Further tests of an exemplar-similarity approach to relating identification and categorization

ROBERT M. NOSOFSKY
Indiana University, Bloomington, Indiana

Further tests were provided of an exemplar-similarity model for relating the identification and categorization of separable-dimension stimuli (Nosofsky, 1986). On the basis of confusion errors in an identification paradigm, a multidimensional scaling (MDS) solution was derived for a set of 16 separable-dimension stimuli. This MDS solution was then used in conjunction with the exemplar-similarity model to accurately predict performance in four separate categorization paradigms with the same stimuli. A key to achieving the accurate quantitative fits was the assumption that a selective attention process systematically modifies similarities among exemplars across different category structures. The tests reported go well beyond earlier ones (Nosofsky, 1986) in demonstrating the generalizability and utility of the theoretical approach. Implications of the results for alternative quantitative models of classification performance, including Ashby and Perrin's (1988) general recognition theory, were also considered.

This article is concerned with the relation between the identification and categorization of multidimensional perceptual stimuli. *Identification* refers to a choice experiment in which there are n stimuli, with each stimulus assigned a unique response; whereas in *categorization*, the n stimuli are partitioned into $m < n$ groups, with each group assigned a unique response.

Nosofsky (1986, 1987) proposed a unified quantitative approach to modeling identification and categorization, which integrated and extended classic theories in the areas of choice and similarity (Carroll & Wish, 1974; Luce, 1963; Medin & Schaffer, 1978; Shepard, 1957, 1958, 1987; Shepard & Chang, 1963; Shepard, Hovland, & Jenkins, 1961). The approach is based on the idea that people represent categories by storing individual exemplars in memory, and classify items according to their similarity to these stored exemplars.

On the basis of confusion errors in an identification paradigm, a multidimensional scaling (MDS) solution for the exemplars is derived. This MDS solution is the one that yields a maximum-likelihood fit to the identification data within the framework of the similarity choice model (Luce, 1963; Nosofsky, 1985b; Shepard, 1958; Smith, 1980; Townsend & Landon, 1982). The framework then allows for the prediction of confusion errors in any given categorization paradigm involving the same stimuli. Similarities between exemplars are computed from the derived MDS solution, with the assumption that a selective attention process may systematically modify distances in the psychological space (Carroll & Wish, 1974). This

notion of selective attention intervening between identification and categorization appears critical for the connection of performance in the two paradigms (Medin & Schaffer, 1978; Nosofsky, 1986, 1987; Shepard et al., 1961).

The present study is concerned with the identification and categorization of highly analyzable stimuli varying along *separable* dimensions, for which selective attention processes are known to play a critical role (Garner, 1974; Shepard, 1964). Nosofsky (1985b, 1986) reported preliminary tests in which the exemplar-similarity model was used to account for the categorization of separable-dimension stimuli, but they were limited in important respects. First, the identification and categorization performance of only two subjects was tested: each subject received extensive experience with the entire ensemble of exemplars (over 12,000 presentation trials). One purpose of the present study was to test the generalizability of the approach by running large groups of subjects, with each subject receiving relatively little training.

More critically, the previous design required Nosofsky (1986) to posit that a "memory-augmentation" process had taken place, in which people added inferred exemplars to their category representations. Specifically, the category structures that were tested included transfer stimuli, which had no experimentally defined correct answers associated with them. The transfer stimuli were included in order to discover how people would generalize from the original training exemplars. Because only two subjects were tested, however, it was necessary to present the transfer stimuli repeatedly during the course of the experiment (in order to obtain a large enough sample size to secure reliable probability estimates). It is reasonable to posit that the subjects' initial decisions about category membership for the transfer stimuli may have influenced subsequent ones. Indeed, Nosofsky's (1986) theoretical

analyses suggested that the subjects augmented their category representations with these repeated transfer stimuli, and based classification decisions on the similarity of probes to the transfer stimuli as well as to the original training exemplars.

Although the memory-augmentation process was viewed as a natural consequence of Nosofsky's (1986) experimental conditions, the need to incorporate it reduced dramatically the power of the theoretical approach. The present design obviates the positing of the memory-augmentation process, because subjects are given only two presentations of each transfer stimulus near the end of the experiment. Thus, the design allows for a more rigorous test of this theory proposed for relating identification and categorization.

EXPERIMENT

In this experiment, the subjects identified or categorized a set of perceptually confusable semicircles varying in size and angle of orientation of a radial line. Previous research indicates that these dimensions are separable (Garner & Felfoldy, 1970; Shepard, 1964). There were four levels of size and four levels of angle, combined orthogonally to yield a 16-member stimulus set. In the identification condition, the subjects simply identified the level of size and level of angle of each stimulus. Four categorization conditions were tested, which are illustrated in Figure 1. In these grids, the columns represent levels of angle and the rows represent levels of size. Cells with a 1 or a 2

Figure 1. Schematic illustration of the four category structures. In each grid, the rows correspond to levels of size and the columns correspond to levels of angle. The cells with a 1 or 2 represent Category 1 or 2 training exemplars. The empty cells represent transfer stimuli.

represent stimuli assigned to Category 1 or 2, whereas empty cells represent unassigned transfer stimuli. Each category structure can be described in terms of a simple rule. In the *size* categorization, small (or large) stimuli are assigned to Category 1 (or 2). Likewise, in the *angle* categorization, stimuli with low (or high) angles are assigned to Category 1 (or 2). The *crisscross* categorization is a biconditional structure: small stimuli with high angles and large stimuli with low angles are assigned to Category 1, and the reverse for Category 2. The fourth condition is called the *diagonal* categorization because the categories can be partitioned by drawing a diagonal line through the stimulus space. Another way of describing the diagonal categorization is to say that it is a "rule-plus-exception" structure. For example, stimuli with high (or low) angles are assigned to Category 1 (or 2), with one exception in each category. The theoretical goal is to predict quantitatively the confusion data observed in the identification condition and all four categorization conditions, using the exemplar-similarity model.

The asymmetric structures associated with the size and angle categorizations (see Figure 1) were used to provide clear evidence of the need to incorporate selective attention components within the framework of the exemplar-similarity model. Consider, for example, the angle categorization. Stimulus 14 is highly similar to a Category 2 training exemplar, namely, Stimulus 15. Thus, an "overall-similarity" exemplar model predicts that Transfer Stimulus 14 will be classified in Category 2. However, angle is the relevant dimension in the angle categorization, and it is expected that subjects will attend selectively to the relevant dimension. Selective attention is represented in the present exemplar model in terms of "stretching" of distances along the attended dimensions, and "shrinking" of distances along the unattended dimensions. To the extent that subjects attend selectively to the relevant angle dimension, Transfer Stimulus 14 will become more similar to the Category 1 training exemplars, and subjects will tend to classify it in Category 1. An analogous situation arises for Transfer Stimulus 9 in the size categorization.

Method

Subjects. The subjects were 308 undergraduates at Indiana University, who participated as part of an introductory psychology course requirement. They were tested individually or in pairs.

