Perceptual Unitization in Part-Whole Judgments

Andrew T. Hendrickson (athendri@indiana.edu)
Department of Psychological and Brain Sciences, 1101 E. 10th Street
Bloomington, IN 47405 USA

Robert L. Goldstone (rgoldsto@indiana.edu)
Department of Psychological and Brain Sciences, 1101 E. 10th Street
Bloomington, IN 47405 USA

Abstract

Categorization relies upon the vocabulary of features that comprise the target objects. Previous theoretical work (Schyns, Goldstone, & Thibaut, 1998) has argued this vocabulary may change through learning and experience. Goldstone (2000) demonstrated this perceptual learning during a categorization task when new features are added that create a single feature unit from multiple existing units. We present two experiments that expand on that work using whole-part judgments (Palmer, 1978) to elicit the feature representation learned through categorization. The implications for different classes of computational models of categorization are discussed.

Keywords: Category Learning, Perceptual Learning, Perceptual Unitization

Introduction

Recent successful computational models in the categorization literature have represented objects as an arrangement of a set of features (Kruschke, 1992; Love, Medin, & Gureckis, 2004; Nosofsky, 1986; Spratling & Johnson, 2006). These models make the assumption that objects are automatically segmented into component features (or feature dimensions, depending on terminology) by the visual system. Kruschke (1992) argues that these features are psychological, not necessarily directly related to any particular physical property of the object, and may be a complex combination of low-level visual properties. The existence of psychological features has been inferred from behavioral measures including response patterns in visual search (Shiffrin & Lightfoot, 1997; Treisman & Gelade, 1980), the speed of classification (Goldstone, 2000), and patterns of classification (Schyns & Rodet, 1997). These features are defined by their function, how their presence influences behavior.

This class of categorization models relies on a set of perceptual features that is fixed at the beginning of category learning and does not change throughout the learning process. These models assume the perceptual system creates representations consisting of stable psychologically separable features that are available for further processing by the categorization system. The viability of the assumption of stable features is consistent with a lack of evidence in some paradigms for the creation of new detectors in primary visual cortex after repeated training (Petrov, Dosher, & Lu, 2005).

Evidence for Flexible Feature Vocabulary Empirical evidence from multiple sources is accumulating that the set of perceptual vocabulary of features does change over the course of learning a new task to include more diagnostic functional features. A flexible set of functional features, in which new features can be learned through experience, may underlie the perceptual vocabulary used in categorization (Schyns et al., 1998). Pevtzow and Goldstone (1994) demonstrated reaction time patterns in whole-part judgments which were consistent with different functional features being learned from the same set of training stimuli depending on the category structure. Similarly, Schyns and Murphy (1994) found error patterns and self-report statements consistent with participants forming stronger feature detectors for diagnostic stimulus fragments than non-diagnostic fragments.

Goldstone (2000) makes the strongest case for a flexible feature vocabulary with results showing reaction time patterns for classification of complex stimuli that cannot be accounted for by models of categorization that rely strictly on independently processing each feature. The results are instead consistent with the formation of new functional features that integrate information from previously separate features into one unit, a learning process referred to as perceptual unitization. In that study the stimuli were constructed by connecting five line segments and assigned to categories such that no individual line segment was predictive of category membership. Sets of segments, varying in size from 2 to 5 segments depending on the condition, must all be processed to correctly categorize each stimulus. The change in reaction times for categorization trials were not accurately predicted by an analytic model in which each necessary segment was processed independently and all information was aggregated after each segment was processed. A model in which the number of independent components that must be processed for a given stimulus decreases through learning more accurately accounts for the decrease in response time found with training. This decrease is proposed to be due to the perceptual unitization of previously independently processed functional features into single functional features that span multiple segments.

Goldstone (2000) goes on to investigate the necessary and sufficient conditions for perceptual unitization. The effects occurred in conditions where the individual segments were separated by blank gaps and not connected to each other as well as in conditions where the unitized segments were interleaved with and connected together by random non-diagnostic segments of the same size. Manipulating the number of segments that must be unitized to form a diagnostic functional feature (which was confounded with size of retinal image) produced a monotonically increasing function between number of segments and number of training trials needed to reach asymptotic reaction times. Evidence for perceptual unitization was not found in any conditions where the order of segments within the object was randomized. Together, these results suggest that a new perceptual unit is created when a stable, image-like pattern is often repeated and is diagnostic for a task. The time required to build this unit is proportional to the complexity of the unit.

