
A probabilistic model of eye movements in concept formation

Jonathan D. Nelson
Cognitive Science Dept., UCSD
La Jolla, CA 92093-0515
jnelson@cogsci.ucsd.edu

Garrison W. Cottrell
Computer Science and Engineering, UCSD
La Jolla, CA 92093-0404
gary@cs.ucsd.edu

Abstract

It has been unclear whether optimal experimental design accounts of data selection may offer insight into evidence acquisition tasks in which the learner's beliefs change greatly during the course of learning. Data from Rehder and Hoffman's (2003, 2005) eye movement version of Shepard, Hovland, and Jenkins's (1961) classic concept learning task provide an opportunity to address these issues. We introduce a principled probabilistic concept-learning model that describes the development of subjects' beliefs on that task. We use that learning model, together with a sampling function inspired by theory of optimal experimental design, to predict subjects' eye movements on the active learning version of that task. Results show that the same rational sampling function can predict eye movements early in learning, when uncertainty is high, as well as late in learning when the learner is certain of the true category.

1 Introduction

Learning about the world and querying the world are both central to competent behavior in many situations. Consider the process by which a physician diagnoses a patient's illness. It is not possible to conduct every relevant medical test. The physician must conduct tests that are expected to provide the best improvement in the ability to correctly diagnose and treat the patient. The new information obtained by the most recent test is used to better select the next test, in an iterative process. During this process, the physician's beliefs about the patient's illness may change drastically. Suppose a patient has a deep cough. The presumed cause of the cough may be nasal drip, caused by an allergy, and anti-allergy medication might be an appropriate treatment. However, suppose a chest X-ray were to show abnormal appearance of a lung. In this case, pneumonia or tuberculosis would become more probable explanations of the illness, and the most appropriate diagnostic tests and treatments would change drastically.

What is known about how human beliefs develop as new information is obtained, and about how people's queries relate to beliefs at the time those queries are made? Can Bayesian statistics help explain how beliefs develop, and can theory of optimal experimental design help explain people's queries? Research on Bayesian reasoning, and on Bayesian theories of concept learning, suggests that at least in situations whose essential character is known in advance, Bayesian statistics may provide a useful language for understanding human belief revision. If beliefs can be described in a Bayesian framework, then principles of optimal experimental design (Savage, 1954, chap. 6, pp. 105-119) can be used to identify potentially

useful queries. A variety of articles suggest that this framework may provide a useful foundation for a theory of human queries, as well (Nelson, in press).

The present article concerns learning and selective attention in Shepard, Hovland and Jenkins's (1961) concept learning task. In this task, subjects learn to categorize objects where each object has three binary dimensions, such as size (large or small), shape (square or circle), and color (black or white). Accounts of learning in this task have largely been based on one or more mechanistic theories, rather than on rational principles (for a review, see Nosofsky, Gluck, Palmeri, McKinley, & Glauthier, 1994; also see Zhang & Cottrell, 2005), with the notable exception of Anderson (1991). Many of these accounts of learning have as a critical component a selective attention mechanism, describing when during learning a subject attends more to the size dimension, when to color, and so on, as a function of the type of concept being learned. However, data on selective attention per se have not been available. In order to directly measure subjects' attentional focus, Rehder and Hoffman (2003, 2005) separated the stimulus dimensions in space, in an eye movement version of the task. In this manner, subjects' eye movements provide direct evidence for what stimulus dimension a subject is attending to at each point in learning.

One interesting property of Shepard et al.'s (1961) task is that learning takes place over an extended period of time, involving dozens or hundreds of trials, and that subjects' beliefs can change drastically over the course of learning. In this article, we introduce a probabilistic model of learning to help explain how beliefs develop as subjects perform this task. One key property of our model of learning is our extension of earlier probabilistic learning and inference models (Edwards, 1968; Tenenbaum, 1999; Tenenbaum & Griffiths, 2001) to cases of imperfect memory or noisy data. We then evaluate whether this probabilistic model of learning can offer insight into eye movements in this task, not by mechanistic (process) theories of selective attention, but rather by theories of optimal experimental design. More specifically, we introduce a model of eye movements (queries) in which the subject's eye movements are optimal at each point in learning, but in which the specific eye movements change during learning as a function of the development of knowledge during the task.

The rest of the article is structured as follows. In the next section, we discuss Shepard et al.'s (1961) classic task, our Bayesian model of subjects' learning on that task, how model performance compares with human data, and novel predictions derived from our model. In the following section we discuss Rehder and Hoffman's (2003, 2005) eye movement version of Shepard et al.'s task. This section introduces our eye movement model, and shows how this eye movement model can provide a rational explanation of much of Rehder and Hoffman's data. The final section of the article summarizes our findings, what we take them to imply, and lays out issues for future research.

2 The task; our probabilistic concept learning model





Shepard et al.'s (1961) task involves eight objects, formed by taking every combination of three binary stimulus dimensions, such as size (large or small), shape (square or circle), and color (white or black). Concepts are comprised of a subset of those eight objects.¹ There are 256 (2^8) possible concepts, although behavioral experiments have mainly considered only the 70 ($8 \text{ choose } 4$) concepts of size 4. (Feldman, 2000, is a notable exception.) Shepard et al. (1961) defined taxonomy of six types of concepts; Table 1 gives illustrative examples. Type I concepts are defined by a single stimulus dimension, for example "large", or "circle". Type II concepts are defined with reference to two stimulus dimensions; an example is "large square or small circle." Types III, IV, and V concepts require attention to all three stimulus dimensions for correct classification, but above-chance performance can be achieved by attending a single dimension. One Type IV concept is "large, but also the small white circle and excluding the large black square." Type VI concepts require attention

¹ An equivalent alternative is to describe concepts as separating two types of objects. Here it is easier to discuss whether an object is consistent with a concept or not.

to all three stimulus dimensions; they are usually learned by memorizing the specific set of objects that is consistent with them.

In an individual trial, an object is drawn at random from the set of eight objects, the learner judges it as consistent or inconsistent with the true, hidden concept, and the learner is given feedback. The objects are usually drawn without replacement to form a *block* of trials, such that objects appear in a random order but with the constraint that each object appears exactly once (or twice) in a block. Shepard et al. (1961) found that across several experimental manipulations, criterion performance was reached most quickly on Type I concepts. The number of trials to criterion performance was ordered as follows: Type I < II < III ≈ IV ≈ V < VI. (Types III, IV, and V were not reliably different from each other). Nosofsky et al. (1994) obtained similar results, with 44.0, 85.4, 121.6, 127.0, 133.8, and 189.2 trials to criterion performance, for concept Types I through VI, respectively.

Table 1. Example concepts

Type	Objects	Err.
I		8.2
II		31.2
IV		36.9
VI		70.6

Note. Err. column is average number of errors in Rehder & Hoffman's (2005) study.

2.1 Introduction to the probabilistic learning model

In this section we introduce a probabilistic model of how the learner's beliefs develop over the course of many trials, in Shepard, Hovland, and Jenkins's (1961) concept learning task. This model is similar in spirit to Anderson's (1991) model, but makes fewer assumptions as to the cognitive processes underlying inference. We explicitly specify a generative model, which is meant to capture the learner's beliefs about the task. Learning, in our model, is optimal Bayesian inference with respect to that generative model.

We use probability theory to describe the learner's prior beliefs, and how those beliefs change in response to new information. We use standard probabilistic notation, in which random variables begin with capital letters, and specific values of random variables begin with lowercase letters. For example, C represents the learner's beliefs about the a priori probability of each of the concepts, represented in a *prior probability distribution* over all concepts c_i . (In our probability model, which includes 70 possible concepts, C is a vector with 70 probabilities that sum to one.) $C=c_i$ represents the event that C takes the value c_i , for instance that the true concept is "small". The learner typically wishes to know $P(C=c_i|X=x)$, the *posterior probability* of a particular concept given a particular observed datum x . For instance, the learner may wish to know the probability that "small" is the true concept, given that in their first trial of learning, they observe that a small black circle is consistent with the true concept. [For concise notation in the article, we will use lowercase letters to denote specific values taken by random variables, for instance c_i , rather than $C=c_i$.] The posterior probability of the concept "small" is not observed directly. Rather, $P(c_i|x)$ is calculated according to Bayes' theorem, with reference to both the prior probability of the concept "small", $P(c_i)$, and the a priori *likelihood* of observing the small black circle as a positive example, if the true concept were "small," $P(x|c_i)$. By Bayes' (1763) theorem, the posterior probability of a particular concept is calculated as follows:

$$P(c_i | x) = \frac{P(x | c_i) * P(c_i)}{P(x)}.$$

The simulated model learner obtains the posterior probability of every concept in this way. $P(x)$, the probability of the observed datum, is a normalizing constant that is computed by summing $P(x | c_i) * P(c_i)$, the numerator of the previous term, over all concepts c_i :

$$P(x) = \sum_{c_i} P(x | c_i) * P(c_i).$$

2.2 Capturing the learner’s intuitions with prior probabilities

If the learner had no intuition that certain concepts were more plausible than others, each concept could be given the same prior probability. However, in the case of Shepard et al.’s (1961) task, some concepts, such as “large”, seem quite a bit more plausible than other concepts, such as random collections of objects. In a probabilistic modeling framework, we seek to capture the learner’s intuitions by giving more plausible concepts higher prior probability, and less plausible concepts lower prior probability.

