Rayner, K. The perceptual span and peripheral cues in reading. *Cognitive Psychology,* 1975, *7,* 65–81.

Rayner, K., & Hagelberg, E. M. Word recognition cues for beginning and skilled readers. *Journal of Experimental Child Psychology,* 1975, *20,* 444–455.

Rosenberg, S., & Schiller, W. Semantic coding and incidental sentence recall. *Journal of Experimental Psychology,* 1971, *90,* 345–346.

Rothkopf, E. Z. Structural text features and the control of processes in learning from written materials. In J. B. Carroll & R. O. Freedle (Eds.), *Language comprehension and the acquisition of knowledge.* Washington, D.C.: V.H. Winston, 1972.

Savin, H. B. What the child knows about speech when he starts to learn to read. In J. G. Kavanaugh & I. G. Mattingly (Eds.), *Language by ear and by eye.* Cambridge, Mass.: MIT Press, 1972.

Schlesinger, I. M. *Sentence structure and the reading process.* New York: Humanities Press, 1969.

Shankweiler, D. E., & Liberman, I. Y. Misreading: A search for causes. In J. F. Kavanagh & I. G. Mattingly (Eds.), *Language by ear and by eye.* Cambridge, Mass.: MIT Press, 1972.

Simon, H. A. How big is a chunk? *Science,* 1974, *183,* 482–488.

Sinclair-de-Zwart, H. Developmental psycholinguistics. In D. Elkind & J. H. Flavell (Eds.), *Studies in cognitive development: Essays in honor of Jean Piaget.* Oxford: Oxford University Press, 1969.

Slobin, D. I. *The ontogenesis of grammar: Facts and theories.* New York: Academic Press, 1971.

Smiley, S. S., Oakley, D. D., Worthen, D., Campione, J. C., & Brown, A. L. Recall of thematically relevant material by adolescent good and poor readers as a function of written versus oral presentation. *Journal of Educational Psychology,* 1978, *69,* 381–387.

Smith, F. *Psycholinguistics and reading.* New York: Holt, Rinehart & Winston, 1973.

Spiro, R. J. Remembering information from text: Theoretical and empirical issues concerning the 'State of Schema' reconstruction hypothesis. In R. C. Anderson, R. J. Spiro, & W. E. Montague (Eds.), *Schooling and the acquisition of knowledge.* Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.

Taplin, J. E., Staudenmayer, H., & Taddonio, J. L. Developmental changes in conditional reasoning: Linguistic or logical? *Journal of Experimental Child Psychology,* 1974, *17,* 360–373.

Till, R. E., Cormak, D. R., & Prince, P. L. Effects of orienting tasks on sentence comprehension and cued recall. *Memory and Cognition,* 1977, *5,* 59–66.

Venezky, R. L. *Theoretical and experimental base for teaching reading.* The Hague, Netherlands: Mouton, 1976.

Venezky, R. L., & Massaro, D. W. *The role of orthographic regularity in word recognition.* Madison: Wisconsin Research and Development Center for Cognitive Learning, 1976.

Wallach, L., Wallach, M. A., Dozier, M. G., & Kaplan, N. E. Poor children learning to read do not have trouble with auditory discrimination but do have trouble with phoneme recognition. *Journal of Educational Psychology,*1977, *69,* 36–39.

Weaver, P. A. *Improving reading comprehension: Effects of sentence organization instruction.* Unpublished manuscript, Harvard University, 1977.

Weber, R. First-graders' use of grammatical context in reading. In H. Levin & J. Williams (Eds.), *Basic Studies in Reading,* New York: Basic Books, 1970.

Weinstein, R., & Rabinovitch, M. Sentence structure and retention in good and poor readers. *Journal of Educational Psychology,* 1971, *62,* 25–30.

Wittrock, M. C., Marks, C., & Doctorow, M. Reading as a generative process. *Journal of Educational Psychology,* 1975, *67,* 484–489.

2 Schemata: The Building Blocks of Cognition

David E. Rumelhart
University of California, San Diego
and
Center for the Study of Reading

The notion of a schema and the related notions of beta structures, frames, scripts, plans, and so on have formed the focus of research in Cognitive Science over the past 3 or 4 years. (cf. Bobrow & Norman, 1975; Chafe, 1976; Minsky 1975; Moore & Newell, 1973; Rumelhart, 1975; Schank & Abelson, 1975; Winograd, 1975). In this paper, I introduce these concepts to those unfamiliar with them and show why so much attention has been paid to them. These various terms have been used by different authors to refer to any of a set of interrelated concepts. These terms are *not all synonymous.* Different authors have different things in mind when they use these different terms. Nevertheless, the various concepts are closely enough related that a discussion of any one of them will serve as an introduction to the others. I thus focus my discussion on the one I know best, *schemata* (the singular is *schema*), as developed in Rumelhart and Ortony (1977).

The term *schema* comes into psychology most directly from Bartlett (1932). Bartlett himself attributes his use of the term to Head (1920). However, it would appear that Kant's (1787/1963) use of the term already anticipated its major conceptual content. The *OED* gives the following definition of the term: "In Kant: Anyone of certain forms of rules of the 'productive imagination' through which the understanding is able to apply its 'categories' to the manifold of sense-perception in the process of realizing knowledge or experience." Some further discussion of Kant's view is given in Rumelhart and Ortony (1977). It is because of this historical precedence that I have chosen to retain the term *schema.*

For all of the aforementioned authors, schemata truly are *the building blocks of cognition.* They are the fundamental elements upon which all information processing depends. Schemata are employed in the process of

interpreting sensory data (both linguistic and nonlinguistic), in retrieving information from memory, in organizing actions, in determining goals and subgoals, in allocating resources, and, generally, in guiding the flow of processing in the system. Clearly, any device capable of all these wondrous things must be powerful indeed. Moreover, because our understanding of none of these tasks that schemata are supposed to carry out has reached maturity, it is little wonder that a definitive explication of schemata does not yet exist and that skeptics view theories based on them with some suspicion. In this chapter, I spell out, as clearly as possible, the nature of schemata and the kinds of problems they were devised to solve. In addition, I present a convincing case that the framework provided by schemata and allied concepts does, in fact, form the basis for a reasonable theory of human information processing.

My discussion through the next several sections of the paper is abstract. Although I do not make direct application of these concepts to a theory of reading until near the end of the chapter, many of the chapters in this volume illustrate the ways in which schemata can lead to insightful analyses of the reading process.

WHAT IS A SCHEMA?

A schema theory is basically a theory about knowledge. It is a theory about how knowledge is represented and about how that representation facilitates the *use* of the knowledge in particular ways. According to schema theories, all knowledge is packaged into units. These units are the schemata. Embedded in these packets of knowledge is, in addition to the knowledge itself, information about how this knowledge is to be used.

A schema, then, is a data structure for representing the generic concepts stored in memory. There are schemata representing our knowledge about all concepts: those underlying objects, situations, events, sequences of events, actions and sequences of actions. A schema contains, as part of its specification, the network of interrelations that is believed to normally hold among the constituents of the concept in question. A schema theory embodies a *prototype* theory of meaning. That is, inasmuch as a schema underlying a concept stored in memory corresponds to the *meaning* of that concept, meanings are encoded in terms of the typical or normal situations or events that instantiate that concept.

Rather than attempting a formal description of schemata and their characteristics at this point, I turn instead to some useful analogies to give the reader a more concrete notion of the nature of schemata as I understand them. I turn first to one of the more fruitful analogies, that of a play.

