Learning and Conditioning Tutorials

Roger D. Ray

Noel Belden

David A. Eckerman

Table of Contents

Learning and Conditioning

When we use simple terms in everyday conversation we don't worry so much about precision. We might talk to a friend about riding in a car but we don't worry about what kind of car our friend is thinking of or whether it's the same car that we're thinking of in every detail. Unless it's important to our story, we don't worry about those details.

For scientists, every detail is important to the story they tell. Precision is much more important for their work than in everyday conversation. Because scientists base their work on understanding the terminology they use, it's very important that terminology in science be clear and precise. Terminology used by scientists is defined as carefully as possible.

When we study concepts like learning we must be sure that we're all talking about the same thing. The goals of science are based on beliefs in reason and logic. Science tries to understand the world around us in objective terms, terms that can help us to understand relationships between events in our world. Because we want to understand ourselves in this world as well, much of science is devoted to the study of human and animal behavior and experience. Thus scientists study learning to understand relationships between stimulus events and behavior.

Learning is one critical process in the behavior of humans and other animals. It is the change in future behavior that occurs after the experience of interacting with certain kinds of environmental events. We all have the ability to learn but our learning is an individual experience. It occurs as a process over our lifespan. Think of the different events of learning in your own life.

When you were very young you went through a period of time where your only way to get around was to be carried by someone else. As you developed the physical capacity, you learned to crawl, and eventually to walk, thus allowing you to behave in new ways. Later, you might have learned to read because of experiences you encountered and a whole new world opened which, again, changed your behavior. As you grew and experienced more and more stimuli in your environment, you learned from these encounters and the process of learning continued. It continues today and will continue throughout the course of your life.

Each time you learn something new it produces a relatively permanent change in your behavior. Once you learn how to walk, you don't forget how unless some kind of traumatic event occurs. Reading is the same way, as is making your bed, riding a bicycle, and driving a car.

The ability to learn is nature's way of equipping us to survive on Earth. Things around us change very quickly so we need a way to adapt to these changes. We don't want to wait until we grow a fur coat to explore the Arctic so we learn to make warm clothes instead. Learning is an adaptation that equips us for survival in many different environments. This is useful because it provides a way for us to survive no matter where we find ourselves. Such an adaptation greatly increases the chances of survival for an individual and a species, which is the whole purpose behind adaptations in general.

In this set of modules you will learn more about the principles of learning and the processes called conditioning. You will learn about the pioneers who made the important discoveries in this field as well as the specifics of their research projects that allowed these discoveries. You will also learn how these discoveries have been applied to real-world problems in education, the home, and even psychological clinics and treatment facilities. And one of those applications involves this computer and internet-based, artificially intelligent, adaptive tutoring system by which you will be able to learn about the principles of learning and conditioning! So let's look at that new form of learning first.

Learning and Conditioning: A Review

When you hear or see the word "learning", what immediately comes to mind? Do you imagine a classroom full of students listening to an instructor lecture on some subject? Do you picture someone puckering and salivating as they think about a fresh lemon being cut into two halves? Does a dolphin act in a marine life park come to mind? How about a child watching a violent movie or a motorist reading a map, do you associate these with the word "learning"? Do you see a rat pressing a bar for food or a pigeon pecking a disk for grain?

Unless you have studied learning in detail, the first example (the classroom lecture) is probably your foremost image when you think about learning. However, by the time you finish studying this chapter, you should easily associate all of the above examples with some type of learning. The chapter first considers the formal definition of learning used by most psychologists who study it. As you will see, this definition includes more than just what takes place in a classroom. Learning is much more general than that, as it includes adaptation of existing behaviors and the acquisition of behavioral skills as well as knowledge (which, to some psychologists, is also a form of behavior). However, while most psychologists agree on a general definition, there are several alternative conceptual and theoretical perspectives on the mechanisms and principles involved in learning a new behavior.

One of these alternative perspectives focuses on classical conditioning processes, procedures, and applications. First identified and labeled by Ivan Pavlov, classical conditioning is based on reflexive behaviors. The procedures for studying classical conditioning illustrate how someone who normally salivates only when food is placed into their mouth (a reflexive action) can learn to salivate at the thought of cutting a lemon in half. There are also important practical applications of classical conditioning, including techniques for treating simple fears or complex phobias in therapeutic settings.

Operant conditioning provides another perspective on the principles of learning new behaviors. Early research on how cats learned to escape from complex puzzle boxes by E.L. Thorndike demonstrated that problem-solving behaviors depend upon consequences in the environment brought about by those behaviors. According to Thorndike's famous Law of Effect, when problems are solved through trial and error, successful behaviors become strengthened or more likely (learned) because they are rewarded with positive outcomes. But what if, in the trial and error process, the correct behavior never comes about?

Another psychologist who studied learning, B.F. Skinner, asked this very question and answered it through experiments with rats and pigeons. Building on the earlier work of Thorndike, Skinner formalized alternative procedures to trial and error for generating new behaviors and labeled these procedures operant conditioning. Like Thorndike, Skinner stressed the importance of behavioral consequences, which he called reinforcements. However, by using a step-by-step procedure known as shaping, Skinner demonstrated how a new behavior could be more quickly developed in an organism through the reinforcement of gradual changes in what the individual already knows how to do. Shaping is how a dolphin learns the complex behaviors often seen in marine park shows.

Operant conditioning has many educational applications besides animal training, however. Operant principles have been applied in the teaching of academic subjects through the use of teaching machines, and these machines have evolved into today's computer assisted instruction. Appropriate social behaviors can also be developed and maintained through operant conditioning using more secondary forms of reinforcement, which may even collectively define what is called a token economy.

Not all psychologists agree on the importance of observable behavior and reinforcements, however. Those working from a cognitive perspective focus more on mental activities and cognitive, or brain, processes that occur in learning. One of the early researchers taking a more cognitive approach was Edward Tolman. Tolman conducted research with rats where he demonstrated that they could learn successful routes through mazes without rewards. He called this phenomenon latent learning. He also performed experiments on learning about space and locations, which he called space learning. Through these studies he developed his concept of the cognitive map, which he viewed as a mental representation of the layout of one's environment. Tolman believed that animals as well as humans learn to navigate through their surroundings by developing such cognitive maps of environmental arrangements.

Another early researcher who laid the foundation for subsequent cognitive perspectives on learning was Wolfgang Kohler. Kohler demonstrated how chimps suddenly appear to solve problems without shaping or trial and error. He called the sudden recognition of a solution to a problem "insight" and felt this occurred as mental processing of abstract and even potential, as opposed to existant, relationships among objects in the environment.

Building on such early cognitive work, Albert Bandura also believed that humans and some animals do not need contact with consequences or reinforcements in order to learn a new behavior. Through his experiments with adults and children, Bandura developed his theory of observational learning. His research showed how people could learn from merely observing another's behaviors and consequences and then imitating that behavior. This explains, for example, how children can learn aggressive behaviors by watching violence on TV especially if violent behaviors are rewarded in what they watch. As we will see in this chapter, cognitive principles have had successful applications in the treatment of depression and anxiety, as in Aaron Beck's cognitive therapy and Albert Ellis' rational emotive therapy.

The final perspective to be discussed is the ecological view of learning. Those working from this perspective don't completely disagree with any of the other views of learning, they simply seek to understand how animals seem to perform some behaviors without having to learn them and how this impacts new behaviors that need to be learned. The migration of birds and the spawning of salmon are examples of highly complex behaviors that seem to have little basis in learning, but may be modified by experience none-the-less.

Martin Seligman reviewed much of the available literature on learning and proposed that animals had a form of biological preparation for easily learning some behaviors, but also had difficulty in learning other behaviors. Building on this concept, John Garcia, who studied the classical conditioning phenomenon of conditioned taste aversion, found that some stimuli are easier for organisms to associate than others. Thus some behaviors are easier to learn because an organism is prepared through evolution to make such associations or to acquire such behaviors. For example, it is much easier (ie., shaping is quicker) to teach a rat to press a lever than to peck a key with its nose for food. It is also easier to train a pigeon to peck a disk than to press a lever for food. This is because evolution has prepared a pigeon to peck and a rat to rear and put its front paws on an object.

Martin Seligman also did research on the interaction of classical conditioning and subsequent operant conditioning. In his research he discovered what he called learned helplessness, which illustrates how an organisms prior learning history can interfere with the later acquisition of new behavior. Seligman's work is ecological in the sense that he focuses on the natural evolution of behavioral processes and how these procedures interact across an organism's individual life span. Work on learned helplessness has been applied to understanding the development of depression in humans and biological preparedness has been applied to animal population control through taste aversion treatments.

The general definition of learning as the study of human and animal behavior and cognition, as well as how each perspective interprets this definition, will be discussed in more detail in the various sections of this chapter. As you will discover, learning reaches beyond the classroom and it is a subject of interest to many psychologists.

Definition of Learning

The formal definition of learning describes the process as "a relatively permanent change in behavior based on an individual's interactional experience with its environment." As such, learning is an important form of personal adaptation. Let's consider each critical element in this definition. Behavioral change occurs in all animals, both human and non-human, and is a process of personal, or ontogenic, adaptation that occurs within the lifespan of each individual to make one's survival more likely. To say that learning is relatively permanent is to emphasize that behavior is flexible and not genetically pre-programmed in form or function. Learned behaviors may exist for a lifetime, but they may also not appear throughout an individual's life.

Experience of, or interaction with, the environment that precedes and follows behavior presents the adaptational requirement and consequence of each interaction. An individual placed in a bubble and kept from any contact with variations in stimuli from the day it is born does not learn many behaviors. The actions of such an organism in this case would be very limited. A living creature may barely survive such an existence.

Let's expand on each of the critical elements in the definition of learning a bit more. Because learning is so intertwined with individual and environment, it is often emphasized as one of the two major forms of biological adaptation. Ontogenic adaptation, the basis of learning, creates behavioral change that is unique for each individual and the process only occurs within the lifespan of that individual based on that individual's experiences with its personal environmental interaction history. This is in contrast to phylogenic adaptation, which creates the shared features that define all members of each species and thus transfers from one generation to the next via genetic transfer and genetic determination. Stressing individual-environmental interaction points out that the environment brings about changes in behavior just as behavior then brings about changes in the environment. Many psychologists believe that organisms learn to adapt to environmental challenges as well as learning to adapt (change) environments to better meet individual survival and comfort requirements. We both create our environments and are created by our environments.

We can observe the process of learning by noting changes in behavior or even the development of new responses through these interactions or experiences with the environment. For example, let's reflect on how you may have learned to ride a bicycle. You may have been very young, and had surely already mastered the various ways of getting from one place to another by first being carried, then learning to crawl, then walk, then even skip or run.

