Chapter 2

Learning Hierarchies[3]

Robert M. Gagné

This article builds upon the theory of cumulative learning that was described in the previous chapter. It blends theoretical explanations with the description of a practical analysis technique for designers to use when determining the appropriate content for a given piece of instruction. This work continues to reflect Gagné's interest in school curriculum applications. Learning hierarchies, originally called hierarchies of knowledge, were not new to Gagné in 1968; he had introduced them earlier in the 1960in his studies of the acquisition of knowledge. His work in this area is interesting in that it reflects an early departure from the behavioral orientation that was predominant. However, it was consistent with the emphasis on behavioral objectives and the systematic analysis of instruction, both of which were then considered quite innovative.

A few years ago, in the course of studies of the learning of tasks resembling those learned in schools (Gagné, 1962), I used the term "learning hierarchy" to refer to a set of specified intellectual capabilities having, according to theoretical considerations, an ordered relationship to each other. It was possible, I stated, beginning with a clear statement of some terminal objective of instruction, to analyze this final capability into subordinate skills in an order such that lower-level ones could be predicted to generate positive transfer to higher-order ones. The entire set of ordered intellectual skills formed a hierarchy that was considered to bear some relation to a plan for effective instruction.

An example of a hierarchy, pertaining to the addition of integers, is shown in Figure 2.1. In the framework of instruction in "modern math," children learn two distinguishable terminal capabilities: one of these, shown on the right, is simply finding sums of positive and negative numbers; a second, shown on the left, constitutes a demonstration of the

---

logical validity of adding any pair of integers, using the properties of the number system to effect this demonstration. For both these tasks, an analysis revealed a set of subordinate capabilities shown in the figure, some in common and some not in common, ranging down to some relatively simple skills which the children were presumed to possess at the beginning of their instruction.

What I should like to do at this point is to tell you some things I have learned about learning hierarchies in the last couple of years. In part, these things have been learned by my research and the research of other people; and in part, from the various reactions I have received about them from many sources. I need to say, surely, that critical comments have most certainly caused me to rethink and clarify, at least in my own mind, what the nature, characteristics, and uses of learning hierarchies may be. Since such hierarchies contain elements of theory, I am most eager to alter or augment these elements to provide improved prediction, if that is possible. What I am likely to be most obstinate about changing, however, is the basic idea or the feasibility of predicting optimal sequences of learning events.

Characteristics of Learning Hierarchies

What are the characteristics of a learning hierarchy? How does one know when he has one, and what precisely can be predicted from it? To find initial answers to such questions, one can review the kind of study which first gave rise to the idea. This was a study derived from an investigation of the learning of a task of constructing formulas for the sums of number series (Gagné, 1962). In the original study using programmed
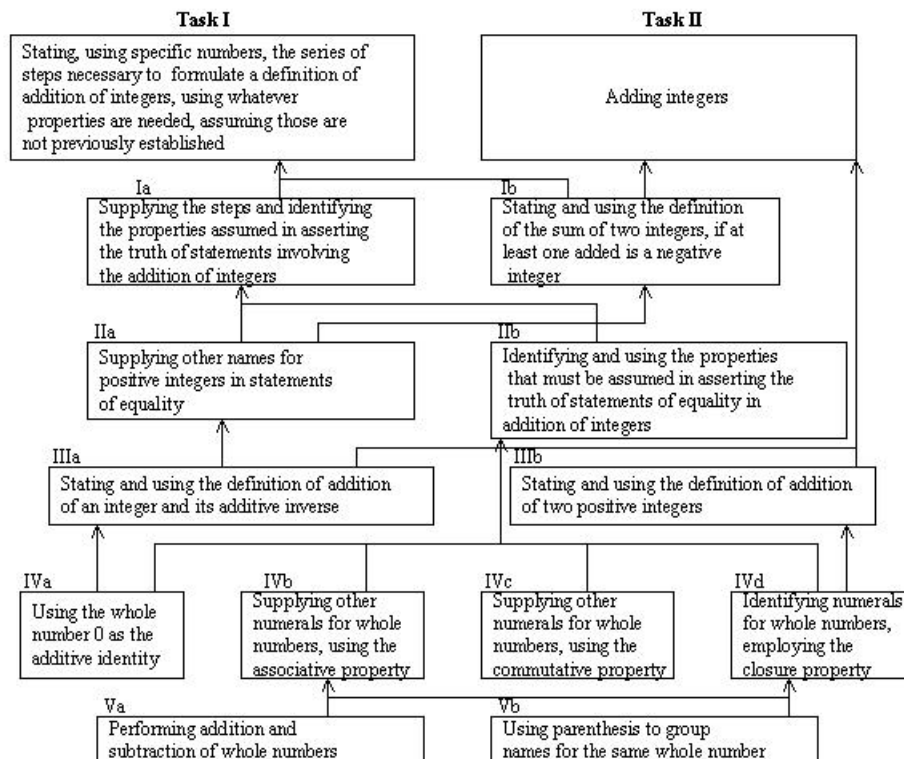

