open the possibility that multiple recodings (possibly cast at different levels of complexity) of the same input remain available for further processing, a property that is undoubtedly central to flexibility in any learning system.

ACKNOWLEDGMENT
The author is a Research Associate of the National Fund for Scientific Research (Belgium).

Flexible feature creation: Child's play?

Gedeon Deák
Department of Psychology and Human Development, Vanderbilt University, Nashville, TN 37203. **deakgo@ctrvax.vanderbilt.edu**
peabody.vanderbilt.edu/psych_and_hd/faculty/deakg

Abstract: Schyns, Goldstone & Thibaut's argument is evaluated from a developmental perspective. Theoretically, feature creation is not necessarily problematic; this view derives from the assumption of innate content (primitive feature sets). Alternative assumptions (e.g., Piaget's theory) are possible. Preschool children readily search for novel features in response to task demands. This is compatible with functionalist approaches, but not the rationalist ones criticized by the authors.

What is innate? Alternative views. Schyns, Goldstone & Thibaut are to be commended for attempting to address a critical problem: if the symbol-grounding problem is solved by primitive feature sets, how can we account for (sporadic) flexible and creative induction based on nonprimitive features? If primitives are used productively, what are the production rules? More radically, when and how do people create new features? The difficulty of the problem is underscored by Fodor's (1975) reductio that all complex concepts are innate. Although the Schyns's et al. approach is more sober, primitive constituents are regarded as a necessary evil (sect. 1.1, para. 4). Are primitive vocabularies necessary? Can the problem be solved without innate content (see Braine 1994)?

The symbol grounding problem is inherently developmental, and developmental analysis yields an alternative formulation: Piaget was aware of the problem, and did not believe in primitive content (i.e., features or concepts) but rather in primitive action patterns. Specifically, a set of innate reflexes evolve into controlled action schemes (later internalized as mental schemes). By analogy to Gibson's differentiation hypothesis (Gibson & Gibson 1955), the infant's ever-finer differentiation of the physical world is reflected in differentiated action responses. Primitives are not content, but structure and *process* – specifically, perceptual learning and motor learning routines. My point is not to advocate Piaget's theory, but to remind us that what is innate might not be primitive feature sets. Whether or not primitive features are assumed, however, feature creation remains to be specified.

The growth of flexibility. Schyns et al. link feature creation to flexibility, a welcome observation. Flexibility is most clearly construed from a functionalist position, by questioning how subjects elect specific features appropriate for different tasks. Subjects might select either previously conceived aspects or novel aspects of the stimulus array. If the latter is typically more effortful and uncertain than the former, there will be a trade-off: feature search will occur only when existing features are ill-suited to a problem. Schyns et al. consider how stimulus characteristics affect feature creation, but they do not address how task and context facilitate or inhibit feature creation. For example, Bransford et al. (1989) suggest that conceptual highlighting of contrasting features might promote feature creation. Deák and Bauer (1996) found that preschoolers search for subtle (and presumably novel) features in certain task contexts, given sufficiently complex stimulus items (this qualification is consistent with Schyns et al.'s argument about stimulus characteristics; see sect. 2.5).

In terms of development, Schyns et al. imply that apprehension of novel features poses particular difficulties for young children (sect. 3.3.2, para. 3–4). However, learning novel feature contrasts is so critical for young children, it would be surprising if they were deficient in it. Schyns et al. seem to know this, and conclude correctly that current popular approaches (i.e., innate theories; perceptual biases) are not enough to specify how children create or select relevant features from an array. Theoretical shortcomings aside (see Deák 1995; in preparation), both popular approaches are empirically wanting. Consider the view that children's feature selection is governed by innate perceptual biases: Deák (1995) qualifies or disconfirms the count noun/shape bias proposed by Landau et al. (1988). A comprehensive review of the literature reveals that preschool children are not *generally* biased to weigh some features over others. Rather, they select stimulus features associated with a particular induction problem (Deák 1995). That is, preschool children are genuinely flexible in shifting attention to features (or combination of features) as task demands change. Moreover, it appears that feature selection and feature creation are closely related, and well-established, by age 3. I will briefly describe data (which inform Schyns et al.'s general position) consistent with this argument.

