Research on Direct Instruction: 25 Years Beyond DISTAR

Gary L. Adams
Siegfried Engelmann

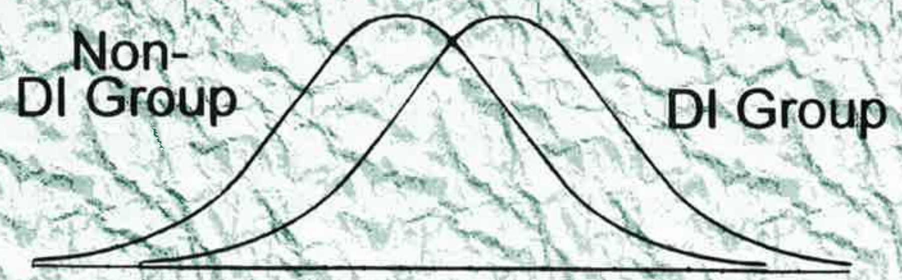

Published by Educational Achievement Systems
319 Nickerson St. - Suite 112
Seattle, WA 98109
Phone or Fax: (206) 820-6111

ISBN: 0-675-21014-3

CHAPTER 2

Features of DI Instructional Programs

Direct Instruction consistently produces superior performance results when it is implemented appropriately. Any instructional system that performs in this manner must have a standardized delivery system that permits the system to accommodate teachers and students of various ability levels and to adapt to a variety of instructional contexts. The seven criteria listed below suggest major features the system must have.

1. Would teachers, even those with below average teaching skills, be able to teach the program successfully after receiving relatively small amounts of training?
2. Does the program permit reliable predictions about how much student progress may be anticipated for a given time period?
3. Is the sum of the "promised" skills relatively substantial compared to the sum of skills that is currently mastered by students during the same period of time?
4. Is there an analytical basis to suggest that these gains are at least plausible?
5. Are there sufficient tests of student performance to serve as a guide for adjusting the rate of presentation to students?
6. Do priority skills receive relatively more instructional attention than trivial skills?
7. Is there consumer protection information to suggest that the outcomes are possible?

These criteria provide a yardstick for evaluating instructional delivery systems. The extent to which they fail to meet the criteria is the extent to which they could not be implemented uniformly across different teachers and different populations of students.

1. Would below-average teachers be able to teach the program successfully after receiving modest training?

In the ideal situation, a teacher would be able to use the program to produce consistent, anticipated results without needing any training. A brief orientation to the material would be sufficient. This ideal is not possible. To teach effectively, a teacher must be an adequate presenter, motivator, and source of feedback and corrections. Because these skills have not been mastered by many certified teachers, initial training requires training on these elements which are not unique to the particular program but that simply involve details of teaching.

In the most efficient system, relatively extensive training should occur before the first year of teaching. The teacher should receive practice with the various skills and an orientation to the program. Adequate training may require one week. Follow-up sessions throughout the year would provide additional practice and information about the instructional system. During the second year, the teacher would generally need far less training.

A program that requires extensive training beyond general teaching techniques would be very costly and would probably fail with a large percentage of teachers. For example, there are very effective practices for teaching students (such as heuristic approaches that require teachers to make precise responses to what students say or write). These programs frequently fail both because they require too much training and monitoring and because a large percentage of teachers fail to master the techniques. Therefore, they are not replicable. On the other hand, if the program provides glib training (a half-day general orientation or activities that stray far from the details of what the teacher will do in the classroom), the program will not be uniformly effective across teachers.

Direct Instruction Solution

The Follow Through Project and other major implementations provided beginning teachers with one week of training, followed by in-class monitoring and inservice sessions throughout the first year of the implementation. The results produced by teachers, who implemented properly (according to the judgments of supervisors), confirmed that this training is adequate for teachers with average to below average teaching proficiency (Meyer, Gersten, & Gutkin, 1983).

2. Does the program permit reliable predictions about how much student progress may be anticipated for a given time period?

This criterion does not imply that all students progress at the same rate or that the program somehow obliterates individual differences. It implies that when the teacher appropriately place students in an instructional sequence, the teacher knows what performance should be expected by the students.

The teacher should be able to say:

> *If I follow the program as written, I will devote 30 minutes to instruction for the following 80 school days. At that time, my highest group (the one that is currently the highest performing) will have gone through more than 80 lessons and will have mastered all the skills taught in that lesson range. My middle group will have gone through 80 lessons and will have mastered all the skills through Lesson 80. My lowest group will have gone about through Lesson 55 and will have mastered all the skills taught through Lesson 55.*

Unless such predictions can be made and validated, the curriculum is not manageable across teachers or students.

If such predictions can be made and validated, the program addresses the question of individual differences. Also, the greater the extent to which different teachers successfully predict student performance, the better the program is in compensating for individual differences among teachers.

Direct Instruction Solutions

This criterion addresses standardization issues: How can the critical aspects of the presentation be standardized so they may be effectively presented by any teacher? How can the program be designed in a manner that permits clear references to clock time and calendar time?

The answers lie in the way students are grouped for instruction, the general design of the lessons, and the directions the teacher follows when presenting the lessons.

