4 On the Classification Problem in Learning and Memory

Endel Tulving
University of Toronto

The purpose of the conference whose proceedings the present volume represents was to discuss and develop a common ground for conceptualizations in animal learning and human memory. This mandate is important and timely. On both historical and rational grounds, the fields of (animal) learning and (human) memory should have a good deal in common, their facts and theories should be mutually relevant, and practitioners in each of the two fields should find it easy to benefit from what those in the other have accomplished and what they are doing. Nevertheless, the two seem to represent separate cultures: they begin with different pretheoretical assumptions, they are concerned with apparently different problems, they employ different methods, and they speak different languages. In the animal learning literature, there are very few references to work in human memory; in the human memory literature there are even fewer mentions of work in animal learning.

A COLLABORATIVE ENTERPRISE

How should we go about meeting the challenge that the objective of the conference presents? How do we change the insularity of the two disciplines? How do we get a typical, busy student of animal learning to take an interest in, say, theoretical speculations surrounding the concept of levels of processing, or in the empirical findings concerning the retention of names and faces of one's fellow high-school students from 40 years ago? How do we make an equally chronically preoccupied student of human memory appreciate the relation between autoshaping and superstitious behavior in pigeons, or the latest "hot data" on the second-order fear conditioning?

Appeals to the good will of the practitioners on the other side of the fence, suggestions to read the other's literature, discussing one's work with the members of the other culture, at specially organized conferences or even in one's own department—these and other obvious devices are unlikely to be successful. At best they can be expected to produce a temporary arousal of curiosity. Even making available for deep study and thought the proceedings of the deliberations of a number of well-meaning practitioners from both cultures at the Umea conference holds only a little more promise of success. A highly motivated and open-minded reader may improve his or her understanding and appreciation of the strange customs and difficult language of the other culture, but it would be a miracle if, as a result of the exercise, a conversion to the unified field of animal learning and human memory took place. Such a field does not yet exist, and even several conferences such as ours will not bring about a rapid change. Any bridging of the current chasm will take time. All we can do now is to try to plant some seeds for the future development of a unified science of learning and memory.

In this chapter, I would like to propose that one possibly fruitful approach to the problem of bridging the two fields may lie in a collaborative enterprise aimed at solving a problem that is:

(a) *new* in the sense that it has not yet been (seriously) approached in either field,

(b) *relevant* in the sense that existing knowledge and previous accomplishments in both fields can be brought to bear on the effort,

(c) *important* in that the problem is perceived as having to do with basic and fundamental objectives of both fields,

(d) *pertinent* in that the pursuit of the problem can be imagined to lead to genuine progress in the understanding of learning and memory, and

(e) *difficult and challenging*, in that it is not at all obvious, at least at the outset, how the problem is to be solved.

A problem that satisfies these criteria is what I refer to as *the classification problem* of learning and memory. It parallels the classification problem in biology, with the difference that biological classification is concerned with living organisms—animals and plants—whereas in our case it would be concerned with varieties of learning and memory. A successful solution of the classification problem of learning and memory would take the form of a classificatory scheme, or system, embodying and expressing the relations among all varieties of learning and memory. Such a classificatory system would constitute an overall framework within which empirical phenomena and theoretical ideas regarding learning and memory can be integrated in a manner neither possible nor conceivable today. Ideally, we would have but a single classificatory system, one that in some real sense corresponds to nature, one that "carves nature at its joints."

Systematic research directed at the classification problem of learning and memory would entail the development of *systematics* as a new branch of the science of learning and memory, analogously with systematics in biology. In the course of such development, many subproblems of taxonomy would emerge and would have to be solved.

Thus, the proposal is for the undertaking of a collaborative venture: a broadly based attack on the classification problem within a new branch of study, the systematics of learning and memory. Such an undertaking would constitute, and bring about, a new perspective on learning and memory, one that would complement and enhance the value of the perspectives that have characterized the field for the better part of the last 100 years.

The chapter expands on the theme of the classification problem, describes the objectives of classification research in the field of learning and memory, discusses some of the problems that are likely to arise in the pursuit of the problem, and makes some tentative suggestions for general rules of procedure. We begin with a brief survey of some of the relevant features of the current scene.

EXPLANATION AND CLASSIFICATION: PRESENT STATE

Explanation in Learning and Memory

There are three major classes of explanation in life sciences, including psychology. They correspond to three kinds of questions that can be asked about biological phenomena: How does it work, what is it like, and where does it come from (cf. Bruce, in press; Solbrig, 1966). One is explanation in terms of underlying *causes*: postulation of mechanisms, processes, intervening variables, hypothetical constructs, and various kinds of hypothetical structures and functions that account for what is observed about the relation between behavior and situation. When we ask a question such as, "What causes forgetting?", or "Why is reinforcement effective?", we usually expect to be given such a causal explanation. The causal explanation has been the norm in the field of learning and memory. Acquisition and extinction of habits, strengthening and weakening of associations, generalization and discrimination, transfer and interference, retention and forgetting have usually been accounted for in terms of hypothetical happenings—functioning structures, mechanisms, processes, intervening variables, hypothetical constructs, and the like.

The second kind of explanation is one where a phenomenon of learning or memory is accounted for in terms of the *properties* of the organism exhibiting the learning or memory. We could refer to these as property explanations. Questions such as "Why can some tasks be learned more readily than others?", or "Why do older children learn Tasks X, Y, and Z faster than younger ones?" usually evoke responses specifying certain (other) characteristics of tasks or individuals under scrutiny.

The third explanation constitutes an extension of a part of the second: The characteristics of learners and rememberers are related to their *history*: previous experiences, individual development through growth and maturation, or the evolution of the species to which the organism belongs. Thus, differential ease of learning of different tasks is accounted for by transfer, differential learning abilities are interpreted in terms of age-related characteristics of the nervous system (neural, conceptual, or otherwise), and species-specific learning abilities are attributed to evolutionary selection pressures. We can refer to this form of explanation as historical. Property explanations and historical explanations have figured less prominently in learning and memory.

Explanations of events, effects, phenomena, or relations between variables, whatever their kind, are usually expressed in terms of more or less formal theories, models, and hypotheses. In principle, each of them is expected to do two things: (a) provide an explanation, and (b) specify the *domain* of the explanation, that is, state *what it is* that the theory, model, or hypothesis accounts for. The many theories, models, and hypotheses that we have in learning and memory have met the first expectation better than the second. Indeed, the domains of our theories, models, and hypotheses usually are either left implicit or specified in terms of arbitrary categories accidentally derived from historical practices of our research field.

A prerequisite for informed specification of domains of explanations is a classificatory system of different varieties of learning and memory. The creation of such a system is the objective of systematics of learning and memory.

Consider a single example of the current state of affairs in the matter of classification, although many others could be given. A recognized problem of fundamental importance in human memory concerns the relation between recall and recognition, attested to by the existence of a flourishing experimental and theoretical literature. (For a representative sample, see the references in Gillund & Shiffrin, 1984, and in Tulving, 1982.) The question of interest in the context of the present discussion is this: What is the domain of theories and models of recall and recognition (e.g., Gillund & Shiffrin, 1984; Tulving, 1982)?

The answer to the question seems to be straightforward: The domain is a collection of phenomena defined in terms of comparisons between recall and recognition in a variety of situations (Gillund & Shiffrin, 1984, pp. 6–7; Tulving, 1982, pp. 140–142). Problems emerge, however, when we ask additional questions: Is the domain of a given model limited to recall and recognition of verbal material, or even just discrete verbal items, or does it extend to other kinds of materials? Is it limited to college-student subjects, or adult humans, or all humans? Do the models have anything to say about recall and recognition in animals? Indeed, can one distinguish between recall and recognition in animals, either in the same way as in humans, or in some comparable fashion? If not, why not? If yes, why do we not have models of recall and recognition in the animal learning or animal memory literature? What kind of learning and memory is involved in comparisons between recall and recognition, and why do we seem to have it in humans but not in animals? If recall and recognition are two forms of retrieval, and if retrieval processes exist in animal learning (e.g., Spear, 1971, 1984), what kinds of models of retrieval processes in animal learning correspond to models of recall and recognition in humans?

