

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Stimulus Generalization in Category Learning

Permalink

<https://escholarship.org/uc/item/80g257zb>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 27(27)

ISSN

1069-7977

Authors

Jones, Matt
Love, Bradley C.
Maddox, W. Todd

Publication Date

2005

Peer reviewed

Stimulus Generalization in Category Learning

Matt Jones, W. Todd Maddox, and Bradley C. Love

[mattj,maddox,love]@psy.utexas.edu

University of Texas, Department of Psychology, 1 University Station A8000
Austin, TX 78712 USA

Abstract

Stimulus generalization is often regarded as a fundamental component of category learning, yet it has not been directly studied in this context. Here we develop a technique for measuring generalization based on sequential effects in subjects' responses. We find that patterns of generalization can adapt to global properties of the task, but only when the category structure is defined by perceptually primitive and separable dimensions. Implications are discussed for attentional learning and the nature of both perceptual and category representations.

Introduction

Perhaps the most fundamental task facing the brain is to use past experience to determine useful behavior in novel situations. For example, in deciding whether a particular snake is poisonous, one might draw on knowledge of other, similar snakes whose toxicity was known. The details of this process can be critical: Basing one's response on other snakes of similar color and markings may be effective, but relying on irrelevant properties such as length could have disastrous consequences. In other words, successful generalization depends critically on knowledge of which variables are relevant to the current prediction.

One task in which stimulus generalization is believed to play an important role is category learning (Medin & Schaffer, 1978). However, in contrast to the rich body of data in conditioning (see Shepard, 1987), stimulus generalization in category learning has yet to be directly investigated. Often it is assumed that generalization operates the same in these two tasks, and generalization functions that have been empirically supported in conditioning studies are incorporated into the similarity functions of categorization models (Kruschke, 1992; Love, Medin, & Gureckis, 2004; Nosofsky, 1986). However, the richer nature of representations involved in category learning (e.g., Maddox & Ashby, 1993; Rosch et al., 1976; Sloman, Love, & Ahn, 1998) suggests that generalization in this domain may be far more complex than is currently assumed.

The primary aim of this paper is to develop and explore a method for directly assessing stimulus generalization in category learning. The technique, described in more detail below, is based on a close connection between generalization and recency effects (Jones & Sieck, 2003). Here we present two experiments designed to validate this approach and to relate it to previous findings on attentional learning. Our results show good support for the approach and illustrate how it can provide insight into perceptual

representations and the distinction between integral and separable dimensions. We conclude by discussing the broader applicability of this new methodology as well as its implications for the nature of perceptual and category representations, attentional learning, and the roles of short- and long-term memory in categorization.

Recency effects and stimulus generalization

Recency effects are a robust phenomenon in repeated judgment tasks. For example, in studies of probability learning (repeated uncued forced-choice tasks), it has been regularly found that subjects are biased to select whichever response was reinforced on the previous trial (see Myers, 1970, for a review). Jones and Sieck (2003) found that this same effect occurs in cued categorization: Once the identity of the current stimulus is controlled for, subjects tend to choose the category that was correct on the previous trial. This marginal effect of learning from the previous trial can be interpreted as generalization from one stimulus to the next, because it reflects the belief that the current stimulus is likely to belong to the same category as the previous stimulus. Consistent with this interpretation, Jones and Sieck found that the magnitude of the recency effect depends on the similarity between the present and previous stimuli, as shown in Figure 1. Stimuli in these experiments were hypothetical medical patients varying in the presence or absence of three symptoms. The recency effect was greatest when successive stimuli were identical and decreased with each cue mismatch, fully disappearing for cases of complete mismatch. The approximately exponential decrease is similar to the functional form of generalization commonly found in conditioning (Shepard, 1987).




Figure 1: Recency effects as a function of number of mismatching cues between present and previous stimuli. (From Jones and Sieck, 2003, Expt. 2, control condition.)

