Methodological Problems in Cognitive Psychology

David Danks
Institute for Human & Machine Cognition
January 10, 2003

(At Least) Three Ways Bayes Nets Can Matter for Cognitive Psych

1. Novel theoretical possibilities
2. More detailed specification of the experiments being performed
3. Novel experimental designs

Novel Theories

- Causal learning community has focused on simple parameter estimation
  - Given a series of cases, people must rate the “causal strength”
  - Leads to relatively simple psychological theories
  - There are often multiple plausible theories that can fit a particular dataset

Novel Theories

- More specifically, causal learning is just parameter estimation for the graph:

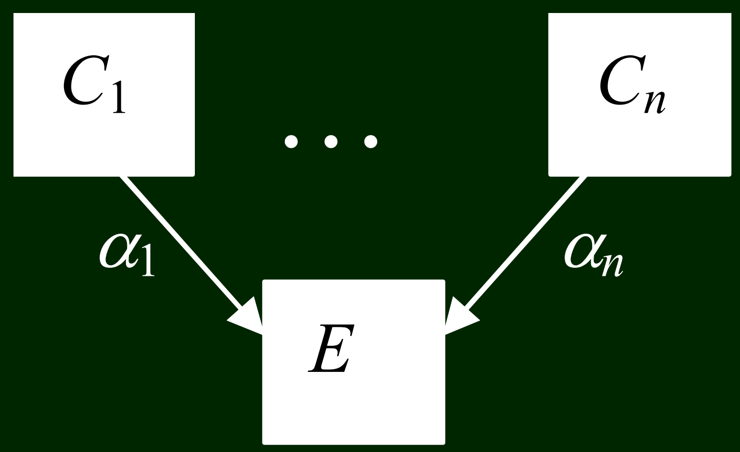

  - Absence of causation is just $\alpha_l = 0$

Novel Theories

- Using Bayes nets suggests novel theories that have not previously been considered in the causal learning community
  - Specifically, consider the possibility that people are also learning structure, not just parameters

- This shift leads to theories such as:
  - Bayesian updating over structures (Tenenbaum)
  - Using interventions to learn structure (Gopnik)

Clarifying Experiments

- Typical causal learning experiments just specify the probability distribution over the causes and effect
- This formulation can obscure the causal structure to be learned

Clarifying Experiments

- Instead, we can use a Bayes net to specify the causal structure (and then make sure the probability distribution fits the graph)

Advantages:

1. Can easily determine whether two experiments are qualitatively different
2. Can better determine the range of situations tested in experiments

Clarifying Experiments

- For example, we find that almost all two-cause experiments used one of these two causal structures:

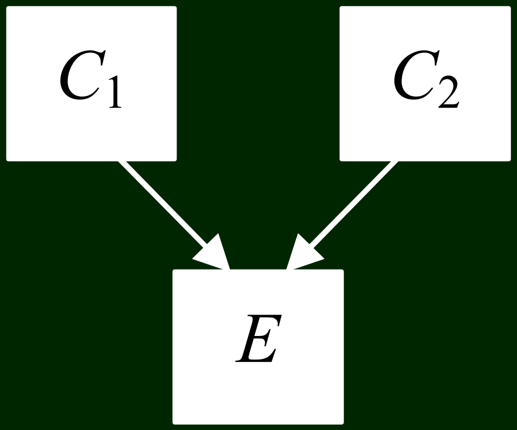

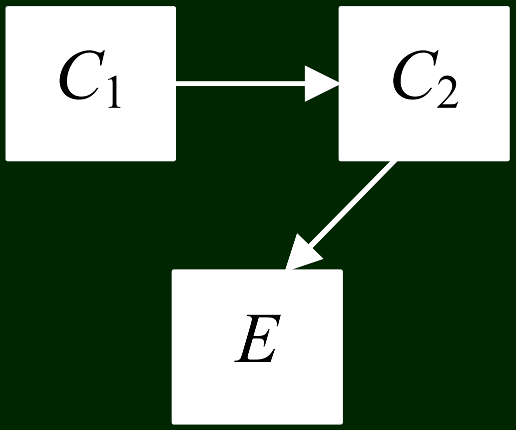

Novel Experiments

- Typical experimental design:
  - Pre-define the causes and effect to fit the probability distribution
  - Provide either a series (or summary) of observations of cases with values for the (earlier defined) causes/effect
  - Ask people to provide a rating of “causal strength”
  - Provide no feedback to the subject until the end of the experiment