We could go on in this way for a long time, asking question after question about the domains of existing theories and models in both (human) memory and in (animal) learning, without much hope for satisfying answers. But this sample makes the point: Existence of explanations does not imply that we always know the boundaries of their domains. If so, establishment of such boundaries would constitute an important research objective. Systematic classification of varieties of learning and memory would contribute to the attainment of that objective: The paramount purpose of classification is to describe the structure of relatedness of objects of interest, a structure within which general statements can be made (Sokal, 1974).

Thus, it could be argued that orthogonal to the three basic kinds of explanation mentioned earlier there exists, or there exists a need for, a fourth kind, that we could refer to as categorical explanation. An event, effect, phenomenon, or relation between situation and behavior is accounted for by identifying it as a member of a particular class or *category* of events, effects, phenomena or relations whose known properties include those requiring explanation in a given instance. For instance, questions such as, "Why is the slope of the forgetting curve flatter for primed fragment completion than for recognition?" and "Why can amnesics do as well as control patients in completing fragmented pictures or words although their recognition performance is very much poorer?" could be answered by identifying the two specific performances (fragment completion and recognition) as manifestations of different memory systems, one whose information is lost relatively rapidly and that is affected by the kind of brain damage that causes amnesia, and another that retains information for a much longer time and that is relatively unimpaired in amnesia (Tulving, Schacter, & Stark, 1982; Warrington & Weiskrantz, 1982). Categorical explanation presupposes the existence of (a) an empirically valid and theoretically sound classificatory system, and (b) theories and models that are known to hold for whole classes, specifiable parts of the system, rather than only for individual phenomena. Thus, classification would provide a basis for categorical explanation, not as a substitute for causal explanations and historical explanations of phenomena of learning and memory, but as a necessary precursor for their generality.

Why Classify?

Classification in all fields of human activity can serve either or both of two purposes. One is practical. Facilitation of communication, reduction of the memory load in handling knowledge, and optimization of storage and retrieval of information provide ready examples of such a pragmatic function of classification.

The second purpose of classification is theoretical. As discussed earlier, it is of much greater interest in our field. A classificatory system of learning and memory would guide functional research on learning and memory, it would help formulate theoretical objectives, aid in the evaluation of conceptual achievements, and provide a basis on which generalizations about learning and memory can be made.

Let us briefly mention a number of specific ways in which both the activity and the results of classification can be of value to the field. In summarizing these payoffs, I will assume a reasonably successful outcome of the classification research, that is, the existence of an acceptable and useful classificatory system, although it goes without saying that anything even resembling such a finished product is unlikely to be available in the foreseeable future.

1. The major contribution of a classificatory system would lie in the enhancement of *comprehensive understanding* of learning and memory. It would constitute one way in which theorists and students could perceive and describe the totality of learning and memory functions in a particular species in relation to the detailed structure and finer components of such a totality. It would facilitate the construction, evaluation, and description of functional models and theories of learning and memory. The classificatory system would represent one kind of manifestation of the overall regularity and order that students of learning have always been seeking.

2. A classificatory system would greatly facilitate selective delimitation of *conclusions* drawn from experiments and other kinds of empirical observation of particular varieties of learning and memory. More often than not, generalizations of empirical facts, and phenomena of learning and memory, to conditions other than those under which the observations were made—if such are explicitly made at all—are based on intuition and private theories.

3. In a closely related vein, the boundary conditions within which *theoretical* claims are assumed to hold, could be specified objectively and more precisely with the aid of a classificatory system. At the present time, neither the writers nor the readers necessarily know the range of situations in which a theoretical statement is expected to hold. In a developed science of learning and memory, we would expect to find statements of the following sort: "The model (or theory) described here applies to all varieties of learning (or memory) in (name of a specific category of learning or memory), and it may also hold, albeit less faithfully, for (names of some other classes). No claims are made that it applies to any other class." Today such statements do not exist in the literature, presumably partly because of the absence of any accepted classificatory system of learning and memory.

4. A classificatory system would provide economical ways of describing *deviations* from specified "norms." For instance, different kinds of amnesic

syndromes (Rozin, 1976; Moscovitch, 1984; Cohen, 1984) could be described in terms of the extent to which functioning of learning and memory of particular classes is impaired. Similarly, age-related changes in learning and memory capabilities in children and in older people may be describable succinctly and economically in terms of levels and branches in a hierarchical taxonomic tree.

5. Reasonably complete *descriptions of learning and memory abilities* of individual organisms or groups of organisms, free from unnecessary redundancies, might become possible by constructing tests and tasks that systematically sample—with any desired degree of thoroughness—the varieties of learning and memory organized by the classificatory system.

6. Existence of an accepted classificatory system would make it easier to construct a more rational *nomenclature* of varieties of learning and memory. In the absence of such a system, we see in the literature anomalies like "cued free recall" and "negative priming effects," expressions comparable to "a six-legged quadruped" and "negative expenditure of energy." While other life sciences devote a good deal of institutionalized and international effort to problems of terminology (e.g., Mayr, Linsley, & Usinger, 1953, *Appendix*; World Health Organization, 1974), students of learning and memory have retained the *laissez faire* attitude that served our science well in its formative years but which is unlikely to remain workable in the longer run. The creation of a useful, stable, and generally accepted terminology would be more likely if we had in existence an acceptable classificatory system.

These, then, are some of the reasons for undertaking the classification enterprise in learning and memory, stated in rather general and somewhat abstract terms. To illustrate the need for a useful classificatory system in more concrete terms, let us contemplate a simple exercise.

Intuitive Decisions About Relatedness

Suppose you are a specialist in learning and memory, and someone asks you the question, "How closely related are X and Y?" where X and Y refer to varieties of learning and memory. Can you answer the question? Always? Sometimes? If yes, on what basis do you do it, and on what sort of a scale do you provide the answer? If you cannot answer the question, why not?

Consider some possible concrete examples of X and Y in the question about the relatedness of X and Y. When in doubt, assume that in the situations described we are talking about performances following appropriate "training."

1. University students recalling words from a studied list versus recognizing words from the same list.

2. Rats turning left in a T-maze versus turning to the brighter side in the same maze.

3. A person recalling the plot of a movie and a bird picking up seeds from various hiding places in which it had cached them earlier.
4. A university student recalling a letter trigram in the Brown-Peterson task and a monkey making the correct choice in a delayed matching-to-sample task.
5. A person imagining the face of her beloved and a pigeon pecking the key with a plus sign on it.
6. A dog jumping over the barrier in a shuttle shock-avoidance task and a mouse staying on the elevated platform in a "step-down" compartment.
7. A professor of psychology recalling the formula of the multiple correlation coefficient and a student of psychology remembering how he failed the first term test in statistics.
8. An amnesic patient recalling the response word to the stimulus word in a paired-associate task with associatively related pairs and a young woman showing perfect adaptation to the "distorted" visual world while wearing Ivo Kohler's prism-lenses.

The list contains descriptions of relatively simple and straightforward situations that are either well known to everybody from real life or represent standard experimental paradigms from learning and memory laboratories. The exercise could be expanded in scope by considering all possible 120 pairings of the 16 memory and learning performances, or a very much larger number of pairings of a much larger number of examples of what humans and animals can learn and remember. And it would not be difficult to add excitement to the game by including descriptions of situations with greater detail. One of the situations might be described as follows: "A girl of 12 years of age, with an M.A. of 15, trying to recall, to extralist semantically related cues, three-word semianomalous sentences presented auditorily at the rate of 5 sec per sentence, on two successive study trials 24 hours earlier, with 24 sentences in the list, under the conditions where her orienting task was to judge the meaningfulness of the sentences on the first trial and the plausibility of the described action on the second." The list of such concoctions, enough of which can be found in the literature, could be made arbitrarily long, and it could include many comparably complicated scenarios from real life outside the laboratory.

It is possible to imagine the existence of a world in which (a) each of the varieties of learning and memory described in our exercise can be reliably *identified* as a member of a particular class or category, (b) each such class ("species") has a name accepted as such by everyone who cares about these things, and (c) the classes of learning and memory are organized into a comprehensive classificatory system. In such a world, the problem posed in the exercise could be solved by a simple look-up procedure, comparable to that used by systematic biologists faced with questions about relatedness of plants or animals.

