What should an active causal learner value?

Neil R. Bramley[1], Jonathan D. Nelson[2], Maarten Speekenbrink[1], Vincenzo Crupi[3], & David A. Lagnado[1]
nrbramley@gmail.com, jonathan.d.nelson@gmail.com, m.speekenbrink@ucl.ac.uk, vincenzo.crupi@unito.it, d.lagnado@ucl.ac.uk
University College London[1] Max Planck Institute for Human Development[2] University of Turin[3]

Max-Planck-Institut für Bildungsforschung 50
50 Years of Research on Human Development

UCL

What is active causal learning?

- Intervening on a causal system provides important information about causal structure (Pearl, 2000; Sprites et al., 1993; Steyvers, 2003; Bramley et al., 2014)
- Interventions render intervened-on variables independent of their normal causes: $P(A|Do[B], C) \neq P(A|B, C)$
- Which interventions are useful depends on the prior and the hypothesis space
- *Goal:* Maximize probability of identifying the correct graph

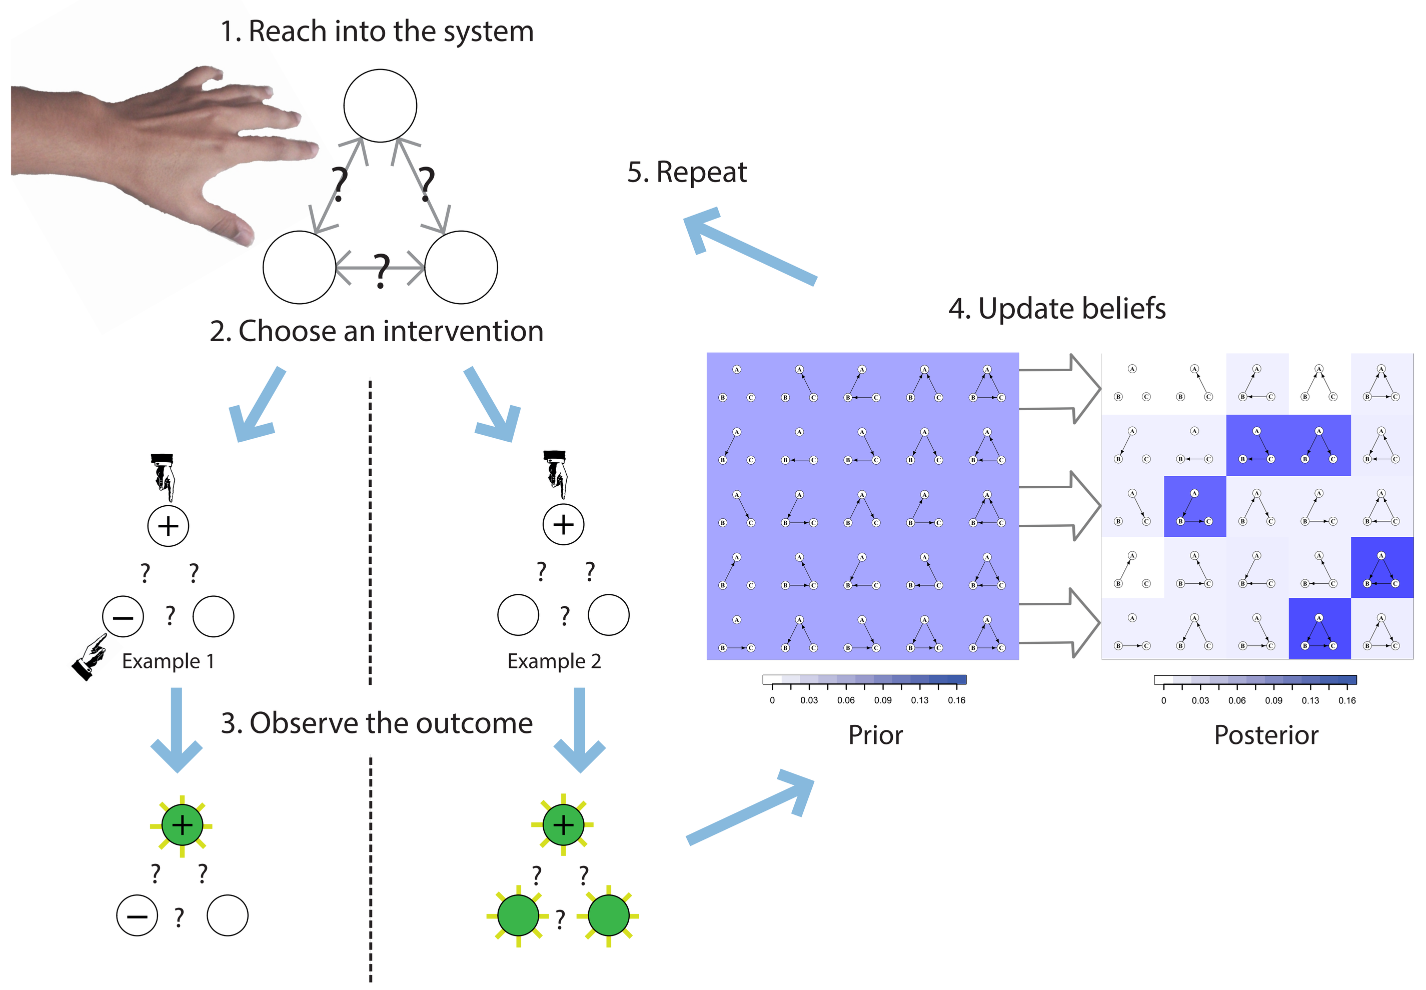

Strategies for choosing interventions myopically

Probability gain Choose intervention $i \in I$ that maximises expected maximum probability ϕ in posterior over causal graphs $G \in G$
$E[\phi(G|i)] = \sum_{o \in O} \phi(G|i, o)p(o|G, i)$ where $\phi(G|i, o) = \max_{g \in G} p(g|i, o)$
Information gain Choose intervention that minimises expected Shannon entropy
$E[H(G|I)] = \sum_{o \in O} H(G|i, o)p(o|G, I)$ where $-\sum_{g \in G} H(g|I, o) = p(g|I, o) \log_2 p(g|i, o)$

Can information gain beat probability gain at probability gain?

Bramley et al. (2014), test performance learning all 25 of the 3-variable graphs in an environment with spontaneous activation rate $p = .1$ and causal power $p = .8$:

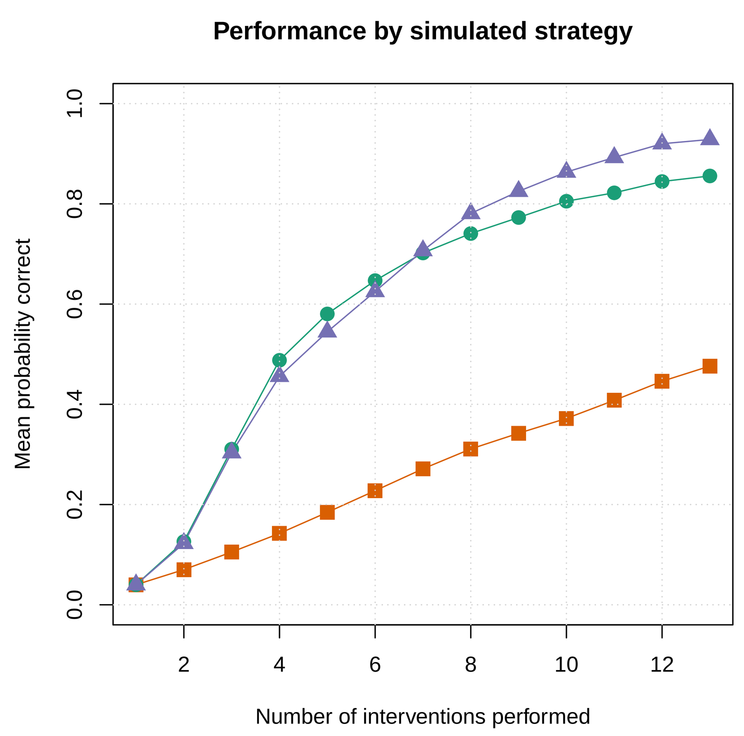

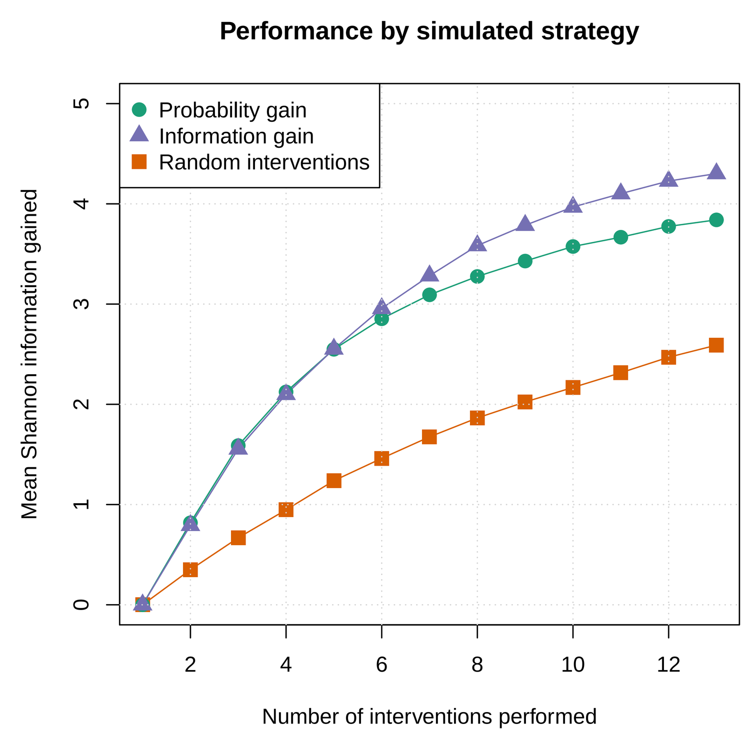

Beyond Shannon: Generalized entropy landscapes

Entropy is *expected surprise* across possible outcomes. Shannon entropy uses the normal mean (e.g. $\sum_x Surprise(x) \times p(x)$) as expectation and $log(1/p(x))$ for surprise. Different types of mean, and different types of surprise, can be used (Crupi and Tentori, 2014). Sharma-Mittal entropy (Nielsen and Nock, 2011) provides a unifying framework:

- Sharma-Mittal entropy: $H_{\alpha,\beta}(p) = \frac{1}{1-\beta}\left(\left(\sum_i p(x)^{\alpha}\right)^{\frac{1-\beta}{1-\alpha}} - 1\right)$, where $\alpha > 0$, $\alpha \neq 1$, and $\beta \neq 1$
- Renyi (1961) entropy if degree $(\beta) = 1$
- Tsallis (1988) entropy if degree $(\beta) =$ order (α)
- Error entropy (probability gain; Nelson et al. (2010)) if degree $(\beta) = 2$, order $(\alpha) = \infty$