Figure 2.1 A learning Hierarchy on the Addition of Integers (From Gagne, R.M., Mayor, J.R., Garstens, H. L. and Paradise, N. E. (1962). Factors in acquiring knowledge of a mathematical task. Psychological Monographs, 76, No. 526).

instruction (Gagné & Brown, 1961), individual differences in learning from the program were of course highly evident. One could have attributed them to differences in "intelligence," and let it go at that. But it seemed to me these differences in learning performance were more likely attributable to certain identifiable skills which were not directly represented in the program itself, but which were needed along the way in doing what the program demanded. They were activities that the learners could or could not do, and which the program was accordingly more or less successful in teaching them to do.

The next step was to figure out what these "subordinate skills" might be. Beginning with the final task, I found it was possible to identify nine subordinate capabilities, related to each other in an ordered way, by successively asking the question concerning each task, "What would the individual already have to know how to do in order to learn this new capability simply by being given verbal instructions?" It is probably of some importance to note that the kinds of capabilities identified in this manner did not directly pertain to number series, but rather included such skills as the following:

a) Identifying the location of numerals in a tabular array by means of letters giving their row and column location
b) Completing statements of equality by supplying missing numbers in equations containing mixed arithmetic operations
c) Identifying the numbers or letters in a tabular array which formed certain spatial patterns represented by lines connecting at 90 degrees and 45 degrees.

I emphasize that the subordinate skills so identified are not related to number series in a logical sense; what they are related to, psychologically, is the kind of behavior the learner has to engage in if the learner is going to be successful at figuring out from a tabular array of number series properties, how to formulate an equation for their sum.

Having identified a hierarchy of capabilities in this way, the next step was to test its validity. First, a test was made on a number of subjects to determine which of these subordinate tasks they already knew how to do. Two subjects could do all of the subordinate tasks but not the final one; two could do all but two; one all but three; and two all but four. Each of these learners was then taught to do whichever subordinate tasks he couldn't initially perform. Then, having completed this learning, each was given verbal directions about how to do the final task, without any practice on it. Six out of these seven subjects then proceeded to execute the final task of making a formula for four number series, which he hadn't seen before. Additional evidence showed that a similarly marked change in capability was brought about in these learners at each "level"

of the hierarchy for which instruction had been given.

Certain patterns of responding to the tests of subordinate tasks also were revealed in this study. Specifically, those who got subordinate skill number 1 correct, also got all the skills lower in the hierarchy correct. Those who got number 2 correct, and missed number 1, got all the skills lower than number 2 correct also. In other words, in these seven learners, there was in fact an ordered relationship (similar to that displayed in a Guttman-type scale)[4] among the subordinate capabilities measured.

On the basis of this brief review, I should like to consider the question of what a learning hierarchy is. What properties of the learning hierarchy were either postulated or revealed in this study?

First, the question by means of which the analysis is begun, namely, "What would the individual have to know how to do...etc.," implies that one is searching for subordinate tasks which will transfer positively to the learning of the task in question. The criterion for such transfer is a stringent one—it is desired that the subordinate skill or skills facilitate the learning to such an extent that it will occur when only verbal instructions, and no further trials of practice, are given. It is evident that choices are being made here, since there are perhaps a number of kinds of subordinate skills which would, under suitable conditions, exhibit some degree of transfer to a given learning task. The method doesn't imply that all of these are searched for, but only those that will meet this stringent criterion. Therefore, it is fair to say that a subordinate capability identified by this method is a skill which is hypothesized to exhibit a substantial amount of positive transfer to the learning of the skill in question.

Second, how does one know if the order assigned to the skills in the hierarchy is correct? To specify this order, one depends first of all on the application of knowledge about transfer of learning, which comes from a great number of sources. A general guide to such ordering is the one I have described (Gagné, 1965), in which simple responses are subordinate to chains or multiple discriminations, which in turn are subordinate to classifying, which in turn is subordinate to using principles or rules. But this of course is rather general guidance, and does not begin to account adequately for the specific choices that must be made in any particular instance. Sometimes one is not sure about the location of a subordinate capability, particularly as to whether it is truly subordinate

---

[4] *Editor's Note:* Louis Guttman had an interest in the development of unidimensional scales - consisting of items unrelated to the characteristic being measured. He developed a technique called scale analysis, a major contribution to the methodology of questionnaire construction and analysis. A Guttman-type scale would be unidimensional in nature. For further information, see Guttman, L. (1944). A basis for scaling quantitative data. *American Sociological Review, 9*, 139-150.

or merely at the same level.

