

## Generalization of color by chickens: experimental observations and a Bayesian model

Article (Published Version)

Baddeley, R J, Osorio, D and Jones, C D (2007) Generalization of color by chickens: experimental observations and a Bayesian model. *American Naturalist*, 169 (S1). S27-S41. ISSN 0003-0147

This version is available from Sussex Research Online: <http://sro.sussex.ac.uk/id/eprint/1819/>

This document is made available in accordance with publisher policies and may differ from the published version or from the version of record. If you wish to cite this item you are advised to consult the publisher's version. Please see the URL above for details on accessing the published version.

### **Copyright and reuse:**

Sussex Research Online is a digital repository of the research output of the University.

Copyright and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable, the material made available in SRO has been checked for eligibility before being made available.

Copies of full text items generally can be reproduced, displayed or performed and given to third parties in any format or medium for personal research or study, educational, or not-for-profit purposes without prior permission or charge, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

# Generalization of Color by Chickens: Experimental Observations and a Bayesian Model

R. J. Baddeley,\* D. Osorio,† and C. D. Jones‡

School of Life Sciences, University of Sussex, Brighton BN1 9QG, United Kingdom

*Online enhancement:* color version of figure 2.

---

**ABSTRACT:** Sensory generalization influences animals' responses to novel stimuli. Because color forms a perceptual continuum, it is a good subject for studying generalization. Moreover, because different causes of variation in spectral signals, such as pigmentation, gloss, and illumination, have differing behavioral significance, it may be beneficial to have adaptable generalization. We report on generalization by poultry chicks following differential training to rewarded ( $T_+$ ) and unrewarded ( $T_-$ ) colors, in particular on the phenomenon of peak shift, which leads to subjects preferring stimuli displaced away from  $T_-$ . The first three experiments test effects of learning either a fine or a coarse discrimination. In experiments 1 and 2, peak shift occurs, but contrary to some predictions, the shift is smaller after the animal learned a fine discrimination than after it learned a coarse discrimination. Experiment 3 finds a similar effect for generalization on a color axis orthogonal to that separating  $T_+$  from  $T_-$ . Experiment 4 shows that generalization is rapidly modified by experience. These results imply that the scale of a "perceptual ruler" is set by experience. We show that the observations are consistent with generalization following principles of Bayesian inference, which forms a powerful framework for understanding this type of behavior.

*Keywords:* chicken, color vision, peak shift, Bayesian, generalization.

---

Responses to novel stimuli are likely to be based on similarity to familiar stimuli, combined with knowledge about how stimuli vary. This variation depends partly on the objects of interest; oranges are edible over a smaller color

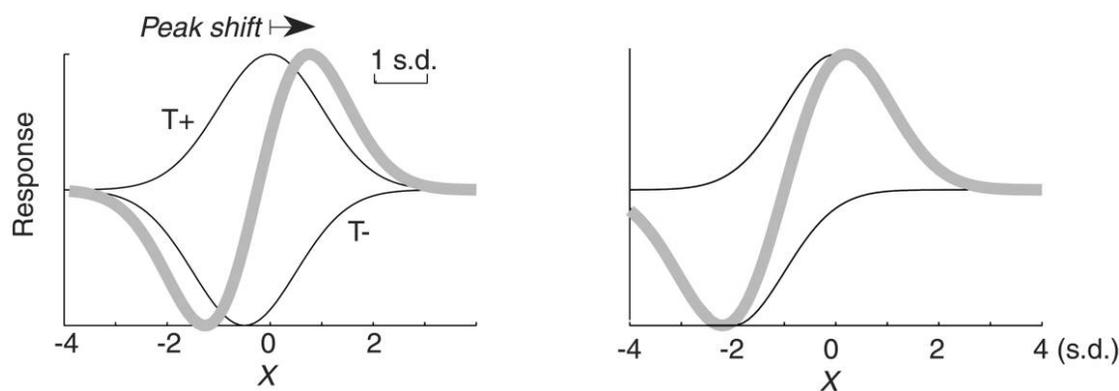
range than apples. More fundamentally, variation is dependent on the physics of the stimulus. Thus, the effect of varying distance from a sound source depends on the acoustic environment, while varying light environments and shadowing distort colors in specific and predictable ways. Handling these problems requires neural "models" of variation, which can be compared to models used in statistics. Ultimately, noise in receptors sets sensory thresholds, so one can ask whether further limits are imposed by subsequent neural processing (Geisler 1989; see Kelber et al. 2003 regarding animal color vision). Whereas discrimination thresholds and receptor noise are fairly easy to define, it is more difficult to give an optimal model to test hypotheses about processes such as generalization and categorization. Bayesian models are increasingly used in psychology and neuroscience to develop testable predictions of optimal performance (Rao et al. 2002; Tenenbaum and Griffiths 2001; Cheng 2002; Cheng and Spetch 2002). Bayesian models have also been applied directly to color perception (Brainard and Freeman 1997), but this work has concerned how representations robust to illumination changes can be optimally derived rather than how we generalize and categorize such stimuli. This article shows how such models might inform our understanding of generalization in color vision.

*Peak shift.* Peak shift is a well-established phenomenon in animal and human (Thomas et al. 1991; Wills and Mackintosh 1998; Ghirlanda and Enquist 2003) perception. After training to two stimuli, subjects give their strongest responses not to the higher-value training stimulus but to a novel stimulus that differs from the neutral (or negatively valued) stimulus in the same way as the former but to a greater extent (fig. 1). The expected outcome of a response may be positive, for example, a food reward, or negative, as with distasteful objects. One can imagine that peak shift may apply not only to learned but also to genetically acquired preferences, such as those involved in mate choice. In evolutionary biology, peak shift is of interest because of its potential relevance to the evolution of "exaggerated" signals, such as those involved in mate choice and warning coloration (Ryan 1990; Weary et al.

\* Present address: Department of Experimental Psychology, University of Bristol, 8 Woodland Road, Bristol BS8 1TN, United Kingdom; e-mail: roland.baddeley@bristol.ac.uk.

† Corresponding author; e-mail: d.osorio@sussex.ac.uk.

‡ Present address: Medical Computing Research Group, School of Computing, University of Glamorgan, Pontypridd, Mid-Glamorgan CF37 1DL, United Kingdom; e-mail: cdjones1@glam.ac.uk.



**Figure 1:** Generalization curves following differential conditioning on a stimulus dimension  $X$ . Horizontal scale is specified as multiples of the standard deviation (SD) of the generalization curve. The model follows Spence (1937) in assuming that where there is  $>1$  known stimulus, the overall generalization curve (gray line) is given by summing curves (black lines) for each known stimulus. Variants on this model that have been applied to human and animal behavior all predict that peak shift should increase as the separation of  $I_+$  and  $I_-$  on  $X$  falls (McLaren and Mackintosh 2002; Ghirlanda and Enquist 2003). Here the plots are for separations of 0.5 (left) and 2 SD (right).

1993; Enquist and Arak 1998; Yachi and Higashi 1998; Lynn et al. 2005). In particular, Enquist and coworkers (Enquist and Johnstone 1997; Enquist and Arak 1998; Ghirlanda and Enquist 1998, 1999) argue that peak shift reflects a mechanistic constraint on perception that is perhaps attributable to the neural network-like properties of the brain (see also McLaren and Mackintosh 2002). A principal aim of our study is to consider the possibility that peak shift in birds is a consequence of generalization behavior following principles of optimal inference. Others have suggested that peak shift may be optimized according to criteria proposed by decision theory (Thomas et al. 1991; Lynn et al. 2005) or Bayesian inference (Cheng 2002). As these papers point out, if perception is indeed optimized for dealing with an uncertain world, it is inappropriate to make any inferences about neural mechanisms.

Previously, we have described how poultry chicks generalize after appetitive conditioning to either a single rewarded color or two rewarded colors (Osorio et al. 1999a; Baddeley et al. 2001; Jones et al. 2001). This work showed that chicks remember color accurately and described how they generalize over different types of color variation, such as hue, saturation, and brightness. Here we investigate differential conditioning and, in particular, the effects of varying the similarity of rewarded ( $T_+$ ) and unrewarded ( $T_-$ ) colors, which results in peak shift.

*Modeling generalization and peak shift.* To describe generalization on some dimension  $X$ , the predicted outcome of responding to a novel stimulus  $x_{\text{novel}}$  can be represented by a plot of response “strength” against stimulus value (fig. 1). With one known stimulus  $x_1$ , generalization curves are typically centered on  $x_1$  and are approximately Gauss-

ian or exponential functions of the stimulus dimension (Spence 1937; Shepard 1987; Ghirlanda and Enquist 2003). Given such a setup, it could be that the width of the generalization curves is dependent either on discrimination thresholds or on the predicted effects of suprathreshold stimulus differences.

If an animal is familiar with more than one stimulus, how should the additional information be used? For two stimuli  $x_1$  and  $x_2$  with profitabilities  $p_1$  and  $p_2$ , one possibility is that the subject sums the separate generalization curves for  $x_1$  and  $x_2$ , weighted by  $p_1$  and  $p_2$  (fig. 1). Spence (1937) proposed this type of model, and it remains influential in theoretical treatments (Pearce 1997; Ghirlanda and Enquist 1998, 2003; Wills and Mackintosh 1998; McLaren and Mackintosh 2002).

