

2

Accretion, Tuning, and Restructuring: Three Modes of Learning

David E. Rumelhart
Donald A. Norman

University of California, San Diego

It is somehow strange that throughout the recent work on semantic memory, the study of learning has been slighted. The term *learning* has fallen into disuse, replaced by vague references to “acquisition of information in memory.” It is easy to fall into the trap of believing that the learning of some topic is no more than the acquisition of the appropriate set of statements about the topic by the memory system. According to this simple view of things, to have learned something well is to be able to retrieve it from memory at an appropriate time. We believe this view is much too simple. Learning can be more than the simple acquisition of statements. We believe it is time to examine learning again, to evaluate just what does happen when people acquire information about a topic and use it appropriately.

The study of learning differs from the study of memory in its emphasis, not necessarily in content. Learning and memory are intimately intertwined, and it is not possible to understand one without understanding the other. But the difference in emphasis is critical. There are many different kinds of learning, and the characterization of the learning process most likely varies according to the type of learning that is taking place. Some forms of learning—especially the learning of relatively simple information—can probably be characterized correctly as a simple accumulation of new information into memory. However, especially when we deal with the learning of complex topics where the learning experience takes periods of time measured in months or even years, learning is much more than the successful storage of increasing amounts of information.

Complex learning appears to have an emergent quality. This learning seems to involve a modification of the organizational structures of memory as well as the

accumulation of facts about the topic under study. At times this modification of the organizational structure seems to be accompanied by a “click of comprehension,” a reasonably strong feeling of insight or understanding of a topic that makes a large body of previously acquired (but ill-structured) information fit into place. Thus the study of the learning of complex topics is related to the study of the understanding of complex topics.

This paper does not satisfy our desire for increased knowledge about the process of learning. Instead we simply hope to whet the appetite of our audience (and of ourselves). We present an analysis of learning and memory, attempting to examine some possible conceptualizations of the learning process, hoping thereby to guide the research of future years. We ourselves are just beginning the study of learning, and the start has proven frustratingly elusive. Indeed, it is the very elusiveness that has given rise to this paper. We now realize that simple characterizations of the learning process will not do. In this paper, we attempt a coherent account of the process of learning within our conceptualizations of a theory of long-term memory—the theory we have called *active structural networks* (cf. Norman, Rumelhart, & LNR, 1975). Our goal is to indicate how different forms of learning might be integrated into one conceptualization of the systems that acquire, interpret, and use information. This paper only sets the stage for development of theories and observations about learning. Hopefully, the stage is new, with useful characterizations that can be used to guide future developments, both of ourselves and of others.

LEARNING AND THE ACQUISITION OF KNOWLEDGE

Accretion, Restructuring, and Tuning

It is possible to learn through the gradual accretion of information, through the fine tuning of conceptualizations we already possess, or through the restructuring of existing knowledge. We find it useful to distinguish between these three qualitatively different modes of learning. Although we are not ready to propose a formal, rigid classification of learning, let us informally talk as if we could indeed classify learning into these three categories: *accretion*, *tuning*, and *restructuring*.

Learning through accretion is the normal kind of fact learning, daily accumulation of information in which most of us engage. The acquisition of memories of the day's events normally involves merely the accumulation of information in memory. A person's knowledge base is merely incremented by a new set of facts. Accretion is the normal learning that has been most studied by the psychologist. The learning of lists, dates, names of presidents, telephone numbers, and related things are examples of learning through accretion. Such learning presumably occurs through appropriate exposure to the concepts to be acquired, with the

normal stages of information processing transforming the information being acquired into some appropriate memory representation, which then is added to the person's data base of knowledge. In this case there are no structural changes in the information-processing system itself.

Learning through tuning is a substantially more significant kind of learning. This involves actual changes in the very categories we use for interpreting new information. Thus tuning involves more than merely an addition to our data base. Upon having developed a set of categories of interpretation (as seen in the following, we call these *schemata*), these categories presumably undergo continual tuning or minor modification to bring them more in congruence with the functional demands placed on these categories. Thus, for example, when we first learn to type, we develop a set of response routines to carry out the task. As we become increasingly better typists, these response routines become tuned to the task and we come to be able to perform the task more easily and effectively. For another example, presumably an analogous phenomenon is going on as a young child learns that not all animals are "doggies." Slowly his "doggie" schema becomes modified into congruence with the actual demands on his interpretation system.