In the current article, our primary focus is on understanding what Rehder & Hoffman’s (2005) subjects believed, so as to most accurately understand subjects’ eye movement behavior, rather than on providing a complete explanation for why subjects had particular beliefs. We therefore set the prior probability of each concept of a particular type to be inversely proportional to the average number of errors committed by Rehder and Hoffman’s study participants (“Err.” column in Table 1) when learning that type of concept, ε_i , raised to a power, λ , a free parameter in our model:

$$P(c_i) \propto (\varepsilon_i)^{-\lambda}$$

Choice of this particular function is relatively arbitrary; it serves to provide spread in the prior distribution. If $\lambda=0$, all concepts have equal prior probability; if λ is large, Type I concepts have most of the prior probability. All concepts of a particular type, for example the Type I concepts “small” and “circle”, receive the same prior probability. Our model only assigns non-zero prior probability to the 70 concepts that include exactly 4 of the 8 objects. Because Type III, IV, and V concepts are not reliably different from each other in the speed at which they are learned (Shepard et al., 1961; Nosofsky et al., 2004), and because Rehder and Hoffman’s (2005) data do not include concepts of Types III or V, our model gives Type III, IV, and V concepts the same prior probability.

Note that our primary goal in setting priors in the present concept learning model is to capture the learner’s beliefs as accurately as possible, so as to predict the learner’s eye movements. While the errors committed by Rehder and Hoffman’s (2005) subjects provide a reasonable basis for understanding what the learner’s prior beliefs are, and serve as a reasonable basis for a model that describes the process of learning and eye movements, those errors do not constitute an explanation of the ultimate source of the learner’s beliefs about each type of concept’s plausibility.

A number of a priori criteria that do not depend on empirical data could be tested as providing a basis for prior probabilities in the model. One such criterion is the number of stimulus dimensions that must be attended, on average, to correctly classify an object, if the true concept were known with certainty. Shepard, Hovland and Jenkins (1961) discussed this criterion. Type I concepts require a single dimension; Type VI concepts require three dimensions, and other types of concepts require intermediate numbers of dimensions. Another criterion, which Feldman (2000, 2003) has investigated in relation to concepts’ learnability, is Boolean complexity. In future work we intend to investigate whether some of these a priori criteria might provide a suitable basis for the prior probability distribution.

2.3 Belief updating: the likelihood function

Although the assignment of prior probabilities is frequently thought to be the most important part of a Bayesian model, the likelihood function—which describes how new information changes beliefs—is equally important. In the case of our concept-learning model, the likelihood function should describe the learner’s beliefs about how the example data, including their labels (in the form of the feedback given after each trial), are generated.

The likelihood functions of classic Bayesian reasoning models (Edwards, 1968), and recent Bayesian concept learning models (Tenenbaum, 1999; Tenenbaum & Griffiths, 2001;

Nelson, Tenenbaum & Movellan, 2001) would appear to be applicable in this task.² However, these models have the property that if an observed datum is incompatible with a particular hypothesis, that hypothesis is eliminated from consideration. Under these models, one exposure to every possible object uniquely identifies the true concept, by eliminating all other concepts. If applied to the concept-learning task, these types of models would predict that after a single block of trials, in which each object is viewed once with feedback, the learner would know the true concept with complete certainty. This result would hold irrespective of the type of concept being learned, and irrespective of the prior probabilities of each concept, such as whether the concept had prior probability 1/10 or 1/1000. Empirical results strongly contradict this implicit claim of those models, however. In fact, no concept is learned after a single block of trials, and there are clear distinctions wherein some types of concept (such as the Type I concept “square”) are learned more quickly than others (such as the Type VI concept “large black square or small white square or large white circle or small black circle”). The limitation underlying the classic likelihood functions is that a single observation completely eliminates incompatible concepts. For instance, observation of the large white circle as a positive example would immediately eliminate concepts such as “square”, “small”, and “black circle or white square.”

To account for the empirical finding that a single example observation does not eliminate incompatible concepts, we assume that the subjects do not encode the examples perfectly in memory. To model this, we include two generators of data. One data generator is uninformative, and labels objects as consistent or inconsistent with the true concept randomly with equal probability. The other data generator is veridical, and always correctly labels example objects. We use a parameter μ , between zero and one, for the probability that the informative data generator generated a particular observation. When $\mu=1$, the model behaves in a manner analogous to earlier probabilistic concept learning models (Tenenbaum, 1999; Tenenbaum & Griffiths, 2001), and learns the task in a single block of trials, irrespective of the type of concept and prior distribution. When $\mu=0$, the likelihood function is uninformative and no learning takes place. Hence, μ can be thought of as a learning rate (Figure 1). From a generative model standpoint, if the veridical data generator is selected in a particular trial, then there are 8 equally probable observations (4 positive and 4 negative); if the noise generator is selected, then there are 16 equally probable observations (8 positive and 8 negative). The likelihood of a particular observation x (a combination of an object and its label), given a particular concept (c_i), where g_v and g_u are the veridical and uninformative data generators, respectively, is given by

$$P(x | c_i) = P(x | c_i, g_v) P(g_v) + P(x | c_i, g_u) P(g_u),$$

with $P(g_v)=\mu$, for the veridical generator and $P(g_u)=1-\mu$, for the noise generator. In the case where all stimulus dimensions are observed at once, as in the classic concept learning task, this reduces to $(1/8)*\mu + (1/16)*(1-\mu)$, if x is consistent with the concept c_i , and $(1/16)*(1-\mu)$ otherwise. To approximate the learner’s beliefs, the likelihood assumes that the observations are conditionally independent of each other, given the true concept. The psychological premise of this assumption is that the learner does not know that the objects




Figure 1. Belief update if $\mu=1$, top; if $\mu<1$, beliefs change less but in same direction, bottom.

² Tenenbaum (1999; Tenenbaum & Griffiths, 2001) focused on situations where the learner encounters only positive examples. However, his models can update beliefs based on negative instances, provided the likelihood function specifies how those instances are generated.

are drawn without replacement in each block of trials, and believes that objects are drawn randomly throughout learning.³ Bayes' theorem is used to update beliefs.

2.3.1 Belief updating illustrated with an example

The likelihood function is illustrated with a simplified example, where the concepts “square” and “large” each have prior probability 0.50, and the other concepts have prior probability zero. (Except for the simplified prior probability distribution, this example is identical to the full concept-learning model.) Let $\mu=0.20$. We assume that the learner is exposed to all stimulus dimensions, as is the case in most versions of Shepard, Hovland and Jenkins's (1961) task, when whole objects, all three attributes of each object, are presented to the subject. The likelihood of a particular observation, if all stimulus dimensions are attended, is 1/8 given a particular concept and the veridical generator, and 1/16 given a particular concept and the noise generator, as in the full-scale model.

Suppose that on a particular trial, the learner sees the large white circle as a positive example of the concept. The learner believes that with probability 0.20, the observation is veridical; and with probability 0.80, the observation is uninformative. In the equations below we use x to denote that the observed object is the large white circle, labeled as a positive example of the unknown true concept. There are 16 possible observations, corresponding to each of the 8 objects as positive and negative examples of the true concept. We use g_v to denote that the veridical generator of data is selected, and g_u to denote that the uninformative generator is selected. We use c_s to denote that the true concept is “square,” and c_l to denote that the true concept is “large.” If the observation is from the veridical data generator, then the posterior probability of “square” is 0 :

$$P(c_s | x, g_v) = \frac{P(x | c_s, g_v) * P(c_s | g_v)}{P(x | g_v)} = \frac{P(x | c_s, g_v) * P(c_s)}{P(x)} = \frac{0 \cdot \frac{1}{2}}{\frac{1}{16}} = 0$$

[$P(c_s | g_v) = P(c_s)$, and $P(x | g_v) = P(x)$, because $P(c_s)$ and $P(x)$ do not depend on G .] Similarly, the probability of the concept “large,” given this observation, would be 1:

$$P(c_l | x, g_v) = \frac{P(x | c_l, g_v) * P(c_l)}{P(x)} = \frac{\frac{1}{8} \cdot \frac{1}{2}}{\frac{1}{16}} = 1$$

If the observation is uninformative, then posterior probabilities of “square” and “large” would each be 0.5, unchanged from the priors. Consider the posterior probability of “large”:

$$P(c_l | x, g_u) = \frac{P(x | c_l, g_u) * P(c_l)}{P(x)} = \frac{\frac{1}{16} \cdot \frac{1}{2}}{\frac{1}{16}} = 0.5$$

The learner, of course, does not know whether the veridical or uninformative generator of labels was selected on this trial, and therefore must average over both possibilities. The posterior probability of the concept “square”, given the observed datum x :

$$P(c_s | x) = \mu P(c_s | g_v) + (1 - \mu) P(c_s | g_u) = 0.2 (0) + 0.8 (0.5) = 0.4.$$

Similarly, $P(c_l | x)$ the posterior probability of “large”, given the observed datum x , is 0.6.

