SECTION I

OVERVIEW OF STRATEGIES

Section I provides the theoretical foundations for the analysis of cognitive skills and the implications that are derived therefrom for how to teach those skills.

The precise analysis of cognitive learning is difficult, if not elusive, because it stands at the juncture of three separate analyses–the analysis of behavior, the analysis of stimuli used as teaching communications, and the analysis of knowledge systems or the content to be taught. (See Figure I.1)

Figure I.1

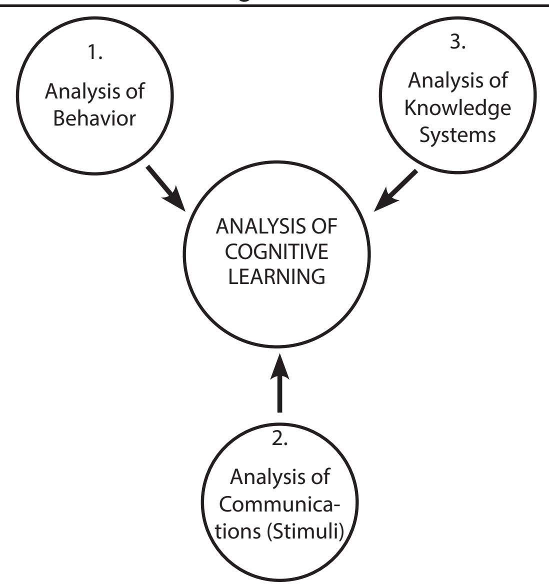

1. *The analysis of behavior* seeks *empirically-based* principles that tell what is universally true about the ways in which the environment influences behavior for different classes of learners.
2. *The analysis of communications* seeks principles for the *logical* design of communications that effectively transmit knowledge. These principles allow one to describe the range of generalizations that should logically occur when the learner receives specific sets of examples. The analysis of communications focuses on the ways in which examples are the *same* and how they *differ*.
3. *The analysis of knowledge systems* is concerned with *logically* organizing knowledge so that relatively efficient communications are possible for related knowledge.

The analysis that has received the most attention from psychological theories is the analysis of behavior (Hilgard & Bower, 1975). Although the other two analyses have received some theoretical attention (e.g. Gagne, 1970; Bloom, 1956; Markle & Tiemann, 1974), there has been little systematic effort to develop *precise* principles of communications used in instruction or to analyze knowledge systems.* This book frames behavior theory within a three-way analysis of human cognitive learning.

The three areas of analysis derive directly from the nature of cognitive learning. The first aspect of cognitive learning is that the *learner learns from the environment,* which means that the environment is somehow capable of communicating concepts or skills to the learner. The analysis of communications provides rules for designing these communications so they are effective transmitters.

Another aspect of cognitive learning is that it *always involves some topic or content.* When we think, we think about something, even if that something is a process. This aspect of cognitive knowledge carries basic implications for designing the communications that we present to the learner. We cannot communicate with the learner without communicating *something.* Conversely, if we are to understand how to communicate a particular bit of knowledge (such as knowledge of the color **red**, or knowledge about the operation of **square root**), we must understand the essential features of the particular concept that we are attempting to convey. Only if we understand what it is and how it differs from related

*Although analyses such as Gagne's deal with learning and with the teaching of concepts, principals, etc., the theoretical development is at best a beginning.

concepts can we design a communication that effectively conveys the concept to the learner.

The final aspect of cognitive learning has to do with the relationship of a given concept to other concepts. The word **large** is related to the word **blue** because both function as adjectives. The color **blue** is related to the color **red** because both have the properties of color. The relatedness of cognitive knowledge suggests that it is possible to develop a classification system for various types of knowledge (circle 3 in Figure I.1). If this classification system is to be of value to the instructional designer, the system should be designed so that the classification of a particular skill carries information about how to communicate that skill to a learner. Concepts that are structurally the same in some respects can be processed through communications that are the same in some respects.

