Causality Beyond Association or Covariation: Hierarchy, Asymmetry, and Developing over Time

Xinyi Yang
Department of Automation
Tsinghua Universtiy
xy-yang21@mails.tsinghua.edu.cn

Abstract

The question "Is causality really different from covariation/association?" runs through the whole study about modeling causal inference. Firstly it asks about the organization of causal inference, which is thought to be hierarchical and convariation computation works further in the process in this essay. Secondly it requires representable difference to make theoretical distinction, which can be offered by asymmetry of cause and effect. Only when difference can be perceived can it really make sense. So it finally implicates the question "If the difference really exists, how do we distinguish them in real life?" Several forms of cues such as temporal order, prior knowledge, statistical relations and intervention can be acquired from the world and help us with causal inference.

1 Introduction

Causality is the abstract notion of cause and effect derived from the perceived environment, which is also the foundation of human understanding.[28] When investigating causality, a basic question about difference or further, relation of causality and association/covariation has been under constant research. There is a variety of situations where association exists between two things but causality does not, *e.g.*, the indicator of a barometer and the storm. According to Sloman and Lagnado [19], a causal relation is not merely an association in the sense that it is not a representation of a mere correlation but rather a representation of something more enduring in nature: Natural laws dictate what causes what regardless of any learning.

In this essay to answer the question "Is causality really different from association/covariation?", the distinction made between *structure* and *strength* will be discussed first in Sec. 2, which aims to explain the implicit hierarchy in causality induction and different focus of associative theories during the inducing process. After figuring out the role of association in hierarchical causality, what makes casual relations distinct from associative relations on most views is that casual relations asymmetrically support intervention (sufficient changes to a cause will also change the effect; the converse is not true), will be elaborated in Sec. 3. However, theoretically we can find obvious difference between association and causality, but young children or even adults sometimes get confused when distinguishing between them. For example, a person reported that his sneeze had caused the New York blackout. But the interactions with the world afford people a variety of cues which can be incorporated to develop their casual inference capability. Sec. 4 will present 4 kinds of main cues. At last, I will combine the characteristics presented before to give out the conclusion and some thoughts about what can be captured by current models and what can not be in Sec. 5.

2 Hierarchical Relation between Structure and Strength

A clear distinction between *structure* and *strength* has been made in Lagnado et al. [13]. "Structure" concerns the qualitative causal relations that hold between variables, *e.g.*, whether smoking causes lung cancer, while "Strength" concerns the quantitative aspects of these relations, *e.g.*, to what degree does smoking cause lung cancer.

Conceptually, the question of structure is more basic than that of strength because one needs to know or assume the existence of a link before one can estimate its strength. This idea receives intuitive support. People often have knowledge about what causes what but little idea about the strength of these relations. Moreover, we seek to establish whether causal relations exist before trying to assess how strong they are. Some researchers have made a natural conjecture that this priority of structure over strength is likewise marked in human cognition[23, 24]. Take the example of smoking again, most of us believe that smoking causes cancer while we know little about the strength of this relation. Furthermore, it is demonstrated that participants in Tenenbaum and Griffiths [21] actually assess the relations to which the evidence supports the existence of a causal link rather than the strength of that link.

The influence of the idea that causal cognition is grounded in qualitative relations and the neglect of structure has led to an overestimation of the importance of statistical data at the expense of other key cues in causal learning.[13] It mainly reflects in associative theories, which focus on learning mechanisms (*e.g.*, Shanks and Dickinson [17]) that encode the strength of covariation between cues and outcomes but are insensitive to the structure distinction between causes and effects. As a consequence, they are incapable of distinguishing between associations that link spurious relations from true causal relations (*e.g.*, the indicator of a barometer and the storm, the atmospheric pressure and the storm), *i.e.*, between direct and indirect causal relations that are generated by hidden causal events[25]. However, the distinction between structure and strength is captured more formally in the causal Bayes net framework[9, 15, 16], in which the structure of a set of variables is represented by a graph [11] (Fig. 1), with the strength of these links captured in the parameterization of the graph.

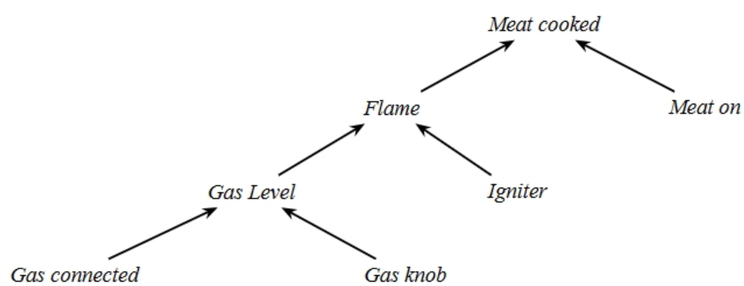

Figure 1: The casual system of a gas grill used to cook meat.

