

Does the utility of information influence sampling behavior?

Doug Markant (doug.markant@nyu.edu)
Todd M. Gureckis (todd.gureckis@nyu.edu)
New York University
Department of Psychology
6 Washington Place, New York, NY 10003 USA

Abstract

A critical aspect of human cognition is the ability to actively query the environment for information. One important (but often overlooked) factor in the decision to gather information is the cost associated with accessing different sources of information. Using a simple sequential information search task, we explore the degree to which human learners are sensitive to variations in the amount of utility related to different potential observations. Across two experiments we find greater support for the idea that people gather information to reduce their uncertainty about the current state of the environment (a “disinterested”, or cost-insensitive, sampling strategy). Implications for theories of rational information collection are discussed.

Keywords: information sampling, active learning, information access costs

Introduction

From controlling the movement of our eyes to determining which sources of news to consult, judging the quality of alternative sources of information is a critical part of adaptive behavior. Research exploring how people make information gathering (or “sampling”) decisions has shown that people can discern subtle differences in the potential information value of various aspects of the environment. For example, measurements of eye movements during object categorization show that people preferentially allocate attention to object features that are most useful for making subsequent classification decisions (c.f., Nelson & Cottrell, 2007; Rehder & Hoffman, 2005).

One aspect that typically complicates the analysis of information sampling behavior is that information rarely comes for free. All natural tasks involve information access costs, even if the only cost is the time required to gather information (Fu, 2011). In addition, different pieces of information may be more useful depending on how one will be tested in the future. Optimal search behavior must weigh the costs of collecting particular bits of information against the benefit it is expected to convey (Edwards, 1965; Juni, Gureckis, & Maloney, 2011; Tversky & Edwards, 1966), a point frequently made in research on animal foraging (Stephens & Krebs, 1986).

Despite its importance, previous work on information sampling has often failed to test whether people take into account costs related to different sources of information. For example, Nelson (2005) provides a comprehensive review of various ways an optimal Bayesian agent might value potential information sources in the absence of task-specific costs (see also Nelson et al., 2010). One conclusion from this line of work is that people make information search decisions that are consistent with normative measures of information value (many of which often make similar predictions). For example, Nelson et al. (2010) studied information sampling in a diagnostic

reasoning task where the predictions of these measures could be readily distinguished. Learners who could query different stimulus features before making classification decisions were found to prefer to learn about features that maximized *probability gain*, a measure of how well a potential observation is expected to improve classification accuracy.

In these studies, however, costs were not explicitly manipulated or controlled. Taking costs into account can alter the optimal strategy in a given task, but it is unclear whether people adjust their behavior in a similar way. The goal of the present paper is to explore the impact of costs on sampling decisions. We begin by evaluating two alternative objectives that people may adopt when deciding what information to gather. Like the models reviewed by Nelson (2005), the first ignores the implications of task-specific costs and casts information sampling strictly in terms of uncertainty reduction (i.e., information gain). The second approach balances the costs and expected benefits of information in the context of the task. We then describe the results of two experiments that manipulate the concordance between these two approaches, in one case creating an environment where the goals of minimizing uncertainty and maximizing utility predict different patterns of information sampling. Our results show that people tend to value information (in terms of the number of hypotheses ruled out by a new observation) over situation-specific costs and benefits. The implications of these results for theories of information sampling are discussed.

The rational analysis of information sampling: comparing “interested” and “disinterested” search

How should a rational agent make information sampling decisions? Existing proposals fall into two broad categories which, borrowing from Chater, Crocker, and Pickering (1998), we call “interested” and “disinterested.” Unlike the distinctions explored by Nelson (2005), these two proposals differ significantly in terms of the overall goal of information sampling.

