Unsupervised category learning with integral-dimension stimuli

Shawn W. Ell [a] , F. Gregory Ashby [b] & Steven Hutchinson [a]

[a] Psychology Department, University of Maine, Orono, ME, USA

[b] Department of Psychological and Brain Sciences, University of California,
Santa Barbara, Santa Barbara, CA, USA

PLEASE SCROLL DOWN FOR ARTICLE

Unsupervised category learning with integral-dimension stimuli

Shawn W. Ell[1], F. Gregory Ashby[2], and Steven Hutchinson[1]

[1]Psychology Department, University of Maine, Orono, ME, USA
[2]Department of Psychological and Brain Sciences, University of California, Santa Barbara, Santa Barbara, CA, USA

Despite the recent surge in research on unsupervised category learning, the majority of studies have focused on unconstrained tasks in which no instructions are provided about the underlying category structure. Relatively little research has focused on constrained tasks in which the goal is to learn predefined stimulus clusters in the absence of feedback. The few studies that have addressed this issue have focused almost exclusively on stimuli for which it is relatively easy to attend selectively to the component dimensions (i.e., separable dimensions). In the present study, we investigated the ability of participants to learn categories constructed from stimuli for which it is difficult, if not impossible, to attend selectively to the component dimensions (i.e., integral dimensions). The experiments demonstrate that individuals are capable of learning categories constructed from the integral dimensions of brightness and saturation, but this ability is generally limited to category structures requiring selective attention to brightness. As might be expected with integral dimensions, participants were often able to integrate brightness and saturation information in the absence of feedback—an ability not observed in previous studies with separable dimensions. Even so, there was a bias to weight brightness more heavily than saturation in the categorization process, suggesting a weak form of selective attention to brightness. These data present an important challenge for the development of models of unsupervised category learning.

Keywords: Categorization; Implicit/explicit memory; Feedback-based learning.

There has been a recent surge in research on unsupervised category learning—that is, the ability to learn categories in the absence of corrective feedback. Studies focusing on unsupervised learning provide an important complement to the studies of supervised learning that have dominated the field as, arguably, much everyday learning occurs in the absence of trial-by-trial feedback. Given the ubiquity of unsupervised category learning, it is not that surprising that individuals can spontaneously construct categories in the absence of feedback (e.g., Medin, Wattenmaker, & Hampson, 1987).

Much of the research on unsupervised category learning, however, has focused on unconstrained tasks where participants have no knowledge that there is an optimal categorization strategy, if one exists at all (Ahn & Medin, 1992; Billman & Knutson, 1996; Clapper & Bower, 1994; Colreavy & Lewandowsky, 2008; Diaz & Ross, 2006; Handel & Imai, 1972; Love, 2002; Medin et al., 1987; Milton, Longmore, & Wills, 2008; Milton & Wills, 2004; Pothos & Chater, 2005; Pothos & Close, 2008; Regehr & Brooks, 1995). In unconstrained tasks, the primary focus is on how participants prefer to construct categories. For instance, in the typical free-sorting task, participants are presented with a number of stimuli (either simultaneously or sequentially) and are asked to place the stimuli into a number of categories in any way they like. The participants are not informed that there is an underlying category structure (if one exists). Therefore, using a common performance measure such as accuracy is problematic because there is no objectively correct response.

Although unconstrained tasks have been important for understanding how characteristics of the stimuli and task influence categorization strategy, it is also important to investigate unsupervised category learning in more constrained tasks in which participants are attempting to learn the optimal categorization strategy (Ashby, Queller, & Berretty, 1999; Zeithamova & Maddox, 2009).[1] In constrained tasks, the primary focus is on what types of category structure individuals are capable of learning. With the exception of feedback, the methodology in constrained tasks closely parallels most supervised category-learning paradigms as participants know that their goal is to learn an underlying category structure. Therefore, accuracy is an appropriate performance measure because there is an objectively correct response.

To our knowledge, all constrained tasks and the majority of unconstrained tasks have used stimuli for which it is relatively easy to attend selectively to the component dimensions (i.e., separable dimensions). Two dimensions are said to be separable if it is possible to attend to one dimension and ignore the other (e.g., hue and shape; Garner, 1974; Imai & Garner, 1965). Conversely, two dimensions are said to be integral if it is impossible to attend to one and ignore irrelevant variations in the other (e.g., brightness and saturation; Garner & Felfoldy, 1970; Torgerson, 1958).[2]

Under supervised conditions, participants can readily learn categories constructed from integral dimensions (Grau & Kemler-Nelson, 1988; McKinley & Nosofsky, 1996; Mounts & Melara, 1995; Nosofsky & Palmeri, 1996; Shepard & Chang, 1963). There is, however, an extensive literature documenting differences in the processing of separable and integral dimensions (Foard & Kemler-Nelson, 1984; Lockhead, 1972). In the context of supervised category learning, selective attention mechanisms operate less efficiently when learning categories constructed from integral, rather than separable, dimensions (Maddox, 2001; Maddox & Dodd, 2003; Nosofsky, 1986, 1987).

A number of studies using unconstrained tasks have shown that the preferred decision strategy varies as a function of whether the stimuli are constructed from separable or integral dimensions (Handel & Imai, 1972; Handel, Imai, & Spottswood, 1980; Imai & Garner, 1965).[3] For

---

[1] See a recent paper by Pothos and colleagues (Pothos, Edwards, & Perlman, 2011) for a related distinction between constrained and unconstrained unsupervised category learning tasks.

[2] More specifically, the observation of (a) a Euclidean metric in multidimensional scaling, (b) interference when the stimuli vary orthogonally, and (c) a redundancy gain when the stimuli are correlated in the speeded classification paradigm are often considered as evidence for dimensional integrality (Garner, 1974), but see Ashby and Maddox (1994).

[3] Colour-naming tasks are a special case of the free-sorting paradigm and have frequently been used to assess people's ability to identify colour categories in a variety of colour spaces (Boynton & Olson, 1987; Sturges & Whitfield, 1995). Data from this paradigm indicate that participants are quite capable of sorting colour stimuli into various categories without feedback. While these studies have explored how variations along the colour dimensions affect the preferred classifications of participants, they have been primarily interested in variations along hue and do not provide strong predictions for stimuli varying along the integral dimensions of brightness and saturation.

example, Handel and colleagues (Handel & Imai, 1972; Handel et al., 1980) compared the separable dimensions of shape and colour with the integral dimensions of brightness and saturation. Separable dimension stimuli were sorted using a one-dimensional strategy, whereas integral dimension stimuli were sorted using a similarity-based strategy. A similar bias to use one-dimensional strategies with separable dimensions has been reported with other unconstrained tasks (e.g., Colreavy & Lewandowsky, 2008; Medin et al., 1987). Importantly, however, research from unconstrained tasks suggests that the bias to use one-dimensional strategies is critically dependent upon the particular category structures (Ahn & Medin, 1992; Pothos & Chater, 2005; Pothos & Close, 2008), spatial configuration of the stimuli (Milton & Wills, 2004), and experimental procedure (Milton et al., 2008; Regehr & Brooks, 1995). For example, simply informing participants of the number of categories has been argued to instil a one-dimensional bias (e.g., Murphy, 2002).

Although the separable-integral distinction is often described as being discrete, such a characterization is likely to be an oversimplification (Ashby & Townsend, 1986). Studies using constrained tasks with stimuli that are strongly separable have consistently demonstrated a bias to use one-dimensional strategies (Ashby et al., 1999; Zeithamova & Maddox, 2009). There are at least two studies using constrained tasks with stimuli that fall somewhere in the middle of the separable-integral continuum (10 × 10 grids of randomly distributed light and dark squares, Fried & Holyoak, 1984; lines connecting nine randomly located dots, Homa & Cultice, 1984). Although learning was evident in both studies, there are several limitations with respect to the question of what individuals are capable of learning under unsupervised conditions on constrained tasks. For instance, the stimuli in the Fried and Holyoak (1984) study varied on 100 physical dimensions, while the Homa and Cultice (1984) stimuli varied along 18 physical dimensions. The dimensionality of the psychological representation of these stimuli is not known, and it is likely that there is no straightforward mapping between the psychological and physical dimensions (Shin

& Nosofsky, 1992). Without knowing the psychological representation of the stimuli, it is impossible to obtain an accurate estimate of the decision strategy that participants were using to perform the task.

To summarize, with separable stimulus dimensions, unsupervised category learning is possible, and, in some cases, there is a bias to use one-dimensional strategies. With integral dimensions, the picture is more complicated. On unconstrained tasks, individuals do not demonstrate a strong preference for one-dimensional strategies. On constrained tasks using stimuli that are likely to have some degree of integrality, unsupervised category learning is possible. We know, however, very little about what types of strategies individuals are capable of learning under unsupervised conditions when the categories are constructed from integral dimensions and whether the bias to use one-dimensional strategies that has been demonstrated on constrained tasks with separable dimensions extends to integral dimensions.