Empirical tryout of the series of hypotheses represented by a hierarchy seems to be a reasonable approach to this problem. On the basis of such a tryout, one can in effect determine whether a particular skill transfers positively to another, or whether they are independent, or whether perhaps they co-vary in their transfer effects. In one paper I have made some suggestions about how these determinations might be made (Gagné, 1967), but I perceive these to be very unsophisticated compared with procedures I can only dimly imagine. An example of a successful tryout of this sort is in a study by Cox and Graham (1966), using a task of elementary mathematics. They were able to show that an initially hypothesized order was incorrect, according to their results. When the hierarchy was rearranged, the existence of an order of subordinate skills was confirmed. Thus it seems to me reasonable to suppose that many individual hypotheses about transfer represented in a hierarchy may have to be checked by some empirical means. If they turn out to be wrong, the conservative conclusion surely is that something is wrong with the specific hierarchy proposed. To the contrary, however, it does not seem reasonable to conclude on the basis of such evidence that *all* hierarchies are wrong.

A third characteristic of hierarchies seems to be of considerable interest. Do they represent a sole learning route to the learning of the final task, or perhaps even a most efficient learning route? Must each individual learner necessarily proceed to acquire each subordinate skill in order to enable him ultimately to learn the final task? By reference again to the method of analysis by means of which the hierarchy is generated, it is quite apparent that the answer to this question is no. Nothing in the method of analysis tells us about the capabilities of the individual learner. A given individual may be able to "skip" one or more of the subordinate tasks, just as a given learner may be able to "skip" parts of an adaptive program of instruction. Another individual may be able to bring to bear on the learning of any given skills some capability which comes from quite a different domain of knowledge, which is not even represented in the hierarchy.

A learning hierarchy, then, in the present state of our knowledge, cannot represent a unique or most efficient route for any given learner. Instead, what it represents is the most probable expectation of greatest positive transfer for an entire sample of learners concerning whom we know nothing more than what specifically relevant skills they start with.

A related point needs to be made about what a learning hierarchy represents, and what it does not represent. Perhaps the best way to say this is that a learning hierarchy does not represent everything that can be learned, nor even everything that is learned, within the domain it attempts to describe. In particular, a diagram of a hierarchy does not

represent what is perhaps the most important result of learning, the potentiality for transfer that is generated. I have spoken of the events reflected in a learning hierarchy as cumulative learning (Gagné, 1968). The cumulative effects of such learning show themselves, in a minimal fashion, by the occurrence of positive transfer from one level of skill to another. But beyond this, each new capability that is learned may generalize to many other situations and domains that cannot possibly be represented on a single chart. Consider, for example, how a child who has learned the skill of volume conservation in rectangular containers, and the skill of conservation in cylindrical containers, may then learn to "conserve" volume in irregularly shaped containers (Gagné, 1968, p. 187). I have pointed out a number of particular subordinate skills from which positive transfer may be expected. The new task can be learned much more quickly than the old, not because the latter is subordinate to it, but because there are many common subordinate skills from which positive transfer may be expected.

According to this reasoning, there are latent consequences of cumulative learning, which are not directly represented on a diagram of a learning hierarchy. Were they to be represented, one would have to draw lines of transfer, somewhat as indicated in Figure 2.2. Depending on particular circumstances in the individual learner, there may be transfer from a lower level, in other words, "skipping." As another possibility, transfer may occur from quite a different domain of knowledge, as when one uses a skill at identifying number series patterns to solve a problem in classifying patterns of letters. Still a third possibility, which should not be overlooked, is the atypical combination of subordinate skills which, while they may seem conceptually very different, may in the case of an individual learner be able to combine to yield a rather unexpected source of learning transfer. A learning hierarchy cannot, in any practical sense, represent all of these possibilities. Yet to deny their existence would be wrong, and in fact quite contrary to the basic conception of what cumulative learning is supposed to accomplish.
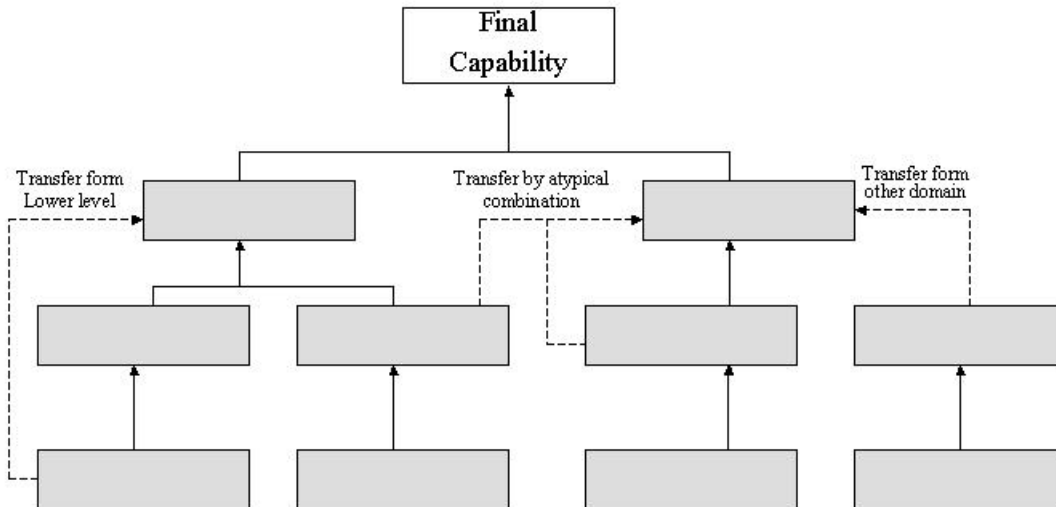
Figure 2.2 The Latent Consequences of Cumulative Learning, Indicated by Sources of Positive Transfer to the Learning of Advanced Capabilities.

