Modeling Discovery Strategies

Felipe Romero
Washington University in St. Louis

Frederick Eberhardt
Washington University in St. Louis

Tamar Kushnir
Cornell University

Overview

While much of the psychological literature on causal learning has focused on learning from *passive observational* data, we here use a simple experimental paradigm adapted from Sobel & Kushnir (2006) where the participant can *actively* intervene to learn the causal relations. The aim is to develop a better understanding of the learning strategies people use to solve causal discovery problems. To this end we consider three different models of how a participant may intervene to test a causal hypothesis: an optimal Bayesian strategy, a simple heuristic, and a random strategy for baseline comparison. Each of these strategies is compared to observed participant behavior. We conclude that human learning in this set-up is better modeled by simple strategies. People's interventions are neither optimal nor consistent with standard models in the psychological literature.

The experiment raises quite general questions about the relation between normative computational accounts of learning and the present descriptive evidence: If observed performance does not follow standard normative accounts, why do participants proceed in the way they do? If their behavior is in some sense rational, what is the nature of this rationality? Alternatively, do we have to conclude that there is no general normative account of the behavior?

The experiment also raises methodological questions: If we want to understand more about people's learning strategies, how can we search the space of possible strategies? The poster presents our results and ideas so far.

Experiment

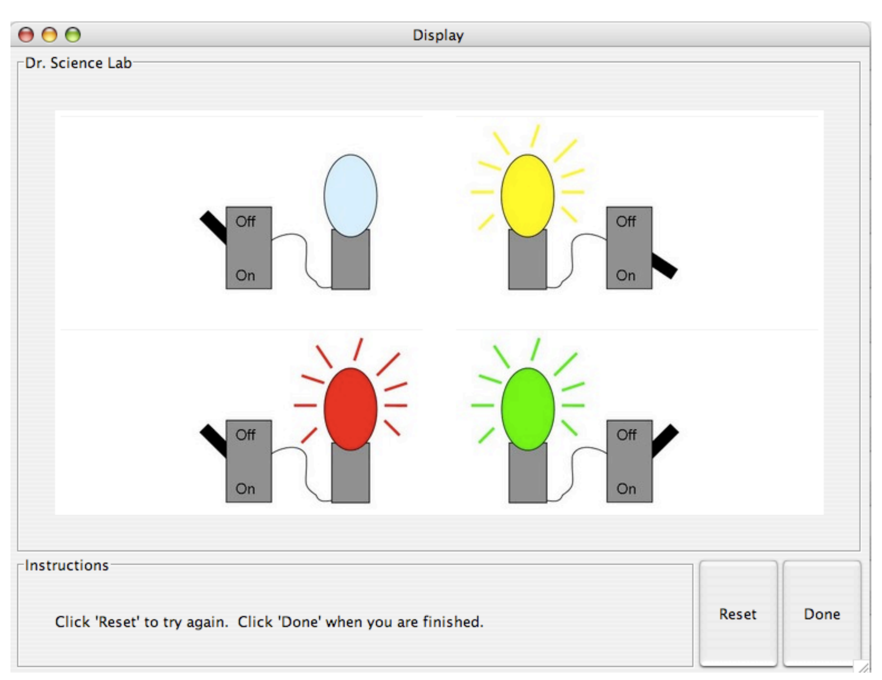

Game (following Sobel & Kushnir, 2006):

•Participants are told that invisible wires connect the light bulbs: e.g. if the yellow light is switched on and it is connected by a wire to the green light, then the green light also comes on. The effects are instantaneous.

•The wires are known to be faulty and only work 80% of the time. If lights are connected by wires in a chain, then the likelihood of the last turning on will be a function of all the connections in the chain working.

Participant's goal: To determine the underlying wiring structure. There are no restrictions on time or number of interventions used.

