

# CYCLES IN CAUSAL LEARNING

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

In the causal learning setting, we wish to learn cause-and-effect relationships between variables such that we can correctly infer the effect of an intervention. While the difference between a cyclic structure and an acyclic structure may be just a single edge, cyclic causal structures have qualitatively different behavior under intervention: cycles cause *feedback loops* when the downstream effect of an intervention propagates back to the source variable. We present three theoretical observations about probability distributions with self-referential factorizations, i.e. distributions that could be graphically represented with a cycle. First, we prove that self-referential distributions in two variables are, in fact, independent. Second, we prove that self-referential distributions in  $N$  variables have zero mutual information. Lastly, we prove that self-referential distributions that factorize in a cycle, also factorize as though the cycle were reversed. These results suggest that cyclic causal *dependence* may exist even where observational data suggest *independence* among variables. Methods based on estimating mutual information, or heuristics based on independent causal mechanisms, are likely to fail to learn cyclic causal structures. We encourage future work in causal learning that carefully considers cycles.

## 1 INTRODUCTION

### 1.1 MOTIVATION FOR CAUSAL LEARNING

Current methods in machine learning struggle to generalize beyond a test set drawn from a distribution identical to the training distribution. We would like to learn models that generalize more robustly: we want models that respond sensibly to adversarial attacks, out-of-distribution samples, and distributional shift. We can think of these tests for robustness as tests for causal learning - we are testing whether the model can adapt to a distribution that is indeed different from the training distribution, but that has the same underlying causal relationships between its high-level variables.

Following Pearl (2009)'s interventional definition of causal graphs, we define the goal in causal learning: we wish to learn enough about the underlying causal relationships between variables to be able to correctly predict the effect of an intervention (Bengio et al., 2019). A model that can predict the outcome of an intervention has learned the stable mechanisms underlying cause and effect, rather than memorizing spurious correlation in the data. Whether we actually perform the intervention is immaterial - a model that can imagine an intervention is performing the counterfactual reasoning necessary for decision making.

Whereas in the statistical learning setting, we can often achieve our objectives using directed acyclic graphical models, in the causal learning setting the need for cyclic models is clear: in any causal setting where intervention on a variable would cause a downstream effect on itself, the model requires a cyclic structure.

### 1.2 CYCLES IN STATISTICAL LEARNING

In the statistical learning setting, the majority of literature on graphical models is based on directed acyclic graphs (DAGs). This confers a significant advantage: in DAGs, the global and local Markov properties are equivalent, even without assuming positivity of the probability distribution (Lauritzen et al., 1990). In contrast, in directed cyclic graphs, the global and local directed Markov properties are *not* equivalent (Richardson, 1997).

However, early work in the field of graphical models did consider cyclic graphs. Spirtes (1994) and Koster et al. (1996) showed that d-separation is valid on cyclic graphs, provided that the equations are linear and all distributions are Gaussian. Pearl & Dechter (2013) showed that d-separation generalizes to all directed cyclic graphs, provided the variables are discrete. We will rely on this result in the proof of Theorem 2.

In certain applications, the nature of observed data is sufficiently cyclic that it is worthwhile to sacrifice the Markov property in order to use cyclical graphical models. For example, feedback loops are common in neuronal spiking (Douglas & Martin, 2007), economic activity (Moe, 1985), and gene transcription regulation networks (Alon, 2007). In time-series data, cycles may represent quantization, in which the time scale that data was collected was too coarse to see the individual events. Models such as Dynamic Bayesian networks handle time-series data by “unrolling” cycles in time, but with a tradeoff of space and data efficiency (Ghahramani, 1997).

### 1.3 CYCLES IN CAUSAL LEARNING

In causal learning, we assume that most data in the real world is generated from the composition of a few independent causal mechanisms (Pearl, 2009) (Schölkopf et al., 2012). Peters et al. (2017) suggest that for this kind of data, there will exist a good factorization of high-level variables that mimics the true causal structure. This motivates the heuristic used in Bengio et al. (2019) to learn a causal graph: assuming independent causal mechanisms, the heuristic independently parameterizes the likelihood of a certain variable being a parent (direct cause) of another variable.

However, in designing loss functions for causal learning, we should note that the difference between learning an acyclic causal structure and a cyclic causal structure, might be the existence of a single edge (representing a single parent-child relationship between variables). While this might trigger a small penalty in a loss function that treats edges independently, failing to learn a cycle in a causal structure is a significant qualitative error: we would be failing to expect a feedback loop in response to intervention on a variable within a cycle. Under intervention, we would like to know whether to expect a positive feedback loop or a negative feedback loop, and what kind of equilibrium, if any, the causal structure will reach.

