Oneonta EPSY 275

Collection edited by: Brian Beitzel
Content authors: Kelvin Seifert, Brian Beitzel, Lisa White-McNulty, and Nathan Gonyea
Online: <http://cnx.org/content/col11446/1.6>

Collection structure revised: 2013/06/11

Table of Contents

Types of Instruction

Because some forms of thinking—critical thinking, creativity and problem solving—are broad and important educationally, it is not surprising that educators have identified strategies to encourage their development. Some of the possibilities are shown in Table 1 and are grouped into two categories: how much the strategy is student-centered versus teacher-centered. It should be emphasized that the classifications in Table 1 are not very precise, but they give a useful framework for understanding some of the options available for planning and implementing instruction.

Table 1

Student-centered	Cooperative learning, Inquiry, Discovery learning, Self-reflection, Independent study
Teacher-centered	Advance organizers, Lectures, Direct instruction, Madeline Hunter's “Effective Teaching”, Taking notes

Table 2 below defines some of the terms used in Table 1:

Table 2 Definitions of Terms in Table 1

Lecture	Telling or explaining previously organized information—usually to a group
Advance organizers	Brief overview, either verbally or graphically, of material about to be covered in a lecture or text
Taking notes	Writing important points of a lecture or reading
Madeline Hunter’s “Effective Teaching”	A set of strategies that emphasizes clear presentation of goals, the explanation and modeling of tasks to students and careful monitoring of students’ progress toward the goals

Facilitating complex thinking: Teacher-directed instruction

As the name implies, teacher-directed instruction includes any strategies initiated and guided primarily by the teacher. A classic example is exposition or lecturing (simply telling or explaining important information to students) combined with assigning reading from texts. But teacher-directed instruction also includes strategies that involve more active response from students, such as encouraging students to elaborate on new knowledge or to explain how new information relates to prior knowledge. Whatever their form, teacher-directed instructional methods normally include the organizing of information on behalf of students, even if teachers also expect students to organize it further on their own. Sometimes, therefore, teacher-directed methods are thought of as transmitting knowledge from teacher to student as clearly and efficiently as possible, even if they also require mental work on the part of the student.

Lectures and readings

Lectures and readings are traditional staples of educators, particularly with older students (including university students). At their best, they pre-organize information so that (at least in theory) the student only has to remember what was said in the lecture or written in the text in order to begin understanding it (Exley & Dennick, 2004). Their limitation is the ambiguity of the responses they require: listening and reading are by nature quiet and stationary, and do not in themselves indicate whether a student is comprehending or even attending to the material. Educators sometimes complain that “students are too passive” during lectures or when reading. But physical quietness is intrinsic to these activities, not to the students who do them. A book just sits still, after all, unless a student makes an effort to read it, and a lecture may not be heard unless a student makes the effort to listen to it.

Advance organizers

In spite of these problems, there are strategies for making lectures and readings effective. A teacher can be especially careful about organizing information *for* students, and she can turn part of the mental work over to students themselves. An example of the first approach is the use of **advance organizers**—brief overviews or introductions to new material before the material itself is presented (Ausubel, 1978). Textbook authors (including ourselves) often try deliberately to insert periodic advance organizers to introduce new sections or chapters in the text. When used in a lecture, advance organizers are usually statements in the form of brief introductory remarks, though sometimes diagrams showing relationships among key ideas can also serve the same purpose (Robinson, et al., 2003). Whatever their form, advance organizers partially organize the material on behalf of the students, so that they know where to put it all, so to speak, as they learn them in more detail. Organizers that are generated by students are not advance organizers because they are creating the organizer while they are learning (i.e., not in advance of instruction).

Recalling and relating prior knowledge

Another strategy for improving teacher-directed instruction is to encourage students to relate the new material to prior familiar knowledge. When one of us (Kelvin) first learned a foreign language (in his case French), for example, he often noticed similarities between French and English vocabulary. A French word for picture, for example, was *image,* spelled exactly as it is in English. The French word for *splendid* was *splendide,* spelled almost the same as in English, though not quite. Relating the French vocabulary to English vocabulary helped in learning and remembering the French.

As children and youth become more experienced in their academics, they tend to relate new information to previously learned information more frequently and automatically (Goodwin, 1999; Oakhill, Hartt, & Samols, 2005). But teachers can also facilitate students' use of this strategy. When presenting new concepts or ideas, the teacher can relate them to previously learned ideas deliberately—essentially modeling a memory strategy that students learn to use for themselves. In a science class, for example, she can say, "This is another example of…, which we studied before"; in social studies she can say, "Remember what we found out last time about the growth of the railroads? We saw that…"

If students are relatively young or are struggling academically, it is especially important to remind them of their prior knowledge. Teachers can periodically ask questions like "What do you already know about this topic?" or "How will your new knowledge about this topic change what you know already?" Whatever the age of students, connecting new with prior knowledge is easier with help from someone more knowledgeable, such as the teacher. When learning algorithms for multiplication, for example, students may not at first see how multiplication is related to addition processes which they probably learned previously (Burns, 2001). But if a teacher takes time to explain the relationship and to give students time to explore it, then the new skill of multiplication may be learned more easily.

Elaborating information

Elaborating new information means asking questions about the new material, inferring ideas and relationships among the new concepts. Such strategies are closely related to the strategy of recalling prior knowledge as discussed above: elaboration enriches the new information and connects it to other knowledge. In this sense elaboration makes the new learning more meaningful and less arbitrary.

A teacher can help students use elaboration by modeling this behavior. The teacher can interrupt his or her explanation of an idea, for example, by asking how it relates to other ideas, or by speculating about where the new concept or idea may lead. He or she can also encourage students to do the same, and even give students questions to guide their thinking. When giving examples of a concept, for example, a teacher can hold back from offering all of the examples, and instead ask students to think of additional examples themselves. The same tactic can work with assigned readings; if the reading includes examples, the teacher can instruct students to find or make up additional examples of their own.