Schemata Are Like Plays

The internal structure of a schema corresponds, in many ways, to the script of a play. Just as a play has characters that can be played by different actors at different times without changing the essential nature of the play, so a schema has *variables* that can be associated with (bound to) different aspects of the environment on different instantiations of the schema. As an example, consider the schema for the concept *buy*. One can imagine a playwright having written a most mundane play in which the entire play consisted of one person purchasing some object from another person. At minimum, such a play must have two people, some merchandise, and some medium of exchange. Whatever else happens, at the outset of the play one character (call him or her the PURCHASER) must possess the medium of exchange (call it the MONEY). The second person, the SELLER must possess the object in question, the MERCHANDISE. Then, by some interaction (BARGAINING) a bargain is struck and the SELLER agrees to give the MERCHANDISE to the PURCHASER in exchange for a quantity of the MONEY. There would, of course, be many ways of playing this little play. The MERCHANDISE could vary from a trinket of little value to an object of incalculable worth. The SELLER and the PURCHASER could vary in status, occupation, sex, nationality, age, and so on; the MONEY could vary in amount and whether it was actually money or clam shells; and the BARGAINING could vary in form. Still, through all of this variation, as long as the fundamental plot remained the same, we could say that the BUY play was being performed.

Now, this little play is very much like the schema that I believe underlies our understanding of the concept *buy* or that for *sell*. There are variables, corresponding to the characters in the play. We have the PURCHASER, the SELLER, the MONEY, the MERCHANDISE, and the BARGAINING. When we understand a situation to be a case of BUYING, we come to associate persons, objects, and subevents with the various variables of our schema. Having made these associations, we can determine to what degree the situation we are observing corresponds to this *prototype* case of BUYING.

Just as a playwright often specifies characteristics of the characters in his play (age, sex, disposition, etc.), so, too, as part of the specification of a schema, we have associated knowledge about the variables of the schema. We know, for example, that the PURCHASER and SELLER are normally people and that the MONEY is normally money. Moreover, we know that the value of the MONEY in question will covary with the value of the MERCHANDISE, and so on. Such knowledge about the typical values of the variables and their interrelationships is called the *variable constraints*.

These constraints serve two important functions in a schema theory. In the first place, variable constraints help in the identification of the various aspects

of the situation with the variables of the schema. If we know that we are observing a case of BUYING, we are not going to map the PURCHASER variable into the object in the world that should serve as the MONEY. We know this, in part, because we know that the PURCHASER is normally an animate being, whereas the MONEY is normally money or some other inanimate object. In the second place, variable constraints can help by serving as *default values* (cf. Minsky, 1975) or initial "guesses" for variables whose values we have not yet observed. Thus, for example, if we take a certain transaction to be one of BUYING but do not notice the MONEY, we can *infer* that there was MONEY and that, in fact, the MONEY probably *was money* amounting in value to about the value of the MERCHANDISE. In this way, the schema can help us make inferences about unobserved aspects of a situation.

It is perhaps useful to note here that variable constraints offer default values for unobserved variables *conditional on the values of the observed variables*. Moreover, the constraints are not *all-or-none* constraints that *require* that certain variables have a fixed range of values. Rather, they are merely specifications of the *normal* range of values for each variable and how this normal range varies with the specification of various combinations of other values on the other variables. Thus, as Rumelhart and Ortony (1977) suggest, it is perhaps most useful to think of variable constraints as forming a kind of *multivariate distribution* with correlations among the several variables.

There is also the notion of an *instantiation of a schema* that corresponds to an *enactment of a play*. A play is enacted whenever particular actors, speaking particular lines, perform at a particular time and place. Similarly, a schema is *instantiated* whenever a particular configuration of values is bound to a particular configuration of variables at a particular moment in time. Interpreting a situation to be an instance of some concept, such as an instance of buying, involves, according to the present view, the instantiation of an appropriate schema, say the BUY schema, by associating the various variables of the schema with the various aspects of the situation. Such a schema, along with its variable bindings, is called an instantiated schema. Just as we could, say, take a movie of an enactment of a play and thereby save for posterity a trace of the enactment, likewise it is the traces of our instantiated schemata that serve as the basis of our recollections.

Before leaving the analogy between the script of a play and a schema, it is useful to note that neither is a complete specification of every detail—both allow room for irrelevant variation and creative interpretation. The script of a play, no matter how meticulous the playwright, allows for an infinity of variations, each of which can properly be considered an enactment of the play. Certain lines composed by the playwright are sometimes changed to suit the interpretation of the director. Nevertheless, within limits, it is the same

play. So it is with schemata. A schema is not so rigidly applied that no variation is allowed. The schema only provides the skeleton around which the situation is interpreted. Variations orthogonal to the specifications of the schema have no bearing on the quality with which the schema is said to account for the situation. Moreover, even minor aspects of the situation that might be considered central to the schema can undergo some variation before we completely reject the interpretation provided by the schema.

Finally, despite all of the ways in which a schema is like a play, there are also numerous ways in which a schema is unlike a play. Perhaps most important of these is degree of abstraction. In our example of the BUY schema, we imagined a play that was more abstract than one any playwright, would ever compose. Normally, the playwright would determine the *kind* of buying involved, as well as more detail about the characters and more constraints on the dialogue. The BUY schema, on the other hand, must be applicable to *any* case of buying and, thus, must necessarily be more abstract that any actual play would ever be. Moreover, whereas a play is normally about people and their actions, a schema may be about events and objects of any sort. Indeed, a schema may merely be about the nature of a wholly inactive object such as a chair. In this case, the schema specifies not action or event sequences but spatial and functional relationships characteristic of chairs. Finally, although a play may contain acts, each with their own structure, a script for a play exists really only on one level. A script does not consist of a configuration of subscripts. A schema, on the other hand, should be viewed as consisting of a configuration of subschemata corresponding to the constituents of the concept being represented. These points are made clearer in the following sections, where I draw analogies between schemata and other familiar concepts.

Schemata Are Like Theories

Perhaps the central function of schemata is in the construction of an interpretation of an event, object, or situation—that is, in the process of comprehension. In all of this, it is useful to think of a schema as a kind of informal, private, unarticulated theory about the nature of the events, objects, or situations that we face. The total set of schemata we have available for interpreting our world in a sense constitutes our private theory of the nature of reality. The total set of schemata instantiated at a particular moment in time constitutes our internal model of the situation we face at that moment in time, or, in the case of reading a text, a model of the situation depicted by the text.

Thus, just as the activity surrounding a theory is often focused on the evaluation of the theory and the comparison of the theory with observations we have made so it is that the primary activity associated with a schema is the

determination of whether it gives an adequate account for some aspect of our current situation. Just as the determination that a particular theory accounts for some observed results involves the determinations of the *parameters of the theory,* so the determination that a particular configuration of schemata accounts for the data presently available at our senses requires the determination of the values of the *variables of the schemata.* If a promising schema fails to account for some aspect of a situation, one has the options of accepting the schema as adequate in spite of its flawed account or of rejecting the schema as inadequate and looking for another possibility. Therefore, the fundamental processes of comprehension are taken to be analogous to hypothesis testing, evaluation of goodness to fit, and parameter estimation. Thus, a reader of a text is presumably constantly evaluating hypotheses about the most plausible interpretation of the text. Readers are said to have understood the text when they are able to find a configuration of hypotheses (schemata) that offers a coherent account for the various aspects of the text. To the degree to which a particular reader fails to find such a configuration, the text will appear disjointed and incomprehensible.

Schemata are like theories in another important respect. Theories, once they are moderately successful, become a source of predictions about unobserved events. Not all experiments are carried out. Not all possible observations are made. Instead, we use our theories to make inferences with some confidence about these unobserved events. So it is with schemata. We need not observe all aspects of a situation before we are willing to assume that some particular configuration of schemata offers a satisfactory account for that situation. Once we have accepted a configuration of schemata, the schemata themselves provide a richness that goes far beyond our observations. On deciding that we have seen an automobile, we assume that it has an engine, headlights, and all of the standard characteristics of an automobile. We do this without the slightest hesitation. We have complete confidence in our little theory. This allows our interpretations to far outstrip our sensory observations. In fact, once we have determined that a particular schema accounts for some event, we may not be able to determine which aspects of our beliefs are based on direct sensory information and which are merely consequences of our interpretation.

