


Supplement to "Experience matters: information acquisition optimizes probability gain", by Nelson, McKenzie, Cottrell, and Sejnowski, to appear May, 2010 in Psychological Science.

Contact: Jonathan Nelson (jonathan.d.nelson@gmail.com, <http://jonathandnelson.com/>)

### *Plankton stimuli*


The actual plankton stimuli appear below (Figs S1-S3). Our plankton stimuli, though hopefully naturalistic in appearance, should not be confused with real copepods. (For instance, the claw feature did not occur in the original images.) The stimuli were designed to have three subtly-varying two-valued features (tail, eye, claw), roughly equidistant from each other. We thank Profs. Jorge Rey and Sheila O'Connell (University of Florida, Medical Etymology Laboratory), for allowing us to base our artificial plankton stimuli on their photographs of real copepod plankton specimens.



*Figure S1.* Example plankton stimuli, from learning phase. Specimen at top has fine tail, blurry eye, and unconnected claw. Specimen at bottom has blunt tail, dotted eye, and connected claw



*Figure S2.* Example plankton stimulus, from information-acquisition phase, with eye and claw obscured.



*Figure S3.* The two versions of each plankton feature: blunt or fine claw (left); blurry or dotted eye (middle), and unconnected or connected claw (right)

### Optimization notes

An example illustrates calculation of  $DStr$ , for Condition 1:

$$\begin{aligned} PStr_{PG} &= 100 * ( eu_{PG}(F) - eu_{PG}(G) ) / maxPStr_{PG} \\ &= 100 * (0.072 - 0) / 0.50 = 14.4 \end{aligned}$$

$$\begin{aligned} PStr_{IG} &= 100 * ( eu_{IG}(F) - eu_{IG}(G) ) / maxPStr_{IG} \\ &= 100 * (0.134 \text{ bits} - 0.280 \text{ bits}) / 1 \text{ bit} = -14.6 \end{aligned}$$

$$DStr = ( | 14.4 | * | -14.6 | )^{0.5} = 14.5, \text{ because } PStr_{IG} * PStr_{PG} < 0.$$

In each optimization, obtained feature likelihoods were rounded to the nearest 0.01 for use in the experiments. In Condition 1 (information gain versus probability gain), the original optimizations produced values such as  $P(f_1/a) = 0.04$ ,  $P(f_1/b) = 0.38$ ,  $P(g_1/a) = 0.57$ , and  $P(g_1/b) = 0$ . These values confounded the possibility of knowing for sure with the desired comparison of information gain and probability gain. (Whereas our desired test was between information gain and probability gain, only  $G$  offered the possibility of a certain result. If participants wished to maximize probability of a certain result, and hence preferred  $G$ , this could have been misinterpreted as a preference to optimize information gain.) We therefore repeated the optimization, requiring  $P(f_1/a) = 0$ , just as  $P(g_1/b) = 0$ . This removed that confound while having negligible effect on strength of disagreement. The same confound appeared in Condition 2, and was also remedied by requiring  $P(f_1/a) = 0$ . In Experiment 3 an environment along these lines where  $P(f_1/a) = 0.04$  was tested; results continue to favor probability gain.

Pairwise optimizations of each OED model vs. the probability of certainty heuristic resulted in virtually identical feature likelihoods. In Condition 4, we therefore optimized the disagreement strength of probability of certainty versus the joint preference of all three OED models. (We defined the joint preference of the OED models as the geometric mean of their individual preference strengths.) A further note is that this optimization produced features for which  $P(f_1/a) = \epsilon$ , and  $P(f_1/b) = 1 - \epsilon$ , where  $\epsilon \approx 0.0001$ . Unfortunately, the difference between  $P(f_1/a) = 0$  and  $P(f_1/a) = 0.0001$ , though important for the probability of certainty model, is not learnable in two hours of experience-based training with natural sampling. We therefore redid this optimization, fixing  $F$  such that  $P(f_1/a) = 0.05$ , and  $P(f_1/b) = 0.95$ .