To allow for a precise quantitative account of the identification and categorization confusions, only subjects who performed relatively well were included in the theoretical analyses. A number of unwanted noise factors could contribute to poor performance—including failure to understand instructions, lack of motivation, and equipment malfunctions. These factors would make it difficult for any model to account accurately for the subjects' classification behavior. The performance criteria that were used are explained in the sections on the identification and categorization conditions below.

Apparatus. The stimuli were 16 semicircles varying orthogonally in size (four levels) and angle of orientation of a radial line (four levels). The four sizes (in terms of the length of the radial line) were 4.51, 4.80, 5.08, and 5.36 mm; the four angles were 50°, 54°, 58°, and 62°. The stimulus set is illustrated in Figure 2. All
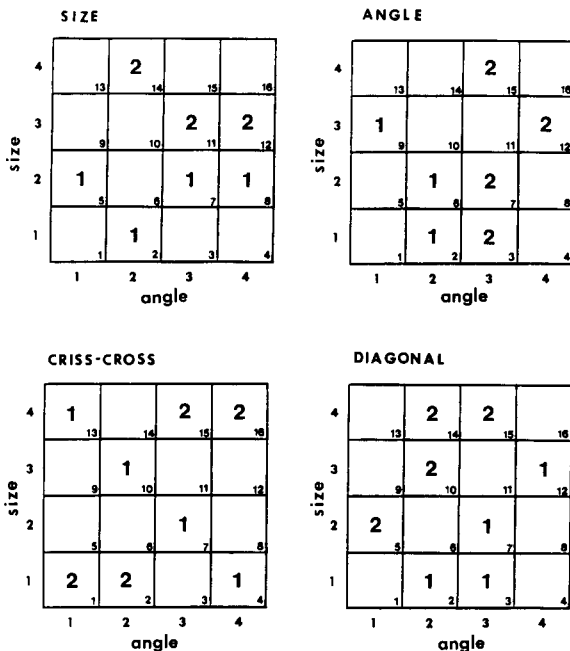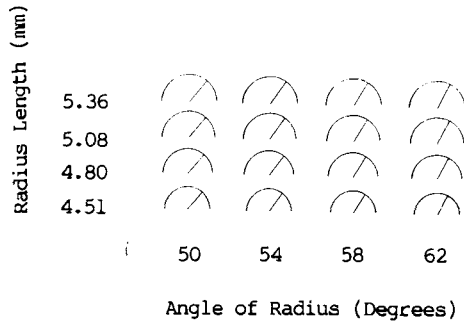
Figure 2. The stimulus set used in this experiment.

the stimuli were equal in luminance (2.73 fL); they appeared as continuous green figures on a black background.

Two VR14 scopes, interfaced with a PDP-11/34 computer, were used to present the stimuli. The subjects sat approximately 2 ft from the scopes in a normally lit room. Thus, the visual angle subtended by the stimuli ranged from approximately .85° to 1.01° at their widest point. The subjects entered their responses by means of response boxes interfaced with the computer.

Identification condition. The identification condition consisted of two phases. In Phase 1, the stimuli were presented with unlimited exposure durations, whereas in Phase 2, short exposure durations (150 msec) were used. The unlimited exposure phase was conducted to facilitate identification training. Phases 1 and 2 included 200 trials each.

On any given trial, a randomly selected stimulus was presented, and the subject identified its size and angle. The size and angle dimension values were each given a label on a scale from 1 (smallest size, lowest angle) through 4 (largest size, highest angle). The subjects entered their responses by pressing one of four consecutively labeled buttons arranged in a row. To enter size i and angle j, the subject first pressed button i and then pressed button j. Following the responses, feedback was provided with lights located directly over the buttons.

Each experimental session lasted approximately 45 min. Of the 57 subjects tested, 37 met a performance criterion of 20% correct overall identification responses during Trials 50–200 of the short exposure phase. (A response was scored as correct only if both size and angle were identified correctly. Chance performance was 6.25% correct.) The identification data analyzed in this article are those obtained during Trials 50–200 of the short exposure phase for the subjects who met the criterion.

Categorization conditions. Four categorization conditions were tested (see Figure 1). Each subject participated in only one condition. None of the subjects who participated in the identification condition participated in the categorization conditions. Each categorization condition consisted of an unlimited exposure duration training phase, a short exposure (150-msec) training phase, and a short exposure (150-msec) transfer phase.

During the training phases, only assigned category exemplars were presented. The subjects classified each stimulus into Category 1 or 2 by pressing one of two labeled response buttons. Following each response, feedback was provided. The unlimited exposure duration training phase was organized into blocks of 150 trials. A single block of unlimited exposure training was used in the size and angle categorizations, whereas three blocks were used in the crisscross and diagonal categorizations. The short exposure duration training phase consisted of a single block of 175 trials in all conditions.

The transfer phase consisted of 65 trials occurring immediately after the short exposure training phase, with no break between phases. During these 65 trials, each unassigned (transfer) stimulus was presented twice. Assigned (training) exemplars were presented on the remaining trials of the transfer phase. Presentations of as-

signed exemplars and transfer stimuli were random within the constraints stated above. Corrective feedback continued to be provided on trials in which assigned training exemplars were presented. Feedback was withheld on trials in which unassigned transfer stimuli were presented, because there were no correct answers on these trials. The subjects were not informed that they would be classifying unassigned stimuli, but they were informed that feedback might be withheld on some trials.

Average percentage correct scores were obtained for each subject during Trials 50–175 of the short exposure duration training phase. The distributions were inspected by eye, and performance criteria for each condition were established. The criteria were 80%, 80%, 60%, and 70% in the size, angle, crisscross, and diagonal conditions, respectively: 37 of 44, 41 of 83, 37 of 63, and 41 of 62 subjects met these criteria in each respective condition. Analysis of the categorization transfer data includes only those subjects who met the criteria.[1]

Theoretical Analysis

Because the theoretical framework has been presented in previous articles (Nosofsky, 1985b, 1986), it will only briefly be reviewed here. According to the exemplar-similarity model of categorization, the "strength" of making a Category J response (R_J) given presentation of Stimulus i (S_i) is found by summing the similarity of Stimulus i to the complete set of presented exemplars of Category J (C_J), and then multiplying by the response bias for Category J. This strength is then divided by the sum of strengths for all categories to determine the categorization probability. Thus, in the present two-category experiments, the probability of classifying Stimulus i into Category 1 is given by:

$$P(R_1 \mid S_i) = \frac{b_1 \Sigma_{j \epsilon C_1} s_{ij}}{b_1 \Sigma_{j \epsilon C_1} s_{ij} + (1-b_1) \Sigma_{k \epsilon C_2} s_{ik}}, \quad (1)$$

where b_1 $(0 \leq b_1 \leq 1)$ is the Category 1 response bias, and s_{ij} denotes the similarity between Exemplars i and j. In the special case in which each stimulus defines its own category, the exemplar model reduces to the classic similarity-choice model (Luce, 1963; Shepard, 1957) for predicting identification confusion data. The probability of identifying Stimulus i as Stimulus j is given by:

$$P(R_j \mid S_i) = \frac{b_j s_{ij}}{\Sigma_k b_k s_{ik}}. \quad (2)$$

The relation between Equations 1 and 2 is one of the keys to predicting categorization from identification. A critical problem, however, is that because of selective attention strategies, the s_{ij} similarity parameters in Equations 1 and 2 may not be invariant across identification and categorization situations (Medin & Schaffer, 1978; Nosofsky, 1986, 1987; Shepard et al., 1961). To characterize attention-based changes in similarities, a theory of similarity is needed. For the continuous-dimension stimuli used in these experiments, it is natural to use the MDS approach. Each exemplar is represented as a point in a two-dimensional psychological space, with the similarity between exemplars being a decreasing function of their distance in the space.