In the world in which we live now, just about the only thing we can do when confronted with problems of relatedness of different varieties of learning and memory is to rely on our intuition. Intuitively, guided by the wisdom distilled from watching ourselves and other people, and by whatever relevant knowledge is found in the memory and learning literature, we do make judgments concerning relatedness of tasks. Many examples could be given, but let us limit ourselves to a few.

Nicholas Mackintosh, in his chapter in this volume, suggested that classical conditioning in animals is like episodic cued recall in humans, and David Olton (1984) has argued that rats have episodic memory of the kind that allows people to remember past events, since rats can respond correctly on delayed conditional matching tasks, which, like episodic remembering, require retrieval of particular past happenings. Fergus Craik (1983) has interpreted the differential impairment in memory capabilities of the aged in terms of the extent to which the performance of tasks require "self-initiated constructive operations" or "activation of conscious operations," properties of tasks that are reflected in their listing in an ordered array. Tasks that are closer to one another on Craik's scale could be regarded as being more closely related to one another than those farther apart. For instance, free recall is more closely related to cued recall than it is to recognition. In a similar vein, the task analysis described by Rönnberg and Ohlsson in this volume can be seen as affording one basis for making judgments about relatedness of varieties of learning and memory. Then there are students of learning and memory (e.g., Baddeley, 1984; Cohen, 1984; Kinsbourne & Wood, 1975; Mishkin, Malamut, & Bachevalier, 1984; Squire, 1982; Tulving, 1983, 1985a, 1985b; Warrington & Weiskrantz, 1982; Weiskrantz, 1982), who believe in the existence of different memory systems, and who are willing to make judgments concerning relatedness of varieties of learning and memory in terms of the extent to which different systems mediate performance of different tasks.

Finally, and perhaps most relevant in the context of the Umeå conference, the many generations of students of animal learning who have pursued their research in the hope that their findings and ideas will throw light on the basic principles of human learning (Jenkins, 1979), must have had in mind assumptions regarding correspondences between particular varieties of animal and human learning and memory. In the same vein, the search for "connecting models," the topic of Estes's paper in the present volume, must be predicated not just on the assumption that common mechanisms exist in (animal) learning and (human) memory, but also on the knowledge, hypothetical or real, concerning correspondences between particular kinds of learning in different species. In the absence of such knowledge, or at least ideas, about these correspondences, the enterprise of looking for parallels between learning and memory in animals and humans could not possibly succeed.

Thus we can say that, despite the strong traditional pretheoretical belief in the unity of learning and memory (Tulving, 1984), students of learning and memory have always been willing, and sometimes obliged, to assume differences in ways in which organisms benefit from their experiences, and to make judgments about relatedness of varieties of learning and memory, both within and between species. In this (limited) sense, they have acted as taxonomists. Their judgments, however, have been based on intuition only. There has been little or no discussion of objective rules for determining relatedness, and no systematic attempts to gather evidence that would either support the intuitive conjectures or correct them.

The pursuit of an explicitly articulated classification problem would supplement pure intuition with objective empirical evidence, explicitly gathered for the purpose, and evaluated and interpreted with the aid of accepted rules of procedure. It would make explicit, general, and systematic the practices that until now have been largely implicit, local, and haphazardly followed. Let us, in the next section of the paper, ponder on the nature of relevant evidence and on the rules of procedure.

OBJECTIVES OF CLASSIFICATION RESEARCH: EVIDENCE AND METHODS

Systematics of Learning and Memory

The distant goal of the classification enterprise in learning and memory would be development of what we might refer to, borrowing the term from biology, as the *systematics* of learning and memory. Paraphrasing Simpson (1961, p. 7), we could define the systematics of learning and memory as the scientific study of the kinds, diversities, and varieties of learning and memory, and natural relations among them. We could further think of systematics as comprising three separate but interrelated specialities: classification, taxonomy, and nomenclature (cf. Savory, 1970).

The subject matter of *classification* would be kinds, diversities, and varieties of learning and memory and its objective would be the ordering of varieties of learning and memory into "natural" groups. The *taxonomy*, or taxometrics (Meehl & Golden, 1982), of learning and memory would comprise the study of the rules, principles, and bases of classifying. *Nomenclature* would be concerned with the development and application of distinctive names and labels for the resulting groups or classes of varieties of learning and memory that would be generally accepted, thereby contributing to the quality of communication between different disciplines concerned with learning and memory.

The history of problems of classification and taxonomy in other fields of science teaches us that the classification problem in learning and memory is going to be complex, difficult, and frustrating. It is also going to be controversial. The whole idea will be rejected out of hand by some, and those who accept the general objective as worthwhile will not always agree on methods, facts, or interpretations. These kinds of problems, however, are usual in any new venture, in science as well as in other spheres of human activity, and they will be overcome in time.

Lessons from Biology

There exists a rich technical literature on classification and its application in many spheres of human activity. Research on principles and procedures of classification crosses traditional subject-matter boundaries; it has benefitted from the fruits of labors of thinkers in a number of disciplines (e.g., Dahlberg, 1982; Engelien, 1971; Felsenstein, 1983; Korner, 1976; Meehl & Golden, 1982; Simpson, 1961; Sneath & Sokal, 1973; Sokal, 1974). Although the application of principles of classification creates unique problems in different disciplines, the nature of the problems, as well as that of solutions, is characterized by a good deal of generality. Thus, as we start contemplating the possibilities and problems of classification in our own field, we can benefit from the experiences accumulated by practitioners in others.

Most relevant for our own purposes are probably the lessons that biologists have learned in their attempts to come to grips with the tremendous variety of living organisms that inhabit the earth (e.g., Huxley, 1940; Mayr, 1982; Simpson, 1961; Solbrig, 1966). Systematic biology, of course, differs from systematic learning and memory in many ways. One of the more conspicuous differences is the fact that there is little difficulty in determining the to-be-classified things in biology: They are concrete things that occupy space, that have boundaries, and that have perceptually identifiable and objectively measurable characteristics. In learning and memory, it is not necessarily clear what the objects of classification are or what they should be. We return to this problem shortly. Another difference pertains to the guidance that phylogenetic knowledge has provided to taxonomists in biology: Since taxonomists in learning and memory will not have such knowledge available, they would have to adopt some other criteria that will help them to construct a natural classificatory scheme. We discuss one such criterion presently.

Although all classificatory systems are based on the idea that like objects are grouped into like classes, and that like classes are further grouped together as superordinate classes, we should probably be sceptical about the possibility of reaching our goal through what have been referred to as numerical phenetics (Mayr, 1982) or numerical taxonomy (Sneath & Sokal, 1973). These terms cover a number of procedures in which the perceptual similarity of objects to be classified is judged by human observers, either experts or others, and the results of such judgments subjected to any one of a large number of statistical techniques that convert the data in similarity or proximity matrices into structured wholes (e.g., Sneath & Sokal, 1973; Felsenstein, 1983).

We would probably not want to follow this route. The number of different classificatory systems that could be produced would be large, and there would be no reasons to believe that any of them would agree with nature. The outcomes of such exercises in some other fields have been disappointing. Thus, for instance, Meehl (1979) has said that in psychopathology, phenetic classification procedures (procedures based on phenotypical descriptions of target entities) coupled with classical techniques of cluster analysis, have *never* led to the discovery of a single taxon (taxonomic unit), and Mayr admits to knowing of not "a single substantial contribution made by numerical phenetics to the classification of any mature groups, or to classification at the level of orders, classes, or phyla" (Mayr, 1982, p. 225). Given such evaluations we would be wise not to rush to embrace these methods, despite their great popularity. Classifying varieties of learning and memory on the basis of their identifiable properties is unlikely to be more successful in our field.

We have always had, and have now, various arbitrarily defined and intuitively constituted categories of learning and memory, categories such as instrumental and classical conditioning, associative and nonassociative memory, and short-term and long-term memory, among others. (See also Gagne, 1977; Melton, 1964; Tolman, 1949). These may be useful for practical purposes, but there is no evidence that they conform to nature. Classification simply on the basis of observed characteristics of varieties of learning and memory is unlikely to be successful.

