Predicting similarity change as a result of categorization

Irene Reppa (i.reppa@swansea.ac.uk)
Department of Psychology, Swansea University, Swansea, SA2 8PP

Emmanuel Pothos (e.m.pothos@gmail.com)
Department of Psychology, City University London, London, EC1V 0HB

Abstract

Learning a particular categorization leads to corresponding changes in the similarity structure of the categorized stimuli. The purpose of the current study was to examine whether different category structures may lead to greater or less similarity change. We created six category structures and examined changes in similarity as a result of categorization in between-participant conditions. The best-supported hypothesis was that the ease of learning a categorization affects change in similarity, with the most change following learning of difficult category structures. There was also support for the hypothesis that similarity change is more likely to occur when the category boundary was not aligned with the physical dimension of variation. Finally, we discuss some methodological challenges in addressing this important research topic.

Keywords: similarity; categorization; learning difficulty; exemplar theory

There is widespread evidence that learning to categorize stimuli in a particular way leads to corresponding changes in the similarity structure of the stimuli (e.g., Gureckis & Goldstone, 2008; Ozgen & Davies, 2002; Schyns & Oliva, 1999; Stevenage, 1998). For instance, stimuli categorized in the same category tend to be perceived as more similar to each other, compared to stimuli categorized in different categories (e.g., Goldstone, 1994; Schyns, Goldstone, & Thibaut, 1997), and stimuli on either side of a category boundary tend to be more discriminable than stimuli on the same side of the boundary (e.g., Harnad, 1987). Similarly, differences have been reported on color perception across different linguistic communities (Roberson et al., 2005).

Research on the influence of categorization on perception has flourished for several reasons. It is theoretically important since it is at the heart of answering core issues regarding representation and the processing of sensory input. Do we perceive a faithful representation of sensory input? Or are our perceptual representations a compromise between constraints from sensory input and whatever categories are useful for the organism? Such research is also important for formal models of categorization, as most of them assume categorization models that assume representations which are stable across learning (e.g., Nosofsky, 1984).

The present research examined the effects of categorization on *similarity.* Changes in similarity might correspond to perceptual changes or changes mediated through the addition of a category label (e.g., Goldstone, Lippa, Shiffrin, 2001; McMurray et al., in press; Roberson & Davidoff, 2000; Sloutsky & Fisher, 2004). Choosing to examine similarity is primarily a methodological simplification, since exploring directly changes relating to perception involves the technical challenge of eliminating (possible) effects from linguistic labeling. However, if across broadly matched category structures, for instance, in terms of learning difficulty, we find similarity changes following learning of some structures but not others, then one can make the additional step of inferring similarity changes over and above changes due to just the category label (see also Roberson et al., 2007).

Despite the numerous reports on the effects of categories on similarity, and perception in particular, there have been some reports of failures of such influence (e.g., Goldstone, 1994; Jiang et al., 2007; Freedman et al., 2003). The aim of present research was to examine possible factors might lead to changes in similarity.

We created six two-dimensional category structures, shown in Figure 1. Two category structures were designed so that the width dimension was diagnostic (Width easy and Width difficult), while the height dimension non-diagnostic, and two more category structures were defined so that height was the diagnostic dimension (Height easy and Height difficult) and width was non-diagnostic. Two versions of each category structure were created, one designed to be easy (e.g., Width easy), and one designed to be more difficult (e.g., Width difficult). Finally, two more category structures were created where both dimensions were relevant: the non-linearly separable (NLS) and the Diagonal structure, explained in more detail later.