Schemata Are Like Procedures

There are at least two inadequacies of the analogies presented earlier. In the first place, plays and theories are passive. Schemata are active processes. In the second place, the relationship between a theory and its constitutent subtheories or between a play and its constitutent subplays are not always evident. Schemata, on the other hand, have a very well-defined constitutent structure.

In both of these ways, schemata resemble procedures or computer programs. Schemata are active computational devices capable of evaluating the quality of their own fit to the available data. That is, a schema should be viewed as a procedure whose function it is to determine whether, and to what degree, it accounts for the pattern of observations. This includes, among other things, associating its variables to the appropriate aspects of its environment—that is, binding its own variables. Thus, to the degree that schemata underlying concepts are identified with *meaning of those concepts,* a schema theory is *both* a *prototype theory* and a *procedural theory* of meaning. Obviously, the degree to which a schema theory of human information processing can work depends on the degree to which procedures can actually be constructed to carry out the tasks I have just assigned to them. I believe they can, and I address this issue in the following sections.

The second characteristic that schemata share with procedures is a structural one. Procedures normally consist of a network (or a tree) of subprocedures. A particular procedure normally carries out its task by invoking a pattern of subprocedures, each of which in turn operates by invoking its subprocedures. Each procedure or subprocedure can return values that can serve as conditions determining which other subprocedures, if any, are to be invoked. So it is with schemata. A schema is a network (or possibly a tree) of subschemata, each of which carries out its assigned task of evaluating its goodness of fit whenever activated. These subschemata represent the conceptual constituents of the concept being represented.

Thus, for example, suppose we had a schema for a FACE. This would consist of a certain configuration of subschemata, each representing a different constituent of a face. For example, there would presumably be a subschema representing the MOUTH, one for the NOSE, and one for each EAR and each EYE. These subschemata would, in turn, consist of a configuration of constituents. The EYE schema, for example, would consist of a configuration of subschemata including, perhaps, an IRIS, EYELASHES, and EYEBROW, and so on.

Just as a procedure uses results produced by its subprocedures to carry out its task, so too a schema uses results produced by its subschemata to carry out its tasks. As I indicated previously, the primary activity of a schema is the evaluation of its goodness of fit. An important mechanism of this evaluation involves the evaluation of the goodness of fit of each of its constituent parts. Thus, if a good EYE is found and a good MOUTH is found, the FACE schema can use this information along with its own evaluation of whether the entire *configuration* is right for a face to generate an overall evaluation of its goodness of fit.

To summarize, then, just as a procedure consists of subprocedures and those subprocedures, in turn, consist of more subprocedures, and so on, so a schema consists of subschemata each of which, in turn, is specified as a

configuration of its subschemata, and so on. One may be struck by the fact that this process must stop somewhere. If each and every schema were merely a configuration of subschemata, the process would never end. The solution to this dilemma for schemata is identical to the solution for procedures. When a computer program is written, this embedding process does not continue indefinitely. Eventually, some subprocedure consists entirely of a configuration of *elementary instructions* for the machine in question. Likewise, with schemata, there must be a set of schemata that are elementary in the sense that they do not consist of a further breakdown in terms of subschemata. Such elementary schemata correspond to what Norman and Rumelhart (1975) call *primitives.*

Schemata Are Like Parsers

A parser is a device that, given a sequence of symbols, determines whether that sequence forms a legal sentence (according to the rule of some grammar) and, if it does, determines the *constituent structure* of the sentence. That is, it determines which symbols in the sequence correspond to which constituents of the sentence. The process of finding and verifying appropriate schemata is thus a kind of parsing process that works with conceptual elements—finding constituents and subconstitutents among the data currently impinging on the system in much the same way that a sentence parser must find the proper parse for the input string of words.

One particularly useful aspect of this analogy is the substantial body of work carried out in computational linguistics on various parsing procedures. I believe that the processing strategies developed for some of the most sophisticated of these carry over nicely in their application to schemata in general. As I discuss later, I have in mind here especially the work of Kaplan (1973) and his development of the general syntactic processor (GSP).

Summary of the Major Features of Schemata

Rumelhart and Ortony (1977) listed four major characteristics of schemata. These were:

1. Schemata have variables.
2. Schemata can embed, one within another.
3. Schemata represent knowledge at all levels of abstraction.
4. Schemata represent knowledge rather than definitions.

The analogies just presented illustrate all of these features. Whereas schemata have variables, so plays have roles, theories have parameters, and procedures have arguments. The embedding characteristic of schemata is best illustrated

by the analogy between schemata and procedures. Schemata consist of subschemata as procedures consist of subprocedures. Just as theories can be about the grand and the small, so schemata can represent knowledge at all levels—from ideologies and cultural truths to knowledge about what constitutes an appropriate sentence in our language, to knowledge about the meaning of a particular word, to knowledge about what patterns of excitations are associated with what letters of the alphabet. We have schemata to represent all levels of our experience, at all levels of abstraction. Finally, our schemata *are* our knowledge. All of our generic knowledge is embedded in schemata.

In addition to these four features, the analogies presented here indicate at least two more general features of schemata:

5. Schemata are active processes.
6. Schemata are recognition devices whose processing is aimed at the evaluation of their goodness of fit to the data being processed.

THE CONTROL STRUCTURE OF SCHEMATA

Perhaps the central questions in the development of a schema-based model of perception and comprehension are: First, how is an adequate configuration of schemata discovered? Second, how is the goodness of fit evaluated? These are largely problems of *control structures.* There are many schemata. Not all of them can be evaluated at once. Somehow, there must be a scheme for activating just those schemata that are most promising. There are two basic sources of activation for schemata. These are usually referred to as *top-down* and *bottom-up* activation. These two directions correspond to what Bobrow and Norman (1975) have called *conceptually driven* and *data-driven* processing. I turn now to a discussion of these two modes of activation.

Conceptually Driven and Data-Driven Processing

A schema may activate a subschema in the way a procedure invokes its subprocedures. This is called *conceptual-driven* processing. In a sense, conceptually driven processing is expectation-driven processing. That is, when a schema is activated and it, in turn, activates its subschemata, the activation of these subschemata derive from a sort of expectation that they will be able to account for some portion of the input data. For example, suppose that, through some mechanism, the FACE schema is considered a promising account for the input and thereby activated and set about evaluating its goodness of fit. The *promise* of the FACE schema is, in a sense, transferred to its MOUTH, NOSE, EYE, EAR, and so on, subschemata.

A second mechanism for schema activation is bottom-up or *data-driven* activation. A schema is said to be activated from the bottom-up whenever a subschema that has been somehow activated causes the various schemata of which it is a part to be activated. If the activation of the FACE schema led to the activation of the PERSON schema, we would say that the activation of the PERSON schema was data driven. Thus, where *conceptually driven* activation goes from *whole to part, data-driven* activation goes from *part to whole.* In schema-directed processing, activation goes in *both* directions.

Schema-directed processing is assumed to proceed in roughtly the following way: Some event occurs at the sensory system. The occurrence of this event "automatically" activates certain "low-level" schemata (such schemata might be called *feature detectors*). These low-level schemata would, in turn, activate (in a data-driven fashion) certain of the "higher level" schemata (the most probable ones) of which they are constituents. These "higher level" schemata would then initiate conceptually driven processing by activating the subschemata not already activated in an attempt to evaluate its goodness of fit.

At some point, when one of these higher level schemata has begun to get further positive results about its goodness of fit (i.e., it has found evidence for other of its constituents), it would activate still higher level schemata, that would look for still larger constituents.

This higher, more abstract schema would then activate, from the top down, still other of its constituent schemata, and this activation would flow through its subschemata back down to lower level schemata. Lower level schemata would eventually either make contact with other schemata that have been activated from the bottom up or would initiate a search for the "predicted" sensory inputs.