Homogeneous Grouping

Predictions of student performance are based on the assumption that students master the material presented during a given time span. The first requirement for mastery is to group students homogeneously for instruction. Even within a homogeneous group, there will be variations in skills. The teacher is directed to teach to the lowest performers in the group, ensuring that *all* students achieve mastery. Groups are not heterogeneous because teaching all students in a heterogeneous group to mastery is impractical. If the teacher teaches to the lowest performers, higher performers suffer because they already know the material being taught and they are not engaged in instruction that is appropriate for their skill level and rate of learning. Teaching to the middle performers in a heterogeneous group overwhelms the lower performers, who will not learn the material as intended but will learn expectations of failure. (These students will have a low self-image because they will receive repeated demonstrations that they fail on the various tasks their classmates are able to master.) Teaching to the lowest performers slows the presentation to a rate that is inappropriate for probably 80% or more of the group.

As a general rule, students who score high on the entry criteria for a DI program are grouped together. Those who perform at the low end of the entry criteria form the low group. On the beginning levels, the classroom may have three homogeneous groups. Note that students are not "tracked," but grouped flexibly and are regrouped as performance changes. Also, students are not intended to be grouped homogeneously during the entire day, but only during specific periods of instruction. For example, a student in the middle group for reading may be in the high group for math. The reason for grouping students is not to deprive them of important experiences, but rather to make the learning experience as appropriate as possible to each student.

The Lesson

In contrast to "lessons" in traditional programs, which may take between one and four class periods, Direct Instruction lessons have a specified time during which students are expected to master all material for that lesson. For some programs, lessons are 30 minutes, for others, 40 to 45 minutes. This time is determined largely by field testing the programs.

The times allotted for lessons translate into expectations about student performance or rate of student mastery. The basic unit of rate is the lesson. The basic measure of rate is the amount of time required for mastery.

To determine how much students will learn in 70 lessons, one can examine Lesson 70 and see what the students will be doing. What they are doing is what they are mastering. One can figure out the number of school days that will be required for a group to progress 70 lessons: high groups master about 1.3 lessons per instructional day; middle groups about 1 lesson a day; and lower groups about .7 lesson per day. The predictions are not perfect because lessons are not perfect. All DI programs tend to have some "long" lessons and some "short" ones. **Despite this less-than-perfect correspondence between clock time and student performance, the anticipated length of time required for lessons is precise enough to permit expectations of mastery that are based on the number of school days.**

These predictions are important for the management of an effective school. If a beginning reading group consists of average performers, the expectation is that by the end of the school year, the group will have completed one level of the beginning reading program, which means that all the students will have mastered the skills presented in the entire program. No students will fail to learn to read. No students will lag seriously behind the rest of the group. Short-term expectations permit a supervisor to quickly check on the progress of the group. The expectation for the number of lessons the group will master has been established. If a teacher is not moving through the program at a rate that is close to the expected rate, there is a problem with the delivery system. If the rate is acceptable but students in the group are not mastering the various skills, either the teacher is not presenting adequately or the students who are lagging are not placed appropriately and should be moved to a group that has a more appropriate lesson pace. In either case, the problem is relatively easy to spot, and the remedy is clearly implied by the performance of the students.

Scripted Presentations

The final feature of DI programs that permits predictions of student performance is the way the programs provide directions to the teacher. The directions are never general. The teacher is not told to conduct a brainstorming session or to "discuss" a broad topic. The teacher is not told general strategies for correcting or placing students. Instead, all directions are specific. Those that address what the teacher is to do when presenting the content of the lessons are in the form of a script that provides the exact wording the teacher is to use.

The rationale for the scripted presentations is that if the teacher presents an adequate set of examples with clear, consistent wording, students will learn the material with less confusion. The program provides the various examples and the wording. Therefore, if the teacher presents the examples and the wording, the students should master the content.

The notion of scripted presentations is not popular because it smacks of "roteness." However, observations of teachers show that variability in teacher wording creates confusion among lower-performing or naive students and that teachers frequently have wording problems when they try to explain details to students or refer to the same details on subsequent days (Carnine, 1980c).

Furthermore, general directions put the teacher at a disadvantage because the teacher must often generate the examples for the students, create the wording, and respond to difficulties or mistakes that the students make. The teacher is being asked to do far more than could be reasonably expected of anybody who is less than an expert teacher.

Another advantage of the scripted presentations is that Direct Instruction teachers do not have to create the details of instruction through elaborate lesson plans; teachers just follow the script. Usually, the result is mastery.

Summary: Progress can be anticipated for appropriately-placed students by providing the specified amount of instruction (30 minutes a day, or whatever the program specifications indicate) and by following the conventions specified by the program. The program structure makes it possible to predict the progress of each group of students for a specific number of school days. Predictions for higher performers will be for the mastery of about 1.3 lessons per day. Middle performers will master 1 lesson per day, and lower performers will master about .7 lesson per day. Students who have completed Lesson 136 will have mastered all the material on that and on all preceding lessons. Given that this teaching is achievable, great quality control and standardization is provided through the lesson structure.

3. Is the sum of the "promised" skills relatively substantial compared to the sum of skills currently mastered by students during the same period of time?

The answer should be "yes." Ideally, both higher- and lower-performing students would learn skills that comparable students have not learned before, learn them in relatively less clock time, and generalize them more accurately and productively. Although lower-performing students would not achieve the same level of achievement as the higher performers, they would tend to show significant "gain."

