Information Selection in Noisy Environments with Large Action Spaces

Pedro Tsividis (tsividis@mit.edu), Samuel J. Gershman (sjgershm@mit.edu),

Joshua B. Tenenbaum (jbt@mit.edu), Laura Schulz (lshulz@mit.edu)

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

77 Massachusetts Ave., Cambridge, MA 02139 USA

Abstract

A critical aspect of human cognition is the ability to effectively query the environment for information. The 'real' world is large and noisy, and therefore designing effective queries involves prioritizing both scope – the range of hypotheses addressed by the query – and reliability – the likelihood of obtaining a correct answer. Here we designed a simple information-search game in which participants had to select an informative query from a large set of queries, trading off scope and reliability. We find that adults are effective information-searchers even in large, noisy environments, and that their information search is best explained by a model that balances scope and reliability by selecting queries proportionately to their expected information gain.

Keywords: exploration; information search; active learning; information gain.

Introduction

As scientists, we sometimes encounter (or conduct) experimental work that is stunning in its breadth but disappointing in its rigor, or work that is categorically decisive but disappointingly narrow in scope. As child or adult intuitive scientists searching for information, we often deal with these epistemic virtues – scope and rigor – as well; we can make general queries that drastically narrow the hypothesis space of answers or make narrower ones, and we can seek information from reliable sources that are more likely to give us correct answers, or from sources that are less so. Our success, whether as professional or intuitive scientists, hinges on our ability to balance these two dimensions in order to produce queries whose answers will be informative.

Early work on information search seemed to show that people fail to make rational decisions when it comes to information acquisition; in the Wason (1968) selection task, with a fairly small (12) action set, only 4% of subjects made the normative information-acquisition selection. However, as Oaksford & Chater (1994) pointed out, the selection made most often by participants in the original Wason task was normative when environmental statistics were taken into account. Specifically, participants' decisions were best explained as maximizing expected information gain in the service of helping them to decide between competing hypotheses.

More recent work on information search has shown that children have strong intuitions about questions' usefulness and search adaptively (Nelson et al., 2013) and that adults can value information over explicit reward when the two are put in opposition (Markant & Gureckis, 2012).

Our interest is in whether these trends persist when people are confronted with the large, noisy information spaces characteristic of the real world. As we explore the world and its affordances, we must select queries that have scope, in that they rule out large numbers of hypotheses at once. And, inasmuch as we can help it, we should minimize noise by querying reliable sources. In the real world, these are often in opposition. So, we ask: how do people trade off scope and reliability when exploring large, noisy information spaces? And when the potential questions are many, do they ask the right ones?

Information-search task

To begin to answer these questions, we designed a novel information-search task. Our goal was for it to be as simple as possible, while having as many features approaching natural exploration as possible. Thus we paired a very simple game - identifying a hidden number on a number line - with a relatively complex search procedure. In playing the game, participants make queries that vary in the abstract features of scope and reliability. Additionally, at each point in the game, participants are faced with a very large number of potential specific queries to choose from. Our task is similar to the Markant & Gureckis (2012) task in that it involves exploring a geometric space to test particular hypotheses, but we wanted to make explicit the abstract features of questions, rather than have these be implicit as a function of the hypotheses at hand.

Participants play by asking 'questions' about the hidden number's location, using 'scanners' that turn blue if the number is under the scanned region and red if the number is not. In each trial, a participant is given four scanners (Fig. 1). In some conditions, the scanners vary in size. Larger scanners can cover larger regions of the number line, ruling out (or in) a larger set of hypotheses than a smaller scanner. Thus, in the context of this study, the 'scope' is directly related to the length of the scanner. However, the scanners are not deterministic; they also vary in their reliability, which is the probability of providing an accurate signal about the presence of the hidden number (false positives and false negatives are equally likely).