Mastery learning

This term refers to an instructional approach in which all students learn material to an identically high level, even if some students require more time than others to do so (Gentile, 2004). In mastery learning, the teacher directs learning, though sometimes only in the sense of finding, writing, and orchestrating specific modules or units for students to learn. In one typical mastery learning program, the teacher introduces a few new concepts or topics through a brief lecture or teacher-led demonstration. Then she gives an ungraded assignment or test immediately in order to assess how well students have learned the material, and which ones still need help. The students who have already learned the unit are given enrichment activities. Those needing more help are provided individual tutoring or additional self-guiding materials that clarify the initial content; they work until they have in fact mastered the content (hence the name *mastery learning*). At that point students take another test or do another assignment to show that they have in fact learned the material to the expected high standard. When the system is working well, all students end up with high scores or grades, although usually some take longer to do so than others.

As you might suspect, mastery learning poses two challenges. The first is ethical: is it really fair to give enrichment only to faster students and remediation only to slower students? This practice could deteriorate into continually providing the fast with an interesting education, while continually providing the slow only with boring, repetitious material. In using the approach, therefore, it is important to make all materials interesting, whether enrichment or remedial. It is also important to make sure that the basic learning goals of each unit are truly important—even crucial—for everyone to learn, so that even slower individuals spend their time well.

The other challenge of mastery learning is more practical: the approach makes strong demands for detailed, highly organized curriculum. If the approach is to work, the teacher must either locate such a curriculum, write one herself, or assemble a suitable mixture of published and self-authored materials. However the curriculum is created, the end result has to be a program filled with small units of study as well as ample enrichment and remedial materials. Sometimes providing these practical requirements can be challenging. But not always: some subjects (like mathematics) lend themselves to detailed, sequential organization especially well. In many cases, too, commercial publishers have produced curricula already organized for use in mastery learning programs (Fox, 2004).

Direct instruction

Although the term *direct instruction* is sometimes a synonym for *teacher-directed instruction*, more often it refers to a version of mastery learning that is highly scripted, meaning that it not only organizes the curriculum into small modules or units as described above, but also dictates *how* teachers should teach and sometimes even the words they should speak (Adams & Engelmann, 1996; Magliaro, Lockee, & Burton, 2005). Direct instruction programs are usually based on a mix of ideas from behaviorism and cognitive theories of learning. In keeping with behaviorism, the teacher is supposed to praise students immediately and explicitly when they give a correct answer. In keeping with cognitive theory, she is supposed to state learning objectives in advance of teaching them (providing a sort of mini-advance organizer), provide frequent reviews of materials, and check deliberately on how well students are learning. Direct instruction usually also introduces material in small, logical steps, and calls for plenty of time for students to practice.

Direct instruction programs share one of the challenges of other mastery learning approaches: because they hold all students to the same high standard of achievement, they must deal with differences in how long students require to reach the standard. But direct instruction has an additional challenge, in that they often rely on small-group interaction more heavily than other mastery learning programs, and use self-guiding materials less. This difference has the benefit that direct instruction works especially well with younger students (especially kindergarten through third grade), who may have limited skills at working alone for extended periods. The challenge is that reliance on small-group interaction can make it impractical to use direct instruction with an entire class or for an entire school day. In spite of these limits, however, research has found direct instruction to be very effective in teaching basic skills such as early reading and arithmetic (Adams & Engelmann, 1996).

Madeline Hunter's effective teaching model

A number of direct instruction strategies have been combined by Madeline Hunter into a single, relatively comprehensive approach that she calls **mastery teaching** (not to be confused with the related term *mastery learning)* or the **effective teaching model** (M. Hunter, 1982; R. Hunter, 2004). Important features of the model are summarized in Table 3. As you can see, the features span all phases of contact with students—before, during, and after lessons.

Table 3 Madeline Hunter's "Effective Teaching Model" *Source*: R. Hunter, 2004

Prepare students to learn. • Make good use of time at the beginning of a lesson or activity, when attention is best • Direct students' attention to what lies ahead in a lesson—for example, by offering "advance organizers" • Explain lesson objectives explicitly
Present information clearly and explicitly. • Set a basic structure to the lesson and stay with it throughout • Use familiar terms and examples • Be concise
Check for understanding and give guided practice. • Ask questions that everyone responds to—for example, "Raise your hand if you think the answer is X" • Invite choral responses—for example, "Is this a correct answer or not?" • Sample individuals' understanding—for example, "Barry, what's your example of X?"
Provide for independent practice. • Work through the first few exercises or problems together • Keep independent practice periods brief and intersperse with discussions that offer feedback

What happens even before a lesson begins? Like many forms of teacher-directed instruction, the effective teaching model requires curricula and learning goals that are tightly organized and divisible into small parts, ideas, or skills. In teaching about photosynthesis, for example, the teacher (or at least her curriculum) needs to identify the basic elements that contribute to this process, and how they relate to each other. With photosynthesis, the elements include the sun, plants, animals, chlorophyll, oxygen produced by plants and consumed by animals, and carbon dioxide that produced by animals and consumed by plants. The roles of these elements need to be identified and expressed at a level appropriate for the students. With advanced science students, oxygen, chlorophyll, and carbon dioxide may be expressed as part of complex chemical reactions; with first-grade students, though, they may be expressed simply as parts of a process akin to breathing or respiration.

Once this analysis of the curriculum has been done, the Hunter's effective teaching model requires making the most of the lesson time by creating an **anticipatory set**, which is an activity that focuses or orients the attention of students to the upcoming content. Creating an anticipatory set may consist, for example, of posing one or more questions about students' everyday knowledge or knowledge of prior lessons. In teaching about differences between fruits and vegetables, the teacher could start by asking: "If you are making a salad strictly of fruit, which of these would be OK to use: apple, tomato, cucumber, or orange?" As the lesson proceeds, information needs to be offered in short, logical pieces, using language as familiar as possible to the students. Examples should be plentiful and varied: if the purpose is to define and distinguish fruits and vegetables, for example, then features defining each group should be presented singularly or at most just a few at a time, with clear-cut examples presented of each feature. Sometimes models or analogies also help to explain examples. A teacher can say: "Think of a fruit as a sort of 'decoration' on the plant, because if you pick it, the plant will go on living." But models can also mislead students if they are not used thoughtfully, since they may contain features that differ from the original concepts. In likening a fruit to a decoration, for example, students may overlook the essential role of fruit in plant reproduction, or think that lettuce qualifies as a fruit, since picking a few lettuce leaves does not usually kill a lettuce plant.