Intellectual Skills

I turn now to one of the most important characteristics of learning hierarchies, and one concerning which I myself have been inconsistent in past writings. The question is, what exactly are these entities, sometimes called capabilities, which make up a learning hierarchy? The answer I would now give is the following. They are *intellectual skills*, which some writers would perhaps call *cognitive strategies*. What they are not is just as important. They are not entities of verbalizable knowledge. I have found that when deriving them, one must carefully record statements of "what the individual can do," and just as carefully avoid statements about "what the individual knows."

I believe that my previous formulation of these entities is misleading, when it deals with what are called "concepts" and "principles." I should prefer to substitute for these, words emphasizing capabilities for action, such as "classifying," and "rule-following." This is more than a nominal change, however. I mean that what learning hierarchies describe is, in computer language, subroutines of a program; what they do not describe is the

facts or propositions retrievable from memory as verbalizable statements.

Why do I emphasize this distinction, and what has led me to make it? First, it is surely noteworthy that the original hierarchies were developed in connection with mathematics tasks. If one stops to think about it, the substance of mathematics is largely a set of skills for manipulating numbers. They differ in complexity, of course, and also in specificity. But they are always intellectual skills, and they are not (and probably should not be) verbalizable knowledge. In the original study using number series (Gagné, 1962), for example, what was being learned was not "knowledge" about number series, but a set of particular *skills* of forming relationships among sets of numbers displayed in a systematic array.

You may recall that I incautiously attempted to generalize the ideas of learning hierarchies to such subject matter fields as the social sciences (cf. Gagné, 1967). In doing this, it is quite easy to fall into the trap of describing "knowledge" entities rather than skill entities. For example, some time ago one of my students worked out a learning hierarchy on weather prediction, with my help and acquiescence. The idea was to teach fourth graders how to predict weather from a weather chart superimposed on a map showing terrain features. The subordinate entities of this hierarchy had a high degree of plausibility, and appeared to describe what the child needed to know if he was going to predict the weather. When a teaching program based upon this chart was tried out, with much good will and persistence, the results can most succinctly be described by saying that it didn't work. The children did not learn much when a sequence of instruction based upon this chart was followed. Under these circumstances, little or no evidence could be seen in the data that positive transfer was occurring from one level of the hierarchy to the next.

I believe that the fundamental reason for this lack of success was that this was not a learning hierarchy for the task of predicting weather. It did not represent the intellectual skills the child needs to possess in tackling the job of figuring out from the "weather chart" how to make a forecast of the weather. I haven't yet made an analysis that satisfies me, but I suspect the intellectual skills that should be included are such things as these: (1) from general descriptions, formulating relevant propositions in syllogistic form; (2) making a systematic review of the effects of specific factors on an air mass; and (3) constructing specific statements describing weather at designated future times. It should be noted that such skills as these were not represented in the original formulation. They represent intellectual operations that the child can do. But they are not descriptions of what he knows (that is, of what he can recall in the sense of non-verbatim verbal propositions).

