The Ontogeny of Transformable Part Representations in Object Concepts

Philippe G. Schyns and **Gregory L. Murphy**
Center for Biological and Computational Learning
Massachusetts Institute of Technology
Cambridge, MA, 02139
schyns@ai.mit.edu

Beckman Institute
University of Illinois
Urbana, IL 61801
gmurphy@s.psych.uiuc.edu

Abstract

Many theories of object categorization assume that objects are represented in terms of components--attributes such as parts, functions and perceptual properties. The origin of these components has been neglected in concept learning theories. We provide further evidence for a theory of part ontogeny in which parts are not simply given by perceptual information but develop over the course of category learning. Two experiments are reported. The experiments showed that subjects could identify a component of unknown stimuli as a part in spite of variation in its shape across exemplars of the category. However, the experiments also revealed perceptual constraints on what variations could be identified as the same part.

Introduction

In the traditional artificial category-learning experiment (e.g., Bourne, 1982; Bruner, Goodnow & Austin, 1956), subjects were often shown very simple stimuli, such as colored geometric stimuli which varied on a few obvious dimensions. Unfortunately, when encountering real-world objects, the set of stimulus dimensions might not be limited in the same way as in such experiments. For example, reading x-rays films, identifying bacteria in a microscope, or the categorization of pre-Cambrian organisms may require the use of new stimulus dimensions for their classifications--stimulus dimensions that were not used prior to the experience with the classification system. Not surprisingly, complete novices who examine chest x-rays have little understanding of what the relevant features are: Lesgold et al. (1988, p. 325) discuss an x-ray that "depicts a patient with multiple densities in the lungs, which represent tumors and other local pathologies... in this film, the novices merged local features (e.g., tumors) and saw a general lung haziness that is probably not present to the extent they thought it was." Although there were perceptual cues in the x-ray that experts noticed, these are not perceptual features that we have pre-existing representations of. If one takes a developmental perspective, it seems clear that young infants are not aware of all the stimulus dimensions that will become important in later category learning (Clark, 1983; Mervis, 1987). In short, we are arguing that it is simplistic to assume that the features that are the basis for concept learning are somewhat clearly given in the environment. Thus,we need an account of how the features themselves are acquired. This paper addresses the nature of the features of object representations that have a privileged role in theories of object recognition and categorization: the object's parts (Biederman, 1987; Hoffman & Richards, 1985; Marr & Nishihara, 1978; Tversky & Hemenway, 1984). We will be interested in the learning of unknown parts, when part shape changes in the object category.

In Schyns and Murphy (1991, 1992), we proposed that an alphabet of primitives of object representation could be learned progressively and developed as the organism categorizes its world. To see how this might work, consider the following thought experiment. Imagine you were a Martian with no experience of Earth objects. The first object you encounter is a sedan car which can be segmented into the parts hood, roof, trunk and wheels. Mars has no car, and you do not have conceptual tools to describe this object. However, supposing your perceptual organization were roughly like ours, you could at least notice "bumps" and "valleys " in the silhouette of the sedan from the side. Since valleys usually occur where parts interpenetrate, it should be straightforward to learn which bump is a hood, a roof, a trunk or a wheel. However, this particular part decomposition is by no means necessary to identify correctly a sedan--you might not know any other object with a similar silhouette. On the second day of your Earthen life, you learn about convertible cars. From the side, the silhouette of a convertible misses a bump: the roof. From now on, you could use the primitive *roof* because it discriminates the two categories of objects: one has a roof and one does not. Before encountering the convertible, you could have represented *sedan* in memory as a single, holistic unit of representation, because there was no need to segment it.

The Martian example illustrates two important constraints on learning part representations in object concepts. As Hoffman and Richards (1985) suggested, perceptual constraints, such as the valleys in an object silhouette--more generally, the loci of minima of Gaussian curvature of a surface--may indicate perceptual breaks for part decomposition. However, as shown in the Martian example, part learning also arises from the need to categorize and to represent objects with an efficient coding scheme. This categorical constraint is summarized in the *Homogeneity Principle* (Schyns & Murphy, 1992):

> The Homogeneity Principle: If a fragment of stimulus plays a consistent role in categorization, the perceptual parts composing

the fragment are instantiated as a single unit in the stimulus representation in memory.