Throughout a lesson, the teacher repeatedly **checks for understanding** by asking questions that call for active thinking on the part of students. One way is to require all students to respond somehow, either with an actual choral response (speaking in unison together), another way with a non-verbal signal like raising hands to indicate answers to questions. In teaching about fruits and vegetables, for example, a teacher can ask, "Here's a list of fruits and vegetables. As I point to each one, raise your hand if it's a fruit, but not if it's a vegetable." Or she can ask: "Here's a list of fruits and vegetables. Say together what each on is as I point to it; you say 'fruit' or 'vegetable', whichever applies." Even though some students may hide their ignorance by letting more knowledgeable classmates do the responding, the general level or quality of response can still give a rough idea of how well students are understanding. These checks can be supplemented, of course, with questions addressed to individuals, or with questions to which individuals must respond briefly in writing. A teacher can ask everyone, "Give me an example of one fruit and one vegetable", and then call on individuals to answer. She can also say: "I want everyone to make a list with two columns, one listing all the fruits you can think of and the other listing all the vegetables you can think of."

As a lesson draws to a close, the teacher arranges for students to have further **independent practice**. The point of the practice is not to explore new material or ideas, but to consolidate or strengthen the recent learning. At the end of a lesson about long division, for example, the teacher can make a transition to independent practice by providing a set of additional problems similar to the ones she explained during the lesson. After working one or two with students, she can turn the rest of the task over to the students to practice on their own. But note that even though the practice is supposedly “independent”, students’ understanding still has be checked frequently. A long set of practice problems therefore needs to be broken up into small subsets of problems, and written or oral feedback offered periodically.

What are the limits of teacher-directed instruction?

Whatever the grade level, most subjects taught in schools have at least some features, skills, or topics that benefit from direct instruction. Even subjects usually considered “creative” can benefit from a direct approach at times: to draw, sing, or write a poem, for example, requires skills that may be easier to learn if presented sequentially in small units with frequent feedback from a teacher. Research supports the usefulness of teacher-directed instruction for a variety of educational contexts when it is designed well and implemented as intended (Rosenshine & Mesister,1995; Good & Brophy, 2004). Teachers themselves also tend to support the approach in principle (Demant & Yates, 2003).

But there are limits to its usefulness. Some are the practical ones are pointed out above. Teacher-directed instruction, whatever the form, requires well-organized units of instruction in advance of when students are to learn. Such units may not always be available, and it may not be realistic to expect busy teachers to devise their own. Other limits of direct instruction have more to do with the very nature of learning. Some critics argue that organizing material on behalf of the students encourages students to be passive—an ironic and undesirable result if true (Kohn, 2000, 2006). According to this criticism, the mere fact that a curriculum or unit of study is constructed by a teacher (or other authority) makes some students think that they should not bother seeking information actively on their own, but wait for it to arrive of its own accord. In support of this argument, critics point to the fact that direct instruction approaches sometimes contradict their own premises by requiring students to do a bit of cognitive organizational work of their own. This happens, for example, when a mastery learning program provides enrichment material to faster students to work on independently; in that case the teacher may be involved in the enrichment activities only minimally.

Criticisms like these have led to additional instructional approaches that rely more fully on students to seek and organize their own learning. In the next section we discuss some of these options. As you will see, student-centered models of learning do solve certain problems of teacher-directed instruction, but they also have problems of their own.

Student-centered models of learning

Student-centered models of learning shift some of the responsibility for directing and organizing learning from the teacher to the student. Being student-centered does not mean, however, that a teacher gives up organizational and leadership responsibilities completely. It only means a relative shift in the teacher’s role, toward one with more emphasis on guiding students’ self-chosen directions. As we explained earlier in this chapter, teacher-directed strategies do not take over responsibility for students’ learning completely; no matter how much a teacher structures or directs learning, the students still have responsibility for working and expending effort to comprehend new material. By the same token, student-centered models of learning do not mean handing over all organizational work of instruction to students. The teacher is still the most knowledgeable member of the class, and still has both the opportunity and the responsibility to guide learning in directions that are productive.

As you might suspect, therefore, teacher-directed and student-centered approaches to instruction may overlap in practice. You can see the overlap clearly, for example, in two instructional strategies commonly thought of as student-centered, *independent study* and *self-reflection.* In **independent study**, as the name implies, a student works alone a good deal of the time, consulting with a teacher only occasionally. Independent study may be student-centered in the sense that the student may be learning a topic or skill—an exotic foreign language, for example—that is personally interesting. But the opposite may also be true: the student may be learning a topic or skill that a teacher or an official school curriculum has directed the student to learn—a basic subject for which the student is missing a credit, for example. Either way, though, the student will probably need guidance, support, and help from a teacher. In this sense even independent study always contain elements of teacher-direction.

Similarly, **self-reflection** refers to thinking about beliefs and experiences in order to clarify their personal meaning and importance. In school it can be practiced in a number of ways: for example by keeping diaries or logs of learning or reading, or by retelling stories of important experiences or incidents in a student’s life, or by creating concept maps like the ones described earlier in this chapter.

Whatever form it takes, self-reflection by definition happens inside a single student's mind, and in this sense is always directed by the student. Yet most research on self-reflection finds that self-reflection only works well when it involves and generates responses and interaction with other students or with a teacher (Seifert, 1999; Kuit, Reay, & Freeman, 2001). To be fully self-reflective, students need to have access to more than their existing base of knowledge and ideas—more than what they know already. In one study about students' self-reflections of cultural and racial prejudices (Gay & Kirkland, 2003), for example, the researchers found that students tended to reflect on these problems in relatively shallow ways if they worked on their own. It was not particularly effective to write about prejudice in a journal that no one read except themselves, or to describe beliefs in a class discussion in which neither the teacher nor classmates commented or challenged the beliefs. Much more effective in both cases was for the teacher to respond thoughtfully to students' reflective comments. In this sense the use of self-reflection, like independent study, required elements of teacher-direction to be successful.

How might a teacher emphasize students' responsibility for directing and organizing their own learning? The alternatives are numerous, as they are for teacher-directed strategies, so we can only sample some of them here. We concentrate on ones that are relatively well known and used most widely, and especially on two: inquiry learning and cooperative learning.