Learning through restructuring is a yet more significant (and difficult) process. Restructuring occurs when new structures are devised for interpreting new information and imposing a new organization on that already stored. These new structures then allow for new interpretations of the knowledge, for different accessibility to that knowledge (usually improved accessibility), and for changes in the interpretation and therefore the acquisition of new knowledge.

Restructuring often takes place only after considerable time and effort. It probably requires some critical mass of information to have been accumulated first: In part, it is the unwieldiness and ill-formedness of this accumulated knowledge that gives rise to the need for restructuring.

We are impressed with the fact that real learning takes place over periods of years, not hours. A good deal of this time can be accounted for by the slow accretion of knowledge. There is an extensive amount of information that must be acquired and elaborate interconnections that must be established among all the information, fitting it into the general web of knowledge being developed within the memory system of the learner (see Norman, in press). But a good deal of time must also be spent in the development of the appropriate memory organizations for the evolution of existing memory structures (tuning) and the creation of new ones (restructuring). This learning requires new structures. Indeed, often the point of the learning is the formation of the new structures not the accumulation of knowledge. Once the appropriate structures exist, the learner can be said to "understand" the material, and that is often a satisfactory end point of the learning process. The accretion of information would appear to be a necessary prerequisite for restructuring; there must be a backlog of experiences and memories on which to base the new structures.

Note the long hours of study that seem to accompany the learning of many tasks. In intellectual domains, we expect students of scholastic topics to spend years of study, from undergraduate instruction, through graduate school, and then afterward, either through postdoctoral studies or as "budding young scholars," acquiring the knowledge and understanding of the field. The acquisition of intellectual knowledge probably continues throughout the lifetime of a scholar.

In skill learning, similar time periods are found. To our mind, the classic result in the literature is Crossman's (1959) study of cigar makers, whose performance continues to improve for at least ten years, with each cigar maker producing some 20 million cigars in that duration. Reaction time tasks in the laboratory have been carried out to at least 75,000 trials, again with continual improvement (Seibel, 1963). Similar figures can be produced for the learning of skills such as language, psychology, chess, and sports. People who are engaged in the serious task of learning a topic, whether it be an intellectual one or a motor skill (the difference is less than one might suspect), appear to show continual improvement even after years of study. As Fitts (1964) put it, "The fact that performance ever levels off at all appears to be due as much to the effects of physiological aging and/or loss of motivation as to the reaching of a true asymptote or limit in capacity for further improvement [p. 268]"

Learning, then, has several different components. In this paper, we concentrate primarily upon the qualitative differences among accretion of knowledge, restructuring of memory, and tuning of existing knowledge structures. Moreover, our discussion is primarily concerned with the latter two modes of learning. Restructuring involves the creation of entirely new memory structures, whereas tuning involves the evolution of old memory structures into new ones. Each of these processes—evolution and creation—can itself be performed in a number of different ways, each way being relevant to a different aspect of the learning process. But, before we can discuss the details of the learning process, we need to discuss our views of the structure of memory and, in particular, the organized memory units: *memory schemata*.

MEMORY SCHEMATA

General Schemata and Particular Instances

Memory contains a record of our experiences. Some of the information is *particular* to the situation that it represents. Other information is more *general*, representing abstraction of the knowledge of particular situations to a class of situations. The memory of eating dinner yesterday represents particular information. Knowledge that people eat meals from plates (using knives, forks, and spoons) represents general information that applies to a large class of situations.

A psychological theory of memory must be capable of representing both general and particular information. We believe that general information is best represented through organized information units that we call *schemata*. To us, a schema is the primary meaning and processing unit of the human information-processing system. We view schemata as active, interrelated knowledge structures, actively engaged in the comprehension of arriving information, guiding the execution of processing operations. In general, a schema consists of a network of interrelations among its constituent parts, which themselves are other schemata.

Generic concepts are represented by schemata. These schemata contain *variables*: references to general classes of concepts that can actually be substituted for the variables in determining the implications of the schema for any particular situation. Particular information is encoded within the memory system when constants—specific values or specific concepts—are substituted for the variables of a general schema. Our representations for specific events are thus instantiations of the general schema for that event¹ type. In some sense, one could consider schemata to represent prototypes of concepts.