Three different hypotheses regarding the effect of category learning on similarity changes were examined. One hypothesis was that category *learning difficulty* would affect the extent of similarity changes. A classification is easy (or more intuitive) if it is more readily obvious to naïve observers (Pothos & Chater, 2002; Pothos et al., 2011). For example, when asked to freely classify a set of stimuli, participants will generate more intuitive classifications more frequently. These classifications will be typically easier to learn than non-intuitive ones. Category learning difficulty might influence similarity ratings in two ways. One possibility is that learning the easy category structures would lead to greater changes in similarity ratings. This is because, for easy category structures participants are able to quickly learn the underlying categorization, perhaps with less emphasis on encoding the individual exemplars (cf. Ashby et al., 1999; Ashy & Ell, 2002). Such inexact initial encoding of the exemplars may mean that exemplar representations end up being developed in a way that is more consistent with the underlying category structure (e.g., Edwards, Pothos, & Perlman, in press). Support for this prediction comes from Folstein, Gauthier, and Palmeri (2010), who manipulated the complexity of the underlying stimulus space (not of category structure, as in the current study). Unlike previous related evidence showing that categorization does not influence similarity (e.g., Jiang et al., 2007; Freedman et al., 2003), they showed significant effects of categorization on perception, when the underlying stimulus space was simple.

The converse prediction, regarding the effect of category learning difficulty on similarity changes, is that learning a difficult category structure might result in more significant and enduring changes in the similarity structure of the stimuli. This possibility is motivated by evidence showing that supervised categorization processes can involve processes of selective attention or other changes in psychological space (e.g., through the sensitivity parameter; Nosofsky, 1984), though such research does not tell us whether such changes are enduring and on the actual stimulus representations.

A second hypothesis is that the linear separability of the learned categories might moderate changes in similarity. Overall, there is quite a lot of controversy regarding the role of linear separability in category learning and perhaps some of this controversy can be ultimately explained in terms of corresponding changes in the similarity structure of the categorized stimuli. Note that connectionist models require that NLS problems are transformed into linearly separable ones at their hidden layer, otherwise learning is not possible (indeed, the inability of perceptrons to learn NLS category structures has been at the heart of the critique of Minsky & Papert, 1969; see also Rumelhart & McClelland, 1986).

To examine possible influences of linear separability in relation to similarity changes, two categories were created to be broadly equal in terms of complexity but differ in whether they were LS or not. One was the NLS and the other was the Diagonal condition (Figure 1). In diagonally separated category structures, the members of one category can only be discriminated from their nearest neighbors in the other category with fine distinctions along both dimensions of variation. They have proved to be challenging for participants to learn (e.g., Ashby, Queller, & Berretty, 1999; Ashby & Ell, 2002).

If the cognitive system shares processing constraints with connectionist systems, maybe it would try to re-represent a NLS classification in an LS way, so that there would be *more similarity change* in learning an NLS classification, compared to an equally complex but LS one (the Diagonal one). Alternatively, it could be the case that more complex classifications are associated with *less similarity change*, if, for example, category learning of such classifications involves rote memorization of the training exemplars (Blair & Homa, 2003). In this case, NLS and Diagonal classification would lead to equivalent similarity changes.

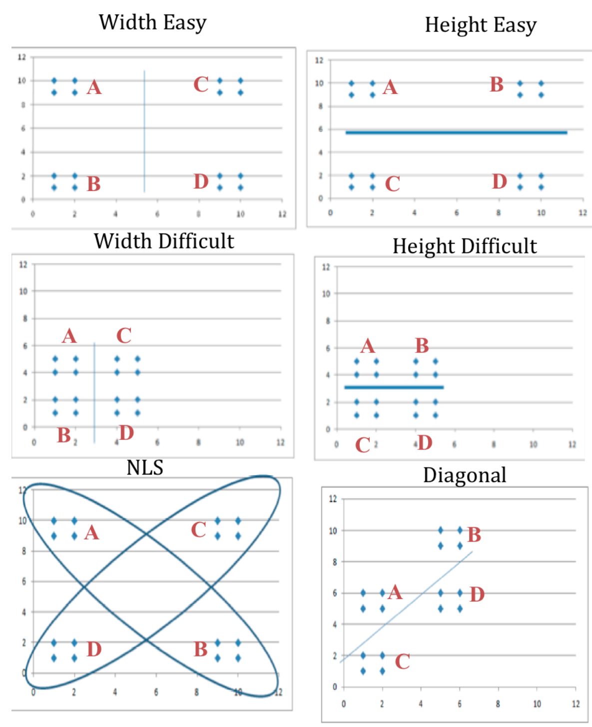

Figure 1. The six category structures employed in the study.