Both the analysis of communications and the analysis of knowledge systems are *logical analyses* that involve assumptions about the learner. The analysis of behavior, however, investigates the learner and how the learner responds to specific communications. Chapter 1 presents an overview of the strategy that we will use to unite the three parts of the analysis; Chapter 2 further develops the analysis of communications; Chapter 3 outlines the organization of knowledge types that will be used throughout the book.

Chapter 1
Theoretical Foundations

A theory of instruction begins with the assumption that the environment is the primary variable in accounting for what the learner learns. The different skills learned by people in different environments suggest that the assumption is reasonable. People who live in primitive societies learn skills quite different from those learned by people who live in urban societies. Although the environment is assumed to be the primary cause of what is learned, it is not assumed to be the total cause. Within any group of people there are individual differences. Also, there are differences that correlate with the age of the learner. Therefore, the learner is also a variable.

To show the relationship between the role of the environment and the learner, we are faced with the basic problem of experimental control. We must control one of these variables (the environment or the learner) before we can make precise observations about the other variable. Ideally, we would rule out or eliminate one of these variables (either the environment or the learner) and observe the remaining variable in a pure state. This solution is not possible. A possible solution is to control one of the variables so that it functioned as if it were ruled out. We cannot readily achieve such control over the learner because we do not know precisely how to do it. However, such control is possible with the environment. We can design communications that are, ideally, faultless. Faultless communications are designed to convey only one interpretation. From a logical standpoint, these communications would be capable of teaching any learner the intended concept or skill. When we present such a communication to the learner, we effectively rule out the environment as a *variable*. The communication is not merely standardized; it is analytically or logically capable of transmitting the concept or skill to *any learner who possesses certain minimal attributes discussed later*. The learner either responds to the faultless communication by learning the intended concept, or the learner fails to learn the intended concept. In either case, the learner's performance is framed as the dependent variable. The extent to which the learner's performance deviates from the performance that would occur if the learner responded perfectly to the communication provides us with precise information about the learner. The deviations indicate the extent to which the learner is not a perfect "mirror" of the environment. Furthermore, these deviations are caused by the learner (not the environment, which has been controlled so that it is faultless).

The strategy of making the communication faultless and then observing the performance of the learner is the basis for the theory of instructions that we will develop. We will use this strategy in designing instructional sequences and in deriving principles for communicating with the learner. The following is a summary of the steps in our strategy, showing where logical analysis is used and where behavior analysis comes into play:

1. Design communications that are faultless using a *logical analysis* of the stimuli, not a *behavioral analysis* of the learner.
2. Predict that the learner will learn the concept conveyed by the faultless presentation. If the communication is logically flawless and if the learner has the capacity to respond to the logic of the presentation, the learner will learn the concept conveyed by the communication.
3. Present the communication to the learner and observe whether the learner actually learns the intended concept or whether the learner has trouble. This information (derived from a behavioral analysis) shows the extent to which the learner does or does not possess the mechanisms necessary to respond to the faultless presentation of the concept.
4. Design instruction for the unsuccessful learner that will modify the learner's capacity to respond to the faultless presentation. This instruction is not based on a logical analysis of the communication, *but on a behavior analysis of the learner.*

Note that the behavioral analysis comes into play only after the communication has been designed so that it is faultless. The faultless presentation rules out the possibility that the learner's inability to respond appropriately to the presentation, or to generalize in the predicted way, is caused by a flawed communication rather than by learner characteristics.

Assumptions About the Learner

The primary problem that we face in pursuing this strategy is that we do not know what constitutes a faultless communication *unless we make some assumptions about the learner.* Stated differently, assumptions about the learner and the communication vary together. The greater the assumed capabilities of the learner, the less the assumed responsibility of the communication. If we assume that the learner will learn from any exposure to the environment, we will provide communications that do not control details of the presentation. If we assume that the learner is not capable of learning from communications that are ambiguous, we will approach the design of communications quite differently. To provide for control of the maximum number of communication variables, we must postulate a *simple* learning mechanism. Also, we must assume that the learner's behavior is lawful, which means the learner who possesses the assumed mechanism will learn what the communication demonstrates or teaches.