Interested (or cost sensitive) sampling The first approach represents a decision-theoretic approach to information sampling. In particular, the agent considers the cost for collecting a piece of evidence and weighs this against the expected benefit it should convey with respect to the goals of the task. For example, a car shopper might decide if the possible savings available from obtaining information contained in a vehicle history report is worth the cost of the report. Similarly, preferentially fixating the features of an object that are diagnostic of its category membership may be a cost-sensitive strategy under the

assumption that additional fixations cost time and the number of fixations needed to reach a correct decision should be minimized. Sampling in this case is “interested” in that information acquisition is focused on some purpose or goal beside acquiring the information itself. In many ways, “interested” sampling is a fully rational strategy and this formulation is often adopted in economics research.

Disinterested (or cost insensitive) sampling The second approach values information to the degree that it reduces our uncertainty about the world. Chater et al. liken this to basic research where the goal is learning without regard for the ultimate utility of this knowledge for society. In their words, “inquiry is valuable for its own sake, because it leads to knowledge” (Chater et al., p.4). Disinterested inquiry can be conveniently expressed as actions which have a high probability of reducing the Shannon entropy over the agent’s beliefs (Lindley, 1956; Mackay, 1992). Critically, disinterested inquiry doesn’t depend on the costs associated with collecting information or using it to make subsequent decisions.

The basic premise of the experiments reported in this paper is that these two strategies or ways of valuing information can be dissociated on the basis of observed choice behavior. We gave participants a simple, intuitive information search task where they were asked to make sequences of decisions to reduce their uncertainty about a hidden target. Mathematical models instantiating each of the two theories just described are then fit to the choice patterns of individual subjects. This fitting procedure (common in the reinforcement learning literature) allows us to evaluate which of the approaches we have described gives the best account of participants’ choices.

The Experimental Task: The Shape Search Game

Participants in our task are presented with a 10 by 10 grid that contains non-overlapping hidden shapes made up of individual grid cells. The hidden targets are randomly drawn from a set that is known to the participant. There are two phases in each game: a *sampling phase* and a *painting phase*. In the sampling phase, the player learns about the form and location of the hidden shapes by choosing squares in the grid to uncover. On each trial, they make an observation at one location, revealing either part of a hidden shape or an empty square. When they think they know the location and form of all the shapes they can stop sampling and enter the painting phase. In the painting phase, the player is tested for their knowledge of the shapes by “painting” any remaining squares they believe belong to one of the shapes in the appropriate color.

The player is penalized one point for every observation made in the sampling phase and two points for every error committed in the painting phase (e.g., failing to fill in a square that belongs to a shape). These costs promote efficient information search in two ways. First, the observation cost discourages sampling in locations whose contents can be inferred from evidence that has already been uncovered. Second, it encourages continued sampling while there is still uncertainty about the hidden shapes, since painting errors are more costly

Experiment 1: Rectangle Search




Figure 1: The generative process underlying the shape search game. A fixed set of possible shapes is specified. A hidden gameboard is created by sampling from this set of shapes and randomly arranging the targets in the grid. During the sampling phase the participants clicks on grid locations to reveal their contents. In the painting phase, the subject draws in the remaining squares using the mouse and is rewarded for accuracy.

than observations. The goal of the game is to finish with the lowest score possible, which is achieved by learning the most about the hidden shapes in the fewest number of observations.

Based on past work with this task (Gureckis & Markant, 2009), we have found that the overall objective of the game is easily understood by the participants. A critical feature of the task (which we exploit in our second experiment) is that it allows for arbitrarily defined targets (i.e., the target shapes may be composed of any configuration of squares) that can be manipulated to vary the complexity of the task.

Formal task analysis

In order to evaluate both “interested” and “disinterested” information search in the task, we compare the search behavior of subjects to that of a rational learner who updates their beliefs about the gameboard in an optimal way. Formally, players make a sequence of observations in order to learn the hidden gameboard, $g_h \in G$, where G is the universe of legal gameboards. Each individual gameboard is defined by an arrangement of N non-overlapping shapes $\{r_1, r_2, \dots, r_N\}$ with unique labels $\{l_1, l_2, \dots, l_N\}$, and each shape consists of a set of squares such that $o_{ij} \in r_n$ (where i and j index the x - and y -coordinates of the square).