Whenever a schema initiates a search for sensory data that are not present, that counts as evidence against that schema and also as evidence against all of those schemata that require the presence of that schema as a constituent subschema. When sufficient evidence is accumulated against a schema, processing of that schema is suspended and processing resources are allocated to other currently more promising schemata. Whenever enough evidence is gained in favor of a schema, that schema is taken as an adequate account for the relevant aspect of the input, and the interpretation offered by that schema is taken as the "correct" interpretation of the relevant event. Later processing on other, higher level schemata may eventually disconfirm a temporarily accepted schema, and we will have the phenomenon of the "double-take."

My discussion of the processing system to this point has been rather abstract. In the following section I examine, in some detail, an example of this mixed initiative processing system. In addition, Woods (Chapter 3, this volume) gives a more complete discussion of these issues.

An Example

Consider the following brief passage:

> Business had been slow since the oil crisis. Nobody seemed to want anything really elegant anymore. Suddenly the door opened and a well-dressed man entered the showroom floor. John put on his friendliest and most sincere expression and walked toward the man.

Although merely a fragment, most people generate a rather clear interpretation of this story. Apparently, John is a car salesman fallen on hard times. He probably sells rather large elegant cars—most likely, Cadillacs. Suddenly a good prospect enters the showroom where John works. John wants to make a sale. To do that he must make a good impression on the man. Therefore, he tries to appear friendly and sincere. He also wants to talk to the man to deliver his sales pitch. Thus, he makes his way over to the man. Presumably, had the story continued John would have made the sales pitch and, if all went well, sold the man a car.

How do people arrive at such an interpretation? Clearly, people do not arrive at it all at once. As the sentences are read, schemata are activated, evaluated, and refined or discarded. When people are asked to describe their various hypotheses as they read through the story, a remarkably consistent pattern of hypothesis generation and evaluation emerges. The first sentence is usually interpreted to mean that business is slow *because* of the oil crisis. Thus, people are inclined to believe that the story is about a business that is suffering as a result of the oil crisis. Frequent hypotheses involve either the selling of cars or of gasoline. A few interpret the sentence as being about the economy in general. The second sentence, about people not wanting elegant things anymore, leads people with the gas-station hypothesis into a quandary. Elegance just does not fit with gas stations. The gas station hypothesis is weakened, but not always rejected. On the other hand, people with hypotheses about the general economy or about cars have no trouble incorporating this sentence into their emerging interpretation. In the former case they conclude that people don't buy luxury items and in the latter they assume that people do not buy large elegant cars—Cadillacs—much anymore. The third sentence clinches the car interpretation for nearly all readers. They are already looking for a business interpretation—that most probably means a SELLING interpretation—and when a *well-dressed man* enters the door he is immediately labeled as someone with MONEY—a prospective BUYER. The phrase *showroom floor* clearly invalidates the gas-station interpretation and strongly implicates automobiles, which are often sold from a showroom. Moreover, the occurrence of a specific event does not fit at all

well with the view that the passage is a general discussion of the state of the economy. Finally, with the introduction of John, we have an ideal candidate for the SELLER. John's actions are clearly those stereotypic of a salesman. John wants to make a sale and his "putting on" is clearly an attempt on his part to "make a good impression." His movement toward the man fits nicely into this interpretation. If he is a salesman, he must make contact with the man and deliver the stereotypic "pitch."

Qualitatively, this little account (which was derived from an analysis of a number of readers describing their current interpretation of the story after each sentence) fits well with the general approach I have been outlining. The process of comprehension is very much like the process of constructing a theory, testing it against the data currently available, and as more data become available, specifying the theory further—that is, refining the default values (as perhaps was the case when those holding the "car hypothesis" from the beginning encountered the sentence about nobody wanting anything elegant anymore). If the account becomes sufficiently strained, it is given up and a new one constructed, or, alternatively, if a new theory presents itself that obviously gives a more cogent account, the old one can be dropped and the new one accepted.

But where do these theories come from? The theories are, of course, schemata. Presumably, through experience, we have built up a vast repertoire of such schemata. We have schemata for salesmen, the kinds of motives they have, and the kinds of techniques they employ. We have schemata for automobiles, including how and where they are sold. We have built up schemata for the "oil crisis"—what kinds of effects it has on what kinds of businesses. We have schemata about business people, the kinds of motives they have, and the kinds of responses they make to these motives. The knowledge embedded in these schemata form the framework for our theories. It is some configuration of these schemata that ultimately forms the basis for our understanding.

But how does a relevant schema suggest itself? It is here that the control structures discussed earlier play an essential role. Presumably, it is the "bottom-up" observation that a certain concept has been referenced that leads to the suggestion of the initial hypotheses. The notion that business was slow suggests schemata about business and the economy. Because the slowness was dated from the occurrence of the oil crisis, it is a natural inference that the oil crisis was the *cause* of the slowness. Thus, a BUSINESS schema is activated. The particular TYPE of business is presumably a variable that must be filled. The information about the oil crisis suggests that it may be an oil-related business. Thus, readers are led to restrict the TYPE variable of the BUSINESS schema to oil-related businesses.

At this point, after the bottom-up activation of the high-level BUSINESS schema has occurred, this schema would generate a top-down activation of

the various possible oil-related businesses. Prime candidates for these are, of course, automobile-related businesses. Of these, selling gasoline and automobiles are the two most salient possibilities.

When the second sentence is encountered, an attempt is made to fit it into the schemata currently considered most promising. As I discussed previously, this information could serve to further restrict the TYPE variable in the automobile BUSINESS schema but does not fit well with the gasoline business schema.

The BUSINESS schema presumably has, as part of its specification, a reference to the BUY or SELL schema discussed previously. Once activated, these schemata search for potential variable bindings. In the case of the automobile business, the MERCHANDISE variable is bound to an automobile. The second sentence suggests an elegant automobile. When the third sentence is encountered, the reader has not yet found a candidate for BUYER or SELLER. The sentence about a well-dressed man immediately suggests a potential BUYER. The phrase "showroom floor" offers additional bottom-up support for the automobile hypothesis. In fact, it is a strong enough clue itself that it can suggest automobile sales to a reader who currently considers an alternative schema more likely. We thus have a BUYER and some MERCHANDISE. The well-dressed quality of the BUYER is consistent with our view that the MERCHANDISE is elegant and therefore expensive—being well-dressed suggests MONEY. We need only a SELLER—that is, an automobile salesman. Readers probably already bring a relatively complete characterization of the "default value" for car salesmen. We need but little additional information to generate a rather detailed description of his goals and motives.

In spite of the length of this example, it should be noted that I have provided only a sketch of the elaborate processing that must occur in the comprehension of even so simple and direct a story as this. The problem is indeed a complex one, and no one yet has been able to construct a model capable of actually carrying out the tasks involved. It is the conviction that the concept of the *schema* is the most promising route to the solution to these problems that has led to its current popularity.

THE MAJOR FUNCTIONS OF SCHEMATA

My intent to this point has been primarily definitional. I have tried to show what schemata were in general and how they generally are supposed to work. In this section I give a few examples, mostly taken from the psychological literature, of phenomena for which schemata appear to offer promising accounts. I first turn to a discussion of *perception,* especially as it relates to reading.