This phenomenon offers a potentially powerful tool for measuring stimulus generalization during category learning. The basic idea is to measure generalization from the previous stimulus by measuring the influence of that stimulus' category membership on the current response. For example, generalization from stimulus X to Y can be defined as the difference in category A responses between trials on which Y follows X and X was in category A and trials on which Y follows X and X was in category B.¹ By determining how this generalization effect depends on the relationship between pairs of stimuli, we can gain important information about the nature of the representations underlying categorization.

To be clear, we do not mean to claim that recency effects and stimulus generalization are the same thing. Presumably stimulus generalization occurs from many or all previous stimuli, but this generalization is far stronger for the stimulus presented most recently. This latter fact is what is meant by the recency effect. The existence of the recency effect is in fact irrelevant to the theoretical issues addressed in this paper. However, it is critical to the practicality of the empirical investigation, as it causes information from the previous trial to account for a large proportion of the variance in subjects' responses, thus allowing for statistically reliable estimates of generalization behavior.

Measuring stimulus generalization in category learning

The present study aims to extend the above findings to a more detailed investigation of stimulus generalization in category learning. We present the results of two experiments designed to verify the viability of the approach by testing hypotheses about how generalization changes with learning. In the concluding section we describe ongoing research using our technique to address a range of other issues.

One important issue in studies of category learning is selective attention. A number of models assume that the similarity metric underlying generalization can adapt, such that certain dimensions receive more weight than others (Kruschke, 1992; Love et al., 2004; Nosofsky, 1986). The standard prediction (e.g., Nosofsky, 1986) is that attention will shift to those dimensions that are most predictive of the category outcome. This implies that the generalization gradient for these dimensions will be sharper; that is, generalization will be weaker between stimuli that differ on a diagnostic dimension as compared to an irrelevant dimension. This adaptive generalization effect makes sense from a normative standpoint, as illustrated by the introductory example. However, empirically it is not entirely clear when adaptive generalization should be expected to occur, and approaches based on fits of the aforementioned models have failed to yield consistent

conclusions (Maddox & Ashby, 1998). The present experiments address this issue using the recency effect-based technique for measuring generalization.

Experiment 1

Experiment 1 investigates stimulus generalization during category learning, and in particular how the pattern of generalization changes with learning. Stimuli were visual images that varied along two continuous and separable dimensions. Three category structures were used: two in which only one stimulus dimension was relevant, and a third in which both dimensions combined additively to predict the outcome (Fig. 2A-C). The principle questions were whether similarity-based generalization occurs with these continuous stimuli, and if so whether generalization adapts to the category structure.

Our primary hypothesis regarding adaptive generalization was that subjects in the unidimensional conditions would weight the diagnostic dimension more heavily, so that generalization between stimuli would be selectively sensitized to discrepancies on this dimension. The prediction for the integration condition was less certain. One possibility was that there would be no effect on generalization because both dimensions must be attended to. This is the prediction made by most attentional learning models, which assume that input dimensions are processed separately each with its own attention weight. However, a second possibility was that subjects would learn to selectively attend to the diagonal dimension; that is, generalization would adapt relative to the category structure just as in the unidimensional conditions. Thus a comparison between the two types of category structures allows a test of how closely generalization is tied to perceptual representation.

Method

Participants. Sixty-five members of the University of Texas, Austin, participated for payment or course credit.

Stimuli. Stimuli were 6-cm square Gabor patterns (sine-wave gratings within a Gaussian envelope), varying in the frequency and orientation of the grating. There were 100 stimuli present in each condition, arranged in a 10×10 grid in stimulus space.

Design. Participants were randomly assigned to one of three conditions. In the Frequency (F) and Orientation (O) conditions, category outcomes depended only on frequency or orientation, respectively. In the Integration (I) condition, both frequency and orientation were predictive of category membership. More precisely, the probability that a stimulus S would belong to category A on any particular presentation was given by $P[S \in A] = [1 + e^{-\sigma f(S)}]^{-1}$, with $f(S)$ defined by frequency (condition F), orientation (O), or the difference (frequency – orientation)/ $\sqrt{2}$ (I). In computing this probability, the two stimulus dimensions were parameterized so as to have equal ranges centered on 0 (between ± 4.5 in conditions F and O and $\pm 4.5\sqrt{2}$ in condition I). The scaling parameter σ was set such that

¹ Note that this approach requires a probabilistic category structure, i.e. one in which every stimulus appears with some non-zero frequency in every category.