Neuropsychological Criterion

If we do not classify on the basis of observable properties of varieties of learning and memory, what do we do? One thing we can do is to take a greater interest in what is known and what will be learned in the future about the neural basis of learning and memory, and pursue the classification problem in terms of known relations between brain activity on the one hand and psychological (behavioral and cognitive) manifestations of learning and memory on the other hand (e.g., Kinsbourne, 1976; Oakley, 1981; Rozin, 1976; Squire, 1982; Weiskrantz, 1968). Thus, instead of requiring that a classificatory scheme be consistent with known facts about the phylogeny of the classified organisms, as is done in systematic biology, we could require that our classification be consistent with what is known about brain activity in learning and memory. Moreover, we could adopt the extent to which two varieties of learning and memory are subserved by the same neural mechanisms or brain processes as one of our criteria for making judgments about their relatedness. We can refer to this rule for making judgments concerning relatedness as the neuropsychological criterion.

Although relatively little is known yet about brain activity in learning and memory, additional knowledge is being gained at an increasingly rapid rate. Even now a great deal can be done with the neuropsychological criterion, if we are willing to adopt it. Changes in brain activity can be effected through lesions or stimulation, as discussed by David Olton in his chapter in this volume, and techniques for observing brain activity *in vivo* are being invented and improved all the time.

The neuropsychological criterion can be also used in situations where changes or differences in brain function accompany, or result from, evolutionary changes of the species (e.g., Bitterman, 1975) or developmental changes of individuals. For instance, it is well known that older organisms can learn things that younger ones cannot (e.g., Amsel & Stanton, 1980; Bachevalier & Mishkin, 1983) and that older ones cannot do what younger ones can (e.g., Craik, 1977; Craik & Rabinowitz, 1983). Varieties of learning and memory thus could be classified on the basis of their presence in different species and the time of their appearance in development (e.g., Schacter & Moscovitch, 1984).

Especially useful for classificatory purposes, at least during the initial stages of the enterprise, are observed dissociations of learning and memory performances. A dissociation is said to have been demonstrated if a treatment that changes brain activity affects learning or memory performance of Type X, but not of Type Y. Dissociation means that a particular structure, or a neural mechanism, constitutes one of the *necessary* neurophysiological conditions for the carrying out of a psychological function. Thus, the observation of the dissociation between performances X and Y, as a result of treatment A, permits the conclusion that the brain mechanism affected by A is necessary for learning or memory manifested in X, and that it is not necessary for Y, and that to this extent X and Y represent different *classes* of learning or memory. Treatments most commonly used in the establishment of dissociations are brain lesions, experimentally produced in animals, and resulting from disease or accidental injury, or from therapeutic surgical procedures, in humans.

Even more compelling are demonstrations of double dissociation: Treatment A affects X but not Y, whereas another treatment B affects Y but not X (Jones, 1983; Shallice, 1979; Weiskrantz, 1968). In this case, the possibility is excluded that the single dissociation reflects simply differential vulnerability of performances X and Y, and the conclusion that X and Y are mediated by different neural systems is much stronger. We will see later in this paper, however, that sometimes it is the *absence* of a double dissociation, in a situation in which a single dissociation has been established, that can provide strong evidence useful for classification.

The general principle in all these criterion procedures is the same: Behavioral expressions of different varieties of learning and memory show systematic correlations with variations in brain activity reflecting species differences, developmental stages, lesions, electrical stimulation, or changes in biochemistry. These correlations constitute one basis on which a "natural" classificatory system of learning and memory can be constructed.

Tasks as Objects of Classification?

In many ways the most difficult part of the classification problem in learning and memory lies in getting started. Before we can begin worrying about the degree of relatedness of varieties X and Y of learning or memory, we must make some preliminary decisions as to what constitutes X and what constitutes Y. Exactly what are the things that we want to compare in order to assign them to categories that then define their "identity"? What are the *objects* of our classification enterprise?

So far in the paper I have used the somewhat fuzzy expressions of "varieties" or "kinds" of learning and memory, or even "performances" and "situations," when referring to the things that we wish to classify, but what exactly do we mean by these terms? How do we delineate a particular variety, kind, performance, or situation of learning or memory? What kind of intellectual operations are necessary to create a situation where we can describe a "variety" of learning and memory, assert that it is the "object" of the classification enterprise, and pose questions about its relatedness to other "objects"?

We could take "phenomena" of learning and memory, or "effects," or "relations between behaviors and situations," or anything else, as objects of classification, and we would always have to face similar kinds of problems: How to delineate them, how to decompose behavior into the chosen units, how to define and describe them. There are no obvious solutions to these problems; they constitute a part of the larger classification problem.

To illustrate the difficulties, uncertainties, and possible frustrations inherent in the selection of *units* of learning and memory, let us consider an obvious candidate for such a unit, namely a learning or memory *task*. We can provisionally define a "task" as a situation in which a particular organism interacts with its environment in a particular way, and, as a consequence, acquires the capacity to behave differently in similar situations in the future. The concept of *task* is widely used in psychological research on learning and memory: The literature is full of descriptions of, and statements about what happens in, all kinds of tasks. Tasks are the most frequently occurring products of decomposition procedures that psychologists have used in extracting units from the stream of behavior and thought. But if we adopted tasks as objects of classification, we would immediately run into problems.

What Are Tasks?

One of the first serious difficulties with tasks as objects of classification lies in the absence of a workable definition of *task*, a definition that would enable its users to decide whether two situations represent the same task or different tasks. Although the concept of *task* has been used widely by all learning and memory theorists and experimenters, few attempts have been made to define it explicitly in relatively abstract terms, in a manner that would not create disagreements with the extant use of the term in different situations. Nor have there been many attempts to specify the rules by which decisions regarding identity and nonidentity of tasks, in different situations, could be made.

One approach to the problem of defining "task" would be through specifying its constituents or properties. Most practitioners would probably agree that such specifications should include statements about (a) the kind of organism that does the learning, (b) the conditions under which learning occurs, (c) the behavior in which the organism engages at the time of learning, and (d) the behavior through which what has been learned is expressed. Each of these aspects of a task can be, and usually is, quite complex. The last three aspects correspond closely to Rescorla's three issues with which theories of learning must be concerned (Rescorla, 1980; also his chapter in this volume). The learning organism may have to be included as a part of the definition of "task" since we cannot always assume that what appear to be similar kinds of learning in different species, or even within a species, are based on the same underlying mechanisms, a point made by Estes in his chapter in this volume (see also Estes, 1978).

But thinking of tasks in this way creates a number of questions: (a) Is the learning organism's motivational, emotional, or pharmacological "state" in any way relevant to the specification of a task? (b) Does the task change as a function of the learner's expertise or relevant prior knowledge? (c) Do different strategies that an organism may bring to bear upon its learning activities change the task? (d) Is the behavior through which the effects of learning are expressed an *obligatory* part of the definition of a task? (e) Is a task a *single* task even if we know that people's performance of it is based on different processes that can be separately identified? (f) How does the fact that identical responses to identical stimuli may reflect quite different underlying processes affect our definition of tasks?

Consider possible answers to some of these questions with the help of two concrete examples of tasks, one that can be used only with humans, and one that can be used with both humans and animals. The first is single-trial free recall: A subject is exposed to a collection of discrete items and he is instructed to recall the names of as many of the presented items as he can (e.g., Waugh, 1961). The second task is the single-object version of delayed matching to sample: A subject is exposed to an object, and after an interval of time has to choose the same object from a test set that consists of the target object and one not encountered on that trial. This task figures prominently in the research that David Gaffan discusses in his chapter in the present volume. (See also Roberts & Grant, 1976.)

What aspects of the free-recall task can we alter before we have to concede that we have changed the *task*? Does it change when we alter the length of the list, or the retention interval, or the activity of the rememberer interpolated between the study and the recall of the list? Does it change when the rememberer changes—say, from a university student to a very young child, or to an Alzheimer patient? Does it change when the to-be-remembered items change, say, from familiar words to 8-letter nonsense strings, or to 16-word sentences, or to photographs of complex scenes? Does the task become another kind of task when the experimenter instructs the subjects to record recalled words in the order in which the items were presented (e.g., Cohen, 1970), or to recall the last few items from the list first (e.g., Craik & Watkins, 1973, Exp. 2), or to recall the items by their initial letters (e.g., Murdock, 1960), or in alphabetical order (e.g., Tulving, 1962)? Is the *task* changed when subjects use the serial or alphabetical recall strategy on their own, without instructions from the experimenter? Is recalling a once-presented list of 16 familiar words 10 seconds after study the same task as recalling the names of one's schoolmates 40 years later (e.g., Bahrick, Bahrick, & Wittlinger, 1975)? What about faces of the same schoolmates? What about recalling names when faces are given as cues? Does the task become a different task when subjects, at the time of recall, are given semantically related words as cues or prompts for the recall of words from a studied list (e.g., Bahrick, 1969)?