3 Asymmetry of Cause and Effect

The biggest difference between casual relations and associative relations is that the former have direction but the latter do not. When A is related to B, we can say that B is also related to A. But when A causes B, we can not say that B can also lead to A. In experiments of Fenker et al. [8], they presented participants with pairs of words one after another, describing events that referred to either a cause (*e.g.*, spark) or an effect (*e.g.*, fire) and manipulated the temporal order of word presentation and the question participants had to respond to. The results revealed that people were faster to verify a causal relation when presented in the cause-to-effect order than when presented in the effect-to-cause order. However, no such asymmetry was observed with questions referring to the associative relation. More evidence has also been found among children. Hong et al. [12] has shown that children performed better on cause-effect inferences than on effect-cause inferences.

This asymmetry points to a problem of Bayesian inference: There is indeed a direction presented in causal maps and variables contain different amounts of information about one another, *i.e.*, P(effect | cause) is different from P(cause | effect) on a particular occasion, but reasoning actually does not systematically favor one over the other[4]. It suggests that causal inference is not merely a way of representing and updating probabilities.

4 Learning from Cues Afforded by the World

People are active agents immersed in a dynamic physical world. Their experiences about events and relations among events when interacting with the world provide a variety of cues to the causal structure. These cues can sometimes mislead people to infer about causal relations, but they can also be combined to construct and update causal models. 4 main kinds of cues will be presented below about their individual and combined influence on our casual inference.

4.1 Temporal Order

Little kids often make strange causal conclusions about things observed. For example, the TV shows pictures and makes sounds every time the mother is standing in front of it, and then the kid walks to the TV trying to make it start to play. A bird flew by while the kid was laughing, so he or she laughs on purpose trying to make a bird appear. In summary, the kid feels the existence of causality between two events when these two events happened very close both in time and space. The temporal order in which events occur provides a fundamental cue to causal structure. The simplest example is that causes occur before (or possibly simultaneously with) their effects, so if one knows that A occurs after B one can be sure that A is not a cause of B. There is also evidence from Shanks et al. [18] that judged causal strength decreases with increased temporal delays unless people have a good reason to expect a delay.

It is doubtless that the temporal order of events is an imperfect cue to causal structure, which often gets us confused when distinguishing between contiguous events' associative relations and casual relations. However, it will often yield a good cue to causal structure especially if it is combined with other cues[13], which will be explained after presenting other cues.

4.2 Prior Knowledge

Coherence with prior knowledge is a potent cue to causal structure. In Griffiths and Tenenbaum [10], the kind of prior knowledge that is relevant to causal induction is divided into three categories, which have been proved great influence on causal induction in several studies:

- Ontology: Information about the types of entities, properties, and relations that arise in a domain. (*e.g.*, Netwon's theory of physics picked out the critical variables for thinking about the motion of objects, like mass, velocity, and acceleration.)
- Plausible Relations: Knowledge of the types of entities in a domain can provide quite specific information about the plausibility of causal relationships. (*e.g.*, Newton precisely laid out the kinds of forces by which the properties of one object can influence those of another.)
- Functional Form: (*e.g.*, how the velocity of one object depends on its mass and the mass and velocity of another object with which it collides.)

Moreover, according to Lagnado et al. [13], prior knowledge may be specific when we have already learned about a causal relation, but prior knowledge can also be abstract and hypothetical. It is consistent with both the tradition of Hume (covariation-based approaches) and Kant (mechanism-based approaches). The former ones characterize human causal induction as the consequence of a domain-general statistical sensitivity to covariation between cause and effect.[5, 6] And the latter ones focus on the role of prior knowledge about the mechanisms by which causal force can be transferred.[1] Several studies show that children do take prior knowledge into account when making causal inferences (*e.g.*, Tenenbaum et al. [22] showed that children made different inferences when they were told beforehand that blickets were rare or common in a backward blocking task).

