Papert, S. (1977). Concepts and Artificial Intelligence. In Macnamara, J. (Ed.). (1977). *Language learning and thought*. Academic Press.

The Katherine Nelson paper Seymour Papert refers to in his chapter follows for context.

9

Concepts and Artificial Intelligence

SEYMOUR PAPERT

Massachusetts Institute of Technology

There are lots of things **I** would like to say about Katherine Nelson's paper, particularly about the fine detail of the observations, but I shall have to pass over them in favor of the central topic of the paper, concepts and their formation. Here it strikes me as curious and disappointing that Nelson does not advert to the whole enterprise that goes under the name of artificial intelligence, because it has been so involved with concepts. So perhaps I shall spend some time pointing out the links with artificial intelligence.

First, I think that Nelson has placed her finger on the important issue in understanding linguistics and certainly the most important in the acquisition of language. Paradoxically, the tremendous development in technical and formal devices for representing language has not been accompanied by a corresponding growth in the devices for representing meanings. So, there is a problem relating language to meaning. I would like to confine myself to the meaning end of that comparison and say something about what a concept is. In doing so, I find myself somewhat hesitant lest I should find myself committed to the "concept of concept." Nevertheless, **I** hope **I** can say something without having to deal with that issue.

What is a concept, then, and what can it be made of? Our attempts to answer this question will resemble discussions of what genes were 50 years ago. People could say a lot about genes, but they had very little idea of what they actually were. Today we know pretty well what genes are and we can define them in terms of DNA, which is probably very different from Mendel's idea of a gene. Our knowledge of concepts is like the earlier knowledge of genes. But I can enhance my knowledge by attempting to build a model of concepts and see what sort of formal entities we will require. In this context one of the questions that can be raised is: Do they have parts? Nelson tells us they do, and in particular that the parts form two major subdivisions. One of these defines the essence of the concept, the other corresponds to the features that we use to recognize an object that exemplifies the concept. Our experience is such that the distinction is impossible to maintain. One might be able to manage it with a concept like ball, but when one tries to carry it through to more complicated concepts, it does not work. Moreover, I think it gratuitous to identify the core concept with a set of functions. How in this account could a child form a concept of the moon?

I would like to illustrate a somewhat different approach with an example from an early program which was drawn up by Patrick Winston of M.I.T. One of the interesting aspects of his program was that it could learn to a limited extent, and in particular it could learn what an *arch* is. Because Winston was dealing with a computer, he could not be satisfied merely to list all the properties of an arch. As an

aside, I get the impression that the features in Nelson's concepts are joined together by some sort of association. That would never do for a computer program. Let us see what Winston's program had.

Winston's machine lives in a world of blocks which will be familiar to those who have read Winograd's thesis. The machine would like to learn what an arch is, and it is being taught by being told that this is an arch and this is not. I will not go into detail, but it somehow recognizes that there are three parts or blocks in an arch, and for each it creates a node in its own internal representation system. In itself this is a very important primitive operation; it is very different, for example, from anything one finds in traditional logic. What the machine does not do is to take some data and combine them to make a new object. Instead, it models the objects in the environment. In addition, the machine can say things about the objects that the nodes represent. It does this by attaching a pointer from the node to one of its concepts. In the example, it attached a pointer from each node to the concept block, and the pointer itself was labeled "is-a." This is a special relation which can be equated with class inclusion. In drawing these pointers, the machine was saying to itself, each of the objects is a block. But there is more to say, because it also recognizes relations among blocks and it sees that one of the blocks is supported by the Other two. So, it draws a line between the appropriate blocks and labels it Supports. Now it has a first approximation to a concept of an arch.

Winston then shows it a structure made from three blocks; in it, one block rests on the other two, but the other two are touching each other. He asks the machine, and it compares the design of this structure with its concept of arch and answers "yes." Then Winston says, "No, it's not an arch." The machine then thinks, and it thinks by comparing its design for the new structure with its design for the structure that was an arch. It notices that in the arch the support blocks do not touch, and so it enters another annotational link between the two supporting blocks which can be paraphrased by "must not touch." It can continue to modify the concept with annotational links to say certain things must be so, certain things may be so, certain things must not be so, and so on.

I think you will agree that all this is different from the inert list that Nelson is proposing as a concept. The machine is fitted with a set of procedures for building nodes which are then interrelated in the appropriate manner. It does not simply make a list of facts and leave it at that. Indeed, the concept of arch will be related to a number of other concepts in such a manner that the machine's entire conceptual system may come into play in constructing a new concept. This is an important point for which I do not see room in Nelson's account. For example, if a child sees an apple rolling, whether he calls it a ball or not is likely to depend on whether he already has a concept of an apple. So, concepts interact in the use we make of them. If Nelson were woodenly to provide for this in her system, she would probably enter with each concept, say of an apple, "is not a ball," "is not an orange," etc., etc. This would, of course, be a rather bizarre way of doing it.

What has work like Winston's done for us, and how is it an advance on what has gone before? To answer this, I would like to draw an analogy with the history of linguistics, though I fear I do not know that history well. However, the way I see it is that the older grammarians, with a good deal of success, knew how to parse a sentence, how

to divide it into its constituent noun phrases and verb phrases; and they knew about nouns and verbs and the other parts of speech. In more recent times, there has been a vast growth of additional machinery which increases the analytic and processing power of grammar. Today we have deep structures and surface structures and transformations and markers, and we can describe syntax in a manner that seems beyond the scope of traditional grammar.

What in "conceptology" corresponds to traditional grammar and what has been added? To my mind, what corresponds to the older grammar is logic, the predicate calculus, and such systems. The key addition to that is the idea of a general data structure which together with an "interpretor" program brings into play an enormous power of computation. The point is subtle, because there is a formal sense in which the information represented in Nelson's concept and Winston's is equivalent. They are equivalent from the point of view of the essential complexity of computation. This is a metric that compares difficulty of computation independent of the mechanisms with which one has to compute with. So functions computed on the basis of Nelson's and Winston's representations of concept may be equivalent from the standpoint of essential computability. This does not mean that they are equivalent for psychology. It seems to me that programs, working with other data structures represent information in a form which is much more flexible, and better organized for psychological purposes than Nelson does. I might add in passing that we in AI do not attempt nearly so clear a distinction between a recognitional part and a functional part of a concept.