Then there is the evidence about the effectiveness of certain kinds of sequences in instruction, or in instructional programs. First I should say that I am not sure a learning hierarchy is supposed to represent a *presentation sequence* for instruction in an entirely uncomplicated way. Presumably, there should be some relation between an ordered set of intellectual skills and an ordering of a sequence of presentation of a set of frames or topics in an instructional program. Results like those of Payne, Krathwohl, and Gordon (1967), however, surely serve to give added emphasis to the distinction between verbalizable knowledge and intellectual skills. The painstaking study conducted by these investigators showed in a most convincing way that sequence of presentation, so far as reasonably mature adult learners are concerned, does not affect what is learned. The authors of this study suggest that, even when frames or topics are presented in scrambled order, the adult learner is able to make them into a coherent and meaningful internal arrangement, and to learn from them. Accordingly, one is led to believe from this study, or others like it, that a learner may acquire certain intellectual skills from a presentation that is quite disorganized when viewed as a sequence of verbalizable knowledge.

It is conceivable that this line of reasoning also applies to the study of Merrill (1965), who found no advantage to review and correction following each topic of an instructional program on imaginary science, as opposed to a condition of no review and correction. While I am by no means highly confident of this interpretation, I believe it might be examined within this general context. To summarize the point, it is that learners can acquire verbalizable knowledge, and even intellectual skills, from sequences of presentation that are altered in various ways from what may he considered "highly organized." The hypothesis I should like to reaffirm, however, is that regardless of presentation sequence, if one is able to identify the intellectual skills that are learned, he will find them to generate positive transfer in an ordered fashion.

Another line of thinking which I believe reinforces the distinction between intellectual skills and verbalizable knowledge comes from an analysis of the kinds of tasks described by Guilford (1967). While I have not undertaken an analysis of all the tasks Guilford describes, I have done some of them, and enough to lead me to believe that in most cases they are sampling both these kinds of entities. The performance being measured, in other words, typically samples the stored verbalizable knowledge the individual has available; and it also samples the intellectual skills that can be brought to bear upon the task. Consider a rather simple example, shown in Figure 2.3. "Which of these letter combinations does not belong with the rest?" The answer is 3, because it contains two vowels.
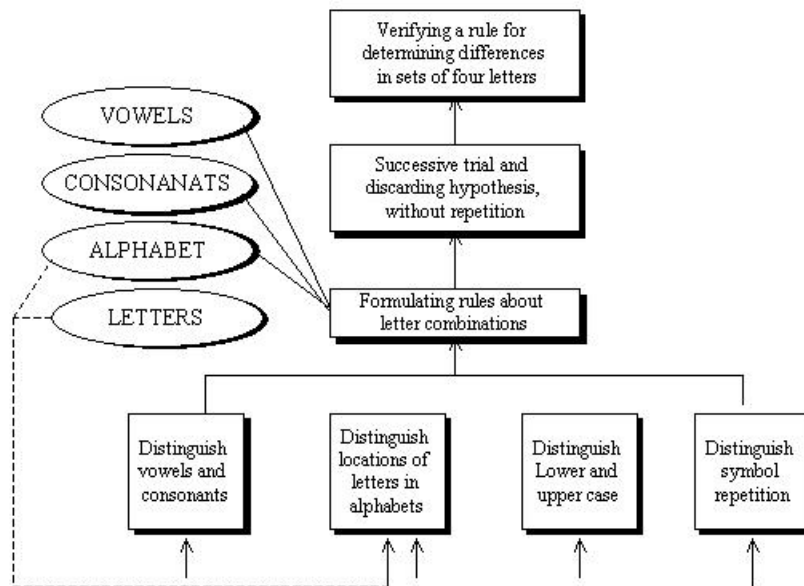
Figure 2.3 A learning Hierarchy for a Task of Letter Combinations, indicating the contribution of Intellectual Skills (boxes) and Verbal Knowledge (circled entities). (Example is from Guildford, J.P. (1968). The Nature of Human Intelligence. New York: McGraw-Hill, p. 42).

What kinds of intellectual skills does the individual bring to bear on such a task? I have suggested what I think they might be, in the hierarchy of boxes in Figure 2.3. They include such things as (1) making hypotheses which are tried and discarded, without repetition; and (2) distinguishing various features of letter combinations, such as vowels and consonants, location of letters in the alphabet, symbol repetition, and so forth. But it is equally apparent, is it not, that the individual who can solve this task also brings to bear some stored verbal entities: he must know what the vowels are, what the consonants are, what the alphabet is, and what the letters are. Both intellectual skills

and an elementary kind of verbalizable knowledge are required in performing the task. But my hypothesis is that they are learned in different ways. The skills have an ordered relation to each other such that subordinate ones contribute positive transfer to superordinate ones. But I do not suppose that the verbalizable entities necessarily have this relationship to each other. Stated in overly simple fashion, one does not have to learn consonants and vowels first in order to insure greatest transfer to learning the entire set of letters; and one does not have to learn the letters first before learning their position in the alphabet.