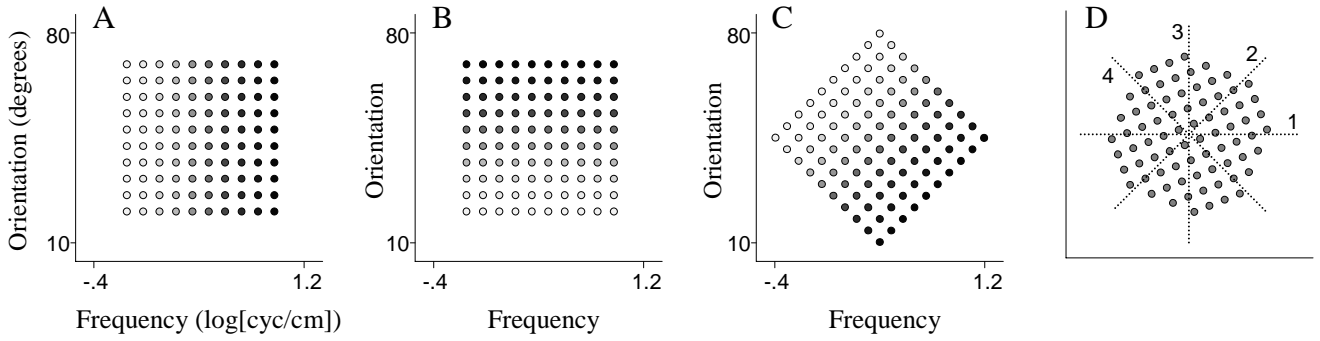


Figure 2: Category structures for both experiments. Panels A-C show Experiment 1 conditions F (frequency), O (orientation), and I (integration), respectively. Each circle represents a stimulus, with shading indicating probability of membership in category A. Panel D shows Experiment 2. Dotted lines indicate category bounds for conditions 1 through 4 (these bounds are probabilistic, as in Panels A-C). Axes are rectangular coordinates derived from the Hue-Chroma polar coordinates of the Munsell system.

$P[S \in A]$ ranged between .05 and .95. The three category structures are illustrated in Figure 2A-C.

Procedure. On each trial, a stimulus was randomly selected and presented in the center of a 43-cm computer monitor. The subject responded by pressing one of two keys on the keyboard. The correct answer for that trial was sampled according to the formula given above and displayed below the stimulus for 1s. Subjects were not informed of the category structure in advance. The experiment consisted of 600 trials.

Analysis

The data from each subject were analyzed separately to obtain measures of long-term cue use as well as patterns of generalization. Specifically, we fit each subject's responses to a logistic regression model with predictors given by the feature values of the current and previous stimuli along with a term representing generalization:

$$\text{logodds}(R_n) \approx \beta_0 + \sum_i w_i S_{n,i} + \sum_i v_i S_{n-1,i} + C_{n-1} \Gamma_{n,n-1}. \quad (1)$$

Here R_n represents the response on trial n ; $S_{n,i}$ gives the value of stimulus n on dimension i ; $\Gamma_{n,n-1}$ is the strength of generalization from S_{n-1} to S_n ; and C_{n-1} is the correct category on trial $n-1$, coded as +1 for A and -1 for B (so that C_{n-1} determines the direction of the generalization effect). The remaining variables are regression weights. The critical variables in this model are w_i , which gives the strength of association from cue i to category A, and $\Gamma_{n,n-1}$, which is discussed below. The purpose of the model is to allow measurement of generalization effects between successive stimuli as a function of their similarity, while controlling for the identity of each stimulus.²

² The functional form of the terms involving S_n and S_{n-1} is intended to be agnostic about the nature of long-term stimulus-category associations. The additive form used here matches that of the category structures, i.e., all 3 conditions can be written as $\text{logodds}(C_n) = \sum_i W_i S_{n,i}$. Controlling for the previous stimulus is necessary due to the possibility of perceptual contrast effects.