We investigate these questions in the present experiment using a constrained task with stimuli constructed from the integral dimensions of brightness and saturation defined in the Munsell colour system (Figure 1). The structure of the Munsell colour system is such that variations along the value dimension correspond to changes in brightness, whereas differences in chroma reflect changes in saturation (Munsell, 1915). For simplicity, the physical dimensions of value and chroma will be referred to by the perceptual labels of brightness and saturation, respectively. Two one-dimensional (vertical and horizontal conditions in the top panels) and two diagonal conditions (positive and negative conditions in the bottom panels) were constructed by randomly sampling from a bivariate uniform distribution defined on the brightness and saturation dimensions. Each of the category structures differs only in the orientation of the optimal decision strategy. As the name implies, to learn the one-dimensional structures, participants should attend to the relevant stimulus dimension (while ignoring the other, irrelevant dimension). To learn the diagonal structures, participants should integrate information from the brightness and saturation dimensions.
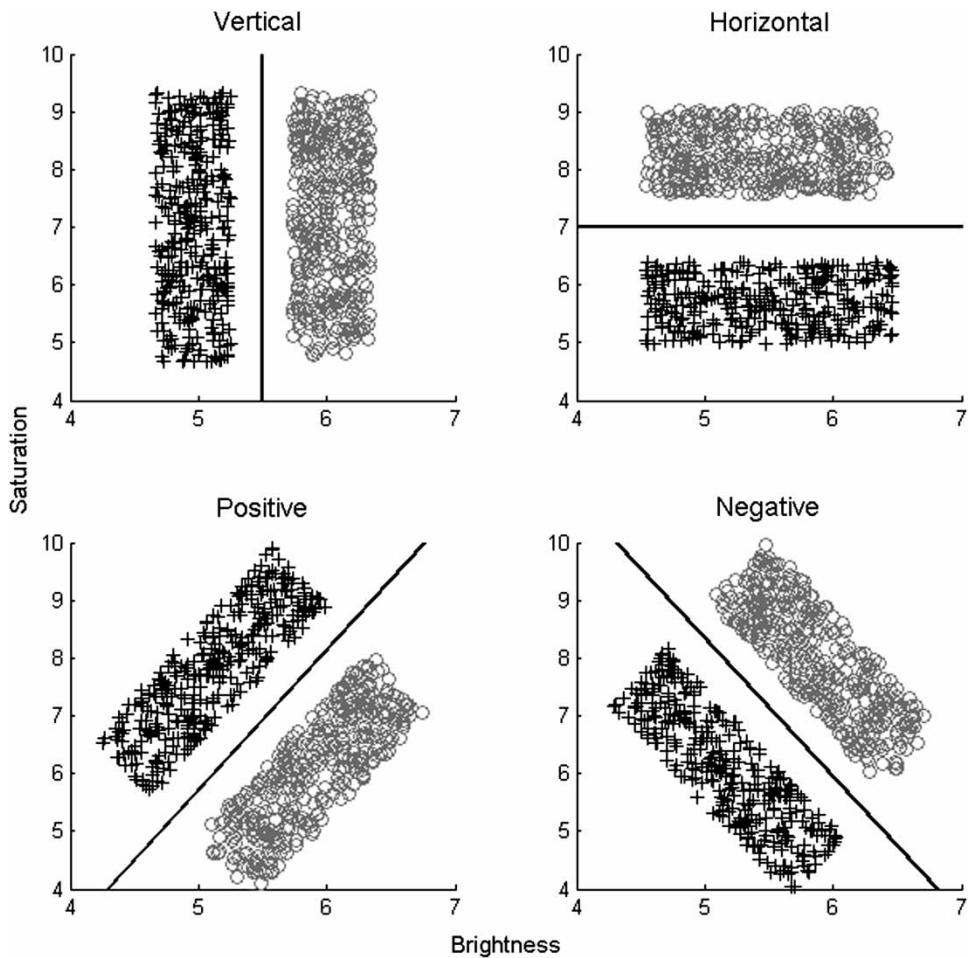
Figure 1. *Scatterplots of the stimuli used in the present Experiments 1 and 2. Each point represents a rectangular, iso-hue colour patch that varied continuously on Munsell value and chroma. Category A and B exemplars are depicted as black plus signs (" +") and grey circles ("o"), respectively. The solid lines are examples of decision strategies that maximize accuracy (i.e., the optimal decision strategies).*

The unsupervised category learning literature makes conflicting predictions for the Figure 1 category structures. Data from constrained tasks with separable dimensions would predict a bias to use one-dimensional decision strategies regardless of the task. As a result, participants should be able to learn the one-dimensional categories, but have difficulty with the diagonal categories. Alternatively, data from unconstrained tasks with integral dimensions would not predict a bias to use one-dimensional decision strategies. For example, similarity-based strategies may be preferred (e.g., Handel & Imai,

1972). This would predict similar performance across the one-dimensional and diagonal category structures because similarity is generally invariant to rotation (Shepard, 1964).

EXPERIMENT 1

Method

Participants and design
Forty participants were recruited from the University of California, Santa Barbara student

community and received partial course credit for participation. Ten participants were randomly assigned to each of four experimental conditions: vertical, horizontal, positive, and negative. No one participated in more than one experimental condition. All participants had normal (20/20) or corrected-to-normal vision and normal colour vision. Each participant completed two sessions of approximately 45 minutes that were separated, on average, by 24 hours. Participants in any condition who were more than three standard deviations away from the average accuracy in that condition during the second day of training were omitted from all subsequent analyses. This criterion for the detection of outliers resulted in the omission of one participant from each of the vertical and negative conditions.

Stimuli and apparatus

The stimuli in all experiments were Munsell colour patches (Munsell, 1915; Newhall, 1940; Newhall, Nickerson, & Judd, 1943) of constant purple–blue hue (10 PB) that varied continuously along the dimensions of value (i.e., brightness) and chroma (i.e., saturation). The complete set of stimuli used in the four different experimental conditions is shown in Figure 1. Each symbol in Figure 1 denotes the value and chroma of a single colour patch. Category A stimuli are denoted by the black "+" signs, and Category B stimuli are denoted by the grey circles. The optimal decision criteria are the vertical and diagonal lines shown in Figure 1.

The experiment used a variation of the randomization technique introduced by Ashby and Gott (1988), in which each category was defined as a bivariate uniform distribution. Each category distribution was specified by the minimum and maximum on each dimension. The exact parameter values for the vertical condition are displayed in Table 1. On each trial, a random sample (x, y) was drawn from the Category A or B distribution, and these values were used to construct a Munsell colour patch of value $x' = 0.0275x + 4.1$ and chroma $y' = 0.055y + 4.3$. While one of the goals of the Munsell system was to equate the perceived difference between equal steps along the value

and chroma dimensions, in practice the perceived difference between two steps on the chroma dimension is approximately perceptually equal to a single step on the value dimension (Newhall, 1940; Newhall et al., 1943; Nickerson, 1940). The choice of scale values in the above transformations was designed to preserve the 2:1 relationship between the dimensions. For each participant, a new sample of 720 stimuli (360 from each category) was generated. All stimuli were generated offline, and a linear transformation was applied to ensure that the sample statistics matched the population parameters.

For the horizontal, positive, and negative conditions, the stimuli for each participant were created by first generating a random sample of 720 stimuli (360 from each category) from the Table 1 distributions and then rotating the resulting stimuli by 90°, −45°, or 45° from vertical (respectively) about the centre of the stimulus space (i.e., the point 50, 50). For the positive and negative category structures, the most accurate one-dimensional rule (i.e., respond A if the stimulus value is less than some criterion, otherwise respond B) yields an accuracy of approximately 85% correct. The presentation order of the stimuli was randomized separately for each participant in every condition and was divided into nine blocks of 80 trials each.

Colour monitor calibration was achieved with a PhotoResearch PR-650 spectral radiometer and the Psychophysics Toolbox software (Brainard, 1997; Pelli, 1997). The transformation from the Munsell colour space to RGB values was performed in three stages. First, the value and chroma coordinates were transformed to CIE xyY chormaticities using the colour lookup table of Newhall et al.

Table 1. *Initial parameters used to generate the vertical category structure before transforming to the Munsell colour space*

Category	Value		Chroma	
	Min	Max	Min	Max
Category A	20	40	7.5	92.5
Category B	60	80	7.5	92.5

(1943) made available by the Carlson Center for Imaging Science (http://www.cis.rit.edu/research/mcsl2/online/munsell_data/real.dat). Those value and chroma coordinates not given in the table were converted to CIE xyY chromaticities using equations given in Wyszecki and Stiles (1982). Second, xyY coordinates were converted to CIE XYZ tristimulus coordinates. Finally, XYZ coordinates were converted to RGB coordinates using the Psychophysics Toolbox software. The experiment was run using the Psychophysics Toolbox in the Matlab computing environment. Each colour patch was presented on a grey background, subtended a visual angle of approximately 6°, and was displayed on a 15-inch CRT with 832 × 624-pixel resolution in a dimly lit room.

Procedure

Each participant was run individually. Participants were told that rectangular colour patches varying in brightness and saturation would be presented one at a time on a monitor, and their task was to learn to categorize the stimuli into two categories.[4] Following Ashby et al. (1999), five response blocks (Blocks 1, 3, 5, 7, and 9) alternated with four observation-only blocks (Blocks 2, 4, 6, and 8). During the observation-only blocks, participants were instructed to look at 80 sequentially presented stimuli and to try and learn about the categories. The stimuli in the observation-only blocks were presented for 1 s with an interstimulus interval of 0.5 s. The observation-only blocks were included in an effort to increase the number of stimuli that the participants were exposed to during an experimental session. The observation-only blocks do not require a response and, thus, take less time to complete than the response blocks (Ashby et al., 1999). During the response blocks, participants were instructed to select a category for each

stimulus and to press a button labelled "A" or a button labelled "B" to show which category had been selected. The participants were told that the category labels were arbitrary, but were warned to be consistent with what they called a member of Category A and what they called a member of Category B. Given that the category labels were arbitrary, it was assumed that participants assigned the stimuli to the two categories in a manner that resulted in the highest accuracy (percentage correct) for each block. Therefore, it was impossible for participants to achieve accuracy below 50% correct in any given block. The participants were told that perfect accuracy was possible, but were never given any feedback about their performance. The stimuli were response terminated (with 5-s maximum exposure duration) in the response blocks, and the response-stimulus interval was 0.5 s. The break between blocks was participant paced.

Results

Accuracy-based analyses

The average learning curves for each of the four experimental conditions are shown in Figure 2. Visual inspection of Figure 2 suggests that accuracy improved across the two days of training only in the vertical and positive conditions, and accuracy was highest in the VERTICAL condition. A 4 (condition) × 10 (response block) mixed analysis of variance (ANOVA; with block as the within-subjects factor) conducted on the accuracy data revealed significant main effects of condition, $F(3, 34) = 5.46$, $p < .01$, $MSE = .35$, $\eta_p^2 = .33$, and block, $F(5.88, 199.95) = 2.22$, $p < .05$, $MSE = .06$, $\eta_p^2 = .06$.[5] However the Condition × Block interaction was not significant, $F(17.65, 199.95) = 1.41$, $p = .13$, $MSE = .06$, $\eta_p^2 = .11$. Further analysis of the main effects revealed that the effect of condition

---

[4] All participants indicated that they understood what brightness was, but several were unfamiliar with saturation. Thus, for all participants, saturation was described as the amount of white in a colour patch, with low levels of saturation indicating a large amount of white in the colour patch. As an example, participants were told that pink is a desaturated red. This explanation was effective (according to verbal report) in eliminating any confusion regarding the saturation dimension.

[5] To meet the assumptions of ANOVA, these data were first subjected to an arcsine transformation. For descriptive purposes, the nontransformed data were presented in all figures and tables. A Huynh–Feldt correction for violation of the sphericity assumption has been applied. All subsequent analyses of accuracy rates used the same transformation and correction. Post hoc comparisons were evaluated using the Student–Newman–Keuls procedure.