To efficiently find the hidden number, participants



Figure 1: A screenshot of the task. The four green rectangles are the scanners. The player has scanned three times and received one positive and two negative answers, as shown by the blue and red scanner outlines on the number line.

have to select scanners that provide a good trade-off between length and reliability, and then place them in informative regions on the number line. The optimal trade-off between these factors (length, reliability, location) can be captured by expected information gain, as described in the next section. Actions that are optimal in the sense of expected information gain minimize uncertainty about the location of the hidden number.

Model

Let *I* denote the number line (in our case, integers ranging from 0 to 100), and let $h \in I$ denote a hypothesis about the hidden number. On each trial, participants choose an action *a* (placing a particular scanner over a portion of the number line) and observe a binary outcome *d* (1 if the number was detected by the scanner, 0 if the number was not detected). There are almost 500 possible actions, since each of the four scanners can be placed anywhere as long as some part overlaps with the number line. The posterior distribution over *h* is updated on each trial according to Bayes' rule:

$$P(h|d,a,D) \propto P(d|h,a)P(h|D), \tag{1}$$

where D denotes the history of actions and outcomes prior to the current trial.

Intuitively, participants should choose actions that maximally reduce their uncertainty about the location of the hidden number; this corresponds to taking actions that maximally reduce posterior uncertainty, which can be quantified by the entropy:

$$H[P(h|d,a,D)] = -\sum_{h} P(h|d,a,D) \log P(h|d,a,D)$$
(2)

Minimizing posterior entropy is equivalent to maximizing information gain (the reduction of entropy after taking action a and observing d):

$$IG(a,d) = H[P(h|D)] - H[P(h|d,a,D)].$$
 (3)

Because the outcome d is not available at the time of choosing a, the best that a participant can do is maximize

expected information gain:

$$\operatorname{EIG}(a) = \sum_{d} P(d|a, D) \operatorname{IG}(a, d), \tag{4}$$

where the posterior predictive distribution is given by $P(d|a,D) = \sum_{h} P(d|a,h)P(h|a,D).$

Let α denote a scanner's error probability and I_a denote the range covered by the scanner. The game generates signals according to:

True positive: $\Pr(d = 1 \mid \{I_a \cap I\} \neq \emptyset) = 1 - \alpha$. **True negative:** $\Pr(d = 0 \mid \{I_a \cap I\} = \emptyset) = 1 - \alpha$. **False positive:** $\Pr(d = 1 \mid \{I_a \cap I\} = \emptyset) = \alpha$. **False negative:** $\Pr(d = 0 \mid \{I_a \cap I\} \neq \emptyset) = \alpha$.

Roughly speaking, conditioned on a scanner choice, participants should bisect the interval that has the highest posterior probability of containing the scanner. This corresponds to the *split-half* heuristic discussed by Navarro & Perfors (2011) and Nelson et al. (2013). However, because in this game signals are stochastic, there may not be a single contiguous interval of highest posterior probability, and thus no reasonable interval on which to precisely perform the split-half heuristic. Nevertheless, actions can still be ranked according to EIG.¹

To help hone our own intuitions about the task and about participants' actions, we created a display (Fig. 2) that in this case shows a vignette of four sequential actions from one particular user's trial. The normative posterior distribution over the location of the hidden number (in grey at the bottom) is displayed for each trial, along with the scaled expected information gain² of each of the 500 available actions (the colored arcs; each dot represents the center of the range at which the scanner could be placed); these are the evaluations of the normative model. The same recommendations are shown as posterior-predictive entropies in the insert (the red dot in the insert shows the posterior-predictive entropy of the action actually chosen by the participant), along with the participant's action at that point (the flat blue or red bar at the bottom), and the result of that action (blue indicates 'yes' and red indicates 'no').

Experiment 1

Method

Participants. 26 participants completed the experiment for pay on Amazon Mechanical Turk.

Materials. We used a base set of scanner reliabilities in the set, $\{0.51, 0.62, 0.75, 0.87\}$, and of lengths in the set, $\{0.0625, 0.125, 0.25, 0.5\}$.³ From these we created scanners as

¹Actions are being ranked according to their one-step information gain; see the discussion section for a comment on this.