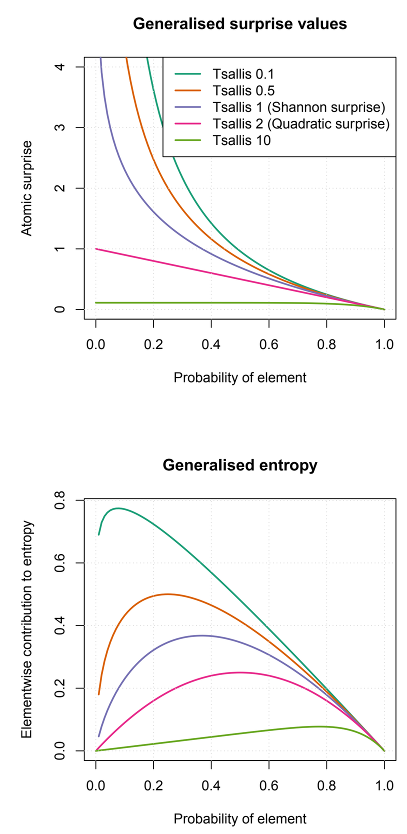

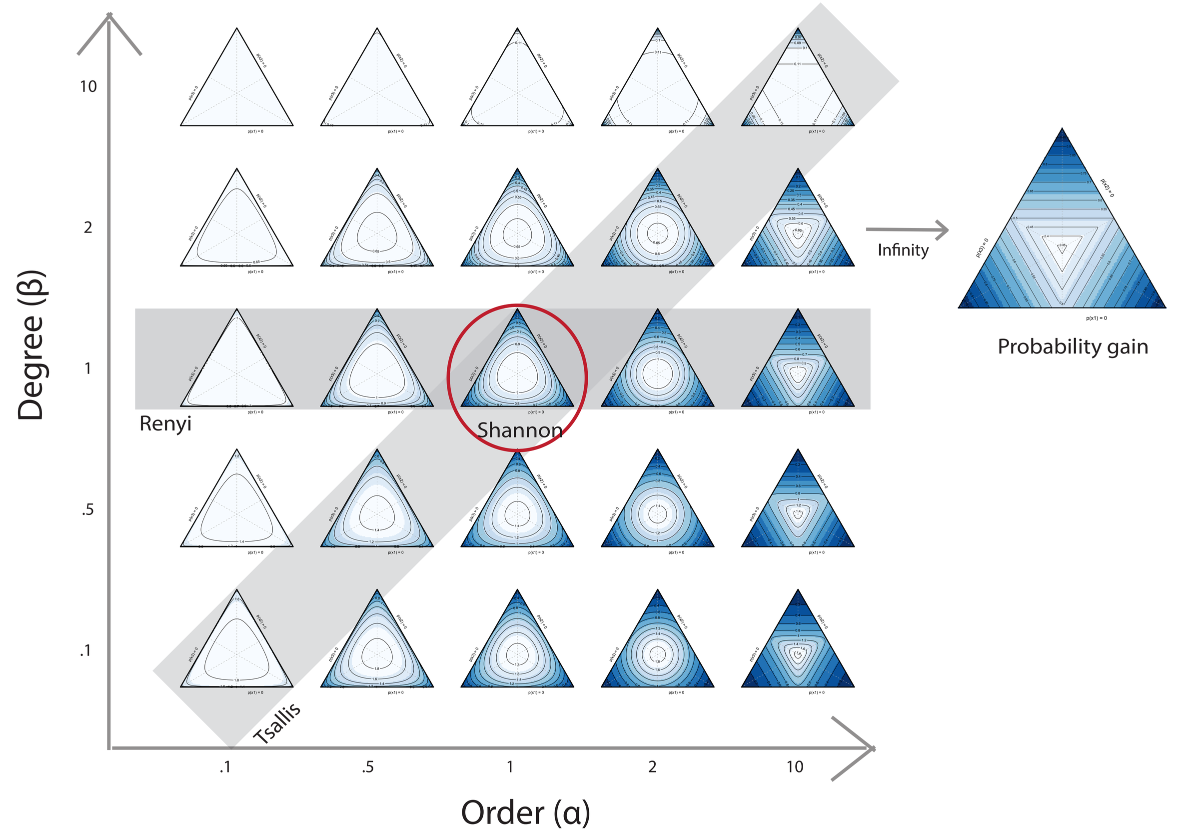

Weird and wonderful (bad) entropy functions

To further broaden the space of entropy functions considered, we first created a variety of atomic surprise functions. From these we derived corresponding entropy functions, always defining entropy as $\sum_x Surprise(x) \times p(x)$. The resulting entropy functions vary in which entropy axioms (Csiszár, 2008) they satisfy.