The learning mechanism that we postulate has two attributes:

1. The capacity to learn any quality that is exemplified through examples (from the quality of **redness** to the quality of **inconsistency**).
2. The capacity to generalize to new examples on the basis of sameness of quality (and only on the basis of sameness).

These attributes suggest the capacities that we would have to build into a computer that functions the way a human does. Note that we are not asserting that these are the only attributes that a human possesses, merely that by assuming the two attributes we can account for nearly all observed *cognitive behavior.*

1. The Capacity to Learn Any Quality from Examples

This assumption indicates *what* the mechanism is capable of learning, not how it learns. A quality is any irreducible feature of the example. The simplest way to identify qualities is to begin with a concrete example. Any example (such as a pencil) has thousands of qualities, which relate to shape, position, parts, color, texture, etc. All differences between a given concrete example and any other concrete example are differences in quality. Also, anything we do to change the example we start with is a change in quality. We can make the pencil shorter, break the point, paint it, change its position, and so forth. Each change is related to a quality of the original example.

The assumption that the learner mechanism learns qualities means simply that if an example possesses a quality, no matter how subtle, the mechanism has the capacity to learn that quality. The only factor that limits the learner mechanism is the acuity of the sensory mechanism that receives information about qualities. This mechanism, however, is capable of learning qualities as subtle as the unique tone of a particular violin or qualities that involve the correlation of events (such as the relationship of events on the sun to weather on the earth).

2. The Capacity to Generalize on the Basis of Sameness of Quality

Attribute 1 above indicates what the learner is capable of learning. Attribute 2 suggests *how* learning occurs. According to this attribute, the learning mechanism somehow "makes up a rule" that indicates which qualities are common to the set of examples presented to teach a concept. By using this rule, the mechanism classifies new examples as either positive examples of the concept or negative examples. A new example is positive if it has the same quality(ies) possessed by all the positive examples presented earlier. It is a negative example if it does not have the same quality(ies).

According to the assumption about the generalization attribute, there is no sharp line between initial learning and generalization. The rule-construction of the learning mechanism is assumed to begin as soon as examples are presented. In formulating a rule, the mechanism does nothing more than "note" *sameness of quality.* Once the mechanism "has determined" what is the same about the examples of a particular concept, generalization occurs. The only possible basis for generalization is sameness of quality. If the example to which the learner is to generalize is not the same as the earlier examples with respect to specific qualities it is impossible for generalization to occur unless the learning mechanism is empowered with magical properties.

A further implication of attribute 2 is that the generalizations the learning mechanism achieves are completely explained in terms of the examples presented to the learner and the qualities that are common to these examples.

Table 1.1 illustrates how the learning of **conservation of substance** is the same as the learning of **red.**

Table 1.1	
Learning of a Cognitive Operation (e.g., Conservation of Substance)	**Learning of Red**
Before exposure to examples, the learner has no knowledge of concept.	Before exposure to examples, the learner has no knowledge of concept.
Only some possible examples are examples of this concept.	Only some possible examples are examples of this concept.
The learner demonstrates mastery of concept by treating selected concrete examples of the concept in specified ways.	The learner demonstrates mastery of concept by treating selected concrete examples of the concept in specified ways.
The learner generalizes to new examples of the concept.	The learner generalizes to new examples of the concept.
The appropriate generalizations are to examples that possess the quality of the concept.	The appropriate generalizations are to examples that possess the quality of the concept.

Both concepts are learned in the same way–through a communication from the environment that shows the nature of the concept. The only difference is *what* is learned. And the “whatness” is the quality that comes from the examples, not the learner. For the learner to learn these diversely different qualities, the learner must have the ability to detect both the quality of **redness** and the quality common to the **conservation** examples (e.g., the relationship between changes in appearance and changes in amount).

The Structural Basis for Generalization

The assumptions about the two-attribute learning mechanism imply the type of structure that we must provide to cause specific generalizations. The two-attribute learning mechanism suggests that the learner operates on qualities and sameness, and that both the qualities and samenesses come from the concrete examples that have the same quality and provide information that these concrete examples are the same in a relevant way.