A General Schema

A schema can represent an entire situation, showing the interrelationships among component events or situations (or subschemata). Thus we might have a schema for a concept such as *farming* that would contain the following information:

*A partial schema for farming.*²

A plot of land is used for the raising of agricultural crops or animals.

Some person cultivates the soil, produces the crops, and raises animals.

Typically farms raise some crops and have a few animals, including cows, horses, chickens, and pigs.

¹This formulation leaves open the question of whether particular representations result from general schemata or general schemata from particular ones. It is possible that our early experiences with some class of events give rise to a set of particular representations of those events. Then we generalize from these experiences by substituting variables for the aspects of the events that seem to vary with situations, leaving constants (particular concepts) in those parts of the representation that are constant across the different events in the class. The result is a general schema for a class of events. Alternatively, we can take a general schema and apply it to a new, particular situation by replacing the variable with constants. We presume that both of these directions continually take place: General schemata are formed through the process of generalization of particular instances; particular knowledge is derived from the principles incorporated within the general schemata.

²Note that this is a personal schema, one relevant to the conceptualizations of one of the authors (D.A.N.), who is horribly ignorant of real farms. This is proper: Schemata within the memory system of a given person reflect (constitute) his beliefs and knowledge. A schema may be wholly inaccurate as a description of the world, but it corresponds to the inaccuracies and misconceptions of the possessor of that schema. Assume that the author of this schema learned about farms through nursery rhymes.

Usually tractors and automated machinery are used to work the fields, and specialized buildings are used to house the products and animals.

... (etc.)

Once we have some general schema for farming, we can use it in a variety of ways. The general schema for farming can be viewed from several different perspectives. In so doing, we learn that:

The land is called a farm.

A farmer is the person who cultivates the land or raises the animals.

Livestock are animals kept on a farm for use or profit.

Farming is the act of cultivating the soil, producing crops, and raising animals.

Agriculture is the science and art of farming.

The barn is the building for housing farm animals.

Variables

The general schema for farming contains variable terms that can be further specified whenever the schema is used. Thus the general schema has the following variable terms:

land

crops or animals

some person

machinery

products

specialized buildings.

The particular values that get substituted for these terms depend upon the purpose for which the schema is being used. On different occasions, different substitutions will be made. If we learned that the Stewards have a carrot farm, then we substitute our concept for the Stewards as the group that plays the role of farmers in the schema, and we substitute carrots for the crops and products. We have substituted constants for these variables; however, some variables, such as *land*, *machinery*, and *buildings*, are still unspecified. Our general knowledge of carrots tells us something of the size of the farm and the kinds of machinery likely to be involved. Our schema for the growing of plants tells us that water and fertilizer are required. Our general schema for farming still has some free variables, but these are not without some constraints: We expect that there will be some animals, probably cows, chickens, horses, and pigs.

Constraints and Defaults

The different variables in a schema are often constrained: We do not expect to find all possible plants or animals on a farm. Tigers, eels, and poison ivy are animals and plants but not within the normal range of possible crops or livestock. Many of the variables in schemata have *default* values associated with

them. These are particular values for the variables that we can expect to apply unless we are told otherwise. Thus we might expect cows, pigs, horses, and chickens to be on a farm, and if nothing is said, we assume their presence. Similarly, we use the schema for *commercial transaction* for interpreting an occasion in which some person *A* has purchased item *O* from some other person *B*; we assume that money was transferred from *A* to *B*. We could be wrong. Money may not have been involved. Or, in the previous example, any particular farm may not have those animals. Nevertheless, these are the default values for our general understanding of the situations in question.

Variables (and their constraints) serve two important functions:

1. They specify what the range of objects is that can fill the positions of the various variables.
2. When specific information about the variables is not available, it is possible to make good guesses about the possible values.

The values for the variables for a schema are interrelated with one another. If a farm raises cattle, we expect a different size for the farm and different machinery and products than if the farm raises wheat, peanuts, or carrots. We would expect the buildings to look different. Similarly, if someone purchases an automobile, we expect a different amount of money to be involved than in the purchase of a pencil.

Schemata and Comprehension

We view a schema as a general model of a situation. A schema specifies the inter-relationships that are believed to exist among the concepts and events that comprise a situation. The act of comprehension can be understood as the selection of appropriate configuration of schemata to account for the situation. This means that there will be some initial selection of schemata and verification or rejection of the choices. A major portion of the processing effort involved in comprehension is directed toward determining the appropriate schemata for representing the situation. Once an appropriate configuration of schemata has been found, the constants of the situation have to be associated with (bound to) the variables of the schema. The schema that is selected will determine the interpretation of the situation. Different schemata will thereby yield different interpretations of the same situation, and different features of a situation will take on more or less importance as a function of that interpretation.

