

Abstract knowledge guides search and prediction in novel situations

Russell E. Warner¹, Patrick Shafto¹, Chris Baker², & Joshua B. Tenenbaum²

¹Department of Psychological and Brain Sciences, University of Louisville

²Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology

Abstract

People combine their abstract knowledge about the world with data they have gathered in order to guide search and prediction in everyday life. We present a Bayesian model that formalizes knowledge transfer. Our model consists of two components: a hierarchical Bayesian model of learning and a Markov Decision Process modeling planning and search. An experiment tests qualitative predictions of the model, showing a strong fit between human data and model predictions. We conclude by discussing relations to previous work and future directions.

People combine their abstract knowledge about the world with data they have gathered in order to guide search and prediction in everyday life. Rather than simply treat every search as a novel event, individuals apply knowledge from previous experiences to make predictions about new situations. Using knowledge about distribution of items across instances, such as a prizes in a cracker jack boxes, individuals have expectations about how many times you will have to reach in to get the prize. Because cracker jacks boxes have more popped corn and peanuts than prizes, you might use that knowledge to predict that any given hand-full will only contain popped corn and peanuts. Now imagine an alternative scenario in which you got your cracker jacks box from a discount bin at the factory. If we found a few boxes contained almost 50% prizes, we could generate a different expectation about the process that created the boxes, and use this knowledge to make predictions about what we will pull out of the box next, and what we might expect from the next box. Here knowledge about boxes, and the processes that generate boxes allow us to make predictions about novel situations. In general, knowledge at multiple levels of abstraction allows us to generate expectations and predictions for novel situations.

Recent research has provided insight into some aspects of this ability to transfer knowledge. Hierarchical Bayesian models (HBMs) have been used to model how people can combine past experience with sparse data in a novel situation to make confident and accurate predictions. For example, Kemp, Perfors, and Tenenbaum (2007) showed that hierarchical Bayesian models can capture how children’s experience with with nouns causes them to infer that objects with the same name tend to have the same shape, allowing generalization of novel names from only one or two examples. A separate line of research has focused on prediction in un-

structured stochastic environments (especially so-called bandit problems) and investigated how people develop strategies in these situations (Steyvers, Lee, & Wagenmakers, 2008; Acuña & Schrater, 2008; Gittins, 1979). However, this work has not addressed how people use abstract knowledge to guide reasoning about novel situations in familiar environments. In fact, Steyvers et al. (2008) specifically suggest the incorporation of a learning element into these problems would be a useful extension of the work. In this paper, we propose and test a framework for just such a purpose; showing how people use experience to guide search and prediction.

A Bayesian model of search and prediction in novel situations

Our approach to modeling how people use previous experience to guide inference and search in novel situations combines the previously mentioned hierarchical Bayesian models with a Markov Decision Process (MDP). Hierarchical Bayesian models provide a representation of how experience can affect beliefs at multiple levels of abstraction, allowing predictions about novel situations (Kemp et al., 2007). Markov decision processes provide a model of search and decision making (Russell & Norvig, 2002). We briefly introduce each approach and then describe how to integrate the two approaches, using a simplified problem to develop intuitions.

Instead of cracker jacks boxes, which have predictable mixture proportions, imagine a discount store that sells boxes of cookies that don’t have predictable mixtures. A clerk at the discount store suggests a game in which, if a boy correctly guesses the next cookie drawn from a box, he gets to keep the box. However, the boy’s mother is about to leave, putting the boy at risk of not getting anything. What should he do? In the case of discount store boxes, which are not well-mixed, if the boy can draw one example before guessing, he can dramatically increase his chance of guessing the next cookie correctly. But if the boy were at a regular grocery store, where the boxes were well-mixed, asking to see one sample would give him no information – his chances of guessing correctly remain at 50/50. With our model, we can represent the boy’s abstract knowledge about mixtures of boxes, which in Figure 1 is Level 3 knowledge. And we can further show how this abstract knowledge should inform his decision to draw or guess.