This example is admittedly an elementary one, and I should not want that fact to obscure what I think to be the generalizability of this distinction. Consider another task, that of solving five-letter anagrams. The work of Mayzner and Tresselt (1965) and others has shown that such a task derives positive transfer from an identifiable set of intellectual skills, pertaining to the formulation of hypotheses regarding probabilities of letter combinations, probabilities of initial letter occurrences, and others. But it is equally evident that an individual learner who is successful at solving anagrams must have a store of verbalizable knowledge to call upon, which in this case are words. In solving a set of anagrams, the individual will show greater success if he knows a large number of words, besides having mastered the intellectual skills involved.

These are the major reasons, then, why I am led to think that learning hierarchies are descriptions of the relationships of positive transfer among intellectual skills, but that they are not descriptions of how one acquires verbalizable knowledge. Obviously, in solving any given problem, both kinds of retained entities must be brought to bear. And it seems equally true that, when a new intellectual skill is being acquired, knowledge must be available to the learner, since the skill cannot be learned "in a vacuum." I do not, in other words, wish to say that either kind of entity is the more important for learning. Both are essential. What seems to me most evident is that they need to be distinguished, and that the conditions governing positive transfer to them are probably very different.

To complete this account of the distinction between intellectual skills and verbalizable knowledge, it is of some importance to point out that this matter has possibly profound implications in its relation to the curriculum. Most educational psychologists, to be sure, recognize the distinction and clearly state that both intellectual skills and verbalizable knowledge must be learned in the schools. Ausubel (1968), for example, acknowledges the difference early in his text on educational psychology. Skinner (1968) draws a distinction between behaviors to be learned for dealing with particular classes of events, and precurrent self-management behaviors, which are more general in their applicability. Rothkopf (1968) distinguishes mathemagenic behaviors from the substance of what is learned. But the importance of each of these types of learned capability for curriculum

design and planning would doubtless be estimated differently by these theorists, and probably still differently by me.

I should be inclined to entertain the notion that the most important things learned in school are intellectual skills, and not verbalizable knowledge. The major reason is, very simply, that one can always look up the knowledge, but the skills have to become "built-in." I can obviously not do justice to this very weighty question at this time. The curriculum implications are such as to lead to a heavy emphasis on what is often referred to as "process," in contrast to content. In elementary science, for example (cf. AAAS Commission on Science Education, 1967), this line of thinking leads one to prefer teaching children the intellectual skills involved in classifying, measuring, and predicting, rather than the verbalizable knowledge of the accomplishments of science.

Evidence Relevant to Learning Hierarchies

Now I must return to the major theme of learning hierarchies. To characterize them briefly, they represent an ordered set of intellectual skills, such that each entity generates a substantial amount of positive transfer to the learning of a not-previously-acquired higher-order capability. The learning of each entity also requires the recall of relevant verbalizable knowledge, which, however, is not itself represented in the hierarchy.

What kinds of evidence should be sought in the attempt to verify the hypotheses represented in a learning hierarchy, and what are the sources of such evidence? First of all, I should be inclined to seek evidence about the transfer of learning from one class of intellectual skill to another—in other words, from studies dealing with two successive levels of a hierarchy, rather than with all the levels at once. The reason for this is a fairly simple one involving consideration of the usual controls of an experiment. If one measures transfer from task A to task B, there will usually be certain proportions of success for task B. But one cannot then take the same groups of subjects, varying in their success with task B, and go on to measure transfer to task C, without violating certain principles of random selection. Thus, the basic experimental method remains one of measuring positive transfer from task A to task B; or alternatively, from task B to task C.

There is quite a good deal of evidence concerning positive transfer from one class of intellectual skill to another. For example, in the verbal paired-associate learning field, the evidence reviewed by Battig (1968) is to the effect that the learning of paired associates is typically facilitated by prior discrimination learning on stimulus-terms and response-terms, as well as by prior learning of stimulus coding responses. When one looks at

categorizing skills (or concepts) like those exhibited by children in performing reversal-shift tasks, recent investigations such as those of Tighe (1965), Smiley and Weir (1966), and Johnson and White (1967) clearly demonstrate the importance of relevant prior learning of dimensional discriminations for transfer to the reversal task. Similarly, the different sort of classifying required in transposition tasks is shown to derive positive transfer from prior discrimination learning in the studies of Beaty and Weir (1966) and Caron (1966).

The importance of prior classification learning for positive transfer to rule learning is shown in a number of studies dealing with conservation tasks of a type derived from the work of Piaget. Beilin, Kagan and Rabinowitz (1966), for example, found prior classification training to transfer to the task of water-level representation in children, to a greater extent than verbal training. In this field of interest, a study of particular relevance to the present discussion is that of Kingsley and Hall (1967). These investigators made a specific analysis to derive a hierarchy of subordinate skills in conservation tasks. They then tested each child to determine which of the subordinate skills he knew, and proceeded to train each of the missing ones. The method, in other words, resembled that employed in "The acquisition of knowledge" (Gagné, 1962), and substantial amounts of positive transfer to the final conservation tasks were obtained.