Direct Instruction Solutions

Because of their design, Direct Instruction programs have a unique feature: **an examination of the material shows all the skills that will be learned by the students.** Although traditional textbook analyses and evaluations of instructional programs assume that this feature is true of all instructional programs, it obviously is not. If it were true of traditional basals, they would represent a paradox, because they tend to present the same material on subsequent levels of instructional programs. In math, for instance, students are taught fraction analysis in grade 2 or 3. Some of the same work is repeated in grades 4, 5, and 6 (with an annual overlap of as much as 70% of the same material). Yet, it is not uncommon to find 7th and 8th graders who are grossly deficient in their understanding of fractions (for example, the ability to indicate whether fractions are more than 1, and what constitutes equivalent fractions and why).

Traditional basals are not designed to teach, rather to serve as something like a resource for the teacher, with activities and general suggestions about how to present the material to the students. Examination of traditional material, therefore, provides only information about the various topics that will be presented to the students, not information about students' learning.

Evaluating the Potential of Material

If Direct Instruction programs are compared with traditional programs through inspection of material (which is what is typically done in textbook adoptions), the conclusion might well be that traditional programs present a larger variety of activities than the DI program, but inspection of the traditional program tells very little about how much of this material students will learn. Any system dedicated to the acceleration of students' performance would not rely on inspection as a basis for judging a program's potential, but on data and hands-on evidence about what works well and what doesn't.

The DI orientation is that progress using Direct Instruction programs can be predicted according to how well the teacher follows the program and presents the material.

Correlations of teacher behavior and student performance have been recorded, and the predictions have been verified. Specifically, the midyear performances of teachers as rated by the supervisors (based on how well teachers followed the program specifications) predicted end-of-year student performance very accurately (correlation between teacher behavior and student achievement of around .8) (Meyer, Gersten, & Gutkin, 1983).

4. Is there an analytical basis to suggest that these gains are at least plausible?

Lesson-by-lesson inspection of a program should identify the opportunities students have to learn and practice each of the targeted skills, concepts, or behaviors. If the program merely exposes students to explanations, problem types, etc. and does not provide systematic review and follow-up over a period of days, there is little basis to believe that the students will uniformly learn the content.

If the design of the program involves traditional units (with a few consecutive periods devoted to a particular topic, which does not reappear for possibly 60 school days or more) the analysis would not support the belief that students have sufficient opportunity to learn the scheduled topics. When the new material is massed in this manner (students working the entire period on the new topic), mastery suffers. Students retain information better when practice in learning and using it is distributed over time. If the same information that is presented in one large block is distributed over four smaller blocks that have a cumulative review, the likelihood of students learning and retaining increases. As a general rule, students will not retain information unless it appears on two consecutive occasions and is then reviewed.

If the presentation involves introduction of many new skills and relationships at the same time (rather than a systematic trickle of new information), the promised skills will not occur for many students, particularly the low performers who are not well practiced at learning new information at a high rate or retaining what they have learned.

The program should also have the support of a theoretical underpinning and empirical evidence.

Direct Instruction Solutions

Designing Direct Instruction materials involves both a unique analysis of the content to be taught and a method of validating or field testing the material. The analytical basis is articulated in the book *Theory of Instruction,* by Engelmann and Carnine (1991). It indicates the basis for inducing generalizations, discriminations, and complex behavioral changes or problem-solving steps. The theory makes only four assumptions about the student:

(a) The student is a capable receiver of information (and is capable of abstracting "features" from examples);

(b) the student generalizes on the basis of sameness of features and only on the basis of sameness;

(c) the student correlates information about sameness of features in perfectly logical ways; and

(d) what the student learns is always consistent with the teaching demonstrations and information the student receives (whether the information is intentional or not).

The student's memory and feature-abstraction capacity improves with practice. The assumption is that poor memory or poor ability to abstract features from concrete examples indicates only that the student is naive (hasn't been taught), not that the student is incapable of learning concepts or discrimination.

The major derivation from the assumptions about the student is the rule about demonstrations that are designed to teach: **if the presentation is consistent with one-and-only-one interpretation, the student will learn the concept that is being presented.** If the presentation is consistent with more than one interpretation, the student will learn one of the interpretations that is consistent with the presentation, but not necessarily the intended one.

This rule means simply that the teacher may think she is teaching one thing, but may actually be teaching others. If the presentation is consistent with three interpretations, the student may learn any one of the three. Furthermore, the "other possibilities" are specifiable, based on what the teacher does.

The single-interpretation rule further implies that the teacher or the teaching presentation is responsible for what the student learns. If the student fails to learn the intended concept, the reason is not a flaw in the student (for example, an orientation to use one side of the brain rather than the other). The cause of student failure is a flawed presentation—one that is consistent with more than one possible interpretation.

The rule that what the student learns is consistent with the teaching presentations the student receives does not mean that all students will learn the same thing from the same flawed presentation. If the presentation is consistent with three interpretations, the student may learn any one of the three.