You probably progressed to various other ways of getting around, such as pedaling a tricycle and later perhaps even a bicycle with two added training wheels to help you learn the balancing difference of pedaling from more of a standing position rather than the lowered seating of your tricycle. On these machines there wasn't much to master other than steering and pedaling correctly. Then one day you were faced with riding a bicycle without training wheels! Suddenly you found balancing was far more challenging than you ever imagined.

But a few repeated efforts and possibly some other-person support to get you moving quickly showed you that balancing on two wheels was more a matter of having the bike moving than anything else. So you soon learned to pedal and mount simultaneously something you had never had to do with your tricycle or with your training wheels. So you now discover that you have a new and relatively permanent skill.

You'll probably always be able to ride a bike so long as you have the physical bodily and balancing requisites. But you may find that as you began driving cars, riding a bike (like riding tricycles) isn't something you have actually done for a very long time. It has been abandoned in favor of an even more adaptive and less strenuous mode of longer distance transportation.

But then one day you discover riding as a sport or exercise! Suddenly it isn't transportation anymore and, in serving quite a different purpose, riding a bicycle may reappear or disappear as life style and recreational opportunities constantly change. So as environments change, so does the use and purpose of the learned behavior of riding a bike. That's the relativity of the persistence or permanence of use of the behavior.

But not all behavior is as obvious as riding a bicycle. Suppose you decide to take a shortcut when biking to school one day and a very large and intimidating dog suddenly barks and chases your bike as you ride by. You become more than a bit aroused and feel the rush of adrenaline immediately as this happens. If you decide not to take this route the next day, one behavioral change is obvious. You altered your path of riding. However if you continue to take the shortcut the next day that behavior has not changed. But you may feel quite anxious and cautiously look for the dog to appear again when you reach the critical point on your path where he appeared yesterday. That is an alternative form of your adaptation, and it is more emotional behavior vs. skill in riding.

These emotional changes are also only relatively permanent because if the dog doesn't appear again over several days, you suddenly find you no longer fear that part of your ride. But don't be surprised if you may one day suddenly find yourself looking for a dog again when you reach that point in the path where you suddenly remember having the initial fearful experience. It may be more permanent than you once thought even though you haven't experienced it for some time, as any veteran soldier is likely to tell you concerning the trauma of experienced war events!

Principles of Phylogenic Adaptation

Learning is often conceptualized as a form of adaptation. But adaptation includes more than learning. Adaptation involves 1) changes in an individual's behavioral repertory that occur in that individual's lifespan (ontogenic adaptation or learning) and 2) changes in species-specific behaviors and anatomical structures that are transferred from one generation to the next through genetics (phylogenic adaptation or evolution). Ontogenic adaptations result from an individual's personal interactions with its environment (Skinner, 1966). Because ontogenic adaptation is the topic of this entire chapter, this section will elaborate mostly on phylogenic adaptation as a contrast to learning as a form of adaptation.

Phylogenic adaptation is the slow process of change in the anatomy of all members of a species. It results in response to biologically important problems posed by the environment. All members of a species can gradually develop to share anatomical as well as behavioral answers to survival problems and this is phylogenic adaptation. Those members of a species that are endowed to successfully survive pass on their genes to new generations, while those who are unsuccessful die off.

Species-specific behaviors, like bird songs or migratory patterns, can also be considered phylogenic because they take generations to develop and are largely determined by shared genetics, rather than individual experiences. Such species-wide adaptations are the foundation for the evolutionary changes of a species. Thus, phylogenic adaptation is not considered learning because it occurs for all members of a species rather than an individual and it deals with genetic transmissions from generation to generation. Phylogenic transformations result mostly from genetic mutations, which determine anatomical or behavioral changes. The changes that succeed are then transferred to the next generation because individual genetic carriers survive environmental challenges long enough to breed, and thus propagate the species. Many psychologists believe that the ability to learn is itself essentially a phylogenic adaptation.

Charles Darwin was a 19th century naturalist who described phylogenic adaptive processes in great detail --a process he called evolution. Darwin's theory of evolution (Darwin, 1859) includes the processes of phylogenic adaptation through natural selection. As the environment changes, genetic mutations in a species result in phylogenic adaptations (anatomical or behavioral) to these changes. Those members of the species who do not develop adaptations do not survive and those with appropriate genetic adaptations continue to reproduce. This is natural selection. As the environment continues to change, natural selection continues to work and the species continues to evolve and change as well.

While his primary focus was on phylogenic adaptations impacting anatomical structure, Darwin also conducted studies on how a species adapts by changing behavior and expressions of emotions (Darwin, 1872), as well as on how individuals of a species adapt to unique biologically important problems. He observed organisms in their natural environment and rarely interfered with the activities of the animals he watched. Instead, he relied mostly on descriptive research methods by taking notes, collecting specimens, and carefully recording his observations. Based on these techniques Darwin developed theories that were not well received in his time, but they have gained enormous popularity and scientific support today.

Later other naturalists, also using largely descriptive methods, began to focus on species-specific behaviors rather than anatomical structures. Ethologists (scientists who focus on studying species-specific behavior) such as Niko Tinbergen (cf., Tinbergen, 1951) and Konrad Lorenz (cf., Lorenz, 1955) were awarded the Nobel Prize in Physiology for their studies of such behavioral patterns specific to a species. They identified complex behavioral sequences, called fixed-action-patterns, which are involved in such species-specific functions as mating and territorial defense. They also discovered imprinting as one form of individualized adaptation to environments that occurs at critical periods of development shared by all members of a given species. Such imprinting processes are important in developing the attachment between young offspring and their mothers that results in the offspring following their mothers from place to place.

Phylogenic adaptation is a slow process because it takes many generations of organisms in order to develop significant changes in anatomy or species-specific behavior. Species-specific behaviors like the fundamentals of a birds songs and migratory patterns took thousands of years to develop into what we see today. The neck of a giraffe and the opposable thumb in humans are phylogenic changes in anatomy that also took millenia to be realized.

Classical Conditioning Overview

Classical conditioning was first identified and developed by a Russian physiologist, Ivan Pavlov (1927/1960). The phenomenon of classical conditioning is widely considered to be the most fundamental form of learning. Even before Pavlov

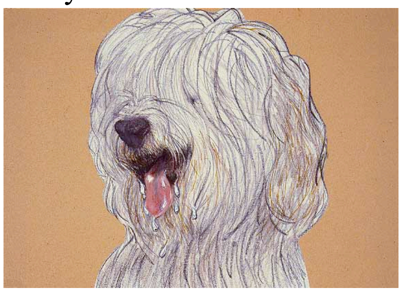

identified the process of conditioning, his work was monumental. In fact, Pavlov was awarded the Nobel Prize in Physiology and Medicine for his research on the digestive system of dogs. To pursue his digestion research, Pavlov developed a procedure for surgically implanting a tube, called a fistula, into living animals. This allowed him to collect and measure digestive secretions, such as those released in the stomach or the mouth. This was a first, because up until Pavlov's innovation, almost everything that was known about physiological processes was obtained from studies involving acute (temporary or sacrificed animals) as opposed to chronic preparations (long-term and in living animals).

Pavlov was interested especially in the mechanisms of reflexive secretions when food was placed into the mouth and as it passed through the other parts of the digestive system, including the stomach. For example, Pavlov or one of his assistants would place meat into the mouth of a dog and then measure the amount of saliva that passed through a salivary fistula implanted for collecting saliva in a test tube pasted onto the outside of a dog's cheek.

With the aid of his fistula preparations Pavlov made a very surprising discovery.

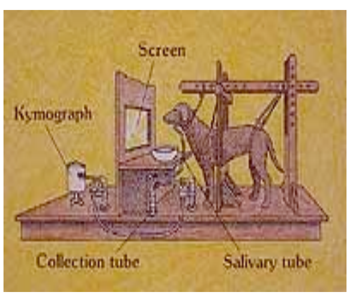

He noticed that his dogs began to salivate upon merely seeing Pavlov's lab assistant entering the room or getting out the food pans, even before any food was presented! Pavlov thought it peculiar that a reflex such as salivation should be present with no apparent stimulus, such as food, to trigger it. Pavlov discontinued his digestion research and focused exclusively on this new and curious phenomenon he originally called "psychic reflexes."

Through further investigation Pavlov discovered that his psychic reflexes developed through a process he could control and manipulate. He called this process "conditioning" because it defined the conditions under which reflexes would occur to previously ineffective, or "reflex-neutral" stimuli. In this process Pavlov began to substitute highly specific manipulated stimuli as alternatives to the less controlled entry of his lab assistants or presentations of empty food pans. Pavlov turned to the use of such specific and controllable stimuli as the sound or even the specific rate, of a metronome's ticking. In Pavlov's classic experiments, from which the process gets its name of Pavlovian conditioning, or classical conditioning, a stimulus is first tested to assure it is "neutral," in that it does not elicit salivation. Then that neutral stimulus is presented along with a stimulus, such as food, that is known to elicit salivation. After a few repetitions of this temporal pairing of the two (neutral and eliciting) stimuli, the previously neutral stimulus is found to be no longer neutral, but now will elicit salivation when presented by itself.

Early philosophers, such as Aristotle, had emphasized the importance of temporal associations for acquiring or learning new actions but a fully developed Associationism philosophy proposing that associations were the form of virtually all learning came in the 17th and 18th century's philosophical movement known as British Associationism. This movement largely began with John Locke (1632-1704) and eventually included other British philosophers such as David Hume and James Mill (Burtt, 1939). The Associationists suggested a set of three principles that they felt established the foundation for the formation of associations in human thought. The Principle of Contiguity stipulated that associations were formed between events that occurred together in time. The Principle of Frequency stated that the more often two events occurred together, the more strongly associated they would be with one another. And finally, the Principle of Intensity referred to the strength of the feelings that accompanied the association.

But it was Pavlov (1927/1960) who was one of the first to study associations objectively and empirically (scientifically) and to give associations an importance in the development of new physiological and emotional reactions as well as mental activities. Classical conditioning is based on reflexive responses and the associations between stimuli that do and do not naturally elicit those reflexive responses. Pavlov investigated many details of how neutral and reflex-eliciting stimuli can be variously paired in time, and thus defined several alternative procedures that are variations on classical conditioning.