Figure 2. *Average accuracy in the four conditions of Experiment 1.*

was driven solely by superior performance in the vertical condition relative to the horizontal, positive, and negative conditions (ps < .05). None of the other pairwise comparisons was significant (ps ≥ .39). The effect of block was driven by an increasing linear trend in accuracy across conditions, as evidenced by a significant linear contrast, $F(1, 34) = 5.75$, $p < .05$, $MSE = .09$, $\eta_p^2 = .15$. In sum, at the group level, a high level of performance was observed only in the vertical condition—the category structure that required participants to attend selectively to brightness while ignoring variations in saturation.

The superior performance in the vertical condition can be seen at the individual participant level as well. Table 2 lists the individual average accuracy rates by block for each participant in each condition during the second day of training. All but one participant in the vertical condition was near optimal (>90%) during the final block of training. In contrast, only eight participants in the remaining conditions (three in the horizontal, four in the positive, and one in the negative) achieved a similar level of accuracy. Furthermore,

two thirds of the participants from the vertical condition maintained an accuracy level >90% during the last four response blocks, whereas only two participants in the other conditions (one in the horizontal and one in the positive) performed at this high level. Interestingly, the individual participant data suggest that while it was certainly more difficult for participants to improve with training in the horizontal, positive, and negative conditions in general, it was not altogether impossible.

In an ideal learning trajectory, accuracy might steadily improve across trials and peak at the completion of training. However, it is possible that participants may have peaked during some training block other than the last. In fact this was true for 29 of the 38 participants. Analysing the accuracy in this way did not change the ordering: vertical ($M = 97.6$, $SD = 4.1$), positive ($M = 91.8$, $SD = 8.5$), negative ($M = 88.2$, $SD = 10.8$), and horizontal ($M = 80.9$, $SD = 13.1$). A one-way ANOVA using each participant's best block generally supported this conclusion. The effect of condition was significant, $F(3, 34) = 5.88$, $p < .01$,

Table 2. *Individual participant accuracy during the second day of training in Experiment 1*

Condition	Block	Participant										Avg	SEM
		1	2	3	4	5	6	7	8	9	10		
Vertical	1	99	58	76	55	53	51	100	95	99		76.1	7.0
	3	99	61	83	94	100	100	100	95	98		92.1	4.1
	5	96	88	76	78	99	100	100	100	99		92.8	3.1
	7	99	71	53	68	99	98	98	96	100		86.6	5.7
	9	99	95	76	91	100	100	95	91	100		94.2	2.4
Horizontal	1	55	68	76	68	62	63	81	70	63	85	69.0	2.8
	3	73	59	78	73	50	54	93	79	63	85	70.5	4.2
	5	70	51	89	81	66	64	98	84	63	90	75.5	4.5
	7	58	62	60	84	58	60	96	74	51	88	69.0	4.6
	9	51	66	90	55	60	59	95	84	54	90	70.4	5.2
Positive	1	74	98	76	80	68	70	69	80	99	73	78.5	3.5
	3	78	100	84	93	71	90	60	78	79	78	80.8	3.6
	5	98	99	88	85	71	56	63	58	74	68	75.8	8.0
	7	98	100	88	88	79	79	68	66	85	75	82.4	3.6
	9	94	100	91	63	60	94	64	74	83	66	78.7	4.8
Negative	1	76	85	51	78	79	51	93	55	88		72.7	4.8
	3	91	51	66	53	70	64	93	59	95		71.3	5.2
	5	85	71	55	89	83	51	81	53	88		72.8	4.7
	7	94	74	70	62	62	51	85	50	91		71.1	4.9
	9	80	81	51	51	67	55	85	54	98		69.1	5.2

Note: Accuracy values, in percentages. Avg = average.

$MSE = .06$, $\eta_p^2 = .34$, with accuracy in the vertical condition being significantly higher than that in the horizontal and negative conditions ($ps < .05$), but only marginally higher than that in the positive condition ($p = .06$). This lack of a significant difference between the best block for the vertical and positive conditions may reflect a ceiling effect in the vertical condition.

Model-based analyses

Analysis of the accuracy data does not directly address the question of what decision strategies were used to perform the Figure 1 tasks. For instance, accuracy for many of the participants in the positive and negative conditions was consistent with both a one-dimensional strategy and a strategy that integrated brightness and saturation (albeit in a suboptimal manner). The following analyses represent a quantitative approach to investigating these questions.

Three different types of decision bound model were fitted to the data of each individual participant, each based on a different assumption concerning the participant's strategy. First, the unidimensional classifiers assume that the participant attends selectively to one dimension (e.g., if the stimulus is bright, respond B; otherwise respond A). For the vertical and horizontal conditions, there were two versions of the unidimensional classifier (UC), one assuming that participants used the optimal decision strategy in the two top panels of Figure 1 (optimal classifier, OC) and one assuming that participants used a UC with a suboptimal intercept on one of the dimensions (UC-brightness and UC-saturation). Second, the conjunctive classifier (CC) assumes that participants make independent decisions about the stimulus on both dimensions (e.g., if the stimulus is bright and saturated, respond B; otherwise respond A). Third, the linear classifier

assumes that participants integrate the stimulus information from both dimensions prior to making a categorization decision. For the positive and negative conditions, there were two versions of the linear classifier, one assuming that participants used the optimal decision strategy in Figure 1 (optimal classifier, OC) and one assuming that participants used a linear classifier with a suboptimal slope and/or intercept (LC).

Each of these models was fitted separately to the data from every response block for all participants using a standard maximum likelihood procedure for parameter estimation (Ashby, 1992b; Wickens, 1982) and the Bayes information criterion for goodness-of-fit (Schwarz, 1978; see the Appendix for a more detailed description of the models and fitting procedure). The data from the first day of training are omitted for brevity.

The primary goal of this analysis was to investigate whether one-dimensional decision strategies dominated under unsupervised training. The distribution of best fitting models in each of the four conditions is listed in Table 3. First consider the one-dimensional conditions. In these conditions, both the optimal and unidimensional classifiers assumed that participants attended selectively to the relevant stimulus dimension. In the vertical condition, 73% were using decision strategies consistent with selective attention to brightness. In contrast, in the horizontal condition, only 8% were using decision strategies consistent with selective attention to saturation. Instead, participants were either attending selectively to brightness or integrating brightness and saturation information to some extent. Given the relatively low accuracy in the horizontal condition, it is important to verify that the linear classifier was not simply better at fitting noisy data. The high percentage of responses accounted for argue strongly against this possibility. Thus, the deficit observed in the horizontal condition was, at least in part, driven by an inability to attend selectively to saturation in the absence of feedback, although this explanation does not account for the fact that average accuracy during blocks where participants were attending selectively to saturation was far less than optimal ($N = 4$, $M = 71.3$, $SD = 17.8$).

Instead, the accuracy deficit during these blocks was due to the use of suboptimal decision criteria.

Next, consider the diagonal conditions. As expected from the accuracy data, participants in both the positive and negative conditions were responding optimally in a relatively small number of blocks. Participants were not constrained to use one-dimensional decision strategies, as evidenced by the relatively small percentage of data sets accounted for by unidimensional models (34% and 22% of the data in the positive and negative conditions, respectively). Instead, as was the case in the horizontal condition, in the majority of blocks, participants were integrating brightness and saturation information, albeit not optimally.

The fact that the linear classifier accounted for a large percentage of the responses in the horizontal, positive, and negative conditions may still be consistent with the use of one-dimensional decision strategies. In some cases, the best fitting linear classifier may deviate only slightly from one-dimensional (e.g., a decision bound rotated 10° from horizontal). Such small deviations may suggest a weak form of selective attention to one of the stimulus dimensions. Plotted in Figure 3 is the distribution of slopes (in degrees) for those data sets that were best accounted for by the linear classifier in the four conditions. There were a substantial number of strategies that would be consistent with a weak form of selective attention to brightness (i.e., near 90°) in the vertical, horizontal, and negative conditions, but not in the positive condition. In addition, in the horizontal condition, the best fitting linear bounds were highly variable and deviated substantially from optimal. In the positive condition, the linear bounds were also highly variable with the majority being positively sloped.

Figure 3 also suggests that when participants were not attending selectively to brightness or saturation, or using the optimal decision strategy, they used decision strategies between 0° and 90°. This preference suggests that there may be some salient decision rule that is a consequence of the integral stimulus dimensions. Indeed, inspection of the entire stimulus space (i.e., plotting all stimuli simultaneously) suggests a strategy that

Table 3. *Percentage of blocks best accounted for by each model across the four conditions of Experiment 1*

	Models						
Condition	OC	UC-B	UC-S	CC	LC	Avg. %RA	SD
Vertical	46.7	26.7	8.9	0	17.8	95.1	7.3
Horizontal	4.0	26.0	4.0	10.0	56.0	88.2	10.9
Positive	14.0	18.0	16.0	12.0	40.0	91.3	8.2
Negative	11.1	20.0	2.2	2.2	64.4	86.9	9.5

Note: OC = optimal classifier. UC = unidimensional classifier. UC-B = unidimensional classifier on the brightness dimension. UC-S = unidimensional classifier on the saturation dimension. CC = conjunctive classifier. LC = linear classifier. Avg = average. %RA = percentage of responses accounted for by the best fitting model. In all conditions, the OC is a special case of one of the other models in which it is assumed that the participant used the optimal decision strategies plotted in Figure 1 (vertical: The OC is a special case of the UC-B; horizontal: The OC is a special case of the UC-S; positive and negative: The OC is a special case of the LC).

could be best described as a "greyness" rule. More specifically, the transition from dim, saturated stimuli to bright, desaturated stimuli could produce such a rule with a slope between 0° and 90°. If participants were truly using such a strategy, it is clear from Figure 3 that its application was highly variable across participants. Consistent with these data, inspection of written descriptions of decision strategies (collected postexperimentally) did not reveal any systematic strategy use or any mention of the word "grey".

Recall that the decision bound models not only provide estimates of the best fitting decision bounds, but estimates of the combined variance of criterial and perceptual noise (σ^2). A comparison of σ^2 estimates across conditions suggests that noise was lower in the vertical condition ($Mdn = 0.26$) than in the horizontal ($Mdn = 2.9$), positive ($Mdn = 0.91$), and negative ($Mdn = 1.27$) conditions. This observation was supported by an analysis of the estimates of σ^2 between all conditions using six separate t tests (Welch's t test, and not ANOVA, was used due to severe violations in homogeneity of variance). All pairwise comparisons were significant ($p < .0085$, following Sidak correction for multiple comparisons).