Like a theory, schemata vary in the adequacy with which they account for any given situation. Schemata both account for existing inputs and predict the values of others. If the account for the early observations is sufficiently good (and no other candidates emerge in subsequent processing), the schema will be accepted, even though there might be no evidence for some of its predictions. These predictions, then, constitute inferences about the situation that are made in the process of comprehension.

When a schema is sufficiently poor at describing the situation, a new schema must be sought. If no single adequate schema can be found, the situation can be understood only in terms of a set of disconnected subsituations—each interpreted in terms of a separate schema.

Schemata Are Active Data Structures

Although this is not the place to go into the details, we believe that the selection and use of schemata is controlled by the schemata themselves. We think of schemata as active processing units, each schema having the processing capability to examine whatever new data are being processed by the perceptual systems and to recognize data that might be relevant to themselves. Schemata activate themselves whenever they are appropriate to an ongoing analysis, and they are capable of guiding the organization of the data according to their structures. Schemata then can control and direct the comprehension process itself. We further suppose that the output of a schema (evidence that the concept represented by the schema is in the input) can then be introduced into the data pile for use by other schemata.

Perhaps the best way to view this is to think of all the data being written on a blackboard, with the schemata examining the blackboard for data relevant to themselves. When a schema sees something, it attempts to integrate the data into its organizational structure and then puts new information onto the blackboard. Other schemata may react to these new data. Thus schemata are *data driven* in the sense that they respond to the existence of relevant data. Schemata perform *conceptually driven* guidance to the processing by using their internal conceptualizations to add new data to the blackboard, thereby guiding the processing of other schemata. Thus each schema is data-driven and provides conceptually guided guidance to others. Further details of this system can be found in a number of sources: The blackboard analogy comes from the work of Reddy (see Reddy & Newell, 1974); active demons are familiar concepts in modern computing systems, from the demons of Selfridge and Neisser (1960), to the actors of Hewitt, Bishop, and Steiger (1973), to the production systems of Newell (1973); descriptions of those concepts relevant to this discussion are to be found in some of our works, in particular Norman and Bobrow (1976), Rumelhart (1977) and Rumelhart and Ortony (1977).

LEARNING

The Accretion of Knowledge

One basic mode of learning is simply the accumulation of new information. We analyze the sensory events of our current experience, match them with some

appropriate set of schemata, form a representation for the experience, and tuck the newly created memory structures away in long-term memory. The newly created data structures are *instantiations* of the previously existing ones, changed only in that the representations for particular aspects of the current situation have been substituted for the variables of the general schema.

This is learning by accretion: learning by adding new data structures to the existing data base of memory, following the organization already present. Learning by accretion is the natural side effect of the comprehension process. In it, we store some interpretation of the actual experience. If later we retrieve the stored information, we use the instantiated schemata to reconstruct the original experience, thereby “remembering” that experience. The schemata guide reconstruction in much the same way that they guide original comprehension.

Accretion, and later retrieval through reconstruction, is the normal process of learning. It is the sort of learning that has traditionally been studied by psychologists, and it is most appropriate to the current developments in the study of memory. Learning through the accumulation of new memories allows the data base of information to be built up. It allows for the acquisition of the large amount of specific knowledge that humans acquire about topics in which they are specialists and about the operation of the world in general. Learning by accretion assumes that the schemata required in the interpretation of new input already exist. Whenever this is not the case, the sheer accretion of knowledge is not effective; there must be a modification of the set of available schemata. This can be brought about either by the evolution of existing schemata (tuning) or the creation of new ones (restructuring). Learning by tuning and by restructuring probably occurs much less frequently than does learning by accretion. But without these other learning processes, new concepts cannot be formed.

Learning by Restructuring

When existing memory structures are not adequate to account for new knowledge, then new structures are required, either by erecting new schemata specifically designed for the troublesome information or by modifying (tuning) old ones.