³ We believe that this discrepancy between our generative model of this task, and the actual task, is minor. However, it would be interesting, in a future experiment with human subjects, to actually draw example objects independently, with replacement, throughout learning.

2.3.2 Comment on μ

We introduced μ as reflecting imperfect encoding of examples in memory on the part of the learner. However, μ could be described in different ways, as well. Note that a perfect inferential agent, with perfect perception and memory, would behave in the same way as our likelihood function, if the agent believed that individual examples (such as the large white circle, if drawn on a particular trial) were correctly labeled with probability μ , and randomly labeled (positive or negative by flip of a fair coin) with probability $(1-\mu)$. If this were the case, situations where $\mu=1$ would reflect cases where the learner has perfect confidence in the way examples are labeled. To the extent that a particular learner has $\mu<1$, that learner is skeptical about the process by which examples are labeled. Classic work on ‘conservatism’ in Bayesian reasoning (Edwards, 1968) would suggest that even in situations with no memory component, μ is likely to be less than 1. Another way of describing our likelihood function would be to state that we do not know whether a particular learner remembers a particular example, but that we believe learners remember examples with probability μ . In that case, our concept learning model might not be describing the beliefs of any particular learner, but rather would be describing the mean of all learners’ beliefs. The computation of posterior probabilities does not depend on which explanation of μ one favors.

Our learning model’s likelihood function illustrates how imperfect encoding or memory, or the learner’s belief that the data are noisy, can be modeled in a generative framework of optimal inference. This likelihood function behaves in the same manner as that of Oaksford and Chater (1998), and a similar manner as *conservative* likelihood functions from classic research on Bayesian reasoning (Edwards, 1968). The key distinction is that whereas previous authors viewed apparent conservatism as an error, or suboptimal property of human inference, our generative formulation shows how in the case where a learner is skeptical that data are noisy, being conservative in belief updating, which our learner is if $\mu<1$, is the optimal Bayesian behavior.

2.4 Simulation of the probability model’s learning

Does our model account for results in the classic non-eye movement concept-learning task, in which all three stimulus dimensions are presented in every trial? We set $\lambda=5$ and $\mu=0.10$ in this article. (We chose those values by hand. The qualitative pattern of results does not strongly depend on the exact parameter values used. In future work we intend to optimize these parameters quantitatively.) Figure 2, top panel, shows how the model’s classification performance improves during the course of learning, for each type of concept. Qualitatively, the most important result is that the ordering observed in empirical investigations (e.g. Shepard et al., 1961; Rehder and Hoffman, 2005) is preserved, with high performance on Type I concepts achieved most quickly, followed, in order, by Types II, IV, and VI. Because the concept-learning model explicitly represents beliefs about the probability of each concept, at each point in learning, it will also prove possible to query the model learner in ways that go beyond available empirical data on human learning. Note that although the probability model that our model learner uses to calculate posterior probabilities includes Type III and V concepts, for easiest comparison to Rehder and Hoffman’s (2005) results, we plot only Types I, II, IV, and VI in the charts in this article.

2.4.1 The model learner’s probability of error

The goal of studying concept learning is largely to know how people’s *beliefs* change through experience. However, evidence for belief change has largely been in the form of *error rates* during learning, so it is important to test the model learner’s error rates. Use of probability matching response functions, rather than optimal winner-take-all response functions, may be common in human subjects (but see Shanks, Tunney, & McCarthy, 2002). We model error rates (Figure 2, top panel) with the assumption that subjects probability match in their responses, picking a response proportional to their estimated probability that the response is correct.

Let x represent an unlabeled observation, such as the large white square, before the learner is given feedback during a particular trial.⁴ Here and elsewhere in the paper, A is a binary random variable that indicates whether or not the correct response is “yes” in a particular trial, e.g. whether the random object in the present trial is a positive ($A=1$) or negative ($A=0$) example of the unknown true concept. $P(\alpha)$ is shorthand for $P(A=1)$. The learner’s belief that a particular unlabeled observation x is consistent with the true concept is:

$$P_t(\alpha | x) = \sum_c P(\alpha | x, c) P_t(c)$$

In a given trial, the learner responds “yes” with probability $P_t(\alpha|x)$. The learner’s probability of error for a particular observation x is the probability of a false positive or miss (Table 2):

$$P_t(\text{error} | x) = 2 P_t(\alpha | x) [1 - P_t(\alpha | x)]$$

The model learner’s overall probability of error is computed by averaging over all possible observations x :

$$P_t(\text{error}) = \sum_x P(x) P_t(\text{error} | x)$$

Figure 2, top panel, gives the model learner’s error rates during learning, for Type I, II, IV, and VI concepts; the bottom panel gives error rates from Rehder and Hoffman’s (2005) data. The model learner shows fastest reduction in error rates for Type I concepts, consistent with Rehder and Hoffman’s data. The model predicts that Type IV concepts will initially have lower error rates than Type II concepts. This phenomenon may result from the high similarity between Type IV and Type I concepts, such that above-chance performance on Type IV concepts can be obtained by responding in accordance with appropriate Type I concepts. It is hard to tell whether this prediction is consistent with the data, which are fairly noisy during early learning. The model predicts that Type II and IV concepts will have relatively similar error rates throughout learning, and that Type II concepts will be mastered first. These predictions appear consistent with the data. Finally, the model learner shows the slowest reduction in error for Type VI concepts. Nonetheless, it appears that the model learner underestimated the relative difficulty of Type VI concepts, at least for Rehder and Hoffman’s subjects. (Nosofsky and Gluck’s, 1989, data show similar error rates for Type IV and VI concepts, even late in learning.)

Table 2. Probability matching illustrated

Response	True state	
	$A=1$ $P_t(\alpha x)$	$A=0$ $1 - P_t(\alpha x)$
“yes” $P_t(\alpha x)$	hit $[P_t(\alpha x)]^2$	false positive $P_t(\alpha x) [1 - P_t(\alpha x)]$
“no” $1 - P_t(\alpha x)$	miss $[1 - P_t(\alpha x)] P_t(\alpha x)$	correct rejection $[1 - P_t(\alpha x)]^2$

Note. In each cell, the top line gives the event of interest; the bottom line gives its probability of occurrence. The learner believes the probability that the observation x is a positive instance of the true concept is $P_t(\alpha|x)$.

⁴ In our model of the classic task, we assume that each observation x is of all three stimulus dimensions of an object. Generally, x could be an observation that results from querying a subset of stimulus dimensions.

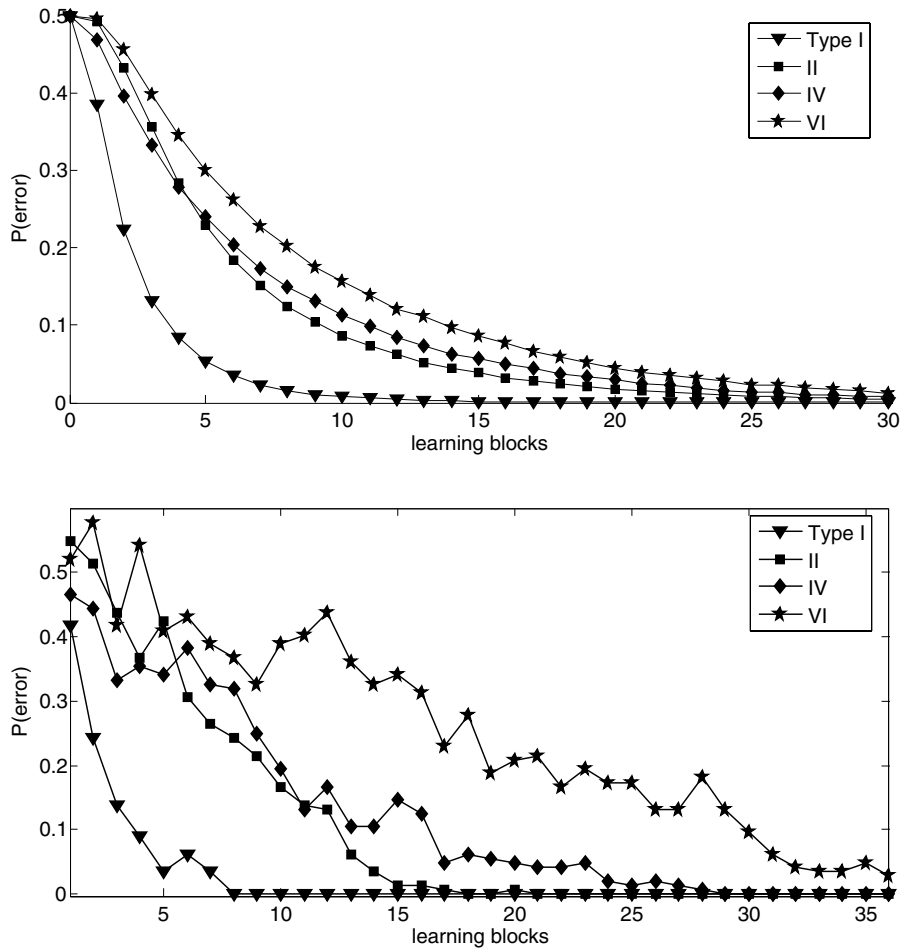


Figure 2. Error rates through learning. Top: model learner. Bottom: data from Rehder & Hoffman (2005), courtesy Bob Rehder and Aaron Hoffman .

