between observed nodes via the information they share, requires computing information with respect to *interventional* probabilities rather than information given in a passive *observational* scenario (recall that this distinction has been irrelevant for us since dependences between observed nodes and noise are always causal and unconfounded). Accordingly, they define the **information flow** from a set \mathbf{X}_A of variables to another set \mathbf{X}_B , imposing that some background variables \mathbf{X}_S are set to \mathbf{x}_S by

$$I(\mathbf{X}_A \rightarrow \mathbf{X}_B | do(\mathbf{x}_S)) := - \sum_{\mathbf{x}_A, \mathbf{x}_B} p(\mathbf{x}_B | do(\mathbf{x}_A, \mathbf{x}_S)) p(\mathbf{x}_A | \mathbf{x}_S) \log \frac{p(\mathbf{x}_B | do(\mathbf{x}_A, \mathbf{x}_S))}{\sum_{\mathbf{x}'_A} p(\mathbf{x}_B | do(\mathbf{x}'_A, \mathbf{x}_S)) p(\mathbf{x}'_A | do(\mathbf{x}_S))}. \quad (42)$$

Here, $p(\cdot | do(\mathbf{x}_A, \mathbf{x}_S))$ is the interventional probability after setting $\mathbf{X}_A, \mathbf{X}_S$ to $\mathbf{x}_A, \mathbf{x}_S$. Moreover, Ay and Polani (2008) define also the average of (42) over all \mathbf{x}_S :

$$I(\mathbf{X}_A \rightarrow \mathbf{X}_B | do(\mathbf{X}_S)) := \sum_{\mathbf{x}_S} I(\mathbf{X}_A \rightarrow \mathbf{X}_B | do(\mathbf{x}_S)) p(\mathbf{x}_S).$$

Note that $I(\mathbf{X}_A \rightarrow \mathbf{X}_B | do(\mathbf{x}_S))$ measures the mutual information of \mathbf{X}_A and \mathbf{X}_B when \mathbf{X}_S is set to \mathbf{x}_S and \mathbf{X}_A is randomized with probability $p(\mathbf{X}_A | do(\mathbf{x}_S))$.

D ADDITIONAL INFORMATION EXPERIMENTS AND MORE DETAILS

For ICC on the Auto MPG data set we have used the publicly available implementation of ICC in DoWhy (Blöbaum et al., 2022), where we have chosen the auto-assign function.

For the river flows, we have worked with an SCM where each variable is a sum of its parents plus noise, that is, all regression coefficients are set to 1. Accordingly, the variance of the target after adjusting all noise variables in S , is given by

$$\text{Var}(X_{Sy} | do(\mathbf{N}_S = \mathbf{n}_S)) = \text{Var}\left(\sum_{j \in \bar{S}} N_j\right), \quad (43)$$

which does not coincide with the sum of the corresponding variances because the noise terms heavily correlate due to confounding. To compute the Shapley value computed for the set function (43) (up to a negative sign) we have used variance-based feature relevance attribution `gcm.parent_relevance` from DoWhy GCM (Blöbaum et al., 2022) with linear regression.

To explicitly show the difference to do-intervention based measures, we consider a modification of do-Shapley using variance as target metric instead of expectation based on the set function

$$\nu(S) := - \sum_{\mathbf{x}_{\bar{S}}} \text{Var}(X_n | do(\mathbf{X}_S = \mathbf{x}_S)) p(\mathbf{x}_S), \quad (44)$$

where $S \subset \{1, \dots, n-1\}$ since conditioning on $do(X_n = x_n)$ is pointless. Whenever the mediator New Jumbles Rock is in the adjustment set, all the variance that originates from the root nodes is screened off. Accordingly, New Jumbles Rock gets quite high contribution, namely around 41%. To compute variance-based do-Shapley for this example, we have used the following rule: whenever New Jumbles Rock is in the adjustment set S , we have $\text{Var}(X_{Sy} | do(\mathbf{N}_S = \mathbf{n}_S)) = \text{Var}(N_{Sy})$, that is the influx at Samlesbury is the only source of uncertainty. Otherwise (that is, the adjustment set contains only root nodes), the variance of Samlesbury is given by the variance of the sum over all noise terms of the non-adjusted nodes and reads $\text{Var}(X_{Sy} | do(\mathbf{N}_S = \mathbf{n}_S)) = \text{Var}(\sum_{j \in \bar{S}} N_j)$.

We also tried ‘causal-DAG agnostic’ Shapley value based feature relevance attribution where the interventional conditional in (44) is replaced with the observational conditional. Then, more than 100% of the variance is attributed to the mediator X_{JR} (with slightly negative contributions of the root nodes), aligning with the fact that all other features were conditionally independent of the target, given the mediator, if there were no confounding effects. There, we have again used the function `gcm.parent_relevance` from DoWhy GCM (Blöbaum et al., 2022), but this time we regress X_{Sy} on the river flows of the upstream nodes, rather than on the noise (i.e. the influx).

The fact that the flow at New Jumbles Rock is mostly explained by the flow of the root nodes is only reflected by the ICC values and not by the other approaches of quantifying influence mentioned here.