Historically, the rule about one-and-only-one interpretation has been very difficult for teachers and researchers to appreciate. This rule implies that if a student fails to learn, the diagnosis of the problem would not focus on the student as much as on what instructional exposure and practice the student received. The currently dominant orientation to student learning assumes that if the student fails to learn, the student has a learning problem. This diagnosis provides no detailed critique of the sequence of activities the teacher provided in attempting to teach the desired skills or operations.

Analyzing Teaching

Any teaching demonstration or activity that is designed to teach a concept must meet certain requirements. It must present examples of the concept and it must provide some sort of label of the concept being presented. Also, it must provide some opportunity for students to respond and give behavioral indications of whether they have learned the label and can appropriately link it with examples of the concept.

A rigorous method permits one to identify interpretations in the teaching demonstration. An interpretation refers to features of the presentation that are present in every example labeled in a particular way and not present in any examples that are not labeled that way. Let's say a teacher taught a naive student "red" by presenting three red triangles on a feltboard. If the naive student, who had no experience with feltboards, repeatedly received this demonstration (and only this exposure to "red"), the student could learn any of the following meanings of the word "red":

1. The student would identify anything on a feltboard as "red" or
2. the student would identify triangles as "red" or
3. the student would identify only red triangles as "red" (not red squares, blue triangles, or any other variation) or
4. the student would identify feltboard triangles as "red," regardless of color or
5. the student would identify red objects on the feltboard as "red," regardless of shape or color or
6. the student would learn various combinations that are consistent with the information provided by the initial teaching demonstration.

(All these interpretations are based on the features shared by all feltboard labeled as "red.")

There are several important points about this orientation:
1. Without feedback from human or "unnatural" sources, the student would never learn the correct meaning of "red." Interactions with the physical environment provide no possible basis for correcting the student. No contradictions result if the student thinks that "red" means triangle, for example.

2. For the student to learn any of the interpretations, the student must have the ability to correlate sameness–the same label with the same observed features of the examples. This is not a vague and unspecified association, nor a response to some unspecified similarity between the examples. It is a precise correlation between the name and observed features that are present in all the examples. While this "correlating" ability may be only modestly developed in students who are "retarded," it characterizes their behavior as well as that of more sophisticated students. Basically, the student plays a "find the rule" game. A verbalized version of the student's procedure goes something like this: "When the teacher provides the same label for more than one object, her behavior signals that the objects are the same in some other way. Find that way. Any sameness that is unique to the labeled objects is a possibility—find the right sameness."

The student, in other words, will always learn what is taught. The only question is whether the teaching is intentional.

If the demonstration is consistent with only one interpretation, the student will learn the intended concept. The one-and-only-one interpretation is an ideal and is probably only approximated in Direct Instruction programs. However, it is attempted and its influences are revealed by the initial teaching of any skill or concept in Direct Instruction programs. When analyzing these presentations, the analyst should always assume the orientation of the naive student who does not know the material that is being taught. The beginning reader, for example, doesn't know how to read and may learn some serious misrules about reading if the presentation is not very clear about what reading is and how it differs from recitation, story telling, or guessing. The presentation that does not make this difference clear will induce various misrules, all of which are observed in the remedial reader.

3. The teacher who uses an inappropriate set of teaching examples receives no information that the student has identified the wrong correlated sameness. For instance, following the flawed teaching of "red" the teacher displays four red triangular shapes on the feltboard and tells Jimmy, "Touch something that is red." The student touches one of the red objects and the teacher says, "Good." The teacher receives no information about whether the student attended to the redness,

the triangularity, the "feltness," or some combination of these features. In effect, the teacher has reinforced whatever interpretation Jimmy derived from the teaching presentation, right or wrong.

Generalizations

According to this interpretation of learning, the student does not first learn something in a concrete singular sense and then generalizes to some larger set. The initial learning is a generalization.

The evidence for this assertion is the fact that the student is able to respond predictably to objects other than those presented in the initial teaching set. For instance, the teacher could follow the flawed feltboard demonstration involving the red triangles by presenting other objects ("red" and "not red") and asking the student to respond to those that are red. "Touch something that is red." The student will respond. This behavior, however, provides a serious challenge to the notion that the student first learns and then generalizes. The student responds to new examples; therefore, the student generalizes. Furthermore, it is logically impossible for the student to be able to correlate the label with the examples without being able to formulate some clear mental (but not verbal) rules about the examples. Even if the student derives the rule that red refers only to the specific objects that teacher presents on the feltboard, the student would have to register the particular features of the objects in the set. (Without this knowledge, the student couldn't possibly know whether an object presented was one that had been presented earlier.)

Initial Learning versus Relearning

Research on relearning shows that students who initially mislearn require substantially more practice to relearn the concepts than they would have if they had learned it correctly. Therefore, a very important premise of sympathetic instruction should be to structure the learning so the student does not have to relearn concepts or operations. Practices that lead to mislearning include reinforcing incorrect responses, not providing clear feedback about the student's performance, and requiring parents to correct learning problems. (For example, homework is often provided to young students. The homework frequently implies that parents must teach what the teacher has failed to teach.)

Clear Demonstrations

The procedure for designing teaching demonstrations that are consistent with only one interpretation is to use negative examples to rule out possible misinterpretations. The ultimate pairs of examples for "ruling out" misinterpretations are exactly the same except for one detail. A pair of examples that show what red is, for example, would be the same shape, same size, but one would not be red. The teacher labels one as "red" and the other as "not red." Since the only difference between the examples is the color, the only possible meaning of red is the particular color property of the example labeled as "red." Other details of the presentation assure that there are enough examples to suggest the range for the concept.

