Exploring fixed-threshold and optimal policies in multi-alternative decision making

Michael Shvartsman[1] Vaibhav Srivastava[2] Jonathan D. Cohen[1]

[1]Princeton Neuroscience Institute, Princeton University [2]Electrical and Computer Engineering, Michigan State University

(1) How do we combine information to make decisions?

Our model of decision dynamics: sequential Bayesian inference

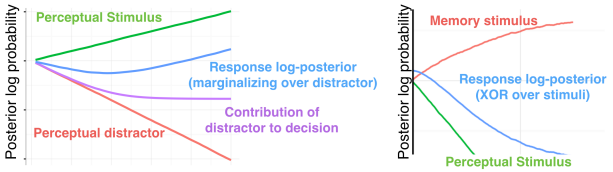

Related formulations of the same problem: POMDP, diffusion

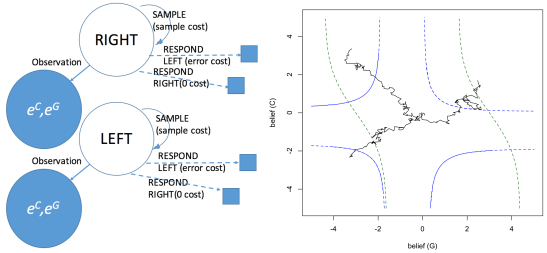

This is also a Partially Observable Markov Decision Process with continuous observations (left), and (in the continuum limit) a two-dimensional diffusion process on the log-posteriors of the two stimuli (right). The dashed lines are equiprobability contours for two different response rules.

(2) Fixed-threshold vs. Bayes-risk-optimal policies

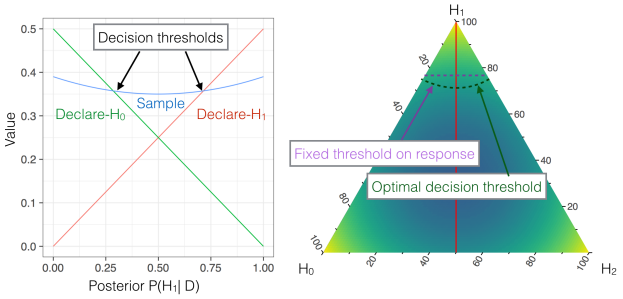

Left: with two hypotheses, linear value of decision action and convex value of sample action intersect at points: optimal policy is fixed threshold w.r.t. the highest posterior (Wald & Wolfowitz, 1948). *Right*: with more than two hypotheses, decision action value is a plane on the simplex increasing parallel to the red line, and the value of the sample action is a 3d paraboloid (shown as heat map). Their intersection is a curve in 2d, rather than the point as in the 2-hypothesis case: the threshold is a function of the full posterior rather than the probability of the likeliest hypothesis. Fixed thresholds are asymptotically optimal (Baum & Veeravalli, 1994).

Contact Information

- Web: http://www.princeton.edu/~ms44
- Email: ms44@princeton.edu

This project was made possible through the support of the John Templeton Foundation, the Swartz Foundation, and Intel Corporation.

(3) Different explanations for flanker effect are optimal under different constraints

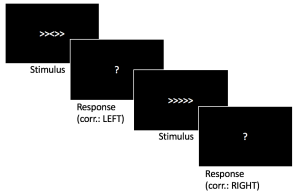

Schematic diagram of flanker task. Subjects respond to the direction of the central stimulus while ignoring the flanking distractors.

The flanker effect: incongruent trials are slower and less accurate

Candidate explanations:

- *Congruence bias (prior)*: subjects expect stimuli to be spatially congruent (Yu, Dayan & Cohen, 2009).
- *Spatial uncertainty (likelihood)*: perceptual evidence is spatially correlated (Yu et al. 2009).
- *Attentional spotlight (inference)*: subjects have an attentional spotlight that shrinks over time (White, Brown, & Ratcliff 2012).

Congruence bias is consistent with either decision rule

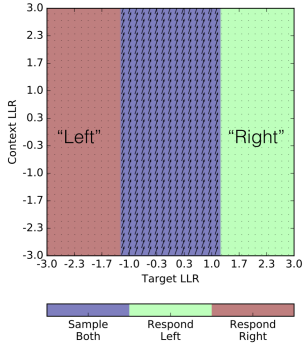

Vector field and optimal policy for congruent >>>>> trial without spatial uncertainty. In this case the fixed threshold policy on the response and the bayes-risk-optimal policy coincide. If the agent can selectively sample from either flanker or targets, the result is identical except that the vector field points right.