Three forms were investigated for the generalization function:

$$\Gamma_{n,n-1} = ke^{-\sum \alpha_i |S_{n,i} - S_{n-1,i}|} \quad (2A)$$

$$\Gamma_{n,n-1} = ke^{-\sum \alpha_i (S_{n,i} - S_{n-1,i})^2} \quad (2B)$$

$$\Gamma_{n,n-1} = k - \sum \alpha_i |S_{n,i} - S_{n-1,i}| \quad (2C)$$

Each of these functions has a value given by k when S_n and S_{n-1} are identical and decreases with increasing dissimilarity between these stimuli. The first corresponds to an exponential function of inter-stimulus distance with distance given by a city-block metric. The second corresponds to a Gaussian function of distance with distance given by a Euclidean metric. Both of these generalization functions have been proposed previously (e.g., Nosofsky, 1986; Shepard, 1987). The third version is a linear function that we consider for sake of generality and because it does not assume a priori that generalization decreases with distance (i.e., the α parameters were allowed to be negative in this model but not in the other two). Because the Gaussian-Euclidean model provided significantly better fits than the exponential-city-block and linear models, the results presented here are based on that model. The other two lead to the same conclusions.

A further property of the generalization functions in Equation 2 is that each stimulus dimension is weighted by the corresponding α parameter. Larger values of α correspond to increased attention and steeper generalization gradients. Thus comparing estimated values of α_1 and α_2 gives a measure of selective generalization. For this purpose a generalization bias parameter, β , is defined as

$$\beta = \frac{\alpha_1}{\alpha_1 + \alpha_2}. \quad (3)$$

This variable measures the relative influence of the two dimensions in determining strength of generalization, and is constrained to lie between 0 and 1 (when $\alpha_1, \alpha_2 \geq 0$).

Table 1: Primary measures for Experiment 1

Condition	$w_{\text{frequency}}$	$w_{\text{orientation}}$	k	β	Performance
F	.735	-.056	1.419	.712	69.0%
O	.024	.377	1.629	.444	61.1
I	.474	-.304	1.070	.652	64.7

Notes: Condition F is frequency-relevant; O is orientation-relevant; I is integration (both relevant). Values of k are medians because of skew; all other values are means.

Results

The generalization model given by Equations 1 and 2B was fit to each subject’s data, with frequency and orientation represented on a common scale as described above. Mean values for primary measures are displayed in Table 1.

Recency effects and similarity-based generalization. The baseline strength of the recency effect, given by the parameter k , was positive for every individual subject. Thus the recency effect is quite robust in this task. To test whether the recency effect declined with stimulus dissimilarity, values of $\alpha_{\text{frequency}}$ and $\alpha_{\text{orientation}}$ were examined from the linear model (Eq. 2C; the Gaussian model is inappropriate for this question because it assumes a negatively sloped generalization function a priori). Estimates were positive for 61 of 65 subjects for $\alpha_{\text{frequency}}$ and 55 subjects for $\alpha_{\text{orientation}}$. Wilcoxon signed-ranks tests (used because both distributions were heavy-tailed) showed both effects to be highly significant, $ps < 10^{-6}$. Therefore generalization depends positively on stimulus similarity.

Selective generalization. Mean values of the generalization bias parameter β indicate that generalization in both unidimensional conditions shifted to depend more heavily on the task-relevant dimension (see Table 1).³ The difference among conditions was confirmed by analysis of variance, $F(2,62) = 3.72, p < .05$. A planned comparison contrasting conditions F and O was also significant, $t(62) = 2.59, p < .02$. Therefore generalization patterns were reliably affected by the category structure.