Brief summary

At least two conclusions can be drawn regarding unsupervised learning of categories constructed from the integral dimensions of brightness and saturation. First, participants demonstrate a relatively limited ability to learn such categories. Only when the category structures required selective attention to brightness were participants able to learn without feedback. Second, these data suggest that participants do not show either the preference for, or a general ability to learn, one-dimensional decision strategies under unsupervised conditions.

EXPERIMENT 2

Given the lack of learning in the horizontal, positive, and negative conditions under unsupervised conditions, it is necessary to demonstrate that participants can, in fact, learn these category structures. The goal of Experiment 2 was to test the ability of participants to learn the Figure 1 categories under supervised conditions. Furthermore, the majority of research on the ability of people to learn categories constructed from integral dimension stimuli has been limited to stimuli constructed from discrete- or binary-valued rather than continuous-valued dimensions. Thus, an added contribution of Experiment 2 is that it extends previous research to continuous-valued dimensions.

Method

Participants and design

Twenty-two participants were recruited from the University of California, Santa Barbara student
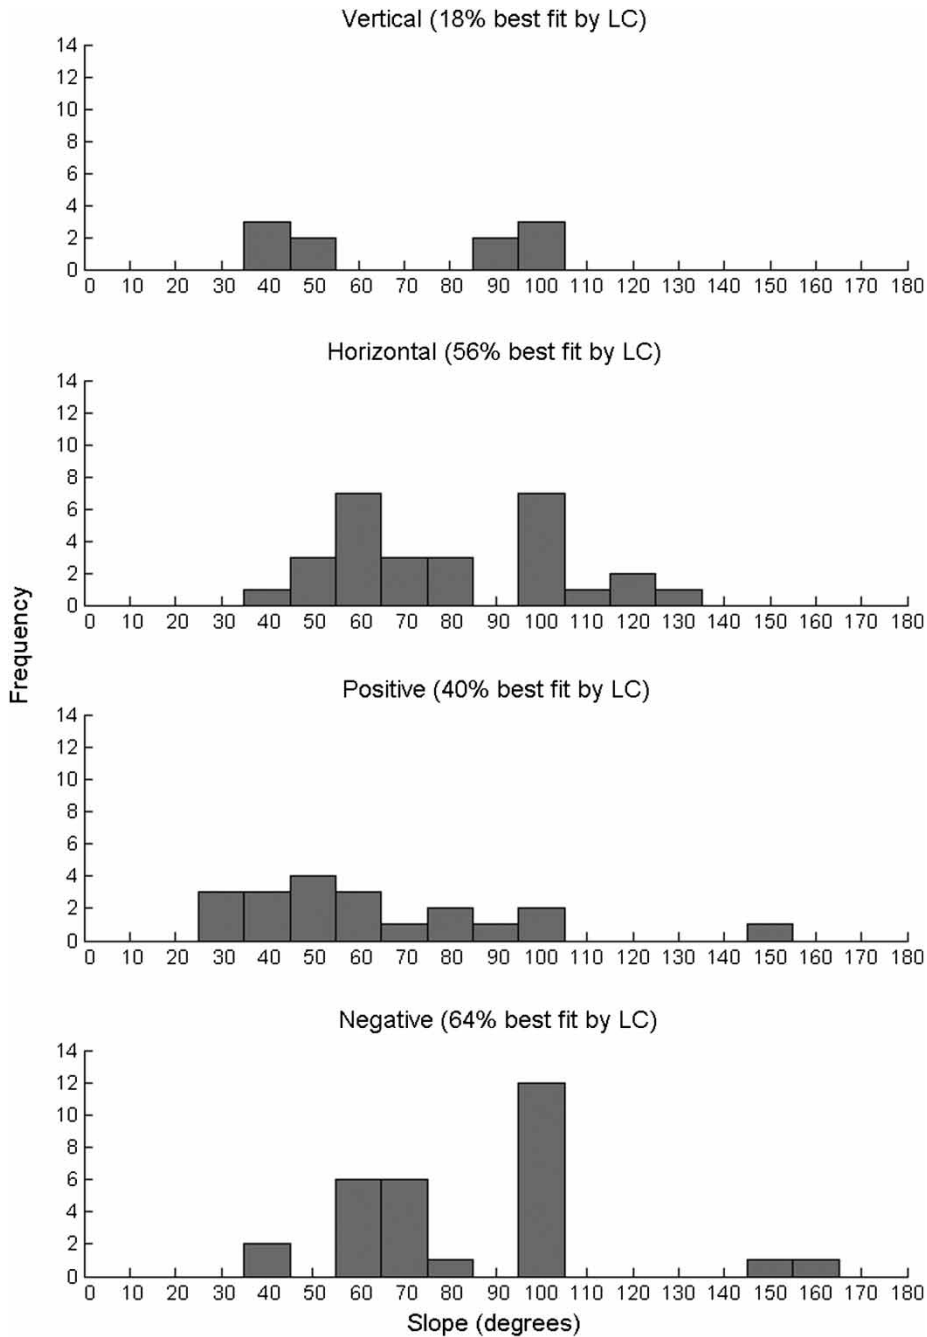
Figure 3. *Distribution of the slope (in degrees) for cases where the linear classifier (LC) provided the best account of the data. The percentage of blocks for which the linear classifier provided the best fit is provided in the figure legends for reference. For descriptive purposes, it was assumed that the slopes ranged from 0° to 180°. The optimal classifier predicts the following slopes: vertical, 90°; horizontal, 0°/180°; positive, 68°; negative, 113°. bin–width = 10°.*

community and received partial course credit for participation. The participants were randomly assigned to the four experimental conditions in the following manner: vertical-6, horizontal-5, positive-6, and negative-5. No one participated in more than one experimental condition. All participants had normal (20/20) or corrected-to-normal vision and normal colour vision. Participants completed one session of approximately 45 min.

Stimuli and apparatus
Stimuli and apparatus were identical to those in Experiment 1.

Procedure
The procedure was identical to that in Experiment 1 with the following exceptions. All participants were presented with nine response blocks comprising 80 trials each. Each stimulus was presented for 1 s followed by a brief (0.5-s) high-pitched tone (500 Hz) if the response was correct and a low-pitched tone (200 Hz) if the response was incorrect. In addition, feedback was given at the end of each block regarding the participant's accuracy during that block.

Results

Accuracy-based analyses
The average learning curves for each of the four experimental conditions are shown in Figure 4. Visual inspection of Figure 4 suggests an ordering of the four conditions by task difficulty early in training similar to that observed at the end of training in Experiment 1. Specifically, the participants in the vertical condition were the most accurate, followed by the positive, negative, and horizontal conditions. These accuracy differences, however, were nonexistent by the end of training. A 4 (condition) × 9 (response block) mixed ANOVA conducted on the accuracy data (with block as the within-subjects factor) largely supported the visual inspection of Figure 4. There was a significant main effect of block, $F(8, 152) = 33.68$, $p < .001$, $MSE = .01$, $\eta_p^2 = .64$, which was qualified by a significant Condition × Block interaction, $F(24, 152) = 2.19$, $p < .01$, $MSE = .01$, $\eta_p^2 = .26$. The

main effect of condition was not significant, $F(3, 19) = 2.50$, $p = .09$, $MSE = .14$, $\eta_p^2 = .28$. A simple main effects analysis revealed a pattern of results consistent with the visual analysis of Figure 4. Specifically, accuracy in the vertical condition during the first response block was higher than accuracy in both the horizontal ($p < .01$) and the negative ($p < .05$) conditions, but not the positive condition ($p = .17$). Accuracy in the vertical condition continued to exceed that of the horizontal condition ($p < .05$), but not that of the negative condition ($p = .48$) during Block 2. The difference between the vertical and horizontal conditions was no longer present during Block 3 ($p = .11$). None of the remaining pairwise comparisons was significant.

Model-based analyses
The same models as those investigated in Experiment 1 were fitted to each participant's responses separately for every block in Experiment 2. The distribution of best fitting models in each of the four conditions is listed in Table 4. The first thing to note is that, in comparison to Experiment 1, the percentage of blocks in which the optimal classifier was the best fitting model greatly increased in all conditions with the addition of feedback. In the one-dimensional conditions, decision strategies assuming participants attended selectively to brightness and saturation provided the best account of the data on 91% (vertical) and 69% (horizontal) of the blocks, respectively. In the diagonal conditions, the use of one-dimensional decision strategies was far less frequent than that in Experiment 1.

In addition, the linear classifier provided the best fit to almost 25% of the blocks in the horizontal, positive, and negative conditions. While this percentage is far less than that observed in Experiment 1, it is still worthwhile to determine whether or not these bounds were consistent with a weak form of selective attention or the integration of brightness and saturation. The distribution of slopes estimated from the linear classifier for those blocks in which the linear classifier provided the best fit is plotted in Figure 5. The results from the horizontal and negative conditions are quite

Figure 4. *Average accuracy in the four conditions of Experiment 2.*

Table 4. Percentage of blocks best accounted for by each model across the four conditions of Experiment 2.

Condition	Models				
	OC	UC-B	UC-S	CC	LC
Vertical	77.8	13.0	0.0	1.9	7.4
Horizontal	62.2	2.2	6.7	6.6	22.2
Positive	59.3	14.8	0.0	0.0	25.9
Negative	60.0	13.3	0.0	6.7	20.0

Note: OC = optimal classifier. UC = unidimensional classifier. UC-B = unidimensional classifier on the brightness dimension. UC-S = unidimensional classifier on the saturation dimension. CC = conjunctive classifier. LC = linear classifier. In all conditions, the OC is a special case of one of the other models in which it is assumed that the participant used the optimal decision strategies plotted in Figure 1 (vertical: The OC is a special case of the UC-B; horizontal: The OC is a special case of the UC-S; positive and negative: The OC is a special case of the LC).

similar to those of Experiment 1. The slopes from the horizontal condition were highly variable and not consistent with one-dimensional decision strategies. In the negative condition, the majority of the slopes were consistent with a weak form of selective attention to brightness (i.e., between 90° and 110°). In contrast to Experiment 1, those blocks that were best fit by the linear classifier in the positive condition were consistent with a weak form of selective attention on brightness.