The distance between Exemplars i and j is given by

$$d_{ij} = \left[\sum_{m=1}^{2} w_m |x_{im} - x_{jm}|^r \right]^{1/r} \quad (3)$$

where x_{im} is the psychological value of Exemplar i on Dimension m, and w_m ($0 \le w_m \le 1$, $\Sigma w_m = 1$) is the attention weight given to Dimension m. The value of r in Equation 3 that provides the best description of psychological distance appears to depend on various experimental factors. In general, the Euclidean metric ($r = 2$ in Equation 3) appears to be appropriate for stimuli varying along relatively unanalyzable, integral dimensions, such as Munsell colors varying in brightness and saturation; whereas something more along the lines of the city-block metric ($r = 1$) is appropriate for highly analyzable, separable-dimension stimuli, such as the present semicircles (Garner, 1974; Shepard, 1964). However, Nosofsky (1985b) found that the Euclidean metric may be appropriate even for separable dimensions when people identify highly confusable stimuli.

The distance d_{ij} is converted to a similarity measure using the transformation

$$s_{ij} = \exp(-c \cdot d_{ij}^p) \quad (4)$$

where c ($0 \le c < \infty$) is a sensitivity parameter reflecting discriminability in the psychological space. An exponential decay function ($p = 1$ in Equation 4) appears to describe accurately the relation between similarity and psychological distance in classification learning experiments using readily discriminable stimuli, whereas a Gaussian function ($p = 2$) appears to operate in perceptual discrimination experiments involving highly confusable stimuli (Nosofsky, 1985a, 1985b; Shepard, 1986, 1987; see Ennis, 1988, and Ennis, Palen, & Mullen, in press, for a theoretical explanation of the changing similarity gradient).

The approach to predicting and relating the present identification and categorization data can now be outlined. The identification confusion data are fitted by means of the similarity choice model (Equation 2), with the assumption that the similarity parameters are functionally related to distances in a multidimensional psychological space (Equations 3 and 4). By fitting this "MDS-choice" model to the identification data, an MDS solution is derived for the stimuli. This MDS solution is then used in conjunction with the exemplar-similarity model (Equations 1, 3, and 4) to predict performance in each of the four categorization conditions. Because the MDS solution (the x_{im} coordinates in Equation 3) will be derived from the identification confusion data, one needs to estimate only the sensitivity parameter c, attention weight w_1, and response bias parameter b_1 in going on to predict categorization performance.

Identification condition. The identification confusion data are reported in Table 1. The MDS-choice model (Equations 1, 3, and 4) was fitted to the data using a maximum-likelihood criterion.[2] The sensitivity and attention weight parameters are nonidentifiable with respect to the stimulus coordinates in the identification condition, and so they were set at default values ($w_1 = .50$, $c = 2.00$).

Exploratory analyses were first conducted to discover the form of the distance metric and similarity gradient (the values of r and p in Equations 3 and 4). In these exploratory analyses, a *constrained* two-dimensional space was assumed in fitting the MDS-choice model to the identification confusion data (Nosofsky, 1985b). In the constrained two-dimensional space, all stimuli sharing a common physical value of angle are constrained to have the same psychological value on Dimension 1, x_{i1}, whereas all stimuli sharing a common physical value of size are constrained to have the same psychological value on Dimension 2, x_{i2}. The assumption of a constrained space greatly reduces the number of freely varying coordinate

Table 1
Identification Confusion Data

Stimulus	Response															
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	116	46	30	11	52	37	10	16	9	10	2	0	1	0	0	0
2	44	59	68	26	21	29	43	17	6	8	3	1	0	0	0	0
3	24	34	87	69	10	27	52	49	1	7	8	7	0	0	1	0
4	6	12	43	121	6	9	21	80	1	3	5	10	1	0	0	2
5	35	15	7	3	89	60	13	9	68	23	5	2	3	4	3	0
6	15	15	15	5	44	58	67	22	30	37	27	10	1	2	0	1
7	4	11	22	12	18	45	78	41	11	28	39	23	0	0	3	1
8	1	6	17	22	12	22	43	99	3	12	41	50	0	0	1	3
9	3	2	0	0	22	17	11	5	101	60	31	5	45	23	9	4
10	1	0	3	0	15	16	24	5	43	73	69	15	10	25	12	8
11	0	1	1	3	7	11	25	28	19	36	86	65	11	17	24	21
12	1	0	3	2	0	7	17	32	8	13	56	116	3	5	19	45
13	1	0	0	0	7	3	2	2	68	30	14	4	135	59	19	5
14	0	0	1	0	0	3	6	1	32	38	36	10	38	71	61	8
15	0	0	0	1	0	3	2	4	17	25	41	30	9	55	88	60
16	0	0	1	0	0	2	3	4	9	15	39	64	5	26	66	121

Note—Dimension levels that correspond to each stimulus number are shown in Figure 1.

Table 2
Summary Fits for MDS-Choice Models

Distance Metric	Similarity Function	SSE	% Var	$-\ln L$
	Constrained Models			
City-Block ($r = 1$)	Exponential ($p = 1$)	16,278	91.2	751.5
City-Block ($r = 1$)	Gaussian ($p = 2$)	21,921	88.1	795.8
Euclidean ($r = 2$)	Exponential ($p = 1$)	21,949	88.1	745.8
Euclidean ($r = 2$)	Gaussian ($p = 2$)	7,458	96.0	587.2
	Unconstrained Model			
Euclidean ($r = 2$)	Gaussian ($p = 2$)	5,511	97.0	557.9

Note—SSE = sum of squared deviations between predicted and observed identification confusion frequencies; % Var = percentage of variance accounted for; $\ln L$ = log-likelihood (optimization criterion was maximum-likelihood).

parameters that need to be estimated, and it reduces the possibility of hitting local minima in searching for the optimal MDS solution. Four versions of the constrained model were fitted to the data, each involving a different combination of distance metric (either city-block or Euclidean) and similarity function (either exponential or Gaussian). The summary fits for each version of the model are reported in Table 2. Among the alternatives considered in the table, the Gaussian/Euclidean model is clearly favored. This result replicates earlier ones reported by Nosofsky (1985a, 1985b) in experiments in which subjects identified highly confusable stimuli.

Having found evidence favoring the Gaussian/Euclidean model, a computer search was conducted to find the *unconstrained* two-dimensional solution that yielded a best fit to the confusion data (that is, all coordinate parameters were allowed to vary freely). The summary fits for this unconstrained model are provided in Table 2. The maximum-likelihood MDS coordinate parameters and response bias parameters are reported in Table 3, with the MDS solution illustrated in Figure 3. The MDS solution has been rotated so that the psychological dimensions correspond to the physical dimensions of size and angle.