On each trial the player makes an observation o_{ij} and receives feedback in the form of a label l_n , where l_0 indicates the observed square is empty, l_1 means it contains part of shape r_1 , and so on. Since each square in the grid is deterministically assigned to either one or zero shapes, we assume that the likelihood of a particular observation o_{ij} belonging to one of the shapes (i.e., $o_{ij} = l_n$ for $n > 0$) for a particular gameboard g is deterministic (i.e., $p(o_{ij} = l_n | g) = 1$ if the location falls within a target and 0 otherwise).

The prior belief about the likelihood of each individual gameboard is represented by $p(g)$. In our experiments, participants were instructed that the shapes were chosen at random and that all legal gameboards were equally likely (i.e., $p(g) = 1/|G|$ for all g , a uniform prior).

Bayes rule can be used to compute the posterior belief

about the identity of the hidden gameboard and to predict the marginal probability of any point in the grid having any particular label l_n (this is a very straightforward Bayesian approach to the problem, see Gureckis and Markant, 2009).

Interested (cost-sensitive) Sampling The objective of the game is to minimize the number of points accumulated, where each individual observation costs C_{obs} points and each error during painting costs C_{error} points. Given these constraints, we can quantify the value of observations with respect to the overall goal of minimizing total costs. We assume that the likelihood of labeling a point o_{ij} with label l_n during the painting phase is simply the marginal probability $p(o_{ij} = l_n | \mathcal{B})$, and the cost associated with that action is tied to the uncertainty about its label when the sampling phase ends (e.g., if $p(o_{ij} = l_n | \mathcal{B}) = 1$, the true label is known with certainty and there is no chance of committing an error during painting)¹. On each trial, the total expected cost $EC(\mathcal{B})$ of ending the sampling phase and entering the painting phase is defined as:

$$EC(\mathcal{B}) = C_{error} \cdot \sum_i \sum_j \sum_n p(o_{ij} = l_n | \mathcal{B}) \cdot [1 - p(o_{ij} = l_n | \mathcal{B})] \quad (1)$$

For a new observation and its observed outcome ($o_{ij} = l_n$) we then calculate the resulting cost *savings*, or reduction in expected costs:

$$S(\mathcal{B}, o_{ij} = l_n) = EC(\mathcal{B}) - EC(\mathcal{B} | o_{ij} = l_n) \quad (2)$$

The savings achieved from feedback is offset by the cost of making the observation (C_{obs}). To account for uncertainty about the true outcome we find the *expected savings* by weighting the savings for each outcome by its likelihood of occurring:

$$ES(\mathcal{B}, o_{ij}) = -C_{obs} + \sum_n p(o_{ij} = l_n | \mathcal{B}) S(\mathcal{B}, o_{ij} = l_n) \quad (3)$$

For each trial, $ES(\mathcal{B}, o_{ij})$ is calculated for all o_{ij} , giving a distribution of the expected saving for remaining observations in the grid. An ideal learner maximizes this value by choosing the location o_{ij} with the highest $ES(o_{ij})$ on each trial.

Disinterested (cost insensitive) Sampling A “disinterested” sampling norm values observations according to their effect on the learner’s beliefs without account for task-specific costs and benefits. This captures the intuition that observations that produce a large change in the agent’s beliefs tend to be more useful than observations that have little or no effect (i.e., nothing new is learned). In our approach this was modeled using *information gain*, which values an observations according to the expected reduction of uncertainty about the hidden gameboard. This uncertainty can be quantified by the Shannon entropy measured over the current belief distribution ($H(\mathcal{B})$).

Entropy is maximized when all hypothesized gameboards are equally likely (as with our initial uniform prior), and minimized when there is only one possible hypothesis. For a given

¹For shorthand, \mathcal{B} represents a vector of probabilities $p(g | o_{ij} = l_n)$, for all $g \in G$, that represents the full posterior distribution over the entire space of gameboards. This is the agents current “belief distribution” about which gameboard is the current target.