Let us turn our attention next to the behavior through which learning is expressed. Does the single-trial free recall task change when we ask the person to (a) reproduce the list items orally, (b) write the names down, (c) recall the words in another language with which the learner is familiar, or (d) tap the names out in the Morse code? Does the negative answer, which most people would give, imply that the overt behavior through which learning is expressed is not a defining feature of *any* task? For instance, if we think of the eyelid conditioning situation as a "task," would we be willing to concede that the exact form of overt behavior in which the learning manifests itself is immaterial? If not, does it mean that sometimes the expressive behavior is, and sometimes it is not, to be regarded as part of the definition of tasks? And if so, what determines when it is and when it is not?

If the reader is gradually becoming, or already has become, sceptical of the possibility of ever finding satisfactory answers to these questions, and hence is starting to have doubts about the usefulness of thinking of tasks as objects of classification, he or she is on the right track: The point of the exercise is to demonstrate that we have difficulty with questions that an outsider might think we answered a long time ago.

It is not only the difficulty of answering these and similar questions that should concern us. We should also be concerned with the rules or criteria that we use when we try to or do come up with answers to the questions. On what basis do we arrive at the answers? Is it possible to articulate such rules and criteria? Can we expect them to be generally acceptable?

Behind these and many other specific questions posed in this paper lurks perhaps the most telling one: How is it possible to ask all these questions that seem quite reasonable and appropriate, and yet have no place in our 100-year-old literature to turn to for the answers, or even for an enlightened discussion of the issues? What kind of perspective of the subject matter has been missing in the study of learning and memory ever since Ebbinghaus, Pavlov, Thorndike, Bryan and Harter, and other early pioneers started it all (Cofer, 1979; Hearst, 1979; Jenkins, 1979)?

Let me ask just two questions about our other task, delayed matching to sample, before we take leave of the problem of what is a task. First, does the task remain the "same" for organisms that have no language and for those that do? If yes, does it mean that possession of language does not affect the identity of *any* task? If not, when does it matter? Second, is the task the same task regardless of the nature of test objects that are used? Consider the difference between two versions of the task. One employs "trial-unique" objects, in the other the objects recur trial after trial. In the "trial-unique" procedure (e.g., Mishkin, 1978), every target object presented for inspection, and every distractor object at test, is always "new": The subject has never encountered them before. In the extreme form of the "repeated-objects" procedure there are only two objects used on successive trials of the experiment, A and B. One is selected randomly on each trial to serve as the inspection stimulus, and both are presented at test. Now, should we think of these two versions of the basic procedure as one and the same task, or do they constitute different tasks? The "same" judgment would obviate the need for any further questions concerning the relatedness of the two versions; the "different" decision would render such questions both important and meaningful.

Tasks As Categories

Although I want to emphasize problems and questions in this paper, rather than declarative ideas and suggestions, it may be useful at this juncture to try to reduce the perplexity that the long string of questions may have created. One way to do so would be to introduce the distinction between individual tasks and classes of tasks. Many a question I have posed here could have been clarified, and answers rendered less uncertain, by the specification that the question referred to a *class* of tasks and not to an individual member of that class.

Biology again provides a useful analogy here, in the form of the concept of "species." Species are taxonomic units at the lowest level of the classificatory hierarchy, being formed into larger categories at higher levels (genera, families, orders, and so on). An important property of species is that it represents a natural *population* of plants or animals. Individual members of any population may vary greatly in many of their characteristics. Only some characters are critical for the determination of whether or not an individual belongs to a species, characters such as ability to interbreed,or chromosomal structure.

When we discuss tasks, we also must distinguish between a single task as such and the population to which it belongs. Like individual members of a species, so individual tasks within a particular population need not at all be identical but may vary in a number of characters. Classes of tasks would be defined in terms of certain critical characters that its members share. The classification problem thus reduces, during the early stages of the game, to the problem of discriminating defining characters from nondefining characters of task categories.

Overt Behavior and Covert Processes

Let us return to the implications of subjects' use of different strategies for the definition of tasks, because it, and some related observations, will lead us to one of the fundamental problems in the classification enterprise. Suppose, for instance, that in a task such as free recall one subject uses a rote learning strategy, implicitly rehearsing individual items by muttering their names over and over during the presentation, whereas another person constructs interacting mental images of the named objects. Would the tasks performed by the two individuals belong to the same "species"?

The general question here has to do with the contrast between observable behavior and underlying processes as characters potentially relevant to classification. Should we go strictly by what the specified requirements are for the learner—by the observable, overt behavior—or should we take into account the covert mental activity in which the learner engages in performing the task?

Consider a slightly different version of the same question. It concerns situations where we know, or have good reasons to believe, that the performance on the task represents a combined effect of separate and experimentally differentiable mechanisms. For instance, we know now that in the single trial free recall task (e.g., Craik & Levy, 1976; Glanzer, 1972; Tulving & Colotla, 1970; Watkins, 1974) primary and secondary memory processes, or systems, contribute to the overall performance on the task that is measured in terms of the number of items recalled (cf. Crowder, 1982). The two components are differentiable in the sense that many independent variables that affect one component have no effect on the other, and vice versa. Given this knowledge, should we still treat the overall procedure as a single "task"?

A third related observation is derived from the common knowledge that one and the same response that a learner makes to one and the same stimulus may represent the effects of rather different underlying processes. For example, if in a paired associate learning task, a subject learns a pair such as *army*-SOLDIER, and later is tested with the stimulus item *army*, the learner can make the correct response SOLDIER either in terms of his remembering seeing the two words in the study list (relying on his episodic memory, as some of us might say) or simply in terms of his general knowledge of the associative structure of words, using a "free-association strategy" (or relying on his semantic memory). How does the fact that there is no unique one-to-one correspondence between underlying mechanisms and processes on the one hand and overt responses on the other affect our definitions and descriptions of tasks? Should it matter in one way or another, or should we ignore it altogether?

All three examples illustrate the almost universally accepted assumption—supported by some factual evidence—that every learning and memory task is composed of a number of components (informational units, relations among them, processes, operations, mechanisms, or whatever), and that different tasks may have both common and non-common components. This assumption creates a fundamental problem: Should we take whatever information we have about composition of tasks into account *before* we begin classifying them, or should we expect that these components will be revealed by the classificatory scheme with which we end up after we have successfully completed the mission? What would a search for *components* of learning and memory mean and entail?

Components of Learning and Memory

Concepts such as "learning" and "memory" do not represent entities that have some sort of existence in the brains or minds of organisms. They are simply broad labels assigned to describe concatenations of neural, behavioral, and mental components whose various combinations serve the function of shaping an organism's knowledge and behavior through its interactions with the world, thereby helping it to adjust and survive.

If varieties of learning and memory represent various concatenations of more elementary components, why should we not try to classify the varieties in terms of these constituent elements? The reason is simple: We do not know what the components are. We do not as yet have even a short list of "basic" processes of learning and memory that can be reliably identified and isolated across different situations. Part of the difficulty is undoubtedly attributable to the intrinsic complexity of the subject matter of our science. Another part of it may lie in the ineffective methods we use, or infertile pretheoretical and theoretical ideas that we have inherited and uncritically accepted from the past, the "curse of an angry god" that Bolles talks about in his chapter in the present volume. Many theorists (e.g., Craik, 1983, also this volume; Nilsson, 1984), indeed, might want to argue that any search for *components* of learning and memory, regardless of how they are conceptualized, is doomed to failure at the outset for the simple reason that there are *no* components of varieties of learning and memory: Learning and memory are manifestations of *interactions* between task demands and the environment, shaped by the capabilities of organisms.

I think that we should not reject ideas before we test them, or at least think hard about them. Although we may not have many compelling facts about components of learning and memory, it is possible to entertain thoughts about how they might be detected should they exist. Consider a simple example.