Deák (1995) found that 3- to 6-year-olds successively attend to different combinations of features in response to different induction tasks: when told that an unfamiliar, complex object "has a toggle," 3- to 6-year-olds extend that fact to another object with the same unfamiliar part. When told that the former object is "made of mylar," children extend that fact to a different object made of the same (unfamiliar) material. How is this relevant to the position stated by Schyns et al., besides being inconsistent with the perceptual bias approach? First, in response to ". . . has a . . ." facts, children apparently search for an unfamiliar part of the object, then seek another object with the same part. Parts were novel, and the objects were complex, with many varying features. Thus, children apprehend and reason about novel features in spite of irrelevant varying information (contrary to Schyns et al., sect. 3.3.2). Second, responses to ". . . made of . . ." facts demonstrate that preschoolers reliably make generalizations about novel kinds of material. This is striking partly because 3-year-olds are believed to lack a coherent concept of material kind (Dickinson 1989). Thus, children "created" a new material feature, although by conventional accounts no feature space existed to be subdivided! This illustrates how task constraints might permit young children to induce novel features or feature combinations. Clues to the nature of the task – for example, the phrases "has a" and "made of" – effectively limit the hypothesis space preschoolers consider.

In sum, assuming that "what is innate" is either nascent theories or rigid perceptual biases (i.e., innate content *qua* primitive features) provides no account of either flexible feature selection or feature creation. A functionalist framework, in contrast, assumes that children have procedures for matching (or learning to match) known or novel features with exigencies of the task at hand. The emergence and nature of these procedures remain to be understood.

Flexible categorization requires the creation of relational features

Peter F. Dominey
Vision et Motricite, Inserm U94 and Institut des Sciences Cognitives CNRS UPR 9075, 69373 Lyon, France. **dominey@lyon151.inserm.fr**

Abstract: Flexible categorization clearly requires an adaptive component, but at what level of representation? We have investigated categorization in sequence learning that requires the extraction of abstract rules, but no modification of sensory primitives. This motivates the need to make explicit the distinction between sensory-level "atomic" features as opposed to concept-level "abstract" features, and the proposal that flexible categorization probably relies on learning at the abstract feature level.

It is clear that a fixed and finite feature set cannot anticipate all possible categorizations. Schyns et al. thus set out to establish a framework in which the feature set can be augmented with new features. Part of the burden of proof on the authors is to demonstrate a case in which new feature creation can clearly be dissociated from the weighting or combination of existing features (sect. 2.6). In this regard, we study a form of categorization of dynamic objects – sequences – in which the defining features for categorization are configurational relations between elements or features, independent of specific features themselves. I develop this point as an example of a case in which the creation of novel features cannot be a simple combination of existing features. I note also that in this and the subsequent points, the dynamic object results generalize to static objects.

We have recently described a dissociation between cognitive processes for learning surface structure and abstract structure of sensorimotor sequences. Surface structure is simply the serial order of the elements in the sequence, whereas abstract structure is defined in terms of positional relations between elements that repeat in a sequence. Thus, the two sequences ABCBAC and DEFEDF share the same abstract structure (123213) but have different surface structures, and are thus isomorphic. We have demonstrated in a serial reaction time task that although surface structure can be learned under implicit conditions, abstract structure can only be learned under explicit conditions (Dominey et al. 1995; 1997; in press). A hallmark of abstract structure learning is the capacity to transfer knowledge of the abstract structure to new, isomorphic sequences. Specifically, when subjects trained with sequences such as ABCBAC are exposed to a new isomorphic sequence DEFEDF, they can transfer knowledge of the shared abstract structure 123213. This transfer yields significant performance benefits for the elements in the new isomorphic sequence that are predictable by the abstract structure. Thus, by means of their training, these subjects have gained category knowledge that allows them to "categorize" sequences as either belonging to the isomorphic set or not. We thus have a condition in which previous experience significantly modifies the immediate appearance of dynamic objects or sequences.