Examination of Direct Instruction programs discloses that new learning is addressed with a presentation that at least approximates the one-interpretation rule and that whatever is introduced appears in subsequent lessons, suggesting sufficient opportunities for the student to practice new skills or discriminations.

5. Are there sufficient tests of student performance to serve as a guide for adjusting the rate of presentation to students?

Ideally, testing is part of teaching. These tests are not formal tests, but tasks that students perform and give the teacher reliable performance information. These tests are important because they form the basis for what the teacher is to do. If students pass the tests, the teacher moves on. If the students fail the tests, the teacher repeats the relevant parts of the program and firms the students' form the basis for what the teacher is to do. If students pass the tests, the teacher moves on. If the students fail the tests, the teacher repeats the relevant parts of the program and firms the students' understanding. An effective sequence must test what is taught, test when the skill or concept is taught, and test in the same manner the skill or concept is taught. For instance, the teacher provides information to students about how to solve a particular problem type or how to decode a particular type of word; then the teacher "leads" or "scaffolds" to provide assistance to students as they follow the steps in the procedure; then the teacher immediately provides a test, in which the students work on a new example and carry out whatever overt steps the teacher modeled for working the problem or reading the word.

With information about students' performance, the teacher is able to identify whether or not students have trouble with any of the overt steps, whether they are able to perform without assistance, which part of the process tends to give them trouble, and generally how well the information the teacher presented has been internalized by students.

Note that the testing requirement tends to create stress on the efficiency requirement (Criterion 2 described on page 8), which demands more teaching than is traditionally provided for a given time period. If the tests are presented to all students in turn, the teacher would not be able to achieve efficiency. The teaching would become so disjointed and laborious that the teach-and-test format would be counterproductive. Therefore, the tests must be designed so they do not take much additional time, yet provide the teacher with sufficient information to determine whether additional instruction is needed before moving on in the program sequence.

Also, the tests should be direct rather than indirect. Direct tests are best because they provide for the best diagnosis of why different students fail items. An example of an indirect test of a beginning reader's ability to decode would be the standardized achievement test, which assumes that the student can read the item, read the choices, select the appropriate choice, and mark it. The test is indirect because the only evidence the teacher has is a mark on a piece of paper. From that sample of the student's work the teacher could not possibly determine whether a student's poor performance resulted from not being able to decode the item, not being able to understand a properly decoded item, not being able to decode all of the choices, not being able to understand all of the choices, or not being able to follow the directions for marking the choice. Given that the poor performance could be the result of any combination of the primary possible causes of failure, the test format is not well designed to reveal information about what the teacher has taught. In contrast, a direct test presents the student with words or sentences to read orally.

Direct Instruction Solutions

Direct Instruction teaching uses students' performance as the criterion for judging whether the material has been mastered. The analytical basis for determining that the program has the potential to teach all the content is not a guarantee that a particular student will learn in so many trials or without any problems. While the goal is not to see how fast it is possible to take the student through the sequence, a faster rate is desirable because it suggests the extent to which the student's performance is being accelerated.

All Direct Instruction programs are designed so that students respond to many tasks that give teachers a great deal of information about their students' learning. The tasks may require verbal or some sort of written response. In either case, they occur immediately after the teacher presents an explanation or demonstration. During each lesson of a beginning program, such as a beginning reading program, the teacher is provided with more than 200 responses from students. If used as diagnostic information, these responses provide the teacher with detailed process information about what each student is learning, which things present some difficulty to various students, and whether each student has a good understanding of the most-recently-presented material.

Easy-to-Difficult Contexts

Because Direct Instruction programs present new material so that it is presented first in an easy context and then in increasingly difficult ones, the teacher has the potential to map the progress of students in great detail and observe how much practice and what kind of practice leads to students becoming "firm" on the various skills that are introduced.

The easiest context for learning new discriminations or operations presents juxtaposed instances of the same operation or task. (The student does a variation of the same thing on subsequent tasks.) More difficult contexts involve less juxtaposition prompting. The final, and most difficult, context is that of the "mixed set" of problems, operations, or tasks. All the earlier-taught skills or concepts are reviewed in an unpredictable order. Good student performance on these tasks indicates that the student has learned all the earlier-taught concepts and is able to respond to them in a setting that provides no clues about the order or range of concepts.

Direct Instruction programs sometimes provide prompts, such as a modified orthography for initial reading; however, all prompts are removed as part of the sequence of making the student completely independent on the material that has been presented.

The cumulative reviews that are provided in the program (earlier-taught skills routinely integrated with newer skills in a context that provides no prompts to the student) provide ongoing information about each student's progress.

Responses

The support or prompts provided during initial learning are derived from a different set of principles than those that drive some behavioral programs. The behavioral program may design prompts with the idea that if the response is produced in the presence of the discriminative stimulus, the response will be strengthened. The Direct Instruction orientation does not consider responses critical. The reason is that any responses in the student's repertoire can be used to teach concepts. (Although the teaching would require a number of inventions, it would be possible to teach a person any concept if the person could reliably produce only one voluntary response, such as moving her eyes.)