Schemata and Perceiving

There are numerous examples in the psychological literature that suggest a schema-like theory to account for them. I mention just a few examples here. Perception, like language comprehension, is an interactive process. Information comes in from our sense organs, which suggest but do not determine appropriate schemata for the interpretation of the sense data. It is often only in the context of the whole that the individual parts of an object can be identified. Similarly, the whole itself cannot be identified apart from its parts. The interpretation of parts and wholes must proceed jointly. Our final interpretation is determined both by the local clues and by consistency among the various levels of analysis. Consider, as an example, Fig. 2.1 taken from Palmer (1975). The object on the left is clearly recognizable as a face, but its parts (series A) are not recognizable out of context. Thus, it cannot be that we first perceive the parts and then construct an interpretation of the whole. Rather, the various shapes of the lines *suggest*, but do not determine, possible interpretations (the wiggly line suggests a possible nose, the acute angle suggests a possible eye, etc.). Lower level NOSE and EYE schemata may be activated, which in turn may activate higher level schemata such as the FACE schema. The FACE schema then activates schemata for all of the parts of the FACE not receiving bottom-up activation. (For example, the lips may not be close enough to LIPS to activate this schema at all out of context. In this case, the LIPS schema would be activated by the FACE schema and find sufficient evidence to serve—in context—to count as LIPS.)

As can be noted from series *B* of Fig. 2.1, it is not that parts of a face cannot *ever* be recognized without the face as a context. But, in order to be

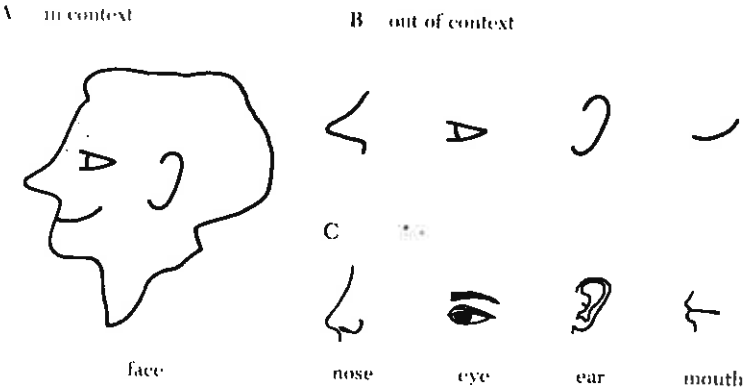

FIG. 2.1. An illustration of part-whole context. Facial features recognizable in the context of a profile (A) are not recongizable out of context (B). When the internal part structure of the facial features is differentiated (C), however, the features become recognizable out of context (Palmer, 1975).

recognized out of context, they, too, must have an internal structure. If enough data are available about its internal structure, a schema, such as the NOSE schema, can serve the function of an organizing whole perfectly well.

There is ample evidence of similar processes in reading. It is well known, for example, that the strings of characters that form words are more easily apprehended than strings that do not form words. The reason for this presumably stems from the fact that we have schemata corresponding to words and none for random letter strings. Just as evidence for a NOSE indirectly constitutes evidence for LIPS through the FACE schema, so too evidence for one letter can constitute evidence for other letters through the schema for the word in question. Thus, for example, evidence favoring a *T* in the first position and an *E* in the third position of a three-letter word indirectly constitutes evidence for an *H* in the second position through activation of the THE schema. The use of such information is presumably the mechanism whereby words are easier to see than random letter strings. Moreover, one of the characteristics that may separate skilled readers from those with less skill is the availability of a greater number of more completely developed word schemata.

It is interesting that schemata not only *contribute* toward the development of an accurate percept, but, by the same token, they can sometimes cause a distortion. An experiment by Bruner and Potter (1964) illustrates the debilitating effect of premature commitment to a particular schema. In their study, subjects were presented with defocused slides of familiar objects. The slides were slowly brought into focus. At each step along the way, as the slides were brought into focus, subjects were to report their best guess of what the slide was. Under these conditions, subjects continued to misidentify the object long after naive subjects (those started with less severe amounts of defocusing) were able to readily identify the object in question.

This result is presumably due to the fact that subjects became committed to their early interpretations of the slide and then required more information to disconfirm their original hypothesis than is normally required.

Schemata and Understanding Discourse

As discussed previously, the process of understanding discourse is the process of finding a configuration of schemata that offers an adequate account of the passage in question. The analysis of the "oil crisis story" given earlier illustrates generally how such a process is supposed to operate. Clues from the story suggest possible interpretations (instantiations of schemata) that are then evaluated against the successive sentences of the story until finally a consistent interpretation is discovered. Sometimes, a reader fails to correctly understand a passage. There are at least three reasons implicit in schema theory as to why this might occur:

1. The reader may not have the appropriate schemata. In this case he or she simply cannot understand the concept being communicated.
2. The reader may have the appropriate schemata, but the clues provided by the author may be insufficient to suggest them. Here again the reader will not understand the text but, with appropriate additional clues, may come to understand it.
3. The reader may find a consistent interpretation of the text but may not find the one intended by the author. In this case, the reader will "understand" the text but will misunderstand the author.

There are numerous examples of these three phenomena in the literature. Perhaps the most interesting set of studies along these lines were carried out by Bransford and Johnson (1973). They studied the comprehension of texts for which subjects could not provide the appropriate schemata, texts in which the schemata were potentially available but there were not sufficient clues to suggest the correct ones as well as texts in which subjects were led to choose a "wrong" interpretation. Consider, as an example, the following paragraph used in one of Bransford and Johnson's (1973) studies:

> The procedure is actually quite simple. First you arrange things into different groups. Of course, one pile may be sufficient depending on how much there is to do. If you have to go somewhere else due to lack of facilities, that is the next step, otherwise you are pretty well set. It is important not to overdo things. That is, it is better to do too few things at once than too many. In the short run this may not seem important but complications can easily arise. A mistake can be expensive as well. At first the whole procedure will seem complicated. Soon, however, it will become just another facet of life. It is difficult to forsee any end to the necessity for this task in the immediate future, but then one can never tell. After the procedure is completed, one arranges the materials into different groups again. Then they can be put into their appropriate places. Eventually they will be used once more and the whole cycle will then have to be repeated. However, that is part of life [p. 400].

Most readers find this passage extremely difficult to understand. However, once they are told that it is about washing clothes, they are able to bring their clothes-washing schema to the fore and make sense out of the story. The difficulty with this passage is, thus, not that readers do not have the appropriate schemata; rather, it stems from the fact that the clues in the story never seem to *suggest* the appropriate schemata in the first place. The "bottom-up" information is inadequate to initiate the comprehension process appropriately. Once the appropriate schemata are suggested, most people have no trouble understanding the text.

Although most readers simply find the passage incomprehensible, some find alternative schemata to account for it and thus render it comprehensible.

Perhaps the most interesting interpretation I have collected was from a Washington bureaucrat who has no difficulty with the passage. He was able to interpret the passage as a clear description of his job. He was, in fact, surprised to find that it was supposed to be about "washing clothes" and not about "pushing papers." Here then, we have an example of the third kind of comprehension failure, "understanding the story" but "misunderstanding the author."

Obviously, a detailed account of the comprehension process requires a detailed description of the schemata readers have available as well as an account of the conditions under which certain of these schemata are activated. A number of researchers have been developing such specific models of specific schemata (cf. Rumelhart, 1975, 1977; Schank & Abelson, 1977).

Schemata and Remembering

In addition to the important role assigned to schemata in comprehension and perception, schemata are assumed to be the guiding forces behind remembering as well. Perhaps the clearest way to see this is to consider some of the commonalities between remembering and apprehension. Rumelhart and Ortony (1977) suggest that the process of remembering is essentially similar to the process of perceiving, except that, in remembering, the data source is no longer sensorial but memorial. To quote:

> There is thus a kind of continuum between understanding and remembering where in the former we have the imposition of an interpretation primarily on incoming "sensory fragments," and in the latter we have the imposition of an interpretation primarily on "memorial fragments." In both cases schemata are employed. It should be emphasized that although remembering can be thought of as perceiving with memory as the modality, the episodic memories on which it is usually based are not merely fragments of the initial sensory input, but a fragmentary representation of our interpretation of that input [p. 27].