These procedural variations on classical conditioning include simultaneous conditioning, delayed conditioning, backward conditioning, trace conditioning, temporal conditioning, and extinction, as well as differential conditioning and its related phenomenon, stimulus generalization. All of these procedures involve different ways of pairing in time classical conditioning variables, which include an unconditional stimulus (UCS) such as food, and a neutral stimulus (NS) such as a ticking metronome that doesn't elicit salivation. As Pavlov discovered, when paired with a UCS like food this neutral stimulus gradually becomes an effective elicitor for salivation, and thus can then be called a conditional stimulus (CS). The effects of such temporal pairing are seen in the development of a conditional response (CR), such as salivation when the metronome alone is ticking. The original salivation to food (UCS) is reflexive and thus requires no pre-conditions to establish its elicitation function. Thus Pavlov called this reflexive form of salivation the unconditional response (UCR).

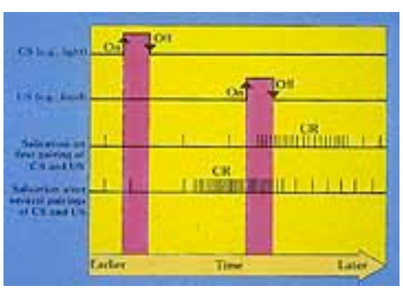

Pavlov's classical conditioning has many applications, including the development of emotions and consumer attitudes through temporal parings of brand names or objects, like cars, with evocative, sexually arousing, or fun-related stimuli in advertising. Classical conditioning also has applications in therapeutic environments. For example, classical conditioning procedures are fundamental in our understanding and treatment of phobias. Research on conditioned emotional responses has led to a better understanding of how phobias and addictions form. Working from this knowledge psychologists have also been able to develop therapies called systematic desensitization, aversion therapy, and counter conditioning, to reduce or eliminate these emotional problems. Farmers use principles of conditioned taste aversion (which stem from classical conditioning procedures) in order to keep predators from attacking their flocks. This is more ecologically sound and humane form of predator control than extermination of the predator population. Classical conditioning has even been applied in the field of medicine where immune responses are conditioned so that a patient takes less medication with the same immune boosting effects.

Variables in Classical Conditioning

There are several variables involved in classical conditioning and it is important to understand how they are labeled and used in the conditioning process. Pavlov was the first to identify and label these variables, and his terms are still widely used. The foremost independent variable is the unconditional stimulus (UCS), such as food. Its associated dependent variable is the response it elicits, called the unconditional response (UCR). In Pavlov's typical research this UCR was salivation.

A second independent variable, as it exists prior to conditioning, is called the neutral stimulus (NS), but when paired successive times with the presentation of the UCS, this NS gradually acquires the function of also eliciting a response similar to the UCR, and in this state the stimulus is called the conditional stimulus (CS). An example of a NS evolving into a CS is when a ticking metronome elicits no salivation prior to conditioning (thus defining the stimulus as "neutral," or a NS), but eventually it comes to elicit salivation (thus becoming a CS) because it has been paired in time with food (the UCS). Such stimulus pairing procedures eventually cause the dependent variable (salivation) to appear as a conditioned response (CR) even if only the CS (metronome) is presented by itself (Pavlov, 1927/1960).

The unconditional stimulus (UCS) is any stimulus that naturally elicits or brings about a specific unconditional response, thus making it reflexive in nature. In Pavlov's research the UCS was typically food, and it naturally brings about the reflexive response of salivation. As another example, a loud sound could be an UCS for a startle reaction as the UCR. The UCS qualifies as an independent variable because Pavlov manipulated its presence in all of his experiments, including in his physiological research of digestive reflexes that predated his work on conditioning.

You may have already noticed that when we describe the effect of the unconditional stimulus (UCS), and eventually the conditional stimulus as well, we use the term "elicit." To elicit a response means to reliably cause that response to occur. A reflexive behavior (UCR) is described as being elicited because it reliably occurs in response to a particular stimulus (UCS). Without this stimulus the response rarely occurs, thus making the UCS necessary and sufficient to produce the response. Thus, for example, you can't startle someone without such a stimulus and try as you may, you cannot bring about the startle response in yourself except by sheer accident (the stimulus is otherwise not unexpected). Because of this, we say that a strong, unexpected stimulus elicits, or causes, the startle response, and this relationship between stimulus and response defines what is meant by the word reflex. In Pavlov's work in classical conditioning, meat elicits salivation.

The amazing finding for Pavlov was, however, that after repeated parings of a metronome, bell, or light NS with the presentation of meat, the NS became a CS and thus began to elicit salivation as well. Thus we say there is a conditional reflex that has been established through the stimulus pairing, or "conditioning," procedure. We learn to react as if reflexively to stimuli that usually do not elicit a reflexive response if the stimuli are contiguous with a stimulus that does elicit a reflexive response. That is the essence of Pavlov's conditioning discovery (Pavlov, 1927/1960).

Use of an UCS automatically elicits the dependent variable in classical conditioning: the unconditioned response (UCR). The unconditioned response is the generic label for the reflexive behavior elicited by the unconditioned stimulus (UCS). In much of Pavlov's conditioning research, the UCR was salivation. However, he also investigated many other forms of reflex relationships beyond food eliciting salivation. Nevertheless, salivation is his most typical response, and may be described as a dependent variable because Pavlov measured salivation using his fistula preparation and because salivary flow is dependent upon the presence of the independent variable, food.

The neutral stimulus is also a very important variable in classical conditioning. A neutral stimulus is any stimulus that does not elicit the reflexive response, or UCR. A metronome does not normally elicit salivation, so in Pavlov's experiments the metronome begins as a neutral stimulus. A metronome could be an UCS for ear pricking behavior, however. So a stimulus is not always an absolutely neutral stimulus, it is only neutral with respect to the response under investigation as the dependent variable, such as not eliciting salivation. So it depends on what reflexive behaviors you are focusing on as to whether a stimulus may be an UCS or a NS.

After appropriate classical conditioning procedures, or CS-UCS pairings, have occurred several times, the neutral stimulus gradually comes to elicit a response that typically resembles the UCR. In this case it is no longer neutral. Thus Pavlov described it as a conditional stimulus (CS) because his experimental conditions had created a new elicitational function for this previously neutral stimulus. A metronome's ticking that elicits salivation after conditioning is a conditional stimulus (CS) for the conditional salivary reflex (CR), since salivation will occur even if the food isn't presented (Pavlov, 1927/1960).

The CS is a classical conditioning label that applies only after conditioning procedures have been used for a sufficient number of trials required to obtain a conditional reaction to that CS. A metronome that is ticking and not eliciting salivation is a neutral stimulus. A metronome that is ticking and subsequently elicits salivation after being paired with a UCS for a few trials is now a conditioned stimulus. It is important to keep in mind as you read about classical conditioning procedures that while the metronome is the same physical stimulus both before and after conditioning, psychologically it is neutral before conditioning and becomes conditional only after conditioning trials are experienced. As such, a stimulus represents two separate functions for the same variable at different stages of the experiment.

Once a CS has the power to bring about a resemblance of the UCR, this new response is labeled a conditional response (CR). Salivation in response to a metronome ticking is a CR, because prior to conditioning, salivation is not elicited by a metronome ticking. As in the case of the CS, it is important to remember that while salivation may appear to be the same response before and after conditioning, it is an unconditioned response before conditioning and a conditioned response after, depending on which type of stimulus elicits it. Hence, the same apparent response serves as two different functional variables and how it is labeled depends upon whether classical conditioning procedures have occurred or not and upon which stimulus (CS or UCS) is eliciting it (Pavlov, 1927/1960).

Later research actually has demonstrated that the CR only appears to resemble the UCR, but even in the case of salivation, the chemical compositions may not be exactly the same for the two forms of saliva. And when the reflexive reaction of the cardiovascular (heart) system to a startling noise is measured, the UCR is a sharp acceleration in heart rate while the CR is actually a deceleration in heart rate! So in this case the CR appears to function more like an "anticipatory reaction" than the actual reflexive response we call the UCR (Deleon, 1964). It is on this as well as other bases that Rescorla, a modern researcher specializing in classical conditioning procedures, has interpreted the critical aspects of Pavlov's procedures to be the existence of an actual stimulus contingency (if-then) relationship between CS and UCS rather than simple associative contiguity in time (Rescorla, 1988). We'll revisit this stimulus-contingency interpretation in a subsequent section on ecological perspectives in learning.

Finally, one of the most important independent variables in classical conditioning is the time element used to define how the two stimuli occur together. That is the NS/CS occurs together with the UCS in time. But the actual timing has been manipulated and explored for its own effects on the conditioning process. Manipulation of the time variable becomes complex, in that there are many variations of how the two stimuli can appear and still be contiguous, or occurring at, or nearly at, the same time (Pavlov, 1927/1960). These variations of timing define alternative classical conditioning procedures, and it is to these various procedures and the role of time in their definitions that we now turn.

Time – based Paradigms in Classical Conditioning

The defining procedure for classical conditioning is that of establishing a stimulus contingency, or if-then relationship between two stimuli which are used as independent variables in the conditioning process. Critical to this definition is the fact that one stimulus at the beginning of the procedure is neutral in function (NS/CS). That is, it does not elicit the reflexive response being investigated for conditioning. The other stimulus, both from the beginning and throughout the procedure, is an effective elicitor (UCS) for the reflexive response (UCR) being conditioned. In his initial research, Pavlov (1927/1960), identified such a procedure as one which involves the temporal pairing of the neutral stimulus with the UCS. Through repetitions of these pairings multiple times (each time constituting a "trial") the neutral stimulus comes to elicit the target reflexive response and is thus transformed functionally into a conditional stimulus (CS). This usually takes repeated trials, as the neutral stimulus rarely elicits a CR after only one pairing.

Pavlov eventually explored many different procedural variations for presenting the stimuli involved in his original classical conditioning process. These procedural variations, which define alternative classical conditioning paradigms, all are based on how time varies in the presentation of the two stimuli, and they include delayed conditioning, simultaneous conditioning, backward conditioning, trace conditioning, temporal conditioning, and extinction. Pavlov also explored paradigms that did not rely on temporal variations between the CS and UCS. These include stimulus generalization and discrimination as well as related effects--all of which will be described in a separate section because they don't use time as the critical defining procedural variable.

Delayed conditioning was actually Pavlov's initial procedure, which was fortunate based on subsequent findings that most other paradigms are not very effective, if at all, for developing a conditional response. In the delayed procedure, Pavlov actually started the ticking of his metronome, the CS, a bit before he presented the food (UCS). The metronome continued to tick from just before the presentation of food and continued ticking throughout the dog's eating. The critical aspect was the slight delay between first presenting the metronome and the subsequent presentation of the food. It is from this temporal delay that the delayed conditioning procedure derives its name (Pavlov, 1927/1960).