The most general implication of the two-attribute mechanism is the nature of the analysis that we must use for cognitive learning. If the only primary difference between such disparate cognitive skills as learning the color **red** and learning **conservation of substance** is the quality that is to be learned, and if the quality comes from concrete examples (and not from the learner), the primary analysis of cognitive learning must be an analysis of *qualities of examples and of the communications that present these qualities to the learner.* This analysis focuses on the stimuli that the learner receives. We refer to this analysis as the *stimulus-locus* analysis (which is developed further in this and subsequent chapters).

More specific implications of the two-attribute learning mechanisms suggest the general parameters of a communication that is capable of inducing a particular generalization. This communication must meet these structural conditions:

1. The set of positive examples presented through the communication must possess one and only one distinguishing quality. If we assume that the learner learns qualities that are presented through examples, we must make sure that the set of examples presented demonstrates *only one identifiable sameness in quality*–not more than one. If every positive example in the set that is presented to the learner possesses two distinct qualities, at least two distinct generalizations are implied by the communication. Since one of these generalizations is inappropriate, the set of examples does not meet the structural conditions necessary for inducing the intended generalization. For instance, if every example of **red** presented to the learner was a circle and every example that was **not-red** was box-shaped, at least two generalizations are implied by the same communication. Possibly the learner will generalize according to sameness in shape (calling any circle “red” regardless of shape). Both generalizations are possible because both are based on the qualities and samenesses shown by the demonstration examples. Since a given learner is assumed to have no preknowledge of the concept and must base the generalization solely on the quality and sameness of demonstrated examples, a given learner may learn an inappropriate generalization from the demonstration of red circles.

To avoid this problem, we must eliminate the inappropriate quality from the demonstration examples. Different techniques are possible for achieving this goal;

however, the simplest is to modify the set of examples so that some of the examples identified by the teacher as "not red" are circles. With this modification, the set does not present *circularity* as a distinguishing quality of the positive examples.

2. The communication must also provide a signal that accompanies each example that has the quality to be generalized. This signal is the only means we have for *treating examples in the same way.* When we present examples that are physically different (such as two examples of **red** that are not the same shade) we must use some form of signal to tell the learner, in effect, that these examples are the same and that the learner must discover how they are the same. The signal, typically a behavior such as saying "red" for all examples that are red, also provides the learner with a basis for communicating with us. The learner can use the same signal, "red," to let us know which generalization examples have the quality of **redness.**

The assumption about the signal accompanying the various examples is necessary because our goal is to induce a particular generalization. However, if we simply present a group of examples that share a particular quality, we cannot guarantee that: (a) the learner will attend to the common quality; or (b) we will be able to communicate about this quality, even if the learner does attend to it. For instance, if we present a group of objects that are red, how do we know that the learner is attending to the sameness in the quality these examples share? We face other problems if we wish to test the learner to see if the generalization was induced. How does the learner indicate which generalization examples have the quality? Unless we use some signal to suggest sameness (such as putting all red objects in one place or calling them "red" or associating some other unique signal with each example), we cannot demonstrate sameness; we cannot test sameness; and we cannot correct the learner who responds inappropriately.

For the most basic type of communication, two signals are implied. One is used for examples that have the quality. Another is used for examples that do not have the quality.

3. The communication must present a range of examples that show the physical variation of the examples that exhibit a common quality. If every example that the communication presents to the learner is exactly the same shade of red, the communication does not provide adequate information about the range of variation in the quality that is to be labeled as "red." Since this demonstration does not imply that other shades of red share the quality of **redness**, the communication is incapable of inducing the appropriate generalization to examples of other shades of **red.**

To show the quality that is to be generalized, the communication must demonstrate (through examples) the range of variation that typifies the concept. In other words, the communication must present positive examples that are physically different, but that share the quality that is to be generalized.

The requirement of showing a range of positive variation derives directly from our assumptions about the learning mechanism. We assume that the learner is capable of learning any quality exemplified through examples. For most concepts, the quality is something that is common to variations that are physically different. We assume that the learner has the capacity to make up a "rule" about this range of variation. We further assume that if we do not show an appropriate range of variation, the learner is not provided with the information that is necessary to formulate the appropriate "rule." Therefore, if the communication fails to demonstrate the range, the learner cannot be expected to generalize appropriately.