Inquiry learning

Inquiry learning stands the usual advice about expository (lecture-style) teaching on its head: instead of presenting well-organized knowledge to students, the teacher (or sometimes fellow students) pose thoughtful questions intended to stimulate discussion and investigation by students. The approach has been described, used, and discussed by educators literally for decades, though sometimes under other names, including *inquiry method* (Postman & Weingartner, 1969), *discovery learning* (Bruner, 1960/2006), or *progressive education* (Dewey, 1933; Martin, 2003). For convenience, we will stay with the term *inquiry learning.*

The questions that begin a cycle of inquiry learning may be posed either by the teacher or by students themselves. Their content depends not only on the general subject area being studied, but also on the interests which students themselves have expressed. In elementary-level science, for example, a question might be "Why do leaves fall off trees when winter comes?" In high school social studies classes, it might be "Why do nations get into conflict?" The teacher avoids answering such questions directly, even if asked to do so. Instead she encourages students to investigate the questions themselves, for example by elaborating on students' ideas and by asking further questions based on students' initial comments. Since students' comments can not be predicted precisely, the approach is by nature flexible. The initial questioning helps students to create and clarify questions which they consider worthy of further investigation. Discussing questions about leaves falling off trees, for example, can prompt students to observe trees in the autumn or to locate books and references that discuss or explain the biology of tress and leaves.

But inquiry is not limited to particular grade levels or topics. If initial questions in a high school social studies class have been about why nations get into conflict, for example, the resulting discussions can lead to investigating the history of past wars and the history of peace-keeping efforts around the world. Whether the topic is high school social studies or elementary school biology, the specific direction of investigations is influenced heavily by students, but with assistance from the teacher to insure that the students' initiatives are productive. When all goes well, the inquiry and resulting investigations benefit students in two ways. The first is that students (perhaps obviously) learn new knowledge from their investigations. The second is that students practice a constructive, motivating way of learning, one applicable to a variety of problems and tasks, both in school and out.

Cooperative learning

Even though inquiry-oriented discussion and investigation benefits when it involves the teacher, it can also be useful for students to work together somewhat independently, relying on a teacher's guidance only indirectly. Working with peers is a major feature of **cooperative learning** (sometimes also called *collaborative learning).* In this approach, students work on a task in groups and often are rewarded either partially or completely for the success of the group as a whole. Aspects of cooperative learning have been part of education for a long time; some form of cooperation has always been necessary to participate on school sports teams, for example, or to produce a student-run school newspaper. What is a bit newer is using cooperative or collaborative activities systematically to facilitate the learning of a range of educational goals central to the academic curriculum (Prince, 2004).

Even though teachers usually value cooperation in students, circumstances at school can sometimes reduce students' incentives to show it. The traditional practice of assessing students individually, for example, can set the stage for competition over grades, and cultural and other forms of diversity can sometimes inhibit individuals from helping each other spontaneously. Strategies exist, however, for reducing such barriers so that students truly benefit from each other's presence, and are more likely to feel like sharing their skills and knowledge. Here, for example, are several key features that make cooperative learning work well (Johnson & Johnson, 1998; Smith, et al., 2005):

- *Students need time and a place to talk and work together.* This may sound obvious, but it can be overlooked if time in class becomes crowded with other tasks and activities, or with interruptions related to school (like assemblies) but not to the classroom. It is never enough simply to tell students to work together, only to leave them wondering how or when they are to do so.
- *Students need skills at working together.* As an adult, you may feel relatively able to work with a variety of partners on a group task. The same assumption cannot be made, however, about younger individuals, whether teenagers or children. Some students may get along with a variety of partners, but others may not. Many will benefit from advice and coaching about how to focus on the tasks at hand, rather than on the personalities of their partners.
- *Assessment of activities should hold both the group and the individuals accountable for success.* If a final mark for a project goes only to the group as a whole, then freeloading is possible: some members may not do their share of the work and may be rewarded more than they deserve. Others may be rewarded less than they deserve. If, on the other hand, a final grade for a group project goes only to each member's individual contribution to a group project, then **overspecialization** can occur: individuals have no real incentive to work together, and cooperative may deteriorate into a set of smaller individual projects (Slavin, 1994).
- *Students need to believe in the value and necessity of cooperation.* Collaboration will not occur if students privately assume that their partners have little to contribute to their personal success. Social prejudices from the wider society—like racial bias or gender sexism, for example—can creep into the operations of cooperative groups, causing some members to be ignored unfairly while others are overvalued. Teachers can help reduce these problems in two ways: first by pointing out and explaining that a diversity of talents is necessary for success on a group project, and second by pointing out to the group how undervalued individuals are contributing to the overall project (Cohen, Brody, & Sapon-Shevin, 2004).

As these comments imply, cooperative learning does not happen automatically, and requires monitoring and support by the teacher. Some activities may not lend themselves to cooperative work, particularly if every member of the group is doing essentially the same task. Giving everyone in a group the same set of arithmetic problems to work on collaboratively, for example, is a formula for cooperative failure: either the most skilled students do the work for others (freeloading) or else members simply divide up the problems among themselves in order to reduce their overall work (overspecialization). A better choice for a cooperative task is one that clearly requires a diversity of skills, what some educators call a *rich group work task* (Cohen, Brody, & Sapon-Shevin, 2004). Preparing a presentation about medieval castles, for example, might require (a) writing skill to create a report, (b) dramatic skill to put on a skit and (c) artistic talent to create a poster. Although a few students may have all of these skills, more are likely to have only one, and they are therefore likely to need and want their fellow group members' participation.

Examples of cooperative and collaborative learning

Although this description may make the requirements for cooperative learning sound somewhat precise, there are actually a variety of ways to implement it in practice. Error: Reference source not found summarizes several of them. As you can see, the strategies vary in the number of how many students they involve, the prior organization or planning provided by the teacher, and the amount of class time they normally require.