Note that no constraints were imposed on the computer search with respect to the underlying psychological dimensions. So, for example, had the subjects' identification confusions been based primarily on the x,y coordinate of the terminal point of the radial line, the search routine was free to converge to such a solution. As is clear from inspection of Figure 3, however, the psychological dimensions are readily interpretable in terms of the size of circle and angle of line.

A scatterplot of observed against predicted confusion probabilities for the MDS-choice model is presented in Figure 4. Although the overall fit of the Gaussian/Euclidean model is fairly good, with 97.0% of the variance accounted for, it should be noted that the model systematically underpredicts correct identification probabilities (see Figure 4). Additional analyses revealed that a value of p in Equation 4 slightly less than two ($p = 1.8$) provided an even better overall fit to the data ($-\ln L = 535.1$), and removed the systematic discrepancies seen in Figure 4.

Ennis (1988; Ennis et al., in press) has proposed that the similarity gradient may be reflecting two components: Gaussian distributed perceptual dispersions associated with the presentation of a stimulus, and a "cognitive" similarity judgment function, posited by Shepard (1987) to be exponential decay in form. The "intermediate" gradient observed under the present experimental conditions may reflect significant contributions from both components. By contrast, in experiments in which there is presumably relatively little perceptual noise associated with stimulus presentations, evidence tends to favor an exponential similarity gradient (Nosofsky, 1987; Shepard, 1958, 1987).

Categorization conditions. The data obtained in the four categorization conditions are reported in Table 4. Given the MDS solution derived from the identification confusions, the exemplar model (Equations 1, 3, and 4) can be used to predict these categorization data. In each condition, three parameters (c, w_1, and b_1) are estimated to predict 16 freely varying data points.

To aid in interpretability, the observed and predicted Category 1 response probabilities for each stimulus in each condition are shown in spatial format in Figure 5.

Table 3
Maximum-Likelihood MDS-Choice Model Parameters

S_i	MDS Coordinates		Bias
	x_{i1}	x_{i2}	b_i
1	.312	-.241	.426
2	.918	-.264	.278
3	1.405	-.187	.416
4	2.062	-.227	.549
5	.228	.640	.522
6	.844	.662	.466
7	1.324	.687	.569
8	1.885	.623	.706
9	.374	1.555	.784
10	.916	1.501	.650
11	1.473	1.544	.817
12	2.128	1.520	.929
13	.135	2.352	.655
14	.889	2.412	.561
15	1.451	2.493	.630
16	2.061	2.382	.687

Note—x_{i1} = psychological value of Stimulus i (S_i) on Dimension 1 (angle); x_{i2} = psychological value of Stimulus i on Dimension 2 (size).
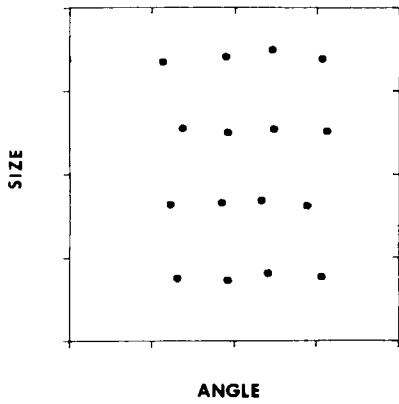
Figure 3. Multidimensional scaling solution derived by fitting the Gaussian/Euclidean MDS-choice model to the identification confusion data.
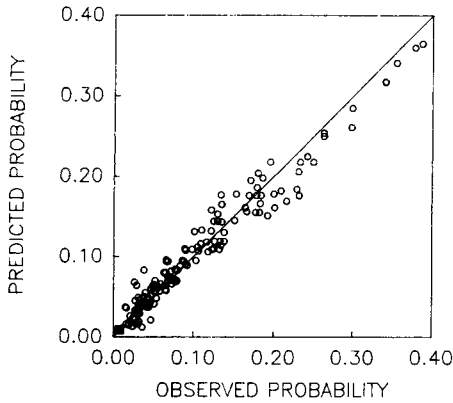

Figure 4. Scatterplot of predicted against observed identification confusion probabilities (MDS-choice model). The solid square in lower left corner of the scatterplot represents the large number of observed and predicted confusion probabilities less than .02.

The maximum-likelihood exemplar model parameters and summary fits are reported in Table 5. The model accounts for 99.4%, 99.6%, 95.2%, and 98.4% of the variance in the size, angle, crisscross, and diagonal categorizations, respectively. A scatterplot of observed against predicted Category 1 response probabilities is presented in Figure 6. The quantitative predictions are accurate across the entire range of the probability space. Likelihood-ratio tests of overall goodness of fit (see Wickens, 1982, chapter 6) do not lead to the rejection of the model in any of the conditions [average $\chi^2(13) = 18.0, p > .05$].

The parameter of greatest interest in this study is the attention weight w_1 (weight given to angle—see Table 5). As expected, the subjects attended selectively to the size dimension ($w_1 = .10$) in the size categorization, and attended selectively to angle ($w_1 = .98$) in the angle categorization. Although the result was not anticipated, the subjects also gave greater weight to the angle dimension in the crisscross and diagonal categorizations ($w_1 = .80$ and $w_1 = .81$, respectively) than they did in the identification condition ($w_1 = .50$).

Note that the estimated attention weights should be interpreted in only a relative sense. In the identification condition, the weights were set at $w_1 = w_2 = .50$, because they were nonidentifiable with respect to the stimulus coordinates. Conceivably, however, the subjects could have weighted size more than angle in the identification condition. Thus, the value $w_1 = .80$ in the crisscross categorization indicates simply that the subjects gave greater weight to angle in the crisscross categorization than in the identification condition. One possible explanation of this finding is that in the identification condition, the subjects reported their size judgment first, followed by their angle judgment. This fixed order of report may have biased the subjects to focus more on the size dimension than on the angle dimension. Such biasing was

Table 4
Observed Frequencies with which Stimuli were Classified in Category 1 or 2 in each Condition

Stimulus	Size C_1	Size C_2	Angle C_1	Angle C_2	Crisscross C_1	Crisscross C_2	Diagonal C_1	Diagonal C_2
1	72	2	79	3	48	168	40	46
2	255	4	155	116	94	138	200	58
3	72	2	48	258	45	29	242	25
4	73	1	2	80	162	49	83	3
5	234	35	81	1	34	40	65	228
6	66	8	190	97	34	40	49	37
7	208	39	60	202	138	102	180	77
8	226	39	2	80	47	27	70	16
9	23	51	262	25	53	21	8	78
10	18	56	47	35	120	103	57	211
11	55	170	11	71	33	41	42	44
12	58	179	4	259	24	50	199	67
13	2	72	76	6	180	52	3	83
14	8	229	47	35	44	30	27	242
15	3	71	24	227	63	160	58	171
16	3	71	2	80	41	195	55	31