To test various predictions of the Homogeneity Principle, we presented subjects with categories of unknown objects called "Martian rocks." A Martian rock was composed of unknown parts separated with perceptual breaks (minima of Gaussian curvature). Using a computer mouse, subjects could interactively delineate the parts they thought were the most relevant to the category. When a conjunction of parts separated by a perceptual break consistently categorized the objects, subjects did *not* instantiate many separate perceptual parts, as a perceptual account of part decomposition would predict. Instead, much like the Martian of the previous example considered *sedan* as a single unit of representation as long as *hood* did not start to categorize other objects, subjects considered as a unified unit of representation the two perceptual parts that performed the relevant categorization, as long as one of the parts did not start to categorize other objects.

These experiments showed that perceptual constraints *and* categorical constraints were necessary to account for the components of objects concepts. However, the learning task was quite unrealistic: the relevant parts were constant in size and shape throughout the category. This situation is not quite realistic. Consider for example the part *leg*. Even in the restricted domain of the animal kingdom, legs vary along a number of dimensions, including length, width and relative proportions of thigh and calf. Within the same category of objects, a part can vary dramatically in shape, but people can identify the part in spite of this variation. Variation in part shape makes identification of a single part across examples difficult. Are certain variations easier to recognize than others? Would our perceptual system see invariants through certain variations of parts but not others? Richards, Dawson and Whittington (1988) suggested that extrema of curvature of an object's 2-D outline convey relevant information about the object's shape. Extrema of curvature are of three types: positive, negative and zero. Formally, extrema of curvature occur where the tangent to the curve rotates at the greatest rate. More concretely, a local minimum of curvature is typically associated with a concavity--a "valley"--and a local maximum with a convexity--a bump. When following the curve from a maximum to a minimum, the rate of change of its tangent goes from a high positive value to a high negative value. The point at which the rate of change equals zero--the inflection point--is a zero of curvature. We saw earlier that minima of curvature could be useful for part segmentation. Maxima and zeros are also interesting for describing the silhouette of objects because extrema of curvature are invariant under rotation, translation and uniform scaling. Because extrema are invariant over such quantitative changes in scale, they form a qualitative description of an object's outline. If people use such a qualitative description to make perceptual judgments, parts that could be subsumed into a single unit of representation should be those with a similar sign of curvature description. To illustrate the point, consider again the silhouette of a sedan seen from the side. Even though the trunk and hood can vary in size in different cars (a quantitative change), the extrema of curvature of its outline (the qualitative description) is constant across most sedans. However, changing the curvature (e.g., making a bump a valley or a straight line as in the sedan vs. convertible example) would change the qualitative description by changing the sign of curvature.

Experiment 1

The following experiment had two main goals. The first was to demonstrate that subjects can use the Homogeneity Principle to abstract features from category exemplars even when their shapes vary. The second was to test the sign of curvature as a perceptual constraint on part learning. The stimuli of the experiment were 2D outlines of "Martian rocks" composed of a target part and a random part separated by a neck (see Figure 1, top panel). There were two phases in the experiment. In the first, subjects viewed different types of transformation of the target part. In the second phase, we tested subjects on previously seen and unseen transformations of the test part. We hoped to find that a transformation preserving the sign of curvature (Figure 1, middle row, left two pictures) facilitates part extraction, whereas a transformation changing the sign (Figure 1, the middle right picture) complicates part extraction. This was accomplished by creating three categories of objects named *concave, convex* and *mixed.* We expected subjects who learned the concave and the convex categories (the ones without a change of sign of curvature in the part) to accurately categorize new and old exemplars, and to demonstrate a prototype effect in the testing phase. Conversely, subjects in the mixed condition should have more difficulty in forming a representation of the target part, and therefore should perform poorly in categorization.

Method

Subjects. Subjects were 18 Brown University students who were paid to participate in the experiment. They were randomly assigned to the three category types with the constraint that the number of subjects be equal in each condition.