A computational representation is not a definition as Nelson's is but an instrument to be employed in manipulating blocks. It carries none of the overtones of logical definitions. It does not replace the external object, it is not an abbreviation of the information in that object, it does not seek to eliminate inessential information. Quite literally, it is to be employed in manipulating objects in all their complexity in their complex environment I think this a qualitative difference which advances our understanding of concepts. It is almost impossible to represent concepts, their development, and functioning by algebraic or logical formulas. One needs a more dynamic representation along the lines suggested by Artificial Intelligence.

8

The Conceptual Basis for Naming

KATHERINE NELSON

Yale University

How are the concepts of the child related to the words of the language? This is the central question to be addressed in this paper, and it involves defining the relation of the conceptualizing process to the learning of language as a continuously interactive and organized system. It will be seen that many mistaken conclusions about the cognitive and linguistic sources of meaning and of errors of meaning in the child's use of language derive from a misspecification of this problem. Although it is understood that the linguistic side of the problem (that is, the semantic or lexical systems) can be described in terms of rules, the present problem is *not* one of describing such rules nor of describing the course taken by the child in learning such rules. Rather it is an exploration of *the development of the child's conceptual system* which is assumed to underly the linguistic system and to determine in large part the limits on understanding that system at any given point.

To approach this problem therefore we need to go beyond the child's words to the conceptual layer underneath. In a previous paper (Nelson, 1974a), I suggested that we could describe the child's concepts in terms of

Preparation of this paper was generously supported by a grant from the Carnegie Corporation of New York.

three parts: a functional core which expresses all of the essential, specific definitional relations of the thing or event in time and space; optional or possible—but not necessary—general, spatiotemporal relations of the thing as the child has previously experienced them; and a set of identificational features which describe how one recognizes that a newly encountered thing or event falls under a previously formed concept. The functional core of the concept in these terms organizes what is centrally important or useful to the individual about that thing based on his/her experience with it. It is for this reason that it is a functional core; that is, what relates the object to the individual is not its perceptual features but its place in his/her life experience. It is for this reason also that the functional core varies from individual to individual, and may particularly vary from adult to child and from child to child. For example, the concept of clock may be quite different for the young child- involving principally noises, people, and locations—and the adult, for whom the essential core function of clock is time-keeping. Yet both adult and child may identify the same instances of clock as clock. That is to say, both adult and child may have the same set of identification features for the concept of clock, although the functional core of the concept is not the same.

In suggesting this concept structure, I have tried to suggest that the *process* of forming concepts normally begins with a functional core to which identification features are added as soon as it becomes necessary to identify new members, and that this process is as true of adults as of young children. The primary developmental difference to be

expected is that children must habitually generate and test more concepts than adults do because they have fewer available from prior experience. I will return to this point later. Here the point is simply that, although there is much to be discovered about the *concept-generating process* in. infancy and early childhood—as well as in adulthood—evidence shows that the infant does in fact operate on the basis of concepts before s/he begins to apply names to them.

I want to develop two points based on this evidence that will provide the structure for the discussion to follow. First, we only get into trouble when we try to explain development in terms of the child learning the correct (standard, adult) way to do things: to say things, to conceive of things, or to categorize things. This kind of view unduly rigidifies adult modes at the same time that it denigrates more open and less developed systems. But more than that, it actually prevents us from seeing the child's system. If we take care to look at the child's system, however, we can see that the child does not make mistakes. The child's system may not match the adult system in various ways, but it is our job to determine what the nature of the system is, and how it develops, not to identify its "errors," which are errors only from a different systematic point of view. The first part of this paper is devoted to an exploration of this proposition in terms of the kinds of concepts the child operates with and their relations to the language s/he uses.

The second and more central point I want to consider is that the young child's efforts to make sense of the world through forming concepts and learning words are a first manifestation of a continuing process of stabilizing an inherently unstable experience in order to operate on it and make predictions about it.[1] The child—like the adult—cannot help trying to make sense of his experience; and learning and using language is but one manifestation of this process. So also are our ways of characterizing what the child is doing, as well as our ways of analyzing language and thought into components. However, in this endeavor, because the tools of a closed and stable system are used to deal with an unstable and open system, we often find that the tools are inappropriate to the task. The second part of this paper will explore the implications of this view.

CONCEPTUAL TYPES: FORMATION, GENERALIZATION, AND REPRESENTATION

The basic *developmental* task is to elaborate the cognitive structure and to learn how to match it to the encountered linguistic structure. In other words the child's task is to *develop conceptions and acquire semantics to match.* Solving this problem involves the utilization of some quite general developmental principles, the most important of which involve stability and predictability, which in turn are related to the central role of context in both conceptualizing and using the language.

Consider first some typical concepts referring to objects named by young children who are first learning to talk; specifically, concepts of particulars, reduplicates, similars, and categories of similars. Most single terms referring to objects learned by young children can be assigned to one of these types. The typology describes a continuum from the individual case to identical things to unlike things which enter into

[1] This point has been made on a general level by both Cassirer (see Bolton, 1972) and Arnheim (1969). It is developed here for the important light it sheds on the function of concepts in the life of the child.

hierarchical relations with higher- order concepts. Thus, each poses a somewhat different problem for the concept-generating child as well as for the theorist. In the course of this discussion, we will focus on two provocative problems in regard to conceptualizing: the basis for generalization, and the function of prototypes in concept structure. It will also become clear that many of the problems previously identified as conceptual errors of the child are in fact linguistic confusions or problems of the cognitive-linguistic matching process.