A prediction of summation models is that if the respective profitabilities  $p_1$  and  $p_2$  of the two stimuli  $x_1$  and  $x_2$  are of opposite sign, for example,  $p_1 = -p_2$ , then the stimulus giving the largest, or peak, response  $x_{\text{max}}$  will not be  $x_1$  but will be displaced away from the less profitable  $x_2$  (fig. 1). The magnitude of this peak shift should increase as the separation between  $x_1$  and  $x_2$  decreases (Wills and Mackintosh 1998; Ghirlanda and Enquist 1999, 2003).

Although the basic model is straightforward and the existence of peak shift is well established, there are relatively few experimental tests that compare possible models, perhaps in part because of the absence of alternative hypotheses. An additional difficulty is that of defining the perceptual continuum necessary to specify the form of generalization curves. One possibility, used for humans (Wills and Mackintosh 1998), is to make an artificial stimulus dimension using discrete features. Here we use our

knowledge of spectral coding and receptor spectral sensitivities to specify a continuum in color space.

Previous studies of pigeons confirm that color peak shift occurs in birds (Guttman and Kalish 1956; Hanson 1959; Weiss and Weissman 1992). That work used the wavelength spectrum as the stimulus parameter, which has difficulties that we discuss below. Hanson (1959) also looked at the effects of varying stimulus separation. Subjects were rewarded for pecking at a key lit by 550-nm light but not one lit by a longer-wavelength light. In tests, the pigeons were allowed to choose colors across the spectrum, and they did indeed respond most strongly to wavelengths below 550 nm. Also, reducing the spectral separation between the rewarded and unrewarded wavelengths, from 40 to 5 nm, increased the magnitude of the peak shift, although, as far as we can tell, the effect is not statistically significant.

*Color spaces and stimulus design.* For investigating high-level sensory processes, such as generalization or categorization, it is desirable to have a measure of the stimuli that is independent of the process itself. In color vision, one might use the wavelength spectrum (where stimuli can be well defined), but discriminability (or perceptual distance) is not simply proportional to wavelength difference, greatly complicating interpretation. Nor is it easy to use monochromatic light in naturalistic conditions. A better starting point is to use receptor excitations (table 1; fig. 2; Osorio et al. 1999b). Noise in receptor signals ultimately limits performance (Geisler 1989; Cohn 2004), and it is parsimonious to account for data in terms of a relatively low-level process. For example, if two different spectra give identical receptor signals (i.e., metameric colors), they will inevitably elicit identical behavioral responses. More important, judgments of similarity can be expected to scale with the magnitudes of differences in receptor signals. Only when judgments do not follow predictions of a receptor model can we draw inferences about (subsequent) neural processing.

It is relatively easy to estimate the expected receptor excitations to measured spectra, and color discrimination thresholds in birds and other animals are well predicted by a model that assumes that discrimination is based on chromatic signals (i.e., excluding “brightness”), whose performance is limited by photoreceptor noise (Vorobyev and Osorio 1998). Goldsmith and Butler (2003, 2005) provide experimental support and a useful discussion of this model with reference to bird vision. For stimulus design, the experiments reported here used the three-dimensional receptor space defined by the outputs of the short-, medium-, and long-wavelength-sensitive single cones (S, M, and L, respectively; table 1; fig. 2). For practical reasons, the very-short/UV-wavelength-sensitive cone was excluded by filtering the illumination (Osorio et al. 1999b). We have no reason to believe that inclusion of a UV signal would have

**Table 1:** Cone excitations and color loci of stimuli used in experiments 1–4 (fig. 2)

Experiment, color	S	M	L	D	$x$	$y$	$\Delta$
Experiment 1:							
$T_+$	.43	.61	.76	.66	.130	.007	
1	.41	.66	.69	.66	.112	.051	.047
2	.43	.63	.72	.66	.115	.025	.023
3	.39	.55	.74	.62	.148	-.007	.023
4	.38	.53	.76	.62	.161	-.020	.041
Experiment 2:							
$T_+$	.23	.21	.55	.34	.229	-.148	
1	.20	.17	.62	.35	.300	-.198	.087
2	.21	.19	.59	.34	.271	-.173	.049
$N_1$	.27	.26	.48	.33	.147	-.093	.099
$N_2$	.40	.33	.33	.33	-.050	-.030	.303
Experiment 3:							
$T_+$	.45	.62	.78	.68	.126	.002	
1	.56	.6	.77	.67	.077	-.028	.058
2	.51	.61	.77	.67	.097	-.013	.033
3	.40	.67	.79	.71	.148	.033	.038
4	.35	.70	.80	.72	.172	.055	.070
$N_1$	.47	.68	.73	.68	.098	.035	.043
$N_2$	.43	.57	.84	.71	.158	-.029	.044
Experiment 4:							
$T_+$	.38	.52	.64	.56	.122	.00	
$T_-$	.39	.54	.78	.63	.159	-.02	.045
$N$	.36	.56	.48	.51	.057	.082	.102

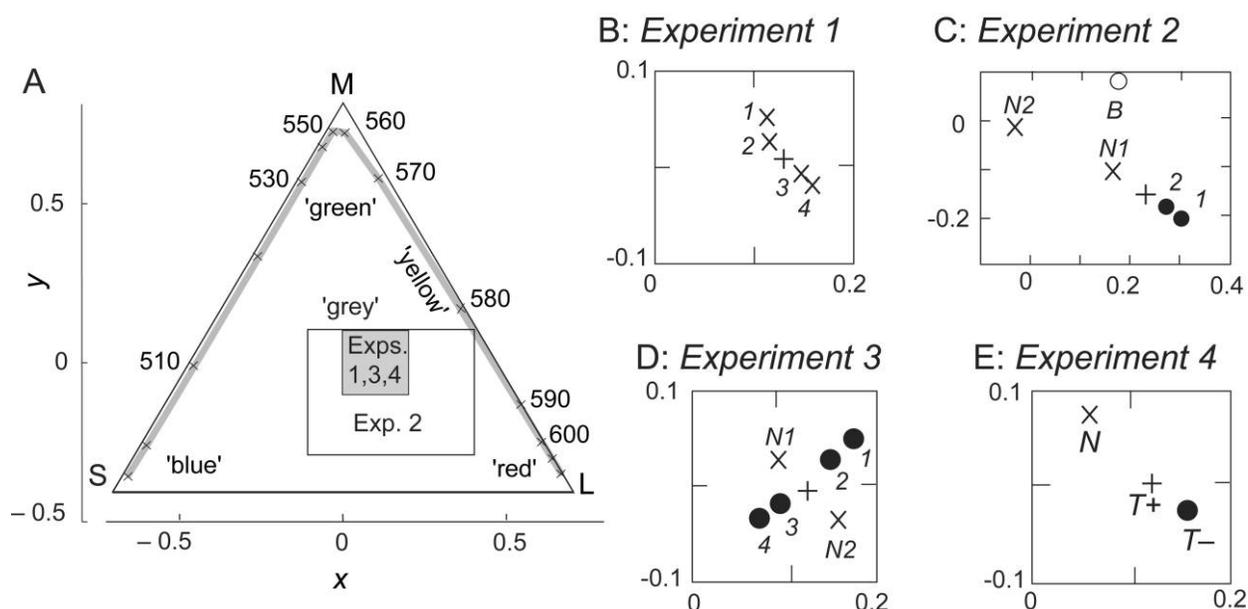
Note: Estimated quantal absorbances of chicken short-, medium-, and long-wavelength-sensitive single cones (S, M, and L, respectively) and double cones (D), relative to a barium sulfate standard. The illumination was filtered to exclude wavelengths visible to the chicken’s UV cones. Here  $x$  and  $y$  are chromaticity values derived from cone excitations according to equation (1), and  $\Delta$  is the Euclidean distance from  $T_+$ . For color differences along any given axis in the  $(x, y)$  chromaticity space, discriminability is proportional to distance. It is not possible to compare color discriminability on different axes without taking account of the relative sensitivities of the three cone mechanisms (Vorobyev and Osorio 1998).

any effect on the substantive findings reported here. The chromaticity coordinates (i.e., the approximate hue and saturation but not brightness) of a color are represented by the projection of its locus in the three-dimensional receptor space onto a unit plane (fig. 2A), given by

$$\begin{aligned}
 x &= (q_l - q_s)/\sqrt{2}, \\
 y &= [q_m - 0.5(q_l + q_s)]/\sqrt{3/2},
 \end{aligned}
 \tag{1}$$

where  $q_s$ ,  $q_m$ , and  $q_l$  are, respectively, the estimated quantum catches of the S, M, and L cones relative to a white adapting background (table 1; Osorio et al. 1999b; Kelber et al. 2003).