2.4.2 The model learner's beliefs about the true concept

Intuitively, although it is needed to assess the model's fidelity to behavioral human performance on the task, a probability matching response function seems a convoluted way to measure the development of beliefs. An advantage of testing our model learner is that beliefs, as measured by the probability distribution $P_t(C)$, per se can be queried at each moment in learning, in addition to probability of correct response. Figure 3 gives the posterior probability of the true concept, over the simulated learning blocks for four different concept classes. Before any learning has taken place, because the prior distribution is relatively peaked ($\lambda=5$), almost all of the model learner's belief is concentrated in the six Type I concepts. (Entropy of prior beliefs in our model is 2.67 bits, barely larger than 2.59 bits, the entropy of a hypothetical model with six equiprobable concepts.) During the course of learning, the probability of the actual concept increases, irrespective of its type, following the ordering of Type I>II>IV>VI.




Figure 3. The model learner’s posterior probability of the correct concept.

2.4.3 The model learner’s uncertainty, throughout learning

While it is difficult to chart the model’s beliefs about the probability of each of the 70 concepts at each point in learning, it is desirable to understand the model learner’s subjective level of uncertainty over all the concepts, at each point. Shannon entropy (Shannon, 1948; Cover & Thomas, 1991) is well suited to quantifying this overall level of uncertainty, which is plotted in Figure 4:

$$H_t(C) = \sum_{c_i} P_t(c_i) \log_2 \frac{1}{P_t(c_i)}$$

An interesting, unanticipated result is that for non-Type I concepts, the model learner’s subjective uncertainty (Shannon entropy) actually increases in early learning. This increase in uncertainty parallels the increase in aggregate probability of the non-Type I concepts. Concepts of Types II, III, IV, V, and VI begin with aggregate probability 0.6%, increasing after five blocks of learning to 75.3%, 63.2%, or 74.2% when the true, hidden concept is Type II, IV, or VI, respectively.




Figure 4. Model learner’s uncertainty (Shannon entropy) about the true concept.

It would be interesting, and potentially feasible, to test our model’s claim that uncertainty actually increases during early learning of non-Type I concepts. These tests should employ a modified task measuring knowledge of the true concept per se, rather than correct

classification performance. With human subjects, these measures could potentially include subjective report, reaction time, or galvanic skin response. Another possibility would be to use fMRI to measure brain metabolic activity that correlates with entropy (or perhaps variance). Preuschoff, Bossaerts, and Quartz (submitted) found brain activity that correlated with uncertainty of monetary reward probability. An fMRI study could address whether uncertainty in a categorization task also correlates with brain activity. Bob Rehder (personal communication, 2005) has noted his impression that subjects initially focus only on Type I concepts, and broaden their search if it later becomes clear that Type I concepts cannot account for the data.

3 Eye movements in concept formation

One of our main reasons for developing a probabilistic model of learning on Shepard, Hovland and Jenkins's (1961) task was to provide a basis for evaluating whether optimal experimental design theories of data acquisition could offer insight into human subjects' eye movements, in Rehder and Hoffman's (2003, 2005) eye movement version of the task. In this section, we describe Rehder and Hoffman's eye movement version of the task; their main results; our eye movement model of the task; and how results from our eye movement model may offer insight into Rehder and Hoffman's data. Formally, our eye movement model is a utility function that quantifies the usefulness of each possible set of eye movements, at the beginning of each trial, throughout the simulated concept learner's learning. Note that although our concept-learning model includes all types of concepts that were considered by Shepard et al., we do not discuss model eye movement results for concept types III and V, because Rehder and Hoffman did not obtain eye movement data from human subjects on concepts of those types.

3.1 Rehder and Hoffman's (2003, 2005) task and eye movement data

Numerous models of category learning describe how selective attention to different stimulus dimensions evolves, yet little direct evidence is available. Rehder and Hoffman (2005) devised a novel experimental design in which eye movements provided direct evidence for how selective attention was deployed at each moment in time. This was accomplished by representing the value of each stimulus dimension with abstract characters (such as \$ or ¢ at the bottom right of Figure 5) at a particular vertex of a large triangle on a computer screen, and recording eye movements throughout learning. A large black circle might be represented by having an "x" in the bottom left, an "!" in the top middle, and a "\$" in the bottom right. A small black circle might be represented by having an "o" in the bottom left, and the other characters as before.




Figure 5. Example stimulus from Rehder & Hoffman's (2005) experiment. Reprinted from Figure 2 in Rehder, B., & Hoffman, A. B. (2005). Eyetracking and selective attention in category learning. *Cognitive Psychology*, 51, 1-41. © 2005 Elsevier. Permission to reprint is pending.

Among 72 subjects, a variety of patterns were observed. For the vast majority of participants, the following findings held:

- (1) Early in learning, all three stimulus dimensions were viewed (fixated).

(2) Improvement in categorization performance occurred before reduction in number of stimulus dimensions viewed to just those necessary for correct classification (“efficient” eye movements). For instance, the modal pattern for Type I learners was for error to drop from 50% (chance) to almost 0% after about 12 trials (1.5 blocks), but for the number of dimensions fixated to remain high until about 16 trials (2.0 blocks). (“Type I learner” means a subject who was assigned to learn a Type I concept.) For Type II learners, errors decreased more gradually, beginning approximately on trial 40 (5 blocks), and decreasing gradually through approximately trial 72 (9 blocks).

(3) Eventually, *efficient* eye movements (to only the necessary stimulus dimensions, given certain knowledge of the true concept) were used. For Type I concepts, efficient eye movements are to a single dimension, such as the size query for the concept “small”. Rehder & Hoffman’s Type I subjects in general fixated only the single relevant dimension, late in learning. For Type II concepts, efficient eye movements are to the relevant two dimensions, such as the size-shape query for the concept “large square or small circle.” Type IV concepts, in the general case, require all three dimensions. In some cases, however, depending on what is learned from the first dimension, Type IV concept objects can be classified with two dimensions. Type IV subjects tended to fixate all three dimensions, although the average number of dimensions fixated by Type IV subjects was slightly less than the number of dimensions fixated by Type VI subjects (Figure 6). The finding that Type IV subjects tended to fixate all three dimensions, even late in learning, could suggest that subjects found it more expeditious to plan a sequence of eye movements in advance (e.g. a remembered scanpath or visual routine), rather than to fixate a single dimension and then consider what dimension to fixate next. Finally, Type VI concepts always require all three stimulus dimensions; the vast majority of subjects always considered all three dimensions.

We will focus on these empirical findings when evaluating our sampling model.

3.2 Quantifying eye movements’ usefulness

We model the history of learning with the assumption that the learner has seen all stimulus dimensions. This assumption corresponds well to the classic task, but is a simplification in modeling Rehder and Hoffman’s (2005) eye movement version. At each point in learning, our concept-learning model is used to describe the model learner’s beliefs. Technically, the eye movement model is a subjective utility function for evidence acquisition, in the sense of Savage (1954). When we refer to the *usefulness* of each possible eye movement, we mean the subjective utility that the model learner expects to obtain, on average, as a result of making that eye movement. In the equations below we refer to queries, rather than eye movements per se, to reflect that a task could be designed so that information is obtained by other means, for instance by asking a question, clicking a mouse pointer on a computer screen, etc.⁵ At each point in learning, the learner picks between 8 possible queries: all

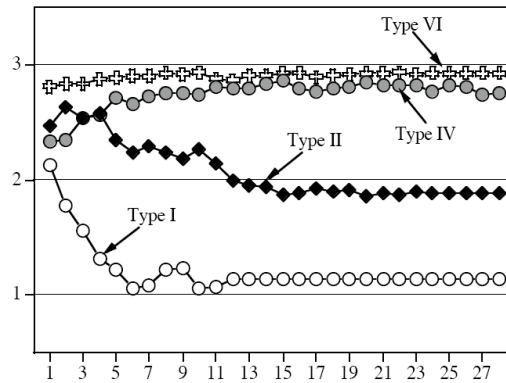


Figure 6. Number of stimulus dimensions fixated (y axis), during different blocks of learning (x axis). Reprinted from Figure 3 in Rehder, B., & Hoffman, A. B. (2005). Eyetracking and selective attention in category learning. *Cognitive Psychology*, 51, 1-41. © 2005 Elsevier. Permission to reprint is pending.