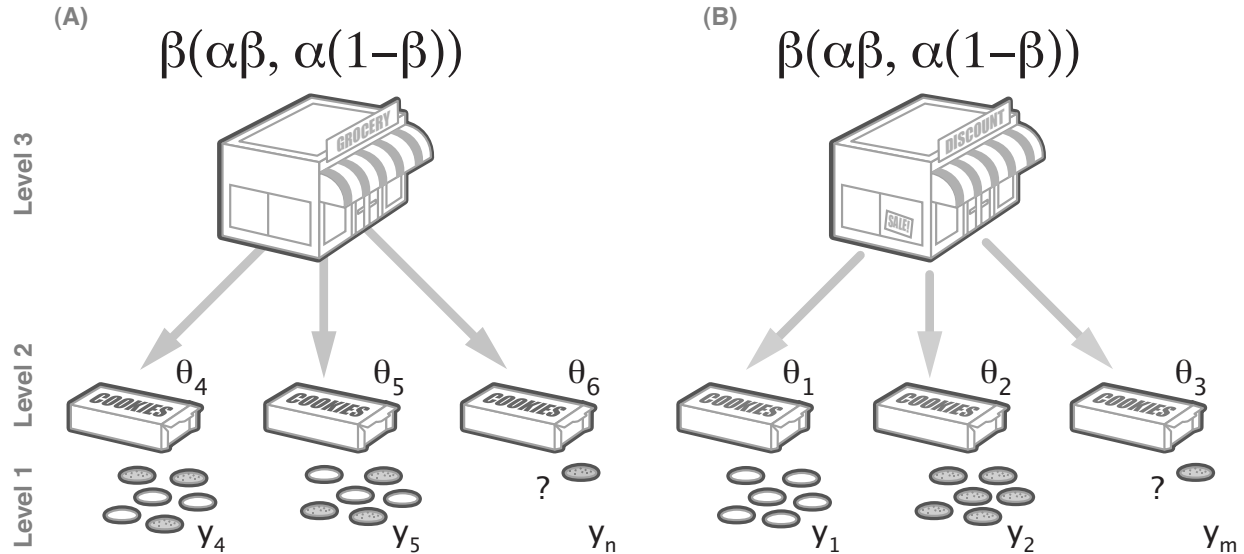


Figure 1: Hierarchical knowledge representation. At the top, the α and β are parameters that represent our confidence about the mixtures of the boxes and the balance between the outcomes. At level 2, we can have knowledge about the proportions in each different box, and at the lowest level, we can have knowledge about specific data (here, cookies).

Modeling abstract knowledge and prediction

To keep our example simple, we will imagine only two kinds of cookies, chocolate chip and sugar, and two places where boxes of cookies can come from. Beliefs about boxes' mixtures are expressed as proportions, and based on samples of cookies from a box, we can make inferences about the proportions of cookies. To incorporate prior experience with boxes of cookies, we imagine that all of the boxes come from some common generating process by which well-mixed boxes go to the grocery store and poorly mixed boxes end up at the discount chain (see Figure 1).

This hierarchical framework allows us to represent knowledge at level 3 which informs and constrains beliefs at level 2. For example, if our previous experience has been with well-mixed boxes, we could infer that boxes from the grocery are most-often well-mixed. Or, based on purchases from the discount store, we conclude that boxes tend to be all one kind or the other. For our example, we will refer to well-mixed boxes as coming from the *even* condition or distribution, and boxes with little variation as coming from the *biased* condition or distribution. If we have a box from the discount store, then the single observation of a cookie allows a strong inference that the box is full of that kind of cookie. However, if we have a box of cookies from the grocery then one observation is unlikely to strongly affect our beliefs.

Formally, we can express beliefs about the contents of the box using a parameter, θ , that represents the proportions. A draw of a single instance, y is then a flip of a coin with the appropriate weight drawn from a Binomial distribution, $p(y|\theta)$. We can express beliefs about the distribution of θ at level 3. The Beta distribution takes two parameters, β which represents the balance between

the two outcomes, and α which represents the strength of those beliefs. Together, these parameters can express a strong belief in the even condition (that produces well-mixed boxes) through a high value of α (> 2), or a strong belief in the biased condition through a low value of α (as it goes to 0).

Our goal is to make a prediction about which cookie is likely to be drawn next from a box, \tilde{y} , given some observed cookies, y , and abstract beliefs, α and β . We can make this prediction by integrating over values of θ ,

$$p(\tilde{y}|y, \alpha, \beta) = \int_{\theta} p(\tilde{y}|\theta) \cdot p(\theta|y, \alpha, \beta) d\theta. \quad (1)$$

Predicting \tilde{y} given θ is like imagining the outcome of the toss of a coin with weight θ . Because of our choice of Binomial and Beta distributions, the integral over θ given y , α , and β is analytically tractable, and turns out to be another Beta distribution, with parameters updated based on the observation y (Gelman, Carlin, Stern, & Rubin, 2003).