Eventually many variations of temporal delay between presenting the NS/CS and the UCS, technically called the inter-stimulus interval (ISI), were investigated and an optimum ISI (or time between the two stimuli) was discovered. Pavlov found that conditioning is most effective if the UCS is presented .5 to 1 second after the presentation of the neutral stimulus (NS). Pavlov found that when more time exists between the two stimuli conditioning is weak, if present at all. Subsequent research has found this rule to vary depending upon what response system is investigated as the dependent variable (conditional and unconditional response). Pavlov was measuring salivation, and his rule of .5 to 1 second optimal delay is true for that form of response.

But if heart rate is being classically conditioned, the .5 to 1 second optimal ISI changes and an ISI of up to 5 seconds or even more delay may be used for highly effective conditioning results (Church |_2 Black, 1958). As we will see in our discussion of ecological perspectives on learning, conditioned taste aversion research by John Garcia has found that delay intervals of several hours may still result in effective conditioned responses being developed (Garcia, Kimeldorf, Hunt, |_2 Davies, 1956). So the optimal time interval between NS/CS and UCS for effective conditioning all depends on what response is being measured and what ecological role that response plays in the physiological functioning of the individual or species.

This discussion of optimal ISI delay periods suggests that several alternative procedures besides delayed conditioning might be investigated. Simultaneous conditioning is one example of a number of these possible variations. It involves the presentation of the neutral stimulus and the unconditional stimulus simultaneously (or at least under the effective requirement of being .5 seconds apart). Essentially, in simultaneous procedures the metronome and the food would be presented at exactly the same time to a dog. Of course, with the food always present, it isn't possible to tell whether salivation is occurring to the metronome or the food. But when Pavlov later tested the NS/CS by presenting the metronome alone, he found that the simultaneous procedure was not very effective for establishing conditioning. The metronome turned out to be generally ineffective in eliciting the UCR of salivation under simultaneous conditioning procedures (Pavlov, 1927/1960).

Backward conditioning is another classical conditioning procedure that is defined from Pavlov's manipulations of the temporal relations between unconditional and neutral stimuli. In the backward procedure the neutral stimulus is presented only after the UCS is presented, usually in the same .5 -1 second time interval that is used in the delayed classical conditioning procedure. As in simultaneous conditioning, even after repeated pairings, the neutral/conditional stimulus is very weak and very unreliable, if effective at all, in its ability to elicit any conditional response as the result of backward pairing procedures.

A procedure, called the trace conditioning procedure, has also been explored whereby the NS/CS is presented and then terminated prior to the presentation of the UCS. Thus there is actually no time where the two stimuli are both present, but rather the UCS comes a short time after the NS/CS has already been terminated. As one might expect, this is not a very effective conditioning procedure, even though the example of lightning and thunder being associated is often mistakenly used as an effective illustration of classical conditioning. For one to come to fear lightning, the thunder clap must occur quite soon after the lightning has disappeared, thus assuring a minimum "trace interval" between the two.

Pavlov (1927/1960) even began to wonder if time itself could be used as if it were a stimulus in creating an effective conditioning procedure. In his many explorations he discovered that highly predictable (that is, regular or equal) time intervals between presentations of food alone would cause the animal to begin salivating when the appropriate interval of time had elapsed, even though no food was presented during such a "test" trial. It appears that the time interval itself is sufficient for the animal to demonstrate conditioning, and thus this procedure is called temporal conditioning. Many farmers are well aware that a regular feeding schedule will create quite a behavioral stirring or restlessness among livestock if a given feeding is a bit late, and this is an example of temporal conditioning.

Extinction and Spontaneous Recovery in Classical Conditioning

Pavlov (1927/1960) investigated what would happen if the CS were presented for a number of trials without the presence of the UCS used for conditioning, but only after successful conditioning had already been established. He quickly discovered that the CR diminishes, and eventually disappears. This procedure is called extinction. The critical elements in the gradual disappearance of the CR is the occurrence of a conditional stimulus (for example, a metronome) that is no longer associated with an unconditioned stimulus (for example food) after successful previous conditioning trials.

In early extinction trials where the conditional stimulus occurs alone, the CS continues to elicit a conditional response. However, after repeated presentations of the CS alone, Pavlov found that the conditional response gradually diminished until it no longer occurred at all. It might reappear briefly on subsequent testing days -- a phenomenon called spontaneous recovery -- but that also quickly disappears. It is the parallel between conditioning and extinction as adaptive behavior suited to changing circumstances and the extinction of species when their behaviors are no longer viable for survival that gives the procedure extinction its name. The diminished conditional response or its total disappearance is extinguished behavior that Pavlov thought was parallel to an extinguished species because it was no longer adaptive (Pavlov, 1927/1960).

Pavlov also found that, following apparently complete extinction of a CR, if the UCS is paired with the CS again that CS quickly (often after only one pairing) regains its ability to elicit the CR again. The reappearance of a CR to the testing presentation of a

CS would occasionally occur even without reconditioning, thus appearing to be a spontaneous recovery of the prior conditioning effect. It is from this reappearance that the phenomenon is called spontaneous recovery.

Other classical conditioning procedures were also explored that did not rely upon time as the critical variable, but instead relied upon similarity of other stimuli to that used as the NS/CS. These procedures pair one stimulus as a potential CS with a UCS, but other somewhat similar stimuli are presented without the UCS, thus incorporating extinction procedures for stimuli that had never been paired with the UCS, but because of their similarities to the CS evoked a CR--a phenomenon Pavlov called stimulus generalization. It is to this unique set of procedures of differential conditioning and the results of their use, called stimulus discrimination, that we shall now turn.

Stimulus Generalization and Discrimination in Classical Conditioning

Pavlov made many interesting discoveries as he continued to explore alternative classical conditioning procedures, including some procedures that did not rely on time as the altered variable. For example, Pavlov (1927/1960) investigated how his dogs might respond to stimuli that should be neutral, because those stimuli had never been paired with an UCS. In one variation of these conditions he noticed that following successful conditioning which established a reliable CS-CR relation, if he presented stimuli that were both different from, yet similiar to, the original CS, these differing stimuli would elicit at least some amount of a conditioned salivation response even though these stimuli had never been present when food was available. For example, Pavlov's dogs responded with salivation to many different rates of a metronome ticking even though only one rate of ticking was used during conditioning. But the more dissimilar the tested rate was from the rate used for original conditioning, the less was the amount of salivation observed. This phenomenon is called stimulus generalization.

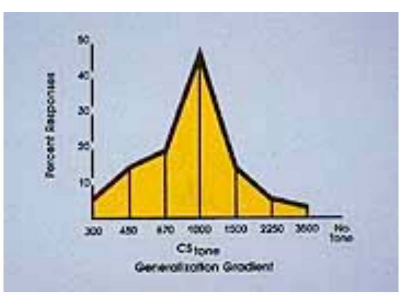

Stimulus generalization testing involves presenting test trials where many variations of stimuli are used, as in many different rates of metronome ticking, and where each test stimulus variation differs somewhat from the original CS. However, these test stimuli are never paired with the UCS as the original CS was. Systematically testing many variations similar to the CS reveals a bell-shaped, or "normal" curve of declining amounts of salivation as the stimuli become more dissimilar from the original CS. This curve is called the stimulus generalization gradient (see illustration).

Pavlov then explored whether or not the animal would extinguish the partial responding to such similar stimuli. Repeated test trials were presented using one, and only one, rate of metronome ticking that was different from the CS and was never paired with food as an UCS. These "extinction" trials alternated with continuing conditioning trials where the original CS was presented and was still temporally paired with food as an UCS. In such conditions the similar CS that is not paired with the UCS is referred to as the CS- (the negative indicating "not paired") and the CS that continues to be paired with the UCS is referred to as the CS|_3 (the plus indicating "is paired").

This procedure defines what Pavlov called a "differential conditioning" procedure. This name comes from the experimenter's intent to test whether the animal can eventually learn to respond "differently" to the two "different" but similar stimuli. Explorations of many different variations of CS- stimuli were used for differential conditioning and revealed that presenting food only in the presence of one rate of metronome ticking (CS|_3 conditions) while another rate was presented several times and always without being paired with food (i.e., CS- conditions) results in extinction of responding to the CS- while continuing to respond to the CS|_3.

With such differential conditioning procedures Pavlov found that stimulus generalization is significantly altered. Both the CS- and other stimuli more similar to the CS- will fail to elicit the CR at all, while the CS|_3 and stimuli very similar to the CS|_3 continue to show conditioned responding. When differential conditioning first begins a CR occurs to many variations similar to the originally paired CS|_3, but after repeated extinction trials for the CS- the CR occurs only to those rates very close to the one being paired with food.

Pavlov noticed something which he considered highly significant during differential conditioning sessions involving a CS- that was extremely similar to the CS|_3, thus presenting the animal with a very difficult stimulus discrimination task. By "very difficult" we mean presenting two stimuli that have only the slightest differences, but only presenting food with one of them --as when an elliptical shape becomes very similar to a true circle shape and only a presentation of the circle is paired with food. During such difficult discrimination training sessions, Pavlov noticed that his dogs would become highly agitated and difficult to handle. Some even develop stomach ulcers and skin sores. Pavlov saw a parallel between this psychological source of physical illness and human psychological abnormalities and thus labeled this phenomenon "experimental neurosis."

Experimental neurosis made research measurements very difficult. The dogs would twist around and try to free themselves from the harnesses. They would also bite and develop painful sores on their bodies that were sensitive to touch. Pavlov realized that such symptoms as rigidity, agitation, skin sores and gastric ulcers were also observed in human individuals labeled as "neurotic" in his time, hence the name experimental neurosis. Outside of the laboratory, the dogs would be inactive and antisocial, just as some "neurotics" were. These problems occurred only during highly difficult discrimination, or "conflicting" tasks however, so this experimental form of conflict became a phenomenon studied extensively by Pavlov in both his animal laboratory and a human clinic he also maintained. This was the first of many subsequent applications of classical conditioning, and it is to some of these other forms of application that emerged following Pavlov's pioneering discoveries that we shall now turn.

Classical Conditioning Applications

Pavlov's detailed investigation of classical conditioning paradigms eventually led him to explore potential implications and applications of classical conditioning for the world outside of the laboratory. For example Pavlov explored whether classical conditioning might be the source of various types of abnormal behavioral problems in humans as well as animals. If he could prove that such problems exist as conditional responses then classical conditioning paradigms, such as extinction, might offer potential treatment strategies for treating these behavioral problems. His work on what he described as experimental neurosis, seen when animals were subjected to extremely difficult discrimination problems during differential conditioning, has already been discussed in the previous section on discrimination and classical conditioning. Later researchers following Pavlov's interest in classical conditioning processes found even more practical applications of the procedure. One of the earliest was the team of John Watson and Rosalie Rayner (1920), who investigated the development of conditioned emotional responses, as sources for everyday fears and, eventually, possibly even more extreme fears called phobias.