The question remains: Has a new feature been created? A response can be provided from simulation studies of abstract structure learning. In these studies, a neural network model is capable of learning and transferring abstract structure between isomorphic sequences (Dominey 1995). In the model, the set of sensory-perceptual inputs or atomic features was fixed. Flexibility came from an adaptive capacity to represent arbitrary relations between atomic features. This type of learning allows it to be recognized that ABCBAC and DEFEDF share a common abstract structure of internal repetition.

In saying that features are both elementary stimulus properties yet little different from concepts, Schyns et al. may have blurred a useful distinction between levels of representation. I thus suggest a slight nuance: "Atomic" features correspond to low-level (sensory) primitives, and "relational" or "abstract" features are defined in terms of configurational relations between atomic features, as in the abstract structure learning already described. In this framework, relational feature creation does not involve the extraction of a new feature dimension explicitly represented in the stimulus at the atomic or perceptual level, but instead entails the extraction of a relation between such features, independent of the feature instances themselves. Such a capability is clearly a key element in a flexible categorization scheme.

Flexible features, connectionism, and computational learning theory

Georg Dorffner

Department of Medical Cybernetics and Artificial Intelligence, University of Vienna, A-1010 Vienna, Austria.
georg@ai.univie.ac.at www.ai.univie.ac.at/oefai/nn/georg.html

Abstract: This commentary is an elaboration on Schyns, Goldstone & Thibaut's proposal for flexible features in categorization in the light of three areas not explicitly discussed by the authors: connectionist models of categorization, computational learning theory, and constructivist theories of the mind. In general, the authors' proposal is strongly supported, paving the way for model extensions and for interesting novel cognitive research. Nor is the authors' proposal incompatible with theories positing some fixed set of features.

Schyns, Goldstone & Thibaut's proposal is remarkable and important in many ways. It is remarkable because it touches on issues that have come up in other domains of cognitive science but have received surprisingly little attention in the literature about categorization. It is thus worth investigating Schyns et al.'s view on features in the context of two areas they have not explicitly discussed: connectionist models of categorization and the theory of learnability.

Connectionist models of categorization can be divided into at least two strands: models based on backpropagation in MLPs (multilayer perceptrons) (Harnad 1987; Rumelhart et al. 1986), and models based on more localized responses in radial-basis function networks (Kruschke 1991) and variations of competitive learning (Dorffner et al. 1996; Grossberg & Stone 1986). MLPs are a variant of neural networks that apply a weighting scheme to features, whereas the other kinds of model compute a distance measure between a prototype (the weight vector) and an input pattern.

With respect to the former, it is interesting to note that hidden units in multilayer perceptrons can be interpreted in exactly the way as Schyns et al. suggest. To arrive at a complex categorization, hidden units develop higher-order combinations of input features that are influenced by the categorization task itself. Put differently, hidden units can be regarded as higher-level features on which subsequent categorization is based (compare Bishop 1995, pp. 226ff). These fulfill Schyns et al.'s requirements, especially the fact that they develop with the help of feedback from the categorization task.

Models using localized responses appear to be more limited in light of the proposal by Schyns et al. Existing models indeed work on a fixed set of input features and apply one level of weighting through the distance measure. Thus, Schyns et al. show the need for a substantial extension. One way to approach hierarchical categorization is to add one or several levels of categories between the input and the final category. In other words, categories cannot be the result of only weighting input features; more complex features themselves can be generated by a categorization process. We have recently argued along similar lines, even though our model is not fully implemented (Dorffner 1997). Our approach agrees with Schyns et al. that "there is little difference between concepts and features" (sect. 4, para. 1). Important in such an approach would be top-down feedback from higher-level categorization to feature-level categorization. This could be implemented in the framework of the adaptive resonance theory of Grossberg (1976) through the recuritment of new feature categories.

Computational learning theory (Valiant 1984) can shed more light on Schyns et al. theory. It is safe to assume that for categorization there must exist a level where features are fixed, the level of peripheral sensory features (e.g., retinal activations). According to computational learning theory, given n such fixed features, if a learner is able to represent all 2^{2^n} dichotomies (assuming that features are binary and the system is learning to distinguish only two categories) then learning is impossible. Thus, substantial bias is necessary to constrain learning, as Schyns et al. acknowledge in