Thus, important aspects of the memory process can be seen as identical to the process of comprehension. I have thus, implicitly alluded to two ways in which schemata affect recollection. First, they are the mechanisms whereby initial interpretations are formed and, as such, they determine the *form* of the memorial fragments. Second, schemata are used to *reinterpret* the stored data in order to *reconstruct* the original interpretation. There is ample evidence for both of these roles. The first point suggests that we remember our *interpretations* of an event or text rather than the text or event itself. Bartlett's (1932) original finding that we remember the *gist* of a story rather than the details suggests this conclusion. Tulving and Thomson's (1973) arguments about encoding specificity would seem to require the same conclusion.

Perhaps the most convincing results, however, come from some recent experiments by Bransford and his associates (cf. Bransford, Barclay, & Franks, 1972; or Barclay, Bransford, Franks, McCarrell, & Nitsch, 1974). In one of their experiments (Bransford, Barclay, & Franks, 1972), they presented subjects with sentences drawn from pairs such as: (1a) The woman stood *on* the stool and the mouse sat on the floor beneath *it*. (1b) The woman stood *on* the stool and the mouse sat on the floor beneath *her*. Or sentences were drawn from pairs such as (2a) The woman stood *beside* the stool and the mouse sat on the floor beneath *it*. (2b) The woman stood *beside* the stool and the mouse sat on the floor beneath *her*. They found that on subsequent recognition tests, subjects could not tell which of the first pair of sentences had been presented to them, but those given sentences drawn from the second pair had no difficulty recognizing which had been presented. Because the two pairs differ by exactly the same number of words, it cannot be a difference in memory for the sentences per se that accounts for the difference in recognizability but rather differences in *interpretation* that account for the differential distinguishability between the two pairs.

The second role of schemata in remembering involves their use in reconstructing the original interpretation. Here, perhaps the best evidence comes from a recent experiment reported by Spiro (1977). Spiro designed his experiment to carefully discriminate between those inferences drawn at the time of comprehension and those drawn at the time of recall. In his experiment, subjects read stories that they were led to believe were true. Then, after reading the stories, subjects were given an additional piece of information that was either *consistent with the implications* of the story or *inconsistent with the implications* of the story. On later recall of the story, there was a general tendency for those with information inconsistent with the thrust of the story to *distort* their recall of the story so as to make it consistent with the later information. Moreover, this cannot entirely be due to a reinterpretation of the story at the time they received the additional information, because the longer the delay between the addition of the information and the recall test, the greater the degree of distortion. In terms of the framework under discussion in this paper, this fact suggests that the longer the amount of time between presentation and recall, the fewer memorial fragments are available and the more the subject must rely on generic knowledge of similar situations—that is, the subject's schemata. The less consistent the original information is from the typical, the more room (and need) there is for distortion.

Thus, just as comprehension is presumed to be identical to the process of selecting and verifying conceptual schemata to account for the situation to be understood, so the process of remembering involves, as a major component, the process of selecting and verifying an appropriate configuration of schemata to account for the memorial fragments found in memory. Such an account constitutes a recollection.

There is an important omission in this account of remembering. It has been tacitly assumed in our account of perception (and therefore in that for remembering) that data present themselves passively to the sensory system, directly exciting certain low-level "features" and that comprehension (or perception) involves the discovery of an adequate account of the spatio-temporal pattern of excitation of these low-level feature detectors. Similarly, in this account of remembering, it appears as if there is, somehow, a set of memorial fragments that are presenting themselves to the memory system for interpretation and that the process of interpretation (selecting potential configurations of schemata and verifying that they are consistent with the stored data—memory fragments) is simply all there is to remembering. *This is quite obviously false for both understanding and remembering.* First, consider the case of perception; then I will show how this extends to remembering.

Perception is goal directed. We do not passively wait for some stimuli to arrive and then at that late date attempt an interpretation. Instead, we actively seek information relevant to our current needs and goals. When we want a phone number, we do not merely interpret our current sensory input as if it were a phone number but we actively *seek* information to be so interpreted. This information-seeking process must go hand in hand with the information interpretation process. Just as expectations (embodied by certain activated schemata) can serve an important function in guiding our process of interpreting input that happens to reach our sensory organs, so these same schemata guide our *information seeking*. Not only do schemata tell us *what to see,* but they also tell us *where to see it*. I have not emphasized this aspect of the comprehension process.

If perception is goal directed, then remembering is even more so. In remembering, we are not merely perusing around our memories, making sense about what "happens to come to mind." The ordinary case of memory is probably more akin to the process of looking for a phone number than watching a TV program or reading a text. Here the "search" problem is severe. Whatever memory probe or question we have available will bring some fragments of memory to mind (activate some of the appropriate memory structures). Perhaps sometimes this is enough—the interpretation of this fragment of memory is enough to respond to the question at hand. More often, however, this is probably not the case. The appropriate fragments are just not "in front of us," impinging on our memory interpretation system. In these cases we must instigate a *search*. Just as in the case of looking for a telephone number, the search is not random. Rather, it must be guided by schemata that represent the layout of our memories in the same way as schemata encode the layout of our homes or our telephone books. When we want a phone number, we realize we might be able to use a phone book. Thus, we might search our house for the phone book. Knowing the layout of our house, we know the typical location of the phone book and will immediately

go there to find it. Note the location of the phone book is a subgoal in the search for the phone number in much the same way that a "retrieval context" is a subgoal in a memory search. When we find the phone book, we use our knowledge of the structure of the phone book to guide us to the appropriate region of the book. Finally, once there, we use our expectations about the structure and meaning of phone numbers to construct an appropriate interpretation of the symbols on the page. Not much work has been done on mapping out the search paths through memory while trying to recollect events after a long delay. Recent work by Williams (1977) would seem to be a promising approach toward the study of these scan paths through memory.

Schemata and Learning

One of the central problems of a schema theory is a specification of the process (or processes) whereby new schemata are developed. Even if it is granted that a set of "hand-crafted" procedures *could* carry out the tasks assigned to them by schema theories, it remains to be shown that there are plausible learning procedures that could result in such a set of schemata. There is currently very little known about what kinds of learning principles would be necessary for such a development. From a logical point of view, there are three basically different modes of learning that are possible in a schema-based system.

1. Traces of comprehension processes form a sort of learning inasmuch as upon having understood some text or perceived some event, we can retrieve stored information about that text or event. Such learning corresponds roughly to "fact learning." Rumelhart and Norman (1977) have called this learning *accretion.*
2. Existing schemata may *evolve* or undergo change to make them more in tune with experience. Our concepts presumably undergo continual change as we gain more experience with new exemplars. This corresponds to the elaboration and refinement of concepts through continued experience. Rumelhart and Norman (1978) have called this sort of learning *tuning.*
3. The third sort of learning involves the *creation* of new schemata. This involves the actual development of new concepts. There are, in a schema theory, at least two ways in which new concepts can be generated: They can be patterned on existing schemata or they can (in principle) be induced from experience. Rumelhart and Norman call learning of new schemata *restructuring.*

I now turn to a discussion of these three modes of learning and the conditions under which they occur.

Accretion. Learning by accretion is probably the most common sort of learning. It is also the sort of learning that has least effect on the operation of

the system. Whenever new information is encountered, there is assumed to be some trace of the comprehension process laid down in memory. This memory trace is the basis for recollections. Generally, these traces are assumed to be partial copies of the original instantiated schemata. Thus, memory traces are assumed to be very much like schemata themselves. They differ only inasmuch as they are fragmentary and they have representations for particular aspects of the original situation in place of the variables of the original schemata. Thus, as we experience the world, we store, as a natural side effect of comprehension, traces that can serve as the basis of future recall. Later, during retrieval, we can use this information to reconstruct an interpretation of the original experience—thereby remembering the experience.

Such an accumulation of knowledge is the normal sort of learning. Although the accumulation of a substantial body of knowledge may be necessary for more fundamental kinds of learning, it causes no new schemata to be formed. Such learning occurs whenever our schemata available at the time of the original experience are deemed adequate for the interpretation of the experience. When we encounter a situation in which currently available schemata do not prove adequate, there is the possibility of schema change and thus a modification of the very devices through which we experience the world.