Brief summary

The results of Experiment 2 indicate that the Figure 1 category structures can be learned with feedback and, more generally, that category structures constructed from continuous-valued, integral (i.e., brightness and saturation) dimensions are easily learned. In all conditions, participants were more accurate, and there was an increase in the percentage of participants using optimal decision strategies with the addition of feedback. Perhaps not surprisingly, the ordering by task difficulty early in training mimicked that observed when feedback was omitted.

Figure 5. *Distribution of the slope (in degrees) for cases where the linear classifier (LC) provided the best account of the data in Experiment 2. The percentage of blocks for which the linear classifier provided the best fit is provided in the figure legends for reference. For descriptive purposes, it was assumed that the slopes ranged from 0° to 180°. The optimal classifier predicts the following slopes: vertical, 90°; horizontal, 0°/180°; positive, 68°; negative, 113°. bin-width = 10°.*

EXPERIMENT 3

The category structures used in Experiments 1 and 2 were designed to equate discriminability across brightness and saturation. Even so, in both experiments, participants were better when the categorization judgement required selective attention to brightness (the vertical condition) than when it required selective attention to saturation (the horizontal condition). This leaves open the possibility that the categories in the horizontal condition were less discriminable than the categories in the vertical condition. To address this question, we ran an unsupervised version of the horizontal condition in which we varied category discriminability along the saturation dimension by increasing the intercategory distance (Figure 6A). If the ability to learn one-dimensional strategies on saturation is dependent upon category discriminability, then accuracy should be higher, and one-dimensional strategies should be used more frequently, in the high discriminability condition than in the low discriminability condition.

One consequence of increasing the intercategory distance is that there is also an increase in the number of qualitatively different decision strategies that predict high accuracy. For example, in the high discriminability condition, a one-dimensional strategy on saturation would be indistinguishable from many strategies that integrate saturation and brightness. To address this issue, we replaced the final response block with a test block using a uniform grid of stimuli that spanned the range of the training stimuli (Figure 6B). Fitting the models to the categorization responses from the test block will provide a stronger test of the use of one-dimensional strategies.

Method

Participants and design

Twenty-six participants were recruited from the University of Maine student community and received partial course credit for participation. Thirteen participants were randomly assigned to each of the two experimental conditions. No one participated in more than one experimental condition. All participants had normal (20/20) or corrected to normal vision and normal colour vision.

Participants completed one session of approximately 45 min.

Stimuli and apparatus

The stimulus generation procedures were identical to those in Experiments 1 and 2, with two exceptions. First, the parameters in Table 5 were used to generate the training stimuli (see Figure 6A). Second, to facilitate the identification of decision strategies in the model-based analyses, the stimuli from the final response block were replaced with a test block that included a uniform grid of stimuli selected to match the range of the training stimuli. In the low discriminability condition, the test stimuli were generated by all possible combinations of 8 equally spaced points on brightness (from 7.5 to 92.5) and 10 equally spaced points on saturation (low discriminability: from 20 to 80; high discriminability: from 5 to 95). The test block is not included in the accuracy-based analysis because there is no objectively correct response for many of the test stimuli with respect to the trained categories.

Procedure

The procedure was identical to that of Experiment 1. The participants were not informed that the stimuli during the final test block differed from the stimuli during the earlier training blocks.

Results

Accuracy-based analyses

Average accuracy during training is plotted in Figure 7. A clear accuracy advantage emerged with training for participants in the high discriminability condition. The results of a 2 (condition) × 4 (response block) mixed ANOVA (with block as the within-subjects factor) were consistent with this claim. Specifically, the main effect of condition, $F(1, 24) = 46.94$, $p < .001$, $MSE = 135.23$, $\eta_p^2 = .66$, and the Condition × Block interaction, $F(1.99, 47.86) = 4.99$, $p < .05$, $MSE = 95.54$, $\eta_p^2 = .17$, were significant. The interaction was driven by superior performance in the high discriminability condition during the last three response blocks (Block 1: $p = .11$; Blocks 3, 5, and 7: $ps < .001$). The main effect of block was not significant, $F(1.99, 47.86) = 0.21$, $p = .81$, $MSE = 95.54$, $\eta_p^2 = .009$.
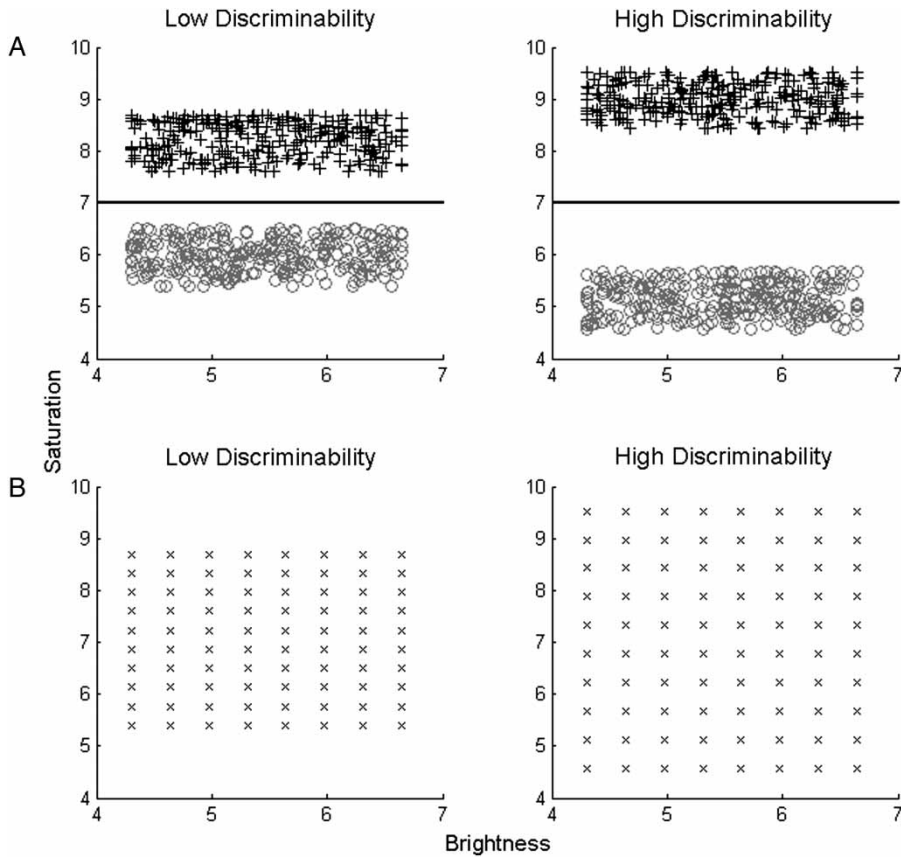
Figure 6. *A. Scatterplots of the training stimuli used in Experiment 3. Each point represents a rectangular, iso-hue colour patch that varied continuously on Munsell value and chroma. Category A and B exemplars are depicted as black plus signs ("+") and grey circles ("o"), respectively. The solid lines are examples of decision strategies that maximize accuracy (i.e., the optimal decision strategies). B. Scatterplots of the test stimuli used during the final response block. The stimulus coordinates are equally spaced and were selected to span the range of the training stimuli.*

Table 5. *Parameters used to generate the low and high discriminability category structures before transforming to the Munsell colour space*

Discriminability	Value		Chroma	
	Min	*Max*	*Min*	*Max*
Low				
Category A	7.5	92.5	60	80
Category B	7.5	92.5	20	40
Category A	7.5	92.5	75	95
Category B	7.5	92.5	5	25

Model-based analyses

Recall that increasing intercategory distance comes with the cost of decreased identifiability of the decision strategy. Thus, the higher accuracy in the high discriminability condition during training may be a consequence of the success of strategies that integrate brightness and saturation rather than an increase in the use of strategies assuming selective attention to saturation. To address this issue, we focused the model-based analyses on the data from the test block in which the stimuli were sampled uniformly from across the range of the
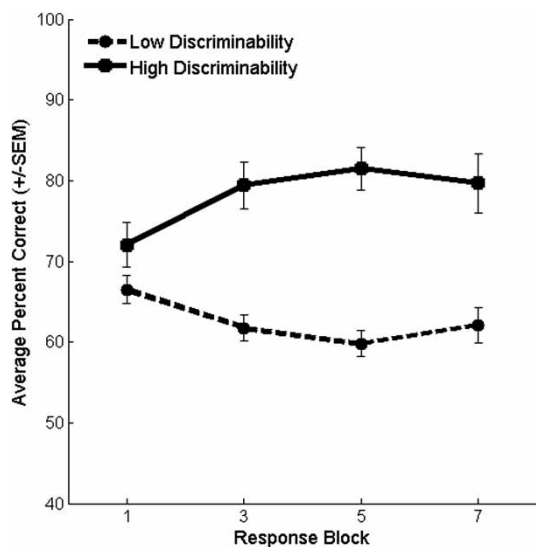
Figure 7. *Average accuracy in the two conditions of Experiment 3.*

training stimuli. The same models as those investigated in Experiments 1 and 2 were fitted to each participant's responses separately for the test block, and the distribution of best fitting models is listed in Table 6. As was the case in Experiment 1, decision strategies consistent with selective attention to saturation were extremely rare ($n_{low} = 2/13$, $n_{high} = 1/13$). This pattern held across the training blocks, with 12.3% and 13.8% of the blocks being best accounted for by decision strategies consistent with selective attention to

saturation in the low and high discriminability conditions, respectively. Importantly, these data suggest that increasing category discriminability for categories defined on saturation was not accompanied by an increase in the use of one-dimensional strategies on saturation.

The preceding analysis implies that the increased accuracy in the high discriminability condition occurred because the more widely separated categories in that condition did not penalize strategies in which selective attention to saturation failed as much as those in the low discriminability condition did. The relative degree to which participants attended to saturation versus brightness can be assessed by examining parameter estimates from the best fitting version of the linear classifier. Recall that in this model, the slope of the decision bound is free to vary. Slopes near 0 or 180° (counterclockwise from horizontal) are consistent with a weak form of selective attention to saturation, whereas slopes near 90° are consistent with a weak form of selective attention to brightness. As shown in Figure 8, the slopes for participants best fitted by the linear classifier were more consistent with a weak form of selective attention to brightness (i.e., 90° counterclockwise from horizontal; $M_{low} = 102.4°$, $SE_{low} = 8.6$; $M_{high} = 95.2°$, $SE_{high} = 8.7$). Thus, despite the higher accuracy of participants in the high discriminability condition, participants in this condition appeared to allocate more attention to the irrelevant brightness dimension than to the relevant saturation dimension.

