Cognitive Psychology 64 (2012) 93–125

Contents lists available at SciVerse ScienceDirect

Cognitive Psychology

journal homepage: www.elsevier.com/locate/cogpsych

Causal structure learning over time: Observations and interventions ☆

Benjamin M. Rottman [a,*], Frank C. Keil [b]

[a] Department of Hospital Medicine, University of Chicago, 5841 S. Maryland Ave., MC5000, Chicago, IL 60637, United States
[b] Department of Psychology, Yale University, P.O. Box 208205, New Haven, CT 06520, United States

A R T I C L E   I N F O

Article history:
Accepted 18 October 2011
Available online 7 December 2011

Keywords:
Causal structure
Causal learning
Time
Observation
Intervention

A B S T R A C T

Seven studies examined how people learn causal relationships in scenarios when the variables are temporally dependent – the states of variables are stable over time. When people intervene on *X*, and *Y* subsequently changes state compared to before the intervention, people infer that *X* influences *Y*. This strategy allows people to learn causal structures quickly and reliably when variables are temporally stable (Experiments 1 and 2). People use this strategy even when the cover story suggests that the trials are independent (Experiment 3). When observing variables over time, people believe that when a cause changes state, its effects likely change state, but an effect may change state due to an exogenous influence in which case its observed cause may not change state at the same time. People used this strategy to learn the direction of causal relations and a wide variety of causal structures (Experiments 4–6). Finally, considering exogenous influences responsible for the observed changes facilitates learning causal directionality (Experiment 7). Temporal reasoning may be the norm rather than the exception for causal learning and may reflect the way most events are experienced naturalistically.

1. Introduction

In our attempts to explain and act on the world around us we almost invariably need to reason about causal structures. Scientists study the causal relationships between the variables involved in

☆ Experiments 1 and 3 were presented at the 32nd Annual Conference of the Cognitive Science Society (Rottman & Keil, 2010). Experiments 4 and 5 were presented at the 33rd Annual Conference of the Cognitive Science Society (Rottman & Keil, 2011).

* Corresponding author. Present address: Department of Hospital Medicine, University of Chicago, 5841 S. Maryland Ave., MC5000, Chicago, IL 60637, United States.

E-mail address: benjaminrottman@uchicago.edu (B.M. Rottman).

doi:10.1016/j.cogpsych.2011.10.003

phenomena ranging from global warming to cognitive processes. Lay people understand phenomena like academic failure, unemployment, and loneliness in terms of complex webs of causes and effects (Lunt, 1988, 1989, 1991). Recently, cognitive scientists have investigated how people build their notions of causal structures. For example, how do people form a belief in the causal relationships between depression, anxiety, and insomnia; what causes what? Most prior work on causal structure learning has used events that are temporally independent, analogous to a between-subjects experimental design. For example, a clinical psychologist may try to learn the causal relationships between depression, anxiety, and insomnia, by observing 100 people who have different combinations of these disorders.

However, many if not most of our learning experiences involve repeatedly learning about one entity over time. For example, one might develop beliefs about the causal relationships between depression, anxiety, and insomnia, by observing a friend who experiences these disorders wax and wane over time, analogous to a within-subjects design. Furthermore, much of our causal learning involves learning about ourselves over time (e.g., do I really have an allergic reaction every time I eat wheat?). Given that we experience most events sequentially over time, this latter type of temporal or "within-subjects" situation may be the norm for causal learning.

To explore these different forms of causal learning, we will first introduce normative causal models for independent vs. temporally dependent scenarios; static and dynamic graphical causal models. Then, we describe different strategies people may use to learn causal structures for these two scenarios, and when manipulating vs. observing variables. Finally, we describe an approach to test whether people use a temporal strategy that is appropriate for learning causal structures in scenarios with temporally dependent events.

1.1. Static graphical causal models

1.1.1. Simulating causal models

Graphical causal models have mainly been developed to represent the causal relationships within a set of independent observations. We call this class of models "static." This section first describes how static causal models can be used to simulate a set of observations (e.g., 100 people who have different combinations of depression, anxiety, or insomnia). The next section describes how one can learn the causal structures that are most likely to have generated a set of observations.

A graphical causal model consists of a set of nodes, which represent events, and arrows between the nodes, which represent causal relationships. In order to simulate how the model functions, one must also know the "parameters" of the model. For every node that does not have any known causes, one needs to know its base rate. Additionally, for every node that does have direct causes, one must know its conditional probability given its direct causes.

Table 1 presents three causal structures representing possible causal relationships between depression, anxiety, and insomnia. Consider the common cause structure, $D \leftarrow A \rightarrow I$, which asserts that anxiety influences depression and insomnia. The parameters for this structure assert that an individual person has a 10% chance of having an anxiety disorder. Out of people who are anxious, 75% are also depressed, and out of people who are not anxious 25% are depressed. Additionally, out of people who are anxious 75% are insomniacs, and out of people who are not anxious 25% are insomniacs.

Knowing the parameters and the causal structure, one can determine the percent of the population that has each of the eight possible combinations of depression, anxiety, and insomnia – the joint probability distribution. For example, the percent of the population who has depression and anxiety but not insomnia, can be determined by taking the product of the base rate of people who have anxiety, 10%, the conditional probability of having depression given that one has anxiety, 75%, and the conditional probability of not having insomnia given that one has depression, 25%. The product, shown on the second line of the joint probability table for observations, is 1.9%.

One can also simulate the causal structure under various interventions. For example, suppose that one was interested in the prevalence of depression and insomnia within an otherwise typical population that was taking anti-anxiety medications so that no one was anxious (A is set to 0 in the bottom section in Table 1). Under the graphical causal model framework, a perfect intervention severs all the relationships from other causes of the manipulated variable. For the common cause structure, after A