Particulars

The child's concepts of particularly valued individuals are obviously related to the learning of proper names for them. It is interesting that Katz, Baker, and Macnamara (1974) have recently demonstrated that even at 17 months many children distinguish linguistically between particular things (dolls) that take proper names and classes of things (also dolls) that take general names, when the only cue to this distinction is in the use or nonuse of the indefinite article. Fodor (1971) pointed out that although both proper names and general terms are rule-bound, proper names are not extendible while general names are. By definition, therefore, proper names cannot be only a special case of general names and the referent relation between word and object cannot be explained on this basis.

It is important to recognize this point, but it does not bear on the notion that a general *concept* can be extended from experience with a particular object or event. Children do refer to particular objects by special names; for example, *Fido* or *Mrs. Brown*, or even *blankie* for the child's favorite blanket. Proper names are invented by adults (and are used by young children) to designate especially valued particular exemplars of more general concepts[2]— objects, people, animals, and so on for which recognizing individuality is especially important. While proper names are understood to refer to only one object and to be unextendible, a valued particular is at the same time a member of a larger set that takes a general name. The general name is extendible to all the exemplars of the concept, but the particular name is not. This linguistic convention is apparently easily understood by most young children who appear readily to comprehend that some objects (people and pets, for example) are so important in their individual selves that they take special identifying names.

According to previous accounts, however, it would seem that not all children are as discriminating in this respect as Katz *et al.'s* (1974) subjects. For example, the common extension of the proper name *Daddy* to other men is frequently referred to in the literature. In this connection, Piaget (1962, p. 255) cites Lucienne at 3 years, 2 months, as follows (L's part of the dialog in italics): "We passed a man: *Is that man a daddy*?—What is a daddy?—*It's a man. He has lots of Luciennes and lots of Jacquelines*.—What are Luciennes?— *The/re little girls and Jacquelines are big girls*." Piaget cites this passage as evidence for his theory that young children are incapable of achieving either true generality or true individuality in their concepts. In the same connection, he cites the following interchange with Jacqueline, also at 3 years, 2 months, when she "could not understand that Lausanne was all the houses together because for her it was her grandmother's house 'Le Cret' that was 'the Lausanne house!' " After giving her this

[2] The terms *Fido* or *Mrs. Brown* actually are also specified by an internal representation as is evident in the fact that they can be named when either no longer exists in actuality (cf. Searle, 1970). This is a somewhat different case from those that we want to call concepts, however, simply because it does lack generality.

explanation, Piaget asks her: "What is Lausanne?—*It's all these houses* (pointing to all the houses round). *All these houses are Le Cret.*—What's Le Cret?—*It's granny's house, it's Lausanne*" (pp. 225-226). Piaget argues that this illustrates that for the young child a single object becomes the representative or prototype for the whole concept, whereas in a true concept all members are equivalent through their common abstract characteristics. Let us set aside the conceptual problems posed by Piaget's faulty definition of *Lausanne* as "all these houses" or the linguistic problems posed by detaching the modifier *Lausanne* from house and giving it an independent status, and examine these examples for what they can tell us about the relation of particulars to general concepts.

It is obvious in the first example that Lucienne has a concept of daddy, of Jacqueline, and of Lucienne (herself) that can be extended—at least in part—to other examples. Actually, it would be more accurate to say that she has a concept of "big girl" that includes as one of its names (probably the most important one) the name *Jacqueline*, while *Jacqueline* has its own representation, an individual one, which is in turn one instance of "big girl." Probably the concept of "big girl" was formed around episodes involving Jacqueline the person (as the concepts of man and little girls involved Lucienne and Daddy). Once the concept was formed on the basis of experience with a particular, it could be extended easily to include other instances that met its defining rules. Thus, the example tells us that a general concept may be centered on (and perhaps derived from) a single exemplar. Lucienne's problem, however, was not confusing Jacqueline as an individual with the general concept (note that she was able to state the conceptual relation adequately), but only substituting the proper name for the general term. It was a linguistic error, not a conceptual one, and one that Katz *et al.'s* much younger children did not apparently make, or at least not frequently, and not in comprehension.

What about the Lausanne example? Behind the truly massive linguistic confusion here (J. has two names for her grandmother's house, one of which uses the town name as a modifier, but she has no independent name for the town itself), there seems to be a question as to whether J. has a concept of town at all. If she believes that Piaget is truly referring to "all these houses" as Lausanne, then she is not incorrect in believing that he is defining some new subclass of houses of which Le Cret is a prime example or prototype, one that can be referred to as either *Lausanne* or *Le Cret*, or even *granny's house.* In this case, we have another example of the formation of a concept on the basis of a particular example, and its extension to others of its kind, that is, to those houses located in Lausanne. This seems a more plausible explanation than Piaget's, which claims that "all these houses … thus constituted a complex object depending upon one of its elements which was seen as representing the whole." The example is unfortunately too confused and complex to disentangle with certainty, however.

What I am suggesting here is that particulars have no special status in the child's conceptual system equivalent to their special status in the linguistic system. On the contrary, a concept can be formed and probably usually is (Rescorla, 1976) on the basis of one exemplar; but the child does not assume thereby that it is unique. Rather, his usual practice seems to be to form a potential category that can include new but like instances. In fact, on the conceptual level, there is some question whether the infant either knows or cares whether he is interacting with the same individual thing from

one time to another during the period when he first begins to name things. I have several times had the experience of asking a 15-month-old child to find a ball that we had been playing with, which had rolled out of sight temporarily, only to have him go to his own toy box and get a completely different ball. This indicated that the particular individual object had no special status with respect either to its name *[ball)* or its conceptual definition (playing, rolling, etc.).