For colors of approximately equal intensity, a model of discrimination thresholds (Vorobyev and Osorio 1998; Kelber et al. 2003) makes the following useful predictions: (1) in the  $(x, y)$  chromaticity space, the shortest distance



**Figure 2:** Loci of stimuli used in a chromaticity diagram based on estimated responses of chicken short-, medium-, and long-wavelength-sensitive cone photoreceptors (table 1). The  $x$  and  $y$  values are calculated from cone responses, as described in “Methods.” Color versions of similar plots are shown in other publications (Osorio et al. 1999a, 1999b; Jones et al. 2001), and a color version of this figure is available in the online edition of the *American Naturalist*. *A*, Chromaticity diagram. The equilateral triangle defines the region given by cone signals  $\geq 0$ , and the gray line indicates the locus of monochromatic lights, with crosses plotted at 10-nm intervals from 490 to 620 nm. The color names “green,” “red,” etc., are those we (normally) give to spectra falling in that part of the chicken’s chromaticity diagram. The two boxes outline the areas of the diagram enlarged in *B–E*. *B*, Loci of stimuli used in experiment 1. The plus sign indicates the rewarded training color, and 1–4 denote the other experimental colors used as unrewarded training stimuli or novel stimuli in tests; see text for details. *C*, Loci of stimuli used in experiment 2. The plus sign indicates the rewarded training color, 1 and 2 denote unrewarded training colors, and *N1* and *N2* denote novel colors used to test generalization. *D*, Loci of stimuli used in experiment 3. The plus sign indicates the rewarded training color, 1–4 denote unrewarded training colors used in the four experimental conditions, and *N1* and *N2* denote novel colors used to test generalization. *E*, Loci of stimuli used in experiment 4.  $T_+$  indicates the rewarded training color,  $T_-$  indicates the unrewarded training color, and *N* indicates the novel color used in tests.

in “just noticeable difference” (jnd) steps (Wyszecki and Stiles 1982) between two color loci is a straight line; (2) two points separated by a given vector are separated by an equal number of jnd steps; and (3) more generally, the model predicts that any pair of colors separated by a given vector are equally discriminable. This last prediction may fail for large color differences but should hold for the fairly similar colors used here (fig. 2; sec. 8.4 of Wyszecki and Stiles 1982). Assuming that predictions 1–3 hold, we can define the following terms for colors in the  $(x, y)$  space: (1) An axis is a straight line; (2) color distance corresponds to the separation of any two points along a given axis;<sup>1</sup> and (3) if a vector  $\mathbf{a}$  separates two colors, the vector  $-\mathbf{a}$  is said to specify the opposite color difference.

This study is similar to Hanson’s (1959) in that birds

were trained to two colors, one of which ( $T_+$ ) yielded a food reward while the other ( $T_-$ ) did not. Tests then recorded the effects of varying the similarity of  $T_+$  and  $T_-$  on responses to novel colors. For experiments 1 and 2, the novel colors lay on the same axis in color space as the two training colors, and so these experiments measured peak shift. For experiment 3, the novel colors lay on an approximately perpendicular axis (fig. 2). Finally, experiment 4 looked at what is learned in a first encounter with a novel stimulus.

## Methods

### *Training and Testing*

Male poultry chicks (ISA-Brown) were kept, trained, and tested in pairs. The experiments followed standard procedures (Osorio et al. 1999b; Jones et al. 2001). Starting a week after hatching, chicks were trained in an arena (0.4 m  $\times$  0.3 m) for two 6-min sessions per day after 2 h

<sup>1</sup> Comparison of color distances along separate axes in this cone space would require assumptions about the relative sensitivities of different cone mechanisms (Vorobyev and Osorio 1998). This is not needed here because we are not comparing the magnitude of generalization on different directions in the receptor space.

of food deprivation. On the arena floor were eight conical paper containers approximately 28 mm in length with a 17-mm base. Four containers of the rewarded color ( $T_+$ ) contained standard chicks crumbs, which were replenished at 1-min intervals (i.e., five times per session); the other four were of the unrewarded color ( $T_-$ ) and were empty. Chicks were trained six times over 3 days and tested on the fourth day.

For experiments 1–3, chicks learned either a fine or a coarse color discrimination (one in which the difference between rewarded and unrewarded stimuli was large or small). In each experiment,  $T_+$  was fixed and  $T_-$  varied according to the treatment. For fine discrimination tasks,  $T_+$  and  $T_-$  were more alike than in the comparable coarse tasks (fig. 2; table 1). Although the chicks eventually made the fine discrimination easily (figs. 3–5), there was a possibility that they would not initially notice the difference. For this reason, the first three training sessions involved a coarse discrimination, and the more difficult task was used for the remaining four training sessions.

Tests were done after 2 h of food deprivation. In tests for experiments 1–3, chicks spent 2 min in the training arena with empty food containers. There was a choice of four stimulus colors:  $T_+$ ,  $T_-$  (fine), and two novel colors,  $N_1$ ,  $N_2$  (fig. 2B–2D; table 1). Tests lasted 2 min and scored the number of times each color was selected by the pair of chicks. We disregarded repeated pecks at the same stimulus or occasions where a chick copied its partner (Osorio et al. 1999a, 1999b). Such copying is readily recognized when one bird observes the other select a stimulus and immediately attempts to take it, normally within a second. The protocol for experiment 4 differed slightly from that for experiments 1–3 and is described later.

### Stimulus Design

As explained in the introduction to this article, colors were specified by their locations in a receptor space based on excitations of the chicken L, M, and S cones (table 1; fig. 2; Osorio et al. 1999b). Stimulus patterns were printed onto the food containers with an Epson inkjet printer. These patterns were tilings of  $6 \times 2$ -mm rectangles (see Osorio et al. 1999a, 1999b) colored at random with either the stimulus color or a background color (gray except in experiment 3), in a ratio of 3 : 7. Chromaticities of the colors were fixed, but their intensities varied at random, with a uniform distribution and a contrast of 0.3 (Osorio et al. 1999a, 1999b). Compared to a uniform color, these tessellated patterns have two advantages: first, they prevent chicks from using brightness rather than chromaticity, and second, they may improve the accuracy of color recognition by providing the background as a reference. The

floor and walls of the area were uniform gray of about the same intensity as the stimuli themselves.

### Statistics

For all experiments, one-tailed *t*-tests were used, either between groups or within groups as appropriate. For experiments 1–3, we used one-tailed tests because the direction of shift was predictable. Tests compared the preferences for the rewarded and novel colors. For experiments 1–3, the null hypothesis was that there was no effect of the training condition on this preference. For experiment 4, the null hypothesis was that there was no difference between the two tests.

## Results

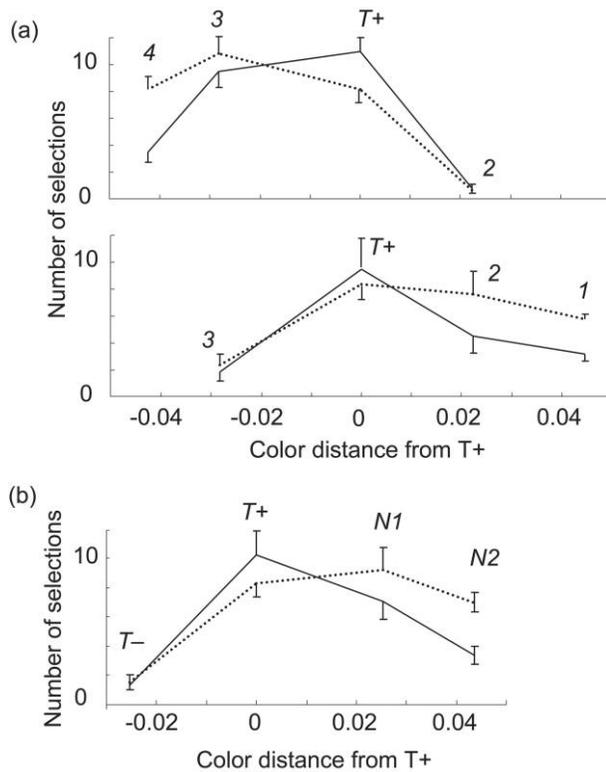
### Experiment 1

This experiment tested how the similarity of rewarded and unrewarded colors affected peak shift. The aim was to test models, such as those of Spence (1937) and Ghirlanda and Enquist (1999, 2003), that predict an increase in peak shift as the separation of the two differentially rewarded stimuli decreases. The experimental colors lay approximately on a line in the chicken ( $x, y$ ) color space (fig. 2A, 2B; table 1).

There were four treatment groups, in all of which the rewarded color  $T_+$  was fixed (orange to the human eye; fig. 2B; table 1), but each had a different  $T_-$ , named  $T_1$ – $T_4$ . For group 1,  $T_1$  was yellower than  $T_+$ ; group 2 used  $T_2$ , which was also yellower but half the color distance from  $T_+$ ; groups 3 and 4 used  $T_3$  and  $T_4$ , which were redder than  $T_+$  and differed from it by approximately equal and opposite color distances as  $T_-$  for groups 2 and 1, respectively. In summary, groups 1 and 4 made coarse discriminations, and groups 2 and 3 made fine discriminations. There were six pairs of chicks in each treatment.

After training over 4 days (see “Methods”), chicks were left without food for 2 h and then given a 2-min test. This scored the number of times each color was selected by the pair of chicks with the following stimuli: groups 1 and 2 had  $T_+$  and  $T_2$ – $T_4$ , while groups 3 and 4 had  $T_+$  and  $T_1$ – $T_3$ . Regardless of training conditions, the “fine”  $T_-$  was used in tests, that is,  $T_2$  for groups 1 and 2 and  $T_3$  for groups 3 and 4.