Watson and Rayner studied not only how fears developed, but also how they generalized to novel stimuli that a person might never have experienced. In their classic demonstration of this phenomenon, they exposed infants to furry animals (as the NS/CS) paired with loud noises (as the UCS) that elicited startle responses and crying (UCR). Their most famous subject, an infant named Little Albert, served as a model for how phobias might develop in all humans. As a result of the stimulus contingency Watson and Rayner established between the furry animal (NS/CS) and the loud noise (UCS), Little Albert began to cry and move away (CR) from furry animals (now a CS) that, prior to conditioning, had been quite neutral (NS) for Little Albert (Watson |_2 Rayner, 1920).

John Watson subsequently went on to become a very successful and significant figure in the American advertising world where he used his knowledge of classical conditioning to change consumer attitudes and purchasing behaviors through stimulus contingencies that appear in various forms of advertising. Thus advertising--where instead of fears, highly positive feelings toward a product or brand are conditioned--has become one of the most pervasive classical conditioning applications in modern societies around the world (Watson, 1936).

Other classical conditioning applications have focused on preventing coyotes from killing livestock. Farmers have used basic laboratory work by John Garcia on conditioned food aversion (Garcia, Kimeldorf, Hunt, |_2 Davies, 1956). In these farming applications chemically laced sheep carcasses are used to make coyotes sick enough to avoid eating similar looking live animals in the future (Gustavson, Garcia, Hawkins, |_2 Rusiniak, 1974). Conversely, doctors have used classical conditioning in conditioned immune response procedures to allow a patient to have an optimally functioning immune system with the least amount of medication possible (Buske-Kirschbaum, Kirschbaum, Stierle, Jabaij, |_2 Hellhammer, 1994).

Classical conditioning procedures have psychotherapeutic value as well. Phobias are often treated with a process stemming from classical conditioning called systematic desensitization (Wolpe, 1958). Alcoholism and other addictions are also treated using a form of classical conditioning (Forrest, 1985). From research on fear to applications in advertising, treating phobias, and keeping coyotes from killing sheep, the processes discovered by Pavlov have proven crucial in our understanding of how organisms learn and respond physiologically and emotionally to events around them. This understanding has led us to better therapies and effective uses of classical conditioning techniques.

Finally, there is another important application for the stimulus contingency that exists between a CS and an UCS in classical conditioning. Virtually any UCS will also have reinforcing functions for another form of conditioning called operant conditioning. And any stimulus that predates an UCS in a reliable manner, as the CS does in classical conditioning procedures, will also come to have a similar reinforcing functionality called conditioned reinforcement (also known as secondary reinforcement). Thus one of the most important applications of classical conditioning is this development of the CS's conditioned reinforcement function. This allows a CS to serve as a reinforcer in operant conditioning and thus broadens the scope and power of that conditioning process very significantly. We will revisit this other application of classical conditioning, especially as it relates to what is sometimes called "magazine training," when we discuss how one can develop a new behavior using a special process in operant conditioning called response shaping (Skinner, 1938). So expect to read more about this unique application of classical conditioning in a subsequent section that explains magazine training in its context of operant response shaping and conditioned operant reinforcement.

Conditioned Emotional Responses

As noted in the introduction of classical conditioning principles, one of the more significant applications of Pavlov's procedures is an understanding how human fears develop and generalize. To demonstrate, empirically, that fear can be conditioned, Watson and Rayner (1920) performed an experiment that continues to be a classic demonstration in the history of psychology.

First, Watson and Rayner allowed an 11-month-old child, nicknamed Little Albert, to play with a white, lab rat. While he was playing with the animal, the researchers produced a loud noise (usually with loud cymbals) behind Little Albert. This startled him and he would cry. After several pairings of the rat (neutral stimulus) and the noise (an UCS for bringing about startle response and crying--the UCR), Little Albert began to cry (now a CR) at the sight of the rat (now a CS) without the noise. Little Albert also cried at the sight of a white rabbit and a furry Mardi Gras mask (an example of stimulus generalization), but not at white paper or white cloth (stimulus discrimination). This experiment was one of the first to demonstrate the role of conditioning in the origin of fear and phobias.

This study was incredibly important to our understanding of fear and how emotions can be conditioned. However, some feel that this knowledge came at too high of a price for Little Albert. After the experiment, Little Albert never came back to the laboratory and no one knows how the experiment effected him as he grew older. Watson and Rayner's study, while very powerful and important, would be considered unethical today and would not be allowed to commence (Watson |_2 Rayner, 1920).

Fear is one of the major emotions studied in the area of conditioned emotional responses, but it is not the only emotion that can be transferred from being elicited by truly unconditional stimuli to being elicited by conditional stimuli because of past pairings between the two. Virtually any emotion can become conditioned. For example, if you were given a gift for a major achievement, you may come to have positive and joyful feelings whenever you look at the gift in the future because of its pairing with such a happy time in your life. While the focus of this section is on fear, it is important to remember that conditioned emotional responses are not limited to fear alone.

Nevertheless, fear is an extremely important example of conditioned emotional responses. When fears become strong or generalized enough to negatively affect one's life, they are labeled phobias. Phobias often interfere with a person's everyday life and can be very maladaptive. True phobias should not be treated lightly and they often call for treatment. Some people will state that they have a phobia, when in fact, they only have a simple aversion or fear. For example, someone who states that they have a phobia of spiders, but simply avoids or kills them probably has a strong dislike or aversion, but not a phobia. Someone who runs out of their house and then begins to cry uncontrollably because they saw a small spider on the floor probably has a phobia, in this case, arachnophobia.

Like simple fears, phobias are conditioned and generalized through experience with pairings of stimuli in one's environment and are included as a conditioned emotional response. Because phobias arise from conditioning, therapy for phobias usually also involves classical conditioning procedures, such as systematic desensitization. These procedures may extinguish phobic responses or even classically condition new and more positive responses to the stimuli. Thus, new emotional responses that are incompatible with and replace the fear can also be classically conditioned (Wolpe, 1958).

Systematic Desensitization

Classical conditioning as it occurs naturally in our everyday lives can, quite by chance, result in strong conditioned emotional responses ranging from simple fears to even more extreme and generalized fears called phobias. Early approaches to treatment for such phobias thus emphasized the importance of facing the originally feared stimulus in the absence of any harmful or feared unconditional stimulus, thus creating a forceful and sometimes highly uncomfortable extinction process called "flooding". In addition to being called flooding, such radical procedures were sometimes referred to as the bronco-busting technique in recognition that this is exactly how early cowboys eliminated a horse's fear of having a rider on its back. In real treatment circumstances, such as the clinic, this approach more often results in the patient failing to return for required successive extinction treatments than in successful treatments!

So one of the more popular conditioning-based treatments for phobias added a more gradual dimension to this extinction process. Such a gradual approach uses what is called systematic desensitization techniques. This is very close to the techniques used by what are sometimes called horse whisperers for dealing with a horse's fear of being ridden or handled using procedures that are quite different from bronco busting or flooding (Wolpe, 1958).

The application of systematic desensitization techniques, as well as its technical name, originated with Joseph Wolpe (1958). Systematic desensitization is a set of methods for eliminating fear through extinction, but it frequently adds another conditioning dimension for actually replacing fearful reactions to stimuli with more positive and adaptive responses. This replacement process is accomplished through the added use of a special form of classical conditioning called counter conditioning, which involves classically conditioning positive reactions that are incompatible with the more negative fear response to feared stimuli.

Systematic desensitization is often very successful in treating a wide variety of phobias. It has even been well received in the treatment of agoraphobia, fear of open and/or public places. This phobia is complex and is often linked to those who suffer general anxiety disorders. When counter conditioning is also used in systematic desensitization, an individual not only experiences extinction of the previously learned fear reaction but also simultaneously acquires a new response to a specific CS (in the case of arachnophobia, this CS is a spider). This new response to be learned is more adaptive than, and is literally incompatible with, the original fear response (Wolpe, 1958). Thus progressive relaxation is often used to compete with phobic reactions.

The process of systematic desensitization involves several successive steps. First, the client and the therapist break up the process of approaching the feared object into many small steps. These steps go from exposing the client at first to the least stressful stimulus, and only gradually moving to the most stressful stimulus that elicits the client's fear. For example, a client may break up handling a spider into several gradual steps that successively approximate the eventual handling behavior. The first step may be as simple as the client merely talking about spiders. After the client can comfortably use the word spider, the therapist may move to having the client imagine seeing a spider from a safe and relatively comfortable distance. Then the client may progress to being able to imagine approaching the spider. Gradually, still photographs, then moving pictures, and perhaps even rubber models of spiders might be used in progression. Throughout each step, patients are taught to evoke relaxation responses using specially taught relaxation techniques to produce parasympathetic responses that are incompatible with fear responses.

Phobic and anxiety reactions involve an activation of the sympathetic nervous system, which includes large secretions of adrenaline into the blood stream along with an elevated cardiovascular arousal where the heart beats more quickly and noticeably. The person may tremble and also break out in a sweat. This sympathetic nervous arousal process occurs simultaneously with an associated decrease in parasympathetic nervous activity. Parasympathetic activity is normally experienced with more positive feelings of mild sexual arousal, relaxation, and heightened digestive activity. Thus the goal of counter conditioning for phobias is not only to decrease sympathetic arousal elicited by feared stimuli, but also to condition an elevation in the parasympathetic nervous activity that is incompatible with this normally elevated sympathetic activity (Notterman, Schoenfeld, |_2 Bersh, 1952).

Gradually, through both the extinction of sympathetic arousal responses plus the pairing of the feared stimulus at each step with stimuli that evoke relaxation, the feared stimulus (i.e. spider) comes to elicit relaxation instead of fear behavior. Eventually, the client can reach the last step, touching or handling the spider. When the individual can complete this reliably, the phobia is fully treated.

While some clients may have handling a spider as their last step, others simply work to a goal of tolerating being in the same room with one and calmly getting someone else to kill it (or if they live alone, they learn to kill it themselves). Phobias are different for everyone and it is important to keep in mind that the point of systematic desensitization is not to get people to love the feared stimulus. It is to condition them to have adaptive behaviors in response to the feared stimulus even is a great dislike or a mild aversion still exists. Of course this is not to say that some who are extremely afraid of getting their face under water because they cant swim don't end up being avid swimmers after overcoming their phobic reactions! It all depends on the life style and desires of the client in overcoming their problems (Wolpe, 1958).