⁵ Some eye movement models, for instance the E-Z Reader models of reading (Reichle, Pollatsek, Fisher, & Rayner, 1998; Reichle, Rayner, Pollatsek, 2003) explicitly account for oculomotor processes, sub-second timing issues, saccade (eye movement) length distributions, and specific properties of visual acuity at different eccentricities (angular distances from center of gaze). Our model is relatively abstract. We disregard issues of peripheral acuity, because Rehder and Hoffman’s (2005) stimuli were designed to render peripheral information uninformative, except with respect to

dimensions, size-shape, size-color, color-shape, size, color, shape, or null. The task for modeling is to see whether a principled utility function, inspired by ideas from theories of optimal experimental design, can offer insight into subjects' eye movements in Rehder & Hoffman's (2005) data. We propose the following function to quantify a particular query's usefulness:

$$u_t(Q) = pg_t(Q) + I_t(Q; C) - j(Q).$$

The subscript t is shorthand to condition on the complete history of learning up to the present trial. Capital letters denote random variables, and lowercase letters denote specific values taken by random variables. For instance, Q denotes a particular query, and each of several q denote particular possible outcomes of that query. Each term in the usefulness function corresponds to a central idea, as introduced below:

1. $pg_t(Q)$. The learner wishes to correctly classify the randomly selected object that is presented in a particular trial (exploitation).
2. $I_t(Q; C)$. The learner wishes to learn the true concept (exploration).
3. $-j(Q)$. The learner wishes to avoid making extraneous eye movements.

The next sections explain the terms of the usefulness function in more detail, and describe the rationale behind each term.

3.2.1 Classifying a random object: $pg_t(Q)$

The learner's immediate task is correct classification on a particular trial. The eye movement model's usefulness function captures this idea (Baron 1981, as cited in Baron, 1985) with its $pg_t(Q)$ term. This term quantifies the amount that a query is expected to improve the probability of correct response on a particular trial, the query's *probability gain*. Because objects are drawn randomly with equal probability, and 4 of the 8 objects are consistent with each concept, chance accuracy is 0.5. Formally:

$$pg_t(Q) = \sum_q P_t(\text{correct} | q) P(q) - 0.5$$

In the above equation, Q is shorthand for a particular query that the learner might make, for example Q_{size} , in which case:

$$pg_t(Q_{size}) = \sum_{q \in \{\text{large}, \text{small}\}} P_t(\text{correct} | q) P(q) - 0.5$$

The probability of correct response, given a particular outcome of a query, is:

$$P(\text{correct} | q) = \max[P(\alpha | q), 1 - P(\alpha | q)],$$

where $P(\alpha | q)$ gives the probability that the correct response on the trial is "yes", given q . The $P(\text{correct} | q)$ term is derived from an optimal (not probability matching) response strategy, in which the learner always responds "yes" if $P(\alpha | q) > 0.5$, and "no" if $P(\alpha | q) \leq 0.5$.

Suppose that at time t , late in learning, the true concept were known to be "large," and that the learner were evaluating the probability gain of the size query (Q_{size}). If the observation (q) were that the size dimension takes the value "small," then $P(\alpha | q) = 0 < 0.5$, so the learner would respond "no," which results in 100% probability of being correct. If the

stimulus position. We disregard the visual system's preference for relatively short eye movements, because the targets—at the vertices of a triangle—are roughly equidistant from the current point of gaze. Timing issues would be interesting to consider in future work, such as whether longer fixations (periods of focused gaze) tend to be more informative features.

observation q were “large,” $P(\alpha | q) = 1 > 0.5$, and the learner would respond “yes,” again with 100% probability of being correct. Before making any eye movements, the learner averages over the two possible query results, each with probability 0.50 of occurrence, and obtains $pg_t(Q_{size}) = 0.50$.

3.2.2 Learning the true concept: $I_t(Q; C)$

The second term in the usefulness function, $I_t(Q; C)$, quantifies a query Q 's value for reducing uncertainty about the true concept C . We use the mutual information between the query Q and the concept C , given learning to date t , to measure this:

$$I_t(Q; C) = H_t(C) - H_t(C|Q)$$

$$= \sum_c P_t(c) \log_2 \frac{1}{P_t(c)} - \sum_q P_t(q) \sum_c P_t(c|q) \log_2 \frac{1}{P_t(c|q)}$$

This term is an exploration term to quantify the anticipated future usefulness of what is learned in the present trial. It also addresses the subject's presumed desire to identify the true concept. Technically, mutual information is the expected reduction in entropy in C from making query Q , or the initial entropy $H_t(C)$ less the conditional entropy $H_t(C|Q)$. Note that $H_t(C|Q)$ is the expected entropy given Q , not the actual entropy given a particular q .

Sometimes called ‘expected information gain,’ the use of mutual information to quantify experiments’ (or queries’) usefulness was proposed by Lindley (1956). Oaksford and Chater (1994, 2003) used information gain in a model of Wason’s (1966, 1968) card selection task. Oaksford and Chater (1996) noted that mutual information is equivalent to the directional Kullback-Liebler distance (Kullback & Liebler, 1951) from prior to posterior distribution, as well. Lee and Yu (1999) suggested using mutual information to predict eye movements. Denzler and Brown (2002) used information gain to control cameras in a robotic vision task. Renninger, Coughlan, Verghese, and Malik (2005) did so in a shape learning task model, as did Movellan (2005) in a social contingency detection model. Nelson (in press) reviews alternate ways an exploration term could be formulated.

3.2.3 Avoiding extraneous eye movements: $j(Q)$

Finally, we include a cost term, $j(Q)$, to represent the subjective cost of fixating one or more stimulus dimensions. This cost does not depend on the history of learning t . We set $j(Q)$ arbitrarily, to 0.04 per dimension fixated. Hence, Q_{all} , the size-shape-color query, has cost 0.12, and Q_{size} has cost 0.04. We have simulated the model with various settings of this parameter. Results do not greatly depend on the specific value chosen, so long as it is small and nonzero.

3.3 Results from the eye movement model

What eye movements does the model predict, at particular stages of learning? In each trial, there are 8 possible queries: size-shape-color, size-shape, size-color, color-shape, size, color, shape, and null. Before learning, the query to all stimulus dimensions ($u=0.880$) is more useful than any query of two ($u=0.585$) or one ($u=0.292$) stimulus dimensions, consistent with Rehder and Hoffman’s (2005) empirical finding that early in learning, learners tend to fixate all stimulus dimensions. This is due to the queries’ different $I_t(C; Q)$, their expected information gain with respect to the learner’s beliefs about the true concept. Because no learning has taken place, this is true irrespective of the true concept. Before learning, $pg_t(Q)$ is zero for every query. Figure 7 illustrates how $I_t(C; Q)$, the open circles, and $pg_t(Q)$, the closed circles, change during learning of the concept “small”, comparing the query to all stimulus dimensions (left panel), and the size query (right panel). Comparing the open circles on the two graphs, the $I_t(C; Q)$ term always favors the query to all stimulus dimensions, relative to the query to the size dimension alone. During learning, however, the $I_t(C; Q)$ terms tend toward zero, such that the difference between $I_t(C; Q_{all})$ and $I_t(C; Q_{size})$ is

miniscule. Because of this, eventually the size query, which has the lower cost, becomes more valuable than the query to all dimensions (Figure 9). In other words, the model initially focuses on learning the correct concept, but as the correct Type I concept quickly becomes known, the model shifts its focus to producing a correct response.

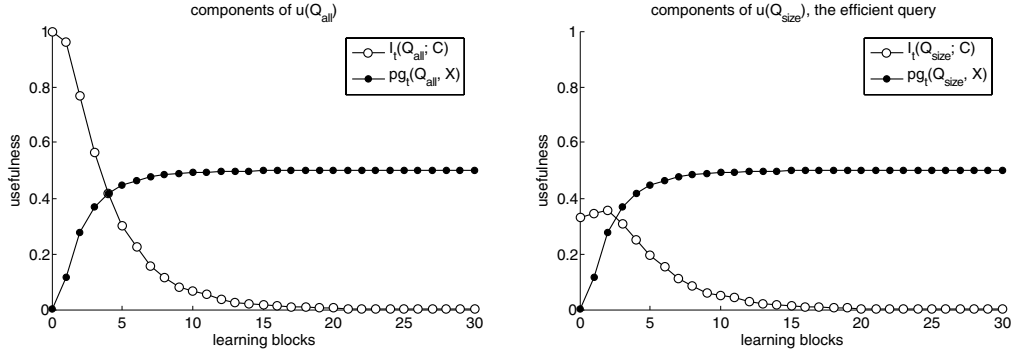


Figure 7. Components of the usefulness movement function for the concept “small.” At left, Q_{all} , the query to all stimulus dimensions; at right, Q_{size} , the efficient size query.