Stimuli. "Martian rocks" were generated using Uniform non-Rational B-Splines (URBS) on a Silicon Graphics computer. A B-Spline is a parametric cubic curve that approximates a set of control points. The shape of a B-Spline is transformed by moving its control points. Each Martian rock was generated with a total of 21 control points. 18 of the points are shown in the top section of Figure 1. (In order to obtain a closed cubic curve, the remaining points, not visible on Figure 1, were always duplicates of the three first points.) Each object consisted of a *random* and a *target* part separated by a "neck." The points controlling the shape of the random part (labelled 16, 17, 0, 1, 2, 3 and 4 on Figure 1, top panel) could vary within a certain region with the constraint that they do not cross regions. To illustrate, imagine that the part was divided into eight pie slices, with each control point at the center of one slice. In learning exemplars, each control point was located according to a random selection of length

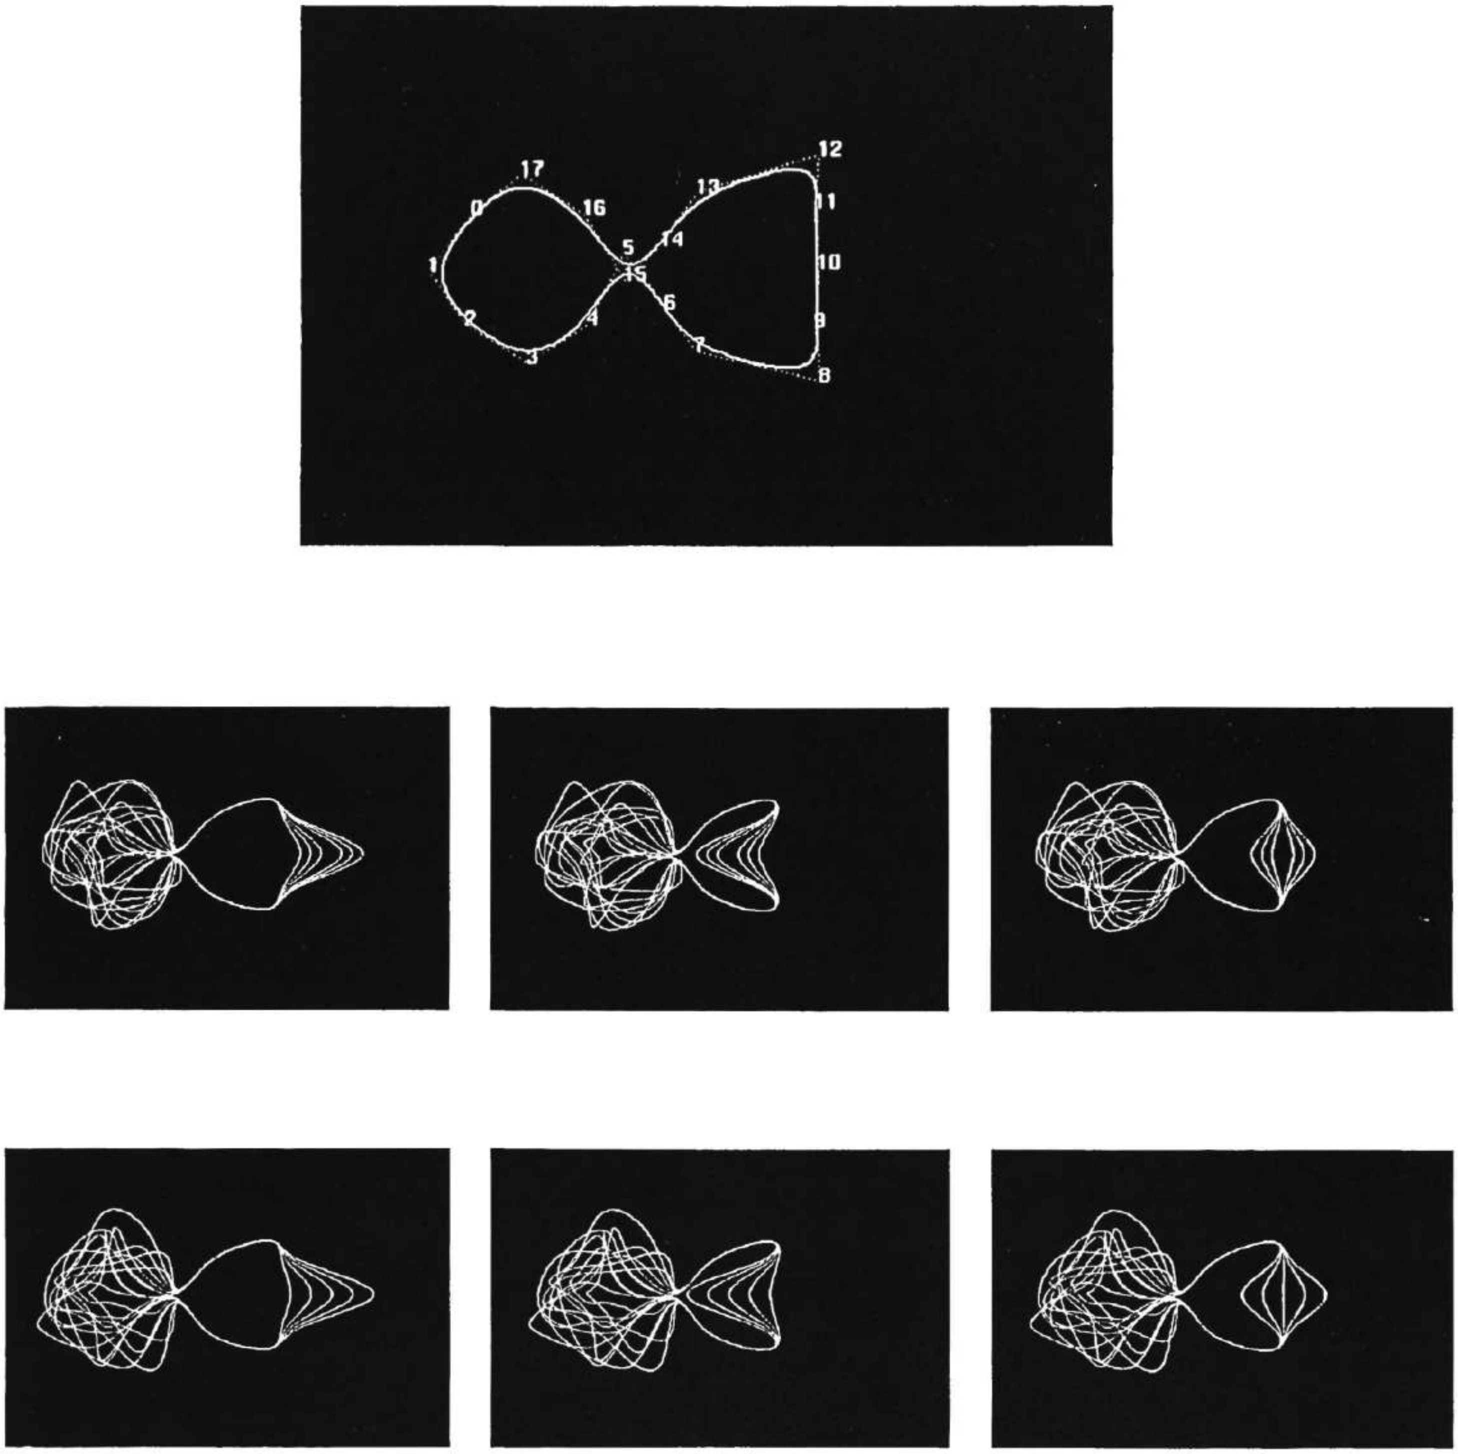

Figure 1. This figure illustrates the stimuli used in Experiment 1 and 2. The top picture shows the cubic B-spline from which the categories were generated. The control points numbered 16, 0, 1, 2, 3 and 4 were used to create the random part, and control point 10 was translated to generate the target part. The middle pictures illustrate the set of transformation experienced during the learning phase of Experiment 1 and 2. The middle pictures represent, from left to right, the training exemplars of conditions convex, concave and mixed. The sets were obtained by superimposing the learning exemplars. The bottom pictures show the corresponding test items used in Experiment 1 and 2.

and angle with the constraint that it does not leave its slice. The point controlling the shape of the target part (labelled 10 in the figure) was slid horizontally inward in the concave condition while it was moved symmetrically outward in exemplars of the convex condition. This point was moved by successive constant increments in the two categories, augmenting the slope of the curves, while keeping their sign of curvature constant. The middle panel of Figure 1 shows the transformations of all exemplars simultaneously by superimposing them. Each real exemplar consisted of only one such transformation. As can be seen, the random (left) part has little apparent regularity across exemplars, whereas the target (right) part varies according to the transformations just described. A similar procedure was applied to generate exemplars in the mixed condition, but the control point was moved both inward and outward, resulting in a change of sign of curvature. In all three categories, the total amount of displacement of the control point was identical. Note that the center of the part transformations (the prototype of the part) was never presented during training.