Brief summary

The goal of Experiment 3 was to determine whether increasing category discriminability improved the ability of participants to learn categories defined by saturation in the absence of feedback. Although participants were more accurate when discriminability was increased, the increased accuracy was not driven by an increased ability to attend selectively to saturation. Instead, consistent with Experiment 1, participants tended to integrate brightness and saturation in a manner that suggested greater weighting of brightness.

Table 6. Number of blocks best accounted for by each model during the final test block of Experiment 3

	Models				
Discriminability	*OC*	*UC-B*	*UC-S*	*CC*	*LC*
Low	1	4	1	0	7
High	0	1	1	0	11

Note: OC = optimal classifier. UC = unidimensional classifier. UC-B = unidimensional classifier on the brightness dimension. UC-S = unidimensional classifier on the saturation dimension. CC = conjunctive classifier. LC = linear classifier. The OC is a special case of the UC-S in which it is assumed that the participant used the optimal decision strategy plotted in Figure 6.
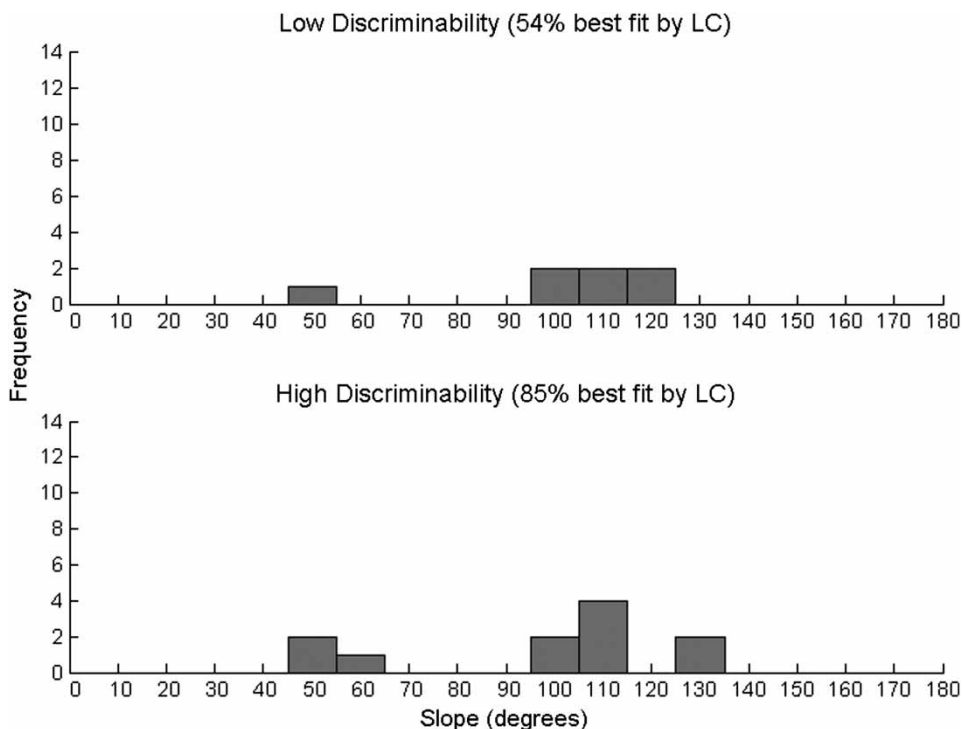
Figure 8. *Distribution of the slope (in degrees) for cases where the linear classifier (LC) provided the best account of the data in Experiment 3. The percentage of blocks for which the linear classifier provided the best fit is provided in the figure legends for reference. For descriptive purposes, it was assumed that the slopes ranged from 0° to 180°. The optimal classifier predicts a slope of 0° or 180°. Bin-width = 10°.*

GENERAL DISCUSSION

Previous research on the unsupervised categorization of stimuli constructed from integral dimensions has focused primarily on categorization preferences in unconstrained tasks where no instructions are provided about the underlying category structure. Although such studies are clearly important, they do not address the ability of individuals to learn in constrained tasks where the goal is to learn a predefined category structure. The present study makes an important contribution to the literature by addressing this question. Experiment 1 showed that the ability to learn categories constructed from brightness and saturation under unsupervised conditions varies as a function of the category structure. Specifically, in the absence of feedback, participants were able to learn only when the category structures required attending selectively to brightness. Experiment 2

demonstrated that these category structures can be learned with feedback and that categorization based on brightness is easier than categorization based on saturation. Experiment 3 demonstrated that unsupervised categorization accuracy for categories defined by saturation can be increased by increasing category discriminability. In contrast to the vertical condition of Experiment 1, however, higher accuracy was not driven by the use of strategies assuming selective attention to the relevant dimension, but rather by the increased accuracy associated with strategies assuming the integration of brightness and saturation.

These data are partially consistent with the predictions motivated by both constrained and unconstrained unsupervised category learning tasks. Data from constrained tasks using separable dimensions predict a bias to use one-dimensional strategies (e.g., Ashby et al., 1999). Strong evidence in support of this bias, however, was only observed

when the categories were defined by brightness. In contrast, data from unconstrained tasks using integral dimensions typically predict that there is a bias to use similarity-based strategies. Indeed, the majority of data from the Experiment 1 and 3 conditions were best accounted for by strategies assuming the integration of brightness and saturation. In each experiment, however, a subset of the participants that integrated brightness and saturation used strategies that weighted brightness more heavily than saturation, suggesting a weak form of selective attention to brightness. Furthermore, the different patterns of strategy use across experimental conditions and the general advantage for the vertical conditions of Experiments 1 and 2 contradict the assumption that there should be a general bias to use similarity-based strategies with integral dimensions. If this was true, then performance should have been equivalent in all conditions since the various conditions were all rotations of each other, and similarity with integral dimensions is generally thought to be rotation invariant (Shepard, 1964).

Is brightness privileged in the Munsell system?

Several aspects of these data suggest that brightness is more efficiently processed than saturation. First, under unsupervised conditions, participants were most accurate when the categories were defined by brightness. Second, under supervised conditions, participants learned categories defined by brightness at a faster rate. Third, participants relied upon one-dimensional rules on brightness much more frequently than those on saturation. Fourth, even when category discriminability along saturation was increased, participants rarely used one-dimensional strategies on saturation. Finally, as might be expected with integral dimensions, many participants integrated brightness and saturation in the absence of feedback. In general, however, these participants did not give brightness and saturation equal weighting. Instead, there was a bias to give brightness more weight in the categorization process.

Indeed, attentional mechanisms operate more efficiently with brightness than saturation under

supervised conditions (Maddox & Dodd, 2003; Nosofsky, 1987). The Maddox and Dodd experiments suggest that the advantage for brightness over saturation is driven by a perceptual bias (i.e., the perceptual representation of brightness is less noisy than the perceptual representation of saturation). In spite of this difference, participants can clearly learn categories defined by saturation under supervised conditions. One possibility is that the perceptual advantage for brightness over saturation is exaggerated under unsupervised conditions.

Furthermore, in everyday life, people are constantly making discriminations based on brightness, but how common are saturation discriminations? In fact, all participants indicated that they understood what brightness was, but several were unfamiliar with saturation and had to be given specific examples. Clearly participants demonstrated an ability to attend selectively to saturation when feedback was provided, but the horizontal condition still required more training to learn than the vertical condition in Experiment 2. Moreover, only 3 of the 10 participants were ever able to achieve near optimal accuracy without the aid of feedback. It may be that with extended training, more participants would have been successful learning categories defined by saturation in the absence of feedback. This, however, seems unlikely given the lack of improvement observed in the horizontal condition of Experiment 1 and the low discriminability condition of Experiment 3. Thus, despite the fact that the Munsell system was designed to equate variation on brightness and saturation, there may be an advantage for making decisions based on brightness (at least when it is paired with saturation).

Experiment 3 was designed to explore whether it is possible to overcome the disadvantage for categories defined by saturation. More specifically, would increasing category discriminability (i.e., increasing intercategory distance) improve accuracy by increasing the use of one-dimensional strategies on saturation? Perhaps not surprisingly, accuracy improved with category discriminability. Critically, however, the improvement in accuracy was not driven by an increase in the use of one-dimensional strategies on saturation. Instead, the

improvement in accuracy was driven by the increased success of strategies assuming the integration of brightness and saturation that resulted from increasing intercategory distance.

We chose to define our stimuli in the Munsell colour system because it provides a more direct connection to previous work on the categorization of stimuli constructed from integral dimensions. Of course, there are many other colour systems. Interestingly, brightness, but not saturation, is generally represented across colour systems. For example, brightness in the Munsell system is monotonically related to luminance in the Natural Color System (e.g., Brainard, 2003). Moreover, it has been suggested that luminance (and not saturation) may be one of the features that guides preattentive visual processing (Wolfe, 2005). Although speculative, these arguments suggest that brightness may be a more fundamental feature than saturation in the representation of colour.

Constrained categorization of separable versus integral dimensions

Data from constrained tasks demonstrate that individuals are capable of learning categories constructed from separable dimensions in the absence of feedback (Ashby et al., 1999; Zeithamova & Maddox, 2009). This capability, however, appears to be limited as individuals were not able to learn when a one-dimensional strategy predicted poor performance (i.e., similar to the diagonal conditions; Ashby et al., 1999). Moreover, even when there is a highly accurate one-dimensional strategy, unsupervised learning with separable dimensions appears to be limited to categories that are highly discriminable (Ell & Ashby, in press).

Given the dominance of one-dimensional rules that had been observed in some previous studies with separable stimulus dimensions (e.g., Ashby et al., 1999), it is surprising that participants in the horizontal condition did not outperform those in the diagonal conditions. One possible explanation relates to our definition of a one-dimensional rule. Here we have operationally defined a one-dimensional rule as a decision bound orthogonal to the physical dimensions of value or

chroma. It is possible that the psychological representation of the decision strategy does not correspond exactly to the dimensional structure we intended (e.g., Melara, Marks, & Potts, 1993). It is also possible that the mapping between the physical space and the perceptual (Brightness × Saturation) space is nonlinear. The Munsell system is based on scaling judgements performed on one colour attribute while the remaining attributes were held constant. Thus, the relations proposed in the Munsell system are not guaranteed to hold when the dimensions are varied orthogonally (Brainard, 2003). Therefore, there is reason to expect that the psychological representation of a rule may not exactly match our operational definition of a rule in the Value × Chroma space.