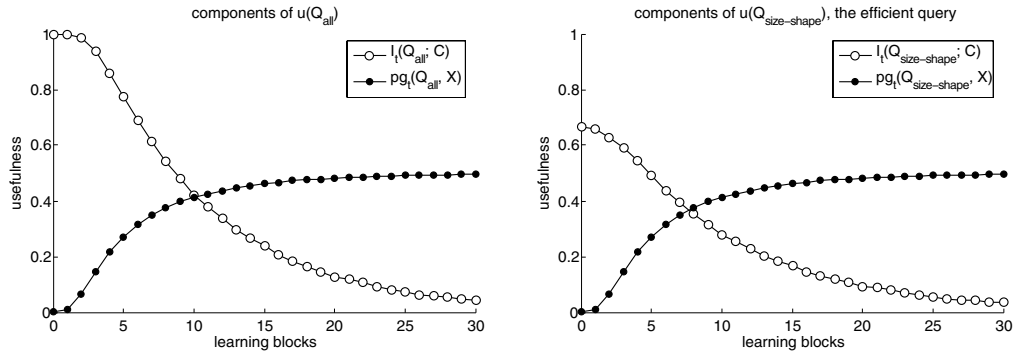


Figure 8. Components of usefulness function for the Type II concept “large square or small circle.” Left panel, the query to all stimulus dimensions; Right, the efficient query, to size and shape.

A similar pattern holds for Type II concepts (Figure 8). We discuss results for the Type II concept “large square or small circle.” For this concept, the efficient query ($Q_{efficient}$) is to the size and shape dimensions. Type II concepts have a smaller initial difference in $I_t(C; Q_{all})$ and $I_t(C; Q_{efficient})$ than do Type I concepts. This is because a query to two stimulus dimensions ($Q_{efficient}$ for Type II concepts) is expected to provide more information than a query to a single stimulus dimension ($Q_{efficient}$ for Type I concepts). The slower learning of Type II concepts, relative to Type I concepts, results in slower decrease of the $I_t(C; Q)$ terms and increase of the $pg_t(Q)$ terms.

For Type IV and VI concepts (not shown), the model always prefers the query to all stimulus dimensions.

3.4 Putting it all together

We now have results from both the concept-learning model and the eye movement model in hand, so we can examine their behavior together. This is the critical test to evaluate whether our model offers insight into Rehder and Hoffman’s (2005) data. Figure 9, for the Type I

concept, “small,” shows that error rate (dashed line) decreases well before the efficient query (x marker) becomes more valuable than the query to all dimensions (+ marker).




Figure 9. Error rate and selected queries’ usefulness, Type I concept, “small..”

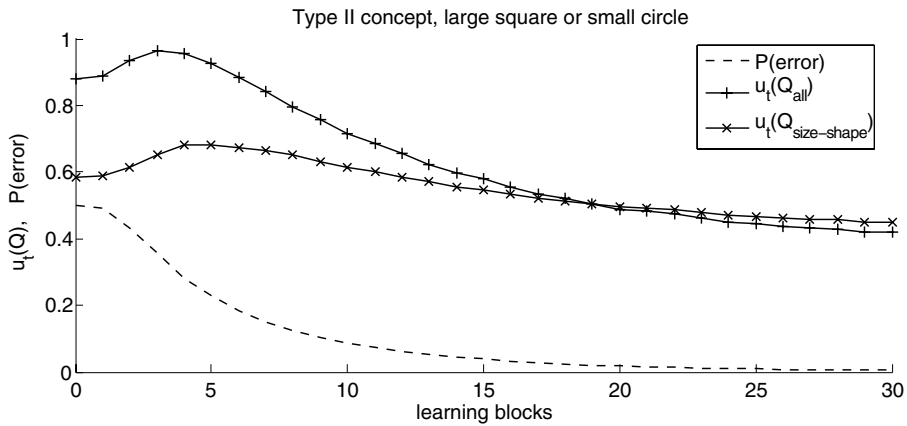


Figure 10. Error rate and selected queries’ usefulness, Type II concept, “large square or small circle.”

Figure 10 shows that the model behaves in a similar manner on Type II concepts. Importantly, however, because learning develops more slowly in the case of Type II concepts, the error rates decrease more slowly, and the point when the efficient query becomes more valuable occurs much later, than in the case of Type I concepts. For Type IV and VI concepts, the model rates the query to all dimensions as having the highest usefulness throughout learning.

Together, these results show that the model produces the main qualitative features that Rehder & Hoffman (2005) noted in their data. In each case, the model offers a principled explanation of the result:

1. Initially, subjects consider all stimulus dimensions, because Q_{all} provides more information than any other query about the true concept, as quantified by the $I_i(Q, C)$ term. The greater usefulness with respect to the $I_i(Q, C)$ term more than overcomes the higher cost of Q_{all} , versus other queries.
2. Error rates decrease during the course of learning, even as subjects continue to fixate all stimulus dimensions, because correct classification performance results from beliefs about several possible concepts, and not just the concept that later proves correct.

- Subjects switch from Q_{all} to $Q_{efficient}$ after achieving high classification performance, because even though high classification performance is being achieved already, it takes additional learning to rule-out plausible alternate hypotheses about the true concept.

3.5 Questions about the usefulness function

A number of points of confusion about our formulation of the usefulness function may arise. Are all the terms necessary? Could some terms be formulated differently? This section discusses those issues.

Isn't there a correct way to quantify queries' usefulness? For the sake of argument, suppose that the task included an explicitly specified set of external rewards. For instance, suppose that the learner is paid \$1 for each trial in which they guess correctly, up to a total of 300 trials. In this case, there is a unique, correct calculation of each query's usefulness, at each point in learning, and dynamic programming could in principle be used to identify what utility function applies at each point in learning, to maximize net total payoff (Javier Movellan, personal communication, 2005). Modeling the subjective utility (the perceived value to the subjects) of money, however, would require separate specification of subjects' utility function for money, as well as of the temporal discount factor that subjects have for future reward. Furthermore, external rewards are not explicitly specified on Rehder and Hoffman's (2005) task, so different people could have different ideas about what function implicitly applies. Similarly, there is no guarantee that introduction of external rewards would eliminate intrinsic goals that the subject might have, such as identifying the true concept. Accordingly, we view the current modeling objective as ascertaining whether ideas from optimal experimental design can help describe subjects' goals on Rehder and Hoffman's task, rather than as ascertaining whether subjects have a theoretically correct model of the value of information on the task.

What if the $pg_i(Q)$ term were eliminated? Intuitively, it may seem that the learner's goal is to learn the true concept, and that this goal is served by the $I_i(C, Q)$ term. Early in learning, if the $pg_i(Q)$ term were eliminated, the model concept learner would behave similarly, because the main driver for the query to all stimulus dimensions is its ability to provide information about the unknown true concept. Late in learning, however, once the true concept is known with high certainty, the model concept learner, rather than relying on the efficient query relative to the true concept, would stop making eye movements altogether (switching to the null query), because the cost of any query would more than outweigh the usefulness of its miniscule expected reduction in uncertainty about the true concept. The paradoxical result would be that as the learner comes to know the true concept with high probability, as attested by the probabilities of each concept in the model, performance on the categorization task would decrease to chance.

What if the $I_i(C, Q)$ term were eliminated? This would cause the learner to have to rely on the $pg_i(Q)$ term to identify useful eye movements. Unfortunately, in the first trial of learning it is impossible to classify the object with greater than chance (50%) accuracy, irrespective of the number of stimulus dimensions fixated. Thus, because of each query's cost of execution, every query (except the null query) would have negative usefulness, and the learner would not fixate any of the stimulus dimensions. In the second learning trial, the learner's beliefs about the concepts would be the same as the learner's initial beliefs, because no learning took place on the previous trial. Hence, the same situation would arise. The problem would iterate indefinitely, and at no point in time would the learner fixate any stimulus dimensions, or achieve greater-than-chance performance on the task.

What if $j(Q)$ term were eliminated, by setting the cost of each query to zero? Consider what would happen when, late in learning, uncertainty is small and the true concept is known with nearly 100% certainty. Even at this point, the learner would fail to switch from fixating all stimulus dimensions to the efficient query, because Q_{all} would be marginally more valuable (Figure 11). This behavior would make sense if eye movements were free. However, every eye movement (saccade) has a biomechanical cost to execute. There is also an opportunity cost in the other eye movements that must be foregone, and the inability to perceive any visual stimuli due to retinal blurring, during each saccadic eye movement's execution.