We present the following theoretical results that suggest that cyclic causal structures may lurk in places we may not think to consider a cyclic structure, given knowledge of the observational distribution. In fact, to find cyclic causal structures, we may need to look for causal *dependence* in places where observational data suggests complete *independence* between variables. In particular, the mutual information on variables in a cycle is zero, suggesting that methods based on estimating mutual information are likely to overlook cycles.

## 2 RESULTS

**Theorem 1.** *A two-variable probability distribution  $P$  that factorizes according to  $p(x, y) = p(x|y)p(y|x)$ , also factorizes according to  $p(x, y) = p(x)p(y)$ .*

*Proof.* Recall Equations (1) and (2), standard definitions in information theory (Cover & Thomas, 2012). From there, we proceed with substitution.

$$H(X, Y) = H(X|Y) + H(Y|X) + I(X; Y) \quad (1)$$

$$\begin{aligned} H(X, Y) &= - \sum_{x,y} p(x, y) \log p(x, y) \quad (2) \\ &= - \sum_{x,y} p(x|y)p(y|x) \log p(x|y)p(y|x) \\ &= - \sum_{x,y} p(x|y)p(y|x) \log p(x|y) - \sum_{x,y} p(x|y)p(y|x) \log p(y|x) \\ &= H(X|Y) + H(Y|X) \end{aligned}$$

Therefore,

$$I(X; Y) = 0, \text{ which implies } p(x, y) = p(x)p(y).$$

□

**Theorem 2.** A probability distribution  $P$  that factorizes as an  $n$ -variable cycle, according to  $p(x_1, x_2, \dots, x_n) = p(x_1|x_2)p(x_2|x_3) \dots p(x_{n-1}|x_n)p(x_n|x_1)$ , has mutual information  $I(x_1; x_2; \dots; x_n) = 0$ .

The full proof appears in APPENDIX A. Here we present a sketch of the proof.

*Proof Sketch.* Our argument will follow a similar structure as the proof of Theorem 1. We first write the joint entropy as a sum of conditional entropy terms, by rewriting a log of products as a sum of logs (Lemma 1). Second, we write the joint entropy using its inclusion-exclusion definition from information theory (Lemma 2). We then show that the first representation is equal to all terms except the final joint mutual information term in the second representation. We can therefore conclude that the joint mutual information is zero.

In the two-variable case, it was straightforward to show this equality because the inclusion-exclusion formula didn't contain any higher-order terms except for the joint mutual information. In the  $n$ -variable case, the inclusion-exclusion formula generates many terms.

To handle the additional terms, we recall that Shannon information obeys set additivity, which means that every information-theoretic identity has an equivalent set-theoretic identity. Following Yeung (1991), we define the information measure  $\mu^*$  to be the unique signed measure that is consistent with Shannon's definitions of entropy and mutual information. This allows us to represent mutual information as the signed measure on the intersection of the sets we are inside of, conditioned on the union of the sets we are outside of. For ease of manipulating expressions, we write the proof using set-theoretic notation.

To finish the  $n$ -variable proof, we consider the set expressions representing the *set difference* of each variable and its parent. We apply the inclusion-exclusion formula to the union of those set differences, and show that the higher-order terms fall into one of two cases. Either the term is the intersection of two adjacent set expressions (Lemma 3), which are disjoint, and therefore zero, or the term represents non-adjacent set expressions (Lemma 4), which are  $d$ -separated, and therefore zero. The only remaining terms are those corresponding to the conditional entropy terms in Lemma 2, leaving only the joint mutual information term which must now be zero.

□

**Theorem 3.** A joint distribution that factorizes in an  $n$ -variable cycle, also factorizes such that the  $n$ -variable cycle is reversed. In other words,  $p(x_1, x_2, \dots, x_n) = p(x_1|x_n)p(x_2|x_1)p(x_3|x_2) \dots p(x_n|x_{n-1})$  implies that  $p(x_1, x_2, \dots, x_n) = p(x_n|x_1)p(x_1|x_2)p(x_2|x_3) \dots p(x_{n-1}|x_n)$ .