Finally, a third hypothesis is that similarity change depends on whether the category boundary is aligned with a dimension of physical variation (e.g. height). According to the COVIS model of categorization (Ashby et al., 1999; Ashby & Ell, 2002), whether the category boundary is along a physical dimension of variation can determine whether the executive (frontal) or the procedural system is engaged. To examine this hypothesis, the Diagonal category structure was compared against the condition best matched for difficulty with it (which turned out to be the Height Difficult category structure).

Participants and Design

One hundred and eighty experimentally naïve participants, all Swansea University students, were tested. There were 20 participants in each of the six experimental groups, each learning one of the six different classifications shown in Figure 1. For all six classifications, successful learning is achieved when participants recognize that items in clusters A and B are in one category and items in clusters C and D are in another category. The dependent variable was changes in similarity as a result of category learning. The procedure for computing changes in similarity is described in the Results section.

Three independent variables were considered, to allow examination of the hypotheses examined. The first was category structure learning difficulty, which was defined ad-hoc in terms of the number of trials to criterion. For this variable the Height Easy and Width Easy conditions were compared with their corresponding Height Difficult and Width Difficult conditions. The second independent variable was whether the category boundary was aligned with the dimension of physical variation or not. The third independent variable was linear separability with two levels, linearly versus non-linearly separable category structures.

For each of the six category structures, there was a corresponding control group providing similarity ratings for the stimuli, but without having gone through the categorization task first. For the Width Easy, Height Easy, and NLS conditions there was a common control group of 20 participants. For the Width Difficult and the Height Difficult conditions a different control of 20 participants, and for the Diagonal group yet another control group of 20 participants. The experiment lasted approximately 50 minutes for the experimental groups and 30 minutes for the control groups.

Materials

We used yellow surface-rendered arrow-like shapes that varied in terms of two dimensions: the width of the arrowhead (horizontal dimension) and the length of the arrow (vertical). The smallest arrow's trunk measured 4.5 centimeters (cm) in height and its head measured 3.0 cm wide. Twenty-four more stimuli were created by incrementing trunk height and head width by 12%. The stimuli employed in the experimental conditions were subsets of this original set of stimuli. The shortest arrow trunk in all six conditions was 4.5cm high and the narrowest arrow head 3.0cm wide. The tallest arrow trunk was 12.5cm in the Width Easy, Height Easy, Diagonal, and NLS conditions and 7.1cm in the Width Difficult and Height Difficult conditions. The widest arrow head was 8.3cm in the Width Easy, Height Easy, and NLS conditions, 4.7cm in the Width Difficult and Height Difficult conditions and 5.3cm in the Diagonal condition.

Procedure

A standard supervised categorization task was employed. A stimulus was presented at the center of a computer screen against a white background, until the participant decided whether it belonged to category A or B, at which point he/she received corrective feedback. Participants continued to categorize stimuli until no mistakes were made for 32 consecutive trials (i.e., all stimuli shown twice) or for a maximum of 256 trials. Five participants failed this criterion (three in the NLS condition and two in the diagonal condition) and these participants were not asked to complete the similarity part of the study. Participants, who completed the categorization task successfully subsequently received the similarity ratings task. In that task, each trial started with a 'Ready?' prompt at the center of the screen. Two stimuli appeared at the screen center for 500ms each, one after the other, with an inter-stimulus interval of 500ms. All possible 16x16=256 stimulus pairs were presented and participants were asked to rate their similarity on a 1-9 scale, such that 1 corresponded to 'very dissimilar' and 9 to 'very similar'. Participants were encouraged to use the entire scale. Participants in the control groups went through the similarity ratings, without having done the categorization task first.