Figure 3 illustrates this selective generalization effect. Shown are the average generalization functions for the diagnostic and irrelevant dimensions, based on combined data from both unidimensional conditions. Curves are based on median values of k , $\alpha_{\text{diagnostic}}$, and $\alpha_{\text{irrelevant}}$, with $\alpha_{\text{diagnostic}}$ equal to $\alpha_{\text{frequency}}$ for condition F and $\alpha_{\text{orientation}}$ for O; $\alpha_{\text{irrelevant}}$ is defined similarly. The graph shows how generalization drops more rapidly with deviations along the diagnostic as compared to the irrelevant dimension.

Selective generalization and long-term cue use. This analysis addressed whether selective generalization is

³ The overall bias towards frequency is just a scaling effect presumably due to greater salience of this dimension given the amount of variation present in this experiment. This salience difference also explains the ordering of performance in the three conditions.




Figure 3: Average generalization curves for relevant and irrelevant dimensions in Experiment 1 (conditions F and O combined). Distance refers to the difference between successive stimuli on the dimension in question.

learned directly or is based on the strength of cue-category associations. A decisional attention measure, analogous to the generalization bias β , was computed for each subject as

$$\gamma = \frac{|w_{\text{frequency}}|}{|w_{\text{frequency}}| + |w_{\text{orientation}}|}. \quad (4)$$

This parameter measures the relative strengths of long-term cue-category associations and is constrained to lie between 0 and 1. Next the ANOVA comparing β across conditions was re-run with γ included as a covariate. The effect of condition remained significant, $F(2,59) = 4.69, p < .05$. The effect of γ was also significant, partial $r = .468, F(1,59) = 10.26, p < .01$. The interaction was nonsignificant, $F(1,59) = .26$. Therefore adaptive generalization is mediated by both the true category structure and actual learning of cue-category associations.

Diagonal selective generalization. In condition I, the “diagonal” dimension $d^- = (\text{frequency} - \text{orientation})/\sqrt{2}$ is maximally diagnostic of category membership and the orthogonal dimension, $d^+ = (\text{frequency} + \text{orientation})/\sqrt{2}$, is irrelevant. Therefore the recency-generalization model was refit to the condition I data using d^- and d^+ in place of frequency and orientation. Analyses based on this model are formally equivalent to analyses presented above for the unidimensional conditions, with the entire design rotated by 45 degrees in stimulus space.

Under the (d^-, d^+) coordinate system there are no scaling concerns, because the two dimensions necessarily have the same perceptual scale (even if frequency and orientation do not). Therefore values of β can be directly compared to .5. The mean value of β obtained under this model was .473 (which is in the direction opposite of that predicted by adaptive generalization) and was not significantly different from .5, $t(22) = .418$. Thus subjects appear unable to adapt their generalization patterns to the diagonal category bound.

Discussion

Recency effects in this experiment were robust and declined with dissimilarity between successive stimuli, consistent with previous findings on stimulus generalization. In addition, comparison of generalization patterns across

conditions showed clear effects of category structure. Specifically, generalization in each unidimensional condition was selectively dependent on the task-relevant dimension. This adaptation effect appears to be due to both the objective category structure and subjects' learning of that structure.

In contrast to the unidimensional conditions, the integration condition showed no evidence of adaptive generalization. When generalization was measured with respect to the diagnostic and irrelevant diagonal dimensions, no difference in the weighting of these two dimensions was found. Therefore it appears that stimulus generalization can adapt to the structure of a categorization task, but that this adaptation is constrained by the nature of the perceptual representations involved.

Experiment 2

The fact that stimulus generalization can become sensitized to primitive perceptual dimensions but not arbitrary combinations of these dimensions suggests a close connection between adaptive generalization and selective attention. Therefore Experiment 2 investigated generalization with stimuli defined by integral dimensions, in which selective attention is known to be difficult (Garner, 1974). Specifically, stimuli in Experiment 2 were color patches varying in hue and saturation. The prediction was that, in contrast to the findings of Experiment 1, subjects would be unable to adapt their generalization behavior so as to selectively attend to either of these dimensions.