This explanation would be more compelling if participants demonstrated some degree of consistency in their decision strategy in the saturation-relevant conditions. In contrast to the vertical condition, there was little agreement in the best fitting decision bounds across participants. If there was some psychological rule that did not correspond to the physical axes, then one would expect some degree of agreement between the participants. Of course, these data do not rule out this interpretation as the different participants may have each been attending selectively to different psychological rules. To the extent that categorization based on saturation is less intuitive than categorization based on brightness, the variability in strategy use is consistent with recent work by Pothos and colleagues (Pothos et al., 2011).

It would be inaccurate to say that unsupervised learning is impossible in the diagonal conditions. Although the optimal classifier was the best fitting model for only a small number of blocks in both the positive and the negative conditions, the data from several other blocks that were best fitted by the linear classifier did not deviate substantially from the optimal decision strategy. Therefore, it appears as though it is possible to successfully integrate information from integral stimulus dimensions in the absence of feedback, but also that there are large individual differences in this ability.

The distribution of best fitting models observed in the diagonal conditions is quite different from that observed with separable dimensions. For example, Ashby et al. (1999) found that participants almost exclusively relied upon one-dimensional decision strategies in category structures similar to those in the present diagonal conditions. This is not the case when the stimuli are constructed from integral dimensions. In the positive condition, participants were integrating brightness and saturation information rather than attending selectively to one of those dimensions. In contrast, in the negative condition, there was more frequent use of decision strategies consistent with one-dimensional rules. However, many participants were also integrating brightness and saturation information.

In addition, in the positive condition of Experiment 1, there were a number of blocks in which decision strategies were more consistent with the optimal classifier (between 40° and 70°) than a one-dimensional strategy. A close inspection of the entire stimulus space suggested that there may have been an alternative "greyness" rule to which participants were sensitive. Inspection of Figure 3 and written descriptions of decision strategies, however, reveal no evidence that participants used a "greyness" strategy. Furthermore, no such bias was observed when feedback was provided in Experiment 2 (Figure 5). Instead, participants whose data were best fit by the linear classifier were more likely to use a strategy consistent with a less accurate one-dimensional rule than the so-called "greyness" rule.

The few available studies on free sorting of integral-dimension stimuli suggest that people prefer to use similarity-based decision strategies (e.g., Handel & Imai, 1972). Although testing between similarity-based and one-dimensional strategies is not the focus of this article, the similarity-based account would seem to predict that performance should be identical across the four category conditions. This is because all of our conditions are simply rotations of each other, and, according to most popular definitions, similarity among integral-dimension stimuli is invariant under rotation of the categories (Nosofsky, 1986).

Implications for models of category learning

Many computational models have been developed that are capable of category learning in the absence of feedback (Ahn & Medin, 1992; Anderson, 1991; Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Billman & Heit, 1988; Carpenter & Grossberg, 1991; Fried & Holyoak, 1984; Love, Medin, & Gureckis, 2004; Pothos & Chater, 2002). The majority of these models, however, predict that at least some aspects of performance on the Figure 1 category structures should be invariant with rotation of the categories, and even fewer make explicit predictions for integral versus separable dimensions. The pattern of results observed in the accuracy- and model-based analyses therefore provides a challenge to future model development.

For example, consider a recent network model of category learning proposed by Love and colleagues (Supervised and Unsupervised Stratified Adaptive Incremental Network, or SUSTAIN; Love, et al., 2004), which has been successfully applied to data from unsupervised tasks. In short, SUSTAIN is a multilayer neural network that maps stimulus representations to the appropriate responses via an intermediate layer of abstract category representations (or clusters). In unsupervised tasks, SUSTAIN assumes that the initial category representation comprises a single cluster and that additional clusters are created as highly dissimilar exemplars are encountered. SUSTAIN was applied to a broader class of unsupervised tasks than that considered here, but it does correctly predict that participants prefer one-dimensional strategies in unconstrained tasks, at least with separable dimensions. Similar to many other models of category learning, SUSTAIN is equipped with a selective attention learning mechanism that, together with a bias to attend to brightness over saturation, would allow it to capture the higher accuracy in the vertical condition in Experiment 1. This attentional bias, however, should be invariant with rotation of the categories, thereby predicting the frequent use of decision strategies on brightness in all conditions—a prediction that is inconsistent with the data.

Pothos and Chater's (2002) simplicity model is also particularly relevant to the issue of unsupervised categorization. Briefly, this model assumes that the preferred categorization strategy will be the simplest one (i.e., in an information-theoretic sense). A strategy's simplicity, or code length, is a function of the similarity structure of the stimuli and the costs and benefits of the constraints imposed by classifying the stimuli. The simplicity model accurately predicts a bias to use one-dimensional strategies over two-dimensional strategies in the absence of feedback and, as a result, the pattern of performance in the vertical condition (Pothos & Close, 2008). The simplicity model, however, would predict a similar bias for the horizontal condition and an approximately equal distribution of one- and two-dimensional strategies in the positive and negative conditions (Pothos & Close, 2008). Neither of these patterns was observed in the present data.

An alternative account is offered by the COVIS (Competition between Verbal and Implicit Systems) model of category learning (Ashby, et al., 1998). COVIS hypothesizes that category learning is a competition between separate hypothesis-testing and procedural-based systems. The hypothesis-testing system is specialized to learn abstract rules (e.g., one-dimensional rules), whereas the procedural-based system is specialized to learn stimulus–response mappings. Because learning in the procedural-based system is highly dependent upon feedback, and there is an initial bias towards the hypothesis-testing system, COVIS predicts that the hypothesis-testing system should dominate in unsupervised tasks. COVIS also predicts that the hypothesis-testing system will experiment with one-dimensional rules only if two conditions are met—that selective attention can be directed to this dimension and that a salient verbal label describes the dimension (e.g., brightness). These conditions are met with most separable dimensions, so COVIS correctly predicts that one-dimensional decision strategies should dominate in unsupervised tasks with separable dimensions (Ashby et al., 1999). Thus, COVIS correctly predicts the high prevalence of one-dimensional strategies on brightness in the vertical condition. COVIS, however, would not predict

that participants would be able to integrate brightness and saturation as evidenced by the high percentage of data sets best fit by the linear classifier in the horizontal, positive, and negative conditions.

SUMMARY

In summary, the present experiments demonstrate that individuals are capable of learning categories constructed from the integral dimensions of brightness and saturation in the absence of feedback. This ability, however, has several limitations. Consistent with the claim that integral dimensions are initially processed holistically (e.g., Kemler Nelson, 1993), participants had some success in conditions that required the integration of brightness and saturation. Consistent with the claim that integral dimensions can subsequently be processed in terms of the individual dimensions given the appropriate task demands (e.g., Garner, 1974), participants were able to learn when a one-dimensional strategy on brightness was highly accurate. In addition, participants demonstrated a general tendency to weight brightness more heavily than saturation across all three experiments, suggesting that brightness may have privileged status relative to saturation. Whether such a pattern holds when brightness is paired with other dimensions is a matter for future research.

Although current models can account for some aspects of these data, to our knowledge no current models can account for the pattern of strategy use observed in Experiment 1. In fairness, although SUSTAIN, the simplicity model, and COVIS all make predictions for unsupervised tasks, none of these models was designed to account for differences between integral and separable stimulus dimensions. Thus, it may be possible to augment these models to account for the present data (e.g., Pothos et al., 2011). Even so, because of the theoretical difficulties posed by our results, these data will provide an important benchmark for the development of theories of unsupervised category learning and may have implications for the application of cognitive science to education and training where constrained tasks are the norm (e.g., the training of medical

professionals to categorize predefined medical conditions).

REFERENCES

Ahn, W., & Medin, D. L. (1992). A two-stage model of category construction. *Cognitive Science*, *16*, 81–121.

Anderson, J. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*, 409–429.

Ashby, F. G. (1992a). Multidimensional models of categorization. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Ashby, F. G. (1992b). Multivariate probability distributions. In F. G Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 1–34). Hillsdale, NJ: Lawrence Erlbaum Associates.

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*, 442–481.

Ashby, F. G., & Gott, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *14*, 33–53.

Ashby, F. G., & Lee, W. W. (1993). Perceptual variability as a fundamental axiom of perceptual science. In S.C. Masin (Ed.), *Foundations of perceptual theory* (pp. 369–399). Amsterdam, The Netherlands: Elsevier.

Ashby, F. G., & Maddox, W. T. (1994). A response time theory of separability and integrality in speeded classification. *Journal of Mathematical Psychology*, *38*, 423–466.

Ashby, F. G., Queller, S., & Berretty, P. M. (1999). On the dominance of unidimensional rules in unsupervised categorization. *Perception & Psychophysics*, *61*, 1178–1199.

Ashby, F. G., & Townsend, J. T. (1986). Varieties of perceptual independence. *Psychological Review*, *93*, 154–179.

Billman, D., & Heit, E. (1988). Observational learning from internal feedback: A simulation of an adaptive learning method. *Cognitive Science*, *12*, 587–625.

Billman, D., & Knutson, J. (1996). Unsupervised concept learning and value systematicity: A complex whole aids learning the parts. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *22*, 458–475.

Boynton, R. M., & Olson, C. X. (1987). Locating basic colors in the OSA space. *Color Research and Application*, *12*(2), 94–105.

Brainard, D. H. (1997). Psychophysics software for use with MATLAB. *Spatial Vision*, *10*, 433–436.

Brainard, D. H. (2003). Color appearance and color difference specification. In S.K. Shevell (Ed.), *The science of color* (2nd ed., pp. 191–216). Washington, DC: Optical Society of America.

Carpenter, G. A., & Grossberg, S. (1991). *Pattern recognition by self-organizing neural networks*. Cambridge, MA: MIT Press.

Clapper, J. P., & Bower, G. H. (1994). Category invention in unsupervised learning. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *20*, 443–460.

Colreavy, E., & Lewandowsky, S. (2008). Strategy development and learning differences in supervised and unsupervised categorization. *Memory & Cognition*, *36*(4), 762–775.

Diaz, M., & Ross, B. H. (2006). Sorting out categories: Incremental learning of category structure. *Psychonomic Bulletin & Review*, *13*(2), 251–256.

Ell, S. W., & Ashby, F. G. (in press). The impact of category separation on unsupervised categorization. *Attention, Perception & Psychophysics*.

Foard, C. F., & Kemler-Nelson, D. G. (1984). Holistic and analytic modes of processing: The multiple determinants of perceptual analysis. *Journal of Experimental Psychology: General*, *113*, 137–159.

Fried, L. S., & Holyoak, K. J. (1984). Induction of category distributions: A framework for classification learning. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *10*, 234–257.