This perspective on responses is important in evaluating a typical Direct Instruction lesson. While many responses are designed to increase facility with new motor skills and verbal operations (such as pronouncing unfamiliar words or writing unfamiliar statement forms), most responses are tests designed more for feedback than for strengthening responses. The student receives information on what sort of mental energies are needed to remember, organize, and produce patterns of responses. The teacher receives feedback from the student's performance and adjusts the presentation accordingly.

Curriculum-Based Assessment

In-program tests or curriculum-based assessments are provided as part of Direct Instruction programs. In most DI programs, formal tests occur every tenth lesson. They focus on the skills or concepts that have become independent during the preceding ten-lesson period. Remedies for poor student performance require the reteaching of specified exercises (which are usually in the most recent ten-lesson span). In Direct Instruction programs, various lessons are also curriculum-based measures of what the students learn.

6. Do priority skills receive relatively more instructional attention than trivial skills?

This question is probably the most difficult for the evaluator to respond to because it involves knowledge of (a) problems that students frequently have with more advanced applications, (b) general properties of these advanced applications, and (c) an understanding of the steps that would lead to efficient preparation for these applications. For instance, advanced math requires students to be quite proficient at copying and rewriting equations. Preparatory math programs that provide no work in copying problems, rewriting them, and transforming them do not serve well to prepare students for higher math. Not everything students do in effective preparatory programs would require copying and rewriting equations, but copying and rewriting should be a clear priority. Note that copying and rewriting equations is currently not a skill that has been identified by the National Council of Teachers of Mathematics (1991), so it is quite possible that the evaluator of a program wouldn't know about the need and possibly would judge procedures that required copying equations of problems as unnecessary and inefficient.

Weak math instruction is characterized by muddled priorities, indicated by a large percentage of activities that either do not involve calculations or that have no serious implications for higher math. Much of the work that students do with graphs, symmetry, and labeling of geometric figures (octagons, etc.) are questionable activities. Knowledge of right triangles is important because hundreds of inferences and problem types derive from understanding properties of this triangle. An equilateral triangle is not even remotely as pregnant. In fact, students could learn virtually everything they need for higher math or future applications without knowing this label. Similarly, labeling the numerator and denominator of fractions is a convention, not a mathematical necessity, particularly for early work with fractions. Any basic fraction operation can be explained by referring to the top number or the bottom number of the fraction. There is some value for the student to learn the traditional labels, but the question is when. The learning of the labels should certainly not become a substitute for learning the operations.

Understanding priorities is especially important for the initial work on a topic or skill because the beginning is where the problems of communication are most severe. The more a teacher's communication is loaded with nonessential vocabulary, the greater the possibility that students will have problems that result not from being unable to perform the manipulations the teacher describes, but from being confused about the meaning of the new words in the teacher's directions. A related problem is that the more a teacher is required (by curriculum-sequence requirements) to teach low-priority skills and knowledge, the greater the stress placed on a good instructional sequence to teach the high-priority skills effectively. Adequate time may not be available for every topic or concept.

Direct Instruction Solutions

One of the more apparent features of DI is the emphasis of priority skills. In the teaching of beginning reading, for example, DI reading programs have as their first priority the teaching of decoding as a rule-governed procedure that involves manipulating sounds. As part of the preparation for decoding, DI programs present extensive exercises in phonemic awareness or phonological awareness.

The first DI reading program—DISTAR Reading (initially published in 1968)—presented verbal exercises in "say it fast," "say the sounds," and "rhyming." This teaching occurred long before there was any interest or general awareness of phonological awareness in the field of reading. For "say it fast," the teacher would present segmented words: "In the picture you'll see a *mmm aaa nnn*. What will you see?" For "say the sounds," the teacher would present a word that students' segment: "Listen: *man*. What word?_ Say the sounds in *man*. Get ready_" For "rhyming," the teacher would say: "You're going to rhyme with *at*. What are you rhyming with?_ Start with *mmm* and rhyme with *at*. Get ready_"

DISTAR was the only commercial program that had such activities, and surprisingly, with all the current interest in phonological awareness, DISTAR is not cited as either a pioneer or an exemplar of how phonological awareness is effectively integrated with the other skills required for students to decode words.

Direct Instruction programs attempt to introduce models that permit generalizable learning of the core or priority skills. For DI, working ratio problems is a priority skill. In the typical math sequence, however, ratio problems are presented, but not through an articulate model. The link between ratios and fractions is not made clear. The precise procedures for analyzing ratio word problems are general and often of very little use in analyzing the problem, and the work does not receive a high priority.

The following is a somewhat detailed illustration of how ratios are taught in the series *Connecting Math Concepts*. Although the description of events may seem somewhat involved, it covers only the broad details and discriminations. Its purpose is to point out the extent to which the DI sequences account for everything that is to be learned and present models that are capable of teaching the skills.

Before students work on ratio problems, they work with equations that involve equivalent fractions. One of the things they do is generate equivalent fractions by multiplying by 1:

$$\frac{2}{3} \times \frac{5}{5} =$$

Note that 1 is expressed as a fraction ($\frac{5}{5}$). The rule about multiplying by 1 is considered essential because the value on one side of the equation must equal the value on the other side. The only way to retain the value of $\frac{2}{3}$ is to multiply by 1.