4. A basic communication must present negative examples to show the limits of the variation in quality that is permissible for a given concept. If we show the learner a range of red examples that differ in shades of redness, the communication may appropriately induce a generalization to new examples that are red. (The learner with the two-attribute learning mechanism should appropriately classify any example that falls within the demonstrated range of variation as "red.") However, this communication does not show the boundaries for the generalization, which means that on a test of generalization, the learner may call pink examples "red."

To show the learner basic concepts, the communication must demonstrate the boundaries for the range of permissible generalization. All negatives presented to demonstrate the limits of permissible variation are the same in that they possess the quality of being "not red." To signal that these negative examples are the same, a common behavior is presented with each example. To assure that the learner does not classify these examples in the same way that the positive examples are classified, the communication presents a different signal for the negatives (for example, “not red”). The basic communication, therefore, presents *two* sets of examples (one for the positives and one for the negatives) and two distinct signals (one to signal each positive and the other to signal each negative).

5. The communication must provide a test to assure that the learner has received the information provided by the communication. The test should present positive examples and negative examples that had not been demonstrated earlier, but that are implied by the range of variation of quality demonstrated for the positives and the negatives. If the learner has formulated an appropriate “rule” for the quality that had been demonstrated through the demonstration examples, the learner should be able to respond appropriately to new examples that fall within the range of variation previously demonstrated. A variation of the same signals that are used to demonstrate positive and negative examples is used when the generalization examples are tested.

In summary, the two-attribute learning mechanism implies that a communication for basic concepts must meet these structural requirements.

1. The communication must present a set of examples that are the same with respect to one and only one distinguishing quality (the quality that is to serve as the basis for generalization).
2. The communication must provide two signals–one for every example that possesses the quality that is to be generalized, the second to signal every example that does not have this quality.
3. The communication must demonstrate a range of variation for the positive examples (to induce a rule that is appropriate for classifying new examples on the basis of sameness).
4. The communication must show the limits of permissible variation by presenting negative examples.
5. The communication must provide a test of generalization that involves new examples that fall within the range of quality variation demonstrated earlier.

Analyzing Whether Communications are Faultless

In addition to serving as guidelines for creating faultless communications, the five points above provide the basis for analyzing communications to determine whether they are faultless. The primary analysis for the communication involves no reference to a particular learner. The analysis does not deal with empirical information, but with the structural basis for generalization that is provided by the communication. A communication is judged faultless if it meets the five structural requirements outlined above. The set of examples presented to the learner must be unambiguous about the quality that is to be generalized. The examples must be designed so that only one quality is unique to all positive examples. The range of positive variation exemplified by the set of demonstration examples must be sufficient to imply the appropriate generalization. The negatives should be precise in demonstrating the boundaries of a permissible generalization. The signals presented with the examples must unambiguously provide the basis for classifying examples as either positives or negatives. The test of generalization that is presented as part of the communication must assure that the learner appropriately responds to new positive and negative examples that are clearly implied by the set of demonstration examples. In summary, the communication is judged faultless if it adequately provides the learner with information about *quality* and *sameness*.

The structural requirements that must be met if a communication is to be judged faultless do not refer to specific techniques that are used to correct an inappropriate communication or to design one efficiently. However, these techniques (which are discussed in later chapters) follow from the structural requirements. If we understand that a communication must show that a particular quality is unique to the positive examples, we will investigate possible techniques that achieve this goal. From the possibilities we will select those that are most efficient and those that show the uniqueness most emphatically. Similarly, the design of the test examples can be reduced to some how-to-do-it formula once we understand what the test examples must do.