This is not to say that some particulars (such as mother and daddy) are not important to the child; in fact, their clear existence for the child demonstrates the establishment of concepts of permanent objects much earlier than claimed by Piaget. Nevertheless, they are not different in kind from members of general categories of objects—either reduplicates or similars as described below. One type of concept usually derives from the other and both can serve as the impetus for concept formation and for naming. The child does not need to know therefore that a particular object is going to reappear tomorrow or that it will continue to exist, because the child does not name particular objects but exemplars of concepts. On the other hand, to give a proper, individual name to something does imply belief in both its individual importance and its continued existence. Thus the relation of naming to the establishment of the permanent object concept in Piaget's terms becomes clear: Individual names do require this concept; general object terms do not.

We must not confuse the conceptual issue with the linguistic one, however; proper names are not generalizable, while general terms are. But the existence of a proper name does *not* mean that its referent is the only instance of a concept, only that as an individual it is important enough to have a name oi its own.

REDUPLICATES

The suggestion that the infant may not have a firm sense of the individual identity of many things leads to a consideration of the place of what I call "reduplicates" in the infant's first concepts. Some collections of things, or serially experienced instances of categories are less individually identifiable than others. A salient example is sheets of typing paper, but of course snowflakes, acorns, paper clips, yellow lead pencils, and spoons, all fall into this category. When we consider the natural ecology of the child's world, it is striking how many of the things he must contact are of this kind; diapers and bottles, for prime examples, and of course many foods, including apples, bananas, toast, endless jars of applesauce and chopped liver. One is as good (or as bad) as the next. Another one of Piaget's favorite examples of the confusion of individuality and generality is relevant to this point. In this example, Jacqueline at 2 years, 6 months, "used the term *'the slug'* for the slugs we went to see every morning along a certain road. At 2:7 (2) she cried: *'There it is!'* on seeing one, and when we saw another ten yards further on she said: *'There's the slug again'* I answered: 'But isn't it another one?' J. then went back to see the first one. 'Is it the same one?—*Yes*—Another slug?—*Yes*—Another or the same?—. . .' The question obviously had no meaning for J." (Piaget, 1962, p. 225)

Here again we see some real linguistic confusions, including the use of the definite *it* and *the* for a newly encountered example of the class. And Piaget obviously was confusing several issues when he insisted on knowing whether it was the same or another one, issues about the child's understanding of the terms *same* or *different* which are only now being disentangled by experimental investigators (e.g., Webb,

Oliveri, & O'Keefe, 1974). But the individuality of slugs, like pieces of paper, acorns, and diapers, is not of interest. One slug is as good as the next. It seems unlikely that Jacqueline is genuinely confused as to whether there is one or more than one slug, or whether one keeps reappearing in different places; rather it is a question of indifference to her since she is not concerned about the identity of any particular slug. She has something to learn about appropriate linguistic usage, but it does not seem proper to view this as a case of conceptual insufficiency—but rather as an example of one special kind of concept, one that does not recognize differences among its individual instances. In these cases—and they are common in the child's concepts, as in his world, no doubt—extension of the name to new exemplars may take place on the basis of perceptual generalization or confusability. The frequency of such cases may predispose the child to expect that things come in many copies, that events will be repeated, and to extend his concepts—and names—to new instances without special training or tuition.

The discussion thus far has shown the basis for the establishment of prototypes in the child's concepts of valued particulars, and for generalization in the child's concepts of confusable reduplicates.

Similars

At the next level, that of similars, we come up against those classically messy cases of dogs, dolls, cars, and cookies named by the child early in the game (Nelson, 1973). In these cases the correct meaning for the word takes in a whole category of objects that do have identifiably different characteristics. In some cases the physical resemblances between members of a class seem less obvious than their differences. For example, the Chihuahua and the Doberman pinscher seem unlikely to fall into the same category from the looks of them. Or take the VW microbus—is it a car? a bus? a truck? (Actually of course it is all three, but then that just forecasts the conclusion that our natural language categories are not hard and fast logical classes but classes with fuzzy boundaries; cf. Labov, 1972.) A great deal has been written about the way children first form and extend word categories of this kind, but we still lack good systematic studies.[3]

This brings us strongly up against the question that has puzzled so many people: How does the child form a concept of such diverse classes and how does he extend it to new members? It is at this point that we need to invoke a *grouping principle*, and the principle that I have suggested (one that is consonant with Piagetian theory of sensorimotor development) is that of function: what things do or what can be done with them. This is a very general use of the term function (culturally defined function is a special case of it), but it is needed to cover both action and all other spontaneous reactive changes in state, which are so salient to the young child, as well as the known relations of things to self and others. Use of this grouping principle can be expressed in the following aphorism: "Categories are formed on the basis of function and generalized on the basis of form." Obviously, generalization principles are needed to identify new members of the category; but they are different from the grouping principles themselves. They are the source of many of the "errors" in the young child's

[3] I am indebted in the discussion that follows to Leslie Rescorla for sharing her data and some of her conclusions from her dissertation research, Rescorla, 1976 which is an attempt to fill this gap.

early naming. When the process is viewed in this way, it is obvious that some naming errors may be errors from the child's point of view as well as the adult's (that is, what he called a ball turns out not to be one), while others may reflect a true difference in concept definition.

In concentrating on so-called errors of generalization, however, some important points have been underemphasized. First, children do very readily extend words learned on the basis of a particular object to new instances of identifiably different but functionally similar objects as well as to perceptually similar objects. This is a strong indication that the child is not operating on the hypothesis that words apply basically to particulars or to things that are indistinguishable from one another, but on a more general principle that words apply to categories of things that, though discriminably different, have something centrally important in common. What I am suggesting here is that children do not make mistakes of generalization but mistakes, when they are made, of category recognition. What these "mistakes" can tell us, then, is not usually what the nature of the child's underlying concepts is, but what his identification rules are, because they occur when the child says, in effect, "Oh, there's another X." In looking at instances of conceptual inclusion, therefore, we need to distinguish carefully between those based on recognition and those based on classification. The child may think he has recognized something but change his mind when he finds that it does not behave as that class of things should (Nelson, 1973, 1974a). When the child's concept is formed on the basis of functional relationships then an error of recognition may occur even when his functional concept is equivalent to the adult's for that word; and many "errors" of word use may occur before the child has had a chance to build up an adequate set of identifiers.