*Proof.*

Suppose

$$\begin{aligned} p(x_1, x_2, \dots, x_n) &= p(x_1|x_n)p(x_2|x_1)p(x_3|x_2)\dots \\ &= p(x_1|x_n) \prod_{i=1}^{n-1} p(x_{i+1}|x_i) \end{aligned}$$

From Bayes' rule, we know that

$$\begin{aligned} p(x_1|x_n) &= \frac{p(x_n|x_1)p(x_1)}{p(x_n)} \\ p(x_{i+1}|x_i) &= \frac{p(x_i|x_{i+1})p(x_{i+1})}{p(x_i)} \end{aligned}$$

Substituting gives us

$$\begin{aligned} p(x_1, x_2, \dots, x_n) &= \frac{p(x_n|x_1)p(x_1)}{p(x_n)} \prod_{i=1}^{n-1} \frac{p(x_i|x_{i+1})p(x_{i+1})}{p(x_i)} \\ &= \left[ p(x_n|x_1) \prod_{i=1}^{n-1} p(x_i|x_{i+1}) \right] \frac{p(x_1)}{p(x_n)} \prod_{i=1}^{n-1} \frac{p(x_{i+1})}{p(x_i)} \xrightarrow{1} \\ &= p(x_n|x_1) \prod_{i=1}^{n-1} p(x_i|x_{i+1}) \end{aligned}$$

□

### 3 FUTURE WORK

The results presented herein apply to graphs where every variable is part of a single cycle: in the general case we are interested in all cyclic directed graphs, which may contain multiple cycles, overlapping cycles, and variables outside of the cycles. We leave these questions for future work.

### 4 CONCLUSION

In the statistical learning setting, we may be motivated to use cyclic graphs when the data overwhelmingly suggests that we should do so: when we are aware of feedback systems, or when it is natural to represent time-series data. In contrast, in the causal learning setting, we are motivated to learn the relationships between underlying variables so that we are able to infer the effect of an intervention. It is therefore imperative to learn a cyclic causal structure whenever one exists: that is, the case where intervening on a variable will have a downstream effect on itself.

However, the presence of cycles in a causal structure may not be obvious from the observation of their joint distribution. In fact, two variables that appear to be completely independent under observation, may indeed have a cyclic dependence present in their causal structure. A group of variables that shares no information between all its members, may contain a cyclic causal structure. And even if we learn that a cycle exists among a group of variables, we may not have distinguished the correct direction of the cycle.

Causal learning is an open area of research, and one that stands to benefit all of machine learning. Heuristics based on assuming independence between causal mechanisms may be a fruitful avenue for learning acyclic causal relationships, but may discourage learning cyclic structures where they exist. As we advance causal learning, we will need methods that treat cycles as first-class citizens.

### REFERENCES

Uri Alon. Network motifs: theory and experimental approaches. *Nature Reviews Genetics*, 8(6): 450–461, 2007.

- Yoshua Bengio, Tristan Deleu, Nasim Rahaman, Rosemary Ke, Sébastien Lachapelle, Olexa Bilaniuk, Anirudh Goyal, and Christopher Pal. A meta-transfer objective for learning to disentangle causal mechanisms. *arXiv preprint arXiv:1901.10912*, 2019.
- Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- Rodney J Douglas and Kevan AC Martin. Recurrent neuronal circuits in the neocortex. *Current biology*, 17(13):R496–R500, 2007.
- Zoubin Ghahramani. Learning dynamic bayesian networks. In *International School on Neural Networks, Initiated by IIASS and EMFCSC*, pp. 168–197. Springer, 1997.
- Jan TA Koster et al. Markov properties of nonrecursive causal models. *The Annals of Statistics*, 24(5):2148–2177, 1996.
- Steffen L Lauritzen, A Philip Dawid, Birgitte N Larsen, and H-G Leimer. Independence properties of directed markov fields. *Networks*, 20(5):491–505, 1990.
- Terry M Moe. Control and feedback in economic regulation: The case of the nlr. *American Political Science Review*, 79(4):1094–1116, 1985.
- Judea Pearl. *Causality*. Cambridge university press, 2009.
- Judea Pearl and Rina Dechter. Identifying independencies in causal graphs with feedback. *arXiv preprint arXiv:1302.3595*, 2013.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. MIT press, 2017.
- Thomas Richardson. A characterization of markov equivalence for directed cyclic graphs. *International Journal of Approximate Reasoning*, 17(2-3):107–162, 1997.
- Bernhard Schölkopf, Dominik Janzing, Jonas Peters, Eleni Sgouritsa, Kun Zhang, and Joris Mooij. On causal and anticausal learning. *arXiv preprint arXiv:1206.6471*, 2012.
- Peter Spirtes. *Conditional independence in directed cyclic graphical models for feedback*. Carnegie Mellon [Department of Philosophy], 1994.
- Raymond W Yeung. A new outlook on shannon’s information measures. *IEEE transactions on information theory*, 37(3):466–474, 1991.