Novel Experiments

- Using Bayes nets suggests changes to these features of the experimental design:
  - "Pre-define the causes and effect to fit the probability distribution" ➔
    Determine the influence of these definitions by using the same probability distribution with different tags for the variables (Waldmann)

Novel Experiments

- Using Bayes nets suggests changes to these features of the experimental design:
  - "Provide either a series (or summary) of observations of cases with values for the (earlier defined) causes/effect" ➔
    Allow the subject (or the experimenter) to intervene on the variables, rather than simply watching them (Gopnik)

Novel Experiments

- Using Bayes nets suggests changes to these features of the experimental design:
  - "Ask people to provide a rating of 'causal strength'" ➔
    Have people choose among graphs that represent causal structure (Steyvers); or
    Have people make a forced choice about whether "*X* causes *Y*" (Danks)

Novel Experiments

- Using Bayes nets suggests changes to these features of the experimental design:
  - “Provide no feedback to the subject until the end of the experiment” →
    Allow the subjects to see the results of their interventions (Gopnik)

Other Applications?

- Use of regression (e.g., in social psychology) Use of latent variable analyses (e.g., PCA)
  - Bayes net search methods are reliable in more contexts than these two methods – how should they be used in practice?
- Application of psychological research
  - Can we use Bayes nets to model the effects of policies (social and personal) based on psychological research?

Other Applications?

- Decision theory that incorporates causation
  - In the standard theory, decisions are made without using causal structure. Actions look like interventions – can Bayes nets inform theories and experiments of decision-making?
- Theory/belief change in children
  - Can this process be modeled using Bayes nets? Are novel experiments suggested by that framework?

Constraint-Based & Bayesian Learning

- Bayesian learning – it's what Chris talked about this morning
  - Establish a probability distribution over graphs (& a distribution over parameters for each graph)
  - Update the distributions based on the observed data using Bayes' Theorem
  - In practice, usually approximated using some type of greedy search algorithm

Constraint-Based & Bayesian Learning

- Constraint-based learning:
  - Examine one's data to determine the independencies and associations in that data
  - Determine the set of graphs that could possibly have produced data with that pattern

Constraint-Based & Bayesian Learning

- Example of constraint-based learning:

$$\left\{ \begin{array}{l} X \cancel{\perp\!\!\!\perp} Y,Z \\ Y \cancel{\perp\!\!\!\perp} Z \\ X \perp\!\!\!\perp Z \mid Y \end{array} \right\} \Rightarrow \left\{ \begin{array}{l} \boxed{X} \to \boxed{Y} \to \boxed{Z} \\ \boxed{X} \leftarrow \boxed{Y} \to \boxed{Z} \\ \boxed{X} \leftarrow \boxed{Y} \leftarrow \boxed{Z} \end{array} \right\}$$

Intervention vs. Observation

- We all agree that causal information is useful because it enables us to predict the outcomes of our interventions
- So for Bayes nets to be good models of causation, they must be good models of the effects of interventions
- How do they do that?

Intervention vs. Observation

- To model an intervention in a Bayes net, you simply remove all edges into the “intervened-upon” variable, and leave all other edges intact
- Thus, to determine whether an intervention on *X* changes *Y*, you just change the graph according to the above rule, and then see whether there is a path from *X* to *Y* *(note: not exactly right, but close enough)*

Intervention vs. Observation

- Consider the simplest case of causation: *Light Switch* → *Light Bulb*
- An intervention on *Light Switch* doesn't change the graph at all (since there are no edges into *Light Switch*), and so that intervention **will** matter for *Light Bulb*
- On the other hand, intervening on *Light Bulb* does change the graph (removing the edge), and so the intervention **doesn't** matter for *LS*

Intervention vs. Observation

- A large (abstract) causal structure:

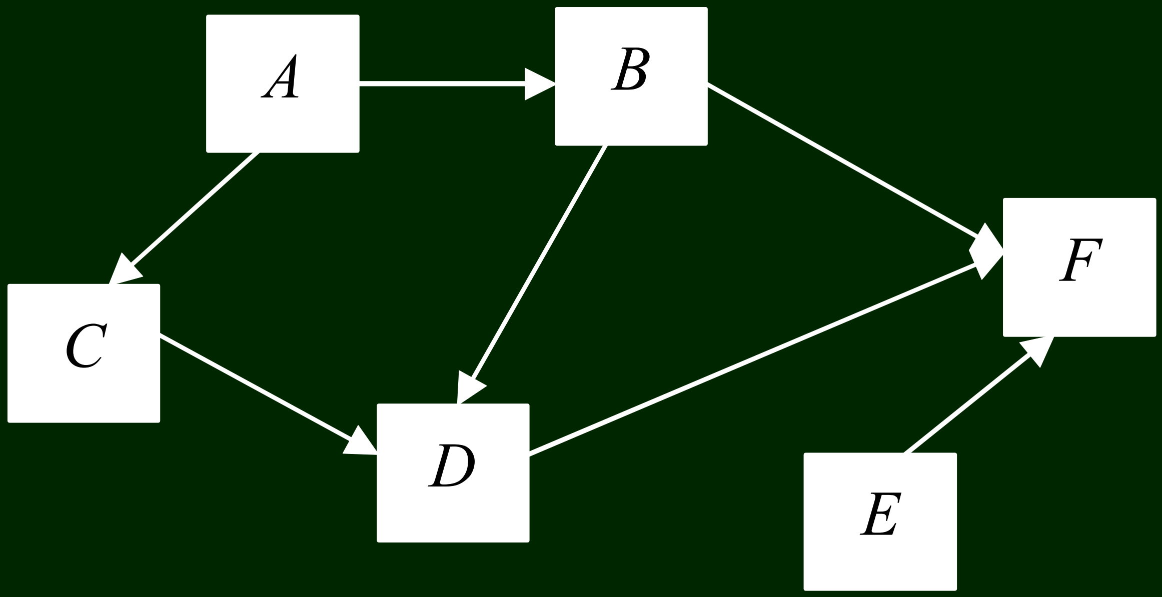

Intervention vs. Observation

- A large (abstract) causal structure:

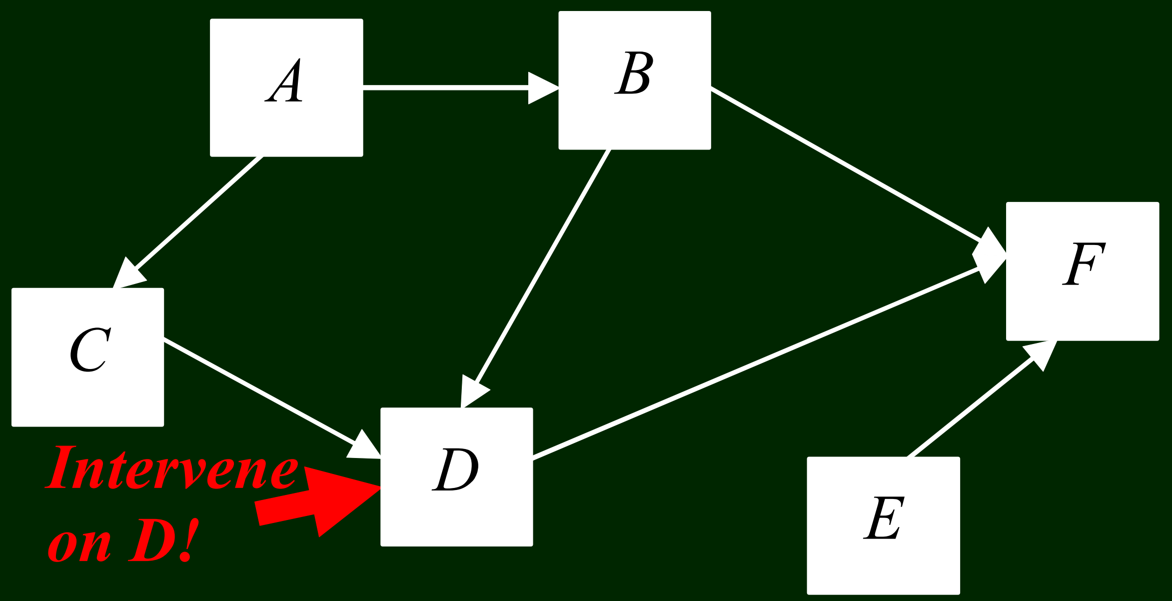

Intervention vs. Observation

- A large (abstract) causal structure:

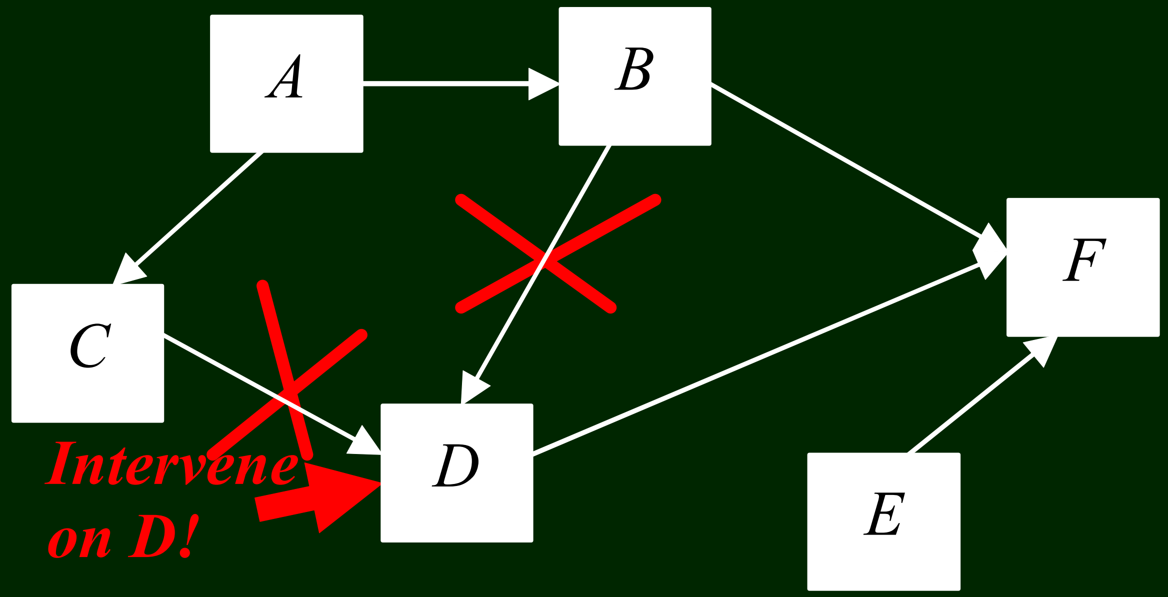

Intervention vs. Observation

- A large (abstract) causal structure:

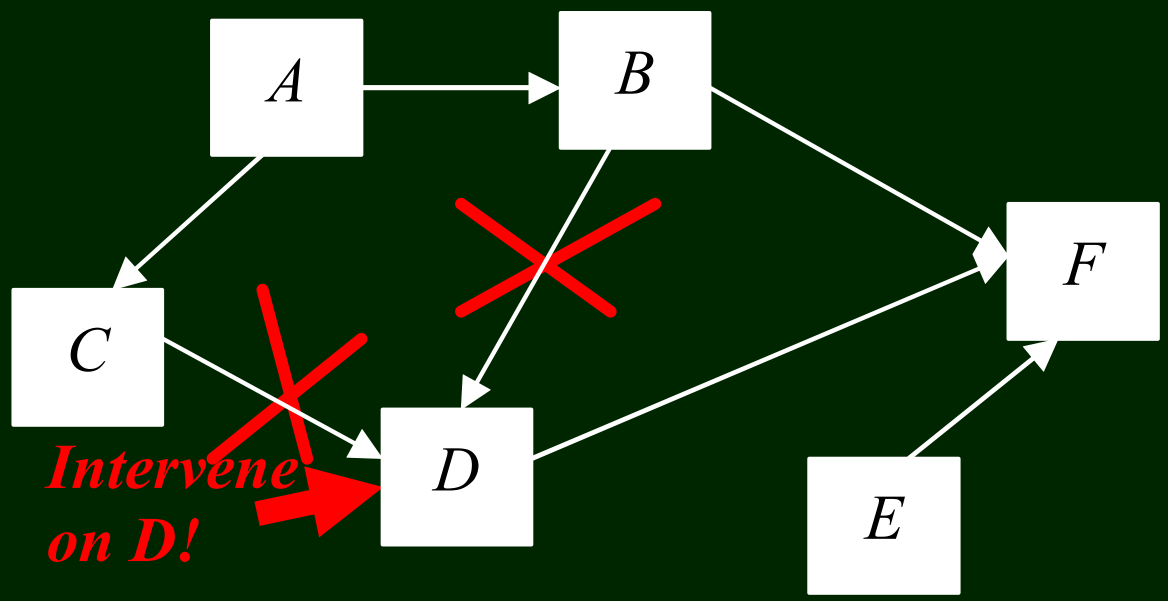

- The intervention on *D* affects only *F*!