Flexible Feature Sets in Categorization Models Models of categorization that incorporate fundamentally different mechanisms for flexible feature sets have been proposed to account for the empirical evidence of unitization and other perceptual learning. CPLUS (Goldstone, 2003) is a connectionist network that performs both categorization and object segmentation; it has been shown to account for the learned segmentation of diagnostic features from a whole object (Pevtzow & Goldstone, 1994) and learning novel complex diagnostic features but has not been extended to incorporate a mechanism for perceptual unitization of existing functional features. Other models, including those by Spratling and Johnson (2006) rely on attention weights, direct competition and lateral inhibitory processes between feature detectors within a hierarchical structure to model results similar to those addressed by the CPLUS model.

The assessment of these models has focused on the learned connections that define the set of features each model has learned rather than direct predictions of reaction time. The whole-part judgment task, first used by Palmer (1978) to assess the naturalness of different decompositions of visual objects into parts and subsequently by Pevtzow and Goldstone (1994) to look at the influence of experience on part decompositions, is ideal for measuring changes or differences in sensitivity to components of objects. This task consists of comparing a whole object to a part probe and asking participants if the part probe is a subset of the whole object. The whole object and the part probe may both be presented at once or in sequence. Correct answers rely on accurately comparing all segments in the part probe to the whole and determining if there is a match for each segment. In trials when the whole object is presented only before the part probe is present, whole-part judgments require a memory component as well as perceptual processing. The logic of whole-part judgment tasks relies on the assumption that decision processes will be more accurate if the part probe aligns with the functional features used to process and identify the whole object. The closer the part probe aligns to the existing functional features that encode the whole object, the more accurate judgments will be. This predicts that changes in the strength or vocabulary of functional features would be reflected in changes in the performance of whole-part judgments involving those features (Pevtzow & Goldstone, 1994). Whole-part judgments may provide complimentary supportive evidence for perceptual unitization that may more tightly constrain models than reaction time measures.

Experiment 1

In Experiment 1, a whole-part judgment task was used to assess the functional features after category training conducive to perceptual unitization. Both familiar segments, which were present during training, as well as unfamiliar segments not presented during training were tested. These were factorially combined with part probes differing in their number of segments and in the presence or absence of segments in locations that were not predictive of category membership during training. During that categorization training, participants learned to correctly assign eight objects composed of three segments into two categories. The category structure was arranged such that no segment was predictive in by itself, but the identity of two of the segments together were perfectly predictive of category membership, while the other segment was never predictive. These constraints produced an exclusive-OR (XOR) category structure in which exactly two segments must be identified correctly to make an accurate category prediction (see Figure 1 below). This XOR category structure differs from the classic Shepard et al. (1961) type II XOR category structure because without training or experience, the segments that compose these objects are not clearly separable features. Learning the structure in this experiment requires many more trials than even the most difficult type IV category structure requiring the memorization of all eight examples.

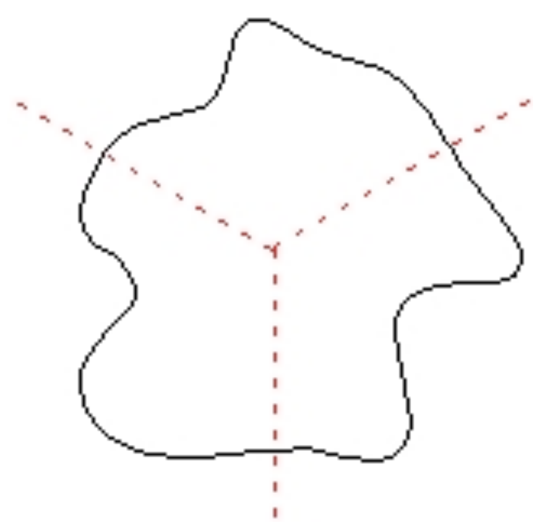

Figure 1: A scaled-down example of stimuli used in both experiments. The object was constructed by connecting three segments at specific locations. Dashed lines were added to indicate the points where the segments join.