For the even distribution, our belief that boxes tend to be 50-50, mitigates the influence of any observation on our beliefs about the contents of the box. However, for the biased distribution, our belief that boxes are all one kind or the other allows a robust inference about the contents of the box, and hence the likely outcome if we were to sample one cookie.

Modeling the search for information

Given different kinds of abstract knowledge, how much information should we obtain before guessing? The answer depends on the rewards (in our example, how much the boy likes sugar cookies), costs (the reduction in reward), and how much information can be gained by continuing to sample. Perhaps the boy's mother is in the

checkout line at the discount store and is leaving with or without the cookies. Not having any cookies at all is a heavy price to pay, so the boy might not sample any.

Markov decision processes formalize the problem of which action to take (i.e. wait for one sample or just guess). We compute the expected value of each decision based on our rewards, costs, and the probability of different possible outcomes. In our example, there may be information to be gained by looking at one or a few cookies, but it has to be traded off against whether the boy will get a box at all if he takes too much time to haggle.




Figure 2: The top row shows two possible distributions that generate boxes of cookies. The left distribution generates boxes that tend to be either all chocolate chip cookies or all sugar cookies, while in the right case, boxes of cookies tend to be evenly mixed. The second row shows typical boxes from each distribution, with the proportion of chocolate chip and sugar cookies in the box. The bottom row shows the possible states of the game. The numbers in each state are the probability of our model guessing at each point with a value of $\rho = 5$. Note that for the left column, the model predicts a strategy of sampling one cookie to see which kind it is, then guessing that kind, while for the right, the model predicts that guessing is the best strategy.

We begin by defining a state $x_{i,j}$ which represents our knowledge about the draws we have already made (see Figure 2). For example, in state $x_{1,0}$ we have drawn one sugar cookie. Because our model of draws is a Binomial distribution with a Beta prior, the probability of guessing correctly $p(\text{correct}|x_{i,j})$ is based on the counts of our

observations smoothed by our prior,

$$P(\text{correct}|x_{i,j}) = \begin{cases} \frac{i+\alpha\beta_1}{i+j+\alpha\beta+\alpha(1-\beta)} & \text{if } i \geq j \\ \frac{j+\alpha\beta_2}{i+j+\alpha\beta+\alpha(1-\beta)} & \text{if } i < j. \end{cases}$$

The expected value of guessing in a state depends on the reward associated with a correct guess. For each state, $V(r|x_{i,j})$ is the total reward minus costs incurred by draws up to that point. The expected value of guessing then is the product of the reward for a correct guess and the probability being correct,

$$V(\text{guess}|x_{i,j}) = p(r|x_{i,j})\text{guess}(r|x_{i,j}). \quad (2)$$

Guessing is valued to the degree that it is both likely to be correct and going to obtain a reward.

The value of sampling can be determined by considering the value of guessing given the possible outcomes of the sampling, weighted by their probability

$$V(\text{sample}|x_{i,j}) = \sum_{\substack{i' \geq i, j' \geq j, \\ i', j' \neq i, j}} V(\text{guess}|x_{i',j'})p(x_{i',j'}|x_{i,j}) \quad (3)$$

where $p(x_{i',j'}|x_{i,j})$ is the probability of going from state $x_{i,j}$ to state $x_{i',j'}$. This is the expected value of guessing in the future.

We formalize the decision to guess or sample using the Luce choice rule (Luce, 1959),

$$p(\text{sample}|x_{i,j}) \propto \left[\frac{V(\text{sample}|x_{i,j})}{V(\text{sample}|x_{i,j}) + V(\text{guess}|x_{i,j})} \right]^\rho \quad (4)$$

where ρ is a parameter that controls how strictly we choose the decision to sample or guess that maximizes the expected reward.

For the biased and even distributions, the model predicts qualitatively different behavior (see Figure 2). In the biased condition, because there is much to be gained from a single observation, the model predicts sampling first. Because this is enough information to be confident about the proportions in the box, it predicts that one sample is enough. Given the even condition, the model predicts that there is no need to search. Because all boxes tend to be well-mixed, the outcome of one or even a few observations is unlikely to affect our beliefs enough to matter.

Learning higher-order knowledge to guide search

In addition to learning about box proportions assuming fixed knowledge about their origins, our approach can be extended to learn about new stores. This extended model captures learning about stores using the priors on the parameters α and β , allowing the model to have uncertainty about what stores' cookie boxes are like. These priors will allow us to make inferences from experience with one store to stores in general.




Figure 3: Inferred values of alpha by trial for two different orders. Model simulations show that when the biased condition is presented before the even condition, the model has greater difficulty in learning in the even condition.