Figure 2: Cumulative frequency of ranks assigned to participants’ samples. A: In Experiment 1, the average rank of participants’ choices is higher than expected from a random sampling strategy, but there is no difference between rankings assigned by the two models. B: In Experiment 2, participants’ samples are more highly ranked according to *EIG* than *ES*.

observation and its observed outcome ($o_{ij} = l_n$) we can calculate information gain as:

$$I(\mathcal{B}, o_{ij} = l_n) = H(\mathcal{B}) - H(\mathcal{B} | o_{ij} = l_n) \quad (4)$$

To account for uncertainty about the true outcome of an observation, information gain for each possible outcome is weighted by the predicted probability of that outcome occurring, giving the *expected information gain* for an observation o_{ij} :

$$EIG(\mathcal{B}, o_{ij}) = \sum_n p(o_{ij} = l_n | \mathcal{B}) I(\mathcal{B}, o_{ij} = l_n) \quad (5)$$

As above, on each trial we compute $EIG(\mathcal{B}, o_{ij})$ for all locations in the grid, and assume that the optimal model chooses the location with the highest value on each trial².

In applying each model to human choice data, the model is “yoked” to the decisions of the player. On each trial, the models assign a value (either *EIG* or *ES*) to each point in the grid. These utilities can be used to compute choice probability of various grid locations. After revealing what the subject actually chose on a given trial, the model updates its posterior beliefs about the current gameboard configuration. These new beliefs then feed into new predictions about the utility of choosing each grid location. The process ends when the participant ends the game.

Experiment 1: Rectangle Search

The first experiment re-analyses a previously published result which introduced explicit task-specific costs (Gureckis &

²It is important to note that this represents a “greedy” policy that chooses the best observation available on any given trial, but this may not reflect the globally optimal solution. The current framework could be extended to account for how participants might estimate the value of sequences of observations. However, due to the computational complexity of finding this solution given the large number of potential observations on any trial, for the present studies we focus our analysis on the greedy model.

Markant, 2009). Six participants played a series of games in which they searched for three rectangular shapes, randomly drawn from the set shown in Figure 1. The set of shapes was displayed on screen throughout the game. Participants were instructed that the three shapes in each gameboard were non-overlapping and were shown a large number of examples gameboard configurations prior to the experiment.

Each observation made by a participant during the sampling phase was ranked according to the predictions of both models (the median rank was used when multiple observations had equal value). Overall, the results show that people consistently sampled points that were assigned a high value by both models, with approximately 50% of their samples falling within the top 10 ranked observations available to them (see Figure 2A). In this experiment, however, the hypothesis set that was used (rectangular shapes of varying shape and size) led to highly similar predictions for both information gain and expected savings, precluding a test of whether people were sensitive to the costs in the task.

It is important to consider why the predictions of the two models converged in this case. As discussed previously, a cost-sensitive learner should value observations that have higher utility—that is, those that will reduce the likelihood of committing errors in the painting phase by the greatest amount. Intuitively, this implies that learning about bigger shapes is especially useful, since it will allow one to correctly label a greater number of squares. This idea is illustrated in Figure 3A for a simple hypothesis set made up of three rectangles in different locations. While observing a “hit” on any shape determines the true hypothesis (middle column), observing a “miss” (righthand column) has different utilities depending on the size of the shape it rules out. For example, ruling out the smallest shape (top row) leaves uncertainty about how to label eight other squares, whereas ruling out the largest shape (bottom row) leaves uncertainty about only four.

While the shape set used in Experiment 1 contained a range of sizes, the fact that the hypotheses were “nested” (i.e., the largest shapes overlapped with a set of progressively smaller ones) meant that learning about larger shapes also tended to rule out many hypotheses. As a result, the predicted choice values according to both models were highly similar. For our second experiment we created an alternative hypothesis space in which there were clearer differences in both the size of alternative hypotheses and the choices that were related to shapes of different size, leading to a greater number of potential observations where the predictions of the two models diverged.