A related point that tends to be overlooked is that young children are often surprisingly adept at differentiating closely related categories, when they are interested in doing so. A number of children at 2 years have a differentiated category of vehicles, for example, that rivals that of people many years their senior, being able to distinguish steamshovels from earth-movers and derricks, for example, whether as real, toy, or pictorial instances. The rudiments of this facility can be seen often at the very beginning of naming. It is a common pattern for a child to learn several closely related words in a general category within a short period of time, for example, *car*, *truck*, and *bus]* or *cake*, *cookie*, and *candy*; or *dog*, *cat*, and *puppy* (all examples taken from my own data; Nelson, 1973). Furthermore, Rescorla (1976) has pointed out that in many cases the child learns to comprehend a large number of differentiated terms while he is still using a single term in production. For example, one child in her study used *car* for a wide range of vehicles, including motorcycle, bike, truck, plane, and helicopter. She was, however, able to pick out each of these correctly in response to their appropriate names. Relying on her productive use of the words, one would conclude that she lacked conceptual differentiation, while testing her comprehension reveals that such differentiation is available.

There are other cases where as adults we use a word analogically or extend it to types not previously included in the category, or to instances of which we are unsure. Often, language as an open creative system not only allows but encourages this. Undoubtedly, the young child learning language, who knows few terms, does more of this. For example, one child in Rescorla's sample extended his term for clock to

include a medallion on the dishwasher, while another used the same term for dripping water which sounded like "tick-tock." We can imagine a fuller specification of what the child was saying as "It looks like clock" or "It sounds like clock." These analogical extensions were an addition to the more common categorical extension of the word to include watches, meters, and timers. This kind of intentional generalization often seems to follow a period when the word is used more narrowly and categorically (cf. Bloom, 1973). This is an early linguistic example of a quite general developmental sequence involving flexibility of application following a period of rigid rule-boundedness (cf. Nelson & Nelson, in press).

In addition to analogical and categorical generalization, there are other generalization "errors," for example) those based on affective principles (Lewis, 1951), where the child seems to extend a term on the basis of his emotional reactions rather than on any objective quality of the object or experience. These have been documented by many observers and some children seem to be more disposed to them than others. Such reports are often difficult to interpret with assurance, but that is not important. After all, adults do extend words metaphorically and on the basis of affective and aesthetic considerations; these uses cannot simply be expunged from the language and they are surely important to the young child. Many of the idiosyncratic expressions of children seem to be related particularly to emotional states. This area is wide open to further investigation, despite the difficulties of interpretation.

Another kind of "error" more likely to lie in the interpretation than in the child is evidenced in cases where the child did not mean to name an object at all, but for some reason used a word in a situation that made its use look like an instance of inappropriate generalization. When a child, just beginning to talk, speaks only in single words and often practices speaking without the intention to communicate, such instances are probably not infrequent. Rescorla found many of these cases which she termed predicate overextensions in her study; for example, a child who pointed to Daddy's shoe in the closet and said *Daddy*. While the possessive relation seems clear in this case, many other cases are ambiguous and may be termed overgeneralizations inappropriately. I know of no way to guard against such cases, but we should be wary of calling them generalization errors.

Thus we can conclude that while children do make what adults view as mistakes of overgeneralization, these mistakes are probably based on valid underlying general principles of conceptualization, and it is the conceptualization skills themselves that are the most important thing to recognize, if we are to understand the general process of naming.

Superordinates

Having considered some of the problems of generalization that emerge from the process of learning names for first object concepts, let us turn to the complexities of combining such concepts into natural language categories, such as food, animals, vehicles, clothes, or furniture. The process involved is to be distinguished from that of extending a single concept term to new exemplars of the type just discussed. It is rather the process of combining two or more concepts into a superordinate category *without sacrificing the identity of the original concepts*. I have already noted that many of the early vocabularies of young children exhibit several representatives from one of these adult categories. Does that indicate that there is any categorical cohesiveness of

hierarchical structure as far as the child is concerned? The early productive use of a single term such as *car* for several classes that are differentiated in comprehension, as discussed above, appears to indicate an early awareness of what we would call a vehicle category. However, few children learn or use superordinate terms at the outset. While Roger Brown (1958) pointed out long ago that parents tend not to use superordinates in talking to children, there are exceptions as in the use of *toys*; and one child in Rescorla's study learned *fruit* as a general term for nectarines, applesauce, and pears.

There is also a good bit of fragmented evidence that even quite young children know and utilize some of the relationships among words, things, and categories. For example, when called upon to remember two objects, 2- and 3-year-olds will remember them better if they are from the same category (animals, utensils, or vehicles) than if they are from different categories (Goldberg, Perlmutter, & Myers, 1974). Similarly, Faulkender, Wright, and Waldron (1974) found that, for 2-year-olds, habituation was faster for category members than for novel items. Moreover, Nancy Katz (preliminary report) reports on the basis of a series of experiments that children of 30 months understand some superordinate terms (e.g., *toy* and *animal)* and are able to categorize objects correctly into these classes, but (consistent with Piaget's observations) they appear to have difficulty coordinating subordinate and superordinate terms.

Unfortunately, most of our present knowledge of hierarchical categorization and its utilization comes from much older children. Both Piaget (Inhelder & Piaget, 1964) and Vygotsky (1962), for example, have emphasized the classification errors of the preoperational or preschool child, and there is, as a result, relatively little work with preschoolers that is concerned with finding what their underlying principles of categorization might be if they are shown to exist. Like the other problems we have considered, however, it seems probable that here the noted deficiencies are more linguistic than conceptual and that even very young children do form rudimentary superordinate categories but do not ordinarily apply special names to them. Certainly by the age of 5, children are able to enumerate many appropriate instances of natural language categories (Nelson, 1974b), and attribute class properties to new members (Harris, 1975).