²Scaled expected information gain is obtained by dividing the expected information gain of an action by the highest expected information gain available at that decision point.

 $^{^{3}}$ The indicated length is as a proportion of the length of the number line.



Figure 2: Analysis of one participant's four sequential actions in the information-search game. All the scanners were of reliability=0.87 and varied only in length. A: The participant uses the length=.5 scanner and receives a response of 'yes'. The arcs correspond to the scaled information gain of each candidate action; note that the participant selects one of the highest-rated actions. This fact is also indicated in the insert, which shows the posteriorpredictive entropies of all candidate actions (lower is better); the red dot indicates the posterior-predictive entropy of the action chosen by the participant. B: The participant employs the 'split-half' heuristic to test half the remaining space. Note that all of the highest-ranked model recommendations involve testing exactly half of the highest-posterior-probability region. C: The participant continues to use 'split-half'. D: The participant re-tests the entire region of highest posterior probability. This is highly rated, but not as much as it would be had they tested half the remaining space.

follows for the following conditions:

Reliabilities: One scanner in each of the four reliabilities above; all are of length 0.25.

Lengths: One scanner in each of the lengths above; all are of reliability 0.87.

Crossed: We wanted to see whether people could choose well when reliability and length were put in opposition, so we crossed them to create the following [length, reliability] pairings: [0.5, 0.51], [0.25, 0.62], [0.125, 0.75], [0.0625, 0.87].

Each participant performed 6 trials: 4 'crossed' trials, 1 'reliabilities' trial, 1 'lengths' trial; the order of these was randomized.

Procedure. Each participant read a short sequence of interactive instructions. They were told that with each round of the game a number would be hidden at a random location on the number line, told that they could use scanners to find the number, introduced to the scanners, shown that scanners glowed blue to indicate that the number was in that region and red to indicate that it was not, and given experience with the fact that the scanners could produce both false positives and false negatives. Finally, participants were told that once they thought they knew where the hidden number was located, they should use a provided set of crosshairs to indicate its location.

Following the end of the 'instructions' session, participants played six rounds of the game. Each game (trial) presented them with four scanners whose lengths and reliabilities were determined by the condition of the trial. As they played each round, the history of their queries was visible to them as the red or blue glow produced by each scanner remained in its place after the scanner was removed; this was to ensure that memory load did not differentially affect participants' decisions. After a scanner was used, it returned to its original position. Using multiple scanners simultaneously was not allowed by the interface. Participants were allowed to make as many queries as they wanted as they went through a trial; trials ended only when participants indicated their believed location of the hidden number by using the provided crosshairs.

Our main interest was to see how participants would confront a trade-off between scope and reliability. Before examining this, however, we tested each separately as proof of concept.

Results

Lengths condition First we tested whether participants were sensitive to the fact that different queries had more or less coverage of the hypothesis space. We gave them four scanners which varied only in their lengths; the scanners all had reliabilities of 0.87, and lengths of [0.5, 0.25, 0.125, 0.0625]. For each decision point reached by a participant (that is, before each scanning action), we calculated the expected information gain afforded by centering each available scanner at the center of the region actually queried by the participant. We then ranked those scanners according to this expected information gain. Participants picked the best scanner most often - 45% of the time - showing that they were sensitive to the imperative to cover as much ground of the hypothesis space as possible. A repeatedmeasures one-way ANOVA showed that choice proportions differed significantly across scanners [F(3,25) =16.55, p < 0.0001]. A post-hoc t-test showed that the difference between the first and second scanner ranks was significant [t(25) = 5.80, p < 0.0001].

Reliability condition Next we tested whether participants were sensitive to the imperative to reduce noise; that is – did participants make reliable queries when they were given a chance to do so?

Participants were given four scanners which varied only in reliability; the provided reliabilities were [0.87, 0.75, 0.62, 0.51], and all scanners were of length 0.25.