Table 4 Strategies for encouraging cooperative learning

Strategy	Type of groups involved:	What the teacher does:	What the students do:
Think-pair-share (Lyman, 1981)	Pairs of students, sometimes linked to one other pair	Teacher poses initial problem or question.	First, students think individually of the answer; second, they share their thinking with partner; third, the partnership shares their thinking with another partnership.
Jigsaw classroom, version #1 (Aronson, et al., 2001)	5-6 students per group, and 5-6 groups overall	Teacher assigns students to groups and assigns one aspect of a complex problem to each group.	Students in each group work together to become experts in their particular aspect of the problem; later the expert groups disband, and form new groups containing one student from each of the former expert groups.
Jigsaw classroom, version #2 (Slavin, 1994)	4-5 students per group, and 4-5 groups overall	Teacher assigns students to groups and assigns each group to study or learn about the same *entire* complex problem.	Students initially work in groups to learn about the entire problem; later the groups disband and reform as expert groups, with each group focusing on a selected aspect of the general problem; still later the expert groups disband and the original general groups reform to learn what the expert students can now add to their general understanding.
STAD (Student-Teams-Achievement Divisions) (Slavin, 1994)	4-5 students per team (or group)	Teacher presents a lesson or unit to the entire class, and later tests them on it; grades individuals based partly on individuals' and the team's improvement, not just on absolute level of performance.	Students work together to insure that team mates improve their performance as much as possible. Students take tests as individuals.
Project-Based Learning (Katz, 2000)	Various numbers of students, depending on the complexity of the project, up to and including the entire class	Teacher or students pose a question or problem of interest to other students; teacher assists students to clarify their interests and to make plans to investigate the question further.	Students work together for extended periods to investigate the original question or problem; project leads eventually to a presentation, written report, or other product.

Instructional strategies: an abundance of choices

Looking broadly at this chapter, you can see that choices among instructional strategies are numerous indeed, and that deciding among them depends on the forms of thinking that you want to encourage, the extent to which ideas or skills need to be organized by you to be understood by students, and the extent to which students need to take responsibility for directing their own learning. Although you may have personal preferences among possible instructional strategies, the choice will also be guided by the uniqueness of each situation of teaching—with its particular students, grade-level, content, and purposes. If you need to develop students' problem solving skills, for example, there are strategies that are especially well suited for this purpose; we described some. If you need to organize complex information so that students do not become confused by it, there are effective ways of doing so. If you want the students to take as much initiative as possible in organizing their own learning, this too can be done.

Yet having this knowledge is still not enough to teach well. What is still needed are ideas or principles for deciding *what* to teach. In this chapter we have still not addressed an obvious question: How do I find or devise goals for my teaching and for my students' learning? And assuming that I can determine the goals, where can I find resources that help students to meet them?

Behavioral View of Learning

Several ideas and priorities, then, affect how we teachers think about learning, including the curriculum, the difference between teaching and learning, sequencing, readiness, and transfer. The ideas form a "screen" through which to understand and evaluate whatever psychology has to offer education. As it turns out, many theories, concepts, and ideas from educational psychology *do* make it through the "screen" of education, meaning that they are consistent with the professional priorities of teachers and helpful in solving important problems of classroom teaching. In the case of issues about classroom learning, for example, educational psychologists have developed a number of theories and concepts that are relevant to classrooms, in that they describe at least *some* of what usually happens there and offer guidance for assisting learning. It is helpful to group the theories according to whether they focus on changes in behavior or in thinking. The distinction is rough and inexact, but a good place to begin. For starters, therefore, consider two perspectives about learning, called behaviorism (learning as changes in overt behavior) and constructivism, (learning as changes in thinking). The second category can be further divided into psychological constructivism (changes in thinking resulting from individual experiences), and social constructivism, (changes in thinking due to assistance from others). The rest of this chapter describes key ideas from each of these viewpoints. As I hope you will see, each describes some aspects of learning not just in general, but as it happens in classrooms in particular. So each perspective suggests things that you might do in your classroom to make students' learning more productive.

Behaviorism: changes in what students do

Behaviorism is a perspective on learning that focuses on changes in individuals' observable behaviors—changes in what people say or do. At some point we all use this perspective, whether we call it "behaviorism" or something else. The first time that I drove a car, for example, I was concerned primarily with whether I could actually do the driving, not with whether I could describe or explain how to drive. For another example: when I reached the point in life where I began cooking meals for myself, I was more focused on whether I could actually produce edible food in a kitchen than with whether I could explain my recipes and cooking procedures to others. And still another example—one often relevant to new teachers: when I began my first year of teaching, I was more focused on doing the job of teaching—on day-to-day survival—than on pausing to reflect on what I was doing.

Note that in all of these examples, focusing attention on behavior instead of on "thoughts" may have been desirable at that moment, but not necessarily desirable indefinitely or all of the time. Even as a beginner, there are times when it *is* more important to be able to describe how to drive or to cook than to actually do these things. And there definitely are many times when reflecting on and thinking about teaching can improve teaching itself. (As a teacher-friend once said to me: "Don't just *do* something; *stand* there!") But neither is focusing on behavior which is not necessarily less desirable than focusing on students' "inner" changes, such as gains in their knowledge or their personal attitudes. If you are teaching, you will need to attend to all forms of learning in students, whether inner or outward.

In classrooms, behaviorism is most useful for identifying relationships between specific actions by a student and the immediate precursors and consequences of the actions. It is less useful for understanding changes in students' thinking; for this purpose we need a more *cognitive* (or thinking-oriented) theory, like the ones described later in this chapter. This fact is not really a criticism of behaviorism as a perspective, but just a clarification of its particular strength or source of usefulness, which is to highlight observable relationships among actions, precursors and consequences.

Behaviorists use particular terms (or "lingo", some might say) for these relationships. They also rely primarily on two basic images or models of behavioral learning, called *classical conditioning* and *operant conditioning*. The names are derived partly from the major learning mechanisms highlighted by each type, which I describe next.

Classical conditioning: learning new associations with prior behaviors

As originally conceived, **classical conditioning** begins with the involuntary responses to particular sights, sounds, or other sensations (Lavond, 2003). When I receive an injection from a nurse or doctor, for example, I cringe, tighten my muscles, and even perspire a bit. Whenever a contented, happy baby looks at me, on the other hand, I invariably smile in response. I cannot help myself in either case; both of the responses are automatic. In humans as well as other animals, there is a repertoire or variety of such specific, involuntary behaviors. At the sound of a sudden loud noise, for example, most of us show a "startle" response—we drop what we are doing (sometimes literally!), our heart rate shoots up temporarily, and we look for the source of the sound. Cats, dogs and many other animals (even fish in an aquarium) show similar or equivalent responses.