Conditioned Taste Aversion

Another application of Pavlov's classical conditioning procedures is that of conditioned taste aversion. Like conditioned emotional responses, taste aversion can be conditioned through everyday experience of chance pairings between neutral and eliciting stimuli in the natural environment. Organisms can come to avoid certain foods/tastes through classical conditioning principles, and almost every human has had such an experience. What favorite food did you get sick on once and now can no longer eat, no matter how hungry you are? If you don't have one you are somewhat unique in your good fortune!

John Garcia, one of the first to study the phenomenon of conditioned taste aversion, demonstrated this with rats in a laboratory. In Garcia's experiments, rats would freely drink water (a NS) in their own cages. During an experiment some of these rats would be transferred to experimental cages also containing water, but also where x-rays were present in these test cages. Exposure to these x-rays (UCS) produced nausea and sickness (UCR) in the rats. With time, the rats would not drink in these test cages (Garcia, Kimeldorf, Hunt, |_2 Davies, 1956).

Garcia et al. noticed that the water bottles in these cages were plastic while those in the home cages were glass. Plastic water bottles are known to give the water in them a distinct taste (a CS), while the water in glass bottles is virtually tasteless. Garcia concluded that the x-ray induced nausea was becoming classically conditioned to the taste of the water in these plastic bottles but not with the tasteless water of the home cage bottles (an inadvertent creation of differential conditioning procedures). Because nausea is an unpleasant response, the rats in the experiment came to avoid the water with plastic taste hence the term conditioned taste aversion. It is thought that conditioned taste aversion has important survival value, as the process prevents organisms from eating potentially dangerous foods that have similar tastes or odors to the ones that have made them ill in the past.

One use of Garcia's conditioned taste aversion findings is that of predator control. Predator control uses food aversion instead of an aversion to a liquid, however. Conditioning as an approach to predator control efforts began with an argument between sheep farmers and environmentalists in the Western United States. The debate was over what to do with coyotes that were appearing on the farmers' land and killing and eating their livestock, especially sheep. This problem cost the farmers tremendous amounts of money and they felt the only solution was to shoot and kill the coyotes when they attacked sheep. The environmentalists felt that this was inhumane and harmful to the environment as the coyotes the farmers were shooting were already endangered. Both sides had reasonable arguments and it was an application of classical conditioning procedures in the form of conditioned taste aversion that began to solve the problem (Gustavson, Garcia, Hawkins, |_2 Rusiniak, 1974).

The farmers were instructed to leave poisoned sheep carcasses around the perimeter of their pastures. The poison used is thiamene, a tasteless drug that causes extreme nausea and vomiting. When coyotes ate the poisoned meat they became very sick. Soon, they came to associate the sight and smell of sheep with the actual taste of the meat that had elicited their illness so they avoided eating sheep. This processes is a more humane way of controlling predators and is a solution, based on classical conditioning, that is considered much more humane than killing the coyotes (Gustavson et. al., 1974).

The effectiveness of food aversion on predator control has had mixed reports, however. Some farmers reported that coyotes didn't eat mutton anymore, but they still continued to kill the sheep. On the other hand, some farmers reported complete success with the process. Of course every new generation of coyotes has to be conditioned in a similar fashion, since learning and conditioning is only an individual adaptation that doesn't extend to new generations. While most farmers agree that the application of taste aversion has lessened the problem, it has yet to be seen as a complete solution. Some farmers are still forced to resort to shooting the coyotes when they attack their livestock (Timberlank |_2 Melcer, 1988).

Taste aversion therapy is an application of classical conditioning procedures that uses aversive, or unpleasant, stimuli to counteract undesirable and even maladaptive behaviors. Addictions are such behaviors in humans. So aversive conditioning has been explored as a treatment for addictions. This form of therapy usually involves conditioned taste aversion principles developed by Garcia, but it is not a necessity. In the treatment of the addiction to alcohol, or alcoholism, the unpleasant feeling of nausea is paired with the consumption of alcohol. What eventually results is an aversion, or avoidance, of alcohol. The process, typically referred to as antabuse treatment, is not perfect, as at times the individual suffering from alcoholism does not follow the procedures 100|_1 of the time (Forrest, 1985).

First, an individual with alcoholism is given a drug that they must take every day. Antabuse drugs have no effects unless the person drinks alcohol. If the person drinks, the antabuse drug reacts with the alcohol to create an extreme feeling of nausea (thus the antabuse drug plus alcohol is an UCS for the UCR of nausea). Following repeated pairings alcohol (now a CS) comes to elicit nausea (now a CR) without the antabuse and the person avoids the taste of alcohol by not drinking. This process has shown to be effective, but it is not a perfect solution. One of the problems with aversion therapy, even though it has proved to be successful in the treatment of addictions, involves its unpleasant nature and the inclination of some individuals to resist or avoid treatment. Likewise, a person addicted to alcohol may stop taking the antabuse drug so that they can drink with no consequences. It takes much effort on the part of the patient to overcome an addiction with aversion therapy (Forrest,1985).

Conditioned Immunity

Applications of Pavlov's classical conditioning principles include a very broad range of real world situations that give rise to personally significant physiological reactions. From conditioned emotional responses such as fear and phobias to taste aversions that almost every individual uniquely has acquired, conditioning can be seen to occur naturally in our everyday world. Likewise, contrived and manipulated stimulus pairings pervade our lives, such as when advertising presents beautiful people having fun or obtaining joyful relief (UCS) because of some product where the brand is a prominent (conditional stimulus). But even the human physiological immune response is subject to classical conditioning.

Thus Pavlov's conditioning procedures have been applied in medicine in order to improve immune functioning. Researchers have shown that after several pairings of a drug that increases the immune response, such as epinephrine, (an UCS) with a placebo of a certain taste or smell (neutral stimulus) the placebo (now a CS) will increase the immune response (now a CR) when presented alone. This is a useful phenomenon, as doctors can give only a placebo in between scheduled drug injections, thus increasing the body's immune response to infection with a minimal amount of drugs.

A placebo is any substance that does not initially have the effects of a particular drug. A sugar pill is often used as a placebo in drug studies. In the case of classically conditioned immune responses, anything with a distinctive taste, sometimes a certain flavor of ice cream, can serve as a placebo or, in other words, a neutral stimulus. When repeatedly paired with an UCS that elicits an immune response (usually epinephrine), the placebo (such as ice cream) can elicit the response when presented alone. This is very powerful, but it can also lead to problems if something were to be conditioned to suppress the immune system instead of boost it (Buske-Kirschbaum, Kirschbaum, Stierle, Jabaij, |_2 Hellhammer, 1994).

Operant Conditioning: An Overview

Classical, or Pavlovian, conditioning is a process by which new emotional and glandular reactions develop in response to previously neutral stimuli in the environment. But classical conditioning doesn't explain how we develop new skills or behaviors instrumental in changing our external environments. That learning process involves what is typically referred to as instrumental, or operant, conditioning. Operant conditioning describes how we develop behaviors that "operate upon the environment" to bring about behavioral consequences in that environment. Operant conditioning applies many techniques and procedures first investigated by E. L. Thorndike (1898) but later refined and extended by B. F. Skinner (Skinner, 1938).

Thorndike was an American psychologist who was one of the first to investigate the effects of behavioral consequences on learning. His work led him to emphasize both the effects of positive as well as negative behavioral consequences. Because behaviors are instrumental in bringing about such consequences by operating upon the environment in some way, this process for developing new skilled behaviors was first called instrumental conditioning. In subsequent literatures, especially in those inspired by the work of Skinner, the term "instrumental conditioning" was replaced by the term "operant conditioning." Nevertheless it was Thorndike who first concluded that positive consequences strengthen behaviors to make them more likely in similar situations in the future; a phenomenon he labeled the Law of Effect.

Inspired by the much earlier work of both Pavlov and Thorndike, another American Psychologist, B.F. Skinner, went on to develop the principles of operant conditioning. Skinner formalized these principles and identified many variables involved in this form of learning. For example, Skinner revised Thorndike's concept of "reward" by emphasizing that it has "positive reinforcement" effects which result in the increased likelihood of a behavior's future occurrences. Even painful consequences can increase the future likelihood of behaviors that eliminate or avoid such consequences, and thus Skinner emphasized their function as "negative reinforcements." According to Skinner, reinforcement, whether positive or negative, is the process of increasing future behavioral probabilities; meaning any response that is followed by a reinforcer will increase in its frequency of occurrence across time (a concept emphasizing the rate of specific ways of behaving). Skinner also discovered that such reinforcing events don't have to happen each and every time. Instead, intermittent reinforcement is also effective, and Skinner described the effects that different "schedules" of reinforcement (the timing or frequency of reinforcement) have on behavior. He also identified the process of punishment whereby behavioral probabilities may be decreased by consequences. Any behavior that is followed by punishment decreases in frequency.

Using the variables controlling operant conditioning as a foundation, B.F. Skinner also investigated several alternative operant procedures. For example, shaping (Peterson, n.d.) is a process of operantly conditioning a new form of behavior by reinforcing successive approximations to the ultimately desired form of behavior, known as the "target" behavior. Shaping is simply a method for conditioning an organism to perform a new behavior by reinforcing small and gradual steps toward the desired form of behavior. The experimenter starts the shaping process by reinforcing what the individual already does, then by gradually reinforcing only selective variations of that behavior that lead to the form of the target behavior, the experiment may gradually reach the point where only the target behavior is the one reinforced.

Behaviors come to occur only within certain antecedent environmental circumstances through a process called stimulus discrimination conditioning, and this process was also investigated in detail by Skinner. Both stimulus discrimination and generalization exist in classical conditioning processes as well, but we will currently focus only on these processes as they are employed in operant conditioning. Likewise, procedures exist for the extinction of operant behaviors as well as the parallel process of extinction of classically conditioned responses. In operant conditioning, extinction is the process wherein experimenters stop reinforcing or punishing a specific behavior, thus resulting in that behavior's return to pre-conditioning rates, or probabilities, of occurrence.

Operant conditioning techniques and procedures have many applications across various circumstances and problem areas. They have been utilized in the classroom environment with a great deal of success. Early forms of teaching machines first elaborated by Skinner have evolved into today's computer assisted instruction programs. Such programs allow students to receive feedback on their progress in mastering various types of subject matter while simultaneously shaping students to become more skilled in more generalized behaviors such as reading comprehension (Ray, 2004). Appropriate and learning-supportive classroom behaviors may also be developed and maintained with operant principles. In a process called the token economy, students are rewarded for good classroom behaviors or even independent study management programs using points or tokens that may be traded for more primary forms of reinforcement. Token economies are also used in psychiatric hospitals and other institutions to maintain and teach appropriate and adaptive behaviors (Ayllon |_2 Azrin, 1968).

