Abstract Knowledge Guides Search and Predictions in Novel Situations

Russell E. Warner • Patrick Shafto
Department of Psychological and Brain Sciences, University of Louisville

Chris Baker • Joshua B. Tenenbaum
Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology

Introduction

People combine abstract knowledge about the world with data they gather in order to guide search and prediction in everyday life.

For example, if you open a new box of cracker jacks and immediately found the prize, you would not expect to find another. That expectation is an abstraction that derives from repeated encounters with cracker jack boxes and eventually tells you something about the manufacturing process.

We formalize and test a Bayesian model of how people represent such abstract knowledge, and use it to guide search and prediction in novel situations.

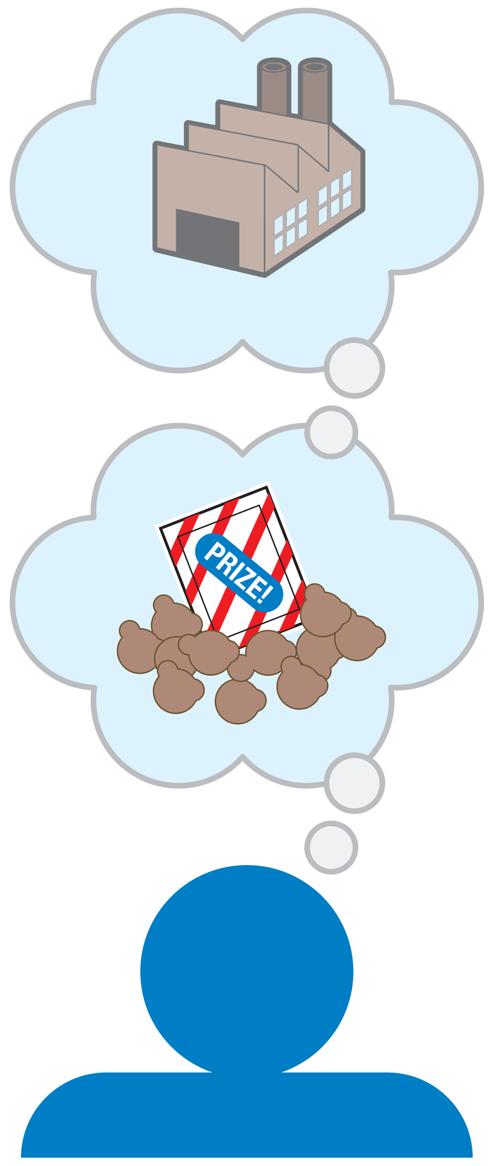

Representing and Reasoning with Abstract Knowledge

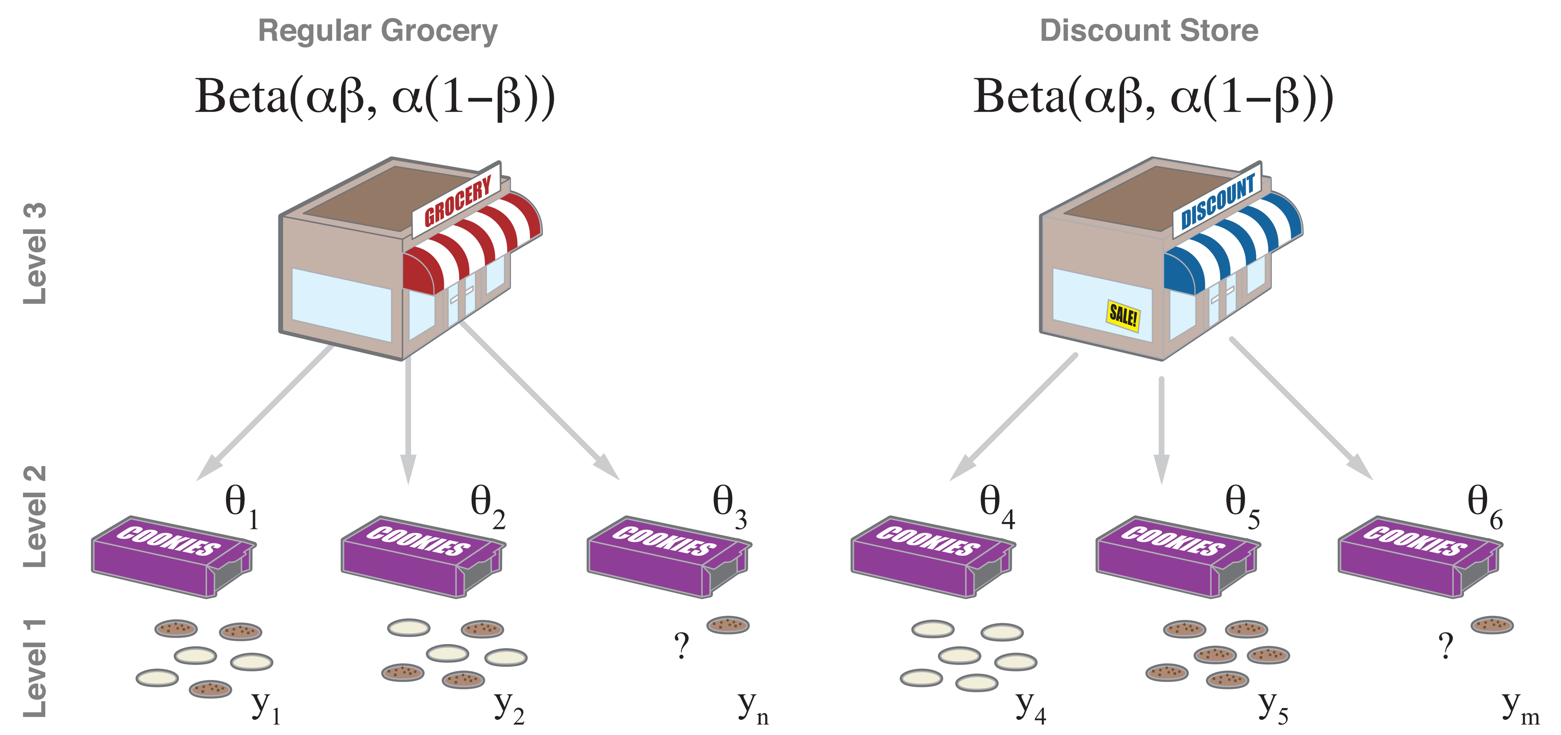

Figure 1

Suppose you usually buy cookies from a regular grocery and know that mixed boxes have as many choclate chip cookies as vanilla wafers. Then one day you find a discount store where the boxes are cheaper, but the cookies are not evenly mixed—perhaps the reason for the discount.

The model represents this knowledge at three levels, individual instances (cookies), the proportion of cookies in a box, and the kinds of boxes that a store has. At the top of Figure 1, the alpha and beta are parameters that represent our confidence about the mixtures of the boxes and the balance between the outcomes. At level 2, we can have knowledge about the proportions in each different box, and at the lowest level, we can have knowledge about specific data (here, cookies). In a discount store, an observation of one cookie informs expectations about a box, since boxes tend to be one kind or the other.

Search and Prediction

Now suppose you are at the discount store and you want a box that contains mostly your favorite kind, but you don't have much time to spend checking. Would you bother to check the box? If so, how many cookies would you need to see? If you were in the regular store, would your strategy change?

We formalize a model utilizing a Markov Decision Process (MDP) to generate strategies for sampling data. The model predicts that the sampling strategy is guided by the beliefs about the store (Top of Figure 2), as well as the implications of possible samples on beliefs. In the discount store (left column) a single sample is enough to change beliefs about the contents, where as in the regular store (right column) no sampling is necessary.