CONCEPTS AND THEIR REPRESENTATION

One of the things revealed by this consideration of the child's concepts of objects and their associated names is that, although s/he is not particularly sensitive to the distinction between one and many, s/he attempts to impose a stable representation on experience, in effect creating "oneness" out of "manyness." Applying names to things is only one manifestation of this disposition to stabilize and make unitary the variability and continuity of experience. As we have seen, however, the infant's attempt to deal with both individual concepts and linguistic terms at the same time often leads to cognitive-linguistic confusions. In turn, our analytical attempts to describe the two systems may also lead to related confusions, because the ever-continuing effort to stabilize experience that is reflected in the infant's disposition to establish "oneness," is also reflected in the theorist's disposition to establish analytical entities. Let us consider some manifestations of this more sophisticated effort to deal

with conceptual variability by applying stabilizing representations to it. As a prime example, consider the persistent notion that the concept can be identified with the object or event from which it derives, that is, the still vivid "Helen Keller insight" that "everthing has a name. Few people now hold this view in its simple form, but it has an extraordinary attraction that cannot be brushed aside. It is the prime example of the reification of words.

Currently more attractive and widespread is the notion that a concept is a set of features or attributes. The notion of semantic features, like the notion of conceptual attributes, has some useful validity. By breaking meaning down into dimensions and components we can reveal something about relations *between* concepts or word meanings. However, we are misled if we believe that the components that we use to stabilize our understanding of certain distinctions actually describe the concept that underlies the word. The child's conceptual problem is to put together the important parts of the ongoing experience into a single representation or concept and only then to find the identifying features that will serve to differentiate among instances of related concepts. It is the analyst who imposes the stabilizing notion of components onto the child's stabilizing concepts; both moves are efforts in the same functional direction.

The notion that a concept may be an image reflects a different approach, but one that again confuses representation with what is represented. The conviction is widespread and growing that children depend more on images than on words in storing their concepts. However, the relation of images to concepts is far from certain, and again, the most probable function of the image is to serve as a stable representation in connection with a concept, functioning at the level of the word, but not equivalent to the concept itself.

The fact is that each of these proposals regarding what a concept *is*—objects, images, sets of features, words—imposes the notion of a static and discrete representation upon a variable continuum of experience with vague boundaries. It is worth considering the proposition that concepts lie somewhere in the middle: less continuous than experience, less discrete than objects. A recent proposal more consonant with this view is that concepts are derived from episodes. It is obvious that much of what is remembered about things is specific to how they were experienced in the past. For example, Schank (1974) has presented a number of examples of conceptual meaning based on the child's first and most salient experiences with the event, as revealed by the spontaneous word associations of a child of about 1 ½ years. One of his examples, including his gloss of her meaning, goes as follows (Schank, 1974, p. 11):

> Situation ….: Question: Hana, what do you do if your hands are cold?
>
> Response: Glove
>
> boots (When she wears gloves she usually wears boots)
>
> snow (the first time she wore her boots, she went to the snow)
>
> crunch (the snow went crunch under her feet)
>
> whee (she went on a sled and I said "whee")

Schank concludes from his material that the "episodes made up of sequences of

actions seem to be the crucial organizing factor in Hana's memory" (p. 12). Further, children's concepts are strongly identified with their first occurrence in memory and such memory is "grouped contextually."

This and other examples clearly indicate that the child remembers previously related episodes and has stored the information extracted from these with the word so that it is available when the word is used again or when she experiences glove or boot again. It must be noted, however, that because she talks about these things does not tell us that this particular episodic stuff *is* the concept. She does not, for example, talk about what boots look like, although she is able to recognize them when she sees them. Personal episodic experience is important to the formation of concepts, but without some further conceptual organization, it will remain simply a disconnected particular episode without general meaning. It is, however, important to note that episodic informadon is extracted, and while such information does not constitute the whole concept, it is included in it.

CONCEPTUALIZATION AS A PROCESS

Thus far it has been suggested that the child succeeds in selecting and stabilizing experiences of different types, using several different mechanisms, and also that we impose another layer of selection and stabilization in describing the resulting concepts. Let us abandon—at least temporarily—the quest for a specification of the form of concepts and consider them as an outcome of the conceptualizing *process.* It can be observed then that conceptualization itself has two outcomes. First, as already stressed, it stabilizes transiently experienced events and, as a result, thought about objects and events becomes possible. What does learning linguistic terms for one's conceptualization add to this? It increases conceptual stability (because of the relatively stable cross-situational use of words) and it also makes communications about one's concepts possible. It increases the *predictability* of experience associated with these events only indirectly, however; for example, the child can react to the word *dog* by activating his expectations about dogs just as he might if he saw a dog. (The child has to learn, of course, which expectations are relevant to words rather than actual events. For example, at the outset, the child reacts to *daddy* or *dog* in a way that indicates he expects daddy or dog to appear. He has to learn that the word does not always implicate an immediate instantiation of the concept.)

Conceptualizing is work and takes time, and the greater the number of situations that can be predicted on the basis of past conceptualizations and their expectations, the less current conceptualizing becomes necessary. Both stability and predictability are assets in dealing with the world. Stable and established expectations are in fact general characteristics of cognition and increase with cognitive development. The child gradually builds up an experiential base of knowledge that increasingly frees him from the necessity of constantly processing new information about novel encounters and from the necessity of forming new concepts. He can increasingly predict what is likely to happen on the basis of his present concepts and he becomes thereby freer to act and to think in a given situation rather than to engage in costly information processing.

For the adult, events are highly predictable in most situations. We look for the deviant event, not to form a new concept or category, but as a cognitive challenge in fitting it into our present conceptual structures. In contrast, the young child, faced with a world

where nearly everything is unfamiliar, must concentrate on those aspects that seem most important, and must try to order them in some way so that the next time they are encountered they will be familiar and therefore predictable.