Figure 11. If cost j is zero, Q_{all} is the most useful query, throughout learning.

Could the $I_t(Q, C)$ term be formalized differently? Several usefulness functions for quantifying the value of information have been proposed. Nelson (in press) provided experimental evidence that some usefulness functions from the literature (Bayesian diagnosticity and log diagnosticity) are not plausible models of human intuitions, and behave poorly in several situations. However, based on a review of the literature, Nelson concluded that several other usefulness functions are roughly equally reasonable as descriptive models of empirical data from human subjects. These functions include mutual information, which we use here, as well as probability gain and impact. If probability gain were to be used as the basis of an exploitation term, it would quantify the expected improvement of correctly identifying the true *concept*, rather than the improvement in probability of correct response (yes or no) on the current trial, which the model's $pg_t(Q)$ term quantifies. Impact, the expected absolute change in beliefs (Wells and Lindsay, 1980; Klayman & Ha, 1987; and Nickerson, 1996) is another candidate. Nelson reported simulations suggesting that these three usefulness functions in general behave similarly, in a two-category, binary-feature scenario.⁶ Other plausible usefulness functions could also be proposed. We are not deeply invested in a particular proposal.

Should a term for intrinsic visual salience (Yamada & Cottrell, 1995; Itti, Koch, & Niebur, 1998; Itti & Baldi, 2005) be included? The justification for a salience term would be that certain kinds of visual stimuli (such as high color contrast areas) might be intrinsically interesting to the visual system, irrespective of one's current task. In Rehder and Hoffman's (2005) task, the values of the stimulus dimensions of the objects were the only illuminated parts of the screen during each trial. A salience-based eye movement model alone could not model Rehder and Hoffman's eye movement data, namely the different patterns of eye movements observed for the different stimulus types, as the visual stimuli were the same for all types of concepts. The other terms are each necessary to model Rehder and Hoffman's eye movement data. For the sake of parsimony, we ignore visual salience in the present model, although it could potentially be important on other tasks.

What likelihood does the model learner think they will use when evaluating each possible query's usefulness? The current model makes the simplifying assumption that when evaluating possible queries' (eye movements') anticipated usefulness, $\mu=1$. The implicit psychological claim here is that when considering what eye movements to make, the learner believes he or she will have complete confidence in the results. In future work we intend to

⁶ Jiri "George" Najemnik (personal communication, 2005) noted that although Najemnik and Geisler (2005) used probability correct, which is equivalent to probability gain, in an ideal observer analysis of a visual search task, they obtain similar results using mutual information, and that the two functions are too similar for determination of which of them best models human data.

explore the implications of using the same likelihood function as in the concept learning model, where $\mu < 1$, when calculating possible queries' usefulness.

4 Summary, discussion, and future work

Shepard, Hovland and Jenkins's (1961) task has received a great deal of study. Like much work in cognitive psychology, process models have largely inspired theories of the task. Many of these models have made claims about the learner's selective attention to different stimulus dimensions throughout learning, in the absence of selective attention data. Rehder and Hoffman's (2003, 2005) insight was to devise an eye movement-based version of the task, so that the viewer's eye movements could provide direct evidence for how selective attention is deployed throughout learning.

Rehder and Hoffman noted that their eye movement data did not appear to fit existing models well. RULEX, a rule-based category-learning model (Nosofsky, Palmeri, & McKinley, 1994) would predict eye movements to one dimension at a time, early in learning. Subjects, however, tended to fixate all stimulus dimensions early in learning. ALCOVE, a connectionist exemplar-based model (Kruschke, 1992) would predict a gradual transition between eye movements to all stimulus dimensions, and eye movements to the relevant dimensions. However, subjects showed a relatively sudden switch to efficient eye movements, after error rates were already low. Rehder and Hoffman (2005) hypothesized that multiple learning modules, some rule-based and some similarity-based, control learning and selective attention on their task. Rehder and Hoffman (2005, pp. 38-39) suggested that perhaps the late switch to efficient eye movements happens because subjects strategically "abandon alternative learning strategies after one module has solved the learning problem," or because of "cognitive limits that influence how objects get prioritized in a visual field, or both." Our results show that recourse to such a complicated set of explanations of Rehder and Hoffman's data may not be necessary, even though some aspects seem consistent with ALCOVE, and other aspects with RULEX.

There were two main goals underlying the research discussed in this article. The first was to evaluate whether a probabilistic model that describes beliefs and their development as optimal Bayesian inference, could help explain the process of learning on Shepard et al.'s task. It appears that the Bayesian concept learning model at least qualitatively matches key properties in data from Shepard et al.'s task, such as rapid reduction in error for Type I concepts, similar error rates for Type II and IV concepts through much of learning, and the expected ordering of concept types in number of trials to criterion performance. The second goal was to evaluate whether, given that Bayesian model, eye movements (selective attention) on Rehder and Hoffman's (2003, 2005) task could be modeled as a process of optimal experimental design. We view our eye movement modeling results as a proof of concept that a usefulness (utility) function that is inspired by principled statistical ideas can reproduce important properties of human behavior. A model with an unchanging set of *goals* (the usefulness function), inspired by ideas from the theory of optimal experimental design, and with *beliefs* that develop in a principled way, modeled with Bayesian statistics, produces the key findings that Rehder & Hoffman (2005) reported: early fixations to all stimulus dimensions, the gradual reduction in classification error, and the late shift to efficient eye movements.

This article reports work in progress. In the future we hope to address several issues:

1. What is the source of prior beliefs? Here we primarily focused on identifying what the learner's beliefs are, at each point in learning, and on how those beliefs change during learning. We therefore used error counts from Rehder and Hoffman's (2005) data as a basis for assigning prior probabilities. Statistical measures, such as Boolean complexity (Feldman, 2000), could perhaps also be used in assignment of prior probabilities. It would be remarkable if subjects' intuitions about different concepts' relative plausibility could be explained in such a principled manner.
2. What range of eye movement (utility) functions can offer insight into Rehder and Hoffman's (2005) data? We chose a usefulness function that seemed plausible, based on

the task and on research using optimal experimental design concepts to understand human intuitions (reviewed by Nelson, in press). A variety of other formulations of the usefulness function could be considered. For instance, the learner's goal of identifying the correct concept, which we quantified with mutual information (information gain), $I_t(C, Q)$, might alternately be quantified with absolute change in beliefs (Wells and Lindsay, 1980; Klayman & Ha, 1987; Nickerson, 1996), or with probability gain (Baron, 1985). It would also be helpful to more thoroughly consider the explore/exploit tradeoff. Exploration, in our model, is quantified by $I_t(C, Q)$, and exploitation is quantified by $pg_t(Q)$. In the present model we simply added those terms, and did not optimize their relative weight. However, subjects may place more weight on one of these terms. In general, we sought to be conservative in not optimizing every aspect of model construction, to avoid overfitting. In the future we hope that minimum description length (Pitt, Myung, & Zhang, 2002) or other techniques may enable us to consider a wider range of models, while insuring against overfitting empirical data.

3. What more can be learned from a more fine-grained analysis of the data? Rehder & Hoffman (2005) reported that their subjects could be clustered according to their eye movements and error data. Optimizing model parameters to subgroups of subjects (see Palmeri & Nosofsky, 1995), or even to individual subjects, could potentially inform discussion of the subjects and the model alike.

The models introduced in this article contribute to a growing body of research in rational analyses of cognitive tasks (Anderson, 1990; 1991, pp. 409-410). This line of research is analogous to ideal observer analyses of perceptual tasks (Knill & Richards, 1996; Movellan & McClelland, 2001). It was presaged by Brunswik's (1952) molar analysis and Marr's (1982) computational analysis; others have recently discussed it as well (Oaksford & Chater, 1999; Movellan & Nelson, 2001). The focus is on the nature of problems that cognition solves, on what would constitute optimal solutions to those problems, and on whether those optimal solutions might approximate human cognition. Our results, of course, do not demonstrate what architecture the brain uses to learn categories, to categorize objects, or to deploy selective attention via eye movements. However, our results demonstrate that a single model, with a small number of parameters, based on Bayesian statistics and theory of optimal experimental design, can (1) provide a parsimonious account of classic concept learning results, and (2) help explain the central findings that Rehder and Hoffman noted in their compelling eye movement data.