It is possible to identify several *logical* components of the delayed matching-to-sample task: (a) inspection of the presentation object, (b) construction of a neural (mental) representation of that object, (c) storage of the representation, (d) retention of the representation over time, (e) observation of the test objects, (f) selection of the test object that matches the representation of the target object, and (g) the making of the appropriate response to the selected test object. This description sketches the sequence of events that have to occur if the animal is to respond consistently and successfully in the task. Failure of any one of the seven processing stages will result in chance performance.

We do not know how these seven "logical" stages correspond to neural and behavioral processes that constitute the performance on the task, although we have good reasons to believe that the situation is in fact very much more complicated than the simple listing of the stages suggests. Any one of the labeled processing stages is probably mediated by a complex pattern of neural events: The apparent unity and simplicity of the operation lies in our description rather than in the underlying processes. The situation is further complicated by the fact that the processes of successive stages must run their course under the general guidance of operating procedures that define the "rules of the game" for the subject of the experiment. These are the rules that Harlow (1959) studied under the rubric of "learning sets"; they may be involved in the establishment of what others have referred to as "reference memory" (Honig, 1978; Olton, Becker, & Handelmann, 1979).

The organism's required knowledge of the "rules of the game" vary with the version of the task. With "trial-unique" objects there is no need for the organism to keep track of the temporal date of the occurrence of the inspection object, whereas in the "repeated-objects" version such temporal dating is indispensable. In the former, general "familiarity" of one of the test objects is sufficient for successful performance, whereas in the latter the "rules of the game" must include a concept like the "trial": test objects must be matched with the inspection object seen on a given trial, not on any trial. Ability to keep a specifically dated record of the appearance of objects on particular trials is mandatory for success with repeated stimuli. In the absence of such a record, performance cannot exceed chance level. According to this analysis, then, both versions of the task have a number of components in common, and they differ in that the repeated-objects version has at least one *additional* component, one that makes it possible for the organism to keep track of the temporal date of occurrence of objects.

We can schematically express the situation as follows. Both versions of the task consist of or require a number of common components, *a, b,* and *c*, and the repeated-objects version consists of or requires an additional component *d*. Given such a situation, it should be possible to find single dissociations of the kind where, say, a particular brain lesion results in the organism's inability to perform the repeated-objects version of the delayed matching-to-sample task without affecting the performance on the version of the task involving trial-unique objects. More important is another expectation derived from our logical analysis: Under the conditions as specified, double dissociation of performance on the two versions of the task would be *precluded*. Given that the components of Task Y entail those of Task X, there is no way of interfering with the operation of X without interfering with Y. The presence of one, and the absence of the other, type of dissociation would thus provide strong evidence for the existence of a particular component of learning or memory, as well as for the necessary involvement of the component in particular tasks.

The same logic applies to the distinction made by Olton and his associates (e.g., Olton, 1983; Olton et al., 1979) between reference and working memory as dissociable systems: Reference memory is required for all the tasks that the Olton group has used in their experiments whereas working memory is required only for some. Such a relation renders possible single dissociations, as demonstrated by the Olton group, and it rules out the possibility of double dissociations. The logic also has definite implications for schemes such as I have described elsewhere (Tulving, 1985a, 1985b) in which three memory systems—procedural, semantic, and episodic—constitute a class-inclusion "monohierarchy."

To what extent *logical* analyses of this kind agree with nature is unknown. Brains and minds of animals and people are complicated, and we should not be surprised to find mismatches between our simple ideas and the more complex world of facts. But such analyses do pinpoint possibilities for finding especially useful kinds of evidence. Mere demonstrations of single or double dissociations of two tasks that vary with respect to several (unknown) components, as a rule, do not tell us very much about anything and do not help us much with the classification problem. However, systematic correlations between brain function and behaviors in carefully analyzed tasks could play a decisive role in the mission of identifying components of learning and memory.

CLASSIFICATION AND EXISTING SYSTEMS

Although the classification problem in learning and memory, and the new perspective on research that it represents, is not yet a part of the scientific agenda of workers in learning and memory, we do already have various speculative proposals regarding different kinds, or systems, of learning and memory. How is the classification problem that I have discussed in this chapter related to current research and notions with respect to these systems? I briefly discuss this issue in relation to a triadic scheme that I have described more fully elsewhere, as well as to other similar schemes.

Triadic Memory

In discussing the relation between episodic and semantic memory not too long ago (Tulving, 1983), I assumed that episodic and semantic memory represent parallel subsystems of propositional (declarative) memory. Because of certain difficulties inherent in such a conceptualization, pointed out by a number of critics (e.g., Kihlstrom, 1984; Lachman & Naus, 1984; Lieury, 1979; Seamon, 1984; McCauley, 1984; Tiberghien, 1984) I have now revised these ideas and have suggested that it may be more appropriate to think of episodic memory as "growing out of but remaining embedded in" the semantic system (Tulving, 1984). The same general idea can also be extended to the relation between procedural memory and propositional memory: rather than thinking of these two as parallel subsystems of the overall memory system, it seems to make more sense to conceptualize semantic memory as subsumed by the procedural memory system (Tulving, 1985a, 1985b).

A learning and memory system refers to a particular set of neural structures, or mechanisms, or both, that subserve different behavioral and cognitive functions, that operate according to different laws and principles, and that have evolved at different times in the phylogeny of the species and evolve at different times in the ontogeny of individuals (Tulving, 1984). The overall arrangement of procedural, semantic, and episodic memory systems could be characterized as "monohierarchical": a hierarchical arrangement of varieties of learning and memory in which pocedural memory contains semantic memory, and semantic memory contains episodic memory, as a specialized subsystem. The arrangement conforms to the principle, found useful in the study of phylogeny of animals (e.g., Carter, 1961; deBeer, 1938), according to which a structure or structural feature that has developed later always must be such that it can be derived by modification of the corresponding earlier structure or structural feature. Thus we assume that in the evolutionary emergence of learning and memory systems, too, systems that have developed later represent extensions and modifications of earlier systems rather than independent, parallel developments of completely different systems.

Other tripartite divisions of memory in animals have been proposed by Ruggiero and Flagg (1976) and Oakley (1981). The category that represents the "simplest" memory in their schemes is essentially identical, or at least homologous, with the procedural system in the triadic classification I have just described, and the other two classes in Oakley's (1981) scheme, "representational" and "abstract" memories, can be regarded as analogous to semantic and episodic memory, respectively.

A number of other theorists have proposed various dichotomous divisions of learning and memory. These schemes, by Cohen and Squire (1980), Mishkin, Malamut, and Bachevalier (1984), O'Keefe and Nadel (1978), and Olton (1983), among others, fit into the triadic or tripartite schemes without undue difficulties.

"Upward" and "Downward" Approaches

There are two major differences between the classification problem as I have discussed it in this chapter and the extant speculations concerning learning and memory systems. First, attempts to subdivide the total subject matter of learning and memory into different systems, exemplified in the schemes just discussed, are reminiscent of the Aristotelian and Linnaean methods of classification of organisms on the basis of the "downward" procedure (Mayr, 1982). One begins with the total population of things to be classified, subdivides them into two classes on the basis of a criterion that defines the "essence" of classes, and then proceeds to repeat the operation at lower levels. Thus, for instance, the distinction between procedural and nonprocedural memory could be based on the presence or absence of symbolic "content" of learning. Similarly, the further division of the nonprocedural memory into episodic and semantic could be based on the presence or absence of personal reference of the symbolic content that is learned and remembered. The fact that sometimes correlations exist among these criterial features (e.g., Tulving, 1983, Table 3.1) does not change the nature of the basic approach.

The classificatory system that I am envisaging as emerging from a systematic application of the principles of taxonomy to varieties of learning and memory, on the other hand, would represent the "upward" procedure (Mayr, 1982), characteristic of the post-Darwinian approach to classification in biology. Here, one begins with the smallest units of the classification, and builds them into larger groupings on the basis of a large number of criterial characters, working from the bottom up.