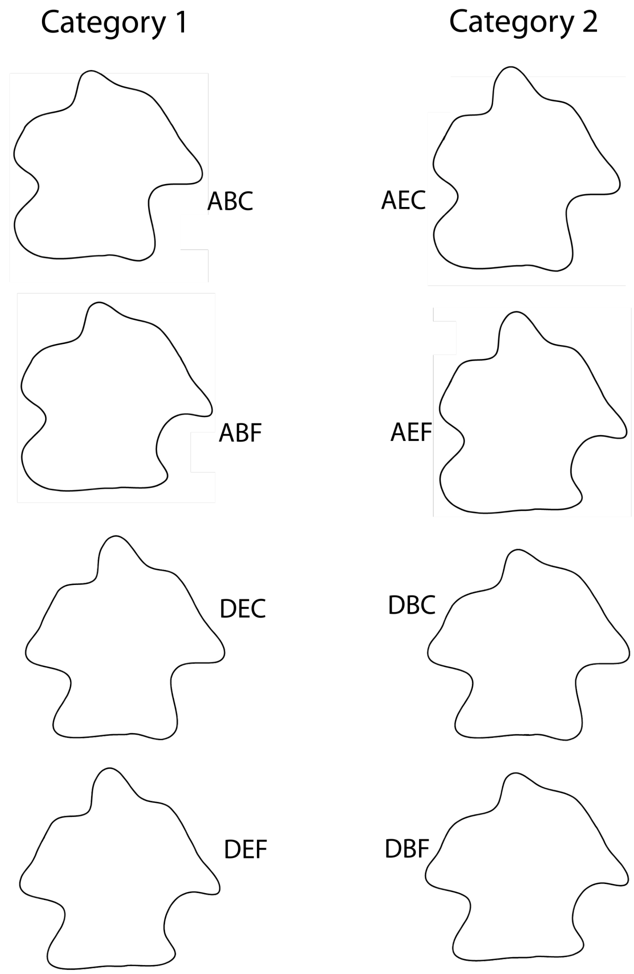

Figure 2: Stimuli and category structure used in Experiment 1. Each letter represents a unique segment and segment appears only in one location across all stimuli (A and D on the left, B and E across the top, and C and F on the right). To correctly categorize a stimulus in category 1, the presence of both segments A and B or D and E must be confirmed. Note that segments C and F provide no information about category membership and the location where they occur is never predictive.

Critical Tests of Model Predictions

Analytic Model with Independent Processing All of the category learning models discussed above that do not incorporate perceptual learning can be characterized at an abstract level as analytical models where each feature is detected independently and mappings are learned between those pre-determined features and the appropriate categories. These models predict improvements in speed and accuracy over the course of training for the processing and recognition of features through the systematic strengthening of association weights or allocation of selective attention to specific features. The crucial characteristic of this class of models is that each feature is processed independently of other features, regardless of connections to categories or attention allocation. If features align closely with the independent segments of the objects then these models, which process features independently, then they will always be more accurate on one-segment part probe trials compared with two-segment part probes. This is because each segment in the part probe must be matched to the corresponding segment in the whole by detecting each feature independently, regardless of association strength or selective attention allocation to the features. The independence of this process produces an overall error rate for the decision process that increases linearly with the number of features that must be matched correctly. This pattern would produce higher sensitivity on one-segment part probe trials compared to two-segment trials. Goldstone (2000) found evidence that this class of models were inconsistent with the pattern of reaction times found during category training using similar stimuli.

Categorization models which include mechanisms of perceptual unitization in which new functional features are learned that are predictive of category membership make the opposite prediction: they predict higher sensitivity for two-segment part probe trials compared to one-segment part probe trials when all segments in the part probe were predictive of category membership during training. The higher sensitivity should be limited to test trials where both segments of the part probe are in the predictive locations because this set of models predict that during categorization training only unitized functional features of the combination of predictive segments should be learned.

Selective Attention to Predictive Locations Both classes of models could be expanded to include a mechanism for assigning selective attention to the locations of predictive segments during learning. If this mechanism drives learning during the categorization phase then during the test phase sensitivity should be higher for all segments in the predictive locations, both for familiar and unfamiliar segments. Both perceptual learning and analytic models that do not have selective attention to location but rely only on learning connections between features will show much larger sensitivity for familiar predictive segments compared to unfamiliar predictive segments.

In summary, Experiment 1 uses a whole-part task to determine if participants learn to represent combinations of independently varying segments as functional features when the combination of those segments is predictive of category membership. By using randomly generated segments and creating arbitrary mappings to categories, it is very unlikely that functional features for these combinations of segments exist before the experiment. Specific contrasting predictions concerning the effect of the number and familiarity of the segments comprising the part probe in the whole-part task differentiate perceptual unitization learning models from analytical learning models, selective attention only models, and models relying on combinations of only independent features and selective attention.