Tuning. Tuning involves the actual modification or evolution of existing schemata. There are essentially three ways in which schemata can evolve. First, our knowledge of the variable constraints and default values can be upgraded continuously as we continue to use the schemata. Whenver we find a case in which we determine that a certain schema offers an adequate account of a particular situation, we can modify the variable constraints and default values in the direction of the current experience. This will make the schema sensitive to slow changes in the population of cases to which the schema is applied. As this process continues, it will continue to sharpen the variables and default values to make the schema better represent the population of situations to which it is applied. Note, however, that this sort of tuning will only occur when the schema is deemed to offer an adequate account of the situation at hand. Thus, because cases that deviate widely from the appropriate variable constraints and default values will not be accommodated by the schema in question, change must be slow.

The second sort of tuning involves replacing a constant portion of a schema with a variable one—that is, adding a new variable to a schema. This sort of schema modification amounts to *concept generalization*—making a schema more generally applicable. Presumably, the occasion for such learning is the discovery, at some point in time, that a particular schema would offer a good account for a particular situation if only some presumably constant feature of the schema were allowed to vary. To the degree that a constant is merely a

variable with very tight constraints, this can be seen as a special case of the previous kind of tuning, namely, a case in which the change is from a variable with highly constrained constraints that becomes one with somewhat more relaxed constraints.

The third sort of tuning is, in a sense, the opposite of the last one, namely, the process of making a variable into a constant or specializing the use of the concept. One occasion for such learning would be the discovery that certain "outlier" situations are better accounted for by other schemata and that the apparent variable is better thought of as a constant. As before, this can also be thought of as a special case of changing variable constraints—in this case tightening them.

Restructuring. If accretion and tuning were the only learning mechanisms, no new schemata could be created. The third learning mode discussed previously involves the creation of new schemata. There are basically two ways in which new schemata can be formed. Rumelhart and Norman (1978) called these *patterned generation* and *schema induction.*

Patterned generation involves the creation of a new schema by copying an old one with a few modifications. Such learning is, in essence, learning by analogy. We learn that a new concept is like an old one except for a few differences. A new schema can differ from an old one by having variables where the old one had constants (a generalization of the old schema), by having constants where an old schema had variables (a further specification of the old schema), or by substituting a new variable or constant for an old variable or constant of the original schema. Once a new schema is created by such processes, the process of tuning will continue to modify the newly created schema to bring it more into line with experience.

The second way in which new schemata can be formed is through the process of schema induction. The notion here is that if a certain spatio-temporal configuration of schemata is repeated, there is reason to assume that the particular configuration forms a meaningful concept and a schema can be formed that consists of just that configuration. This, of course, is the classical contiguity learning. It is interesting that, in spite of the ubiquity of the notion of contiguity learning in learning theories of the past, there is no real *need* for it in a schema-based system. Provided we begin with a sufficiently general set of schemata, the processes of tuning, accretion, and patterned generation can carry us a long way. Schema induction does cause some difficulty for the notion of schemata as I have outlined them. In order for schema induction to work properly, we must posit some aspect of the system sensitive to the recurrence of configurations of schemata that do not, at the time they occur, match any existing schemata. Such a system is not a natural part of a schema-based system.

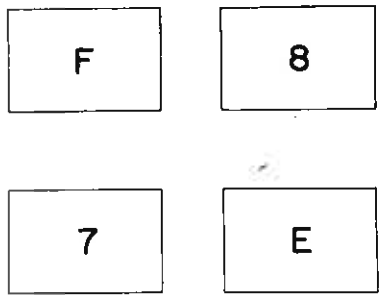

FIG. 2.2. Stimuli presented to the "Label Factory" subjects (from D'Andrade—see footnote 1).

Schemata and Solving Problems

Schemata play a central role in all of our reasoning processes. Most of the reasoning we do apparently *does not* involve the application of general-purpose reasoning skills. Rather, it seems that most of our reasoning ability is tied to particular schemata related to particular bodies of knowledge.

One of my favorite demonstrations of the critical role of schemata in reasoning comes from the work of Wason and Johnson-Laird (1972) and some more recent replications and extensions of their work carried out by Roy D'Andrade.[1] Subjects in D'Andrade's experiments were given one of two formally equivalent problems to solve. Half of the subjects were given the task illustrated in Fig. 2.2.

Subjects were shown the four cards illustrated in the figure and told that they were to imagine themselves as quality control experts in a label-making factory whose task it is to determine whether any labels are incorrectly constructed. A label is properly constructed if, when there is a vowel on one side of the label, there is an odd number on the other side. Then, they were asked to indicate which of the cards must be turned over to assure that the rule is being appropriately followed. Only 13% of the subjects correctly indicated that cards 2 and 4 must be checked. Card 4 must, of course, be checked because it may have an even number on the back. Card 2 must also be checked because it might have a vowel on the back and thus violate the rule. No other cards must be checked. However, about 70% of the subjects felt that card 3 should also be checked.

Here we have a classical failure of human reasoning. It has been argued that this is a case of our interpreting the simple conditional as a biconditional. In any case, results like these are often used to illustrate the weakness of the

[1]Roy D'Andrade has kindly given me access to the data from his as yet unpublished experiment.

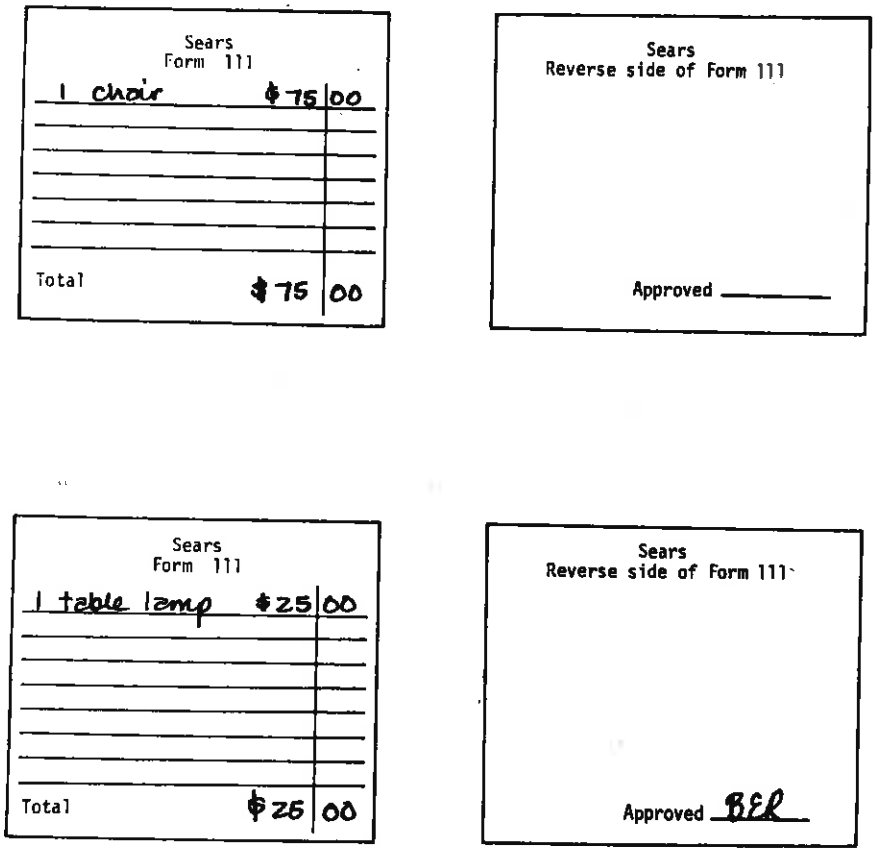

FIG. 2.3. Stimuli for the "Sears" subjects (from D'Andrade—see footnote 1).