Across all subjects and all trials of this condition, participants were strikingly sensitive to reliability, as they selected the best scanner 89% of the time. A repeatedmeasures one-way ANOVA showed that choice proportions differed significantly across scanners [F(3,25) =52.88, p < 0.0001]. A post-hoc t-test showed that the difference between the first and second scanner ranks was significant [t(25) = 8.88, p < 0.0001].

Crossed condition When confronted with a choice of questions that might either greatly reduce the size of the hypothesis space or provide reliable answers, how did participants choose?

We constructed a condition in which these two dimensions were in direct opposition: we offered an array of scanners such that the more coverage of the hypothesis space a scanner provided, the less reliable it would be. We used four scanners whose [lengths, reliabilities] were [0.0625, 0.87], [0.125, 0.75], [0.25, 0.62], [0.5, 0.51]. Participants were clearly sensitive to length/reliability tradeoffs, as they selected the best scanner 57% of the time. A repeated-measures one-way ANOVA showed that choice proportions differed significantly across scanners [F(3,25) = 56.91, p < 0.0001]. A post-hoc t-test showed that the difference between the first and second scanner ranks was significant [t(25) = 5.82, p < 0.0001].

While we found these results encouraging, we were concerned that the particular length/reliability trade-offs we provided were such that one dimension might have contributed significantly more to expected information gain than the other. If this were the case, participants might have made what looked like normative decisions driven by both length and reliability, even though in reality they were driven only by length (or by reliability).

To examine this possibility, we ran a second experiment in which we provided participants with a large array of scanners of different [length, reliability] pairings. Length and reliability would trade off in terms of their relative contributions to a scanner's expected information gain, and therefore good performance in this condition would indicate an actual sensitivity to both dimensions.

Experiment 2

Method

Participants. 26 participants completed the experiment for pay on Amazon Mechanical Turk.

Materials. The same base set of lengths and reliabilities described in Experiment 1 was used to construct a new 'mixed' condition: We selected 8 of the 16 possible pairings of the 4 lengths and 4 reliabilities so as to cover a reasonable range of possibilities for available scanners. Specifically, participants were provided with scanners whose lengths and reliabilities were semi-randomly paired using the original [0.5, 0.25, 0.125, 0.0625] lengths and [0.87, 0.75, 0.62, 0.51] reliabilities. They then played the game in the same way as in Experiment 1.

We also ran a 'deterministic' condition in which scanners varied in length but were all of reliability = 1, randomly interleaved with the 'mixed' conditions. The results were similar to those reported in Experiments 1 and 2, and are left out for the sake of brevity.

Participants performed 6 trials (4 'mixed', 2 'deterministic'); the order of these was randomized for each subject.

Procedure. The procedure was identical to Experiment 1, except that in the instruction phase, participants were additionally provided an opportunity to use a scanner of each reliability as many times as they wanted, on a number line in which the target number was not hidden, in order to become fully familiar with the scanners they would later use. The intent was to have participants learn about the reliabilities, so for this familiarization stage, all scanners were of equal length.

Results

Scanner choice: length + reliability First we examine whether participants are able to select questions with the best abstract features – that is, do they select scanners whose length and reliability combine to provide the highest expected information gain?

As in Experiment 1, we computed the expected information gain for each of the four scanners, conditioned on the placement actually chosen by each participant; we then ranked the scanners according to this EIG and examined where each participant's scanner choice fell in these rankings. A repeated-measures one-way ANOVA showed that choice proportions differed significantly across scanners [F(3,25) = 61.66, p < 0.0001]. A post-hoc t-test showed that the difference between the first and second scanner ranks was significant [t(25) =7.94, p < 0.0001].

Figure 3A shows the distribution over all scanner choices for the 'mixed' condition, together with the predictions of a softmax version of the EIG model, to be explained later. Participants frequently chose the best scanner, roughly 60% of the time across a wide range of length/reliability trade-offs, which suggests that they are sensitive to the expected information gain of both the scope and reliability of queries rather than to either of these factors alone. Across the 8 sub-conditions that comprised the 'mixed' condition, the probability that participants chose the best scanner also correlated with the EIG model's choice probabilities (Fig. 3B, $\rho = 0.606$).