Involuntary stimuli and responses were first studied systematically early in the twentieth-century by the Russian scientist Ivan Pavlov (1927). Pavlov's most well-known work did not involve humans, but dogs, and specifically their involuntary tendency to salivate when eating. He attached a small tube to the side of dogs' mouths that allowed him to measure how much the dogs salivated when fed (Table 1 shows a photograph of one of Pavlov's dogs). But he soon noticed a "problem" with the procedure: as the dogs gained experience with the experiment, they often salivated *before* they began eating. In fact the most experienced dogs sometimes began salivating before they even saw any food, simply when Pavlov himself entered the room! The sight of the experimenter, which had originally been a neutral experience for the dogs, became associated with the dogs' original salivation response. Eventually, in fact, the dogs would salivate at the sight of Pavlov even if he did *not* feed them.

This *change* in the dogs' involuntary response, and especially its growing independence from the food as stimulus, eventually became the focus of Pavlov's research. Psychologists named the process *classical conditioning* because it describes changes in *responses* to stimuli. Classical conditioning has several elements, each with a special name. To understand these, look at and imagine a dog (perhaps even mine, named Ginger) prior to any conditioning. At the beginning Ginger salivates (an **unconditioned response (UR)**) only when she actually tastes her dinner (an **unconditioned stimulus (US)**). As time goes by, however, a neutral stimulus—such as the sound of opening a bag containing fresh dog food—is continually paired with the eating/tasting experience. Eventually the neutral stimulus becomes able to elicit salivation even *before* any dog food is offered to Ginger, or even if the bag of food is empty! At this point the neutral stimulus is called a **conditioned stimulus (UCS)** and the original response is renamed as a **conditioned response (CR)**. Now, after conditioning, Ginger salivates merely at the sound of opening *any* large bag, regardless of its contents. (I might add that Ginger also engages in other conditioned responses, such as looking hopeful and following me around the house at dinner time.)

Table 1 Classical conditioning of Ginger, the dog. ***Before conditioning, Ginger salivates only to the taste of food and the bell has no effect. After conditioning, she salivates even when the bell is presented by itself.***

Before Conditioning: (UCS) **Food** → **Salivation** (UR) (UCS) **Bell** → **No response** (UR) During Conditioning: **Bell + Food** → **Salivation** After Conditioning: (CS) **Bell only** → **Salivation** (CR)	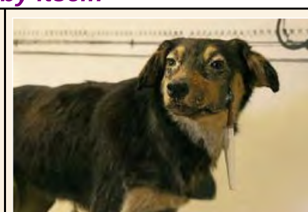

Classical Conditioning and Students

"OK," you may be thinking, "Classical conditioning may happen to animals. But does anything like it happen in classrooms?" It might seem like not much would, since teaching is usually about influencing students' conscious words and thoughts, and not their involuntary behaviors. But remember that schooling is not just about encouraging thinking and talking. Teachers, like parents and the public, also seek positive changes in students' attitudes and feelings—attitudes like a love for learning, for example, and feelings like self-confidence. It turns out that classical conditioning describes these kinds of changes relatively well.

Consider, for example, a child who responds happily whenever meeting a new person who is warm and friendly, but who also responds cautiously or at least neutrally in any new situation. Suppose further that the "new, friendly person" in question is you, his teacher. Initially the child's response to you is like an unconditioned stimulus: you smile (the unconditioned stimulus) and in response he perks up, breathes easier, and smiles (the unconditioned response). This exchange is not the whole story, however, but merely the setting for an important bit of behavior change: suppose you smile at him while standing in your classroom, a "new situation" and therefore one to which he normally responds cautiously. Now classical-conditioning learning can occur. The initially neutral stimulus (your classroom) becomes associated repeatedly with the original unconditioned stimulus (your smile) and the child's unconditioned response (his smile). Eventually, if all goes well, the classroom becomes a conditioned stimulus in its own right: it can elicit the child's smiles and other "happy behaviors" even without your immediate presence or stimulus. Table 2 diagrams the situation graphically. When the change in behavior happens, you might say that the child has "learned" to like being in your classroom. Truly a pleasing outcome for both of you!

Table 2 Classical conditioning of student to classroom. ***Before conditioning, the student smiles only when he sees the teacher smile, and the sight of the classroom has no effect. After conditioning, the student smiles at the sight of the classroom even without the teacher present.***

Before Conditioning:
(UCS) **Seeing Teacher Smile → Student Smiles** (UR)
(UCS) **Seeing Classroom → No response** (UR)
During Conditioning:
Seeing Teaching Smile + Seeing Classroom → Student Smiles
After Conditioning:
(CS) **Seeing Classroom → Student Smiles** (CR)

But less positive or desirable examples of classical conditioning also can happen. Consider a modification of the example that I just gave. Suppose the child that I just mentioned did *not* have the good fortune of being placed in *your* classroom. Instead he found himself with a less likeable teacher, whom we could simply call Mr Horrible. Instead of smiling a lot and eliciting the child's unconditioned "happy response", Mr Horrible often frowns and scowls at the child. In this case, therefore, the child's initial *unconditioned* response is negative: whenever Mr Horrible directs a frown or scowl at the child, the child automatically cringes a little, his eyes widen in fear, and his heart beat races. If the child sees Mr Horrible doing most of his frowning and scowling *in* the classroom, eventually the classroom itself will acquire power as a negative conditioned stimulus. Eventually, that is, the child will not need Mr Horrible to be present in order to feel apprehensive; simply being in the classroom will be enough. Table 3 diagrams this unfortunate situation. Obviously it is an outcome to be avoided, and in fact does not usually happen in such an extreme way. But hopefully it makes the point: any stimulus that is initially neutral, but that gets associated with an unconditioned stimulus and response, can eventually acquire the ability to elicit the response by itself. *Anything*—whether it is desirable or not.