human reasoning system. However, the results of the second part of D'Andrade's experiment point out the fallacy in this conclusion. Figure 2.3 illustrates the stimuli for these subjects. Subjects were told that they were to imagine themselves as store managers for Sears and that they were to inspect receipts at the end of the day to be certain that they were properly filled out. The rule is that, if any purchase exceeded $30, the receipt must have the signature of the department manager on the back. Again, subjects were asked which of the cards were supposed to be turned over; in this case nearly 70% indicated cards 1 and 2. Formally, these two problems are identical. Yet, when phrased in terms of the familiar setting of the Sears store, over five times as many subjects were able to correctly solve the problem. What is the difference here? Why do people appear not to understand the meaning of "if" in the first case and understand it nearly perfectly in the second? In terms of schema theory, the answer is rather straightforward. The first case is

unfamiliar and subjects have no schemata in which to incorporate the problem and, therefore, can bring only very general problem-solving strategies to bear on the problem. The second case more nearly approximates our "real life" problem-solving situations. Once we can "understand" the situation by encoding it in terms of a relatively rich set of schemata, the conceptual constraints of the schemata can be brought into play and the problem readily solved. It is as if the schema already contains all of the reasoning mechanism ordinarily required in the use of the schemata. Thus, understanding the problem and solving it is nearly the same thing.

CONCLUSION

It was my intent in this chapter to give the reader who is unfamiliar with schemata an intuition through which he or she could interpret the rest of this volume and the increasing number of studies employing these conceptualizations. I have aimed for generality rather than specificity in my account. I have tried to show the many domains to which the concept of a schema has been applied and the heuristic value of thinking about psychological and educational problems in terms of schemata. Many of the remaining chapters in this volume as well as some of the papers cited previously provide a more detailed account of schemata and more definite examples of their applicability. Although the development of schema-based theories such as the ones I mentioned previously is yet in its infancy and these ideas have not yet proved their usefulness, I believe that they offer the most promising leads for those of us interested in the difficult problems posed when we try to apply psychological theories directly to domains relevant to education.

ACKNOWLEDGMENT

This paper was written through the support of the National Institute for Education to the Center for the Study of Reading under Contract No. US-NIE-C-400-76-0116.

REFERENCES

Bartlett, F. C. *Remembering.* Cambridge: The Cambridge University Press, 1932.

Barclay, J. R., Bransford, J. D., Franks, J. J., McCarrell, N. S., & Nitsch, K. Comprehension and semantic flexibility. *Journal of Verbal Learning and Verbal Behavior,* 1974, *13,* 471–481.

Bobrow, D. G., & Norman, D. A. Some principles of memory schemata. In D. G. Bobrow & A. M. Collins (Eds.), *Representation and understanding: Studies in cognitive science.* New York: Academic Press, 1975.

Bransford, J. D., Barclay, J. R., & Franks, J. J. Sentence memory: A constructive versus interpretive approach. *Cognitive Psychology,* 1972, *3,* 193–209.

Bransford, J. D., & Johnson, M. K. Consideration of some problems of comprehension. In W. Chase (Ed.), *Visual information processing.* New York: Academic Press, 1973.

Bruner, J. S., & Potter, M. C. Interference in visual recognition. *Science,* 1964, *144,* 424–425.

Chafe, W. L. Givenness, contrastiveness, definiteness, subjects, topics, and point of view. In C. N. Li (Ed.), *Subject and topic.* New York: Academic Press, 1976.

Head, H. *Studies in neurology* (2 vols.). London: Frowde, 1920.

Kant, E. *Critique of pure reason* (2nd ed.; N. Kemp Smith, trans.). London: Macmillan, 1963 (originally published, 1787).

Kaplan, R. M. A general syntactic processor. In R. Rustin (Ed.), *Natural language processing.* New York: Algorithmics Press, 1973.

Minsky, M. A framework for representing knowledge. In P. H. Winston (Ed.), *The psychology of computer vision.* New York: McGraw-Hill, 1975.

Moore, J., & Newell, A. How can Merlin understand? In L. W. Gregg (Ed.), *Knowledge and cognition.* Potomac, Md.: Lawrence Erlbaum Associates, 1973.

Norman, D. A., Rumelhart, D. E., & The LNR Research Group. *Explorations in cognition.* San Francisco: Freeman, 1975.

Palmer, S. E. Visual perception and world knowledge: Notes on a model of sensory–cognitive interaction. In D. A. Norman, D. E. Rumelhart, & The LNR Research Group, *Explorations in cognition.* San Francisco: Freeman, 1975.

Rumelhart, D. E. Notes on a schema for stories. In D. G. Bobrow & A. M. Collins (Eds.), *Representation and understanding: Studies in cognitive science.* New York: Academic Press, 1975.

Rumelhart, D. E. Understanding and summarizing brief stories. In D. LaBerge & S. J. Samuels (Eds.), *Basic processes in reading: Perception and comprehension.* Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.

Rumelhart, D. E., & Norman, D. A. Accretion, tuning and restructuring: Three modes of learning. In J. W. Cotton & R. L. Klatzky (Eds.), *Semantic factors in cognition.* Hillsdale, N.J.: Lawrence Erlbaum Associates 1978.

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.

Schank, R. C., & Abelson, R. P. Scripts, plans, and knowledge. *Advance Papers of the Fourth International Joint Conference on Artificial Intelligence.* Tbilisi, Georgia, USSR, 1975.

Schank, R. C., & Abelson, R. P. *Scripts, plans, goals, and understanding.* Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.

Spiro, R. J. Remembering information from text: The "state of schema" approach. In R. C. Anderson, R. J. Spiro, & W. E. Montague (Eds.), *Schooling and the acquisition of knowledge.* Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.

Tulving, E., & Thomson, D. M. Encoding specificity and retrieval processes in episodic memory. *Psychological Review,* 1973, *80,* 352–373.

Wason, P. C., & Johnson-Laird, P. N. *Psychology of reasoning.* Cambridge, Mass.: Harvard University Press, 1972.

Williams, M. D. *The process of retrieval from very long term memory.* Unpublished doctoral dissertation, University of California, San Diego, 1977.

Winograd, T. Frame representations and the declarative–procedural controversy. In D. G. Bobrow & A. M. Collins (Eds.), *Representation and understanding: Studies in cognitive science.* New York: Academic Press, 1975.

3 Multiple Theory Formation in Speech and Reading

William A. Woods
Center for the Study of Reading,
Bolt Beranek and Newman Inc.

OVERVIEW

High-level perceptual tasks such as reading, speech understanding, and visual scene interpretation are characterized by the need to dicover a structured interpretation that accounts for the stimuli present. This process is prerequisite to deciding what has been perceived and thus precedes whatever process decides what to do with the resulting perception—what significance to attach to it, whether to remember it, how to incorporate it into the knowledge base of the perceiver, and so on. In this chapter, I attempt to make the case that the process of arriving at an interpretation of the input involves the formation and evaluation of many alternative partial hypotheses about the input and that this process goes on largely below the level of introspective awareness of the perceiver. Even though skilled reading involves a variety of "metacognitive" strategies (Brown, Chapter 19, this volume), in normal reading these processes are themselves invoked without conscious attention to the process of doing so.

I focus on the problem of reading and draw on insights and analogies from work in Natural Language Parsing and Continuous Speech Understanding. Because I do not have space here to give an adequate introduction to all of the background material that I would like to use, I will instead refer the reader to three previous papers: Woods, 1973a; Woods, 1975a; and Woods, 1975b. The most recent material on speech understanding is unfortunately only contained in technical reports. I recommend in particular the BBN final report (Woods, Bates, Brown, Bruce, Cook, Klovstad, Makhoul, Nash-Webber, Schwartz, Wolf, & Zue, 1976), vols. I, III, IV, and V. For a brief overview of the BBN speech understanding system, see Wolf and Woods (1977).