The second major difference between existing schemes and the one envisaged here is one that we have already discussed: the currently existing large dichotomous and trichotomous categories are mostly products of individual scientists' intuition, albeit tempered with some empirical evidence. The approach is largely intuitive because of the absence of agreed-upon, or even reasonably widely discussed, principles and rules on which the divisions are based. There are only general rules of thumb, such as those that are used in connection with functional dissociations or double dissociations. Many relevant questions—such as those concerning components of tasks, and their weights in different tasks—can be answered only in terms of implicit rules and hunches.

The classification problem envisaged here would be accompanied and guided by the development of useful explicit rules and procedures that are worked out according to generally accepted methods of science, and it would entail the application of these rules and principles to making decisions about the relatedness of tasks, or whatever units are to be classified.

The fact that the existing classificatory schemes are "downward" and based on intuition does not necessarily mean that they are wrong, or that they will necessarily be replaced by different divisions at top levels of the hierarchical scheme that might emerge from the "upward" classification activities. It is quite possible that the present large categories, or systems, may map quite well onto the classificatory structure that will be constructed from the bottom up.

Nevertheless, an approach more systematic than the one that has been used to date is clearly called for. Even those of us who are willing to pontificate as to the nature and correspondence of components of different tasks and different memory systems, and who are willing to interpret evidence from experiments and other empirical observations in terms of our own hypothetical classificatory schemes, have very little to say about subsystems within the larger postulated systems. The classificatory structure that will be produced by the systematists in learning and memory will undoubtedly turn out to be very much richer in detail than are the existing schemes.

CONCLUSIONS

I have proposed in this chapter that students of learning and memory should devote a part of their experimental and theoretical efforts to the construction of a natural classificatory system of varieties of learning and memory. In elaborating this proposal, I pointed to gaps in our knowledge, gaps revealed by many difficult questions. Many questions I posed are of the kind that would have to be dealt with and answered in the course of the classification enterprise. The central suggestion was to adopt the neuropsychological criterion as an important basis for a natural system of classification, and as a potential deterrent to the proliferation of many arbitrary classificatory schemes.

At the beginning of the chapter I said that work on the classification problem constitutes a kind of a collaborative venture of students of (animal) learning and (human) memory that might help to build bridges between them. But the worthwhile nature of the enterprise should be clear even if we did not have to worry about the present separation of the two fields. If the systematics of learning developed independently of the systematics of memory, and if instead of collaboration we were to find ourselves engaged in a friendly but vigorous competition as to who "gets there first," nothing much would be lost.

Indeed, we must not overlook the possibility that, despite our nobler motives and aspirations, progress in classification might turn out to be easier to achieve if we initially restricted the domains of our activities by species of learners, and by broad categories of learning and memory. Many subproblems of the classification problem may become more tractable if we thought of them in connection with a particular population of learners, and a particular class of learning.

On the day when I completed the final version of this chapter, a new book entitled *Taxonomies of Human Performance* by Fleishman and Quaintance (1984) came to my attention. Although Fleishman and Quaintance are primarily concerned with taxonomic problems in those parts of psychology that are of particular interest to specalists in engineering psychology and research on human factors, their book should be required reading for all aspiring taxonomists in learning and memory. Chapter 2, for instance, provides an excellent summary of basic concepts of classification and their application in biological and psychological sciences, and elsewhere one finds a thorough discussion of the concept of *task*, the source of so much difficulty in this paper. I was pleased to see a good deal of overlap between Fleishman and Quaintance's ideas and those discussed here. It looks as if classification is coming into its own, creating new perspectives in psychological research, and that its pursuit in learning and memory is but another manifestation of the everpresent *Zeitgeist*.

ACKNOWLEDGMENT

The preparation of this paper has been supported by the Natural Sciences and Engineering Research Council of Canada Grant No. A8632, and by a Special Research Program Grant from the Connaught Fund, University of Toronto. I am grateful to Marcel Kinsbourne for the idea of the classification problem (Kinsbourne, in press), to Fergus Craik and David Olton for advice, and to Janine Law for collaboration in library research and the production of the manuscript.

REFERENCES

Amsel, A., & Stanton, M. (1980). Ontogeny and phylogeny of paradoxical reward effects. In J. S. Rosenblatt, R. A. Hinde, C. Beer, & M. Busnel (Eds.), *Advances in the study of behavior*. New York: Academic Press.

Bachevalier, J., & Mishkin, M. (1983). The development of memories vs. habits in infant monkeys. *International Journal of Psychophysiology, 1*, 116.

Baddeley, A. (1984). Neuropsychological evidence and the semantic/episodic distinction. *Behavioral and Brain Sciences, 7*, 238–239.

Bahrick, H. P. (1969) Measurement of memory by prompted recall. *Journal of Experimental Psychology, 79*, 213–239.

Bahrick, H. P., Bahrick, P. O., & Wittlinger, R. P. (1975). Fifty years of memory for names and faces: A cross-sectional approach. *Journal of Experimental Psychology: General, 104*, 54–75.

Bitterman, M. E. (1975). The comparative analysis of learning. *Science, 188*, 699–709.

Bruce, D. (1985). The how and why of ecological memory. *Journal of Experimental Psychology: General, 114*, 78–90.

Carter, G. S. (1961). *A general zoology of the invertebrates, Ch. 23*. London: Sidgwick & Jackson.

Cofer, C. N. (1979). Human learning and memory. In E. Hearst (Ed.), *The first century of experimental psychology*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Cohen, N. J. (1984). Preserved learning capacity in amnesia: Evidence for multiple memory systems. In L. R. Squire & N. Butters (Eds.), *Neuropsychology of memory*. New York: Guilford Press.

Cohen, R. L. (1970). Recency effects in long-term recall and recognition. *Journal of Verbal Learning and Verbal Behavior, 9*, 672–678.

Cohen, N. J., & Squire, L. R. (1980). Preserved learning and retention of pattern analyzing skill in amnesia: Dissociation of knowing how and knowing that. *Science, 210*, 207–209.

Craik, F. I. M. (1977). Age differences in human memory. In J. E. Birren & K. W. Schaie (Eds.), *Handbook of the psychology of aging*. New York: van Nostrand Reinhold.

Craik, F. I. M. (1983). On the transfer of information from temporary to permanent memory. In D. E. Broadbent (Ed.), *Philosophical Transactions of the Royal Society of London, B*. London: The Royal Society, *302*, 341–359.

Craik, F. I. M., & Levy, B. A. (1976). The concept of primary memory. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes. Vol. 4: Attention and memory*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Craik, F. I. M., & Rabinowitz, J. C. (1983). Age differences in the acquisition and use of verbal information. In H. Bouma & D. G. Bouwhuis (Eds.), *Attention and performance X*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Craik, F. I. M., & Watkins, M. J. (1973). The role of rehearsal in short-term memory. *Journal of Verbal Learning and Verbal Behavior, 12*, 599–607.

Crowder, R. G. (1982). The demise of short-term memory. *Acta Psychologica, 50*, 291–323.

Dahlberg, I. (1976). Classification theory, yesterday and today. *International Classification, 3*, 85–90.

Dahlberg, I. (1982). Editorial: Classification science—a true scientific discipline? *International Classification, 9*, 63.

deBeer, G. R. (1938). *Evolution: Essays on aspects of evolutionary biology*. Oxford: Clarendon Press.

Engelien, G. (1971). *Der Begriff der Klassifikation*. Hamburg: Helmut Buske Verlag.

Estes, W. K. (1978). On the organization ands core concepts of learning theory and cognitive psychology. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes. Vol. 6: Linguistic functions in cognitive theory*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Felsenstein, J. (Ed.). (1983). *Numerical taxonomy*. Berlin: Springer Verlag.

Fleishman, E. A., & Quaintance, M. K. (1984). *Taxonomies of human performance*. Orlando, FL: Academic Press.

Gagne, R. M. (1977). *The conditions of learning*. New York: Holt, Rinehart, & Winston.

Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review, 91*, 1–67.

Glanzer, M. (1972). Storage mechanisms in recall. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory*. New York: Academic Press.

Harlow, H. F. (1959). Learning set and error factor theory. In S. Koch (Ed.), *Psychology: A study of science* (Vol. 2). New York: McGraw-Hill.