Table 3 Classical conditioning of student to classroom. ***Before conditioning, the student cringes only when he sees Mr Horrible smile, and the sight of the classroom has no effect. After conditioning, the student cringes at the sight of the classroom even without Mr Horrible present.***

Before Conditioning:
(UCS) **Mr Horrible Frowns → Student Cringes** (UCR)
Mr Horrible's Classroom → No response
During Conditioning:
Mr Horrible Frowns + Sight of Classroom → Student Cringes
After Conditioning:
(CS) **Seeing Classroom → Student Cringes** (CR)

The changes described in these two examples are important because they can affect students' attitude about school, and therefore also their *motivation* to learn. In the positive case, the child becomes more inclined to please the teacher and to attend to what he or she has to offer; in the negative case, the opposite occurs. Since the changes in attitude happen "inside" the child, they are best thought of as one way that a child can acquire *i* **intrinsic motivation**, meaning a desire or tendency to direct attention and energy in a particular way that originates from the child himself or herself. Intrinsic motivation is sometimes contrasted to **extrinsic motivation**, a tendency to direct attention and energy that originates from *outside* of the child. As we will see, classical conditioning can influence students' intrinsic motivation in directions that are either positive or negative. As you might suspect, there are other ways to influence motivation as well. Many of these will be described when we discuss the topic of student motivation later in the course. First, though, let us look at three other features of classical conditioning that complicate the picture a bit, but also render conditioning a bit more accurate, an appropriate description of students' learning.

Three key ideas about classical conditioning

Extinction: This term does not refer to the fate of dinosaurs, but to the *disappearance* of a link between the conditioned stimulus and the conditioned response. Imagine a third variation on the conditioning "story" described above. Suppose, as I suggested above, that the child begins by associating your happy behaviors—your smiles—to his being present in the classroom, so that the classroom itself becomes enough to elicit his own smiles. But now suppose there is a sad turn of events: you become sick and must therefore leave the classroom in the middle of the school year. A substitute is called in who is not Mr Horrible, but simply someone who is not very expressive, someone we can call Ms Neutral. At first the child continues to feel good (that is, to smile) whenever present in the classroom. But because the link between the classroom and your particular smile is no longer repeated or associated, the child's response gradually *extinguishes*, or fades until it has disappeared entirely. In a sense the child's initial learning is "unlearned".

Extinction can also happen with negative examples of classical conditioning. If Mr Horrible leaves mid-year (perhaps because no one could stand working with him any longer!), then the child's negative responses (cringing, eyes widening, heart beat racing, and so on) will also extinguish eventually. Note, though, that whether the conditioned stimulus is positive or negative, extinction does not happen suddenly or immediately, but unfolds over time. This fact can sometimes obscure the process if you are a busy teacher attending to many students.

Generalization: When Pavlov studied conditioning in dogs, he noticed that the original conditioned stimulus was not the only neutral stimulus that elicited the conditioned response. If he paired a particular bell with the sight of food, for example, so that the bell became a conditioned stimulus for salivation, then it turned out that *other* bells, perhaps with a different pitch or type or sound, also acquired some ability to trigger salivation—though not as much as the original bell. Psychologists call this process generalization, or the tendency for similar stimuli to elicit a conditioned response. The child being conditioned to your smile, for example, might learn to associate your smile not only with being present in *your* classroom, but also to being present in other, similar classrooms. His conditioned smiles may be strongest where he learned them initially (that is, in your own room), but nonetheless visible to a significant extent in other teachers' classrooms. To the extent that this happens, he has *generalized* his learning. It is of course good news; it means that we can say that the child is beginning to "learn to like school" in general, and not just your particular room. Unfortunately, the opposite can also happen: if a child learns negative associations from Mr Horrible, the child's fear, caution, and stress might generalize to other classrooms as well. The lesson for teachers is therefore clear: we have a responsibility, wherever possible, to make classrooms pleasant places to be.

Discrimination: Generalization among similar stimuli can be reduced if only one of the similar stimuli is associated consistently with the unconditioned response, while the others are not. When this happens, psychologists say that **discrimination learning** has occurred, meaning that the individual has learned to distinguish or respond differently to one stimulus than to another. From an educational point of view, discrimination learning can be either desirable or not, depending on the particulars of the situation. Imagine again (for the fourth time!) the child who learns to associate your classroom with your smiles, so that he eventually produces smiles of his own whenever present in your room. But now imagine yet another variation on his story: the child is old enough to attend *middle school*, and therefore has several teachers across the day. You—with your smiles—are one, but so are Mr Horrible and Ms Neutral. At first the child may generalize his classically conditioned smiles to the other teachers' classrooms. But the other teachers do not smile like you do, and this fact causes the child's smiling to extinguish somewhat in their rooms. Meanwhile, you keep smiling in your room. Eventually the child is smiling *only* in your room and not in the other rooms. When this happens, we say that **discrimination** has occurred, meaning that the conditioned associations happen only to a single version of the unconditioned stimuli—in this case, only to *your* smiles, and not to the (rather rare) occurrences of smiles in the other classrooms. Judging by his behavior, the child is making a distinction between your room and others.

In one sense the discrimination in this story is unfortunate in that it prevents the child from acquiring a liking for school that is generalized. But notice that an opposing, more desirable process is happening at the same time: the child is also *prevented* from acquiring a generalized *dislike* of school. The fear-producing stimuli from Mr Horrible, in particular, become discriminated from the happiness-producing smiles from you, so the child's learns to confine his fearful responses to that particular classroom, and does not generalize them to other "innocent" classrooms, including your own. This is still not an ideal situation for the student, but maybe it is more desirable than disliking school altogether.

Operant conditioning: new behaviors because of new consequences

Instead of focusing on associations between stimuli and responses, **operant conditioning** focuses on how the effects of consequences on behaviors. The operant model of learning begins with the idea that certain consequences tend to make certain behaviors happen more frequently. If I compliment a student for a good comment during a discussion, there is more of a chance that I will hear comments from the student more often in the future (and hopefully they will also be good ones!). If a student tells a joke to several classmates and they laugh at it, then the student is more likely to tell additional jokes in the future and so on.