Test items were constructed from items seen during learning and new part transformations. One type of transformation, called *compatible*, had the same sign of curvature as the transformations in the learning set (i.e., convex transformations in the convex category, concave in the concave category). The range of compatible transformations was greater than that of the transformations experienced during learning: The relevant control point was slid inward or outward by one increment from the innermost or outermost learning transformations, so that it would be outside the range of the learning transformations (i.e., one transformation was smaller and one was larger than in learning). Another type of transformation was *incompatible* with the transformations that generated the parts in the learning set (i.e., convex parts in the concave category, and vice versa). The mixed category only had consistent transformations, as it contained both convex and concave parts. The final test item type was the category prototype, consisting of the center of all the transformations used to create the learning items. The prototypes were not shown during learning.

To illustrate the test stimuli, consider the convex condition (the leftmost picture of the bottom panel of Figure 1). Five test exemplars are shown, each corresponding to a different location of the control point. The leftmost and rightmost outlines are extrapolations compatible with the training set . Next to them are old exemplars, and finally, the center outline is the interpolation of the training set, the prototypical part never shown during learning.

There were 40 test items, 20 true M2 rocks and 20 false items. The 20 true items consisted of 8 old, 8 consistent and 4 prototype items. For the concave and convex categories, the false items consisted of 4 incompatible transformations of the part, plus 16 other items in which both the target part and the rest of the rock was randomized. For the mixed items, there were no incompatible transformations, and so all 20 false items consisted of the latter type.

Procedure: Learning Phase. Subjects were instructed that they would see 48 members of the same category of Martian rocks called M2 and that they should study the rocks and to learn what M2 rocks were like in general. Each of the six possible training variations appeared eight times in a random order. Each of these rocks appeared one at a time for 2 seconds on the computer monitor.

Testing Phase. In order to measure what subjects had learned about the category, a categorization task with 40 exemplars was conducted. One test rock at a time was randomly selected and presented on the computer screen. Subjects had to decide, by pressing the "yes" or "no" key on the computer keyboard, whether the rock was an M2. No feedback was given during the testing phase, and subjects were instructed to respond as quickly and as accurately as possible.

Results

The quality of part representation should be reflected in categorization performance in the test phase, since the categories were defined by the presence of the target part. On average, subjects from the *concave*, *convex* and *mixed* conditions categorized their stimuli with 97%, 99% and 79% accuracy, respectively. A one-way analysis of variance showed a significant effect of condition on the mean percent of correct response ($F(2, 15) = 11.65$, $p < .01$). A post-hoc Fisher protected t-test ($LSD = 12.97$) revealed significant differences between the mean percent correct of the mixed and the other two conditions (p's < .01), but no significant difference between the convex and concave conditions. No subject categorized the extrapolations with change of sign of curvature as M2 rocks.

To test for a prototype effect, we compared reaction times for each type of true test item. Mean correct RTs to old, new and prototypical exemplars are summarized in Table 1. A two-way analysis of variance revealed a main effect of condition ($F(2, 15) = 8.0$, $p < .01$), a main effect of exemplar type ($F(2, 30) = 8.18$, $p < .01$) and a significant interaction of condition and exemplar type ($F(4, 30) = 4.19$, $p < .01$). As shown in Table 1, the mixed condition is slowest overall. But more importantly, the prototype was the fastest item type for the convex and concave conditions, but was the slowest type for the mixed condition.

Table 1. Mean reaction times in milliseconds for the correct categorization of exemplars in the tree conditions of Experiment 1.

	Convex	Concave	Mixed	Mean
Old	593	664	818	692
New	760	803	967	843
Prototype	592	603	1054	750
Mean	648	690	946	

Discussion

The most important result from this study is the finding that people could identify outlines that varied in shape as being the same part. Subjects had no difficulty identifying what was common to the category when the part varied "in the same direction," either concave or convex. But when the part crossed the perceptual boundary (i.e., when the sign of curvature changed), subjects were much less accurate and were slower at identifying the object as being in the same category. Since category membership depended on the presence of this part, these measures indicate how easy subjects found it to identify the part.