Joint choices of scanner and placement Having determined that participants correctly weigh the trade-off between the abstract features of scope and reliability to select the best-available scanners, we can examine how



Figure 3: **A**: The 'mixed' condition reveals that participants are able to select the best scanner across a wide variety of reliability/length trade-offs. **B**: The proportion of times participants chose the most informative scanner (x axis) versus predictions of the *EIG* model (y axis), for each of the 8 sub-conditions in the 'mixed' condition.

well they select the specific queries they make.

Given a choice of scanner, how sensitive are participants to the expected information gain of where they place the scanner along the number line? For each decision, we ranked the expected information gain of each possible placement of the scanner chosen by participants. Participants chose a placement in the top 10th percentile 38% of the time, and in the top 20% almost 60% of the time. A repeated-measures one-way ANOVA showed that choice proportions differed significantly across percentile bins [F(9,25) = 55.43, p < 0.0001]. A post-hoc t-test showed that the difference between the first and second bins was significant [t(25) = 7.11, p < 0.0001].

We can also ask about the overall quality of the queries made, over all possible choices of query scope, reliability, and location. Participants are highly sensitive to expected information gain in this space, selecting queries in the top 10% of the available set more than 50% of the time (Fig. 4). A repeated-measures one-way ANOVA showed that choice proportions differed significantly across percentile bins [F(9,25) = 71.4, p < 0.0001]. A post-hoc t-test showed that the difference between the first and second bins was significant [t(25) = 7.24, p <0.0001]. They are also decreasingly likely to select an action in each of the subsequent percentiles, with a profile very well fit by the *EIG* model ($\rho = 0.994$).



Figure 4: In their overall query choices (combining scanner scope, reliability, and location), participants are highly sensitive to expected information gain; note also the tight fit of the *EIG* model ($\rho = 0.994$).

Model comparisons

The above results suggest that participants are sensitive to expected information gain (EIG), rather than making selections on the basis of reliability or length alone. To test this claim more rigorously, we fit five alternative models in addition to *EIG*, which make decisions according to the following criteria: *length*, *reliability*, *length* + *reliability*, *EIG-length*, *EIG-reliability*. Each model computes a 'value' of an action, V(a), as a linear function of the stated parameters: $V(a) = \sum_i \beta_i f_i$, where f_i is some feature of the current trial/action (i.e., EIG, length, or reliability) and β_i is a coefficient fit to each participant by maximum likelihood (when all features share the same coefficient, β is often referred to as the *inverse temperature*). This value is then transformed to a choice probability according to the softmax function, $P(a) \propto \exp\{V(a)\}$.

The *reliability*, *length*, and *reliability+length* models test the hypothesis that participants are sensitive to these features of scanners without using them to compute expected information gain. The *EIG-reliability* and *EIG-length* models test the hypothesis that participants are sensitive to expected information gain but insensitive to one or another feature of queries. Specifically, in the *EIG-reliability* model, decisions are made according to the expected information gain that arises from considering the reliabilities of the scanners, but ignoring their lengths.⁴ The *EIG-length* model, on the other hand, takes the correct scanner lengths into account but assumes a reliability of 1.0 for each scanner. The original *EIG* model tests the hypothesis that participants are sensitive to expected information gain with no qualifications.

For each model, we computed the (participantspecific) Bayesian Information Criterion approximation to the marginal likelihood, and then submitted the models to the Bayesian model selection algorithm of Stephan et al. (2009), which estimates the group-level exceedance probability for each model (the probability that a given model is more likely than all other models considered). The exceedance probabilities for the models are as follows: *EIG*: 0.9982, *EIG-reliability*: 0.0000, *EIG-length*: 0.0004, *reliability*: 0.0068, *reliability+length*: 0.0000, *length*: 0.0000.