In the computational level, a few cases where formal accounts of the integration of prior knowledge and data have been explored (*e.g.*, Alloy and Tabachnik [2]), but they have focused on just one aspect of prior knowledge. Constraint-based structure-learning algorithms are particularly limited in their use of prior knowledge. As these algorithms are defined, they use only a weak form of prior knowledge that particular causal relationships do or do not exist.[10] However, a strategy through a prior distribution on the parameters in Bayesian models can be used to incorporate the effects of prior knowledge in parameter estimation, allowing the expectations of learners to influence their inferences about the strength of causal relationships (*e.g.*, Lu et al. [14]).

4.3 Statistical Relations

The key idea of Hume's analysis of causation is that people are exposed to patterns of data and directly, values of variables, and they can learn statistical relations from observation and imitation. One of the simplest forms of statistical relations is the covariation between two events (*e.g.*, smoking increases the probability of lung cancer.) But as stated before, the existence of a stable covariation between two events indicates a possible causal relation exists, but does not reveal whether one causes another. Bayesian networks provide a straightforward representation of relations between variables and also provide a conditional one.[13]

4.4 Intervention

Informally, intervention involves imposing a change on a variable in a causal system from outside the system. It is through our actions and manipulations of the environment around us that we acquire our basic sense of causality, which is called learning from exploration and experimentation. Several studies have compared learning through intervention with learning through observation, with both adults and children. As the first experiment in Sobel and Kushnir [20] has shown, learners were better at learning causal models when they observed intervention data that they had generated, as opposed to observing data generated by another learner. Several studies are trying to incorporate intervention with causal inference, for example, Zhang et al. [27] presents a causal inference framework to improve Weakly-Supervised Semantic Segmentation through intervention.

4.5 Developing Casual Inference from Combinations of Several Cues

These cues described above usually do not work separately, but there is an obvious relation among them and they can also be combined to construct and update causal models. Firstly, as for the relation among them, on the one hand, information acquired during the process of learning from statistical relations and intervention can always be included in prior knowledge. On the other hand, temporal order and prior knowledge can assist us in making appropriate interventions for more apparent changes. Secondly, there are various ways to combine these cues in causal inference. The most common is the combination of temporal order and statistical data that is always taken as basic inputs to the inference process of psychological models. Moreover, prior knowledge can address the possible mismatch between causal order and learning order. When people find a good reason according to prior instructions or knowledge to expect a time delay, judged causal strength will not decrease just because of increased temporal delays.[3]

Furthermore, the combinations of these cues not only affect our casual judgment in the moment, but also develop our casual inference systems over time. The intuitive evidence can be found just from our own growing process, such as not taking standing in front of the TV as the cause of TV's working after we have realized the existence of remote control or seen TV playing with a person staying on the couch all the time. Using some clever paradigms, Cummins et al. [7] verified some intuitions: The number of disablers influenced the acceptability of modus ponenes inference (allows one to infer a conclusion from a premise and its implication), and the number of alternative causes influenced the acceptability of affirming the consequent inference (a formal fallacy of inferring the converse from the original statement). It reminds me that the reason why little children or even adults get confused when distinguishing between associative relations and casual relations is that they do not have enough knowledge reserve about some disablers and alternative causes. And most disablers and alternative causes are acquired from the cues afforded by the world.

5 Discussion

In summary, the hierarchical relation between structure and strength and the focus of association/covariation decide its place in casual inference. Causality's characteristic of asymmetry distincts it from simple associative relations. There are several kinds of cues afforded by the world from which our casual inference systems develop to better perceive causality and distinguish between associative relations and causal relations. Current models like Bayesian inference models can capture some key points like separate representation of structure and strength, incorporation with temporal order, prior knowledge and statistical data, but they are still struggling in representing asymmetry. On top of that, recent RL models have produced a wide body of research about agents learning how to play, but the majority of model-free RL methods still have great difficulty transferring learned policies to new environments with consistent underlying mechanics but some dissimilar surface features [26].

Therefore, there is still a long way towards a computational model that captures causal relation and causal structure. The distinct advantage of intervention revealed in studies occurs to me that, except incorporating cues to casual structure in models, modeling the ability to actively create valid intervention, make intervention and learn cues from the results of intervention may be a considerable option.

References

[1] W-k Ahn and C Kalish. The role of covariation vs. mechanism information in causal attribution. *Cognition and explanation*, pages 227–253, 2000.

[2] Lauren B Alloy and Naomi Tabachnik. Assessment of covariation by humans and animals: the joint influence of prior expectations and current situational information. *Psychological review*, 91(1):112, 1984.