Optimal threshold compensates for spatial uncertainty

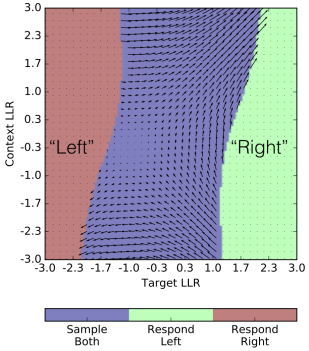

Vector field and optimal policy for congruent >>>>> trial with spatial uncertainty. Due to spatial uncertainty, the vector field accellerates in the congruent quadrants (top right and bottom left). Because of the increased value of the 'sample' action relative to the 'respond' action in more congruent states, the threshold moves out, compensating for the spatial uncertainty effect. The flanker effect due to spatial congruence is reduced or diminshed in the BRO policy relative to a fixed-threshold policy.

Narrowing attentional spotlight naturally falls out of spatial congruence

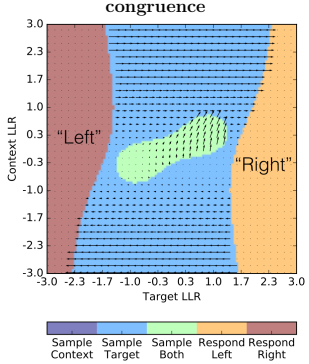

Vector field and optimal policy for congruent >>>>> trial with spatial uncertainty and ability to selectively sample from either flanker or targets. The transition from the light green area (sample-both) to the blue area (sample-target) is a discrete narrowing of the attentional spotlight and happens in a way that is sensitive to congruence.

(4) Under Bayes-risk-optimal policies, response criterion increases as stimulus is identified

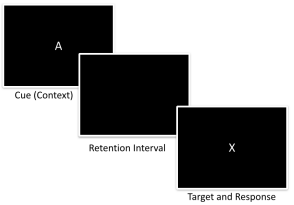

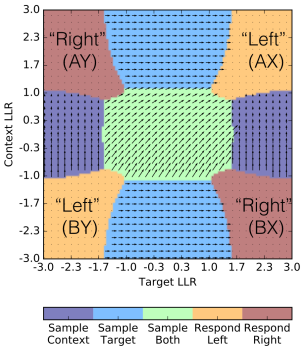

Top: schematic diagram of the AX-CPT. Subjects compute an exclusive-or response based on two stimuli presented sequentially in time. *Right*: vector field and optimal policy for AX trial of AX-CPT, showing the same characteristic pattern: because of the increased value of the 'sample' action relative to the 'respond' action in more congruent states, the threshold is higher with greater confidence in the other stimulus. Those belief states are reachable in discrete time, or in continuous time via noisy prior bias.

(5) The intention superiority effect may be a signature of the optimal decision rule

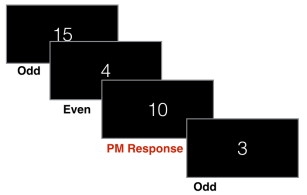

Left: schematic diagram of the prospective memory paradigm. Subjects perform a two-alternative forced choice task on some stimulus dimension (e.g. even-odd judgment) while remaining vigilant for a 'special' stimulus for which a third response is given (e.g. the number '10'). The *intention superiority effect* is when subjects are quicker to perform the original task to the special stimulus (e.g. quicker to identify that '10' is even than that '4' is even).

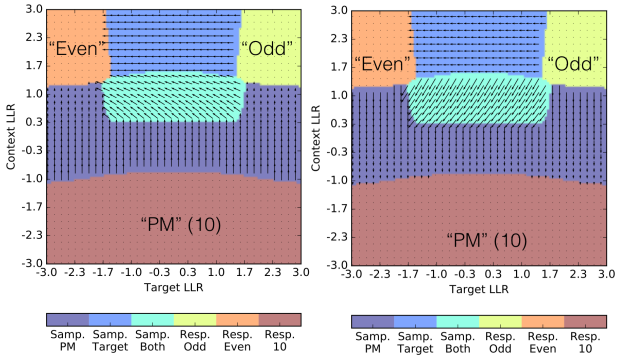

Optimal policy for the PM task (assuming the PM cue involves a memory retrieval, and attention can be separately assigned to the PM and non-PM dimensions). *Left*: vector field when the stimulus is even but not '10', i.e. the 'even' response is appropriate. *Right*: vector field when the stimulus is '10', i.e. the 'PM' response is appropriate. Notice that greater certainty that this is a non-PM trial will raise the threshold on the non-PM response, and therefore erroneous 'even' responses on PM trials are likelier to happen at lower levels of belief, and therefore faster.