Experiment 2: Letter Search

In Experiment 2, we simplified the task to involve searching for a single target in the grid. For the hypothesis space we created a set of simple “letter” shapes (see Figure 3B). There were two types of games: L/D games, where the hidden letter could be an L or D, and C/U games, where the hidden letter could be a C or U. In each game a single point belonging to the hidden shape was revealed before the participant began sampling.

These modifications resulted in a much less complex hypothesis space (e.g., only 15 hypotheses were possible at the beginning of an L/D game, and 14 for C/U games). Importantly, because the two shapes involved in each type of game differed in area (for example, the ‘D’ shape contained a greater number of filled squares than the ‘L’), the predictions of information gain and expected savings diverged in the task.

Methods

Participants Sixteen NYU undergraduates completed the study for course credit or for a \$10 payment. The experiment was presented on a standard Macintosh computer.

Materials A gameboard consisted of one of the four letter shapes seen in Figure 3 placed in any location on the grid. All possible gameboards were generated, resulting in 56 gameboards for each letter, or 112 for each game combination (L/D or C/U). For L/D games, 15 gameboards were randomly selected from the pool of L and D gameboards. This was repeated for C/U games. The resulting total of 30 unique games were used for all participants. The order of games played was randomized for each person.

Procedure Many aspects of the design were identical to Experiment 1 (described in Gureckis and Markant, 2009), so we highlight differences below. In Experiment 2 we sought to reduce any reliance on the visual display for recalling the specific shapes being used. Participants began the experiment with two training phases to memorize the four shapes. In the first, a letter cue (e.g., the character ‘L’ in a standard computer font) was presented at the top of the screen along with its corresponding shape, which appeared inside a 4x3 grid. The participant was asked to copy the same shape onto an empty 4x3 grid below. This was done twice for each letter (L, D, C, U) in randomized order. In the second training phase, they were presented only with the letter cue and an empty 4x3 grid, and asked to fill in the correct letter shape from memory. This was repeated three times for each letter in random order. In order to progress from one training trial to the next, the participant was required to successfully reproduce the correct shape. Training was followed by on-screen instructions which were modified to reflect the new hypothesis spaces.

Before the sampling phase began, a 2-second cue was displayed on the screen that indicated the type of game about to be played (the characters “L D” or “C U”). This cue was also displayed on the right side of the display during the game. This ensured that participants were aware of the shapes that were possible but that they had to use their memory of the actual shapes to guide their observations.

Sampling and painting phases proceeded in the same way as Experiment 1. The final score was then displayed, including how many points were the result of sampling and how many were the result of painting errors. The lowest score obtained by the participant in any game so far was shown to provide motivation and a means to evaluate their performance over time. Each participant played 30 games at their own pace, resulting in a total of 480 games collected.

Results

Sample rank On each sampling trial, the ideal models were used to compute the expected information gain and expected savings for all remaining observations available. A participant’s decisions were ranked according to EIG and ES (if multiple observations had the same value, the median rank was used). The relative frequency of each sample rank was computed for each participant across all trials, and averaged across participants (see Figure 2B). The rank frequency shows that participants’ choices were more highly ranked on average according to EIG than ES. Participants’ samples were ranked within the top 5 observations according to EIG on approximately 57% of trials, whereas according to ES only 35% of samples fell in the same range.




Figure 3: **A:** Illustration of the divergence between information gain and expected savings. **Left:** A hypothesis space is comprised of three possible rectangles, h_1 , h_2 , and h_3 . **Middle:** A hit on any one rectangle leads to the same number of hypotheses being ruled out, and no uncertainty about the label of any square in the grid. **Right:** A miss in any of the three locations rules out a single hypothesis, but the *predictive utility* of the sample differs based on its location. The labels for 8 squares are uncertain following a miss that rules out rectangle h_1 , 6 squares following a miss ruling out h_2 , and 4 squares following a miss ruling out h_3 . **B:** Gameboard design and example model predictions in Experiment 2. Two types of games were possible: L/D games, where the hidden shape could be an L or D, and C/U games, where the hidden shape could be a C or U. Predicted value distributions for EIG and ES are shown for the first sampling trial in each kind of game, with a darker value indicating a higher value according to the model.