Both the creation and tuning of schemata go hand-in-hand in the learning process. Thus in learning a skill such as typing, new schemata for the appropriate actions must be developed. But once the basic motor schemata have been developed, then further increases in proficiency will come about through the tuning of the existing schemata. Similarly, in the learning of some complex topic matter, probably the first step is the accretion of a reasonable body of knowledge about the topic, followed by the creation of new schemata to organize that knowledge appropriately. Then, continued learning consists of further tuning of those schemata (as well as continued accretion of knowledge and possibly creation of other new schemata, which in turn then have to be tuned).

If the only learning processes were memory accretion and tuning, one could never increase the number of conceptual categories over those initially given. Thus it is essential that new schemata be created. Logically, there are two ways in which new schemata can be formed. First, a new schema can be patterned on an old one, consisting of a copy with modifications. We call this process *patterned generation* of schemata. Second, new schemata can be induced from regularities in the temporal and/or spatial configurations of old schemata. We call this process *schema induction*. It is a kind of contiguity learning.

Patterned generation of schemata is doubtless the source of a good deal of ordinary concept formation.³ Perhaps the simplest form of patterned generation occurs through the use of analogies. Thus, even if we never had direct experience with a *rhombus*, we could develop a schema for one by being instructed that a rhombus has the same relationship to a square that a parallelogram has to a rectangle. The rhombus schema can be created by patterning it on the square schema, modifying it in just the way the parallelogram schema differs from the rectangle schema.

Note that this is creation of a new schema by generalizing an old one. The modification involves replacing a constant term of the square schema (the right angles at the corner) with variables to produce a new, more general schema. Patterned schema generation can also occur through modifying old schemata, replacing some of the variable components of a schema with constants. Thus, for example, we might very well form the concept of a "cocker spaniel" by modifying the schema for "dog." In this case, we would pattern the cocker spaniel schema on the dog schema but with certain variables much more tightly specified.

Schema induction is a form of learning by contiguity. If certain configurations of schemata tend to co-occur either spatially or temporally, a new schema can be created, formed from the co-occurring configuration. Learning of this kind is probably the least frequent mode of learning (or equivalently the most difficult). Yet it is an important procedure for learning. The difficulty with induction is in the discovery of the regularities. We suspect that most schema creation occurs through patterned generation. Experienced teachers find that analogies, metaphors, and models are effective teaching devices. We do not often (ever) see temporal contiguity as an effective teaching tool in the classroom or in the acquisition of most complex topics. Temporal contiguity is the fundamental principle of most theories of learning, but it seems to have amazingly little direct application in the learning of complex material. As far as we can determine,

³Note that we are *not* referring to the concept identification tasks that have been studied within the laboratory. The normal experiments on concept formation probably involve very little learning. Probably these tasks have been more concerned with problem solving, where the subjects are asked to discover the rules that will properly classify the particular stimulus set under study.

most complex concepts are learned because the instructor either explicitly introduces an appropriate analogy, metaphor, or model, or because the learner happens across one. We believe that most learning through the creation of new schemata takes place through patterned generation, not through schema induction.

Schema Tuning

Existing schemata can often serve as the base for the development of new ones by minor changes: by “fine tuning” of their structure. We call this process *tuning*. We restrict the use of the term *tuning* to those cases where the basic relational structure of the schema remains unchanged, and only the constant and variable terms referred to by the schema are modified. These terms can be changed in four ways:

1. *Improving the accuracy.* The constraints of the variable terms of the schema can be improved to specify the concepts that fit the variables with more accuracy.

2. *Generalizing the applicability.* The range of a given variable can be generalized to extend its range of applicability. Either the constraints on a variable can be relaxed, or a constant term can be replaced with an appropriately constrained variable term.

3. *Specializing the applicability.* The range of a given variable can be constrained by adding to the constraints of the variable, in the extreme, by effectively replacing the variable with a constant term.

4. *Determining the default values.* The values of the variable that normally apply can be discovered and added to the specification of the schema. Whenever a particular variable is not specified, the default values provide intelligent guesses that can be used in making inferences and guiding further processing.