Method

Participants Undergraduate students from Indiana University participated in this experiment to fulfill course credit. 47 participants completed the experiment within the allotted 60 minutes. All participants who did not reach the accuracy criteria within the allotted time during training were not included in any analysis.

Materials Stimuli were formed by combining three curved segments randomly without replacement from a set of nine segments. The angle of curvature was 120 degrees with a radius of 6.6 cm at the endpoints. Two segments in each stimulus were rotated such that the endpoints of all three segments aligned to create a closed object with length and height of 13.2 cm. The position and rotation angle of each segment was randomized across participants but constant for a participant. Participants viewed the display from a distance of approximately 45 cm, resulting in a viewing angle of 3 degrees for each object. Curved segments were connected at or near local maxima of curvature along each object because the ends of each segment were constrained to be locally convex curves. The locations at which segments were connected remained constant across all stimuli and participants.

Design Category membership for each object was determined by randomly selecting two segments not occurring at the same position for each subject. Each object that contained both or neither of those segments was assigned to one category and all other objects were assigned to another category; creating an exclusive-OR category structure with 2 of the 3 positions predictive of category membership (see Figure 2).

Categorization Procedure On each categorization trial, an object was presented in the center of the screen and participants were instructed to press one of two keys to indicate category membership for the object. Feedback indicating if the participant's response was correct was displayed while the object remained on the screen until the participant pressed a button to move to the next trial. Feedback was presented for 500 ms and was followed by a prompt for the participant to proceed to the next trial. A blank screen inter-trial interval of 750 ms preceded the next trial. Categorization trials were grouped into blocks of eight trials in which each unique object appeared once in a random order. Participants remained in the categorization phase until their accuracy was above 85% on four consecutive blocks.

Whole-part Procedure Participants were given two blocks of 192 trials, resulting in a total of 384 whole-part judgment trials. On each trial, an object (the whole) was presented for 1000 ms, followed by a blank screen for 750 ms, before a set of segments (the part probe) appeared and participants were instructed to determine if all segments in the part probe were present in the whole by pressing one of two keys for "match" and "do not match." All whole object or part probe stimuli were presented in the center of the screen with random jitter of up to 0.5 cm in any direction. Participants were not provided with any feedback concerning their response. A blank screen inter-trial interval of 750 ms preceded the subsequent trial. Participants were instructed to proceed as quickly as possible without sacrificing accuracy in their responses; accuracy and response times were collected. After every 50 trials, a short break was provided.

Four factors were manipulated independently within each block of trials to determine the composition of the part probe: 1) Number of segments in the part probe: on half the trials the part consisted of one segment, otherwise it consisted of two segments, 2) Familiarity of segments: on half the trials the segments in both the whole object and part probe were from the set of segments in category training, the other half of trials used no familiar segments in either the whole or the part, 3) Location of segments: segments in the part probe appeared at all three positions of an object equally often but consistently in the same location. 4) Correct answer: on half of the trials all segments in the part probe matched the segments in the whole and the correct answer was yes, otherwise one segment was replaced with a non-matching segment of the same familiarity and assigned to the same location to create "no match" trials. The order of trials was randomized within each block.

Results

A 2 (both the part probe and whole object consisted of familiar vs. unfamiliar segments) X 2 (number of segments comprising the part probe: one or two segments) X 2 (all segments comprising the part probe in category-predictive locations vs. at least one segment of the part probe in the non-predictive location) analysis of variance (ANOVA) was conducted with sensitivity (d-prime) from signal detection theory as the dependent measure. Sensitivity combines information from trials where the correct answer is yes and no.

A significant main effect of number of segments in the part probe was found, $F(1,46) = 4.64$ ($p = 0.036$) with mean sensitivity of one segment trials was 1.23 and mean of two segment trials 1.36. A main effect of category-predictive location was found, $F(1,46) = 5.43$ ($p = 0.024$) with mean sensitivity on trials in which all segments in the part probe were in predictive locations was 1.39 and the mean where the part probe contained a segment in the non-predictive location was 1.21. A non-significant trend toward a main effect of segment familiarity was found, $F(1,46) = 3.41$ ($p = 0.07$) with mean sensitivity of 1.41 for familiar segments and 1.18 for unfamiliar segments.

The main effect of category-predictive location was modulated by a two-way interaction with segment familiarity, $F(1,46) = 16.58$ ($p < 0.001$) (see Figure 3). The main effect of number of segments in the part probe was also modulated by a two-way interaction with segment familiarity, $F(1,46) = 4.09$ ($p = 0.049$) (see Figure 4). There was no significant interaction between category-predictive location and number of segments $F(1,46) = 1.36$ ($p = 0.25$). The three-way interaction was non-significant $F(1,46) = 2.52$ ($p = 0.12$).