Results

Data cleaning

There were two simple checks that the participants were sufficiently attentive during the similarity ratings task. Participants who did not use the whole similarity rating scale (1-9) and those who did not rate two identical stimuli as identical (by giving them an average rating of seven or above) were excluded from the data. This procedure led to the elimination of 3 participants from the Width Easy group, 3 from the Height Easy group, 2 from the Width Difficult, 1 from the NLS group, and 3 participants from the control groups.

Learning Results

Trials to criterion and errors correlated highly with each other (r=.84, p<.0005). Both the trials to criterion and the errors varied across category structures [$F(5,105)$=15.25, p=.0005 and $F(5,105)$=11.75, p=.0005, respectively]. The 'easy' versions of category structures were easier to learn than the 'difficult' versions of the classifications. Also, participants found it easier to learn the 'width' classifications than the 'height' ones, a result showing that the perceptual salience of the two dimensions was not equivalent. Category structures defined along a single dimension (Width easy/difficult & Height easy/difficult) were easier to learn than those defined along two dimensions (NLS and Diagonal), $t(109)$=5.94, p=.0001. There was no difference in ease of learning between LS category structures and the NLS one, $t(109)$=.48, p>.05. Finally, and as expected, the NLS and Diagonal classifications were the most difficult ones to learn, with no difference between them (p>.05).

Similarity measures

Change in similarity as a result of learning could be quantified in various ways. The measures typically employed in studies of changes of *perception,* as a result of categorization, emphasize discriminability along diagnostic vs. non-diagnostic dimensions (e.g., Folstein, Gauthier, & Palmeri, 2010; Goldstone, 1994). However, in the present study, any putative similarity changes as a result of categorization would relate to the categorization objective, that is, learning the different category structures. Therefore, it is more appropriate to consider a measure of similarity change, which is informed by the category structures in each case. Following theory on the determinants of category structure (Rosch and Mervis, 1975; Love, Medin, & Gureckis, 2004; Pothos & Chater, 2002; Pothos & Bailey, 2009) and categorization work in general (e.g., Mathy et al., in press), we employed two dependent variables for how the similarity structure might change as a result of categorization: within category and between category similarity change. Note that similarity in these definitions is empirical similarity from participant ratings. Within and between category similarity change allow us to directly explore the circumstances when the similarity structure for a set of stimuli becomes more consistent with a learned categorization. Within category similarity was the average similarity from all possible pairs in the same category and between category similarity was the average similarity ratings for all pairs across different categories.

Table 1. Trials to criterion, errors, and change in within and between similarity values, as a result of learning, for the six category structures employed in this study. The category structures have been ordered in terms of difficulty of learning. Asterisks indicate that the difference between experimental and control groups for each condition, revealed by independent-samples t-tests, was significant. Positive values indicate that the mean similarity value was higher for the experimental group compared to the control group, while negative values indicate the opposite.

			Similarity change	
Category structure	Trials to criterion	Errors	Within	Between
Width Easy	45.50	7.2	.28*	-.13
Width Difficult	71.00	5.6	.45*	.11
Height Easy	71.50	15.0	.19	-.07
Height Difficult	101.35	21.1	.59*	.11
NLS	102.60	34.6	.57*	.14
Diagonal	172.00	35.9	.41	-.43*

To provide baseline similarity values, within and between category similarity was calculated for the control participants, following the calculation procedure for their respective experimental groups.

Once similarity values were computed for all groups (both experimental and control), similarity *change* values were computed for each experimental group. Clearly, any changes in similarity as a result of categorization are only meaningful compared to a pre-learning baseline. For example, suppose a participant provided similarity ratings for the stimuli after learning the Width Easy classification. We would then compute her, e.g., *between similarity change* value as the between similarity value from her similarity ratings minus the average between similarity value of all corresponding control participants. Henceforth, when we refer to change in similarity values we imply similarity values computed in this way from the similarity ratings of the control participants, for each category structure. Adopting this analytical approach considerably simplifies comparisons of similarity changes across different category structures. Within and between similarity change can be understood as acquired equivalence and distinctiveness, respectively, but defined in terms of the learned categorizations, rather than stimulus dimensions.