Operant conditioning also has made significant contributions in therapeutic settings. For example, anxiety and other similar physiological disorders can be treated with a technique known as biofeedback. Many of the earliest demonstrations of biofeedback came from Neal Miller's laboratory (Miller, 1969), and typically involved the monitoring of a patient's vital body functions (heart and breathing rate, blood pressure, etc.) while also displaying their status via some type of feedback device, such as a computer display. Patients may be trained to relax or otherwise behave in ways that keep their vital physiological processes within a more healthy range of functioning. Feedback telling individuals that they are being successful reinforces their efforts to control their own physiological functions. Patients may, for example, wear portable monitoring and feedback devices to learn to relax in usually anxiety provoking situations.

Procedures like shaping and chaining are also applied in the training of performance animals, in obedience classes for family pet, and in the training of animals as personal assistants for blind or paralyzed humans (Pryor, 1985). Performance animals like these, as well as the ones seen in marine parks and circuses, learn complex chains of behaviors through operant conditioning procedures involving reinforcement and antecedent stimulus discrimination. For example dogs in obedience classes are taught to behave to simple commands that offer visual and auditory cues antecedent to desired behaviors as well as positive reinforcement as consequences for performing those behaviors in response to those cues.

Origins of Operant Conditioning: Instrumental Learning and the Law of Effect

Edward Lee Thorndike was an American psychologist studying animal learning while a graduate student at Harvard University in the late 1890's. He was especially interested in how animals learn to engage in new behaviors that are instrumental in solving problems, such as escaping from a confined space. The instrumental character of behavior in changing an animals' circumstances led subsequent authors to refer to Thorndike's form of learning as instrumental learning, although Thorndike preferred to describe it as "trial and success" learning (Chance, 1999). Pretty much these same behavioral processes were renamed operant conditioning by a much later researcher, B.F. Skinner (Skinner, 1938), who was also interested in how such skills "operate upon environments" (hence his more descriptive term "operant") to bring about significant consequences for the individual.

Thorndike designed many ingenious experiments into study such behavior. In one series of investigations Thorndike placed hungry cats into an apparatus called a puzzle box, from which the animals learned to escape to obtain rewards of food. At first Thorndike's cats seemed to behave almost randomly, using trial and error to find their way out of the puzzle box. Thorndike graphed the time it took an animal to escape from the puzzle box for each successive trial he gave the animal. He quickly discovered that the time for escape gradually declined over several repeated trials, with each successive trial typically taking less and less time. He called this a learning curve and proposed that the slope of this curve reflected the rate at which learning occurred (Chance, 1999). From such studies Thorndike proposed his Law of Effect, which states that if successful behaviors in a trial and error situation are followed by pleasurable consequences, those behaviors become strengthened, or "stamped in" and will thus be more quickly performed in future trials (Thorndike, 1898).

As noted above, in order to study the problem-solving behavior of cats using trial and error procedures Thorndike developed a special puzzle box apparatus. Various forms of puzzle boxes were constructed, but a typical one was a wooden cage equipped with a door held by a weighted loop of string holding, and a pedal, and a bar. A cat had to press the pedal, pull the string, then push the bar to unlatch the door to the box. This allowed the animal to then escape from the box and obtain food as a consequence.

The term instrumental conditioning is used to describe Thorndike's procedures for animal learning because the term ties behaviors to the generation of their consequences in learning-that is, the behavior is instrumental in obtaining important consequential outcomes in the environment. Thorndike's procedures involved what many refer to as "trial and error" procedures. For example, when Thorndike placed a hungry cat into his puzzle box, the cat would produce many behaviors in its attempts to escape the confinement. Eventually, the animal would produce the correct behavior quite by chance, usually clawing a string and then stepping on a pedal to open the door. This correct behavior had consequences because Thorndike would leave a plate of food just outside the box that the cat would eat from once it escaped. Thorndike's Law of Effect proposed that such rewards strengthened the behaviors that obtained the reward, thus making that behavior more quickly performed with fewer errors on future trials.

Thorndike's Law of Effect took two forms: the "strong" form and the "weak" form. Food as consequences represented the strong, or behavioral strengthening, form. The "weak" side of the Law of Effect describes what happens when a behavior fails to accomplish such pleasurable consequences, thus leading to a weakened, or "stamped out" impulse to behave in a similar fashion in similar situations in the future. Thorndike's studies were among the first to demonstrate and precisely measure the power of consequences in the environment (especially rewards) and their ability to control behavior, and thus Thorndike's work laid the foundation for the subsequent development of a more behavioral perspective on the learning process.

Operant Conditioning Principles

Another American Psychologist working at Harvard, B. F. Skinner, also studied the behavior of animals with a focus on consequences. Although Skinner's work came much later than that of Thorndike (Skinner began his work on operant conditioning in the 1930s), his research was based on the principles Thorndike had identified. Skinner (1938) believed that in order to understand psychology you had to focus only on observable behaviors.

Because observable behaviors and the role environments play in developing and controlling those behaviors are the focus of operant conditioning, Skinner and the field of operant conditioning is often considered to represent the most radical form of the perspective on learning called "behaviorism." Thorndike's work anticipates this movement as well, but Thorndike predated the philosophical emphasis on observable behaviors as the exclusive outcomes in learning. Throughout our discussion of operant conditioning, you will read terms such as "behavioral," "behaviorism," and "behaviorist." These terms typically refer to the work of Pavlov and Thorndike as the foundations of the perspective, but it was John Watson (1913) who described the perspective in most detail, and Skinner (1938) who most completely illustrated the power of the approach in what he called "radical behaviorism". These researchers emphasized the importance of observable behavior in understanding psychology and generally excluded mental activities in their studies. Because of this focus on behavior, their work is deemed "behavioral" and their conceptualization of learning is labeled "behaviorism." Keep in mind that this term does not include the cognitive or ecological perspectives.

Through his research, Skinner's radical behaviorism (1938) identified variables and formalized procedures using those variables in a conceptualization to learning called "operant conditioning." This term comes from Skinner's emphasis on the fact that behaviors operate (thus being an "operant") on the environment in order to gain certain consequential stimuli and to avoid others. Unlike classical conditioning, which Skinner called Respondent Conditioning because it focuses on the processes of learning in reflexive responses, operant conditioning focuses on how organisms learn totally new behaviors through experience with consequences in the environment. Skinner's operant conditioning is founded on Thorndikes' instrumental conditioning, but Operant Conditioning involves a wider variety of processes and labels consequences quite differently.

Skinner used rats as subjects for much of his work, but he is even more famous for his later work with pigeons. Dissatisfied with mazes or Thorndike's puzzle box, Skinner designed an apparatus to study animal behavior in a slightly different fashion. The operant chamber, or Skinner box as it came to be known, was designed to prevent human interruption of the experimental session and to allow the study of behavior as a continuous process, rather than in separated trial-by-trial procedures.

In Thorndike's puzzle box, the animal would have to be physically placed back into the box after each rewarded escape trial. Skinner felt that such procedures interfere with behavior as a "stream of events". For rats an operant chamber has a lever (technically called an operandum) that can be pressed over and over to deliver food pellets, with each press counting as a single occurrence of the behavior. For pigeons, one or more disks can be pecked as the operanda to deliver reinforcement for this behavior, usually in the form of food grain. The disks are often lighted for stimulus discrimination and generalization training. After an animal receives a reinforcement for pressing a bar or pecking a disk, there is no need to reset the system; the chamber is ready to deliver more reinforcements as soon as the animal responds again.

The cumulative recorder was another innovation introduced by Skinner to automatically graph response rates (that is, it shows an accumulation of the number of total responses as this total is distributed across time). in its original form, this machine recorded the number and timing of an operant behavior by using a continuously rolling piece of paper with a fixed ink-pen to mark time across a continuous X axis, as well as another pen that advanced one step up the Y axis each time a bar was pressed or key was pecked. Skinner was able to study animal behavior for as long as he deemed necessary without ever having to interfere with or even observe his animal.

Almost all of what Skinner (1938) discovered about operant conditioning principles came from his use of the operant conditioning chamber and its cumulative recorder-produced data. One procedure and its associated variables that Skinner identified was that of reinforcement. According to Skinner reinforcement involves the presentation or removal of stimulus consequences that increase the future rate of any specific class of operant behaviors, such as bar pressing or key pecking. The consequential stimulus variable is considered to be a reinforcer only if it's presentation or removal as a consequence for a behavior increases the future rate or probability of that form of behavior.

Skinner felt that when the presentation of a stimulus results in an increase in behavioral probability, positive reinforcement has occurred. Skinner also identified two types of positive reinforcers; primary (usually biological) and secondary or conditioned (must be classically conditioned to acquire reinforcing functions like the primary stimulus has). When the removal of a stimulus as a consequence for a behavior increases the likelihood of that form of behavior, negative reinforcement has occurred. Escape and avoidance learning are how we often describe changes in behavior rates that increase because of negative reinforcement.

Skinner also studied the procedure of punishment. Punishment is the opposite of reinforcement. It occurs when the probability of a behavior decreases with the presentation or removal of a stimulus. If presentation of an aversive stimulus decreases the likelihood of behavior occurring again, positive punishment has taken place. If the removal of a positively valued stimulus decreases the chances that a behavior will occur again, negative punishment, also called time out, has occurred. Skinner noted that punishment is often an inefficient way of controlling behavior, and in order to work at all it must be applied immediately after the behavior, it must be consistent and follow after every instance of the behavior, and it must be fairly strong.

So Skinner (1938) developed his procedures for operant conditioning through the manipulation of the operant variables of reinforcing and punishing consequences. But Skinner noticed that when he presented a reinforcement every time a behavior occurred, the rat or pigeon would become satiated quickly and would stop producing a certain behavior in high rates. Skinner labeled this procedure continuous reinforcement. Eventually he tried reinforcing behaviors using a non-continuous procedure -- a process he called partial, or intermittent, reinforcement. There are several types of partial reinforcement, each with different rules for applying one or another of type of consequences (reinforcements or punishments). These rules for "scheduling" reinforcement intermittently either rely on counting behaviors, such as fixed or variable ratio rules, or adding a time interval to the behavioral rule, such as fixed or variable interval schedules. Each type of schedule rule effects behavior in different and unique ways.

Skinner was also one of the first to seriously consider a fundamental flaw in Thorndike's trial and error learning procedure. Instead of using Thorndike's vocabulary which described an animal as random "trying" to solve a problem, Skinner preferred to talk about different activities as alternative forms of emitted behaviors. And all responses or behaviors that look alike or act upon the environment in a similar fashion form a "class" of related emitted responses, or an "operant class" of behavior. This contrasts with Pavlov's elicited behaviors, such as salivation, where known unconditional stimuli are used to bring about the "respondent" behaviors, as Skinner referred to them. But what if the animal never emits the correct behavior in a trial and error situation? Having made only errors, nothing could be reinforced and thus no learning (relatively permanent change in behavior) would take place either.

