Inferring relevance in a changing world

Robert C. Wilson, Jonathan D. Cohen and Yael Niv

Models of learning usually concentrate on how we learn the relationship between different stimuli or actions and rewards. However, in real world situations "stimuli" are ill-defined. On the one hand, our immediate environment is extremely multi-dimensional. On the other hand, in every decision making scenario only some (usually few) aspects of the environment are relevant for obtaining reward, while most are irrelevant. Thus a key question for neuroscience is how the brain *learns* these relevant dimensions, adapting to different tasks and a changing world. In this work we describe experiments and modeling of an experimental paradigm designed to study this question directly.

On each trial of our "dimensions task", subjects are presented with three stimuli differing on three feature dimensions. Two of the stimuli are associated with a low probability of reward, while one is highly rewarding. Importantly, the identity of the more rewarding stimulus is determined by the value of only one of the feature dimensions. Thus, to maximize reward, subjects must learn what is the relevant dimension, as well as what is the correct feature within this dimension. Importantly, they must infer the relevant dimension from indirect feedback via rewards.

What makes this task even more difficult is that the identity of the rewarding dimension and/or rewarding feature changes abruptly and in an unsignaled manner every 15-30 trials. Thus subjects must constantly revise their estimates of the rewarding dimension and deal with the triple uncertainty of probabilistic rewards, unknown rewarding dimension identities and unknown change-point locations.

We analyze human performance in this task relative to a Bayesian ideal observer model that can optimally integrate past information (i.e. stimuli and their associated rewards) in order to compute the joint-distribution over dimension and value. Although we find a strong correlation between subject's choices and reward probability of the ideal observer, a more detailed look at the trial by trial data shows that subjects can often be far from optimal. We thus expand on this analysis, first using the behavioral data to arbitrate between different possible models, and then using the best-fitting model to determine the possible computations and types of suboptimal behaviors that the brain might be using to solve the problem. Our work provides theoretical and experimental insight into how humans can implement or approximate Bayesian inference in a highly non-trivial task that is representative of real-world decision making.