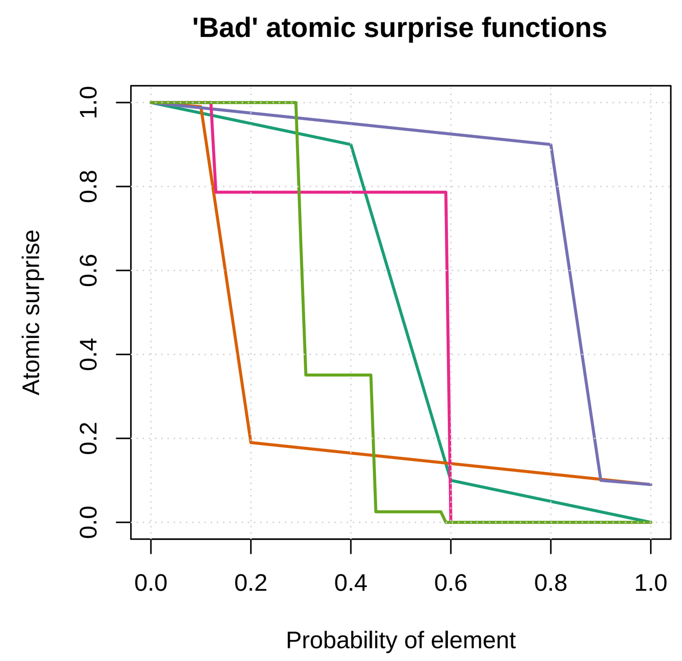

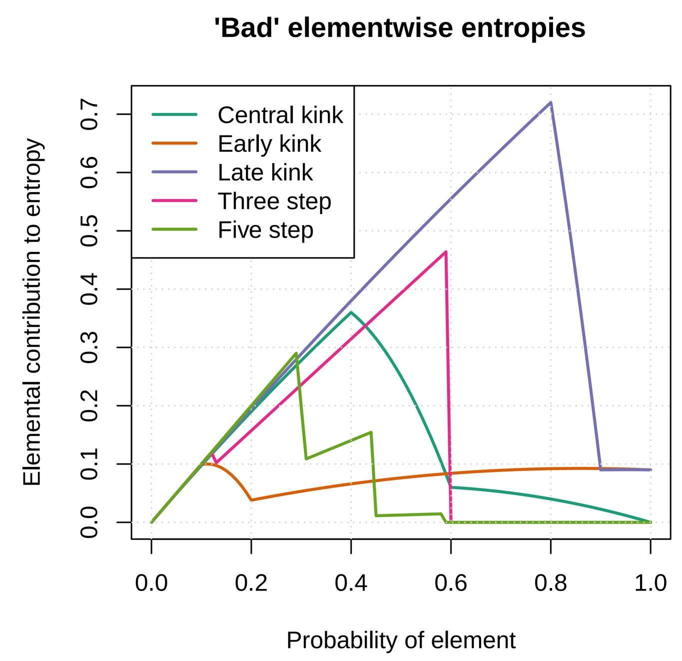

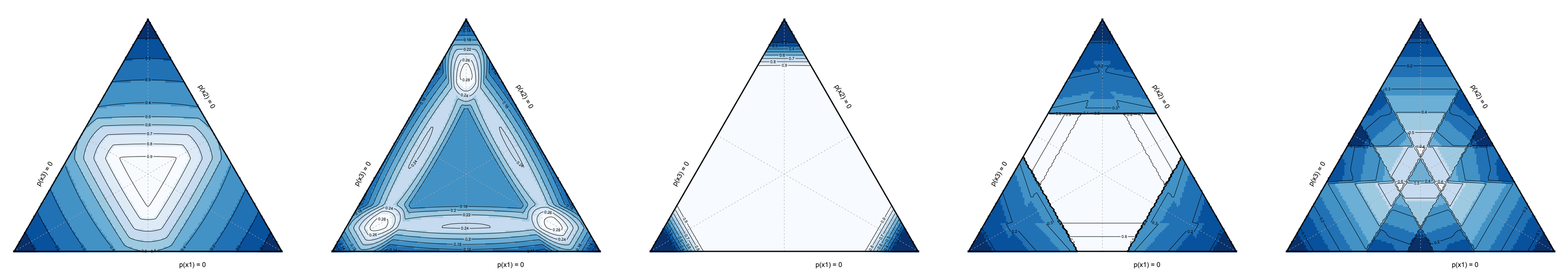

Central kink Early kink Late kink Three random steps Five random steps

Simulating many strategies in many environments

- Learners based on 5×5 entropies in Sharma-mittal space $\alpha, \beta \in [1/10, 1/2, 1, 2, 10]$, plus our 5 weird and wonderful entropy functions.
- Test cases: all 25 of the 3-variable causal graphs, through 8 sequentially chosen interventions. Repeated simulations 10 times each in 3×3 environments:
  1. Spontaneous activation rates of $[.2, .4, .6]$
  2. Causal powers of $[.4, .6, .8]$