As with classical conditioning, the original research about this model of learning was not done with people, but with animals. One of the pioneers in the field was a Harvard professor named B. F. Skinner, who published numerous books and articles about the details of the process and who pointed out many parallels between operant conditioning in animals and operant conditioning in humans (1938, 1948, 1988). Skinner observed the behavior of rather tame laboratory rats (not the unpleasant kind that sometimes live in garbage dumps). He or his assistants would put them in a cage that contained little except a lever and a small tray just big enough to hold a small amount of food. (Table 4 shows the basic set-up, which is sometimes nicknamed a "Skinner box".) At first the rat would sniff and "putter around" the cage at random, but sooner or later it would happen upon the lever and eventually happen to press it. Presto! The lever released a small pellet of food, which the rat would promptly eat. Gradually the rat would spend more time near the lever and press the lever more frequently, getting food more frequently. Eventually it would spend *most* of its time at the lever and eating its fill of food. The rat had "discovered" that the consequence of pressing the level was to receive food. Skinner called the changes in the rat's behavior an example of **operant conditioning**, and gave special names to the different parts of the process. He called the food pellets the **reinforcement** and the lever-pressing the **operant** (because it "operated" on the rat's environment). See below.

Table 4 Operant conditioning with a laboratory rat

Operant → Reinforcement Press lever → Food pellet	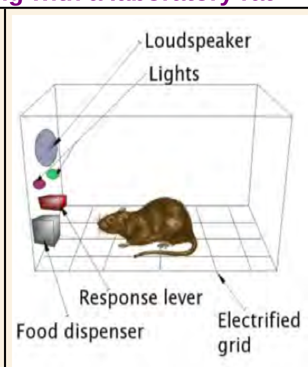

Skinner and other behavioral psychologists experimented with using various reinforcers and operants. They also experimented with various patterns of reinforcement (or **schedules of reinforcement**), as well as with various **cues** or signals to the animal about when reinforcement was available. It turned out that all of these factors—the operant, the reinforcement, the schedule, and the cues—affected how easily and thoroughly operant conditioning occurred. For example, reinforcement was more effective if it came immediately after the crucial operant behavior, rather than being delayed, and reinforcements that happened intermittently (only part of the time) caused learning to take longer, but also caused it to last longer.

Operant conditioning and students' learning: As with classical conditioning, it is important to ask whether operant conditioning also describes learning in human beings, and especially in students in classrooms. On this point the answer seems to be clearly "yes". There are countless classroom examples of consequences affecting students' behavior in ways that resemble operant conditioning, although the process certainly does not account for all forms of student learning (Alberto & Troutman, 2005). Consider the following examples. In most of them the operant behavior tends to become more frequent on repeated occasions:

- A seventh-grade boy makes a silly face (the operant) at the girl sitting next to him. Classmates sitting around them giggle in response (the reinforcement).
- A kindergarten child raises her hand in response to the teacher's question about a story (the operant). The teacher calls on her and she makes her comment (the reinforcement).
- Another kindergarten child blurts out her comment without being called on (the operant). The teacher frowns, ignores this behavior, but before the teacher calls on a different student, classmates are listening attentively (the reinforcement) to the student even though he did not raise his hand as he should have.
- A twelfth-grade student—a member of the track team—runs one mile during practice (the operant). He notes the time it takes him as well as his increase in speed since joining the team (the reinforcement).
- A child who is usually very restless sits for five minutes doing an assignment (the operant). The teaching assistant compliments him for working hard (the reinforcement).
- A sixth-grader takes home a book from the classroom library to read overnight (the operant). When she returns the book the next morning, her teacher puts a gold star by her name on a chart posted in the room (the reinforcement).

Hopefully these examples are enough to make four points about operant conditioning. First, the process is widespread in classrooms—probably more widespread than classical conditioning. This fact makes sense, given the nature of public education: to a large extent, teaching is about making certain consequences for students (like praise or marks) depend on students' engaging in certain activities (like reading certain material or doing assignments). Second, learning by operant conditioning is not confined to any particular grade, subject area, or style of teaching, but by nature happens in nearly every imaginable classroom. Third, teachers are not the only persons controlling reinforcements. Sometimes they are controlled by the activity itself (as in the track team example), or by classmates (as in the "giggling" example). A result of all of the above points is the fourth: that multiple examples of operant conditioning often happen at the same time. The skill builder for this chapter (*The decline and fall of Jane Gladstone*) suggests how this happened to someone completing student teaching.

Because operant conditioning happens so widely, its effects on motivation are a bit more complex than the effects of classical conditioning. As in classical conditioning, operant conditioning can encourage **intrinsic motivation** to the extent that the reinforcement for an activity can sometimes be the activity itself. When a student reads a book for the sheer enjoyment of reading, for example, he is reinforced by the reading itself; then we often say that his reading is "intrinsically motivated". More often, however, operant conditioning stimulates *both* intrinsic *and* extrinsic motivation at the same time. The combining of both is noticeable in the examples that I listed above. In each example, it is reasonable to assume that the student felt intrinsically motivated to some partial extent, even when reward came from outside the student as well. This was because *part* of what reinforced their behavior was the behavior itself—whether it was making faces, running a mile, or contributing to a discussion. At the same time, though, note that each student probably was also **extrinsically motivated**, meaning that another part of the reinforcement came from consequences or experiences *not* inherently part of the activity or behavior itself. The boy who made a face was reinforced not only by the pleasure of making a face, for example, but *also* by the giggles of classmates. The track student was reinforced not only by the pleasure of running itself, but *also* by knowledge of his improved times and speeds. Even the usually restless child sitting still for five minutes may have been reinforced partly by this brief experience of unusually focused activity, even if he was *also* reinforced by the teacher aide's compliment. Note that the extrinsic part of the reinforcement may sometimes be more easily observed or noticed than the intrinsic part, which by definition may sometimes only be experienced within the individual and not also displayed outwardly. This latter fact may contribute to an impression that sometimes occurs, that operant conditioning is really just "bribery in disguise", that only the *external* reinforcements operate on students' behavior. It is true that external reinforcement may sometimes alter the nature or strength of internal (or intrinsic) reinforcement, but this is not the same as saying that it destroys or replaces intrinsic reinforcement. But more about this issue later!

Comparing operant conditioning and classical conditioning: Operant conditioning is made more complicated, but also more realistic, by many of the same concepts as used in classical conditioning. In most cases, however, the additional concepts have slightly different meanings in each model of learning. Since this circumstance can make the terms confusing, let me explain the differences for three major concepts used in both models—extinction, generalization, and discrimination. Then I will comment on two additional concepts—schedules of reinforcement and cues—that are sometimes also used in talking about both forms of conditioning, but that are important primarily for understanding operant conditioning.