Two conditions: *active* (participants decide which light to intervene on and when to stop) and *forced* (participants are told to follow a sequence of interventions corresponding to a participant in the active condition)

Results

21 adult participants, each solved 10 structures (randomly selected). Mean number of interventions 21.4 (std: 13.0)

Underlying Structures		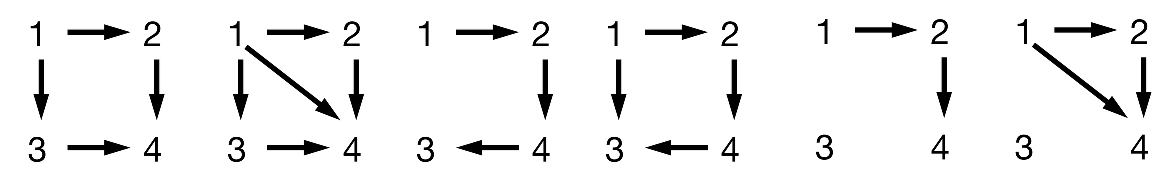					
% of edges discovered correctly	**Active**	87%	93%	66%	76%	89%	95%
	Forced	91%	87%	78%	79%	94%	90%

The difference in correctness between the two conditions was not significant ($p < 0.2$ in two-sample paired t-test), but the overall accuracy suggests that in most cases participants successfully learned the causal structure.

Learning Strategies

Two assumptions to model the learning strategies:

1)A model of a learning strategy must **infer the next choice of intervention** given the information so far (past interventions and observed light patterns).

2)If the model has a very high **predictive accuracy**, then the model has at least captured some of the core computational aspects of the human learning strategy.

Random Strategy

Select at each step randomly which light to turn on next.

Naïve Heuristic

1.Begin by selecting the Top Left light.

2.Select the light four times (to obtain a small sample given that wires are indeterministic)

3.Proceed clockwise to the next light, and repeat from 2.

Optimal Strategy

1. Construct a distribution *G* representing uniform degrees of belief over each possible wiring diagram

2. For each light, consider switching it on: intervention *s*

2.1 Consider each possible resulting light pattern L_j (the evidence)

2.2 For each such light pattern update your distribution over the possible wiring diagrams G_i using Bayes rule

$$P_s(G_i \mid L_j) = \frac{P_s(L_j \mid G_i)P(G_i)}{P_s(L_j)}$$

2.3 Compare the entropy *H* between prior and posterior to obtain the expected information gain *IG* of turning on a particular light s

$$H(p) = -\sum p_i \log_2 p_i \qquad IG(s) = H(G) - H_s(G \mid L)$$

3. Choose the light that maximizes expected information gain *IG*, observe the resulting light pattern and update the distribution *G*.

Predictive Accuracy

Strategy	Accuracy (std)
Random	24.5% (2%)
Naïve	66.9% (11%)
Optimal	32.0% (7%)