In the optimizations (see Table 1 in the article), a feature where  $P(f_1/a) = 4/7 \approx 0.57$ , and  $P(f_1/b) = 0$ , occurred frequently. This may be because, holding  $P(a) = 0.70$  and  $P(f|b) = 0$  constant,  $P(f|a) = 4/7$  is the highest feature likelihood such that the feature has zero probability gain. In Condition 1 and Condition 2,  $F$  is rarely  $f_1$  (7% or 9% of the time); but if  $F=f_1$ , the probability of species  $b$  changes from 30% to 100%. If  $F=f_2$ , the probability of species  $a$  increases (from 70% to 75% or 77%). If  $G=g_1$ , it is species  $a$  for sure. However, if  $G=g_2$ , it is a 50/50 chance whether the species is  $a$  or  $b$ . These possibilities cancel each other out, such that the overall probability of correct guess is not improved by querying  $G$ , despite  $G$ 's higher information gain and impact. In Condition 3,  $F$  is  $f_1$  12% of the time; if  $F=f_1$  uncertainty is eliminated; information gain prefers  $F$ . If  $F=f_2$  the probability of species  $a$  goes from 70% to 80%, which also reduces uncertainty. Impact depends on the absolute difference in feature likelihoods, which favors  $G$  ( $0.73 - 0.22 = 0.51$ ) over  $F$  ( $0.40 - 0 = 0.40$ ). In Condition 4, all the OED

models, which were jointly optimized versus probability of certainty, prefer  $F$ , which leads to always knowing the true category with high probability, but never for sure.  $G$  leads to knowing the true category for sure 40% of the time, but to lower overall probability correct, to higher uncertainty, and to lesser absolute change in beliefs

#### *Experiment notes*

Between 6% and 22% of participants did not reach criterion performance in each condition of Experiment 1. Condition 1 had 13% nonlearners (4/32); Condition 2, 7% (2/30); Condition 3, 22% (8/36); and Condition 4, 6% (2/31). Condition 3 was difficult because one of its stimulus items, which occurred less than 1/3 of the time, led to only 57% posterior probability of the most-probable category, and thus took a great deal of experience to learn.

Did subjects learn both features  $F$  and  $G$ , as intended, or only marginal probabilities involving a single feature? In some conditions, it is theoretically possible to only learn  $F$ , and yet to achieve the performance criterion. We therefore analyzed the proportion of optimal responses for each configuration of features. (Optimal is choosing the more-probable species, irrespective of how close the posterior probability is to 50%, given a particular configuration. This is true irrespective of which utility a person wishes to optimize in the information-acquisition phase.) We present data for Experiment 1, Condition 1, below; this is representative of the conditions where it is theoretically possible to only learn the  $F$  feature.

If subjects only learned the  $F$  feature, then the green line ('certain-a config,' f2,g1) and the blue line ('uncertain-a config,' f2,g2) would be overlaid, except for random jitter, throughout learning, as these configurations differ only along the  $F$  feature. The results, however, show that subjects differentiated these configurations, quickly mastering the certain-a configuration, yet struggling with the uncertain-a configuration until very late (e.g. the last 4% of learning trials) in the learning process.




Figure S5. Aggregate learning data for Experiment 1, Condition 1.

The difference between the green line (top), for the certain-a configuration ( $f2, g1$ ), and the blue line (bottom), for the uncertain-a configuration ( $f2, g2$ ), demonstrate that subjects learned configurally. The red line depicts the certain-b ( $f1, g2$ ) configuration.

Because different subjects learned in different numbers of trials, and because different configurations of stimuli occurred with different frequencies, the data below are normalized so that the first 1/25th (4%) of trials on a particular configuration is plotted first, the second 4% of trials on a particular configuration is plotted second, etc., for each subject. In this way, rare stimuli and frequent stimuli, and subjects who learned quickly and slowly, contribute equally to the proportion of optimal responses denoted at each point in learning. (Note that the figure requires color.)

What do individual subjects data show? Figure S4 shows every learning trial for each subject in Experiment 1, Condition 1. Each of the 28 rows represents a single subject.