Formally, to learn at this more abstract level, we place a distribution on our beliefs about the stores. We want to express a belief about the kinds of boxes a store will have – does it have one kind or the other, or both? To express uncertainty about the distribution of boxes, we must place distributions on α and β , and learn the values of these parameters for each factory, and beliefs about factories in general. The strength parameter is α and a natural prior for this is an exponential distribution. The exponential distribution has a single parameter, γ , and to express weak neutral beliefs we place a normal prior on $\log(\gamma)$ with mean 1 and variance 5. Similarly, we must place a prior on β and learn about the value of this parameter. A natural choice is another Beta distribution, with parameters κ expressing the strength, and δ expressing the balance. Because both of the situations we consider here are balanced (symmetric), we fix $\delta = 0.5$. As previously with α , we place an exponential distribution on κ , with the parameter set to 1, to express weakly neutral beliefs about the value of β . Together, κ , δ , and γ express neutral beliefs about stores.

To transfer knowledge across stores, we make inferences about the parameters γ and κ at level 4. In the extended model, the probability of a box with proportion θ is,

$$p(\theta|\gamma, \kappa, \delta, \alpha, \beta) \propto p(\theta|\alpha, \beta)p(\alpha|\gamma)p(\beta|\kappa, \delta)p(\gamma)p(\kappa). \quad (5)$$

Given some observed mixture proportions θ , we would like to make inferences about γ and κ . This can be done using Bayes rule; however, because α and β are both unknown, we must integrate over possible values,

$$p(\gamma, \kappa|\theta, \delta) \propto \int_{\alpha} \int_{\beta} p(\theta|\alpha, \beta)p(\alpha|\gamma)p(\beta|\kappa, \delta)p(\gamma)p(\kappa). \quad (6)$$

In this paper, we approximate the integral using a coarse grid approximation over the values of α , β , γ , and κ .

The import of this extended model is that it allows us to investigate how experience with one factory affects

beliefs about, and therefore search and prediction in another. For instance, we may ask how experience with a store that has homogeneous boxes would affect learning about a new store that has well-mixed boxes, and vice versa. Figure 3 shows the results of model simulations for two different presentation orders: even-biased, and biased-even. The model predicts that the even condition is more difficult when the biased condition is presented first. This is because it is a larger coincidence to find boxes of cookies that happen to be all one kind or the another, while evenly mixed boxes can happen in a variety of ways. This coincidence creates a strong bias that novel experiences with evenly mixed boxes has difficulty overcoming.

Experiment: The Big Urn

In the following experiment, we will test the two model predictions discussed above. First, people should be able to use knowledge about even and biased distributions to guide their search. In the case where people experience a biased distribution (as seen in Figure 2, column 1), their search strategies will converge on drawing once and then guessing. In the case where they experience an even distribution, they should converge on a strategy based simply on guessing. Second, the model predicts greater difficulty in transferring from the biased to the even condition than vice versa.

Methods

Participants. Twenty-seven students participated in this experiment in exchange for partial course credit or chits.

Procedure. Participants were asked to play a game on a computer. The object was to guess the color of a 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.”) and then shown a new urn.

One *game* consisted of working through a course of 20 urns with proportions drawn from a beta distribution with the same parameters. Scores were tracked so that participants could note their performance and they were asked to achieve a score that was better than an “escape” score at the end of a sequence of urns. These escape scores were derived from the expected value at the end of the 20 urns using the correct strategy, minus one-half of a standard deviation: 33 for the biased condition, 26 for the even condition. Once a participant beat this escape score, she was transferred to a new condition in which the mixtures of the urns were generated by a new distribution. Participants were told that these conditions would be different. The two conditions, biased (Beta(1,1)) and even (Beta(15,15)) were presented in random order.

We note that by setting the escape score at the expected value of the optimal strategy minus one-half of




Figure 4: The probability of guessing in the even condition (A,B) and biased conditions (C,D) calculated for subjects and the model. Each graph shows the probability of guessing during the first turn or after one observation. Both people and the model tend to guess right away in the even condition, while both tend to draw once, then guess in the biased condition.

the standard deviation, it became possible for participants to escape games by chance. For this reason, participants who drew less than three samples total over the course of at least 60 urns or participants who draw the maximum sample over 60 urns or more were thrown out. There lack of any adjustment to their strategies belied a misunderstanding about the nature of the game. Fourteen participants were removed due to this consideration.