Students also learn to complete pairs of equivalent fractions. For instance:

$$\frac{2}{3} \qquad = \frac{8}{\square}$$

or

$$\frac{2}{3} \qquad = \frac{\square}{15}$$

Solving this type of problem involves generating the fraction that equals 1. For the top problem, the student figures what to multiply 2 by to get 8. The answer is 4, so the fraction that equals 1 is $\frac{4}{4}$:

$$\frac{2}{3} \times \frac{4}{4} = \frac{8}{\boxed{12}}$$

For the other problem, the student figures out what to multiply 3 by to get 15 The answer is 5, so the fraction that equals 1 is $\frac{5}{5}$.

$$\frac{2}{3} \times \frac{5}{5} = \frac{\boxed{10}}{15}$$

Once students have a firm understanding of the basic operation, they work problems in which the missing factor is not a simple multiple. For instance:

$$\frac{2}{3} \qquad = \frac{5}{\square}$$

Before working this type of problem, students work on in-place multiplication, which permits them to express the missing factor as a fraction composed of the two numbers shown in the problem:

$$2 \times \frac{5}{2} = 5$$

The fraction is composed of the other two numbers in the problem—5 and 2. With the skill of in-place multiplication, students express the fraction that equals 1 as a fraction over a fraction:

$$\frac{2}{3}\left(\frac{\frac{5}{2}}{\frac{5}{2}}\right) = \frac{5}{\boxed{\frac{15}{2}}}$$

Students also work with tables, the rows of which have the same function:

	3	6
	2	4
Total	5	□

Note that the column values are added, so the bottom number is the sum of the other two. The missing value (10) can be determined either by adding the 6 and 4 or by multiplying the bottom row by 2 (the factor for all rows).

Table problems involving ratios are introduced by presenting two missing numbers in the second column.

	3	
	8	10
Total	11	

Students understand that each row has the same factor. They can find that factor for the row that has two values (8 and 10). The factor is $\frac{10}{8}$. Students multiply either of the other rows by that factor to get a second value in the column:

	3	$3\frac{6}{8}$
	8	10
Total	11	

The final missing number ($13\frac{6}{8}$) is calculated by adding $3\frac{6}{8}$ and 10.

Once the model has been introduced, students work a wide range of word problems, including those that express ratios as fractions and decimal values. The general discrimination students learn is that **if the problem refers to three names, make a table. If the problem presents two names, work a simple ratio problem.**

The ratio of boys to girls is 3 to 4. If there is a total of 350 children, how many are boys? How many are girls?

(The problem presents three names—boys, girls, children. So a table is required.)

The sale price of sofas is 20% less than the regular price. If the regular price of a sofa is $196, how much would a person save by buying the sofa on sale?

(This problem implies a table.)

	%	$
difference	20	
sale price	80	
regular price	100	196

As noted earlier, the goal of Direct Instruction programs is not to teach students a variety of ways to solve complicated problems, but rather a way that works for the full range of problems.

In the case of ratio-table problems, the model works for problems that involve probability. The problems show why probability problems represent a mixture of additive elements and factors.

There are 52 cards in a deck. 13 of them are spades. If a person took 60 trials at drawing cards from the deck, what is the expected number of spades the person would draw from the deck?

	Cards	Trials
spades	13	15
cards that are not spades	39	45
total	52	60

The table shows that the expected numbers for spades and other cards (second column) are based on the ratio properties of $\frac{60}{52}$. At the same time, the table shows that the basic logic of probability trials is retained: the trials for spades plus trials for cards that are not spades equals the total trials. A student with this template or model for understanding simple probability understands that probability is mathematical.

Non-Traditional Skill Sequences

The traditional practice provides students with very general guidelines for working application problems (read, think, try, and check). These rules show students, after the fact, how some problems are worked. The DI approach is to teach analyses of word problems that permit before-the-fact problem-solving steps.

First, students are taught the relevant calculations. Next, students are taught a type of word problem that uses the particular calculation taught. Finally, the word-problem type is practiced in an unstructured context. Problems are initially grouped according to relevant features. The ratio word problems, for example, fall into two main categories: those that involve "classification" and those that involve "difference." For all classification problems, the total is 100% or all there is. If there are blue birds and other birds in a swamp, the total for all the birds is 100%. This whole is divided into blue birds and others.

The word problem involving differences refers to comparisons. If train A is 20% longer than train B, there must be a difference number. Furthermore, there may be a value that is more than 100%. In the sentence above, the comparison is made with train B. Therefore, train B is 100%. Train A is 120%.

Within each major problem type (classification problems and difference problems) are many subtypes—problems that refer to fractions, problems that require manipulation of mixed numbers, and problems that require intermediate steps before basic calculations are possible. The DI sequence addresses all the subtypes of attempts to provide articulate teaching. The same procedure applies to all content in DI programs.