Note the greatly higher rates of suboptimal responding to the uncertain-a configuration (left column), versus the certain-a configuration (middle column), which differ only according to the  $G$  feature. This demonstrates that individual subjects separately (configurally) learned each stimulus item, and did not only learn marginal probabilities associated with the  $F$  feature. Some subjects vacillate between periods of correct and incorrect responding on the uncertain-a configuration, further evidence that they perceive the difference between the configurations.

Could the subjects, once they learned probabilities involving both features and each configuration of features, have forgotten those configural probabilities late in learning, before the information-

acquisition phase?<sup>1</sup> It was possible to debrief the vast majority of subjects following the experiment; the vast majority of these subjects showed high familiarity with environmental probabilities, including the fact that various configurations (though both pointing to species a, for instance) had widely varying levels of certainty.

To more systematically evaluate this qualitative result, we subsequently obtained data from an additional 13 subjects in the Experiment 1, Condition 1, environment. (There was one additional nonlearner.) Eleven of thirteen subjects preferentially viewed the *F* feature, consistent with earlier information-acquisition results. This replication experiment included a new knowledge test page (following the information-acquisition phase) in which subjects were explicitly asked, for each kind of specimen that appeared, the percent of instances in which it had been species a and b. Subjects were also asked which percent of specimens, overall, were species a and b. Analysis of individual subjects' results (Table S1) shows that the vast majority of subjects were qualitatively very close in their beliefs, identifying the more probable species overall, the more probable species given each configuration of features, and the approximate certainty induced by each configuration of features. Thus, subjects preferred the *F* feature given their knowledge of configural environmental probabilities, not because it was the only feature that they learned.

Additional data, describing corresponding analyses of other conditions, are available from the first author. These data show configural learning throughout.

---

<sup>1</sup> Note that this concern is not a theoretical possibility in some conditions, in which responding optimally to all configurations unequivocally implies that a subject effectively differentiates the two features, and not just a single feature. This a theoretical possibility in Experiment 1, Conditions 1 and 2—though it is implausible: note from Fig. S5 that such forgetting would have to have occurred in the last 4% or so of learning trials.

*Figure S4**(at right; notes below)*

Data for learning phase from Experiment 1, Condition 1, from each of 28 individual subjects who obtained criterion performance.

Key: trials are ordered from top to bottom, and left to right, in each rectangle.

Each subject appears on one row; each configuration in one column. Optimal responses are depicted in white; suboptimal responses are depicted in black.

Left column:

uncertain-a (f2,g2; 56.9% are Species a);

Middle column:

certain-a (f2,g1; 100% are Species a);


Right column:

certain-b (f1,g2; 100% are Species b).

The f1,g1 configuration does not occur in this environment.

The higher suboptimal response rates for the uncertain-a configuration (left) than for the certain-a configuration (middle) show that subjects learned configurations of features, and not merely the higher probability gain feature. Suboptimal response rates are statistically greater for the uncertain-a configuration than the certain-a configuration in 26 of 28 subjects, by both difference-of-proportions and bootstrap tests.

In each rectangle, trials are ordered from top to bottom, left to right:



Optimal responses in white; suboptimal in black. One subject per row. Uncertain a (left column), certain a (middle column), certain b (right)



Table S1. Subjects show high calibration to the environmental probabilities.

Item	True percent	Median rating	Mean rating	Individual subjects' probability ratings:												
				#1	#2	#3	#4	#5	#6	#7	#8	#9	#10	#11	#12	#13
P(a f2,g2)	57	65	67	50	55	99	80	65	54	55	75	75	67	77	55	65
P(a f2,g1)	100	100	100	100	100	100	100	95	100	100	100	100	100	100	100	100
P(a f1,g2)	0	0	9	0	0	0	0	20	0	0	0	0	0	95	0	0
P(a)	70	82	80	48	73	90	90	90	79	85	82	75	62	79	94	92

*Note.* The item being judged is in the left column; its true percent next; and the median and mean of subjects' estimated percentages next. Individual subjects (columns #1 to #13, at right) in most cases showed very good learning of environmental probabilities. Whether species a or b was more probable was randomized across subjects. In this table, 'a' denotes whichever species was more probable in a particular subject's randomization.