Table 2. The F-tests examining the three hypotheses regarding similarity changes as a result of category learning.

	Similarity change	
Hypothesis	Within	Between
Learning difficulty (Height vs. Width, easy and difficult)	$F(1,68)=5.40$, $p=.02$	$F(1,68))=.92$, $p=.89$
Category boundary aligned with physical variation (Height Difficult vs. Diagonal)	$F(1,36)=.48$, $p=.49$	$F(1,36)=5.66$, $p=.02$
Linear separability (Height difficult & Diagonal vs. NLS)	$F(1,55)=.13$, $p=.72$	$F(1,55)=.04$, $p=.84$

Similarity Analyses

Similarity change for the six category structures we employed are shown in Table 1. For within similarity changes, positive values indicate changes in the similarity structure of the items more consistent with the learned classification. For between similarity changes, it is the other way round; between similarity is defined in terms of the similarity of items in different categories, so that if between similarity *is negative* this means that items in different categories become less similar (and therefore consistent with the learned classification).

The hypothesis that learning difficulty influences similarity change was examined in a 2 (Dimension: width vs. height) x 2 (Difficulty: easy vs. difficult) ANOVA. For within similarity changes, there was a significant main effect of Difficulty (shown in Table 2), with greater similarity change for difficult category structures than easy ones. There was no main effect of Dimension, $F(1, 68)=.03$, $p>.05$, nor a significant interaction, $F(1, 68)=.90$, $p>.05$. For between similarity change, there was no main effect of Difficulty, $F(1, 68)=.92$, p>.05, or Dimension, $F(1, 68)=.01$, $p>.05$, and no significant interaction, $F(1, 68)=.02$, $p>.05$.

The hypothesis that similarity changes would be determined by whether the category boundary is aligned with the dimension of physical variation, or not was examined in a one-way ANOVA (Height Difficult VS. Diagonal category structures). As shown in Table 2 this hypothesis was supported for between but not for within similarity changes. Finally, linear separability did not predict within or between similarity changes (Table 2).

Discussion

There has been considerable interest in changes in similarity (and perception) induced as a result of categorization, though few researchers have attempted a systematic study of the factors that make such changes likely (for an exception see Folstein et al., 2010). The overarching question in this research was whether the nature of the category structure is a relevant factor in trying to understand changes in similarity as a result of categorization. Three main possibilities were considered. The first possibility was that category difficulty would influence similarity change. We suggested that in cases where there are well-separated categories, similarity change may correspond more to within similarity change (cf. Chin-Parker & Ross, 2004), but for more poorly separated categories between similarity change may be more pronounced. In either case, more difficult category structures were expected to lead to greater similarity change. Our findings supported this hypothesis, but only partially. Difficulty of learning a category structure predicted changes in within category similarity, with stimuli in the same categories becoming more similar for more difficult, compared to the easier category structures.

The second possibility was that similarity changes are influenced by whether the category boundary was aligned with a dimension of physical variation. Indeed, this hypothesis was supported only for between category similarity change: when the category boundary was not aligned with a dimension of physical variation, then stimuli in different categories became less similar following categorization training. Although the influence of this factor could not be anticipated by prior work on similarity changes (e.g., work on categorical perception; Harnad, 1987), its role can be predicted within modern categorization theory (e.g., Ashby et al., 1998; Ashby & Ell, 2002). For instance, when the category boundary is aligned with a dimension of physical variation, even when the categories are poorly separated, participants focus on within category information, rather than on between category contrasts. Work on the COVIS model of categorization shows that category boundaries aligned with a dimension of physical variation are simpler than ones which are not, even for poorly separated categories (Ashby et al., 1999). Therefore, the complexity of the category boundary instead of the actual difficulty of the category structures (as defined in this study), may be a factor driving between similarity change. This possibility needs further work to be fully supported.