The results were as follows (fig. 3). First, the familiar  $T_-$  was less attractive than the novel color ( $T_2$  or  $T_3$ ) that differed equally from  $T_+$  ( $t = 7.432$ ,  $df = 22$ ,  $P < .001$ ), and second, there was no effect of the training condition, fine or coarse, on the preference for  $T_-$  ( $t = 0.212$ ,  $df = 21$ ,  $P = .417$ ). These observations showed that there was a peak shift away from  $T_-$  and also that the selectivity



**Figure 3:** Results of experiment 1, which compared color generalization of subjects trained to make coarse (dotted lines) or fine (solid lines) discriminations. The rewarded color ( $T_+$ ) was orange, and experimental stimuli varied mainly in hue (fig. 2A, 2B; table 1). The plot shows the mean number of selections  $\pm$  SEM of each color in a 2-min test versus the color distance from  $T_+$ . *a*, Top, data for subjects trained with “yellower”  $T_-$  colors: 1 (dotted line) and 2 (solid line); bottom, data for subjects trained with redder  $T_-$  colors: 4 (dotted line) and 3 (solid line). On these plots, colors “redder” than  $T_+$  are assigned negative values, and “yellower” colors are assigned positive values. *b*, Pooled data from *a* for subjects trained to make either coarse (dotted line) or fine (solid line) discriminations.  $T_-$  is assigned a negative value relative to  $T_+$ , and the novel colors are assigned positive values. Plotted color distances are the means from the two sets of experimental treatments illustrated in *a*.

in the direction of  $T_-$  from  $T_+$  was unaffected by the separation of the training colors,  $T_+$  and  $T_-$ . The key observation was that the peak shift (i.e., the preference for the novel colors) compared to  $T_+$  was smaller for groups 2 and 3, which had learned a fine discrimination, than for those in groups 1 and 4, which had learned a coarse discrimination ( $t = -4.192$ ,  $df = 21$ ,  $P < .001$ ). In summary, the experiment showed a peak shift and that this shift was smaller after learning a fine discrimination than after learning a coarse discrimination (cf. fig. 1).

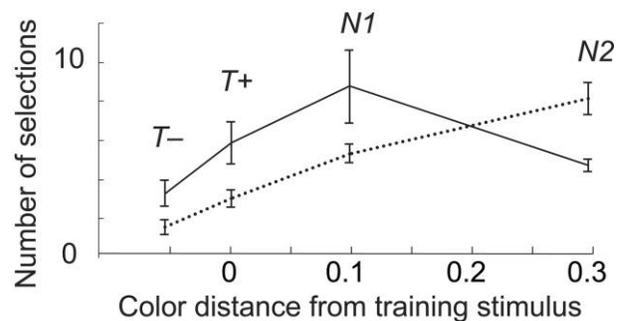
Because the results were not fully consistent with predictions of models of peak shift, we repeated the experiment and obtained very similar results. (1) The novel col-

ors were preferred to the unrewarded color that differed equally from  $T_+$  ( $t = 6.060$ ,  $df = 18$ ,  $P < .001$ ). (2) There was no effect of the training condition on the preference for the unrewarded color ( $t = -0.019$ ,  $df = 17$ ,  $P = .493$ ). (3) The preference for the novel colors compared to  $T_+$  was lower for groups 2 and 3, which had learned a fine discrimination, than for groups 1 and 4, which had learned a coarse discrimination ( $t = -3.161$ ,  $df = 17$ ,  $P = .002$ ).

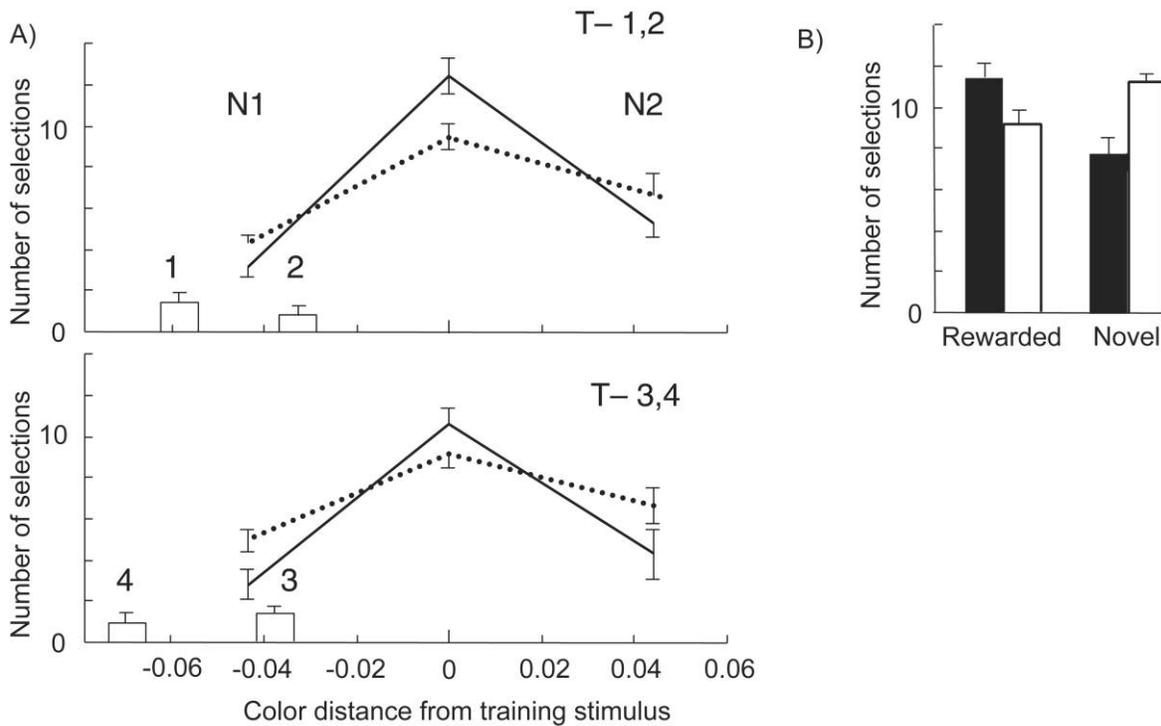
### Experiment 2

This experiment resembled experiment 1 but tested reddish colors (fig. 2C; table 1) that varied in saturation (i.e., difference from gray). Because a spontaneous preference for more saturated colors might occur without training, we tested peak shift toward gray. For group 1, the unrewarded color  $T_1$  was more saturated than  $T_+$ ; for group 2,  $T_2$  was also more saturated but half the color distance from  $T_+$ . The novel colors lay on the same axis:  $N_2$  was nearly achromatic, and  $N_1$  was intermediate between  $T_+$  and  $N_2$ . The background was not gray but yellow, so that the chromatic contrast between all stimulus colors and the background was approximately constant (fig. 2C). Group 1 was said to be making coarse and group 2 fine discriminations. Subjects were trained and tested on the same schedule as for experiment 1. The test colors were  $T_+$ ,  $T_2$ ,  $N_1$ , and  $N_2$  (fig. 2C; table 1). There were 10 pairs of chicks in each treatment.

The results resembled those for experiment 1 (fig. 4). (1) The novel colors were preferred to  $T_2$  ( $t = 7.252$ ,  $df = 17$ ,  $P < .001$ ), which is indicative of a peak shift. (2) The preference for the novel colors compared to  $T_+$  was



**Figure 4:** Results of experiment 2, which compared color generalization of subjects trained to make coarse (dotted line) or fine (solid line) discriminations. The rewarded color ( $T_+$ ) was red, and unrewarded training colors were more saturated and novel test colors less saturated than  $T_+$  (fig. 2A, 2C; table 1). The plot shows the mean number of selections  $\pm$  SEM of each color in a 2-min test versus the color distance from  $T_+$ .



**Figure 5:** Results of experiment 3, which compared color generalization of subjects trained to make coarse (dotted lines) or fine (solid lines) discriminations on one axis in color space and tested generalization on a different axis (fig. 2A, 2D; table 1). The rewarded color ( $T_+$ ) was orange. The plot shows the mean number of selections  $\pm$  SEM of each color in a 2-min test versus the color distance from  $T_+$ . A, Top, data for subjects trained with  $T_-$  colors 1 (dotted line) and 2 (solid line); bottom, data for subjects trained with  $T_-$  colors 4 (dotted line) and 3 (solid line). Histograms show the responses of each group to  $T_-$ . The position of the bar represents the color distance from  $T_+$ . However, because  $T_-$  colors and novel colors lie on different axes, their discriminabilities cannot be directly compared (see “Methods”). B, Histogram of pooled data for subjects trained to coarse (white bars) or fine (black bars) discrimination tasks, comparing responses to  $T_+$  and novel colors.

lower for group 2, which had learned a fine discrimination, than for group 1 ( $t = -4.519$ ,  $df = 16$ ,  $P < .001$ ). Thus, the effect of distance between  $T_+$  and  $T_-$  on peak shift reported for experiment 1 applied for a peak shift toward gray with a more highly saturated  $T_+$ .

### Experiment 3

This experiment resembled experiments 1 and 2, but now the training and test colors lay on different axes in color space (fig. 2A, 2D; table 1). There were four experimental groups, all of which were rewarded with the same  $T_+$  (fig. 2D). For groups 1 and 2, the unrewarded colors  $T_1$  and  $T_2$  differed from  $T_+$  in the same direction, with  $T_1$  offering a larger color difference than  $T_2$ . For groups 3 and 4, the unrewarded colors  $T_3$  and  $T_4$  were roughly equal and opposite color distances from  $T_+$  as colors 2 and 1, respectively. Training and testing were on the same schedule as for experiment 1. There were five or six pairs of chicks in each treatment.