SIZE

03 03	02 03	01 04	02 04
	2		
13	14	15	16
31 31	33 24	28 24	28 25
	2	2	
9	10	11	12
86 87	84 89	81 84	83 85
1	1	1	
5	6	7	8
99 97	99 99	99 97	99 99
	1		
1	2	3	4

ANGLE

98 93	56 57	13 10	01 02
	2		
13	14	15	16
92 91	55 57	14 13	01 02
1		2	
9	10	11	12
96 99	62 66	23 23	03 02
1		2	
5	6	7	8
94 96	56 57	19 16	01 02
1		2	
1	2	3	4

CRISS-CROSS

81 78	56 60	33 28	19 17
1		2	2
13	14	15	16
64 72	57 54	48 45	36 32
	1		
9	10	11	12
40 46	49 46	56 58	61 64
	1		
5	6	7	8
22 22	37 41	55 61	76 77
2	2		1
1	2	3	4

DIAGONAL

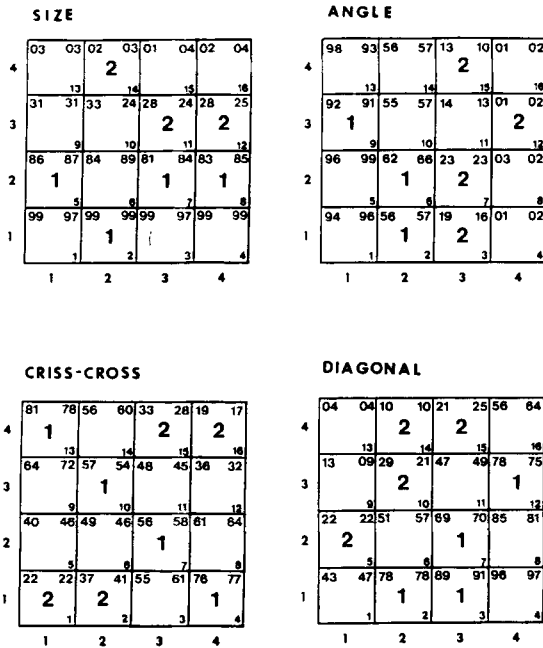|			
04 04	10 10	21 25	56 64
	2	2	
13	14	15	16
13 09	29 21	47 49	78 75
	2		1
9	10	11	12
22 22	51 57	69 70	85 81
2		1	
5	6	7	8
43 47	78 78	89 91	96 97
	1	1	
1	2	3	4

Figure 5. Predicted and observed Category 1 response probabilities for each stimulus in each condition. The top left value in each cell = predicted probability; the top right value in each cell = observed probability.

not inherent in the categorization conditions, where the processing of dimensions was more at the subjects' option.

To illustrate the important role of the attention weight parameter, a restricted version of the exemplar model was fitted to the data in which the attention weight was set at $w_1 = .50$ (that is, a constant distribution of attention across identification and categorization). The maximum-likelihood parameters and summary fits for this restricted model are presented along with those of the full model in Table 5. In all cases, the fits of this restricted model

are dramatically worse than those of the full model. Likelihood-ratio tests (Wickens, 1982) indicate that the weight parameters are significantly different ($p < .001$) from .50 in all conditions [$\chi^2(1) = 62.4, 240.0, 33.8,$ and $122.2,$ in the size, angle, crisscross, and diagonal categorizations, respectively].[3]

Two stimuli that are useful for bringing out the importance of the attention weight parameter are Stimulus 14 in the angle categorization and Stimulus 9 in the size categorization (see Figure 5). In the angle categorization, the restricted model with $w_1 = .50$ predicts that Stimulus 14 will be classified in Category 1 with probability .18. By contrast, with selective attention incorporated, the predicted value is .56, essentially identical to the observed value of .57. With nondifferential selective attention, Stimulus 14 is more similar overall to the Category 2 exemplars than to the Category 1 exemplars. With selective attention to angle, however, the psychological space is "stretched" along the horizontal dimension and "shrunk" along the vertical dimension, making Stimulus 14 more similar to the exemplars of Category 1. An analogous situation arises for Stimulus 9 in the size categorization [restricted model: $P(R_1|S_9) = .48;$ full model: $P(R_1|S_9) = .31;$ observed: $P(R_1|S_9) = .31$].

A restricted version of the exemplar model was also tested in which the bias parameter was set at $b_1 = .50$ (nondifferential response bias). As indicated in Table 5, the bias parameter did not contribute to the quantitative fits in nearly as dramatic a fashion as did the attention weight parameter. Indeed, likelihood-ratio tests revealed that the bias parameter did not differ significantly from .50 in the size or diagonal categorizations [average $\chi^2(1) = 0.80, p > .05$], although it did in the crisscross and angle categorizations [average $\chi^2(1) = 21.5, p < .01$]. Presumably, the bias parameter would play a more critical role in conditions in which differential payoffs for alternative categorization responses were used, category frequencies were manipulated, and so forth.

Table 5
Maximum-Likelihood Parameters and Summary Fits for Full and Restricted Versions
of the Exemplar-Similarity Categorization Model

Condition	Model	Parameters			Fits		
		c	w_1	b_1	SSE	% Var	$-\ln L$
Size	Unconstrained	1.60	.10	.50	.015	99.4	40.8
	Equal attention	2.38	(.50)	.49	.077	97.0	72.0
	Equal bias	1.60	.10	(.50)	.015	99.4	40.8
Angle	Unconstrained	3.20	.98	.43	.010	99.6	44.3
	Equal attention	3.57	(.50)	.45	.305	86.4	164.3
	Equal bias	3.09	1.00	(.50)	.029	98.7	56.8
Crisscross	Unconstrained	1.62	.80	.45	.025	95.2	47.7
	Equal attention	1.23	(.50)	.45	.087	83.1	64.6
	Equal bias	3.00	.93	(.50)	.046	91.1	56.7
Diagonal	Unconstrained	2.42	.81	.49	.023	98.4	48.3
	Equal attention	1.81	(.50)	.48	.217	85.0	109.4
	Equal bias	2.42	.81	(.50)	.021	98.6	49.1

Note—Values in parentheses were constrained a priori. The parameter w_1 gives the attention weight for angle, and $1-w_1$ the attention weight for size. SSE = sum of squared deviations between predicted and observed Category 1 probabilities; % Var = percentage of variance accounted for; $\ln L$ = log likelihood.
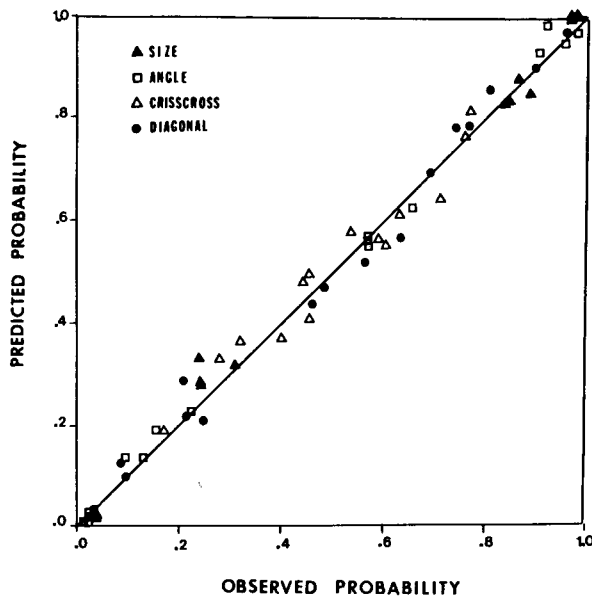
Figure 6. Scatterplot of predicted against observed Category 1 response probabilities for each stimulus in each condition.