Hearst, E. (1979). One hundred years: Themes and perspectives. In E. Hearst (Ed.), *The first century of experimental psychology*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Honig, W. K. (1978). Studies of working memory in the pigeon. In S. H. Hulse, H. Fowler, & W. K. Honig (Eds.), *Cognitive processes in animal behavior*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Huxley, J. S. (1940). *The new systematics*. Oxford: Claredon Press.

Jenkins, H. M. (1979). Animal learning and behavior theory. In E. Hearst (Ed.), *The first century of experimental psychology*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Jones, G. V. (1983). Note on double dissociation of function. *Neuropsychologia, 21*, 397–400.

Kihlstrom, J. F. (1984). A fact is a fact is a fact. *Behavioral and Brain Sciences. 7*, 243–244.

Kinsbourne, M. (1976). The neuropsychological analysis of cognitive deficit. In R. G. Grenell & S. Gabay (Eds.), *Biological foundations of psychiatry*. New York: Raven Press.

Kinsbourne, M. (in press). Systematizing cognitive psychology. *Behavioral and Brain Sciences*.

Kinsbourne, M., & Wood, F. (1975). Short-term memory processes and the amnesic syndrome. In D. Deutsch & J. A. Deutsch (Eds.), *Short-term memory*. New York: Academic Press.

Korner, S. (1976). Classification theory. *International Classification, 3*, 3–6.

Lachman, R., & Naus, M. J. (1984). The episodic/semantic continuum in an evolved machine. *Behavioral and Brain Sciences, 7*, 244–246.

Lieury, A. (1979). La memoire episodique est-elle emboitee dans la memoire semantique? *L'Annee Psychologique, 79*, 123–142.

Mayr, E. (1982). *The growth of biological thought*. Cambridge, MA: Belknap Press.

Mayr, E., Linsley, E. G., & Usinger, R. L. (1953). *Methods and principles of systematic zoology*. New York: McGraw-Hill.

McCauley, R. N. (1984). Inference and temporal coding in episodic memory. *Behavioral and Brain Sciences, 7*, 246–247.

Meehl, P. (1979). A funny thing happened to us on the way to the latent entities. *Journal of Personality Assessment, 43*, 563–581.

Meehl, P., & Golden, R. R. (1982). Taxometric methods. In P. C. Kendall & J. N. Butcher (Eds.), *Handbook of research methods in clinical psychology* (pp. 127–181). New York: Wiley.

Melton, A. W. (Ed.). (1964). *Categories of human learning*. New York: Academic Press.

Mishkin, M. (1978). Memory in monkeys severely impaired by combined but not by separate removal of amygdala and hippocampus. *Nature, 273*, 297–298.

Mishkin, M., Malamut, B., & Bachevalier, J. (1984). Memories and habits: Two neural systems. In G. Lynch, J. L. McGaugh, & G. Weinberger (Eds.), *The neurobiology of learning and memory*. New York: Guilford Press.

Moscovitch, M. (1984). The sufficient conditions for demonstrating preserved memory in amnesia. In L. R. Squire & N. Butters (Eds.), *Neuropsychology of memory*. New York: Guilford Press.

Murdock, B. B. Jr. (1960). The immediate retention of unrelated words. *Journal of Experimental Psychology, 60*, 222–234.

Nilsson, L.-G. (1984). New functionalism in memory research. In K. M. J. Lagerspetz & P. Niemi (Eds.), *Psychology in the 1990's*. Elsevier Science Publishers B.V. (North Holland).

Oakley, D. A. (1981). Brain mechanisms of mammalian memory. *British Medical Bulletin, 37*, 175–180.

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon Press.

Olton, D. S. (1983). Memory functions and the hippocampus. In W. Seifert (Ed.), *Neurobiology of the hippocampus*. New York: Academic Press.

Olton, D. S. (1984). Comparative analysis of episodic memory. *Behavioral and Brain Sciences, 7*, 250–251.

Olton, D. S., Becker, J. T., & Handelmann, G. E. (1979). Hippocampus, space and memory. *The Behavioral and Brain Sciences, 2*, 313–365.

Rescorla, R. A. (1980). *Second-order conditioning*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Roberts, W. A., & Grant, D. S. (1976). Studies of short-term memory in the pigeon using the delayed matching to sample procedure. In D. L. Medin, W. A. Roberts, & R. T. Davis (Eds.), *Processes of animal memory*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Rozin, P. (1976). The psychobiological approach to human memory. In M. R. Rosenzweig & E. L. Bennett (Eds.), *Neural mechanisms of learning and memory*. Cambridge, MA: MIT Press.

Ruggiero, F. T., & Flagg, S. F. (1976). Do animals have memory? In D. L. Medin, W. A. Roberts, & R. T. Davis (Eds.), *Processes of animal memory*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Savory, T. (1970). *Animal taxonomy*. London: Heinemann.

Schacter, D. L., & Moscovitch, M. (1984). Infants, amnesics, and dissociable memory systems. In M. Moscovitch (Ed.), *Infant memory*. New York: Plenum Press.

Seamon, J. G. (1984). The ontogeny of episodic and semantic memory. *Behavioral and Brain Sciences, 7*, 254.

Shallice, T. (1979). Neuropsychological research and the fractionation of memory systems. In L. G. Nilsson (Ed.), *Perspectives on memory research*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Simpson, J. (1961). *Principles of animal taxonomy*. New York: Columbia University Press.

Sneath, P. H. A., & Sokal, R. R. (1973). *Numerical taxonomy*. San Francisco: Freeman.

Sokal, R. R. (1974). Classification: Purposes, principles, progress, prospects. *Science, 185*, 115–123.

Solbrig, O. T. (1966). *Evolution and systematics*. New York: Macmillan.

Spear, N. E. (1971). Forgetting as retrieval failure. In W. K. Honig & P. H. R. James (Eds.), *Animal memory*. New York: Academic Press.

Spear, N. E. (1984). Behaviours that indicate memory: Levels of expression. *Canadian Journal of Psychology, 38*, 348–367.

Squire, L. R. (1982). The neuropsychology of human memory. *Annual Review of Neuroscience, 5,* 241–273.

Tiberghien, G. (1984). Just how does ecphory work? *Behavioral and Brain Sciences, 7,* 255–256.

Tolman, E. C. (1949). There is more than one kind of learning. *Psychological Review, 56,* 144–155.

Tulving, E. (1962). The effect of alphabetical subjective organization on memorizing unrelated words. *Canadian Journal of Psychology, 16,* 185–191.

Tulving, E. (1982). Synergistic ecphory in recall and recognition. *Canadian Journal of Psychology, 36,* 130–147.

Tulving, E. (1983). *Elements of episodic memory.* Oxford: Clarendon Press.

Tulving, E. (1984). Multiple learning and memory systems. In K. M. J. Lagerspetz & P. Niemi (Eds.), *Psychology in the 1990's* (pp. 163–184). North Holland: Elsevier Science Publishers B.V.

Tulving, E. (1985a). How many memory systems are there? *American Psychologist,* in press.

Tulving, E. (1985b). Memory and consciousness. *Canadian Journal of Psychology, 26,* 1–12.

Tulving, E., & Colotla, V. (1970). Free recall of trilingual lists. *Cognitive Psychology, 1,* 86–98.

Tulving, E., Schacter, D. L., & Stark, H. A. (1982). Priming effects in word-fragment completion are independent of recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition. 8,* 336–342.

Warrington, E. K., & Weiskrantz, L. (1982). Amnesia: A disconnection syndrome? *Neuropsychologia, 20,* 233–248.

Watkins, M. J. (1974). Concept and measurement of primary memory. *Psychological Bulletin, 81,* 695–711.

Waugh, N. C. (1961). Free versus serial recall. *Journal of Experimental Psychology, 62,* 496–502.

Weiskrantz, L. (1968). Some traps and pontifications. In L. Weiskrantz (Ed.), *Analysis of behavioral change.* New York: Harper & Row.

Weiskrantz, L. (1982). Comparative aspects of studies of amnesia. In D. E. Broadbent & L. Weiskrantz (Eds.), *Philosophical Transactions of the Royal Society of London, B.* London: The Royal Society, 298, 97–109.

World Health Organization (1974). *Glossary of mental disorders and guide to their classification.* Geneva: WHO.

In L.-G. Nilsson & T. Archer
(Eds.), Perspectives on Learning
and Memory, 1985.
Hillsdale, N.J.: Erlbaum