Tests used two novel colors  $N_1$ ,  $N_2$  that lay on a different axis in the ( $x$ ,  $y$ ) color space from the training colors. Colors  $N_1$  and  $N_2$  differed from  $T_+$  by equal and opposite vectors (fig. 2D). Tests included colors  $T_+$ ,  $T_-$ ,  $N_1$ , and  $N_2$ . Color  $T_-$  was the unrewarded training color for the group in question.

In the generalization tests, the preference for  $T_+$  relative to  $N_1$  and  $N_2$  was higher after learning a fine discrimination than after learning a coarse discrimination (fig. 5). With pooled data for the two coarse groups (1 and 4) and the two fine groups (2 and 3; fig. 5B), the difference between coarse and fine treatments was significant ( $t = -3.362$ ,  $df = 11$ ,  $P = .003$ ).

### Experiment 4

Experiments 1 and 3 recorded initial preferences of the chicks for novel colors. Experiment 4 dealt with the stability of color generalization by asking how this initial exposure affects subsequent preferences. Chicks learned a

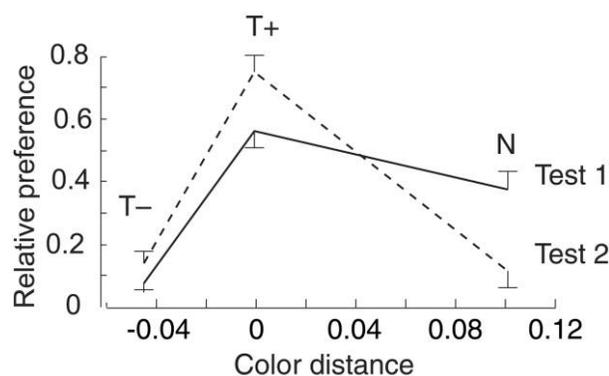
color discrimination task with a yellow  $T_+$  and a redder  $T_-$  (fig. 2A, 2E; table 1). Tests used  $T_+$ ,  $T_-$ , and a novel color  $N$ . In the initial test, preferences were given by the proportion of the first 10 pecks directed at each of the three colors (fig. 6). Chicks were then left in the test arena until they had selected both the rewarded and the novel colors eight times. We ensured equal numbers of selections simply by removing the relevant stimuli when they had been selected eight times. This gave them equal experience of the  $N$  and  $T_+$  in the unrewarded test condition. A second test 1 h later scored the first 10 selections by the pair of chicks.

In the initial test, there was a peak shift (fig. 6), with  $N$  preferred to  $T_-$  ( $t = -3.968$ ,  $df = 7$ ,  $P = .002$ ). In the second, test this shift was much reduced; the preference for  $N$  relative to  $T_+$  dropped significantly ( $t = -2.931$ ,  $df = 7$ ,  $P = .010$ ). Because the chicks had equal experience of  $N$  and  $T_+$  in the test condition (i.e., an average of four unrewarded selections per subject), this showed that they were able to modify generalization after a brief exposure to the novel stimuli.

### Discussion

Poultry chicks learn color quickly and accurately (Osorio et al. 1999a), which makes them good subjects for studying stimulus generalization. Experiments 1 and 2 show that after differential training there was a peak shift, as reported elsewhere for bird color vision (Hanson 1959; Weiss and Weissman 1992). However, this shift was smaller after a fine discrimination was learned than after a coarse discrimination task (figs. 2–4). To our knowledge, the only previous report of a comparable effect on peak shift is for human judgments of visual angle (Thomas et al. 1991). Experiment 3 did not test peak shift but instead tested the effects of differential training on generalization along roughly orthogonal axes in the receptor color space (figs. 2D, 5). The result was consistent with experiments 1 and 2 in that learning a fine discrimination gave a narrower generalization curve. Finally, experiment 4 showed that when tested “in extinction,” the relative preferences for novel and previously rewarded colors are unstable. Chicks modify stimulus preferences rapidly, without differential rewards on novel and familiar stimuli. The implication is that generalization is labile and is unlikely to be due to a process analogous to blurring or filtering of receptor outputs (Spence 1937; Blough 1975); we return to this subject in “Peak Shift within a Bayesian Framework.”

When Hanson (1959) trained pigeons to discriminate monochromatic lights, he found that peak shift along the optical spectrum increased as the separation of  $T_+$  and  $T_-$  fell (see also Pearce 1997; Blough 2001; Ghirlanda and Enquist 2003). Similar effects are reported on other stim-



**Figure 6:** Results of experiment 4, which tested how a brief exposure to an unrewarded novel stimulus modified generalization. The plot shows relative preference for  $T_+$ ,  $T_-$ , and the novel color  $N$  (fig. 2E; table 1) in two successive tests 1 h apart. Between test 1 (solid line) and test 2 (dashed line), there is a substantial decline in the preference for  $N$  relative to  $T_+$ .

ulus dimensions and from other species (Pearce 1997; Wills and Mackintosh 1998; Ghirlanda and Enquist 2003). The peak shifts observed in these studies are consistent with Spence’s (1937) model for stimulus generalization on a sensory continuum, which assumes that generalization curves to separate stimuli are summed (fig. 1), and variants on it (Ghirlanda and Enquist 1998, 1999, 2003; McLaren and Mackintosh 2002). The findings reported here and elsewhere (Jones et al. 2001) on color generalization do not follow the predictions of these types of model, raising two questions. First, why should empirical data differ between studies, and second, what does this say about underlying mechanisms?

Hanson’s (1959) experiment on pigeons was rather similar to experiments 1 and 2 but used monochromatic lights with a minimum color difference of 5 nm. By comparison, in experiment 1 the difference between  $T_+$  and  $T_-$  in the coarse training condition was equivalent to about 1 nm between the monochromatic lights in the 570–590-nm range (fig. 2). Also, Hanson interspersed probe tests in the sequence of training stimuli. This procedure is likely to reveal mainly the steady state response for a given experimental schedule. Experiments 1–3 tested only the initial response, and as experiment 4 shows, responses to novel stimuli are rapidly modified by experience. Finally, Hanson did not know the spectral sensitivities of bird photoreceptors and instead used the optical spectrum as the stimulus dimension on which to investigate generalization. Unfortunately, the spectral line need not equate to a line in receptor space (e.g., fig. 2A). Thus, Hanson used 550 nm as  $T_+$  and longer wavelengths as  $T_-$ , but 550 nm is at the apex of the chicken cone space (and would also be so for

pigeons), which means that the effects are not easy to interpret as peak shifts. This illustrates the importance of taking account of low-level perceptual mechanisms in interpreting processes such as generalization and categorization (Bennett et al. 1994; Wills and Mackintosh 1998; Blough 2001). By starting with receptor signals (table 1; Osorio et al. 1999b) and using relatively small stimulus differences, the experiments reported here should minimize distortions owing to errors in the mapping of physical stimuli onto perceptual spaces.

#### *What Kind of Model Is Required to Explain This Kind of Phenomenon? Perceptual Rulers*

Turning to models, a feature of Spence (1937) is the proposal that if an animal is familiar with multiple stimuli, it will have a separate generalization curve for each stimulus (e.g.,  $T_+$  and  $T_-$ ; fig. 1). Preferences for novel stimuli are predicted by the overlap between the  $T_+$  and  $T_-$  curves. Such models cannot predict the effects observed in experiments 1–3 if it is assumed (1) that stimulus generalization is done on a perceptual space that preserves the topology of the sensory (color) space defined by receptor signals (fig. 2A; table 1) and (2) that in this perceptual space generalization curves have single peaks or troughs located at the values of the familiar stimuli, for example,  $T_+$ ,  $T_-$ .

In fact, a much larger class of models cannot account for the observed data. One way of understanding what kind of models can and cannot account for this pattern of data is in terms of the concept of a perceptual ruler. This is a mechanism whose scale varies to match the currently relevant signal statistics. In most models of perception at threshold, the measurements are made relative to the noise level: if a signal change is significantly larger than the background noise, it will be detected; otherwise, it will not. In this case, the perceptual ruler is determined by the noise level, and, except for changes in test conditions (caused, for example, by variation in illumination intensity or background noise), this level will be fixed independent of the stimuli. Any model where the ruler that is used to measure stimuli is fixed cannot account for our data.

In contrast to more usual models of perception, if you ask whether a dog is large or small, you may expect different answers if the last two concepts asked about were “whale” and “elephant” than if they were “mouse” and “flea.” In this case, it is clear that the perceptual scale is determined by the other relevant categories. We believe that only models that have this characteristic of a flexible perceptual ruler can account for our pattern of data in experiments 1–3 (see Newell 1998 for another visual task that requires such an explanation).

#### *A Bayesian Account of the Peak Shift*

How then to provide a more formal model of these phenomena? Rather than framing the results in terms of classical animal learning theory (although it may be possible to do so), an alternative is to model phenomena as Bayesian inference. An increasing body of work on perception is Bayesian (Rao et al. 2002; Kersten and Yuille 2003). This has become the predominant method in machine learning, and neural networks are often viewed as implementing Bayesian principles (Kinouchi and Caticha 1992). This is not the first time that this probabilistic framework has been used to understand phenomena traditionally dealt with by learning theory (e.g., Dayan et al. 2000), but here we show that it naturally predicts peak shift and related phenomena (see also Cheng and Spetch 2002).