The adjustment of variable constraints must be an important mechanism of learning.⁴ We must learn over what ranges variables vary; we must learn how the various variables co-vary. Our processing increases in efficiency if a schema specification is accurate and if we are not wasting time attempting to fit it to improper situations. Moreover, our understanding of a situation is more com-

⁴Note that there is really very little difference between constrained variables and constants. Schemata refer to terms with differing amounts of constraints upon the concepts that can be used in those terms. When the constraints are minimal, we have a free variable: Any concept can be substituted. Usually, the constraints specify some reasonable range of alternative concepts that can be used, excluding certain classes and allowing others. When the constraints are so restrictive that only a single unique concept can be used, then this is the equivalent of having a constant rather than a variable. In the normal case, schemata take variables that are partially constrained and thus provide some structure while at the same time represent a reasonable degree of generality.

plete if we account for it by a more, rather than less, specific schema. With more experience we can determine the typical values for the terms, providing information about default values to be used in the absence of further specification. The literature of language acquisition provides good illustrations of the role of variable adjustment. Let us look briefly at them.

Tuning to improve accuracy. The child must learn the range of conditions over which particular syntactic rules are applicable. Consider the child who can count and who realizes that the adjective meaning of the i -th element of a sequence can be formed by adding the suffix *th* to a number i . The child will correctly generate such words as *fourth*, *sixth*, *seventh*, etc. The child will, however, also generate words like *oneth*, *twoth*, *threeth*, *fiveeth*, etc. The child has too broad a rule: the rule is over regularized. The child must tune the general rule so that it has the correct constraints on its applicability. The process whereby the restrictions are learned involves adjusting the variables of the schema to permit its invocation only for the appropriate conditions. The schema must be tuned to improve its accuracy of application.

Tuning to generalize the applicability. Bowerman (in press) reports that young children use action words first only about themselves, then later generalize them to other people and animals, and finally use them for inanimate objects as well. This would appear to be a case where the schema must be tuned by loosening the variable constraints to make it more generally applicable.

Generalization of schemata occurs when an existing schema is modified so as to apply to a wider range. One example is when the meaning of a term is *extended* to cover other cases. This process, called *metaphorical extension* by Gentner (1975), was illustrated by her use of the word "have" in the following examples:

1. Sam *has* a large kettle.
2. Sam *has* a nice apartment.
3. The kettle *has* an enamel coating.
4. Sam *has* good times.

Presumably the verb "have" gets a primary meaning of something like "own." By extension, aspects of the owning relationship become inessential to the application of the concept of "having." Originally "have" would seem to require the owner to be one with complete control over the object in question. As the usage gets extended, the requirement of having complete control is loosened until finally, by sentence (4), it appears to require only that the object in question be strongly associated, in some way, with the subject.

Although it is much more common in language acquisition to find cases of children overgeneralizing a concept, which then must be restricted in its range of application, there are cases reported in which children first over-restrict the application of a term and then must *generalize* its use to the entire conceptual

category. Thus, Dale (1976) reports a case in which a child first applied the word "muffin" to only blueberries and blueberry muffins but not to other muffins. The process whereby the word comes to be extended to other muffins involves generalization of schemata.

In general, reasoning by analogy would seem to involve the generalization of a schema. In this case, a schema that is applicable in one domain is extended to a new domain by modifying one or more of its elements but maintaining the bulk of its internal structure. Thus, for example, when we consider fog "creeping on little cat's paws," the "creep" schema must somehow be extended to fog. Although this extension probably doesn't involve much learning, it follows the same principles that we have in mind.

Tuning to specialize applicability. A common occurrence in the child's acquisition of language is to overgeneralize the words, to use one word for a much larger set of circumstances than is appropriate. Thus a child may call all small animals "doggie" or all humans "mamma." Clark (1973) summarizes much of the literature on this phenomenon. Overgeneralization probably occurs because the child has selected too few features to identify the concept, when so many things will satisfy the definition. The child must specialize its understanding of the schema by either restricting the range of the variable terms or by adding more terms that must be followed before the schema is acceptable. Specialization by the first method fits our notion of tuning. Specialization by the second actually would be a form of patterned generation of schemata: forming a new schema based upon the old but modified by adding a few more terms.

Children may learn to use the term "ball" to apply to all small objects. They must learn to restrict the class of objects to which the term applies. Similar examples have been reported with the use of relational terms like "more-less," "long-short," "big-wee," etc. (cf. Donaldson & Wales, 1970). Children first learn to apply either term when the appropriate dimension is in question and then learn to restrict the application of the concepts to the appropriate direction on the dimension. Again, additional structure is inserted into the relevant schemata.