[3] Marc J Buehner and Jon May. Knowledge mediates the timeframe of covariation assessment in human causal induction. *Thinking & Reasoning*, 8(4):269–295, 2002.

[4] Nancy Cartwright. What is wrong with bayes nets? *The monist*, 84(2):242–264, 2001.

[5] Patricia W Cheng and Laura R Novick. A probabilistic contrast model of causal induction. *Journal of personality and social psychology*, 58(4):545, 1990.

[6] Patricia W Cheng and Laura R Novick. Covariation in natural causal induction. *Psychological Review*, 99(2):365, 1992.

[7] Denise D Cummins, Todd Lubart, Olaf Alksnis, and Robert Rist. Conditional reasoning and causation. *Memory & cognition*, 19:274–282, 1991.

[8] Daniela B Fenker, Michael R Waldmann, and Keith J Holyoak. Accessing causal relations in semantic memory. *Memory & cognition*, 33:1036–1046, 2005.

[9] Clark N Glymour. *The mind's arrows: Bayes nets and graphical causal models in psychology*. MIT press, 2001.

[10] Thomas L Griffiths and Joshua B Tenenbaum. Theory-based causal induction. *Psychological review*, 116(4):661, 2009.

[11] Christopher Hitchcock. Causal models. 2018.

[12] Li Hong, Zheng Chijun, Gao Xuemei, Gao Shan, and Lin Chongde. The influence of complexity and reasoning direction on children's causal reasoning. *Cognitive Development*, 20(1): 87–101, 2005.

[13] David A Lagnado, Michael R Waldmann, York Hagmayer, and Steven A Sloman. Beyond covariation. *Causal learning: Psychology, philosophy, and computation*, pages 154–172, 2007.

[14] Hongjing Lu, Alan L Yuille, Mimi Liljeholm, Patricia W Cheng, and Keith J Holyoak. Bayesian generic priors for causal learning. *Psychological review*, 115(4):955, 2008.

[15] Judea Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan kaufmann, 1988.

[16] Judea Pearl et al. Models, reasoning and inference. *Cambridge, UK: CambridgeUniversityPress*, 19(2):3, 2000.

[17] David R Shanks and Anthony Dickinson. Associative accounts of causality judgment. In *Psychology of learning and motivation*, volume 21, pages 229–261. Elsevier, 1988.

[18] David R Shanks, Susan M Pearson, and Anthony Dickinson. Temporal contiguity and the judgement of causality by human subjects. *The Quarterly Journal of Experimental Psychology*, 41(2):139–159, 1989.

[19] Steven A Sloman and David Lagnado. Causality in thought. *Annual review of psychology*, 66: 223–247, 2015.

[20] David M Sobel and Tamar Kushnir. The importance of decision making in causal learning from interventions. *Memory & Cognition*, 34:411–419, 2006.

[21] Joshua Tenenbaum and Thomas Griffiths. Structure learning in human causal induction. *Advances in neural information processing systems*, 13, 2000.

[22] Joshua B Tenenbaum, Charles Kemp, Thomas L Griffiths, and Noah D Goodman. How to grow a mind: Statistics, structure, and abstraction. *science*, 331(6022):1279–1285, 2011.

[23] Michael R Waldmann. Knowledge-based causal induction. In *Psychology of Learning and Motivation*, volume 34, pages 47–88. Elsevier, 1996.

[24] Michael R Waldmann and Laura Martignon. Bayesian network models of causal learning. In *Twentieth Annual Conference of the Cognitive Science Society*, pages 1102–1107. Erlbaum, 1998.

[25] Michael R Waldmann and Jessica M Walker. Competence and performance in causal learning. *Learning & Behavior*, 33(2):211–229, 2005.

[26] Chiyuan Zhang, Oriol Vinyals, Remi Munos, and Samy Bengio. A study on overfitting in deep reinforcement learning. *arXiv preprint arXiv:1804.06893*, 2018.

[27] Dong Zhang, Hanwang Zhang, Jinhui Tang, Xian-Sheng Hua, and Qianru Sun. Causal intervention for weakly-supervised semantic segmentation. *Advances in Neural Information Processing Systems*, 33:655–666, 2020.

[28] Yixin Zhu, Tao Gao, Lifeng Fan, Siyuan Huang, Mark Edmonds, Hangxin Liu, Feng Gao, Chi Zhang, Siyuan Qi, Ying Nian Wu, et al. Dark, beyond deep: A paradigm shift to cognitive ai with humanlike common sense. *Engineering*, 6(3):310–345, 2020.