#### *Modeling Generalization as Bayesian Inference*

The Bayesian approach is based on three observations (Jaynes and Bretthorst 2003). (1) Many problems are characterized by uncertainty. While it is possible to deal with the possible using logic, greater insight is gained by modeling the plausible and probable. This is particularly true of human and animal behavior, where uncertainty and ambiguity are the rule. (2) Probability theory is the correct basis for dealing with uncertainty. Any system of inference that does not map to probabilistic inference will by necessity produce inconsistent inferences (Jaynes and Bretthorst 2003). (3) Finally, on a practical level, many problems associated with uncertain inference deal with mapping from some observed data to some inferred and unobserved state of the world. These problems are often difficult to even formalize. In contrast, if one knows the unobserved state, then constructing models of the probability of the observed is often far simpler. Bayes’s rule is simply a probabilistic identity that allows one to map from  $P(X|Y)$ —for example, the probability of  $X$ , given  $Y$ —to a probability of the form  $P(Y|X)$ . While in the abstract this may not seem useful, in practice it makes many complicated problems tractable.

Formally, Bayes’s rule consists of the equation  $P(X|\text{data}) = P(\text{data}|X)P(X)/P(\text{data})$ . The first term is known as the posterior, or the probability of  $X$ , given the data, and is usually the term of interest. Here it is the probability of a reward, given the observed stimulus color. The first term on the right-hand side of the equation,  $P(\text{data}|X)$ , is known as the likelihood function. This is the term that captures our intuitions about how the data were generated. It is often convenient to think of  $P(\text{data}|X)$  as a generative model that maps between the unobserved cause and the observed effect. The third term,  $P(X)$ , is known as the prior. For those unfamiliar with Bayesian

statistics, this term causes concern, but in many practical situations (such as the one to hand), it is irrelevant. In contrast, the last term,  $P(\text{data})$ , which resembles a simple normalization constant, often does much of the mathematical work. This term, sometimes known as the partition function (from statistical mechanics), implicitly requires summation over all possible causes of the data because  $P(\text{data}) = \sum P(\text{data}|X')P(X')$ . In our situation, where there can be only a few plausible causes, this term is simple to evaluate, and we will show that it can account for the observed peak shift.

#### *Peak Shift within a Bayesian Framework*

In this study, stimuli vary on a continuous dimension  $X$ , and after a limited number of examples, the animal generalizes to the infinite number of possible alternative values on this dimension. This requires a model that extrapolates from single examples to predictions for a range of values. The first step is to consider the possible inferences from a single example  $x_1$  of the concept  $S$ . On the intuition that similar stimuli are liable to have similar implications (such as rewards), the predicted benefit of responding to a novel stimulus  $x_{\text{novel}}$  can be represented by a generalization curve, which plots response strength against stimulus value (fig. 1). With one example  $x_1$ , these generalization curves are typically centered on  $x_1$  and assumed to be Gaussian (Spence 1937).

Viewing such curves simply as generalization curves is informative, but a more useful interpretation is within the general theory of inference under uncertainty. In this Bayesian framework, following others (Shepard 1987; Tenenbaum and Griffiths 2001), one very natural interpretation of a generalization curve is as a likelihood function, a model of the probability of observing a novel measurement  $x_{\text{novel}}$ , given that category  $S$  is present and having observed  $x_1$ , for example,  $P(x_{\text{novel}} \in S|x_1, S)$ . This function specifies the relationship between inputs and categories, as proposed by Spence (1937), but by identifying it with a statement about a probability (specifically, a likelihood) rather than simply a strength. We can later invoke Bayes's rule to perform any required inference in more complicated situations where there is ambiguity about the cause of a given sensory signal.

It has been argued that under this probabilistic interpretation, generalization curves should be approximately exponential, decaying as a function of the distance from the example (Shepard 1987; Tenenbaum and Griffiths 2001). Here we assume Gaussian generalization curves from a single example because, first, this more closely resembles Spence's (1937) original suggestion, second, we do not manipulate variables that would affect it (see Tenenbaum and Griffiths 2001 for discussion on what var-

iables may affect this generalization curve), and third, using exponential curves would change none of the conclusions.

Two additional mechanisms are then required to model behavioral generalization and discrimination: first, a mapping between the animal's internal model and observed pecking behavior, and second, a means to deal with multiple categories, some rewarded and others not. Bayesian inference estimates the probability of a reward being associated with any given color, but what is observed is pecking behavior. To allow comparison between the model and behavior, we need to propose some relationship between these two quantities. The most natural relationship is to assume that animals (approximately) obey Herrnstein's (1961) matching rule. This rule states that the frequency of selections is proportional to the estimated probability of reward. The matching rule generates suboptimal behavior in an unchanging world but is (potentially) more efficient in a changing world. Much experimental evidence indicates that foraging animals obey the matching rule (Kacelnik 1984; Stephens and Krebs 1986; Brito e Abreu and Kacelnik 1999).

This leads to the second complication: dealing with multiple possible causes of the observed color. In our experiments, rather than having only the single rewarded category  $S$ , we conceptually have four categories. Two are obvious, namely, trained-rewarded,  $S_+$ , and trained-unrewarded,  $S_-$ . Two other conditions are also relevant: the stimulus is neither  $S_+$  nor  $S_-$  but is associated with reward nonetheless ( $SO_+$ ), or the stimulus was neither  $S_+$  nor  $S_-$  but is not associated with reward ( $SO_-$ ). These four conditions cover all possible situations: a stimulus was either encountered before or not and is rewarded or not rewarded. Given these four potential stimuli ( $S_+$ ,  $S_-$ ,  $SO_+$ , and  $SO_-$ ), we can generate a model of the pecking behavior by assuming that the birds use the matching law to translate between generalization curves (models of the conditional probability of the novel stimulus, given a specific category) and the category probability (and hence the probability of reward). Bayes's rule converts models of probabilities that given colors belong to the four categories to a model of the probability of reward. Bayes's rule states

$$P(\text{reward}|x_{\text{novel}}) = P(\text{reward})P(x_{\text{novel}}|\text{reward})/P(x_{\text{novel}}). \quad (2)$$

To take the terms on the right-hand side in order,  $P(\text{reward})$  is the relatively uninteresting stimulus-independent term that represents the stimulus-independent chance of reward. This value depends on overall motivation, and it will affect only the total response rate. We therefore assume

that, in our experimental condition,  $P(\text{reward})$  is constant, and we ignore it.

The second term,  $P(x_{\text{novel}}|\text{reward})$ , is important. This specifies the probability of the novel stimulus, given that there was a reward. In this study, there were only two ways that a reward could occur, either as a result of the trained stimulus  $S_+$  or as a result of another rewarded (but unknown) category  $SO_+$ . Therefore,  $P(x_{\text{novel}}|\text{reward})$  is the sum of these two probabilities, weighted by the prior probability of their occurrence:

$$P(x_{\text{novel}}|\text{reward}) = P(x_{\text{novel}}|S_+)P(S_+) + P(x_{\text{novel}}|SO_+)P(SO_+).$$

In this equation,  $P(x_{\text{novel}}|S_+)$  is the sum of ordinary (Gaussian) generalization curves centered on the known training examples (fig. 1), and  $P(S_+)$  is the stimulus-independent prior probability of the known rewarded stimulus occurring determined by both its historical frequency of occurrence and the recency of previous stimulations. Stimuli that have occurred often and recently are more likely to occur again. Here  $P(x_{\text{novel}}|SO_+)$  is the generalization curve for all unknown but rewarded objects. Since it is unknown, we can assume that all stimuli are equally likely to belong to it, so this generalization curve is flat over the possible range of stimuli. Finally,  $P(SO_+)$  is the probability of an unknown but rewarded stimulus occurring.

However, in our experiment, there were no cases of rewards not associated with  $S_+$ , so  $P(SO_+) = 0$ , and hence the form of  $P(x_{\text{novel}}|SO_+)$  is irrelevant. Here we have assumed that all rewarded stimuli are rewarded with 100% probability. A full account of generalization would require an additional level of inference, but this is not required for the data here.