Methods

Participants. Sixty members of the University of Texas, Austin, participated for payment or course credit.

Stimuli. Stimuli were 5-cm circular color patches. The same 76 stimuli were used in all conditions. These colors formed a regular grid in Munsell color space under the rectangular coordinate system derived from the polar coordinates of Hue and Chroma (saturation), as depicted in Figure 2D. Hue ranged from 4.3RP to 1.4R and Chroma from 12.9 to 20.5; Value (luminance) was constant at 7.

Design. Participants were randomly assigned to one of four conditions. The category structure for each condition was defined analogously to the structures in Experiment 1, with outcome probabilities for the individual stimuli again ranging from 5 to 95%. Orientations of the four category structures were all separated by 45 degrees, with each bound offset by 22.5 degrees from the stimulus grid (see Fig. 2D).

Procedure. The procedure mirrored that of Experiment 1 and consisted of 500 trials.

Results

Data were again analyzed by fitting the generalization model (Eqs. 1 & 2B) to each subject's data. Because there are no canonical perceptual axes for color space, the model for each subject was fit using the diagnostic and irrelevant dimensions for that subject's category structure.

The recency-effect parameter k was positive for 55 of the 60 subjects, indicating a robust recency effect. Mean values of α obtained from the linear version of the model were significantly negative for both the diagnostic and irrelevant dimensions (Wilcoxon signed-ranks test, $ps < 10^{-9}$).

Because the model was fit using the category-specific axes, the adaptive generalization hypothesis predicts a mean value of β greater than .5. Contrary to this prediction, the mean β was .430, with the difference from .5 non-significant, $t(59) = 1.64$, $p > .1$. A more direct test of adaptive generalization was obtained by comparing pairs of conditions with orthogonal category structures (1 vs. 3 and 2 vs. 4). The models for these conditions were based on the same axes with their labels reversed; thus a direct contrast of β between groups was obtained by subtracting one group's values from 1 (e.g., β for condition 1 was compared to $1-\beta$ for condition 3). This contrast is the same as that performed in Experiment 1 between conditions F and O, which provided the primary evidence for adaptive generalization in that experiment. In the present experiment, both contrasts were in the direction opposite of that predicted, and neither was significantly different from zero: $t(28) = 1.57$, $p > .1$ for conditions 1 vs. 3; $t(28) = .77$, $p > .4$ for conditions 2 vs. 4. Therefore generalization appears to have been unaffected by category structure.

A final analysis compared generalization to long-term cue use, as defined by the decisional attention parameter γ (Eq. 4). The correlation between γ and β across subjects was .074, which is nonsignificant, $p > .5$.

Discussion

Subjects in Experiment 2 exhibited recency effects and similarity-dependent generalization comparable to what was found in Experiment 1. Average performance was also matched (64.9% in Experiment 1, 64.3% in Experiment 2). However, this time there was no evidence for adaptive generalization. Analysis of weights in the similarity metric showed no effect of either objective category structure or actual cue use, both of which were seen to have significant effects in Experiment 1. Our use of four category structures all varying by 45 degrees eliminates the possibility that selective generalization is possible along some unspecified perceptual axes. Whatever these axes might be, they would have to be within 22.5 degrees of one of the structures used here, in which case that condition should have exhibited some degree of selective generalization. Therefore it appears that for the integral dimensions of hue and saturation people are unable to selectively attend to any one dimension for the purposes of adaptive generalization.

General Discussion

Stimulus generalization has long been acknowledged as an important component of category learning, but has not previously been studied directly. The present experiments demonstrate how variability in sequential effects can be used to obtain a straightforward measure of generalization from one stimulus to the next. Subjects' tendency to extend

the previous category to the current case was seen to decrease as a function of the dissimilarity between successive stimuli, in a manner similar to established findings in studies of conditioning (Shepard, 1987). Furthermore, generalization was seen to adapt to the category structure, such that task-relevant features were weighted more heavily. However, when the category discrimination was based on combining information from multiple dimensions or filtering one integral dimension from another, selective generalization did not occur.