Finally, linear separability predicted neither within nor between similarity changes, even though this factor was manipulated across conditions, which were broadly matched for learning difficulty.

One debate in the literature concerns the extent to which similarity changes reflect perceptual changes, changes in item representation, changes in the category's internal structure, the addition of a label as a feature in determining similarity, or simply task demands. This is an important issue that is beyond the scope of the current research. It is important to note, however, that our finding that task difficulty influences the magnitude of similarity changes, is inconsistent with the view that similarity changes are due to the addition of category label to stimulus representations. That is, if a category label was added to stimulus representations in all cases, we should not have observed different degrees of similarity change for different category structures.

The current research revealed several methodological challenges in the study of changes in similarity as a result of categorization. First, several kinds of category structures are needed. Second, it is clearly of crucial importance to specify an appropriate index of similarity change, which takes into account possible differences between category structures. Indeed, in the present study, we did not observe equivalent results across the measures we introduced within and between similarity change. Researchers specifically interested in perception often consider acquired distinctiveness or equivalence, as a result of categorization (e.g., Goldstone, 1994; Harnad, 1987). Such measures are suitable when there are stimuli on either sides of a category boundary, but they are perhaps less suitable when the nearest neighbor stimuli on either side of a category boundary may be distant from each other in psychological space. This will often be the case for category structures that are meant to correspond to naturalistic ones (cf. Pothos et al., 2011).

A major methodological challenge was comparing effects of categorization on similarity for the different category structures. To do this we computed similarity values on the basis of similarity ratings, after they have been potentially modified by category learning (experimental participants) and without any categorization learning (control participants). Consequently, the dependent variables corresponded to the change of similarity ratings as a result of categorization. While we believe our solution to this problem to be adequate, it would be worthwhile to explore alternative approaches in future research.

Overall, the issue of whether some category structures are more or less likely to lead to corresponding changes in the similarity structure of categorized stimuli is a novel and exciting one. Here we presented a promising design to address it and a range of preliminary conclusions.

References

Ashby, G. F. & Ell, S. W. (2002). Single versus multiple systems of category learning: Reply to Nosofsky and Kruschke (2002). Psychonomic Bulletin & Review, 9, 175-180.

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

Blair, M. & Homa, D. (2001). Expanding the search for a linear separability constraint on category learning. Memory & Cognition, 29, 1153-1164.

Blair, M. & Homa, D. (2003). As easy to memorize as they are to classify: The 5-4 categories and the category advantage. Memory & Cognition, 31, 1293-1301.

Chin-Parker, S. & Ross, B. H. (2004). Diagnosticity and prototypicality in category learning: A comparison of inference learning and classification learning. Journal of Experimental Psychology: Learning, Memory, and Cognition, 30, 216-226.

Edwards, D. J., Pothos, E. M., & Perlman, A. (in press). Relational vs. absolute representation in categorization. American Journal of Psychology.

Folstein, J.R., Gauthier, I., & Palmeri, T. J. (2010). Not all spaces stretch alike: How the structure of morphspaces constrains the effect of category learning on shape perception. Vision Sciences Society, 10th annual meeting.

Freedman, D. J., Riesenhuber, M., Poggio, T., and Miller, E. K. (2003). A comparison of primate prefrontal and inferior temporal cortices during visual categorization. Journal of Neuroscience, 23, 5235-5246.

Goldstone, R. L. (1994). Influences of Categorization on Perceptual Discrimination. Journal of Experimental Psychology: General, 123, 178-200.

Goldstone, R. L., Lippa, Y., Shiffrin, R. M. (2001). Altering object representations thought category learning. Cognition, 78, 27-43.