To summarize, the generalization from the positively rewarded stimulus resembles that of Spence (1937). The difference in the Bayesian account centers on how information from the unrewarded stimuli is incorporated. In Spence's account, information from the unrewarded stimuli is treated as having a negative reward, and phenomena such as peak shift occur because of the difference between the generalization curves of rewarded and unrewarded stimuli (fig. 1). In the Bayesian account, unrewarded stimuli influence generalization through the denominator of equation (2) and operate in a divisive manner: in Spence's scheme, reinforcing the unrewarded color subtracts from the probability of responding; in the Bayesian scheme, it increases the  $P(x_{\text{novel}})$  and hence scales the probability of responding. As is clear from equation (2), calculation of the probability of reward also requires the last term,  $P(x_{\text{novel}})$ . Since our four events are mutually exclusive and

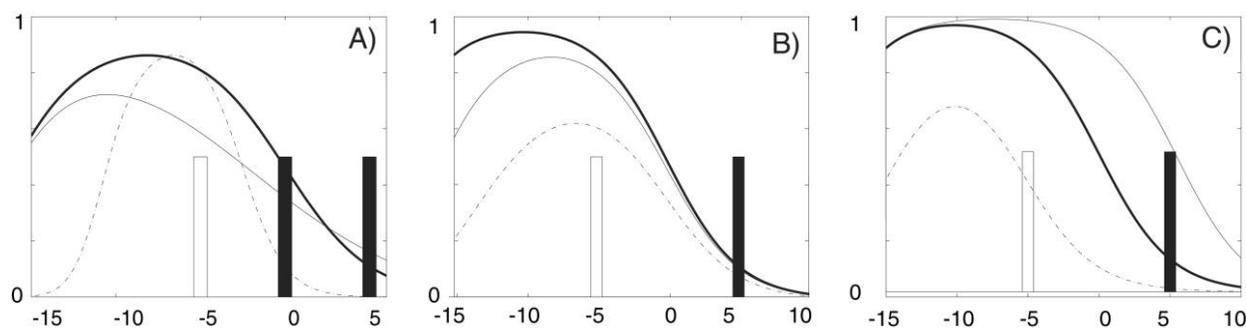
exhaustive,  $P(x_{\text{novel}})$  can be calculated by summing all four situations, weighted by their probability of occurrence:

$$P(x_{\text{novel}}) = P(x_{\text{novel}}|S_+)P(S_+) + P(x_{\text{novel}}|SO_+)P(SO_+) + P(x_{\text{novel}}|S_-)P(S_-) + P(x_{\text{novel}}|SO_-)P(SO_-).$$

It is this stage that represents the influence of unrewarded stimuli on generalization. As can be seen from figure 7B and equation (2), the probability of reward is decreased divisively by  $P(x_{\text{novel}}|SO_+)$ . Increasing the probability of an unrewarded stimulus increases the denominator in equation (2),  $P(x_{\text{novel}})$ , while not increasing the numerator. Probability theory states that the probabilities of an exhaustive set of possibilities have to sum to 1. This means that if  $x_{\text{novel}}$  could potentially be the result of either a rewarded stimulus or an unrewarded one, these two stimuli compete to represent the stimulus, and therefore the probability of  $x_{\text{novel}}$  being the rewarded stimulus decreases. This can be seen graphically in figure 7B. Of the four terms in  $P(x_{\text{novel}})$ , the first two are simply the previously encountered reward-based terms. The third and fourth terms represent the probability that  $x_{\text{novel}}$  was a previously encountered unrewarded stimulus (with a Gaussian generalization curve centered on the stimulus value) or a novel but unrewarded stimulus (with a flat generalization curve). The term  $P(SO_-)$ , the probability of the stimulus being unknown and unrewarded, cannot be ignored, since because of the way the experiment is conducted, the animals encounter previously rewarded stimuli that have now been eaten and hence are unrewarded. The simplest way to model this is a catchall unspecific, unrewarded category.

#### *Implications of a Bayesian Account of the Peak Shift*

As has been shown, interpreting the generalization curves as likelihood functions leads to an extremely general Bayesian account of peak shift. Before considering the actual experiments, a few words are in order on the differences between our account and that of Spence (1937) and related models. First, in the latter accounts, the unrewarded stimulus must be regarded as negatively rewarded, and its effect is subtractive (fig. 1). This presents a minor difficulty in specifying the perceived value of unrewarded stimuli. By comparison, in a Bayesian account, the effects of the unrewarded stimuli are competitive and divisive. The effects on generalization are well specified, being proportional to the chance that the unrewarded stimulus caused the observation. Second, the theory predicts that manipulations that affect the stimulus-independent probability of  $T_+$  and  $T_-$  (such as biased presentation rates) will affect gener-



**Figure 7:** Predictions of a Bayesian model of generalization. *A*, Need for adaptive generalization curves. The thin solid line shows a typical peak shift that we would expect when animals are trained with rewards at  $-5$  and not rewarded at  $+5$ . This curve was generated using the proposed Bayesian model, but similar curves could be generated by models based on Spence's (1937) principles. The thick solid line shows the generalization expected if the unrewarded stimuli were not at  $+5$  but at  $0$ , a condition corresponding to our fine condition. As can be seen, if generalization curves from the positive and negative examples have the same extent, then peak shift should increase. This is not what we found. If, instead, the extent of generalization from an individual example is moderated by the nearest counterexample, the observed decrease of peak shift is predicted (*dashed line*). *B*, Effects of time can change generalization curves. The thick solid line shows a generalization curve due to training on a positively rewarded stimulus at  $-5$  and an unrewarded stimulus at  $+5$ . If a long interval is inserted before testing, then the probabilities  $P_+$  and  $P_-$  will decay. This will cause the changes in generalization shown by the thin solid line (a fivefold decay) and the dashed line (a further fivefold decay). *C*, Effects of the relative frequency of the rewarded and unrewarded stimuli. One simple manipulation is to alter the relative number of rewarded and unrewarded stimuli, and this plot shows the effect of this manipulation for one set of parameters. We illustrate effects of equal numbers of rewarded and unrewarded stimuli (*thick solid line*), five times as many unrewarded as rewarded stimuli (*dashed line*), and five times as many rewarded as unrewarded stimuli (*thin solid line*). The main effect is to shift the generalization curves.

alization, essentially shifting curves to the left and right. Notice also that in this model, if only  $T_+$  stimuli are ever presented, this will result in a flat generalization as a function of  $x_{\text{novel}}$ : even if a stimulus is very unlikely to have been generated by the observed positively rewarded class, given no evidence for a negatively rewarded class, it is still more likely to have been generated by this class.

Finally, the effects of the two “unobserved” classes are of interest. If  $P(S0_-)$ , the chance that a stimulus was generated by an observed and unrewarded class, is  $0$ , then open-ended generalization will occur (see also Ghirlanda and Enquist 1999, 2003). If, in contrast,  $P(S0_+)$  is large, then the generalization will form a peak. In summary, peak shift phenomena can be simply mapped to the mechanics of Bayesian probabilistic inference. Doing so generates many of the same predictions as Spence's (1937) account, but (1) the phenomenon is placed within the general framework of probabilistic inference that is used to explain many other aspects of perception and learning, (2) arbitrary parameters (such as the level of negative reward for unrewarded stimuli) become principled, (3) a new family of interpretable and principled parameters with specific effect on generalization is produced, and (4) there is no need to assume that performance is limited by neural mechanisms beyond the receptors (Enquist and Arak 1993, 1998). These effects are shown in figure 7. We therefore argue that this framework is worth developing as a basis for understanding sensory generalization.

#### *Variable Widths of Generalization Curves*

Finally, we return to the key experimental observation here, namely, the effect of separation of  $T_+$  and  $T_-$ . Are generalization curves to single stimuli fixed or plastic and perhaps determined by the inferred stimulus variability? Models that use fixed generalization curves may not make optimal use of available information, because fixed-width curves disregard any information about the stimulus variability. When making inferences about objects that show great variability (the size of dogs or the color of apples), common sense suggests that one may extrapolate or generalize from a single example over a larger range than when one encounters an example from a category that shows less variability (size of cats or the color of ripe oranges).

The practical advantages of taking account of stimulus variability are well known in machine learning (Hastie et al. 2001). A common and powerful technique for both classification and density estimation are so-called kernel methods. Conceptually, these methods resemble Spence's (1937) account in that a model of the data is constructed by superimposing multiple generalization functions (known in the literature as “kernels”). What is relevant here is that these methods work well but only if the scale of the generalization function is well matched to the data; otherwise, performance can be poor. As is often the case, the reason for this can be understood from the extreme cases. If the generalization from a given example extended

only infinitesimally, then there would essentially be no generalization: all novel stimuli would be classed as different from all others. In the other case, if generalization extended from a given example infinitely, all novel stimuli would be classed as equally similar to a novel one. The method usually used in machine learning to choose an appropriate generalization curve width is cross validation (the width is chosen by keeping some percentage of the data back and testing how well a particular generalization curve width allows prediction of these “unseen” data). Cross validation in its simplest form is unrealistic as a biological model. Despite this, because practical performance for realistic problems has been found to be so highly dependent on adapting the generalization widths to the data, it is of interest to explore whether such a mechanism is operating in color category learning.

Returning to the experimental results reported here, for two stimuli separated by a distance  $x$ , the simple summation and normalization of the generalization curves predicts the observed peak shift. The critical manipulation is when the distance between the two stimuli is increased or decreased. For fixed generalization curves, there is only one possible prediction: the stimulus of maximum re-

sponse gets shifted away from the negative response. This contrasts with a system that adapts its generalization curves to one stimulus in response to estimated stimulus variability. In the first situation, the peak shift is the same. The differences are observed when one changes the distance between  $T_+$  and  $T_-$ . If the generalization from  $T_+$  is affected by the inferred variability (estimated from the range spanned by  $T_+$  and  $T_-$ ) and scales with it in some way, then decreasing the distance between  $T_+$  and  $T_-$  will decrease the degree of generalization from  $T_+$ . Hence, we can observe that decreasing the distance between  $T_+$  and  $T_-$  can shift the peak associated with  $T_+$  toward  $T_-$ .

One simple way of understanding this effect is that the range of stimuli scales the perceptual space. For this reason, we call this effect a “perceptual ruler,” a concept that can be understood independent of the underlying algorithm. Our data suggest that such an effect is used by chicks to help them take advantage of information about stimulus variability.

#### Acknowledgment

This work was supported by the Biotechnology and Biological Sciences Research Council.

## APPENDIX

### The Bayesian Model

The body of the article describes a Bayesian model, and the predictions of this model are shown in figure 7 for a number of different situations. Although it possible to reconstruct the model from text, this appendix contains an explicit description of the model and gives the parameters that were used to generate the three panels of figure 7.

