Mechanisms of active causal learning

Neil Bramley, David Lagnado & Maarten Speekenbrink
University College London

Introduction

- Intervention is key to causal learning (Pearl, 2000).
- People benefit from the ability to intervene (Lagnado and Sloman, 2004; Steyvers et al., 2003) but how do people select interventions and update their beliefs?
- In this experiment participants were incentivised to learn the structure of probabilistic causal systems through free selection of multiple interventions. This design allows fine grained analysis of peoples' intervention choices and structure judgements.

Learning through interventions

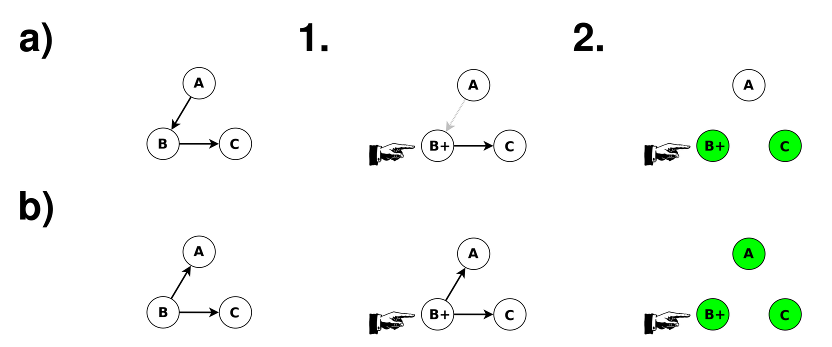

Figure 1. Two possible structures a) $A \to B \to C$ and b) $A \leftarrow B \to C$. 1. An intervention clamps B *on* (+ symbol), this severs any incoming causal links (dotted in grey). 2. If we observe that C also turns on (highlit in green) but A does not, then we have evidence for structure a). If A also comes on then this is evidence for structure b).

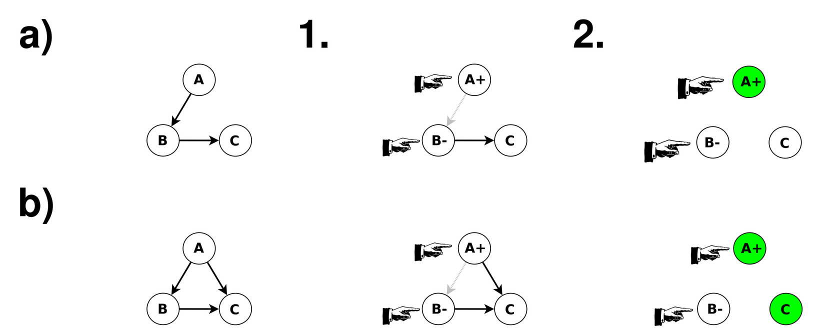

Figure 2. To distinguish a) $A \to B \to C$ from b) $A \to B \to C, A \to C$ you can manipulate A while holding B constant. 1. Achieved by clamping A *on* (+ symbol) and clamping B *off* (- symbol). 2. Then, if C still activates this is evidence for b).

Models

We assume actions are softly maximised over intervention values v: $p(\text{Int}_t = i) = \frac{e^{\alpha\langle v_i\rangle_t}}{\sum_{k=1}^{n} e^{\alpha\langle v_k\rangle_t}}$
and beliefs are soft maxed over posterior distribution $q(g)$: $p(\text{Stated-beliefs}_t = j) = \frac{\sum_{l \in j} e^{\beta q(g=l)_t}}{\sum_{k=1}^{m} e^{\beta q(g=k)_t}}$.
But what are the intervention values v?

The scholar Assumes that participants choose actions to maximise their expected *information*, or equivalently minimise their expected *uncertainty*, about the true structure.

$$-\sum_k p(g_k)\, log_2\, p(g_k)$$

The gambler Assumes participants choose actions which maximise the expected posterior probability of the most likely structure (equivalent to minimising the probability of error).

$$max_k p(g_k)$$

The utilitarian Assumes that participants choose actions which maximimise their expected payout.

$$max_k \sum_j L(g_k, g_j) p(g_j)$$

Bounded active learning

Assume people *forget* $\gamma \times 100\%$ of their posterior after each trial and/or are *conservatively* biased toward models consistent with the links they have already marked, by factor η.

Heuristic active learning

Simple endorsement Assumes people turn nodes on one at a time, and endorse links to anything that activates as a result (Fernbach and Sloman, 2009).

Disambiguation Assumes people also perform some "controlled experiments", (e.g. A^+B^-) to test disambiguate between chain and fully-connected structures.