GENERAL DISCUSSION

The main contribution of this research involved the demonstration that precise quantitative accounts of both identification and categorization confusion data could be achieved within the framework of the proposed exemplar-similarity model. On the basis of identification confusions, an MDS solution for the exemplars was derived. This MDS solution was then used to predict performance in four separate categorization conditions with the same set of stimuli. A key to achieving the accurate quantitative predictions was the use of selective attention weights that systematically modified similarities between exemplars across the different categorization conditions. Whether the weights are reflecting changes in "perceptual" similarity or some form of "cognitive" or "judgmental" similarity is an important but as yet unexplored question.

The present tests of the exemplar-similarity model go beyond earlier ones reported by Nosofsky (1986) in several respects. Whereas Nosofsky (1986) tested only two subjects, each of whom received extensive discrimination training on the entire ensemble of exemplars, the present tests involved large numbers of subjects, each of whom received relatively little training. Also, whereas Nosofsky's (1986) design required him to posit that subjects had augmented their memory representations with "inferred" category exemplars, the present tests did not require recourse to this memory-augmentation process. Thus, the present research represents a substantial advance in demonstrating the generalizability and utility of the proposed theoretical approach to relating identification and categorization.

In the present study, a "similarity-based" model was used to accurately predict classification performance, despite the fact that separable-dimension stimuli were in-

volved. Although there is widespread support for the idea that similarity in psychological space predicts classification performance for integral-dimension stimuli (e.g., Monahan & Lockhead, 1977; Nosofsky, 1987; Shepard, 1958; Shepard & Chang, 1963), the role of similarity in determining the classification of separable-dimension stimuli has been called into question (e.g., by Garner, 1974; Shepard et al., 1961). The present research suggests that if assumptions are made concerning the role of selective attention in influencing similarity, then a similarity-based approach to predicting classification performance using separable-dimension stimuli may be tenable. L. B. Smith (in press) has recently extended the present modeling approach to characterize developmental changes in the free classification of separable-dimension stimuli.

Consideration of Alternative Models

Although the exemplar-similarity model provided an accurate quantitative account of the present identification and categorization data, its predictions have not been contrasted with those of alternative classification models. Unfortunately, many extant competing models have simply not been developed to the point at which they can directly be applied to predict categorization performance for the present continuous-dimension stimuli. For example, Gluck and Bower (1988a, 1988b) proposed and tested an adaptive network model of classification learning that has been applied in paradigms involving stimuli varying along discrete, binary-valued dimensions (see also Estes, Campbell, Hatsopoulos, & Hurwitz, in press). The nodes in the network correspond to the individual dimension values and, in an extended version of the model, to combinations of these values. Activation of the nodes is unambiguous and all-or-none. For example, presentation of a large red stimulus would lead to activation of the large, red, and large-red nodes. There are several problems regarding application of this adaptive network model to the present categorization data. First, it is unclear how to represent continuous-dimension stimuli in terms of the node-network architecture. Furthermore, it seems likely that additional assumptions would be necessary regarding node activation. For example, presentation of a stimulus with Size 3 (S3) and Angle 2 (A2) would presumably lead to activation of an S3-A2 node, but might also lead to activation of nodes that represent similar values (e.g., S4-A2 and S3-A3 nodes, etc.). The details of such similarity-related node-activation functions would need to be worked out. This problem of how to precisely represent similarity relations among confusable, continuous-dimension stimuli currently exists for a variety of connectionist, distributed-memory, and even some alternative exemplar models of classification (e.g., Carpenter & Grossberg, 1987; Hintzman, 1986; Knapp & Anderson, 1984; McClelland & Rumelhart, 1985; Metcalfe-Eich, 1982; Murdock, 1982).

One theoretical approach which in its present form is applicable to continuous-dimension stimuli is the general recognition theory (GRT) currently being developed by

Ashby and his associates (Ashby & Gott, 1988; Ashby & Perrin, 1988; Ashby & Townsend, 1986). The GRT is a multidimensional generalization of Thurstonian models of classification. Presentation of a stimulus is assumed to give rise to a "perceptual effect" in a multidimensional psychological space. Because of noise in the system, the same perceptual effect will not be yielded on each trial; rather, there is assumed to be a distribution of perceptual effects associated with each stimulus. The subject is assumed to establish decision boundaries that partition the multidimensional space into response regions. Any perceptual effect falling into Region A would result in a Category A response.

To use the GRT to make quantitative predictions of classification performance, assumptions need to be made regarding the form of the distributions of perceptual effects as well as the types of decision boundaries that are adopted. Most applications of the GRT have assumed multivariate Gaussian distributions. A variety of decision bounds have been investigated, including independent-decision bounds, general linear bounds, and likelihood-ratio bounds (Ashby & Gott, 1988).

The procedure for using the GRT to predict the present identification and categorization data would be analogous to the one used for fitting the exemplar-similarity model. On the basis of the identification confusion data, the parameters of the underlying distributions of perceptual effects would be estimated (e.g., see Ashby & Perrin, 1988). Then, for a given choice of decision boundary, predictions would be made regarding performance in each categorization condition. The probability of a Category 1 response, given presentation of Stimulus i, would be found by integrating over the portion of the Stimulus i distribution that falls in the Category 1 response region. In general, this would involve multiple integration of Gaussian distributions over what might be fairly complex spatial regions. Such computations go beyond the scope of the present research.

Although quantitative tests of the GRT are not made in this article, the observed patterns of identification and categorization confusions can be shown to have important implications regarding its use. In the following discussion, I assume decision boundaries that are based on likelihood-ratio. Likelihood-ratio boundaries deserve special consideration for two reasons. First, optimal classification strategies can essentially always be couched in terms of likelihood-ratio (Green & Swets, 1966, chapter 1). Also, Ashby and Perrin (1988) assumed likelihood-ratio boundaries in applying the GRT to the modeling of similarity judgments.