Comparison of Strategies

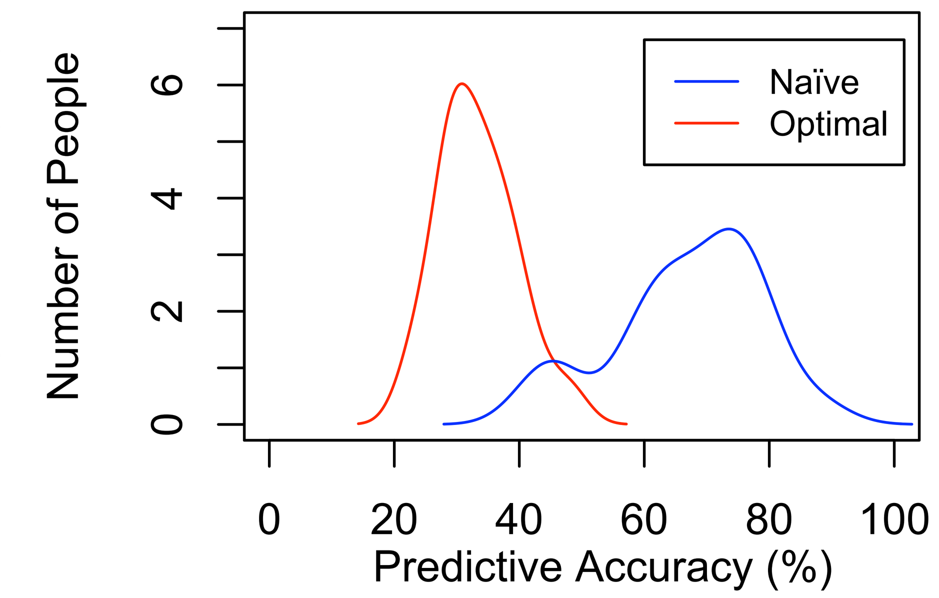

Probability Matching?

Probability matching has broad support in the literature as describing real choice behavior in psychology and economics (cf. Luce Choice Rule). A participant is said to be probability matching when they select a hypothesis H using a method that is equivalent to taking a (weighted) random sample from their (posterior) distribution over hypotheses (assuming constant utilities). For our set-up probability matching is normatively suboptimal, but we test whether it is descriptively accurate of how the choice is performed once the information gain for each light has been computed. We proceed following Goodman et al. (2007):

		Distribution of Expected Information Gain			
Trial	**Choice**	**Red**	**Green**	**Yellow**	**Blue**
1	**Red**	**0.4**	0.3	0.3	0
2	**Red**	0.2	0.2	**0.4**	0.2
3	**Blue**	0.1	0.3	0.2	**0.4**
4	**Blue**	0.1	0.2	**0.4**	0.3
5	**Green**	0	0.3	**0.4**	0.3

When a light provides 0.4 of information gain, the participant chooses it 40% of the times

Prediction

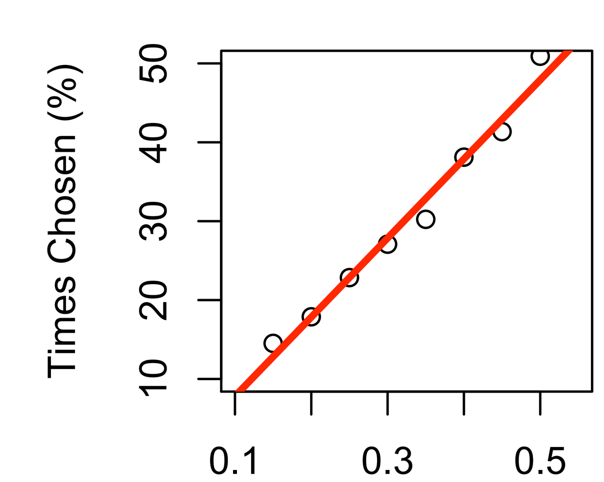

Probability of Choice

Actual Choices

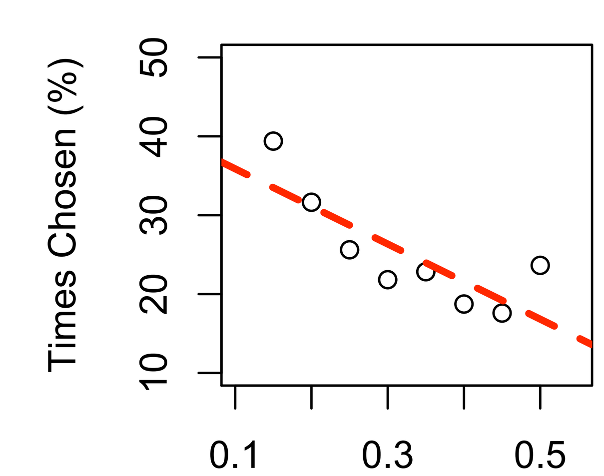

Probability of Choice

Further Issues

Under what circumstances does the observed behavior make sense?

- Confirmation bias? – but how should that be tested?
- Issues of memory and processing? – what is the nature of these processes?
- Irrationality? – but why is learning then so successful?
- Exploration vs. exploitation? – what search strategy would this entail?