Gureckis, T. M. & Goldstone, R. L. (2008). The effect of the internal structure of categories on perception. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), Proceedings of the 30th Annual Conference of the Cognitive Science Society (pp. 843). Austin, TX: Cognitive Science Society.

Harnad, S. (Ed.) (1987). Categorical Perception. Cambridge: Cambridge University Press.

Jiang, X., Bradley, E., Rini, R.A., Zeffiro, T., Vanmeter, J., and Riesenhuber, M. (2007). Categorization training results in shape- and category-selective human neural plasticity. Neuron 53, 891-903.

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

Mathy, F., Haladjian, H. H., Laurent, E., & Goldstone, R. L. (under review). Similarity-dissimilarity competition in disjunctive classification tasks.

McMurray, B., Aslin, R. N., Tanenhaus, M.K., Spivey, M. J., & Subik, D. (in press). Gradient sensitivity to within-category variation in words and syllables. Journal of Experimental Psychology: Human Perception and Performance.

Medin, D. L. & Schwanenflugel, P. J. (1981). Linear Separability in Classification Learning. Journal of Experimental Psychology: Human Learning and Memory, 7, 355-368.

Minsky, M. L. & Papert, S. A. (1988). Perceptrons. 3rd Edition. MIT Press: Cambridge, MA.

Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. Journal of Experimental Psychology: Learning, Memory, and Cognition, 10, 104-114.

Ozgen, E. & Davies, I. R. L. (2002). Acquisition of categorical color perception: A perceptual learning approach to the linguistic relativity hypothesis. Journal of Experimental Psychology: General, 131, 477-493

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

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

Pothos, E. M. & Bailey, T. M. (2009). Predicting category intuitiveness with the rational model, the simplicity model, and the Generalized Context Model. Journal of Experimental Psychology: Learning, Memory, and Cognition, 35, 1062-1080.

Pothos, E. M., Chater, N., & Stewart, A. J. (2004). Information about the logical structure of a category affects generalization. British Journal of Psychology, 95, 371-386.

Pothos, E. M., Perlman, A., Bailey, T. M., Kurtz, K., Edwards, D. J., Hines, P., & McDonnell, J. V. (2011). Measuring category intuitiveness in unconstrained categorization tasks. Cognition, 121, 83-100.

Roberson, D. & Davidoff, J. (2000). The categorical perception of colors and facial expressions: The effect of verbal interference. Memory & Cognition, 28, 977-986.

Roberson, D., Davidoff, J., Davies, I. R. L., & Shapiro, L. R. (2005). Color categories: Evidence for the cultural relativity hypothesis. Cognitive Psychology, 50, 378-411.

Rosch, E. & Mervis, C. B. (1975). Family Resemblances: Studies in the Internal Structure of Categories. Cognitive Psychology, 7, 573-605.

Rumelhart, D. E. & McClelland, J. L. (1986). Parallel distributed processing: explorations in the microstructure of cognition. Cambridge: MIT Press.

Ruts, W., Storms, G., & Hampton, J. (2004). Linear separability in superordinate natural language concepts. Memory & Cognition, 32, 83-95.

Schyns, P. G. & Oliva, A. (1999). Dr. Angry and Mr. Smile: When categorization flexibly modifies the perception of faces in rapid visual presentations. Cognition, 69, 243-265.

Schyns, P., Goldstone, R. L., & Thibaut, J. (1997). The development of Features in object concepts. Behavioral and Brain Sciences, 21, 1-54.

Sloutsky, V. M. & Fisher, A. V. (2004). Induction and categorization in young children: A similarity-based model. Journal of Experimental Psychology: General, 133, 166-188.

Smith, J. D., Murray, M. J. Jr.,& Minda, J. P. (1997). Straight talk about linear separability. Journal of Experimental Psychology: Learning, Memory, and Cognition, 23, 659-680.

Stevenage, S. V. (1998). Which twin are you? A demonstration of induced categorical perception of identical twin faces. British Journal of Psychology, 89, 39-57.