## A PROOF OF THEOREM 2

For intuition, see the proof sketch appearing in the main paper. We start by stating and proving the following lemmas.

**Lemma 1.**  $H(X_1, X_2, \dots, X_n) = H(X_1|X_2) + H(X_2|X_3) + \dots + H(X_{n-1}|X_n) + H(X_n|X_1)$

*Proof.*

$$\begin{aligned}
 H(X_1, X_2, \dots, X_n) &= - \sum_{x_1, \dots, x_n} p(x_1, \dots, x_n) \log p(x_1, x_2, \dots, x_n) \\
 &= - \sum_{x_1, \dots, x_n} p(x_1, \dots, x_n) \log p(x_1|x_2)p(x_2|x_3) \dots p(x_n|x_1) \\
 &= - \sum_{x_1, \dots, x_n} p(x_1, \dots, x_n) \log p(x_1|x_2) - \dots - \sum_{x_1, \dots, x_n} p(x_1, \dots, x_n) \log p(x_n|x_1) \\
 &= H(X_1|X_2) + H(X_2|X_3) + \dots + H(X_n|X_1)
 \end{aligned}$$

□

**Lemma 2.**  $H(X_1, X_2, \dots, X_n) = \mu^* \left( (X_1 - X_2) \cup (X_2 - X_3) \cup \dots \cup (X_{n-1} - X_n) \cup (X_n - X_1) \right) + I(X_1; X_2; \dots; X_n)$

*Proof.* We note that  $X_1 \cap X_2 \cap \dots \cap X_n$  is disjoint from  $(X_1 - X_2) \cup (X_2 - X_3) \cup \dots \cup (X_{n-1} - X_n) \cup (X_n - X_1)$ .

That allows us to write

$$X_1 \cup X_2 \cup \dots \cup X_n = ((X_1 - X_2) \cup (X_2 - X_3) \cup \dots \cup (X_{n-1} - X_n) \cup (X_n - X_1)) \cup (X_1 \cap X_2 \cap \dots \cap X_n).$$

We apply the information measure to both sides of this equation, because applying a measure to equal sets will give equal measures. Therefore, we have  $H(X_1, X_2, \dots, X_n) = \mu^* \left( (X_1 - X_2) \cup (X_2 - X_3) \cup \dots \cup (X_{n-1} - X_n) \cup (X_n - X_1) \right) + I(X_1; X_2; \dots; X_n)$ .

□

**Lemma 3.** *Adjacent set expressions are disjoint: If  $k = j \oplus 1$  or  $j = k \oplus 1$ , then  $\mu^* \left( \bigcap_{i \in W} (X_i - X_{i \oplus 1}) \right) = 0$*

*Proof.* First,  $j, k \in W$  implies

$$\bigcap_{i \in W} (X_i - X_{i \oplus 1}) \subseteq (X_j - X_{j \oplus 1}) \cap (X_k - X_{k \oplus 1})$$

And, for  $k = j \oplus 1$

$$\begin{aligned} (X_j - X_{j \oplus 1}) \cap (X_k - X_{k \oplus 1}) &= (X_j - X_k) \cap (X_k - X_{k \oplus 1}) \\ &= X_j \cap X_k^c \cap X_k \cap X_{k \oplus 1}^c \\ &= \emptyset \end{aligned}$$

with a symmetric argument holding for  $j = k \oplus 1$ .

Therefore,

$$\begin{aligned} \bigcap_{i \in W} (X_i - X_{i \oplus 1}) &\subseteq \emptyset, \text{ so} \\ \bigcap_{i \in W} (X_i - X_{i \oplus 1}) &= \emptyset, \text{ so} \\ \mu^* \left( \bigcap_{i \in W} (X_i - X_{i \oplus 1}) \right) &= \mu^*(\emptyset) = 0. \end{aligned}$$

□

**Lemma 4.** *Non-adjacent set expressions are d-separated: If  $k \neq j \oplus 1$  and  $j \neq k \oplus 1$ , then  $\mu^* \left( \bigcap_{i \in W} (X_i - X_{i \oplus 1}) \right) = 0$ .*

*Proof.* If  $k \neq j \oplus 1$  and  $j \neq k \oplus 1$ , it implies that  $j, k, j \oplus 1$ , and  $k \oplus 1$  are four distinct values, corresponding to four distinct variables in the cycle.

We are interested in whether  $X_j$  and  $X_k$  are d-separated by  $X_{j \oplus 1}, X_{k \oplus 1}$ .