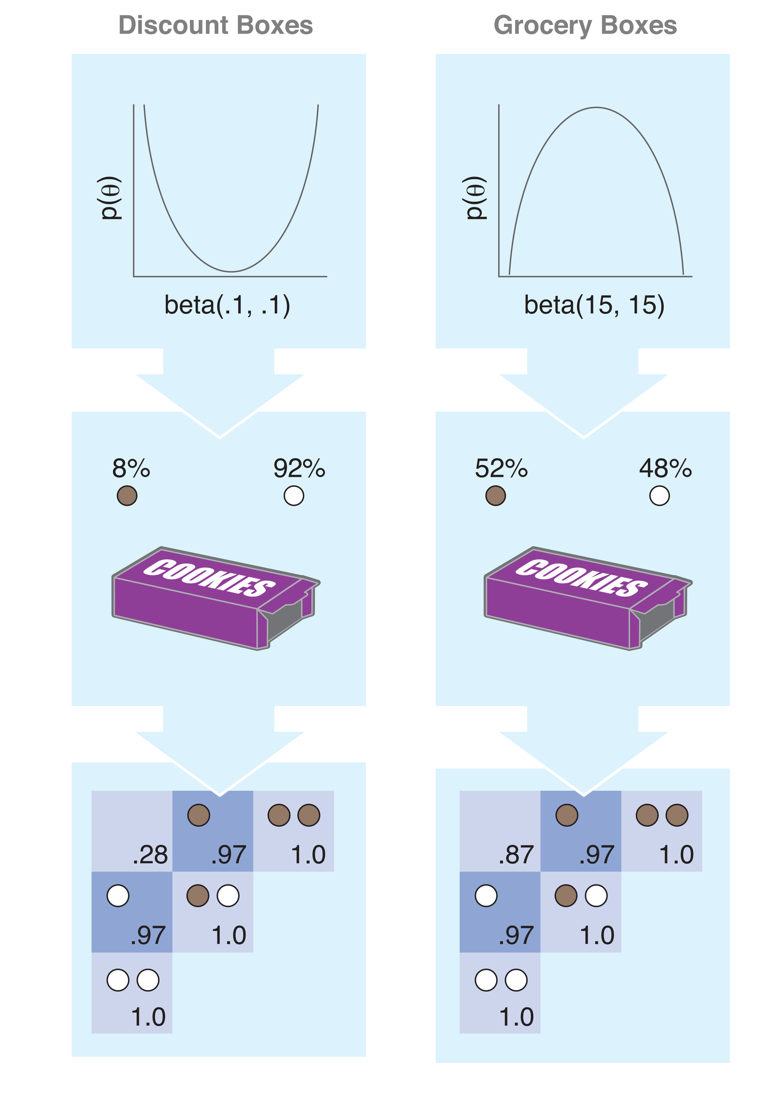

Figure 2

Higher Order Knowledge

This basic framework for search and prediction model assumes the learner already knows the situation they are in; however, people can learn what kinds of situation they are in. Learning is captured by updating beliefs about stores (the parameters of the beta distribution). With this extended model we can ask how experience in one situation affects learning in a subsequent situation. The simulation in Figure 3 shows alpha being updated as new evidence is encountered. When the unmixed (biased) condition is presented before the well-mixed (even) conditions, the model predicts greater difficulty in learning in the even condition.

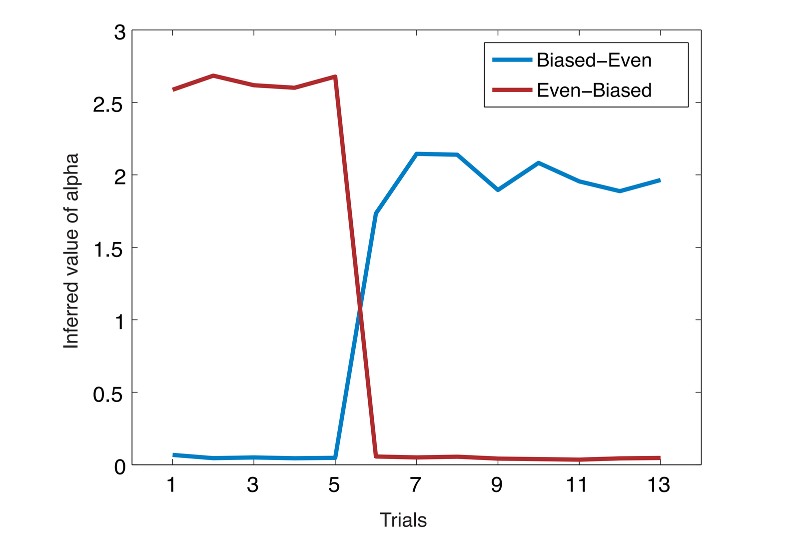

Figure 3

The Big Urn Game

Twenty-seven participants played a game on a computer. The object was to guess the color of the next ball drawn from an (virtual) urn filled with red and blue balls. Participants could score three points for a correct guess or draw a ball from the urn for the cost of one point.