A similar process may very well be involved in becoming skillful at a motor task. At first when we learn to carry out a complex motor task, there is broad variation in the movements used to accomplish the task, but with experience in the situation, the variability of the movements is reduced. Consider, as an example, learning to juggle. At first we have great difficulty. We often toss the ball too high or too low. Our catching hand has to reach for the balls as they fall. With practice, our throws become increasingly precise. We come to be able to anticipate where the ball will fall with increasing accuracy. It would thus seem that at the early stages of learning to juggle, the appropriate schemata are only loosely interrelated—any of a variety of components may be configured together. With practice new constraints are added to our juggle

schema, and it becomes an increasingly precise, well-tuned schema (see Norman, 1976).

LEARNING IS NOT A UNITARY PROCESS

One major point of this paper is that learning is not a unitary process: No single mental activity corresponding to learning exists. Learning takes place whenever people modify their knowledge base, and no single theoretical description will account for the multitude of ways by which learning might occur. Indeed, we do not believe that we have necessarily described all the varieties of learning in this short classification. But we have attempted to demonstrate a reasonable variety of the classes of learning that might occur, with a description of the mechanisms that might be responsible for them. The classification is summarized in Figure 2.1.

It is interesting to note that the different kinds of learning occur in complementary circumstances. Memory accretion is most efficiently done when the incoming information is consistent with the schemata currently available. In this case, the information is easily assimilated. The more discrepant the arriving information from that described by the available schemata, the greater the necessity for change. If the information is only mildly discrepant, tuning of the schemata may be sufficient. If the material is more discrepant, schema creation is probably required. Of course, in order for restructuring to occur, there must be recognition of the discrepancy. But when mismatched by the available schemata, the learner may so misinterpret (misunderstand) the material that the discrepancies might not even be noted. The need for restructuring might only be noted with mild discrepancies when the misfit is glaring.

This discussion has concentrated on descriptions of the changes that take place in the memory schemata during learning. We have not discussed the mechanisms that might operate to cause these changes. The mechanisms for accretion are reasonably well developed: This is the process most frequently studied, most capable of being described by most theories of memory. We suspect that schema tuning is also a relatively straightforward operation, one that might not require much different mechanisms than already exist in theories of memory. But the restructuring of memory through the creation of new schemata is quite a different story. Here we know little of the process whereby this might take place. Moreover, we suspect that the occasions of schema creation are not frequent. Reorganization of the memory system is not something that should be accomplished lightly. The new structure that should be formed is not easy to determine: The entire literature on "insightful" learning and problem solving, on creativity, on discovery learning, etc., can probably be considered as reports of studies of how new schemata get created. We do not believe that the human memory system simply reorganizes itself whenever new patterns are discovered:

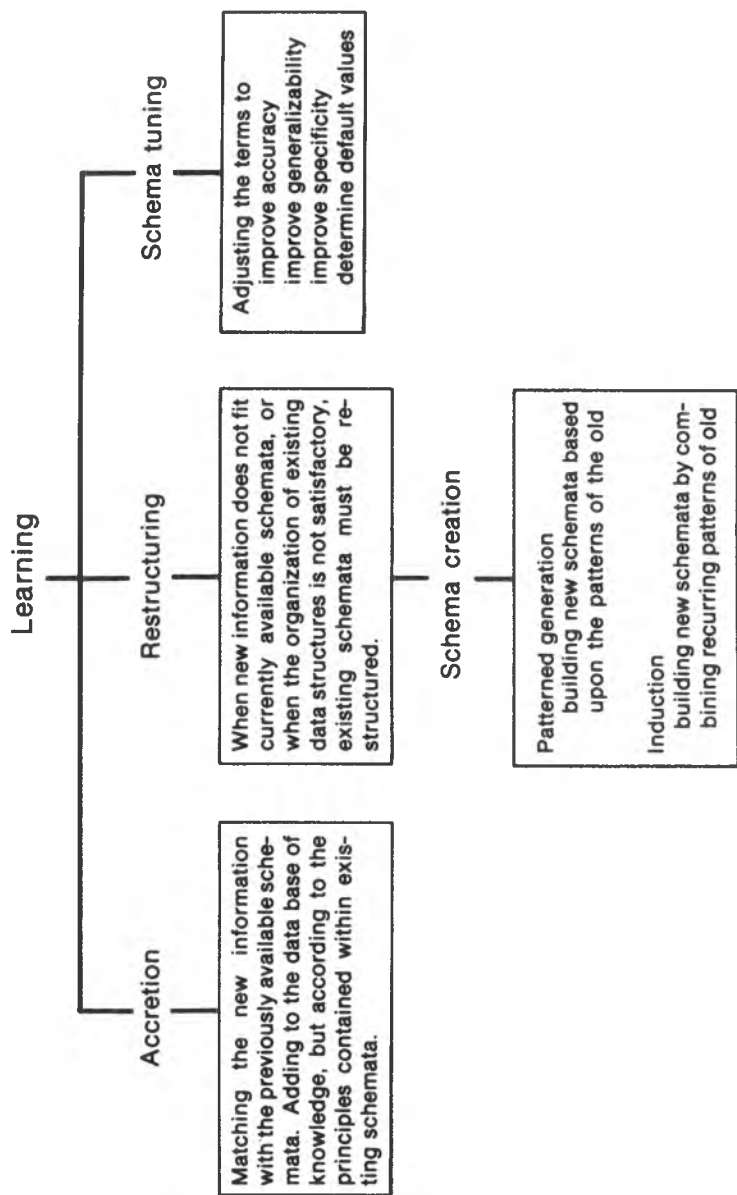


FIG. 2.1. A classification of the mechanisms by which learning might occur.

the discovery of patterns, the matching of analogous schemata to the current situation, must probably require considerable analysis. This is the area that we believe requires the most study in the future.

ACKNOWLEDGMENTS

This research was supported by the Advanced Research Projects Agency and the Office of Naval Research of the Department of Defense and was monitored by ONR under Contract No: N00014-76-C-0628.

REFERENCES

- Bowerman, M. Semantic factors in the acquisition of rules for word use and sentence construction. In D. Morehead & A. Morehead (Eds.), *Language deficiency in children: Selected readings*. Baltimore: University Park Press, in press.
- Clark, E. V. What's in a word: On the child's acquisition of semantics in his first language. In T. E. Moore (Ed.), *Cognitive development and the acquisition of language*. New York: Academic Press, 1973.
- Crossman, E. R. F. W. A Theory of the acquisition of speed-skill. *Ergonomics*, 1959, 2, 153-166.
- Dale, P. S. *Language development: Structure and function*. New York: Holt, Rinehart & Winston, 1976.
- Donaldson, M., & Wales, R. J. On the acquisition of some relational terms. In R. Hayes (Ed.), *Cognition and the development of language*. New York: Wiley, 1970.
- Fitts, P. M. Perceptual-motor skill learning. In A. W. Melton (Ed.), *Categories of human learning*. New York: Academic Press, 1964.
- Gentner, D. Evidence for the psychological reality of semantic components: The verbs of possession. In D. A. Norman, D. E. Rumelhart, & the LNR Research Group, *Explorations in cognition*. San Francisco: W. H. Freeman, 1975.
- Hewitt, C., Bishop, P., & Steiger, R. *A universal modular ACTOR formalism for artificial intelligence*. Proceedings of the Third International Conference on Artificial Intelligence, Stanford, California, 1973.
- Newell, A. Production systems: Models of control structures. In W. G. Chase (Ed.), *Visual information processing*. New York: Academic Press, 1973.
- Norman, D. A. *Memory and attention* (2nd ed.). New York: Wiley, 1976.
- Norman, D. A. Learning complex topics. *American Scientist*, in press.
- Norman, D. A., & Bobrow, D. G. On the role of active memory processes in perception and cognition. In C. N. Cofer (Ed.), *The structure of human memory*. San Francisco: W. H. Freeman, 1976.
- Norman, D. A., Rumelhart, D. E., & the LNR Research Group. *Explorations in cognition*. San Francisco: W. H. Freeman, 1975.
- Reddy, R., & Newell, A. Knowledge and its representation in a speech understanding system. In L. W. Gregg (Ed.), *Knowledge and cognition*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1974.
- Rumelhart, D. E. Toward an interactive model of reading. In S. Dornic (Ed.) *Attention and performance, VI*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.
- Rumelhart, D. E., & Ortony, A. The representation of knowledge in memory. In R. C.

- Anderson, R. J. Spiro, & W. E. Montague (Eds.) *Schooling and the acquisition of knowledge*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.
- Seibel, R. Discrimination reaction time for a 1,023-alternative task. *Journal of Experimental Psychology*, 1963, 66, 215-226.
- Selfridge, O. G., & Neisser, U. Pattern recognition by machine. *Scientific American*, 1960, 203, 60-68.