According to an unbiased likelihood-ratio decision rule, in a two-category experiment a perceptual effect t will be classified into Category 1 whenever

$$L(C_1 | t) > L(C_2 | t), \qquad (5)$$

where $L(C_1 | t)$ denotes the likelihood of Category 1 given perceptual effect t. Using Bayes' Theorem, Inequality 5 is satisfied whenever

$$L(t | C_1) \cdot L(C_1) > L(t | C_2) \cdot L(C_2). \qquad (6)$$

Inequality 6 can be expanded in terms of the individual stimulus distributions that compose the overall category distributions of perceptual effects. In particular, a Category 1 response is made if

$$\sum_{j \in C_1} L(t | S_j) \cdot L(S_j) > \sum_{k \in C_2} L(t | S_k) \cdot L(S_k), \qquad (7)$$

where $L(t | S_j)$ denotes the likelihood of perceptual effect t, given presentation of Stimulus j. When the individual stimuli have equal presentation rates, as in the present experiments, the decision rule is simply to respond "Category 1" if

$$\sum_{j \in C_1} L(t | S_j) > \sum_{k \in C_2} L(t | S_k). \qquad (8)$$

Assume that the distributions of perceptual effects associated with the stimuli S_j are bivariate Gaussian, with mean on dimension m, μ_{jm}, variance on dimension m, σ_m^2, and covariance zero. (Note that we assume that variability along each dimension is constant across stimuli, although the variance along Dimension 1 is not necessarily the same as the variance along Dimension 2). Then the likelihood of perceptual effect $t = (t_1, t_2)$, given presentation of Stimulus j, is

$$L(t | S_j)$$

$$= \frac{1}{2\pi\sigma_1\sigma_2} \exp\{-[(t_1 - \mu_{j1})^2/2\sigma_1^2 + (t_2 - \mu_{j2})^2/2\sigma_2^2]\}. \qquad (9)$$

Substituting into Inequality 8 and dividing all terms by the factor $1/(2\pi\sigma_1\sigma_2)$, the decision rule is to respond "Category 1" if

$$\sum_{j \in C_1} \exp\{-[(t_1 - \mu_{j1})^2/2\sigma_1^2 + (t_2 - \mu_{j2})^2/2\sigma_2^2]\}$$

$$> \sum_{k \in C_2} \exp\{-[(t_1 - \mu_{k1})^2/2\sigma_1^2 + (t_2 - \mu_{k2})^2/2\sigma_2^2]\}. \qquad (10)$$

As noted by Nosofsky (1988), this likelihood-based decision rule is formally identical to one produced by summing similarities to individual exemplars (using a Gaussian similarity function and a Euclidean distance metric). In particular, suppose that the decision rule is to respond "Category 1" if the summed similarity of t to all exemplars of Category 1 is greater than the summed similarity of t to all exemplars of Category 2:

$$\sum_{j \in C_1} s(t, S_j) > \sum_{k \in C_2} s(t, S_k). \qquad (11)$$

In the Gaussian/Euclidean exemplar model, the similarity of t to Stimulus j is given by

$$s(t, S_j) = \exp(-c \cdot d_{tj}^2)$$

$$= \exp\left\{-c\sqrt{w_1(t_1-x_{j1})^2 + w_2(t_2-x_{j2})^2}^2\right\}$$

$$= \exp\left\{-[cw_1(t_1-x_{j1})^2 + cw_2(t_2-x_{j2})^2]\right\}, \quad (12)$$

and so the decision rule is to respond Category 1 if

$$\sum_{j\epsilon C_1} \exp\{-[cw_1(t_1-x_{j1})^2 + cw_2(t_2-x_{j2})^2]\}$$

$$> \sum_{k\epsilon C_2} \exp\{-[cw_1(t_1-x_{k1})^2 + cw_2(t_2-x_{k2})^2]\}. \quad (13)$$

Letting the coordinates for Stimulus j in the exemplar model correspond to the mean of the distribution of perceptual effects associated with Stimulus j in the GRT, $x_{jm} = \mu_{jm}$, and letting the product of the sensitivity and weight parameters in the exemplar model be inversely related to the variance parameters in the GRT, $c \cdot w_m = 1/2\sigma_m^2$, it can be seen that the decision boundaries defined by these versions of the exemplar model and the GRT are formally identical (i.e., compare Inequalities 10 and 13).

Now in the exemplar-similarity model tested in this article, classification predictions were made using a response-ratio rule (Equation 1), whereas the GRT makes classification predictions by integrating over distributions of perceptual effects that fall in particular response regions. However, given the preceding theoretical development, it is clear that the models will often make similar qualitative predictions. As an illustration, Figure 7 shows an approximate representation of the structure of the present stimulus set in terms of the GRT. (The precise structure would be found by fitting the GRT to the identification confusion data.) In the figure, the stimulus distributions are represented by ellipses, with the width of the major and minor axes of the ellipses representing variability along each dimension. In Panel A of Figure 7, variability along each dimension is presumed to be equal, so the distributions of perceptual effects are represented by circles. This equal variability assumption parallels the equal weighting assumption that was made earlier when fitting the exemplar model to the identification confusion data.

Panel A of Figure 7 also illustrates the angle categorization embedded in the stimulus structure, and shows the boundary of equal likelihood for Categories 1 and 2. This boundary was computed by solving for individual (t_1, t_2) pairs for which $\Sigma_{j\epsilon C_1}L(t|S_j) = \Sigma_{k\epsilon C_2}L(t|S_k)$ for the stimulus-distribution parameters that are assumed in the figure. According to the GRT, any perceptual effect falling to the upper right of this decision boundary would result in a Category 2 response. It is clear from inspection that the bulk of the distribution for Transfer Stimulus 14 falls in the Category 2 response region, and so the likelihood-based GRT would incorrectly predict classification for this critical transfer stimulus, just as did the overall-similarity exemplar model.

Panel B of Figure 7 illustrates a modified representation of the structure of the stimulus space for the angle categorization. Here it is assumed that the subjects have

attended selectively to the relevant angle dimension. This selective attention strategy is represented in terms of reduced variability of perceptual effects along the angle dimension, and increased variability along the size dimension. A new boundary of equal likelihood has been computed, and it is clear that the GRT would now make more accurate predictions of classification for Transfer Stimulus 14.

In summary, there is a close formal relation between the sensitivity and attention weight parameters in the exemplar model and the variance parameters in the likelihood-based GRT. Furthermore, the selective attention assumption that was needed for connecting the identification and categorization data in terms of the exemplar-similarity model also seems necessary for the likelihood-based GRT. Of course, an alternative approach to predicting the categorization data in terms of the GRT would involve the assumption that the distributions of perceptual effects are invariant across identification and categorization, but that different decision rules are used. For example, independent-decision boundaries may be used for some category structures, and likelihood-ratio boundaries for others. It is unclear at present which approach is preferable.