The five structural requirements derive directly from our assumptions about the learner. We can appreciate the implications of the two-attribute learning mechanism by considering how the structural basis for a generalization would change if we changed our assumptions about the learning mechanism. For instance, if we assumed that the learner generalized on the basis of *similarity,* not sameness, we would not be provided with a strict standard about whether the communication that we design presents examples that are *"similar."* The notion of similarity is not precise and it begs the question of how the *"similar"* examples are the same. If examples are similar, they must be the same with respect to some quality, but the notion of similarity does not require us to identify this qualitative sameness. Therefore, similarity leads to an imprecise standard for evaluating our communication. By assuming that the learner generalizes only on the basis of sameness, we are required to create examples that are the *same* in some identifiable way, and the standard we use is objectively stronger.

If we assumed that the learner's generalizations are not clearly determined by the common quality of the concrete examples of a concept, we would not be provided with a standard for judging whether a communication adequately shows both the quality and the range of variation in the quality across various examples. We might assume that the generalization would occur simply if the learner received some "exposure" to the concept. But we would not have any analytical yardsticks for determining whether the "exposure" presented through a particular communication was adequate.

With the assumed two-attribute learning mechanism, however, we are provided both with general guidelines for creating structures that will induce specific generalizations, and with more specific implications about what the communication must do and what it must avoid doing.

Predictions of Generalizations

The procedure for determining flaws in a communication is a logical one, based on observable details of the communication. The procedure therefore permits us to make predictions about what the learner will learn. These predictions are independent of the learner. The basic form that these predictions take is that if the communication is flawless (adequately meets the five structural requirements), the learner will learn the generalization that is conveyed through the communication. The learner will respond appropriately to the examples that test the generalization and will respond to additional examples that are implied by the demonstration examples. Conversely, if the communication has flaws, some learners who receive this communication will learn the inappropriate quality demonstrated by the flawed aspect of the communication.

Equally important, the development of procedures for determining whether a communication is faultless permits us to engage in a very precise study of the learner. A faultless communication serves as a *standard* against which we compare the learner's performance. If this communication is analytically faultless (with respect to clarity in communicating one and only one possible generalization), any learner who possesses the two-attribute learning mechanism will learn the concept that is presented by the communication. If a learner does not perform in the predicted manner, we immediately know three things about that learner:

1. We know that the learner does not have (or is not using) the two-attribute mechanism.
2. We know the precise ways that the learner's performance deviated from the predicted performance.
3. Because we know that the problem resides with the learner and not with the communication (which is judged faultless), and because we know precisely how the learner has deviated from the predicted standard, we know how we must modify the learner so that the learner is capable of performing acceptably in response to the communication.

We are able to make these strong inferences about the learner because we have ruled out the possibility that the learner's poor performance can be accounted for by the presentation. Furthermore (as we observed earlier), we would not be able to draw precise conclusions about the learner unless we ruled out the possibility that the communication has flaws and that the learner is responding in a logically reasonable way to the flawed communication. If the learner generalized to circles following the communication that presented circularity as a quality common to all positive examples of **red**, we would be presumptuous if we interpreted this generalization as an indication of a "faulty" learning mechanism. Only if the communication is faultless can we make strong inferences about the learner.

Stimulus-Locus and Response-Locus Analyses

Although the major goal of this book is to describe procedures for designing effective instructional communications, not to study the learner's behavior, the procedure that we use parallels the one that we would use to study the learner. We use two analyses. The primary analysis is a *stimulus-locus analysis,* which deals with an analysis of the stimuli or communications the learner receives. The second analysis is the *response-locus analysis,* which focuses on the learner. This analysis comes into play if the learner is unable (for whatever reason) to produce the responses that are called for by the communication. The response-locus analysis consists of techniques for modifying the learner's capacity to produce responses. If the learner does not respond in the predicted manner to a faultless communication, the assumed "fault" lies not with the communication, but with the learner. Therefore, we must switch our focus. This switch involves a complete change in orientation, from a concern with the analyses of communicating quality and sameness in a precise manner, to the laws of behavior. These laws provide us with specific guides about the amount of practice, the massing and distribution of trials, the schedules of reinforcement, and other variables that cause the growth or strengthening of the learner's response to take place. For example, if the learner apparently forgets the word **red** and cannot respond to various examples in a faultless presentation that asks the learner, "What color is this?", we modify the learner's capacity to "remember" how to produce the name. When the learner reliably remembers the words, we return to the original communication. The learner is now assumed to be an adequate receiver, capable of responding according to the predictions of the stimulus-locus analysis.