Garner, W. R. (1974). *The processing of information and structure*. New York, NY: Wiley.

Garner, W. R., & Felfoldy, G. L. (1970). Integrality of stimulus dimensions in various types of information processing. *Cognitive Psychology*, *1*, 225–241.

Grau, J. W., & Kemler-Nelson, D. G. (1988). The distinction between integral and separable dimensions: Evidence for the integrality of pitch and loudness. *Journal of Experimental Psychology: General*, *117*, 347–370.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York, NY: Wiley.

Handel, S., & Imai, S. (1972). The free classification of analyzable and unanalyzable stimuli. *Perception & Psychophysics*, *12*, 108–116.

Handel, S., Imai, S., & Spottswood, P. (1980). Dimensional, similarity, and configural classification of integral and separable stimuli. *Perception & Psychophysics*, *28*, 205–212.

Homa, D., & Cultice, J. (1984). Role of feedback, category size, and stimulus distortion on the acquisition and utilization of ill-defined categories. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *10*, 83–94.

Imai, S., & Garner, W. R. (1965). Discriminability and preference for attributes in free and constrained classification. *Journal of Experimental Psychology*, *69*, 596–608.

Kemler Nelson, D. G. (1993). Processing integral dimensions: The whole view. *Journal of Experimental Psychology: Human Perception & Performance*, *19*, 1105–1113.

Lockhead, G. R. (1972). Processing dimensional stimuli: A note. *Psychological Review*, *79*, 410–419.

Love, B. C. (2002). Comparing supervised and unsupervised category learning. *Psychonomic Bulletin & Review*, *9*(4), 829–835.

Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A network model of category learning. *Psychological Review*, *111*, 309–332.

Maddox, W. T. (2001). Separating perceptual processes from decisional processes in identification and categorization. *Perception & Psychophysics*, *63*, 1183–1200.

Maddox, W. T., & Ashby, F. G. (1993). Comparing decision bound and exemplar models of categorization. *Perception & Psychophysics*, *53*, 49–70.

Maddox, W. T., & Dodd, J. L. (2003). Separating perceptual and decisional attention processes in the identification and categorization of integral-dimension stimuli. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *29*, 467–480.

McKinley, S. C., & Nosofsky, R. M. (1996). Selective attention and the formation of linear decision boundaries. *Journal of Experimental Psychology: Human Perception & Performance*, *22*, 294–317.

Medin, D. L., Wattenmaker, W. D., & Hampson, S. E. (1987). Family resemblance, conceptual cohesiveness, and category construction. *Cognitive Psychology*, *19*, 242–279.

Melara, R. D., Marks, L. E., & Potts, B. C. (1993). Early-holistic processing or dimensional similarity? *Journal of Experimental Psychology: Human Perception & Performance*, *19*, 1114–1120.

Milton, F., Longmore, C. A., & Wills, A. J. (2008). Processes of overall similarity sorting in free classification. *Journal of Experimental Psychology: Human Perception and Performance*, *34*(3), 676–692, doi:10.1037/0096-1523.34.3.676

Milton, F., & Wills, A. J. (2004). The influence of stimulus properties on category construction. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *30*(2), 407–415. doi:10.1037/0278-7393.30.2.407

Mounts, J. R. W., & Melara, R. D. (1995). Classification of color dimensions in multiple contexts. *Journal of Experimental Psychology: Human Perception & Performance*, *21*, 257–274.

Munsell, A. H. (1915). *Atlas of the Munsell Color System*. Malden, MA: Wadsworth-Howland & Company.

Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.

Newhall, S. M. (1940). Preliminary report of the O.S.A. subcommittee on the spacing of the Munsell colors. *Journal of the Optical Society of America*, *30*, 617–645.

Newhall, S. M., Nickerson, D., & Judd, D. B. (1943). Final report of the O.S.A. subcommittee on the spacing of the Munsell colors. *Journal of the Optical Society of America*, *33*, 385–412.

Nickerson, D. (1940). History of the Munsell color system and its scientific application. *Journal of the Optical Society of America*, *30*, 575–586.

Nosofsky, R. M. (1986). Attention, similarity, and the identification categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–57.

Nosofsky, R. M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *13*, 87–108.

Nosofsky, R. M., & Palmeri, T. J. (1996). Learning to classify integral-dimension stimuli. *Psychonomic Bulletin & Review*, *3*, 222–226.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.

Pothos, E. M., & Chater, N. (2002). A simplicity principle in unsupervised human categorization. *Cognitive Science*, *26*, 303–343.

Pothos, E. M., & Chater, N. (2005). Unsupervised categorization and category learning. *Quarterly Journal of Experimental Psychology A*, *58*(4), 733–752.

Pothos, E. M., & Close, J. (2008). One or two dimensions in spontaneous classification: A simplicity approach. *Cognition*, *107*(2), 581–602. doi:10.1016/j.cognition.2007.11.007

Pothos, E. M., Edwards, D. J., & Perlman, A. (2011). Supervised vs. unsupervised categorization: Two

sides of the same coin? *Quarterly Journal of Experimental Psychology*, *64*, 1692–1713.

Pothos, E. M., Perlman, A., Bailey, T. M., Kurtz, K., Edwards, D. J., Hines, P., et al. (2011). Measuring category intuitiveness in unconstrained categorization tasks. *Cognition*, *121*(1), 83–100. doi:10.1016/j. cognition.2011.06.002.

Regehr, G., & Brooks, L. R. (1995). Category organization in free classification: The organizing effect of an array of stimuli. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *21*, 347–363.

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, *6*(2), 461–464.

Shepard, R. N. (1964). Attention and the metric structure of the stimulus space. *Journal of Mathematical Psychology*, *1*, 54–87.

Shepard, R. N., & Chang, J. (1963). Stimulus generalization in the learning of classifications. *Journal of Experimental Psychology*, *65*, 94–102.

Shin, H. J., & Nosofsky, R. M. (1992). Similarity-scaling studies of dot-pattern classification and recognition.

Journal of Experimental Psychology: General, *121*(3), 278–304.

Sturges, J., & Whitfield, T. W. (1995). Locating basic colours in the Munsell space. *Color Research and Application*, *20*, 364–376.

Torgerson, W. S. (1958). The nature of measurement. *Theory and methods of scaling*. New York, NY: Wiley.

Wickens, T. D. (1982). *Models for behavior: Stochastic processes in psychology*. San Francisco, CA: W. H. Freeman.

Wolfe, J. M. (2005). Guidance of visual search by preattentive information. In L. Itti, G. Rees, & J. Tsotsos (Eds.), *Neurobiology of attention*. San Diego, CA: Academic Press.

Wyszecki, G., & Stiles, W. S. (1982). *Color science: Concepts and methods, quantitative data and formulae* (2nd ed.). New York, NY: John Wiley & Sons.

Zeithamova, D., & Maddox, W. T. (2009). Learning mode and exemplar sequencing in unsupervised category learning. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *35*(3), 731–741 doi:10.1037/a0015005

APPENDIX

Description of decision bound models and fitting procedure

To get a more detailed description of how participants categorized the stimuli, a number of different decision bound models (Ashby, 1992a; Maddox & Ashby, 1993) were fitted separately to the data for each participant from every block. Decision bound models are derived from general recognition theory (Ashby & Townsend, 1986), a multivariate generalization of signal detection theory (Green & Swets, 1966). It is assumed that, on each trial, the percept can be represented as a point in a multidimensional psychological space and that each participant constructs a decision bound to partition the perceptual space into response regions. The participant determines which region the percept is in and then makes the corresponding response. While this decision strategy is deterministic, decision bound models predict probabilistic responding because of trial-by-trial perceptual and criterial noise (Ashby & Lee, 1993).

The appendix briefly describes the decision bound models. For more details, see Ashby (1992a) or Maddox and Ashby (1993).

Unidimensional classifier

This model assumes that the stimulus space is partitioned into two regions by setting a criterion on one of the stimulus dimensions. Two versions of the unidimensional classifier (UC) were fitted to these data: One assumed that participants attended selectively to brightness (UC-B), and the other assumed that participants attended selectively to saturation (UC-S). The unidimensional classifier has two free parameters: a decision criterion on the relevant perceptual dimension and the variance of internal (perceptual and criterial) noise (i.e., σ^2). In the vertical and horizontal conditions, a special case of the unidimensional classifier, the optimal unidimensional classifier (OC), assumes that participants use the unidimensional decision bound that maximizes accuracy (Figure 1). This special case has one free parameter (σ^2)

Conjunctive classifier

Another possibility is that the participant uses a conjunction rule involving separate decisions about the stimulus value on the two dimensions, with the response assignment based on the outcome of these two decisions (Ashby & Gott, 1988). The conjunctive classifier (CC) assumes that the participant partitions the stimulus space into four regions. Based upon inspection of the data from the individual participants, two versions of the CC were fitted to these data. The first assumed that individuals assigned a stimulus to Category A if it was high on brightness and low on saturation; otherwise the stimulus was assigned to Category B. The second assumed that a stimulus was assigned to Category A if it was low on brightness and low on saturation; otherwise the stimulus was assigned to Category B. The CC has three free parameters: the decision criteria on the two dimensions and a common value of σ^2 for the two dimensions.

Linear classifier

This model assumes that a linear decision bound partitions the stimulus space into two regions. The linear classifier (LC) differs from the CC in that the LC does not assume decisional selective attention (Ashby & Townsend, 1986). Instead, the LC requires integration of the perceived values on the stimulus dimensions. The LC has three parameters: slope and intercept of the linear bound, and σ^2. In the positive and negative conditions, a special case of the LC, the optimal linear classifier (OC), assumes that participants use the linear decision bound that maximizes accuracy (Figure 1). This special case has one free parameter (σ^2).

Model fitting

The model parameters were estimated using maximum likelihood (Ashby, 1992b; Wickens, 1982), and the goodness-of-fit statistic was

$$\text{BIC} = r \ln N - 2 \ln L,$$

where BIC is the Bayesian information criterion, N is the sample size, r is the number of free parameters, and L is the likelihood of the model given the data (Schwarz, 1978). The BIC statistic penalizes a model for poor fit and for extra free parameters. To find the best model among a set of competitors, one simply computes a BIC value for each model and then chooses the model with the smallest BIC. To assess the absolute fit of the models, the percentage of responses accounted for by the best fitting model was computed for each data set. This statistic ranges from 0% to 100% with the latter implying that the model perfectly accounted for all of the participant's responses.