But there is another side to this process: Concepts must remain open and changeable as new experiences with novel instances are encountered. It is not enough to categorize an animal as a dog, because each new dog is different. Each new experience that is not fully predicted must be able to modify one's concept to some extent: It adds a variation to the set of possible identifying features, or it adds a new relation to the set of probable relationships. Thus, to stability must be added flexibility and modifiability. Not only that, but we must be able to differentiate and recombine old concepts in order to form new concepts on the basis of old ones, as well as on the basis of new experience. Conceptualizing is a process that never stops, and although it may result in structures that are stable enough to be manipulated, the process itself is a fluid one.

Development of Word-Concept Relations

Previously, we saw that, although they did not always match, the child's concepts were similar to the meaning of linguistic terms. Now, however, it is clear that the conceptual characteristics just sketched—that is, flexibility and modifiability—are unlike those of linguistic concepts, that is, word meanings. Although words can be combined to form new conceptualizations and new terms invented, word meanings themselves change only slowly over generations compared to the rapid and continually shifting changes that are possible for the individual's concepts, and indeed are necessary throughout childhood. How then can the two systems of concepts and words be coordinated?

In order to deal with these questions, it is helpful to introduce a distinction made by Katz (1972) between narrow and broad semantic concepts. Broad concepts, according to Katz, include not only definitive information but incidental information, while narrow concepts include only the definitional. He gives as an example the concept of Martian, which as a narrow concept means a sentient creature inhabiting the planet Mars. As a broad concept, however, it may include information such as green with long antennae, travels in flying saucers.

Now it is obvious that in talking about how to characterize *cognitive* concepts we have throughout been considering broad concepts, that is, concepts that have organized *all* of the information that one has experienced and filed away in regard to people, things, events, or properties. Some of this information is more crucial in the sense of predicting new cases than others. We noted that some perceptual features, for example, help to identify new instances of the concept (if Martians really *were* green with long antennae, these would rightly be included as identifying features; instead they are simply incidental information). But at the core of any individual concept is an essential functional meaning which can be used to test concept membership; for example, "If it's a ball it will roll and bounce." If the test is violated, we know that, however much the thing *looks* like our prototype of the concept, it is not one of them.[4]

[4] This is not to deny the interdependence of form and function characteristics (Labov, 1972). Rather, it is to assert that function takes precedence both ontogenetically (we remember what things look like *because* of what they do) and definitionally (the crucial test tends to be the functional one).

Like the child naming dogs, however, we can also make mistakes in recognition. If we think we have an anteater, but it eats grass instead of ants, we may conclude that we were mistaken and inquire as to its true identity. We will not conclude that we actually have a different kind of anteater, because the conceptual test of an anteater is that it eats ants. We may of course imagine that since it looks so much like an anteater it must be a close relative of the anteater, but we do not insist that its outward appearance is defining. That is, we rarely rely on the maxim "clothes make the man" but rather on "handsome is as handsome does."

When a word is attached to a concept, the core meaning is presumably what ultimately defines the word for the individual, and it is therefore the individual equivalent of the narrow linguistic concept, that is, the stable linguistic or socially agreed upon definition. However, the core meaning of the concept to which the word is attached for any given individual may or may not match in any or all respects the narrow linguistic concept or word definition. Moreover, as we have noted, the individual's concept—in contrast to the word meaning or narrow linguistic concept—is less stable in the sense that it may change over time and in different contexts. An individual's concept can be reorganized at the core, can be differentiated, merged into a higher order concept, and so on. These changes may take place both in the direction of the linguistically accepted definition and away from it.

To recapitulate, we have two levels of conceptual organization, one social and linguistic, the other personal and cognitive. There are, in addition, two layers of concepts at each level—one stable and defining, the other flexible and inclusive. To learn the meaning of a word, the child must eventually match his own core concept meaning to the narrow linguistic concept. If these do not match, the word may be used inappropriately. Obviously, there are many opportunities for partial matching as well as overspecification, and these misspecifications are likely to derive from what the child has considered important and defining in his prior experiences with concept instances.

However, there is some evidence that for the young child the word implies the whole concept, not only its central core. This means that the child attaches a lot of excess baggage to the meaning of words, even when his core concept and his identificational features match those of the adult's. Saltz (1971) provides an example from his work with considerably older children—during the early school years—who appear to have a broad concept of daddy attached to the term *daddy,* such that daddies must be good. A daddy cannot therefore be a robber, and children will vociferously deny the possibility that daddies can be robbers. Very early, the child must be able to detach the core concept from its prior relations in order to consider it in new relations and thus to understand and produce novel sentences, for example (see Nelson, 1974a). However, it appears that for the young child these relations must not contradict previously established ones. Thus, the child may be able to conceptualize and talk about daddy building a house or climbing a mountain or flying to Mars, although he has not previously experienced any of these, but he will not admit to daddy robbing a bank because that is in conflict with the prior established relation that daddy is good. That is, although the core concept may enter into new relations, it cannot enter into relations that conflict with any prior information already included in the inclusive concept, which is, roughly, everything I know about X.

For quite a long time the child may attach the word to all of the expectancies in the inclusive concept (that is the concept that organizes all of the information the child has about that object or event). This is apparent in the kinds of definitions young children give as well as in their word associations. For example, the following definitions (from Krauss, 1972) appear not to refer to what the essential meaning of the term is for the child, but to something about its associated expectations.

> "Dogs are to kiss people."
>
> "A party is to say how-do-you-do and shake hands."
>
> "Buttons are to keep people warm."
>
> "Dishes are to do."
>
> "A watch is to hear it tick."

In a similar way, the word associations of young children, unlike those of older children, reflect the functions and perceptual attributes of the stimulus word as well as idiosyncratic contextual association. Thus the word seems to activate intraconceptual relations of a wide-ranging (noncore) variety. Later, the word appears to activate interconceptual relations such as coordinates and contrasts (Hall & Halperin, 1972; Nelson, 1977; Riegel, 1970).