There are also a number of recent studies verifying the general idea of positive transfer to problem-solving situations from prior learning on subordinate relevant rules. DiVesta and Walls (1967), for example, demonstrated positive transfer from relevant "pre-utilization" training to the Maier two-string problem. Davis (1967) showed the effectiveness for transfer of previously learned verbal rules to switch-light problems, and a similar theme is developed by Overing and Travers (1966, 1967) in their studies of the problem of hitting an underwater target. In problems concerning mathematical groups and combinatorial topology, Scandura and Wells (1967) demonstrated positive transfer effects from prior learning in concrete situations involving relevant rules.

In this brief sampling of relatively recent studies, one can see repeated many times the general affirmation of the hypothesis that the learning of each particular category of intellectual skill depends substantially, in a positive transfer sense, on the previous learning of another particular category of intellectual skill. In brief, problem-solving draws positive transfer from prior rule learning, which is contributed to in the same sense by prior classification learning, which is in turn strongly affected by prior discrimination learning, and so on. I should say, therefore, that I look for verification of the learning hierarchy idea in studies of positive transfer from one intellectual skill to another. In studies of this sort over the past few years, there is a good deal of confirming evidence.

The other major type of study from which evidence about learning hierarchies may be derived is one that attempts to try out a total hierarchy, applicable perhaps to a limited topic, but in which the various levels of intellectual skill are to be learned in a single instructional sequence. The collaborative studies I did on the learning of algebraic equation-solving and adding integers (Gagné & Paradise, 1961; Gagné, Mayor, Garstens, & Paradise, 1962), for example, are of this sort. The results one first obtains from such studies may indicate that some incorrect hypotheses were made concerning predictions about positive transfer. Specifically, a capability thought to be subordinate to another may turn out to be superordinate, or even coordinate. Such a finding calls for the rearrangement of the hierarchy, as was done, for example, in the previously mentioned study by Cox and Graham (1966) dealing with the addition of two-place numbers. Following such a step, the new hierarchy can then be tried out, in order to seek evidence of positive transfer from one "level" to the next.

I need to mention that methods of analyzing data from such a study are not at all clear. Various possibilities have been tried beginning with Guttman scaling techniques, but none seem entirely satisfactory as yet. It is highly encouraging to know, however, that the measurement techniques needed for such analyses are apparently being worked on by a number of highly competent people. Hopefully, these will contrast with the rather crude methods used in the study from which Figure 2.1 was taken. Just to remind you what these were, by reference to Figure 2.1, what we attempted to do was to find the probability of achieving Task 1 for those learners who had learned an immediately subordinate capability, Ia, and to contrast this probability with that for learners who had *not* learned the same subordinate capability. The findings were 73% for the first set of learners, and only 9% for the second set. In other words, there was indeed substantial positive transfer. Similar confirming findings were reported for all of the comparisons possible within the learning hierarchy.

My present estimate then is that there are two major kinds of study which are likely to provide evidence concerning learning hierarchies. One investigates only two "levels" of a hierarchy at a time, and in effect becomes a more-or-less traditional study of positive transfer between categories of intellectual skill. The other type attempts to construct a hierarchy which applies to longer sequences of instruction, and which after first establishing a suitable order for the capabilities to be learned, seeks a measure of the dependence (in the positive transfer sense) of one learned entity on another.

Concluding Statement

It will surely be apparent from this restatement and possible clarification of a theoretical view, that in one sense the notion of a learning hierarchy reduces itself to the notion of positive transfer. The question remains, what transfers to what? My answer has been,

and still is, that the "what" of this question can be answered in terms of different varieties of learned capabilities. In particular, specific responses transfer to discriminations, which transfer to classifications, which transfer to rules, which in turn may transfer to more complex forms of rule-governed behavior, such as that exhibited in problem solving.

The entities that are affected by positive transfer in this manner deserve to be called intellectual skills or strategies. But it seems important to distinguish these from verbalizable knowledge. While the learning and retention of the latter entities must surely have a theoretical rationale, for example, Ausubel's (1968), it seems to me to differ in respect to the properties of positive transfer which are applicable to intellectual skills.

When one says, therefore, as I am inclined to say, that we need more evidence about learning hierarchies, he may simply be repeating something that has surely been said before: we need more evidence about positive transfer. Despite the encouraging signs from recent studies I have mentioned, it appears that there is an enormous amount still to be known about this subject. Perhaps reducing "learning hierarchies" to such familiar terms will encourage more investigation and more systematic thinking about this phenomenon, which is so obviously of central importance to education.

