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 event[1] 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.

[1] This 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.

[2] Note 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 cultives 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.