In order to conceptualize freely and use the language appropriately and creatively, however, the core concept must become identified with the narrow linguistic concept and the inclusive concept must retain its predictive function while at the same time giving up its validating function. The child must become able to think about bad daddies as well as good daddies. But if bad daddies are possible, does it follow that everything is? How can the child avoid the conceptual chaos implicit in the realization that any relation can be posited even if it conflicts with what is known? To deal with this problem the child must come to use words only in their narrow socially agreed-upon definition while at the same time he retains all of the experiential material in the broad concept. The broad concept, therefore, serves as a source of predictions about future experiences but not as a restrictor of acceptable relations. This is a move that takes place much later than the initial word-learning period examined here, however, and its dimensions remain to be explored. These very sketchy suggestions are only indications that there is a great deal yet to be learned about the relation between conceptualizing and naming by the young child long after the period of first word learning.

CONCLUSION

This paper has attempted to show that the young child engages in a general conceptualizing process forming different kinds of concepts (such as particulars, reduplicates, similars) all of which involve similar operations designed to stabilize transient experience and to provide a basis for prediction about future events. To understand the acquisition of linguistic terms and their meanings we must understand the conceptual development process and how it interacts with language learning. Prototypes, images, and words can all be viewed functionally as part of the individual's attempt to make stable and discrete an inherently transient and continuous experience. It is suggested that while conceptualizing serves to order experience, the conceptual product is inherently less stable and discrete than language

itself or the representations and systematic descriptions we apply to it. The young child's conceptualizations, because they are less differentiated by context, and because they must be more open to new experience, are in turn less stable than the adult's. New analytic tools—less rigid than those currently used for semantic analysis—are needed if we are to understand the continuously interacting systems involved in this development and to ultimately answer the question asked at the outset of this paper.

ACKNOWLEDGMENTS

I would like to thank William Kessen for insightful comments on previous drafts; and for their helpful comments and reactions, I would like to thank the participants of the Conference on Language Learning and Thought in Infants (McGill University, May 1975).

REFERENCES

Arnheim, R. *Visual thinking.* London: Faber & Faber, 1969.

Bloom, L. *One word at a time.* The Hague: Mouton, 1973.

Bolton, N. *The psychology of thinking.* London: Methuen, 1972.

Brown, R. How shall a thing be called? *Psychological Review,* 1958, *65,* 14-21.

Jaulkender, P. J., Wright, J. C., & Waldron, A. Generalized habituation of concept stimuli in toddlers. *Child Development,* 1974, *45,* 1002-1010.

Fodor, J. A. The ontogenesis of the problem of reference: A review of some theories of linguistic symbols. In C. E. Reed (Ed.), *The learning of language.* New York: Appleton, 1971.

Goldberg, S., Perlmutter, M., & Myers, N. A. Recall of categorically related and unrelated lists by two-year-olds. *Journal of Experimental Psychology,* 1974, *18*, 1-8.

Hall, J. W., & Halperin, M. S. The development of memory-encoding processes in young children. *Developmental Psychology,* 1972, *6,* 181.

Harris, P. L. Inferences and semantic development. *Journal of Child Language*, 1975, *2,* 143-152.

Inhelder, B., & Piaget, J. *The early growth of logic in the child.* London: Routledge & Kegan Paul, 1964.

Katz, J. J. *Semantic theory.* New York: Harper & Row, 1972.

Katz, N. Superordinate naming and hierarchical classification of young children. Preliminary report, 1975.

Katz, N., Baker, E., & Macnamara, J. What's in a name? A study of how children learn common and proper names. *Child Development*, 1974, *45,* 469-473.

Krauss, R. *A hole is to dig.* New York: Harper & Row, 1952.

Labov, W. The boundaries of words and their meanings. Paper presented at the Conference on New Ways of Analyzing Variation in English, Washington, D.C., October, 1972.

Lewis, M. M. *Infant speech* (2nd ed.). London: Kegan Paul, 1951.

Nelson, K. Structure and strategy in learning to talk. *Monographs of the Society for Research in Child Development,* 1973, *38* (1-2, Serial No. 149).

Nelson, K. Concept, word and sentence: Interrelations in acquisition and development. *Psychological Review,* 1974, *81* (4), 267-285. (a)

Nelson, K. Variations in children's concepts by age and category. *Child Development,* 1974, *45,* 577-584. (b)

Nelson, K. The syntagmatic-paradigmatic shift revisited. *Psychological Bulletin,* 1977, *84,* 93-116.

Nelson, K., & Nelson, K. Cognitive pendulums and their linguistic realization. In K. E. Nelson (Ed.), *Children's language.* New York: Gardner Press, in press.

Piaget, J. *Play, dreams and imitation* (Translated by C. Gattegno and F. M. Hodgson). New York: Norton, 1962.

Riegel, K. F. The language acquisition process: A reinterpretation of related research findings. In L. R. Goulet and P. B. Baltes (Eds.), *Theory and research in life-span developmental psychology.* New York: Academic Press, 1970.

Rescorla, L. Concept formation in early word learning. Ph.D. dissertation, Yale University, 1976.

Saltz, E. *The cognitive basis of human learning.* Homewood, 111.: Dorsey Press, 1971.
Schank, R. C. Is there a semantic memory? Unpublished manuscript. Instituto per gli Studi Semantici Cognitivi, Castognola, Suitz., 1974.

Searle, J. *Speech acts.* Cambridge: Cambridge University Press, 1969.

Vygotsky, L. *Thought and language* (Translated by E. Haufman and G. Vakar). Cambridge: M.I.T. Press, 1962.

Webb, R. A., Oliveri, M. E., & O'Keefe, L. Investigations of the meaning of "different;" in the language of young children. *Child Development*, 1974, *45*, 984—991