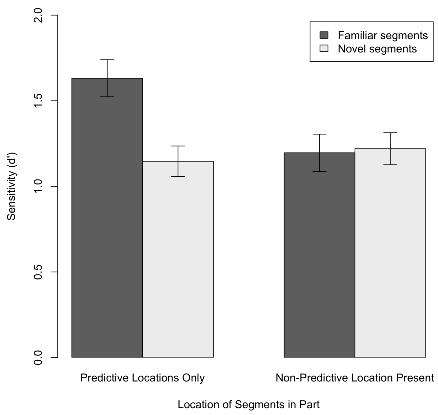

Figure 3: Interaction between the presence of a segment in the non-predictive location of the part and the familiarity of the segments on sensitivity (d-prime). The sensitivity was much higher for parts with familiar segments when all segments were in predictive locations.

Experiment 1 Discussion

The results of Experiment 1 show strong evidence that unitized features comprised of two segments were learned during the category-training phase. Figure 3, showing the strong interaction between segment familiarity and the location of segments in the part probe indicates that sensitivity for familiar segments in predictive locations was much higher than any other combination. This pattern of results is not consistent with the predictions of a learning model based entirely on the reallocation of spatial attention to the category-predictive locations of objects. Instead, a higher sensitivity was only shown for familiar segments in those locations, indicating that something about those segments was learned.

The significantly higher sensitivity to two-segment part probes compared to one-segment, indicates that what is being learned is not individual features for each independently-varying segment, as suggested by an analytical model. Instead, these results are more consistent with an account where pairs of independently varying segments are processed as unitized features.

One potential objection to the conclusions from Experiment 1 is that the pattern of results was perhaps a function not of changes in processing during the category-learning phase, but of specific properties of the stimuli themselves or the whole-part protocol. To address this concern, Experiment 2 presents the whole-part judgment task from Experiment 1 to participants who have not experienced the category-learning phase.

Experiment 2 was a control condition for Experiment 1. All the experimental methods from Experiment 1 were repeated in 2 except participants did not participate in the category training phase and proceeded directly to the whole-part judgment phase.

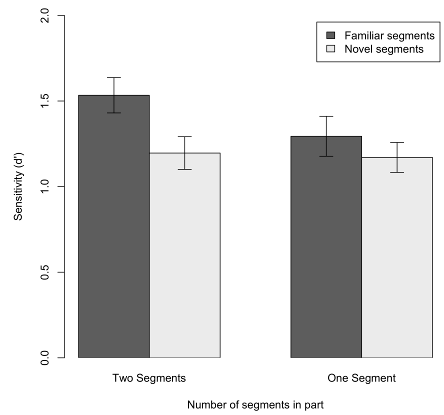

Figure 4: Interaction between the number of segments in the part and the familiarity of the segments on sensitivity (d-prime). The sensitivity was much higher for parts with familiar segments when the part consisted of two segments.

Experiment 2

Without the category training phase no effect of category-predictive location should be found. Additionally, without any exposure during category learning, no difference should be found between familiar and unfamiliar segments. Without the opportunity to learn unitized features in training, the trend from Experiment 1 that higher sensitivity was shown for parts consisting of two segments compared to one segment should be reversed in Experiment 2.

Method

Participants 47 undergraduate students from Indiana University participated in order to partially fulfill course credit.

Materials The exact stimuli from Experiment 1 were used. Despite no category training, two segment locations were randomly assigned to be predictive of category membership and six segments were assigned to be the set of familiar segments, as in Experiment 1.

Categorization Procedure No category-learning phase occurred. Participants proceeded directly to the whole-part procedure.

Whole-Part Procedure Participants were given at least two blocks of 192 trials, resulting in a total of 384 whole-part judgment trials. The whole-part procedure from Experiment 1 was duplicated for the first 384 trials. If time permitted, participants did further blocks but those results are not included in any analysis.

Results

A 2 (familiar vs. unfamiliar) X 2 (part probe size) X 2 (all category-predictive segments vs. at least one non-predictive) ANOVA was conducted with d-prime as the dependent measure. No significant main effect of number of segments in the part probe was found, $F(1,46) = 0.37$ ($p = 0.54$). No significant main effect of category-predictive location was found, $F(1,46) = 0.01$ ($p = 0.96$). No significant main effect of segment familiarity was found, F(1,46) = 2.07 (p = 0.16). No two-way or three-way interactions were significant using a criterion of 0.05.