The preceding discussion has been intended to illustrate implications of the present results for one particular alternative classification model. More generally, it can be stated that the present research has furnished a rich set of identification and categorization confusion data that should provide a fertile testing ground for a wide variety of quantitative models of classification performance.
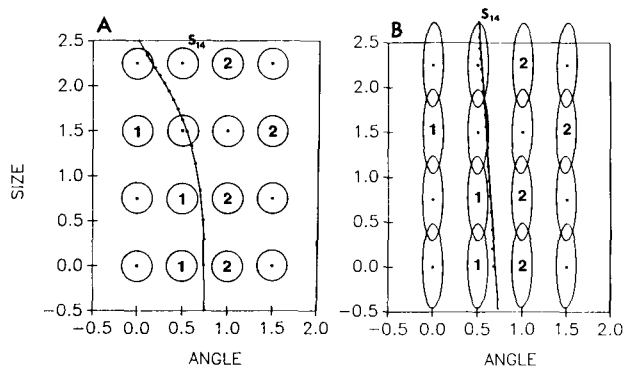

Figure 7. Schematic illustration of the structure of the angle categorization in terms of the general recognition theory. For simplicity, a constrained two-dimensional space is assumed: on Dimension 1 (angle) the stimulus means are located at 0, .5, 1, and 1.5, whereas on Dimension 2 (size) the means are located at 0, .75, 1.5, and 2.25. In Panel A, there is equal attention to the size and angle dimensions, with $\sigma_1^2 = \sigma_2^2 = .50$. (These variances correspond to $c = 2.0$ and $w_1 = .5$ in the exemplar-similarity model.) In Panel B, there is selective attention to the angle dimension, represented by reduced variability along Dimension 1 ($\sigma_1^2 = .28$) and increased variability along Dimension 2 ($\sigma_2^2 = 2.50$). (These variances correspond to $c = 2.0$ and $w_1 = .95$ in the exemplar-similarity model.) The length of the radius along the major and minor axes of the ellipses is approximately 1/4 of a standard deviation (σ). Boundaries of equal likelihood for Categories 1 and 2 are also drawn. Note the contrasting predictions of classification for Transfer Stimulus 14 (S_{14}) across Panels A and B.

REFERENCES

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., & Perrin, N. A. (1988). Toward a unified theory of similarity and recognition. *Psychological Review*, **95**, 124-150.

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

Carpenter, G. A., & Grossberg, S. (1987). Neural dynamics of category learning and recognition: Attention, memory consolidation, and amnesia. In J. Davis, R. Newgurgh, & E. Wegman (Eds.), *Brain structure, learning, and memory*, AAAS Symposium Series, 1987.

Carroll, J. D., & Wish, M. (1974). Models and methods for three-way multidimensional scaling. In D. H. Krantz, R. C. Atkinson, R. D. Luce, & P. Suppes (Eds.), *Contemporary developments in mathematical psychology* (Vol. 2, pp. 57-105). San Francisco: W. H. Freeman.

Ennis, D. M. (1988). Confusable and discriminable stimuli: Comments on Nosofsky and Shepard. *Journal of Experimental Psychology: General*, **117**, 408-411.

Ennis, D. M., Palen, J., & Mullen, K. (in press). A multidimensional theory of similarity. *Journal of Mathematical Psychology*.

Estes, W. K., Campbell, J. A., Hatsopoulos, N., & Hurwitz, J. B. (in press). Base-rate effects in category learning: A comparison of parallel network and memory storage-retrieval models. *Journal of Experimental Psychology: Learning, Memory, & Cognition*.

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

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

Gluck, M. A., & Bower, G. H. (1988a). Evaluating an adaptive network model of human learning. *Journal of Memory & Language*, **27**, 166-195.

Gluck, M. A., & Bower, G. H. (1988b). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, **117**, 227-247.

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

Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, **93**, 411-428.

Knapp, A. G., & Anderson, J. A. (1984). Theory of categorization based on distributed memory storage. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **10**, 616-637.

Luce, R. D. (1963). Detection and recognition. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (pp. 103-189). New York: Wiley.

McClelland, J. I., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, **114**, 159-188.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, **85**, 207-238.

Metcalfe-Eich, J. (1982). Composite holographic associative recall model. *Psychological Review*, **92**, 1-37.

Monahan, J. S., & Lockhead, G. R. (1977). Identification of integral stimuli. *Journal of Experimental Psychology: General*, **106**, 94-110.

Murdock, B. B., Jr. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, **89**, 609-626.

Nosofsky, R. M. (1985a). Luce's choice model and Thurstone's categorical judgment model compared: Kornbrot's data revisited. *Perception & Psychophysics*, **37**, 89-91.

Nosofsky, R. M. (1985b). Overall similarity and the identification of separable-dimension stimuli: A choice model analysis. *Perception & Psychophysics*, **38**, 415-432.

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. (1988). *Relations between exemplar-similarity and likelihood models of classification*. Manuscript submitted for publication.

Shepard, R. N. (1957). Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space. *Psychometrika*, **22**, 325-345.

Shepard, R. N. (1958). Stimulus and response generalization: Tests of a model relating generalization to distance in psychological space. *Journal of Experimental Psychology*, **55**, 509-523.

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

Shepard, R. N. (1986). Discrimination and generalization in identification and classification: Comment on Nosofsky. *Journal of Experimental Psychology: General*, **115**, 58-61.

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, **237**, 1317-1323.

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

Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs*, **75**(13, Whole No. 517).

Smith, J. E. K. (1980). Models of identification. In R. Nickerson (Ed.), *Attention and performance VIII* (pp. 129-158). Hillsdale, NJ: Erlbaum.

Smith, L. B. (in press). A model of perceptual classification in children and adults. *Psychological Review*.

Townsend, J. T., & Landon, D. E. (1982). An experimental and theoretical investigation of the constant-ratio rule and other models of visual letter confusion. *Journal of Mathematical Psychology*, **25**, 119-162.

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

NOTES

1. Performance criteria vary across conditions because some categories were harder to discriminate than others. The approach to establishing criteria was to inspect the percentage correct distributions by eye to look for natural breakpoints. The criteria were established prior to the model analyses. Note that severe subject selection took place in the angle categorization, where only the top 50% of the subjects were included in the analyses. This occurred because the subjects in the angle condition received only one block of training, whereas the subjects in the crisscross and diagonal conditions received three blocks of training. One block of training was sufficient in the size categorization, because size turned out to be more highly discriminable than angle for the dimension levels tested in this experiment (see Figure 3).

2. A computer search was used to find the parameter values that maximized the log-likelihood function

$$\ln L = \sum_i \ln N_i! - \sum_i \sum_j \ln f_{ij}! + \sum_i \sum_j f_{ij} \cdot \ln p_{ij}$$

where N_i is the frequency with which Stimulus i was presented, f_{ij} is the observed frequency with which Stimulus i was identified as Stimulus j, and p_{ij} is the predicted probability with which Stimulus i is identified as Stimulus j. (Note that only the final term in the log-likelihood function changes with variations in the model parameters.) This likelihood function assumes that the responses for each stimulus are multinomially distributed, and that the distributions for each stimulus are independent. The same likelihood function is used for fitting the exemplar model to the categorization matrices, except that here f_{ij} (or p_{ij}) corresponds to the observed frequency (or the predicted probability) with which Stimulus i is classified in Category j.

3. A restricted version of a model arises whenever some of the parameters in the model are constrained on a priori grounds. Let L be

the likelihood of a data set given the maximum-likelihood parameters in the full model, and let L^* be the corresponding likelihood for a restricted version of the model. Assuming that the full model is correct, then for large sample size, the statistic $-2(\ln L - \ln L^*)$ is distributed approximately as a chi-square random variable with degrees of freedom equal to the number of parameters that were constrained. Thus, this likelihood-ratio statistic can be used to test whether the fit of the restricted model is significantly worse than that of the full model. Tests of overall goodness of fit of a model are special cases of likelihood-ratio testing. One compares the fit of the model to that of a fully saturated model that fits the data perfectly.