In life, virtually every interesting learning or inference scenario involves some active contribution of the learner in deciding where to look, what questions to ask or experiments to conduct, or what data to attend to. Optimal experimental design-inspired models of evidence acquisition, however, have seldom considered tasks where beliefs change so drastically, as information is sought out, decisions are made, and feedback is given, over dozens or hundreds of trials. The present model suggests that unified, rational accounts of learning and information acquisition may apply to a wide range of behavior. Our model contributes to a growing body of work using information- and decision-theoretic concepts to model perceptual tasks (Lee & Yu, 1999; Sprague & Ballard, 2003; Trommershäuser, Maloney & Landy, 2003; Najemnik & Geisler, 2005; Movellan, 2005; Renninger, et al., 2005). We hope it may also spur work to explore the relationship between active information acquisition in perceptual, such as eye movement, and other learning and inference tasks.

Author note

Bob Rehder and Aaron Hoffman kindly corresponded about their experiment and findings, and provided their data to us. Bob Rehder, Javier Movellan, Michael Lee, Tim Marks, Flavia Filimon, and two anonymous reviewers provided feedback on a draft of this manuscript, and numerous ideas on the research and its explication. Nathaniel Smith, Josh Tenenbaum, Eric Wiewiora, and Danke Xie provided helpful suggestions on this work. JDN was funded by NIMH grant 5T32MH020002-05 to T. Sejnowski and by NSF grant DGE 0333451 to GWC. GWC is supported by NIH grant MH57075. An earlier version of this

work was presented at ICDL 2004 as Nelson, Cottrell, and Movellan (2004). New developments in this research will be posted at <http://www.jonathandnelson.com/>

References

- Anderson, J. R. (1990). *The Adaptive Character of Thought*. Hillsdale NJ: Erlbaum.
- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98, 409-429.
- Baron, J. (1985). *Rationality and Intelligence*. Cambridge: Cambridge University Press.
- Bayes, T. (1763). An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London*, 53, 370-418.
- Brunswik, E. (1952). *The conceptual framework of psychology*. Chicago: University of Chicago Press.
- Chater N, Oaksford M (1999). Ten years of the rational analysis of cognition. *Trends in Cognitive Sciences*, 3(2), 57-65.
- Cover, T. M., & Thomas, J. A. (1991). *Elements of Information Theory*. New York: Wiley
- Denzler, J., & Brown, C. M. (2002). Information theoretic sensor data selection for active object recognition and state estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24, 145- 157.
- Edwards, W. (1968). Conservatism in human information processing. In B. Kleinmuntz (Ed.), *Formal Representation of Human Judgment*, (pp. 17-52). New York: John Wiley
- Feldman, J. (2000). Minimization of Boolean complexity in human concept learning. *Nature*, 407, 630-633.
- Feldman, J. (2003). The simplicity principle in human concept learning. *Current Directions in Psychological Science*, 6, 227-232.
- Itti, L. & Baldi, P. (2005). A principled approach to detecting surprising events in video. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*.
- Itti, L., Koch, C., & Niebur, E. (1998). A Model of Saliency-Based Visual Attention for Rapid Scene Analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11), 1254-1259.
- Johansen, M. K., & Palmeri, T. (2002). Are there representational shifts during category learning? *Cognitive Psychology*, 45, 482-553.
- Klayman, J. & Ha, Y.-W. (1987). Confirmation, disconfirmation, and information. *Psychological Review*, 94, 211-228.
- Knill, D. C., & Richards, W. (Eds.), (1996). *Perception as Bayesian inference*. Cambridge, U.K.: Cambridge University Press.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*. 99, 22-44.
- Kullback, S., & Liebler, R. A. (1951). Information and sufficiency. *Annals of Mathematical Statistics*, 22, 79-86.
- Lee, T. S., & Yu, S. X. (2000). An information-theoretic framework for understanding saccadic eye movements. In S. A. Solla, T. K. Leen, & K.-R. Müller (Eds.), *Advances in Neural Information Processing Systems*, 12, 834-840. Cambridge, MA: MIT Press.
- Lindley, D. V. (1956). On a measure of the information provided by an experiment. *Annals of Mathematical Statistics*, 27, 986-1005.
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: W. H. Freeman.
- Movellan, J. R. (2005). An infomax controller for real time detection of contingency. Proceedings of the International Conference on Development and Learning, Osaka, Japan, July, 2005.
- Movellan, J. R., & McClelland, J. L. (2001). The Morton-Massaro law of information integration: Implications for models of perception. *Psychological Review*, 108(1), 113-148, 2001.
- Movellan, J. R., & Nelson, J. D. (2001). Probabilistic functionalism: a unifying paradigm for the cognitive sciences. *Behavioral and Brain Sciences*, 24, 690-692.

- Najemnik, J., & Geisler, W. S. (2005, March 17). Optimal eye movement strategies in visual search. *Nature*, 434, 387–391.
- Nelson, J. D. (in press). Finding useful questions: on Bayesian diagnosticity, probability, impact, and information gain. *Psychological Review*. Preprint available at <http://www.jonathandnelson.com/>
- Nelson, J. D., Tenenbaum, J. B., & Movellan, J. R. (2001). Active inference in concept learning. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the 23rd Conference of the Cognitive Science Society*, 692-697. Mahwah, NJ: Erlbaum.
- Nelson, J. D.; Cottrell, G. W.; Movellan, J. R. (2004). Explaining eye movements during learning as an active sampling process. International Conference on Development and Learning, Oct, 2004.
- Nickerson, R. S. (1996). Hempel's paradox and Wason's selection task: Logical and psychological puzzles of confirmation. *Thinking and Reasoning*, 2, 1-32.
- Nosofsky, R. M., & Gluck, M. A. Adaptive networks, exemplars, and classification rule learning. Paper presented at the 30th annual meeting of the Psychonomic Society, Atlanta. [Cited in Anderson, 1991].
- Nosofsky, R. M., Gluck, M., Palmeri, T. J., McKinley, S. C., & Glauthier, P. (1994). Comparing models of rule-based classification learning: A replication and extension of Shepard, Hovland, and Jenkins (1961). *Memory and Cognition*, 22, 352-369.
- Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101, 53-79.
- Oaksford, M., & Chater, N. (1996). Rational explanation of the selection task. *Psychological Review*, 103, 381-391.
- Oaksford, M., & Chater, N. (1998). A revised rational analysis of the selection task: exceptions and sequential sampling. In M. Oaksford, & N. Chater (Eds.), *Rational Models of Cognition* (pp. 372-393). Oxford: Oxford University Press.
- Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review*, 109, 472-491.
- Preuschhoff, K.; Bossaerts, P.; Quartz, S. R. (submitted). Encoding expected reward and risk in economic decision making.
- Rehder, B., & Hoffman, A. B. (2003). Eyetracking and selective attention in category learning. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th Annual Conference of the Cognitive Science Society*. Boston, MA: Cognitive Science Society.
- Rehder, B., & Hoffman, A. B. (2005). Eyetracking and selective attention in category learning. *Cognitive Psychology*, 51, 1-41.
- Reichle, E. D., Pollatsek, A., Fisher, D. L., & Rayner, K. (1998). Toward a model of eye movement control in reading. *Psychological Review*, 105, 125-157.
- Reichle, E. D., Rayner, K., & Pollatsek, A. (2003). The E-Z reader model of eye-movement control in reading: comparisons to other models. *Behavioral and Brain Sciences*, 26, 445-526.
- Renninger, L. Walker, Coughlan, J., Verghese, P., & Malik, J. (2005). An information maximization model of eye movements. In L. K. Saul, Y. Weiss, & L. Bottou (Eds.), *Advances in Neural Information Processing Systems*, 17, (pp. 1121–1128). Cambridge, MA: MIT Press.
- Savage, L. J. (1954). *The Foundations of Statistics*. New York: Wiley.
- Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, 15, 233-250.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27, 379–423, 623–656
- Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs: General and Applied*, 75(13), 1-42.
- Sprague, N., & Ballard, D. (2003). Eye Movements for Reward Maximization. In S. Thrun, L. K. Saul, and B. Scholkopf (Eds.), *Advances in Neural Information Processing Systems*, 16. Cambridge, MA: MIT Press.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J. & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27, 453-489.

- Tenenbaum, J. B. (1999). *A Bayesian Framework for Concept Learning*. Ph.D. Thesis, MIT
- Tenenbaum, J. B. (2000). Rules and similarity in concept learning. In S. A. Solla, T. K. Leen, & K.-R. Müller (Eds.), *Advances in Neural Information Processing Systems, 12*, 59–65. Cambridge, MA: MIT Press.
- Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences, 24(4)*, 629-640.
- Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003). Statistical decision theory and the selection of rapid, goal-directed movements. *Journal of the Optical Society of America A, 20(7)*, 1419-1433.
- Yamada, K., & Cottrell, G. W. (1995). A model of scan paths applied to face recognition. In *Proceedings of the Seventeenth Annual Cognitive Science Conference*, Pittsburgh, PA, pp. 55-60. Mahwah, NJ: Erlbaum.
- Zhang, L., & Cottrell, G. W. (2005). A computational model which learns to selectively attend in category learning. Paper presented at the International Conference on Development and Learning, Osaka, Japan, July, 2005.