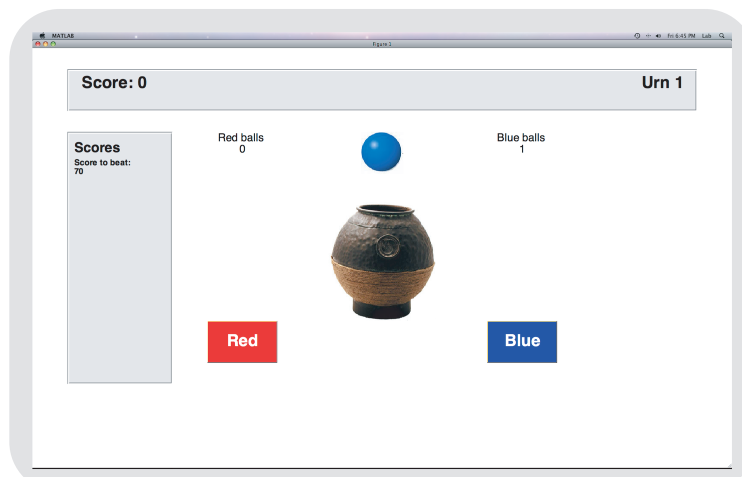

After guessing the color of a ball from an urn, correctly or incorrectly, the players were informed about the mixture of the colored balls in that urn (i.e. "This urn contained 52% blue balls and 48% red balls.") They were then shown a new urn. Players worked with urns generated from one Beta distribution until they sucessfully beat a criterion score calculated from our model's optimal strategy. They were then told "You have completed 20 urns, and succeeded in predicting enough of them correctly to move on to the next condition."

Results

Do people learn different strategies for different situations?

We analyzed people's behavior in the last half of their first condition to assess the strategies that people converged on in each condition. Figure 4 shows the probability of guessing on the first chance, and after one draw. The probability of the initial draws are significantly greater in the biased condition than in the even condition, consistent with the model's prediction that a draw in the biased condition presents useful information.

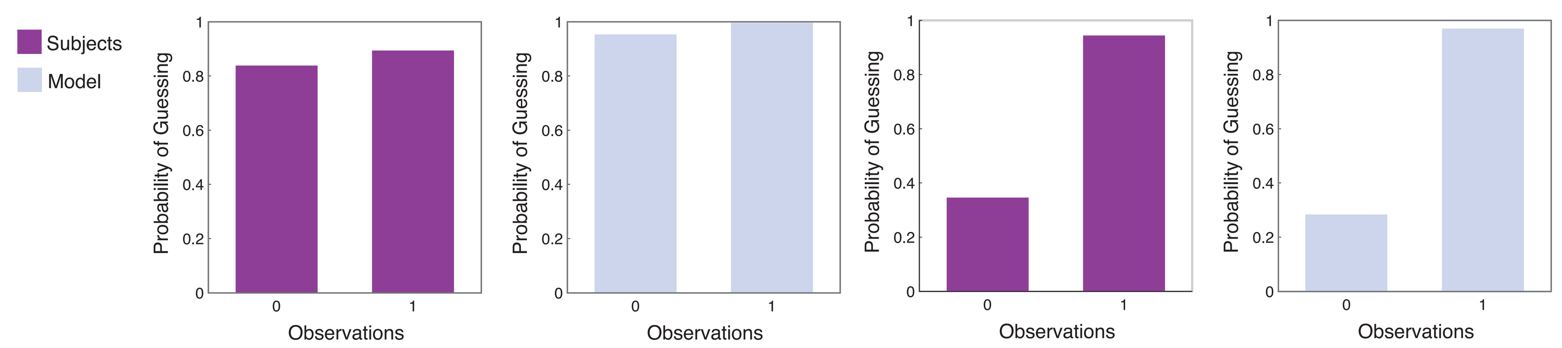

Figure 4

How does experience with one situation affect learning in a subsequent situation?

In the initial condition, people should not have strong expectations about the urn generating process. By the time they arrive in the second condition, their experience should bias their expectations (see Figure 5). To test the prediction that learning should be more difficult in the biased-even condition, individuals in that condition should play, on average, a larger number of games, indicating that it was more difficult to meet the criterion score. Figure 6 shows this to be the case.