General Discussion

The significant effects and trends in the results from Experiment 1 are not replicated in Experiment 2, indicating that nothing in the stimuli or testing procedure is able to account for the results in Experiment 1. Consistent with Goldstone (2000), analytic models that rely exclusively on functional features that do not span multiple segments are not able to account for the greater performance on trials containing two predictive-segments in the part probe. The addition of spatial locations-based selective attention mechanisms in these analytic models does not address this shortcoming because of the strong interaction of the observed effects with familiarity.

One class of analytic models that can account for the pattern of results in Experiment 1 are those that do not create new unitized features but initially include functional features that span the small areas where segments connect. Slowly learning to heavily weight existing features that span those locations would produce results that cannot be discriminated from a perceptual unitization process in this data. The lack of main effect in Experiment 2 between parts of size one and two lend slight support to the argument that participants may not naturally decompose these objects into the independently varying segments. However, this class of models would fail to account for the perceptual unitization of up to five connected segments into one functional feature found in Goldstone (2000). Separating this account from the predictions of perceptual unitization models within the whole-part framework will require the manipulation of category structure or multiple phases of whole-part judgments over the course of category learning in future experiments.

Further work on identifying individual differences within the degree to which functional features are strongly represented is also suggested by this framework. Future directions should include the identification of which specific functional features an individual is most sensitive and relating systematic differences in performance in the category learning phase to the test phase performance. This work also provides clear results for the application of cognitive models of categorization and perceptual learning. The class of models that represent the independently varying segments of these objects as separate functional features will not adequately capture the pattern of results found, even with the inclusion of selective attention to those features or to spatial locations. Models of categorization that learn vocabularies of functional features that span familiar predictive segments and are learned during category training are consistent with these results. Combined with the results of Goldstone (2000), these experiments strongly support models of categorization that include mechanisms for the perceptual unitization of smaller functional features into larger features during category learning.

References

Goldstone, R. L. (2000). Unitization During Category learning. *Journal of Experimental Psychology: Human Perception and Performance, 26*(1), 86-112.

Goldstone, R. L. (2003). Learning to Perceive While Perceiving to Learn. In R. Kimchi, M. Behrmann & C. Olson (Eds.), *Perceptual organization in vision: Behavioral and neural perspectives* (pp. 233-278). New Jersey: Lawrence Erlbaum Associates.

Kruschke, J. K. (1992). ALCOVE: An Exemplar-Based Connectionist Model of Category Learning. *Psychological Review, 99*(1), 22-44.

Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A Network Model of Category Learning. *Psychological Review, 111*(2), 309-332.

Nosofsky, R. M. (1986). Attention, similarity, and the indentification-categorization relationship. *Journal of Experimental Psychology: General, 115*, 39-61.

Palmer, S. E. (1978). Structural aspects of visual similarity. *Memory & Cognition, 6*(91-97).

Petrov, A., Dosher, B., & Lu, Z.-L. (2005). The dynamics of perceptual learning: an incremental reweighting model. *Psychological Review, 112*(4), 715-743.

Pevtzow, R., & Goldstone, R. L. (1994). *Categorization and the parsing of objects.* Paper presented at the Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society.

Schyns, P. G., Goldstone, R. L., & Thibaut, J.-P. (1998). The development of features in object concepts. *Behavioral and Brain Sciences, 21*, 1-54.

Schyns, P. G., & Murphy, G. L. (1994). The ontogeny of part representation in object concepts. In D. L. Medin (Ed.), *The Psychology of Learning and Motivation* (Vol. 31, pp. 301-349).

Schyns, P. G., & Rodet, L. (1997). Categorization creates functional features. *Journal of Experimental Psychology: Learning, Memory & Cognition, 23*(33), 681-691.

Shepard, R. N., Hovland, C. L., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs: General and Applied, 75*(13), 1-42.

Shiffrin, R. M., & Lightfoot, N. (1997). Perceptual learning of alphanumeric-like characters. In R. L. Goldstone (Ed.), *Perceptual Learning: The Psychology of Learning and Motivation* (Vol. 36, pp. 45-81): Academic Press.

Spratling, M. W., Johnson, M. H. . (2006). A feedback Model of Perceptual Learning and Categorization. *Visual Cognition, 13*(2), 129-165.

Treisman, A., & Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology, 12*(97-136).