Model fits We next computed the likelihood of participants’ decisions under the two alternative models. For each trial, the value of available observations was transformed into choice probabilities using the softmax function:

$$P(o_{ij}) = \frac{e^{\beta \cdot V(o_{ij})}}{\sum_{x,y} e^{\beta \cdot V(o_{xy})}} \quad (6)$$

The parameter β was fit on an individual basis for each model by maximizing the log-likelihood summed across all observations made by a participant. In all cases, EIG provided a better fit to participants’ data than ES (Table 1).

Stopping decisions Our final analysis focused on participant’s decisions to stop sampling. While EIG and ES make the same prediction as to when sampling should stop³, we were interested in whether people showed any sensitivity to the cost of collecting information. If people uncovered more squares than necessary it would suggest a failure to account for the cost of new observations (either in terms of ES or EIG). We classified each game according to whether the person decided to stop sampling before the trial predicted by the model (“early”), on the same trial (“optimal”), or after that trial (“late”). On average, participants ended sampling early ($M = 0.46$, $SD = 0.14$) or on the optimal trial ($M = 0.50$, $SD = 0.14$) on a similar proportion of games. In contrast, participants oversampled very rarely ($M = 0.04$, $SD = 0.01$).

³This convergence was due to the cost structure we used, in which the penalty for stopping before the hidden target was known was greater than the cost of an additional sample. Increasing the cost of sampling relative to the cost of errors would lead to ES predicting earlier stopping decisions than EIG.

Discussion

The results of our second experiment show that people performed well compared to the ideal searcher model, frequently choosing highly ranked observations and consistently performing better than expected by a random search strategy. In addition, participants rarely gathered more information than necessary, which is consistent with prior work showing that people are sensitive to costs incurred by oversampling (Fu & Gray, 2006). Most importantly, we found that participants’ sampling choices were better described by information gain than a cost-sensitive utility measure (expected savings).

Prior studies of human information collection have focused to a large extent on “disinterested” accounts of sampling decisions, showing that people are sensitive to the amount of information conveyed by different sources of information (Nelson, McKenzie, Cottrell, & Sejnowski, 2010). However, this line of work has often failed to consider whether people account for variations in task-specific utility when making sampling decisions. In our task, we introduced penalties for information access and uncertainty that altered the optimal sampling strategy. By manipulating the set of hypotheses in Exp. 2, we showed how sampling based on an information-maximizing strategy can be distinguished from cost-sensitive sampling. With respect to the distinction we began this paper with, our results suggest that people (at least in this task) prefer to gather information according to a “disinterested” measure of value.

Notably, differences in value between choices were not directly observable by the participant (as opposed to when some observations are more costly or difficult to make than others). Whether a given observation was valued differently according to EIG and ES depended upon the set of hypotheses remaining, and this divergence could change from trial to trial in response to new data. As a result, establishing which model pro-

vided a better account required fitting them to participants' decisions across a variety of choice contexts. This highlights a feature of our approach in that we could evaluate different sampling strategies using a set of highly variable choice sequences. Through the use of a well-defined hypothesis set and explicit cost structure, the "shape-search" task provides a useful framework for studying how information search decisions and task demands interact over the course of learning.

Of course, one potential caveat of the current study is that (due to computational demands) we evaluated a greedy decision policy such that the predicted value of a new observation does not take into account the consequence or utility of subsequent actions. It is possible that fully accounting for sequential dependencies in the search problem may alter the optimal utilities computed by the model. However, one might reasonably question if the computational demands of such a multi-step planning process are within reach of human reasoners. In addition, it is unclear that accounting for multi-step planning strategies would alter the choice utilities in a way that would bias for or against the results we report. Also note that comparing Exp. 1 and 2 illustrates that expected saving is not always at an advantage (i.e., in some choice environments the models become less distinguishable).