How do the models predict participants' actions, in general? We computed the average log likelihood for participants' joint choices of scanner (length and reliability) and placement (the specific location on the number line queried) under each model (Fig. 5) and find that the full *EIG* model, sensitive to all three factors, performs best. The *EIG-length* and *EIG-reliability* models do not fare well, as they operate on incorrect assumptions (that all scanners were deterministic, or that all scanners were of the same average length, respectively) which dis-

⁴There is no way to ignore length in these calculations while still computing EIG, so to to model ignorance in this case we set the length of each scanner equal to the average length of the available scanners.



Figure 5: Log likelihoods assigned by each model to participants' joint choices of scanner type and placement.

tort the EIG calculation. The *reliability* and *reliability* + *length* models do fairly well here, but looking at scanner placements conditioned on scanner choices, they predict no better than chance; this is the main reason they fare worse than the full *EIG* model in predicting overall actions choice.

Discussion

The ability to search efficiently for information in large and noisy environments is critical for real-world learning and discovery. Constructing scientific theories or their intuitive analogues hinges on being able to successfully test competing hypotheses, which often requires balancing the epistemic virtues of scope - to deal with the world's enormity - and reliability - to deal with the world's noise. In this paper we have shown that adults can appropriately balance these trade-offs, effectively conducting the most informative tests when conducting information search with large action spaces. This result is consistent with previous work on information search (cf. Nelson et al., 2013), but more closely resembles natural exploration because of the large action spaces used and the fact that participants had to confront unreliable information sources and the scope-reliability tradeoff.

In our attempt to understand how humans deal with this trade-off, we compared several alternative computational models. Three were sensitive to the abstractions of scope and reliability without any regard for the way in which these factored into a query's expected information gain; three others calculated expected information gain and used scope and reliability either explicitly or implicitly. We found that human actions were best explained by a model that selects actions according to their expected information gain, using both scope and reliability along with the precise location of a proposed test in the EIG calculation.

While this study was designed to capture abstract features of information search in many natural environments, one might be concerned that our results are specific to visuospatial or physical domains in which the target exists at some location and participants might be able to deploy geometric intuitions that may not hold for more abstract forms of information search. To some extent our study is less prone to this objection than other recent work using spatially organized tasks (Markant & Gureckis, 2012), because of the added complexity introduced by varying reliability, which does not have an immediate spatial component. Still, in future work we plan to extend both our modeling and our experimental paradigm to entirely nonspatial domains, where scope and location have more abstract interpretations.

One salient feature of our computational model is that it evaluates actions only on their one-step information gain; it is 'myopic', in AI terms. Perhaps human information search can effectively look further ahead than one step, considering EIG over the many possible outcomes of multiple sequential choices. It is also possible that myopic EIG describes human search best. We plan to test this in future work.

Recently Markant & Gureckis (2012) showed that adults engage in 'disinterested', rather than 'interested', search; their actions are more consistent with a model that maximizes information gain than with one that aims to maximize tangible rewards. But in both their work and ours, the information sought is still arguably tied to some reward: in both cases, participants were given the goal of obtaining some specific chunk of information. Our ultimate goal is to explain pure curiosity as rational information search, and so we are exploring extensions of our current studies and models in environments that trigger curiosity *naturally*, without the need for any explicitly posed search goal, that appeal to children as much as to adults.

References

- Markant, D., & Gureckis, T. (2012). Does the utility of information influence sampling behavior? In Proceedings of the 34th annual conference of the cognitive science society (pp. 719–724). Austin, TX.
- Navarro, D. J., & Perfors, A. F. (2011). Hypothesis generation, sparse categories, and the positive test strategy. *Psychological Review*, 118, 120-134.
- Nelson, J. D., Divjak, B., Gudmunsdottir, G., Martignon, L. F., & Meder, B. (2013). Children's sequential information search is sensitive to environmental probabilities. *Cognition*, 130, 74–80.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101.
- Stephan, K., Penny, W., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *Neuroimage*, 46, 1004-1017.
- Wason, P. (1968). Reasoning about a rule. *Quarterly* Journal of Experimental Psychology, 20, 273–281.