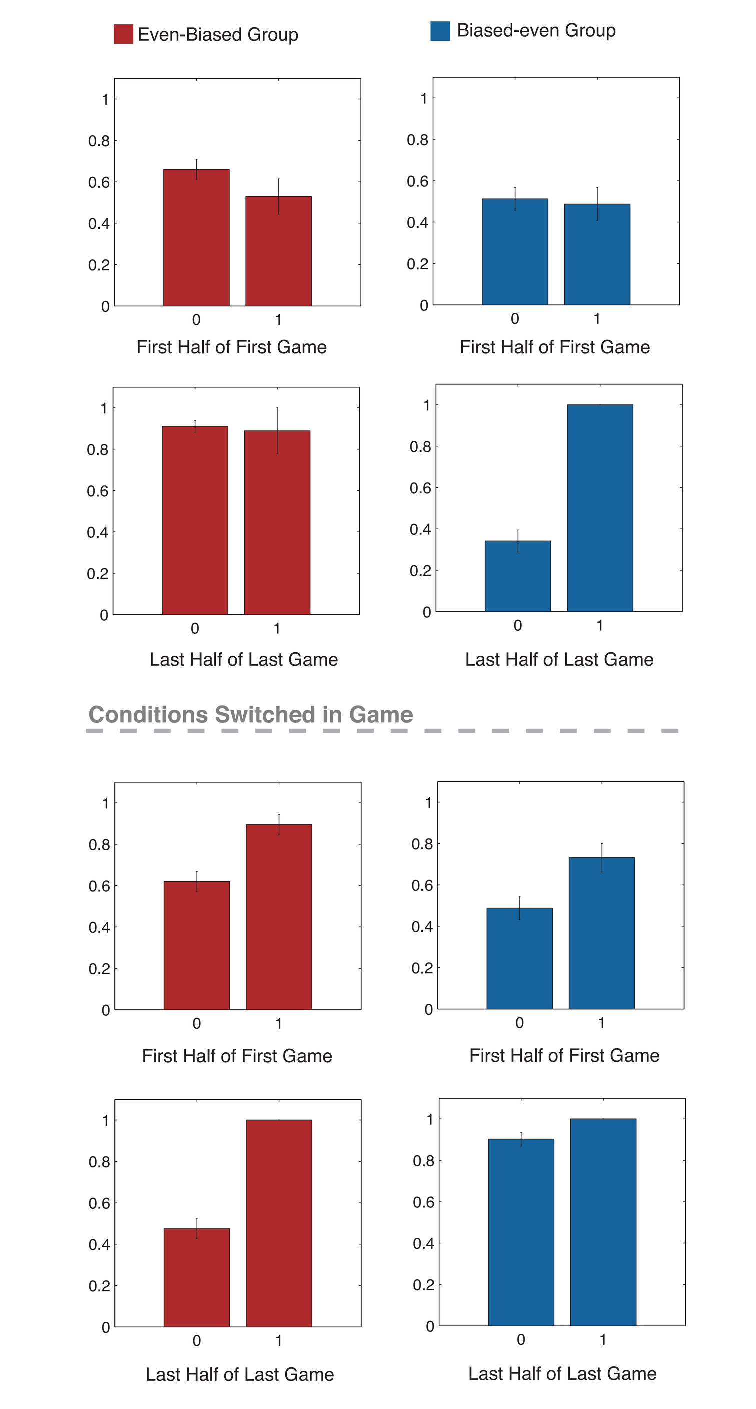

Figure 5

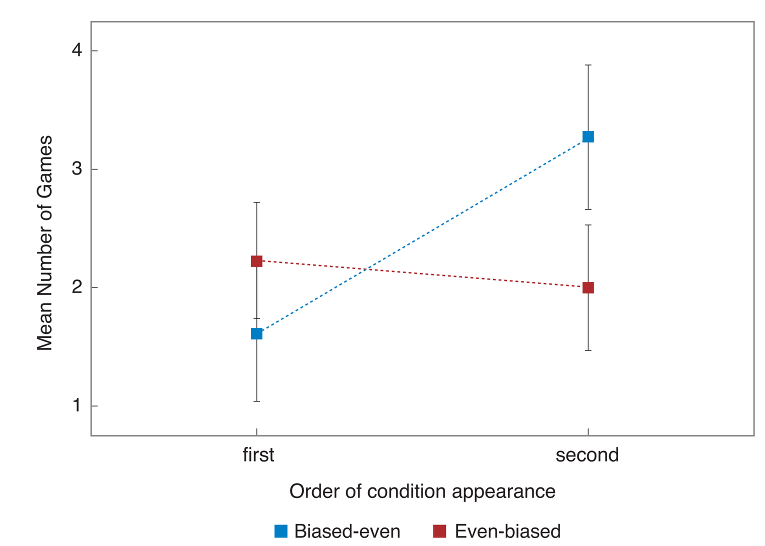

Figure 6

Conclusions

We have shown that people differentiate between structured and unstructured situations, they use this knowledge to guide search, and that prediction problems affect subsequent learning differently. The knowledge that we have modeled is much simpler than that which guides everyday searches, but we are hopeful that our work will provide a step in this direction.