Method

- 79 people recruited from MTurk for Flash-based online active learning task.
- Participants learned 5 causal structures, performing 12 interventions per structure (see Figure 4), taking 29.4 minutes ($SD = 16.4$).
- Paid \$1–\$4 depending on performance ($M = \$2.80$). They received 1 point per correct link (20 cents per point).
- Structures were probabilistic: nodes randomly activated with $p = .1$ and causes failed to produce effects with $p = .2$.

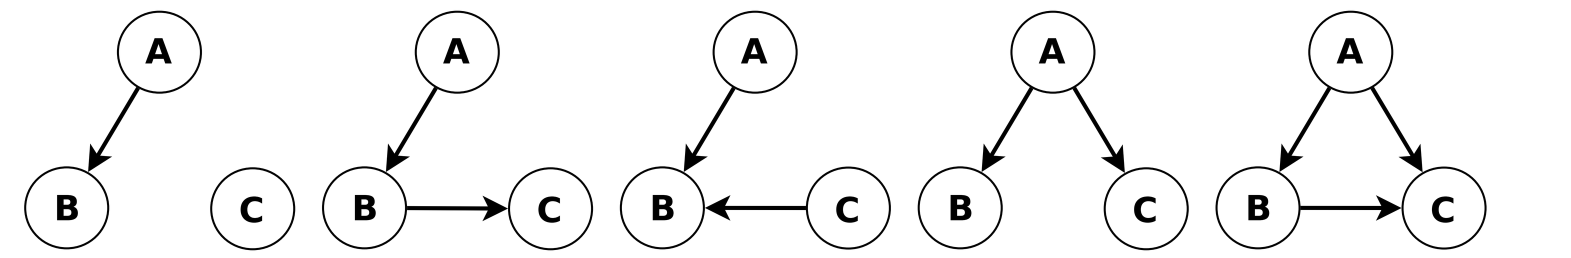

Figure 3. Causal test models. From left to right: 1. Single link $A \to B$, 2. Chain $A \to B, B \to C$, 3. Common Effect $A \to B, A \to C$, 4. Common Cause $A \to B, A \to C$, 5. Fully-connected (chain + direct link) $A \to B, A \to C, B \to C$.