References

AAAS Commission on Science Education. (1967). Science—A process approach, purposes, accomplishments, expectations. Washington, DC: American Association for the Advancement of Science.

Ausubel, D. P. (1968). Educational psychology: A cognitive view. New York: Holt, Rinehart and Winston.

Battig, W. F. (1968). Paired-associate learning. In T.R. Dixon and D.L. Horton (Eds.), Verbal behavior and general behavior theory (pp. 149–171). Engle wood Cliffs, NJ: Prentice-Hall.

Beaty, W. E., & Weir, M. W. (1966). Children's performance on the intermediate-size transposition problem as a function of two different training procedures. Journal of Experimental Child Psychology, 4, 332–340.

Beilin, H., Kagan, J., & Rabinowitz, R. (1966). Effects of verbal and perceptual training on water level representation. Child Development, 37, 317–330.

Caron, A. J. (1966). Far transposition of intermediate-size in preverbal children. Journal of Experimental Child Psychology, 3, 296–311.

Cox, R. C., & Graham, G. T. (1966). The development of a sequentially scaled achievement test. Paper presented at the annual meeting, American Educational Research Association, Chicago. (ED 010 206)

Davis, G. A. (1967). Detrimental effects of distraction, additional response alternatives, and

longer response chains in solving switch-light problems. Journal of Experimental Psychology, 73, 45–55.

DiVesta, F. J., & Walls, R. T. (1967). Transfer of object-function in problem solving. American Educational Research Journal, 4, 207–216.

Gagné, R. M. (1962). The acquisition of knowledge. Psychological Review, 69, 355–365.

Gagné, R. M. (1965). The conditions of learning. New York: Holt, Rinehart and Winston.

Gagné, R. M. (1967). Curriculum research and the promotion of learning. In R. Tyler; R. Gagné, & M. Scriven (Eds.), Perspectives of Curriculum Evaluation, AERA Monograph Series on Curriculum Evaluation, No. I, (pp. 19–38). Chicago: Rand-McNally.

Gagné, R. M. (1968). Contributions of learning to human development. Psychological Review, 75, 177–191.

Gagné, R. M., & Brown, L. T. (1961). Some factors in the programming of conceptual learning. Journal of Experimental Psychology, 62, 313–321.

Gagné, R. M., Major, J. R., Garstens, H. L., & Paradise, N. E. (1962). Factors in acquiring knowledge of a mathematical task. Psychological Monographs, 76, No. 526.

Gagné, R. M., & Paradise, N. E. (1961). Abilities and learning sets in knowledge acquisition. Psychological Monographs, 75, No. 518.

Guilford, J. P. (1967). The nature of human intelligence. New York: McGraw-Hill.

Johnson, P. J., & White, P. M., Jr. (1967). Concept of dimensionality and reversal shift performance in children. Journal of Experimental Child Psychology, 5, 223–227.

Kingsley, R. C., & Hall, V. C. (1967). Training conservation through the use of learning sets. Child Development, 38, 1111–1126.

Mayzner, M. S., & Tresselt, M. E. (1965). Solving words as anagrams: An issue reexamined. Psychonomic Science, 3, 363–364.

Merrill, M. D. (1965). Correction and review on successive parts in learning a hierarchical task. Journal of Educational Psychology, 56, 225–234.

Overing, R. L. R., & Travers, R. M. W. (1966). Effect upon transfer of variations in training conditions. Journal of Educational Psychology, 57, 179–188.

Overing, R. L. R., & Travers, R. M. W. (1967). Variation in the amount of irrelevant cues in training and test conditions and the effect upon transfer. Journal of Educational Psychology, 58, 62–68.

Payne, D. A., Krathwohl, D. R., & Gordon, J. (1967). The effect of sequence on programmed instruction. American Educational Research Journal, 4, 125–132.

Rothkopf, E. Z. (1968). Two scientific approaches to the management of instruction. In R. M. Gagné & W. J. Gephart (Eds.), Learning research and school subjects. Itasca, IL: Peacock.

Scandura, J. M., & Wells, J. N. (1967). Advance organizers in learning abstract mathematics.

American Educational Research Journal, 4, 295–301.

Skinner, B. F. (1968). The technology of teaching. The Century Psychology Series. New York: Appleton-Century-Crofts. (ED 018 996)

Smiley, S. S., & Weir, M. W. (1966). Role of dimensional dominance in reversal and non-reversal shift behavior. Journal of Experimental Child Psychology, 4, 296–307.

Tighe, L. S. (1965). Effect of perceptual pretraining on reversal and non-reversal shifts. Journal of Experimental Psychology, 70, 379–385.