This difficulty of learning the part can be indirectly observed in subjects' responses to the prototype. In the convex and concave conditions, prototypical exemplars never experienced before were recognized as quickly as old exemplars. This result is compatible with a prototype representation of the part. Subjects in the mixed condition responded more slowly to their prototype, suggesting that they might have formed not one, but two distinct representations of the part: one with a negative sign of curvature, one with a positive sign. Since the prototype in the mixed condition is neither concave nor convex, it is difficult to categorize, even though it is the middle item of the presented parts.

In the concave and convex conditions, subjects never categorized test exemplars with different signs of curvature as M2 rocks. In contrast, extrapolations that preserved the sign of curvature were accepted as category members. Thus, although variation in the part did not prevent subjects from noticing the constant part, not all variation was perceived as the same part. It appears that the category context created one set of constraints, and the perceptual restrictions on changing sign of curvature created another.

Experiment 2

The previous experiment suggested that when the common part of a category did not follow a perceptual constraint (keeping the same sign of curvature), subjects found it difficult to learn the category. However, it is not entirely clear how this was related to the extraction of a part or parts in the stimuli. One possibility is that when the sign of curvature was respected (in the convex or concave conditions), subjects simply learned that there was a single part, which varied in the amount of its curvature. In the mixed conditions, subjects may have learned that the category had two different parts (the convex and concave versions), each of which also varied in the amount of its curvature. If this is the case, then after category learning, subjects should differ in whether different versions of the part are perceived as the same or different, and this should have effects on whether these versions can be used in learning new categories. In particular, if the subjects in the mixed condition encoded two different parts, then they could have easily learned two new categories that differed in the presence of these parts. In contrast, subjects in the concave or convex conditions would not have formed such a distinction and so would have difficulty learning new categories based on differences in part variation.

In Experiment 2, subjects first learned a single category of concave, convex or mixed stimuli, as in the previous experiment. In a second phase, they then learned to divide these stimuli into two subcategories: the three "leftmost" transformations of the part vs. the three "rightmost" transformations of the part (see Figure 1). In the case of the mixed categories, this corresponded to parts with different signs of curvature (i.e., one subcategory had the convex and the other had the concave transformations). If subjects had already discriminated these two transformations into two conceptual parts, then it would be very easy to learn the new categories, since the categories could be distinguished on the basis of these two features. If they had not distinguished the transformations, then learning new categories should be difficult. We expected that subjects in the mixed category would have formed such parts, whereas subjects who had learned concave or convex transformations would not have already formed such parts, and so should find learning the subcategories more difficult.

Method

Subjects. Subjects were 18 Brown University students who were paid for participating. They were randomly assigned to the concave, convex and mixed conditions.

Stimuli. The stimuli were the same as in Experiment 1. Uniform non-rational B-Splines composed of a target and a random part were transformed as explained earlier to generate the concave, convex and mixed categories. Stimuli were displayed on the screen of a Silicon Graphics computer.

Procedure: First learning phase. Subjects were exposed to exemplars of Martian rocks, some of them forming a category called M2. Subjects were told to study the rocks and to learn what M2 rocks were like in general in order to categorize them correctly. The learning phase was divided into blocks, consisting of 24 rocks, divided into 12 M2 and 12 random rocks. The M2 rocks were two repetitions of the six possible transformations of the first Phase of Experiment 1 (see Figure 1, the middle pictures). The random rocks were generated as explained earlier, by randomizing the control points of the target part. By pressing the "yes" or "no" key on the computer keyboard, subjects had to indicate the category membership of each rock. To reach criterion, they had to be correct on all 24 exemplars of a learning block. If they made a mistake, they continued the learning phase until reaching criterion.

Second learning phase. The second learning Phase started upon completion of the first learning phase. Subjects were instructed that in fact, M2 rocks could be subdivided into two subcategories M2-*a* and M2-*b* and that they would learn to categorize M2-*a* and M2-*b* rocks correctly. The learning procedure was similar to the one of the first learning Phase, but subjects saw only M2 rocks. A learning block was composed of 12 exemplars of M2 rocks, divided into 6 M2-*a* and 6 M2-*b*. In the *mixed* condition, M2-*a* rocks were two repetitions of the three concave transformations, and M2-*b* rocks were two repetitions of the three convex transformations. In the concave and convex conditions, the 3 leftmost transformations were M2-*a*, and the 3 rightmost transformations were M2-*b*. By pressing keys labelled "M2-*a*" or "M2-*b*" on the computer keyboard, subjects indicated their subcategorizations of M2 stimuli. As in the first learning phase, subjects had to go through as many learning blocks as needed to categorize the exemplars perfectly. Each block took about 5 minutes to complete.