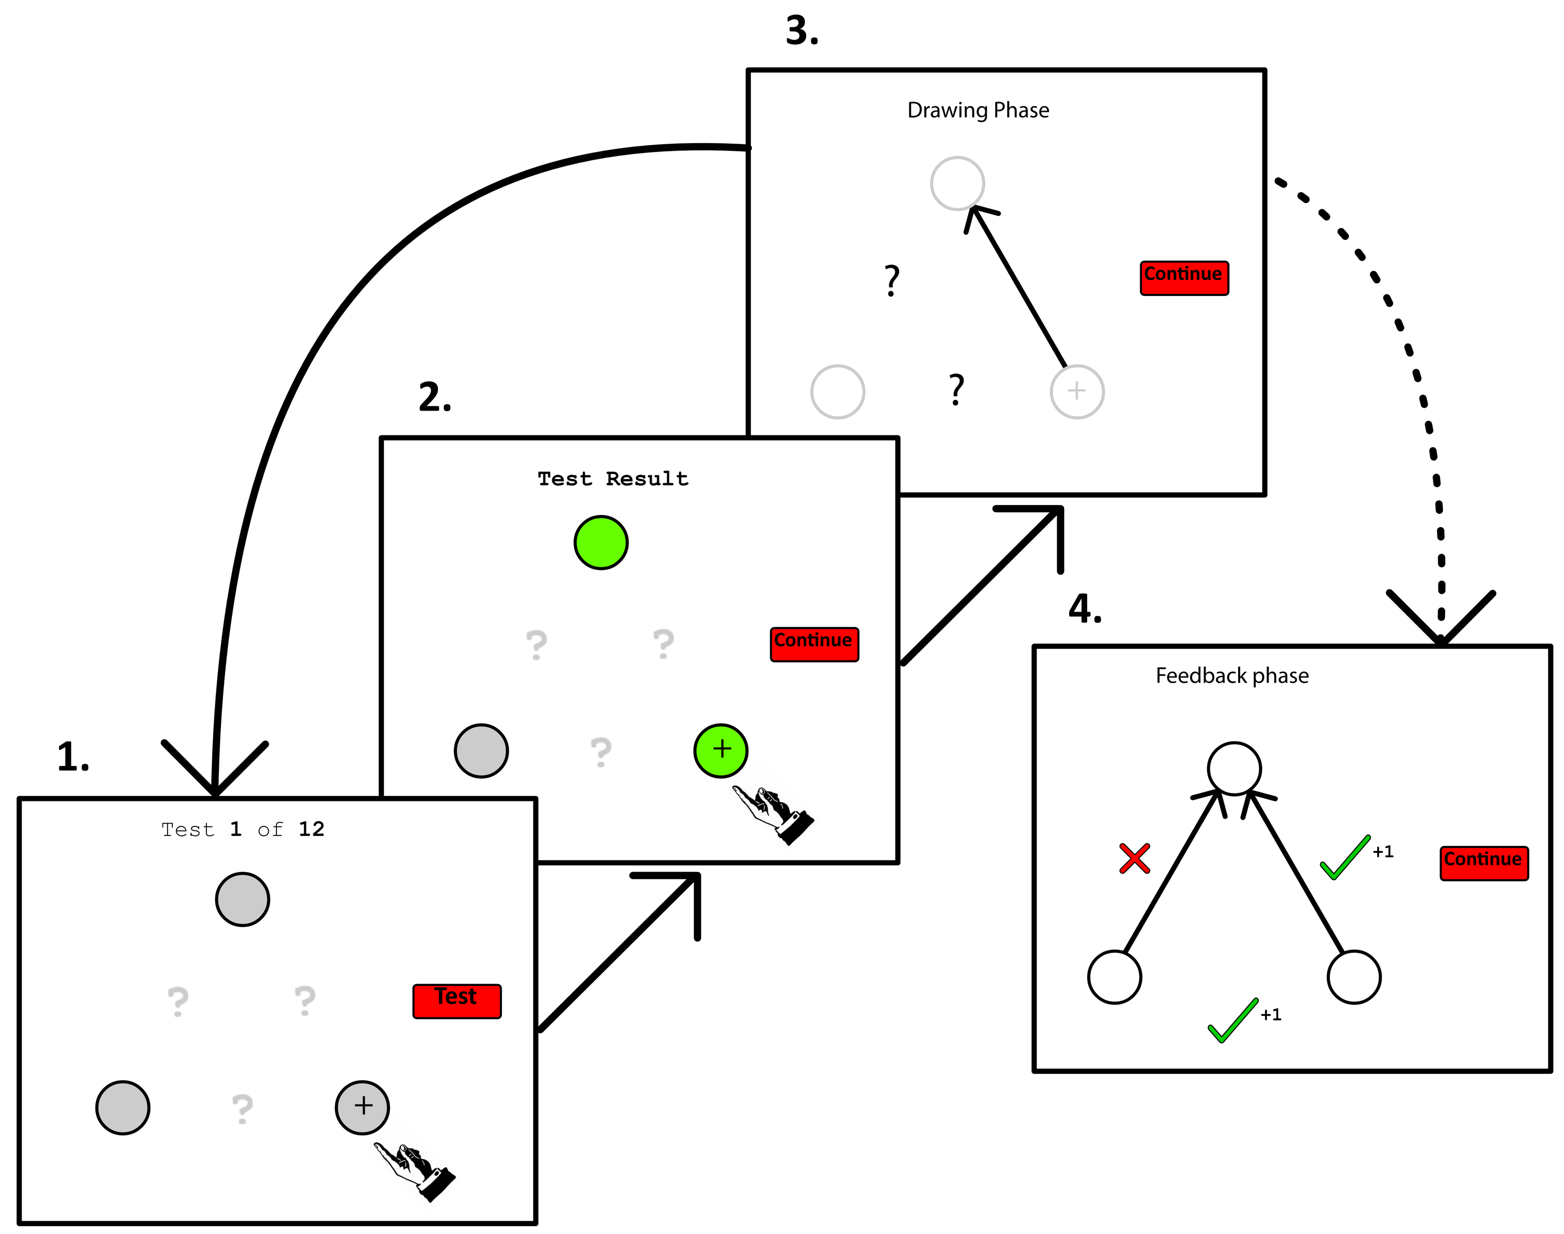

Figure 4. The procedure for a problem. 1. Choosing an intervention, 2. Observing the result. 3. Updating causal links. And, after 12 trials, 4. Getting feedback and a score for the chosen graph.

Results

Performances far above chance ($t(78) = 8.60; p < .0001$), with 8.97/15 ($SD = 4.09$) causal links identified on average (chance = 5) and 33.9% of the models completely correct on average (chance = 3.7%). Interventions on average 2.77 times as informative as randomly chosen ones, but .50 as informative as optimal choices. Better performers selected more disambiguation steps ($F_{1,77} = 8.053, p = 0.0058$). Most common error was to mistake chain for fully connected structure 18/79 compared to just 20/79 who got it right.