There are exactly two paths between  $X_j$  and  $X_k$ , neither path containing a collider.  $X_{j \oplus 1}$  is on the opposite path from  $X_{k \oplus 1}$ . Therefore, conditioning on  $X_{j \oplus 1}$  and  $X_{k \oplus 1}$  will d-separate  $X_j$  and  $X_k$ .

More formally, this d-separation condition applies to  $\mu^* \left( \bigcap_{a \in A} a - \bigcap_{b \in B} b \right) = 0$  for any  $A, B$  such that

$$\{X_j, X_k\} \subseteq A \text{ and } \{X_{j \oplus 1}, X_{k \oplus 1}\} \subseteq B.$$

We can rewrite  $\mu^* \left( \bigcap_{i \in W} (X_i - X_{i \oplus 1}) \right) = \mu^* \left( \bigcap_{i \in W} X_i - \bigcap_{i \in W} X_{i \oplus 1} \right)$ , where  $\{X_j, X_k\} \subseteq \bigcap_{i \in W} X_i$  and  $\{X_{j \oplus 1}, X_{k \oplus 1}\} \subseteq \bigcap_{i \in W} X_{i \oplus 1}$ .

That allows us to apply the d-separation condition, giving  $\mu^*\left(\bigcap_{i \in W} (X_i - X_{i \oplus 1})\right) = 0$ .

□

Now we can give the proof of Theorem 2, which we restate for convenience:

**Theorem 2.** *A probability distribution  $P$  that factorizes as an  $n$ -variable cycle, according to  $p(x_1, x_2, \dots, x_n) = p(x_1|x_2)p(x_2|x_3) \dots p(x_{n-1}|x_n)p(x_n|x_1)$ , has mutual information  $I(x_1; x_2; \dots; x_n) = 0$ .*

*Proof.* We consider a probability distribution  $P$  that factorizes according to  $p(x_1, x_2, \dots, x_n) = p(x_1|x_2)p(x_2|x_3) \dots p(x_{n-1}|x_n)p(x_n|x_1)$ . We will write the joint entropy in two forms, in Lemma 1 and Lemma 2.

It remains to show that  $\mu^*((X_1 - X_2) \cup (X_2 - X_3) \cup \dots \cup (X_{n-1} - X_n) \cup (X_n - X_1)) = H(X_1|X_2) + H(X_2|X_3) + \dots + H(X_{n-1}|X_n) + H(X_n|X_1)$ , which will allow us to conclude that  $I(X_1; X_2; \dots; X_n) = 0$ .

To do this, we will apply the inclusion-exclusion formula to  $\mu^*((X_1 - X_2) \cup (X_2 - X_3) \cup \dots \cup (X_{n-1} - X_n) \cup (X_n - X_1))$ , which we can do because it is a measure on a union of sets.

Informally, the inclusion-exclusion formula will produce many terms. The terms containing one set difference, correspond to the conditional entropy terms. The higher-order terms containing an intersection of at least two set differences, will fall into one of two cases, but in both cases the term will equal zero. The first case we handle in Lemma 3 and the second case we handle in Lemma 4.

To be more formal, let  $\mathbb{W}$  be the power set on  $\{i \in \mathbb{N}, i < n\}$ . The inclusion-exclusion formula gives us a sum of terms, where each term is  $\mu^*\left(\bigcap_{i \in W} (X_i - X_{i \oplus 1})\right)$  for some  $W \in \mathbb{W}$ . We can ignore the negative signs on higher-order terms because we will soon show these terms are all zero.

The higher-order terms are those where  $|W| \geq 2$ . For these terms, we consider  $\mu^*\left(\bigcap_{i \in W} (X_i - X_{i \oplus 1})\right)$

where  $\oplus$  indicates addition mod  $n$ . Since  $|W| \geq 2$ , without loss generality we pick any two distinct elements of  $W$  and call them  $j, k$ . If  $k = j \oplus 1$  or  $j = k \oplus 1$ , we apply Lemma 3 to show the term is zero. If  $k \neq j \oplus 1$  and  $j \neq k \oplus 1$ , we apply Lemma 4 to show the term is zero.

We have now shown that all the terms with  $|W| \geq 2$  are zero, which leaves only the terms with  $|W| = 1$ , giving  $\mu^*((X_1 - X_2) \cup (X_2 - X_3) \cup \dots \cup (X_{n-1} - X_n) \cup (X_n - X_1)) = H(X_1|X_2) + H(X_2|X_3) + \dots + H(X_{n-1}|X_n) + H(X_n|X_1)$ . Applying Lemma 1 and Lemma 2 leaves us with  $I(X_1; X_2; \dots; X_n) = 0$ , as desired. □