The goal is to make a prediction of the probability of reward as a function of the stimulus value, for example,  $P(\text{reward}|x_{\text{novel}})$ . Assuming  $P(\text{reward}) = \text{constant}$ , this results in the following equations:

$$P(\text{reward}|x_{\text{novel}}) = P(x_{\text{novel}}|\text{reward})P(\text{reward})/P(x_{\text{novel}}), \quad (\text{Bayes's rule}),$$

$$P(x_{\text{novel}}|\text{reward}) = P(x_{\text{novel}}|S_+)P(S_+) + P(x_{\text{novel}}|SO_+)P(SO_+),$$

$$P(x_{\text{novel}}) = P(x_{\text{novel}}|S_+)P(S_+) + P(x_{\text{novel}}|SO_+)P(SO_+) + P(x_{\text{novel}}|S_-)P(S_-) + P(x_{\text{novel}}|SO_-)P(SO_-).$$

Assuming Gaussian-shaped generalization functions around both the positively and negatively rewarded stimuli, we have

$$P(x_{\text{novel}}|S_+) = \exp [-(x_{\text{novel}} - \text{loc\_positive})^2/(2 \times \text{width}^2)]$$

and

$$P(x_{\text{novel}}|S_-) = \exp [-(x_{\text{novel}} - \text{loc\_negative})^2/(2 \times \text{width}^2)],$$

where in figure 7 the positive stimuli were always were located at  $-5$  (i.e.,  $\text{loc\_positive} = -5$ ) and the negative stimuli had values of  $+5$  or  $0$ , as indicated in the text ( $\text{loc\_negative} = 5$  or  $0$ ).

The two unobserved stimuli classes have an even simpler form:

$$P(x_{\text{novel}}|S0_+) = 1,$$

$$P(x_{\text{novel}}|S0_-) = 1.$$

This leaves the four prior probability parameters,  $P(S_+)$ ,  $P(S0_+)$ ,  $P(S_-)$ , and  $P(S0_-)$ , together the width of the generalization curve.

### Literature Cited

- Baddeley, R., D. Osorio, and C. D. Jones. 2001. Colour generalization by domestic chicks. *Behavioral and Brain Sciences* 24:654–655.
- Bennett, A. T. D., I. C. Cuthill, and K. Norris. 1994. Sexual selection and the mismeasure of color. *American Naturalist* 144:848–860.
- Blough, D. S. 1975. Steady state data and a quantitative model of generalization and discrimination. *Journal of Experimental Psychology Animal Behavior Processes* 1:3–21.
- . 2001. The perception of similarity. In R. G. Cook, ed. *Avian visual cognition*. <http://www.pigeon.psy.tufts.edu/avc/toc.htm>.
- Brainard, D. H., and W. T. Freeman. 1997. Bayesian color constancy. *Journal of the Optical Society of America A* 14:1393–1411.
- Brito e Abreu, F., and A. Kacelnik. 1999. Energy budgets and risk-sensitive foraging in starlings. *Behavioral Ecology* 10:338–345.
- Cheng, K. 2002. Generalisation: mechanistic and functional explanations. *Animal Cognition* 5:33–40.
- Cheng, K., and M. L. Spetch. 2002. Spatial generalization and peak shift in humans. *Learning and Motivation* 33:358–389.
- Cohn, T. E. 2004. Detection and sampling. Pages 811–824 in L. M. Chalupa and J. S. Werner, eds. *The visual neurosciences*. MIT Press, Cambridge, MA.
- Dayan, P., S. Kakade, and P. R. Montague. 2000. Learning and selective attention. *Nature Neuroscience* 3(suppl.):1218–1223.
- Enquist, M., and A. Arak. 1993. Selection of exaggerated male traits by female aesthetic senses. *Nature* 361:446–448.
- . 1998. Neural representation and the evolution of signal form. Pages 21–87 in R. Dukas, ed. *Cognitive ethology*. University of Chicago Press, Chicago.
- Enquist, M., and R. Johnstone. 1997. Generalization and the evolution of symmetry preferences. *Proceedings of the Royal Society B: Biological Sciences* 264:1345–1348.
- Geisler, W. S. 1989. Sequential ideal-observer analysis of visual discrimination. *Psychological Review* 96:267–314.
- Ghirlanda, S., and M. Enquist. 1998. Artificial neural networks as models of stimulus control. *Animal Behaviour* 56:1383–1389.
- . 1999. The geometry of stimulus control. *Animal Behaviour* 58:695–706.
- . 2003. A century of generalization. *Animal Behaviour* 66:15–36.
- Goldsmith, T. H., and B. K. Butler. 2003. The roles of receptor noise and cone oil droplets in the photopic spectral sensitivity of the budgerigar, *Melopsittacus undulatus*. *Journal of Comparative Physiology A* 189:135–142.
- . 2005. Color vision of the budgerigar (*Melopsittacus undulatus*): hue matches, tetrachromacy, and intensity discrimination. *Journal of Comparative Physiology A* 191:933–951.
- Guttman, N., and H. I. Kalish. 1956. Discriminability and stimulus generalization. *Journal of Experimental Psychology* 51:79–88.
- Hanson, H. M. 1959. Effects of discrimination training on stimulus generalization. *Journal of Experimental Psychology* 58:321–334.
- Hastie, T., R. Tibshirani, and J. Friedman. 2001. *Elements of statistical learning: data mining*. Springer, Berlin.
- Herrnstein, R. J. 1961. Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of Experimental Analysis of Behavior* 4:267–272.
- Jaynes, E. T., and G. L. Bretthorst. 2003. *Probability theory: the logic of science*. Cambridge University Press, Cambridge.
- Jones, C. D., D. Osorio, and R. J. Baddeley. 2001. Color categorization by domestic chicks. *Proceedings of the Royal Society B: Biological Sciences* 268:2077–2084.
- Kacelnik, A. 1984. Central place foraging in starlings (*Sturnus vulgaris*). I. Patch residence time. *Journal of Animal Ecology* 53:283–299.
- Kelber, A., M. Vorobyev, and D. Osorio. 2003. Animal colour vision: behavioural tests and physiological concepts. *Biological Reviews* 78:81–118.
- Kersten, D., and A. Yuille. 2003. Bayesian models of object perception. *Current Opinion in Neurobiology* 13:150–158.
- Kinouchi, O., and N. Caticha. 1992. Optimal generalization in perceptrons. *Journal of Physics A* 25:6243–6250.
- Lynn, S. K., J. Cnaani, and D. R. Papaj. 2005. Peak shift discrimination learning as a mechanism of signal evolution. *Evolution* 59:1300–1305.
- McLaren, I. P. L., and N. J. Mackintosh. 2002. Associative learning and elemental representation. II. Generalization and discrimination. *Animal Learning and Behaviour* 30:177–200.
- Newell, F. 1998. Stimulus context and view dependence in object recognition. *Perception* 27:47–68.
- Osorio, D., C. D. Jones, and M. Vorobyev. 1999a. Accurate memory for colour but not pattern contrast in chicks. *Current Biology* 9:199–202.
- Osorio, D., M. Vorobyev, and C. D. Jones. 1999b. Colour vision in domestic chicks. *Journal of Experimental Biology* 202:2951–2959.
- Pearce, J. M. 1997. *Animal learning and cognition*. Erlbaum, Hove.
- Rao, P. N., B. A. Olshausen, and M. S. Lewicki. 2002. *Probabilistic models of the brain: perception and neural function*. MIT Press, Cambridge, MA.
- Ryan, M. J. 1990. Sexual selection, sensory systems, and sensory exploitation. *Oxford Surveys in Evolutionary Biology* 7:156–195.
- Shepard, R. N. 1987. Towards a universal law of generalization for psychological science. *Science* 237:1317–1323.
- Spence, K. W. 1937. The differential response in animals to stimuli varying within a single dimension. *Psychological Review* 44:430–444.
- Stephens, D. W., and J. R. Krebs. 1986. *Foraging theory*. Princeton University Press, Princeton, NJ.
- Tenenbaum, J. B., and T. L. Griffiths. 2001. Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences* 24:629–641.
- Thomas, D. R., K. Mood, S. Morrison, and E. Wiertelak. 1991. Peak shift revisited: a test of alternative interpretations. *Journal of Experimental Psychology Animal Behavior Processes* 17:130–140.

- Vorobyev, M., and D. Osorio. 1998. Receptor noise as a determinant of colour thresholds. *Proceedings of the Royal Society B: Biological Sciences* 265:351–358.
- Weary, D. M., T. C. Guilford, and R. G. Weisman. 1993. A product of discriminative learning may lead to female preferences for elaborate males. *Evolution* 47:333–336.
- Weiss, S. J., and R. D. Weissman. 1992. Generalization peak shift for autoshaped and operant key pecks. *Journal of Experimental Analysis of Behavior* 57:127–143.
- Wills, S. J., and N. J. Mackintosh. 1998. Peak shift on an artificial dimension. *Quarterly Journal of Experimental Psychology B* 51: 1–31.
- Wyszecki, G., and W. S. Stiles. 1982. *Color science: concepts and methods, quantitative data and formulae*. 2nd ed. Wiley, New York.
- Yachi, S., and M. Higashi. 1998. The evolution of warning signals. *Nature* 394:882–884.

Symposium Editors: Andrew T. D. Bennett and Marc Théry