Conclusion

So, why might human reasoners preferentially adopt "disinterested" sampling over "interested" sampling? One possibility is that sampling based on information gain (or other "disinterested" norms) may reflect a general purpose strategy that is useful in a variety of contexts. In particular, information gain can still be computed even when the cost of uncertainty (i.e., not knowing which hypothesis is true at the end of sampling) is difficult to predict. In addition, in many natural environments it may be consistent with the predictions of a cost-sensitive utility function, as illustrated by our first experiment. At the very least, our results highlight the need to understand the kinds of problems that lead people to adapt to task-specific costs in lieu of a general-purpose, "disinterested" approach to information search.

Table 1: Model fits

Subj	β_{EIG}	β_{ES}	-LLH(EIG)	-LLH(ES)
1	17.544	0.95	215.05	334.65
2	5.18	0.88	391.77	426.18
3	5.93	0.77	326.45	388.79
4	6.96	0.75	325.81	403.61
5	10.15	0.89	241.04	324.30
6	12.34	1.06	250.17	329.72
7	9.64	1.20	276.05	339.62
8	7.63	0.98	308.63	373.03
9	5.36	0.90	354.49	387.43
10	7.79	0.91	315.92	387.56
11	5.85	0.85	340.22	388.05
12	6.39	0.92	328.24	373.22
13	6.97	1.03	319.96	362.27
14	7.02	0.86	283.32	341.79
15	8.94	0.99	293.38	371.22
16	10.0	0.76	293.65	403.22

Acknowledgments

The authors would like to thank Dan Navarro, Jonathan Nelson, and the other reviewers for their helpful comments. This work was supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of the Interior (DOI) contract D10PC20023. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI, or the U.S. Government.

References

- Chater, N., Crocker, M., & Pickering, M. (1998). The rational analysis of inquiry: the case of parsing. In M. Oaskford & N. Chater (Eds.), *Rational models of cognition* (pp. 441–468). Oxford University Press.
- Edwards, W. (1965). Optimal strategies for seeking information: models for statistics, choice reaction times, and human information processing. *Journal of Mathematical Psychology*, 2, 312–329.
- Fu, W., & Gray, W. (2006). Suboptimal tradeoffs in information seeking. *Cognitive Psychology*, 52(3), 195–242.
- Fu, W. (2011). A dynamic context model of interactive behavior. *Cognitive Science*, 35(5), 874–904.
- Gureckis, T., & Markant, D. (2009). Active learning strategies in a spatial concept learning game. In N. Taatgen & H. van Rijn (Eds.), *Proc of the 31st annual conference of the cognitive science society* (pp. 3145–3150). Cognitive Science Society. Austin, TX.
- Juni, M., Gureckis, T., & Maloney, L. (2011). Don't stop 'till you get enough: adaptive information sampling in a visuomotor estimation task. In L. Carlson, C. Hölscher & T. Shipley (Eds.), *Proceedings of the 33rd annual conference of the cognitive science society*. Cognitive Science Society. Austin, TX.
- Lindley, D. (1956). On a measure of the information provided by an experiment. *Annals of Mathematical Statistics*, 986–1005.
- Mackay, D. (1992). Information-based objective functions for active data selection. *Neural Computation*, 4, 590–604.
- Nelson, J. (2005). Finding useful questions: on bayesian diagnosticity, probability, impact, and information gain. *Psychological Review*, 114(3), 677.
- Nelson, J., & Cottrell, G. (2007). A probabilistic model of eye movements in concept formation. *Neurocomputing*, 70(13–15), 2256–2272.
- Nelson, J., McKenzie, C., Cottrell, G., & Sejnowski, T. (2010). Experience matters: information acquisition optimizing probability gain. *Psychological Science*, 21(7), 960–969.
- Rehder, B., & Hoffman, A. (2005). Eyetracking and selective attention in category learning. *Cognitive Psychology*, 51, 1–41.
- Stephens, D., & Krebs, J. (1986). *Foraging theory*. Princeton, NJ: Princeton University Press.
- Tversky, A., & Edwards, W. (1966). Information versus reward in binary choices. *Journal of Experimental Psychology*, 71(5), 680–683.