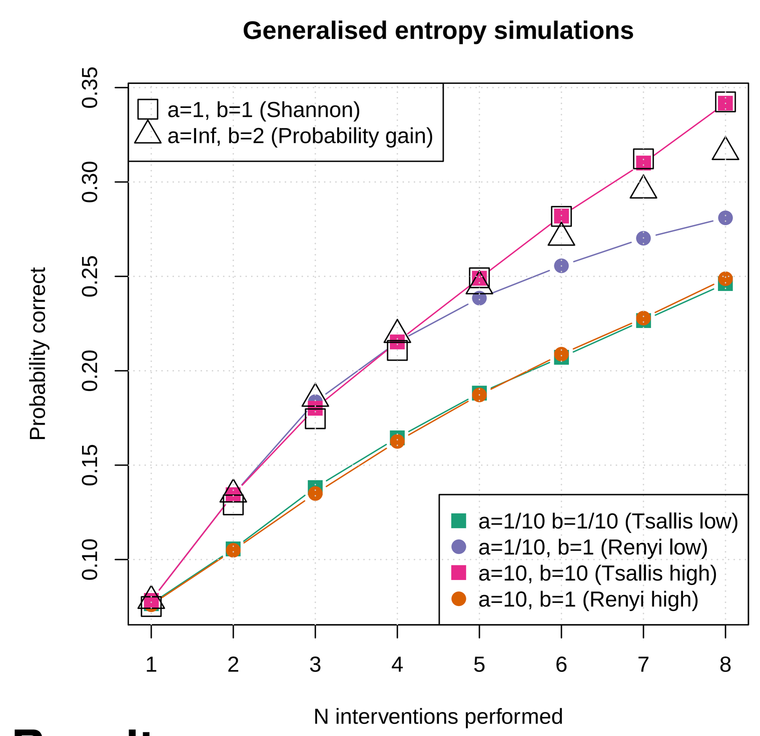

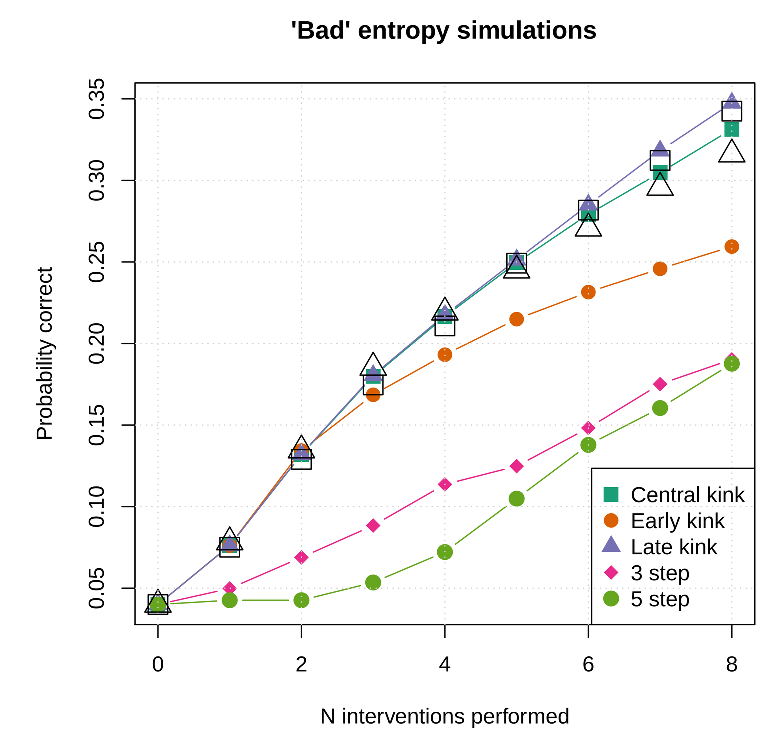

Results

- Shannon entropy achieves higher accuracy than probability gain
- many entropy functions achieve similar accuracy as Shannon entropy
- but the interventions chosen vary widely:

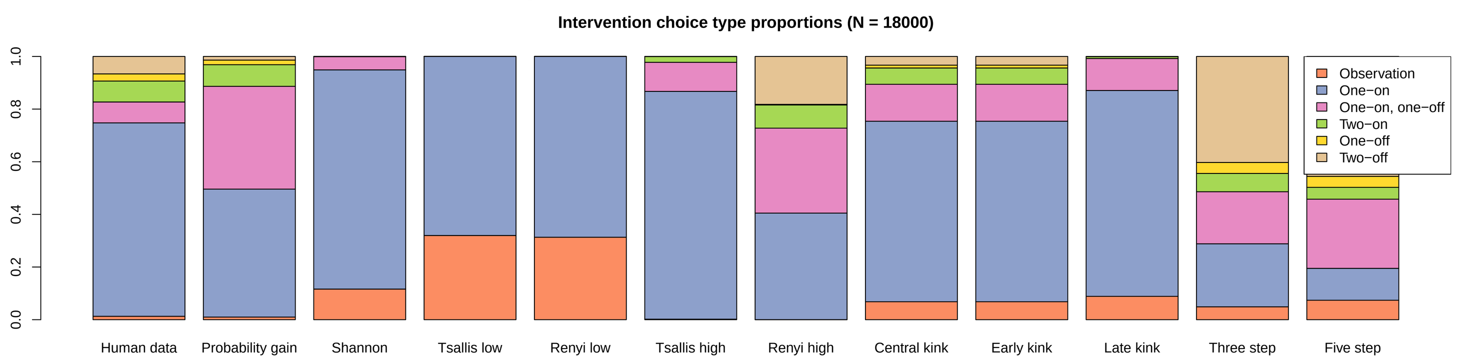

Questions:

- in which instances do different strategies strongly contradict each other?
- which properties of entropy functions are important for particular situations?
- when are particular entropy functions cognitively plausible?

References

Bramley, N. R., Lagnado, D. A., and Speekenbrink, M. (2014). Forgetful conservative scholars - how people learn causal structure through interventions. *Journal of Experimental Psychology: Learning, Memory and Cognition*, doi:10.1037/xlm0000061.

Crupi, V. and Tentori, K. (2014). State of the field: Measuring information and confirmation. *Studies in History and Philosophy of Science Part A*, 47:81–90.

Csiszár, I. (2008). Axiomatic characterizations of information measures. *Entropy*, 10(3):261–273.

Nelson, J. D., McKenzie, C. R., Cottrell, G. W., and Sejnowski, T. J. (2010). Experience matters information acquisition optimizes probability gain. *Psychological science*, 21(7):960–969.

Nielsen, F. and Nock, R. (2011). A closed-form expression for the sharma-mittal entropy of exponential families. *arXiv preprint arXiv:1112.4221*.

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

Renyi, A. (1961). On measures of entropy and information. In *Fourth Berkeley Symposium on Mathematical Statistics and Probability*, pages 547–561.

Sprites, P., Glymour, C., and Scheines, R. (1993). *Causation, Prediction, and Search (Springer Lecture Notes in Statistics).* MIT Press. (2nd edition, 2000), Cambridge, MA.

Steyvers, M. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27(3):453–489.

Tsallis, C. (1988). Possible generalization of boltzmann-gibbs statistics. *Journal of statistical physics*, 52(1-2):479–487.