The pattern of generalization behavior found here is largely in agreement with many models of category learning, including the GCM (Nosofsky, 1986), ALCOVE (Kruschke, 1992), and SUSTAIN (Love et al., 2004). This correspondence helps to validate our approach and also provides empirical support for these models. However, the present results bear only on generalization from the most recent stimulus. Considering categorization tasks in light of the distinction between short- and long-term memory raises the possibility that information from trials further back is used in a qualitatively different way. Therefore questions about the nature of long-term category representations (e.g., exemplars, prototypes, decision bounds, clusters) might be more fruitfully answered if the contributions of short-term memory (i.e., stimulus generalization from recent trials) were also taken into account. We are currently working to expand the present methodology to address this issue.

Another important question raised by our results concerns the learning process underlying adaptive generalization. Process models of attentional learning such as ALCOVE and SUSTAIN assume that attention weights are learned via error correction, and thus depend directly on the category structure. The results of Experiment 1 partially support this hypothesis, but they also indicate that attentional learning is related to learning of cue-category associations. Unfortunately the relation found here is only correlational and could be due to some third variable such as individual variation in relative feature salience. Therefore a more direct test is needed to determine the relative contributions of direct learning versus perceived feature relevance to adaptation of generalization. Research currently in progress that manipulates sequential dependencies between cues and categories will hopefully shed light on this question.

A final conclusion from this study regards the close relationship between adaptive generalization and perceptual representation. It might have been reasonable to expect that subjects in the integration condition of Experiment 1 would recognize when changes in frequency and orientation were mutually reinforcing versus offsetting and generalize more in the latter case. The fact that generalization did not adapt in this way implies that generalization is closely tied to and constrained by the nature of perceptual representations.

Adaptive generalization thus provides a useful indicator of what types of stimulus variation constitute primitive perceptual dimensions. Under this interpretation, Experiment 2 suggests viewing integral dimensions as single dimensions from a perceptual standpoint, even

though physically they have multiple degrees of freedom. It also seems possible that some spaces, for example those defined by multiple dimensions of spatial location, admit selective generalization in any direction. We are currently evaluating this possibility. Through this sort of systematic investigation we hope to develop a better theory of the varieties of perceptual spaces, and address one of the more fundamental questions in cognitive science of what constitutes a perceptual dimension.

Acknowledgments

This work was supported by NRSA F32-MH068965 from the NIH to MJ, NIH Grant R01-MH59196 to WTM, and AFOSR Grant FA9550-04-1-0226 and NSF CAREER Grant 0349101 to BCL.

References

- Garner, W. R. (1974). *The processing of information and structure*. Oxford, England: Lawrence Erlbaum.
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, *117*, 227–247.
- Jones, M., & Sieck, W. R. (2003). Learning myopia: An adaptive recency effect in category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *29*, 626–640.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, *99*, 22–44.
- 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., & Ashby, F.G. (1993). Comparing decision bound and exemplar models of categorization. *Perception & Psychophysics*, *53*, 49–70.
- Maddox, W.T., & Ashby, F.G., (1998). Selective attention and the formation of linear decision bounds. Commentary on McKinley and Nosofsky (1996). *Journal of Experimental Psychology: Learning, Memory and Cognition*, *24*, 301–321.
- Medin, D. L., & Schaffer, M.M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207–238.
- Myers, J. L. (1970). Sequential choice behavior. In G. H. Bower & J. T. Spence (Eds.), *The psychology of learning and motivation: Vol. 4*. New York: Academic Press.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–57.
- Rosch, E., Mervis, C.B., Gray, W.D., Johnson, D.M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, *8*, 382–349.
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*, 1317–1323.
- Solman, S. A., Love, B. C., & Ahn, W. K. (1998). Feature Centrality and Conceptual Coherence. *Cognitive Science*, *22*, 189–228.