Results

The mean number of learning blocks to reach criterion in the first and second learning phases are reported in Table 2. A two-way analysis of variance revealed a main effect of condition ($F(2, 15) = 11.46$, $p < .01$), a main effect of learning phase ($F(1, 15) = 9.84$, $p < .01$), and a significant interaction of condition and learning phase ($F(4, 15) = 25.18$, $p < .01$). All three effects reflect the finding that subjects in the mixed condition required very many blocks in the first learning phase, which decreased dramatically in the second learning phase while increasing in the other two conditions.

Table 2. Mean number of learning blocks necessary to reach criterion in Phase 1 and Phase 2 of Experiment 6.

	Convex	Concave	Mixed
Phase 1	3.67	3.17	13
Phase 2	5.83	4.17	1.83

Discussion

In Experiment 2, the difficulty of forming a part representation in the *mixed* condition was assessed by the accuracy of category judgments. The present experiment showed a similar effect with a different procedure. As expected, the average number of learning blocks required to learn the category was much higher in the mixed condition than in the concave or the convex conditions. It is worth re-emphasizing that the total size of transformations in the three conditions was identical, so the difficulty of learning the mixed categories must be due to the change in the sign of curvature of its transformations.

But the main goal of this experiment was to test whether a change of sign of curvature would give rise to two distinct units of representation for the category. If this were true, subjects in the mixed group would enter the second learning phase with a category already segmented into two subcategories: one with the convex part, and one with the concave part. Therefore, reaching criterion in this phase would be faster for the mixed group than for the other groups. This prediction was confirmed by the data. Although the mixed condition was by far the slowest in the first phase, in the second phase this group learned the category in 3.1 fewer blocks on average than the other groups, which was a reliable difference, $t(16) = 2.15$, $p < .05$.

To summarize, the results give evidence of one kind of perceptual constraint on part construction when the parts are not constant across exemplars. Although the Homogeneity Principle predicts that potential parts present on all objects of a category should be glued into one unit of representation, these results show that not all variations of a component result in a single potential part. People seem to expect that some aspects of the part will remain constant in order for it to be the "same" part encoded into memory; in particular, the extrema of curvature of 2-dimensional outlines of objects are expected to remain constant. If there is a change in the sign of curvature, subjects will instantiate two units, though with more difficulty than they would instantiate a single part. However, those two units then become available for learning other categories.

References

Biederman, I. 1987. Recognition-by-components: A theory of image understanding. *Psychological Review* 94:115-147.

Bourne, L. E., Jr. 1982. Typicality effects in logically defined categories. *Memory & Cognition* 10:3-9.

Bruner, J. S., Goodnow, J., & Austin, G. A. 1956. *A study of thinking*. New York: Wiley.

Clark, E. V. (983. Meaning and concepts. In J. H. Flavell & E. M. Markman (Eds.), *Manual of child psychology, Vol. 3: Cognitive development* (pp. 787-840). New York: Wiley.

Hoffman, D. D., & Richards, W. 1985. Parts in recognition. *Cognition* 18:65-96.

Lesgold, A., Glaser, R., Rubinson, H., Klopfer, D., Feltovich, P., & Wang, Y. 1988. Expertise in a complex skill: Diagnosing x-ray pictures. In M. T. H. Chi, R. Glaser, M. J. Farr (Eds.), *The nature of expertise*. Hillsdale, NJ: Erlbaum.

Mervis, C. B. 1987. Child-basic object categories and early lexical development. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual bases of categorization* (pp.201-233). Cambridge: Cambridge University Press.

Richards, W., Dawson, B., & Whittington, D. 1988. Encoding contour shape by curvature extrema. In W. Richards (Ed.), *Natural Computation..* Cambridge, MA: MIT Press.

Schyns, P. G., & Murphy, G. L. 1991. The ontogeny of units in object categories. *Proceeding of the XIII Meeting of the Cognitive Science Society*, 197-202. Hilldsale, NJ: Erlbaum.

Schyns, P. G., & Murphy, G. L. 1992. The representation of parts in object concepts. Manuscript submitted for publication.

Tversky, B., & Hemenway, K. 1984. Objects, parts, and categories. *Journal of Experimental Psychology: General,* 113:169-193.