Results

Our first set of analyses focus on people’s search strategies once they have learned the distribution. When the distribution of the urns is known, the model predicts different search and prediction strategies (see Figure 2). For the even condition, the model predicts no draws. For the biased condition, the model predicts drawing one ball and guessing whichever color was drawn. When we analyze people’s behavior in the last half of their first condition to assess whether people in biased and even conditions learned different strategies. Figure 4 shows the probability of guessing on the first chance, and after drawing. The model predictions shown are based on a value of $\rho = 5$, but the effects are robust across a range of values. The probability of the initial draws are significantly greater in the biased condition than in the even condition, as predicted. These results suggest that people learned which strategy was appropriate for each condition.




Figure 5: In agreement with the model, individuals who were shown the biased condition first took longer to adjust their strategy in the even condition.

By breaking games up into four stages, we can see the development of strategies (see Figure 6). Both groups initially behaved without commitment, guessing or drawing alternately. Following that initial stage, both groups shifted more dramatically toward distinct strategies for their first condition. Once subjects achieved their escape score and switched conditions, they appear to initially continue to use their first strategy before shifting toward a more appropriate strategy for their second condition.

Our second set of analyses focuses on learning at the third level of the hierarchy: learning about the process that is generating urns. In the initial condition, people should not have strong expectations about the third level of the hierarchy. However, by the time they arrive in the second condition, their experience should bias their expectations. Figure 5 shows that individuals coming into the even condition second played, on average, a larger number of games, indicating that it was more difficult to escape this second condition $t(17) = 5.956, p < .001$. One possible explanation for the difference in the mean number of games played in the second condition is policy transfer. Subjects simply continued to apply their initial strategy to the novel condition, and probability dictated that they were less likely to escape the even condition by chance. Policy transfer would not, however, explain why there is an evident shift in strategies after the change in conditions for both groups as seen in Figure 6, after the dashed line. Our model indicates that it is more difficult for subjects to shift from the biased condition, because it is a much larger coincidence that urns would be of mostly one color than evenly mixed. This provides a more accurate explanation of the data.




Figure 6: Development of strategies over the course of the experiment. Initially, all participants started out roughly even between drawing and guessing after zero and one observation. At the end of the first game, people had learned the predicted strategy for each condition. After participants switched conditions, they initially reflect the strategy they started with, but eventually converge toward the strategy predicted for that condition.

Discussion

Everyday reasoning requires systematic searches for information in novel situations, and prior experience plays a critical role in guiding this process. We have presented a computational framework for modeling how abstract knowledge guides search in novel situations. In our experiment, we showed that people use abstract knowledge about distributions to guide search and prediction about a simple guessing game scenario. Further, the results suggest that learning about some situations hinders development of appropriate search strategies in other situations.

Much research has been dedicated to the question of how people make decisions in uncertain environments. Research on decision making in highly stochastic environments has suggested that people may not tend to choose options that maximize rewards. We believe a more interesting avenue lies not in unstructured environments, but in structured scenarios like the ones we face everyday. Searching for a reference on the internet, for instance, goes on in a much more structured environment than that which is captured by bandit problems, and it is in these structured environments that people succeed.

We believe that our approach using hierarchically structured knowledge to guide search and prediction presents an interesting step toward capturing people’s abilities to transfer knowledge across situations to guide search. However, important issues remain for modeling real-world behavior. The knowledge that we have modeled is much simpler than that which guides everyday searches. Similarly, our approach to planning requires that we have full knowledge about possible states and

associated rewards in the world. These represent two important challenges for future work, but we are hopeful that our work will provide a step in this direction.

References

- Acuña, D., & Schrater, P. (2008). Bayesian modeling of human sequential decision-making on the multi-armed bandit problem. *Proceedings of the 30th Annual Conference of the Cognitive Science Society*.
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2003). *Bayesian data analysis, second edition*. New York: Chapman & Hall/CRC. Hardcover.
- Gittins, J. (1979, Jan). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)*, 41(2), 148–177.
- Kemp, C., Perfors, A., & Tenenbaum, J. B. (2007, May). Learning overhypotheses with hierarchical bayesian models. *Developmental Science*, 10(3), 307–321.
- Luce, R. D. (1959). *Individual choice theory: A theoretical analysis*. New York: Wiley.
- Russell, S. J., & Norvig, P. (2002). *Artificial intelligence: A modern approach (2nd edition)*. New Jersey: Prentice Hall.
- Steyvers, M., Lee, M., & Wagenmakers, E. (2008, Jan). A bayesian analysis of human decision-making on bandit problems. *Manuscript in press for Journal of Mathematical Psychology*.