The basic difference between the response-locus analysis and the stimulus-locus analysis is that the stimulus-locus analysis does not involve the learner. It involves the logic of *ruling out all the possibilities but the one* to be conveyed through a teaching communication. The response-locus analysis is based on empirical findings on learning.

When instruction in skills involves teaching new responses (those the learner has never produced before in response to any signal), we use the stimulus-locus analysis to design the sequence of skills to minimize possible conceptual confusion. We also use response-locus techniques to assure that the new responses are induced efficiently. However, even for the teaching of "motor skills" (such as shoe-tying, ball-throwing, etc.), the stimulus-locus analysis is the primary one. The reason is that the communications must be clear and must be organized so that the appropriate generalizations are induced and the appropriate response generalizations are implied. These communications, however, rely heavily on the application of behavioral principles.

Extending the Stimulus-Locus Analysis to Types of Knowledge

If we follow the stimulus-locus assumptions to their conclusion, we discover that knowledge may be classified according to the *samenesses of communications* used to teach various concepts. The samenesses in features of the communication parallel samenesses in the concepts that are to be taught. Viewed differently, the extent to which concepts are the same provides a precise measure of the extent to which faultless communications for these concepts may have the same features or attributes. Let's say that we design a faultless presentation for a particular concept. The communication isolates the quality presented to the learner, unambiguously signals the quality through examples, and provides additional examples for testing the learner's generalizations. To design a faultless communication for a concept that is highly similar to the original one, we would create a communication that is highly similar to the original one. The close logical parallel between the structure of the concepts we wish to teach and the structure of the communications that convey these concepts faultlessly results because the two concepts are the same with respect to many qualities. The samenesses in quality of the concepts is reflected in the samenesses in the communications that convey these qualities. Conversely, if two concepts differ in many ways, the faultless communications that communicate them will have many differences.

By extending the notion of the parallel between the structure of concepts and the structure of communications that convey them faultlessly, we are provided with general guidelines for creating *classes of cognitive skills.* For this classification, each category consists of concepts or skills that are the same with respect to important *structural* features. *Since the concepts in each category share samenesses, all concepts within a given category can be processed* *through simple variations or transformations of the same basic communication or form*. To classify a concept within this system is to be provided with an algorithm for communicating the concept to a learner who is assumed to possess the two-attribute learning mechanism.

Summary

The design and analysis of communications are based on assumptions about the kind of information the learner is capable of extracting from the communication. For analytical purposes, we postulated a learning mechanism that has these attributes: the capacity to learn any stimulus quality shown through examples, and the capacity to generalize a sameness of quality to new examples. This assumed mechanism implies that the primary analysis of cognitive learning must focus on quality and sameness of the examples presented to the learner. Further implications suggest the structural criteria that must be met by a communication if the communication is to induce a generalization for a basic concept.

1. The positive examples of the concept must be distinguished by one and only one quality.
2. An unambiguous signal must accompany each positive example, and a different signal must accompany each negative example.
3. The examples must demonstrate the range of variation to which the learner will be expected to generalize.
4. Negative examples must clearly show the boundaries of permissible positive variation.
5. Test examples, different from those presented to demonstrate the concept, assure that the generalization has occurred.

These criteria serve as guidelines for designing faultless communications and for determining whether a particular communication is faultless. The analysis of communications according to the structural features of the communication is the stimulus-locus analysis. The stimulus-locus analysis assumes that the learner is a “receiver” capable of attending to the information presented through a “faultless” communication. However, a particular learner may not learn in the predicted manner. The difference between the learner’s actual performance and that predicted by the stimulus-locus analysis suggests the extent to which the learner does not respond to the basic logic of the communication (the logic of quality and sameness). If the learner is incapable of producing responses that are implied by the stimulus-locus analysis, our focus shifts from the stimulus-locus analysis to the response-locus analysis. Behavioral principles are used to induce new responses and to maintain responses.