The judgment of what is relatively important and what should be prioritized in early levels of instruction derives from an analysis of what we expect the student to do later. Although this principle seems simple enough, its application leads to a greatly different set of priorities and skill sequences than those provided by traditional educational practices. Currently, for instance, the emphasis on reading instruction revolves around "literature." While there may be some benefits in introducing young students to "literature," a main priority is for students to learn to read textbooks and non-fictional material, to analyze arguments, and to apply rules—activities that students will do when reading to learn subjects like science and geography. The Direct Instruction reading sequence extensively develops skills in working in these areas.

In the final analysis, the DI sequences do not teach to achievement tests because they do not use the same response forms, and tests often fail to measure much of what the DI programs teach. The fact that students do relatively well on these tests implies that these students have learned a relatively broad range of skills.

7. Is there consumer protection information to suggest that the outcomes are possible?

Major instructional programs are routinely published without first being subjected to any field testing to determine the extent to which students actually master the material that is presented and the extent to which teachers are able to follow the program's presentation specifications. These programs provide no evidence that the material is "teachable" or "learnable." See the NCTE Report Card on Basal Readers (Goodman, Shannon, Freeman & Murphy, 1988). The ideal program-development sequence would require extensive field testing and revisions based on student performance.

The ideal field test would address these questions:

(a) Are teachers uniformly able to present the material as suggested by the teacher's guide or program?

(b) Is the specified presentation capable of being delivered in a "reasonable" time frame to both higher-performing and lower-performing students?

(c) Is the mastery percentage nearly 100% on all the skills and concepts that are taught?

As the questions indicate, the concern is with time, fidelity of presentation across teachers, and student performance. Ideally, every activity presented in the field-test version of the program would be assessed. There would be rules about the percentage of students who miss a particular type of task that would signal a need to revise the sequence. (For instance, if 6 out of 25 students have trouble with a problem type or specific problem, the sequence would be revised, either by replacing the problem or by providing additional preparation for that problem type.)

There would also be rules about time. If the teacher could not present a lesson in approximately the time allotted for the lesson, the program would not be able to predict student performance accurately. If the program indicates that the structured teaching part of the lessons should require 45 minutes, most lessons should be capable of being delivered by the teacher in 45 minutes. (This requirement is based on the assumption that students who are placed in the program have the skills needed to perform in the program and that the teacher has presentation skills that permit lessons to be reasonably paced.) The revised program or at least the parts of it that had generated problems during the initial field tryout would be tried out again to verify that the problems have been solved.

Programs should be shaped primarily by facts about student and teacher performance and not by authors' prejudices. Given that the prototype sequence covers the subject or topic, what happens to shape the program occurs as a function of the performance of teachers and students. If the tasks and activities do not meet performance standards, they are revised or dropped, regardless of the philosophy that generated them. Authors who follow this process construct programs that are quite far from their original ideas of what would work well or what kind of activities are appropriate and manageable.

Direct Instruction Solutions

The main form of consumer protection information that is available for Direct Instruction products comes through valid empirical studies that compare DI outcomes with those of other products or approaches. Although this information is available, it is not readily accessible to the consumer because the various studies are not featured in school teacher and administrator publications. These publications generally promote practices that are currently in vogue, such as teaching reading through literature and whole language, using "developmentally appropriate practices," and treating math and science as "ill-structured" disciplines.

A related problem is that impressive research and data are frequently presented to support favored approaches; however, the research is not rigorous and is often contradicted by later research. For example, over the past decade in California, dozens of school district and state tests have shown that students were performing at a reasonable level; however, independent assessments, such as the National Assessment of Educational Progress, have vigorously contradicted these outcomes and have shown that student performance is extremely low.

Direct Instruction programs attempt to follow sound consumer protection practices. The main result is that the programs are shaped and designed for use by the teacher. The programs attempt to teach the content and skills that are assumed by sophisticated academic applications so that students who complete the program will be prepared for future work. This pursuit serves teachers who will work with the students in the higher grades.

The details and components of DI programs are designed so they are efficient, which means that everything taught has a specific role with respect to future activities. The practices that shape the prototype program and goals tend to identify problems with the early versions of the program and suggest how they could be solved. If inert activities are included in the early versions of the products, they will probably be removed from the program early on analytical grounds (because they have no obvious function) or later during the field test (usually because they require training, but do not affect student performance on future activities).

The field testing of the program results in a product that should be technically sound or at least thoughtful. The developers should be able to answer questions about any detail procedure or convention. Their answers should reflect information gained from the extensive analysis of components and field testing data. Questions include the following: Why is the wording of the task as it is? Why does this task appear on the number of lessons it appears on, rather than some other number? Why was the particular response format that appears in the task selected over others? Why is the specific sequence of activities designed as it is during each lesson or from one lesson to the next? Direct Instruction programs are certainly not without weak spots or glitches; however, they are generally reliable.

Direct Instruction assumes that the teacher is analogous to the driver of a car. Just as a driver does not have to understand the intricate details of the fuel injection system, the teacher doesn't have to understand the engineering details of the program to teach it effectively. Any program that is shaped by the criteria listed in this section would have strong analytical underpinnings; however, those details would be quite invisible to the teacher because they are not greatly relevant to the teacher's role.

What is central to the teacher's role is the performance of the students. The program provides the teacher with a frame-by-frame view of student growth, the problems that students tend to have, how much practice they need, and how the various generalizations are induced by the program. From this view, the teacher will be both an engineer and a diagnostician of learning, in a very real sense.