

# Causal learning and inference in humans and machines

Friday December 7, 2001

Alison Gopnik, Tamar Kushnir, Kevin Murphy, *UC Berkeley*

Josh Tenenbaum, Tom Griffiths, *Stanford University*

<http://www-psych.stanford.edu/~jbt/causal-workshop.html>

Both cognitive science and artificial intelligence have recently begun to confront questions of causality: How do we learn the causal texture of our environment? What role should causality play in models of reasoning and conceptual structure? This workshop aims to explore the potential connections between research on causality in these two fields. There has already been some productive cross-fertilization: the development of causal Bayes nets in the AI community has often had a strong psychological motivation, and recent work by several groups in cognitive science has argued that fundamental aspects of how people learn and reason about causes may be best explained by theories based on Bayes nets. Here we will focus on three specific research areas where the state of the art in AI and cognitive science meet:

1. The roles of observation and intervention in causal learning
2. Discovering hidden causes
3. The causal structure of theories and concepts

Each of these three topics will be the focus of a dedicated workshop session, following an introductory session laying out some general issues. Each session will feature presentations by researchers from artificial intelligence and cognitive science, followed by a liberal discussion period. Time has also been allotted at the end of the day for brief, informal presentations by other workshop participants.

In parallel with the specific focus for each session, discussions will consider some of the general challenges facing interdisciplinary research on causality:

- The behavior of simple Bayes net models often seems to correspond qualitatively to patterns of human reasoning (e.g., the explaining-away phenomenon). But to what extent can Bayes nets provide specific insights into the mechanisms of human causal learning and inference, particularly in large, complex domains?
- Do approximation methods for inference and learning on large Bayes nets have anything to do with human cognitive processes?
- Can the way that people reason about causality guide the development of more sophisticated computational tools for causal discovery?

# Causal learning and inference in humans and machines

Friday December 7, 2001

Organizers: Alison Gopnik, Tamar Kushnir, Kevin Murphy, Josh Tenenbaum & Tom Griffiths

Friday morning session: 7:45am–10:30am

## Morning session 1: Introduction and overview

7:45am **Introduction**, *Alison Gopnik and Josh Tenenbaum*

8:00am **Human causal learning and inference: a graphical models perspective**, *Josh Tenenbaum*

8:30am Discussion

8:45am *break*

## Morning session 2: The roles of observation and intervention in causal learning

9:00am **A decision theoretic approach to active learning of causal structure**,  
*Kevin Murphy and Mark Paskin*

9:25am **Spinning wheels and bossy ones: Children, causal structure and the calculus of intervention**, *Laura Schulz*

9:50am **Models of human causal reasoning in on-line and active learning settings**, *Mark Steyvers*

10:15am Discussion

Friday afternoon session: 4:00pm–7:00pm

## Afternoon session 1: Discovering hidden causes

4:00pm **Improved learning models via discovery of hidden variables**,  
*Gal Elidan, Noam Lotner, Nir Friedman and Daphne Koller*

4:20pm **Human learning of causal structure with hidden variables**, *David Danks*

4:40pm **Causes, coincidences, and theories**, *Tom Griffiths*

5:00pm Discussion

5:30pm *break*

## Afternoon session 2: Causal structure of theories and concepts

5:45pm **A Bayes net theory of categorization and feature prediction**, *Bob Rehder*

6:05pm **Reasoning from factual and counterfactual causes**, *Dan Yarlett*

6:25pm Discussion