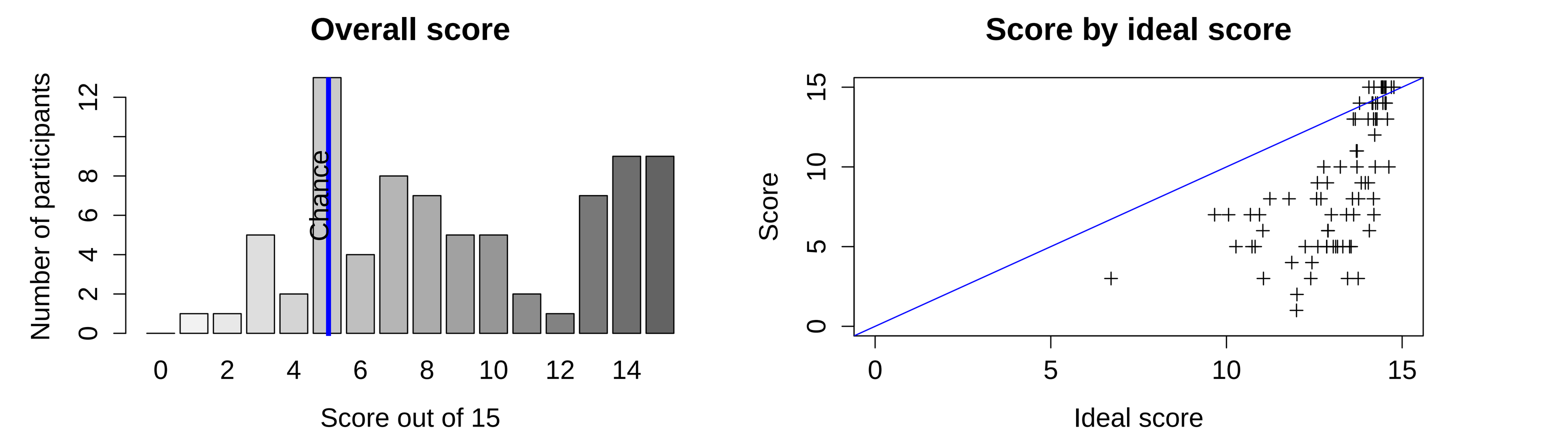

Figure 5. Left: Histogram of scores, (5 is chance, 15 is ceiling). Right: Score achieved by participants against their ideal score (if they integrated and maximised over all information). Given a set of interventions, optimal performance would lie on the blue line. The better someone's interventions, the further to the right their point lies.

Model Fits

Table 1. Ideal learner models. Overall BIC's for initial models (parameter values excluded for space). Number of participants best fit by each model. Best fitting overall model according to BIC in bold.

Model	Steps-ahead	BIC	♯ participants
Scholar	**1**	**45047**	**68**
Gambler	1	47917	2
Utilitarian	1	47464	1
Scholar	2	45525	6
Gambler	2	46877	2
Utilitarian	2	46778	0

Table 2. Bounded models, with forgetting and conservatism parameters. Overall BICs. ♯ best described, and ♯ excluding those who did not score above chance (5 points).

Model	BIC	♯ participants (/79)	♯ participants$_{>5}$ (/57)
Scholar	**31145**	**31**	**29**
Gambler	31575	29	19
Utilitarian	31630	14	8

Table 3. BIC's for heuristic models. Number of participants best fit by each model.

Model	BIC	♯ participants
Simple endorser	37334	40
Disambiguator	**36938**	**39**

The scholar model was clearly favoured out of the ideal models. Simple heuristic models do well, beating ideal learner models. The bounded scholar model was the best overall with a McFadden pseudo-R^2 of $.44$. Less forgetful people are more likely to be disambiguators ($F_{1,77} = 6.02, p = .016$). And forgetting rate predicts performance ($F_{1,77} = 21.4, p < .001, +.25$ points per each 10% less forgetful) suggesting that memory (or concentration) drove the poor performances of some participants.

Conclusions

- Many people are highly effective active causal learners, able to learn fully probabilistic structures over multiple interventions.
- Most people are *forgetful conservative scholars* - querying the environment to reduce their *uncertainty*, rather than their *probability of being correct* or *expected payout*, forgetting old evidence as they go along, but mitigating this by being conservative about their existing beliefs.
- Heuristic models can capture these patterns with a mixture of simple endorsement and disambiguation steps.

References

Fernbach, P. M. and Sloman, S. A. (2009). Causal learning with local computations. *Journal of experimental psychology: Learning, memory, and cognition*, 35(3):678.

Lagnado, D. and Sloman, S. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 30:856–876.

Pearl, J. (2000). *Causality*. Cambridge University Press (2nd edition), New York.

Steyvers, M., Tenenbaum, J. B., Wagenmakers, E., and Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27:453–489.