Skinner believed that by manipulating consequences in a certain systematic way, an organism could be led to the correct behavior much faster than if one simply waited for the animal to happen upon the response by chance. The procedure he developed for accomplishing this step-by-step process is called shaping, and it's purpose it to reinforce behavior in gradual steps that begin with only rough approximations to the eventual "target" that one has as the goal of learning. In Skinner's research a behavior he often shaped was a lever press by a rat in an operant chamber (Skinner, 1951; Peterson, n.d.). He would first reinforce the animal for being in the vicinity of the bar, then for sniffing the bar, then touching the bar with a paw, then standing over the bar and eventually pressing the bar, all in successive approximations or gradual steps to the final bar pressing he wanted the animal to learn. Chaining is yet another procedure that is based on shaping, but it is used to condition a whole complex series of different responses, not just one.

Extinction, stimulus discrimination and stimulus generalization also exist in operant conditioning just as in classical conditioning. Extinction occurs when reinforcement or punishment no longer occurs as a consequence for a given behavior. Spontaneous recovery can also occur in operant conditioning if extinction is tested again later, and rapid reacquisition occurs if reinforcement or punishment is again the consequence for behavior. Stimulus discrimination involves presenting reinforcement or punishment only under certain antecedent stimulus conditions and not others until the organism only produces the behavior under the given antecedent settings. Generalization is the opposite: reinforcement or punishment is the consequence of behavior in many antecedent settings and the organism produces the behavior across these many different circumstances. The procedures developed by Skinner have been tested in many different applied settings and are very commonly used today.

Reinforcement in Operant Conditioning

Thorndike's studies of instrumental learning where cats learned to escape from puzzle boxes led to his conclusion that behaviors are controlled by their consequences, which was stated as his Law of Effect (Thorndike, 1898). In his studies of operant conditioning Skinner (1938) also stressed the importance of behavioral consequences, which he referred to as reinforcement and punishment. Reinforcement occurs when the probability of a certain behavior increases due to the presentation of a stimulus as a behavioral consequence (positive reinforcement) or the removal of a stimulus as a behavioral consequence (negative reinforcement).

It is important to keep in mind that reinforcement is a process and occurs only if behavioral probability increases. Thus a consequential stimulus is not a reinforcer if it's presentation (positive reinforcement) or removal (negative reinforcement) does not increase the likelihood that the behavior will occur again. We often assume that something will reinforce behavior, but until the behavior has shown an increase in probability, you cannot be sure. For instance, you may think candy would reinforce a child for studying, but if the child doesn't study more often when given candy upon doing so, candy is not a positive reinforcer.

There are two kinds of positively reinforcing stimuli (stimuli that are generally reinforcing when presented to an individual) known as primary reinforcers and secondary (or conditioned) reinforcers. Both types can be delivered following various rules for delivery, thus defining various schedules of reinforcement. Often some type of procedure, such as deprivation, is required to establish that a certain stimulus will function as an effective reinforcer. According to Skinner (1938), reinforcement is much better at controlling behavior than punishment, which is defined by a decrease in the probability of any behavior that causes the punishing stimulus to be presented (negative punishment) or removed (positive punishment, or time-out).

Another way to positively reinforce behavior is to rely upon Premack's Principle (Premack, 1959, 1971). According to the Premack Principle, a normally higher frequency behavior can be used to positively reinforce a desired behavior that is normally lower in frequency. A parent is more likely to positively reinforce a child for studying (a low frequency behavior without intervention) by allowing the child to watch TV (a high frequency behavior without intervention) after studying for some specified time. In this case, allowing the consequential behavior of watching TV causes the probability of studying to increase. The Premack Principle has also been utilized in operant conditioning research on rats. Rats can be successfully reinforced for bar pressing (very low frequency behavior without intervention) by allowing the rat to run in a running wheel (normally high frequency behavior in rats).

Skinner (1938) also found that consequences resulting in the removing an aversive (painful, uncomfortable, or undesired) stimulus that was already present could also increase the probability that a certain behavior would occur. He called this process negative reinforcement. Crucial to negative reinforcement is: 1) the pre-existing presence of an aversive stimulus, 2) then a specific form of behavior that 3) has the consequence of terminating or removing that aversive stimulus. A parent who wants to reinforce the studying behavior of the child could use negative reinforcement by removing normally required chores for a week. It is important to remember that negative reinforcement is labeled "negative" because it relies upon the removal of an aversive stimulus, not because it is a "negative" way to reinforce behavior. And it is reinforcement because the behavior that removes the stimulus increases in probability.

Frequent use of negative reinforcement, inside or outside of the laboratory, will lead to what is often referred to as escape and/or avoidance behavior, as when you have an increased probability of taking an aspirin to escape a headache or to avoid developing muscle pain after strenuous exercise. Escape is the first of two phases of behavioral development involving the use of negative reinforcement. Avoidance is the second phase. If the floor of an operant conditioning chamber is electrified to deliver a mild electrical shock, a rat's bar press in the presence of this shock is negatively reinforced when the bar pressing turns off the shock. The rat will always experience the shock, but through negative reinforcement it learns to escape this aversive stimulus by pressing the lever that terminates the shock.

A child is negatively reinforced for whining about doing chores when someone reduces the time the child spends doing those chores. In this case whining becomes a means for escaping chores. But the child still has to come into contact with the aversive event (chores) before he/she can escape them by whining. As noted above, taking aspirin for a headache is a classic example of escape learning. The reduction or elimination of pain negatively reinforces taking the medicine. The headache must be experienced for this to occur, but the individual escapes the pain through pill taking behavior.

Avoidance is also a term that refers to increasing the likelihood of behaviors by the use of negative reinforcement. Avoidance typically appears as a second phase of development following the phase of escape. If a rat learns to press a lever by escaping a brief shock, eventually that rat begins to press even before the shock is delivered if pressing delays the next onset of shock (i.e., keeps the shock from occurring for a while). In this case, the rat may never again come into contact with shock, but bar pressing continues because it has been negatively reinforced. This is the essence of avoidance learning. A child who's whining is always reinforced by the removal of chores may learn to avoid doing chores altogether by whining even before starting chores. It would be far better to establish studying as a means by which the child can avoid chores!

Sometimes avoidance learning is facilitated by using some sort of antecedent stimulus signal for the impending shock. If, for example, a light in the chamber signals that a bar press by a rat may prevent the occurrence of an electric shock, the rat's bar press will be negatively reinforced by the termination of the light (escape behavior). Of course, at the same time the rat must also be avoiding any contact with shock because shock was prevented from coming on by the bar being pressed. After only a few experiences with actual shocks following such a light signal, the rat will learn to prevent shocks altogether by pressing the bar as soon as the light turns on. The bar press is avoidance behavior that is under the control of a discriminative antecedent stimulus (the warning light), and is thus called discriminative avoidance.

Reinforcement is a naturally occurring process, and doesn't have to be managed by someone. For example, can you think of any superstitions? Many people believe that walking under a ladder will give you bad luck or finding a four-leaf clover will bring you good luck. Well, in operant conditioning, superstitious behavior is a behavior that increases in probability because it happened to be reinforced merely by chance (Skinner, 1948). This happens especially when reinforcement occurs based on rules that are independent of a specific behavior, such as time since last reinforcement, rather than on what behavior was occurring. In pigeons, superstitious behavior may include shaking wings or other unusual behaviors before pecking a disk for reinforcement. The pigeon may have shaken its wings before pecking for food when it was first reinforced. That wing-shaking behavior is said to be superstitious because it has nothing to do with gaining reinforcement, yet it has increased in likelihood none-the-less.

In humans, blowing on dice before rolling them may be a form of superstitious behavior. A gambler may have once blown on a pair of die and then won the jackpot after he/she rolled the right numbers. The gambler may now believe that this blowing behavior led to the winning and will continue to do so on every roll. This behavior is superstitious because blowing on dice has nothing to do with the numbers you roll or the winnings you obtain. This can happen in a punishment situation as well. If blowing on the die resulted in a bad roll and the gambler lost everything, blowing on die will greatly decrease in frequency if it even occurs again. This decrease in behavior is superstitious because again, blowing on the die did not determine the result of the roll or the loss of money.

Conditioned Reinforcement and Operant Conditioning

Skinner (1938) described two types of reinforcing stimuli: primary and conditioned (or secondary) reinforcers. A primary reinforcer is anything that has the power to increase behavioral probabilities because it is involved with a biological need of the organism. Food, sex and temperature stabilities are often used as reinforcement because we need them as a species. Escape from pain and social acceptance/contact can also be considered as primary reinforcers due to their evolutionary importance to humans and certain other organisms. Primary reinforcers also provide a powerful source of motivation when an organism is deprived of them.

In operant conditioning, primary reinforcers are often used because of their immediate power to modify behavior. This power comes from the fact that they are of direct biological importance to the organism. Food, water, exercise, and escape from pain can be considered as primary reinforcers because a lack of these can be physiologically harmful and/or painful. Skinner used primary reinforcers, usually food, in most of his laboratory studies with pigeons and rats.

A conditioned or secondary reinforcer is anything that can increase the probability of behavior because of its reliable association with primary reinforcers. Classical (Pavlovian) conditioning is at work here, as can be seen in the case of money. Money can be a powerful reinforcer, although it has no real use to us unless it can buy the things that meet more primary needs, such as food, shelter, and entertaining stimulation. This is classical conditioning with money as the CS and food, etc. as the UCS. A small piece of metal (such as coins) or piece of paper (such as dollar bills) has little value per se, as illustrated by play money. However, food, social contact, relief from pain, and even relief from boredom (all primary reinforcers) can be obtained with a sufficient amount of money that is legal tender for things we need. Therefore, because money has been paired with these primary reinforcers so often, it takes on the power to increase the probability of behavior in and of itself. This is why it is labeled as a secondary, or conditioned, reinforcer.

Skinner's work illustrated that deprivation is a common procedure for effectively changing the nature of a reinforcer in operant conditioning (Skinner, 1938). Such procedures are described by some researchers as establishing motivation, and are thus referred to as "establishing operations" (Michael, 1982, 1993; Dougher |_2 Hackbert, 2000). A pigeon will not press a lever very frequently for food if it is satiated (full). Depriving the animal of this primary reinforcer (usually experimental deprivation involves food or water) will motivate the animal to perform, because now the reinforcer satisfies a biological need. In operant research, animals are not deprived of food or water to a point that is dangerous or very distressing (all research with animals must follow strict ethical guidelines in any discipline). The animals are usually made just hungry or thirsty enough so that food or water works as an effective primary reinforcer. But almost anyone who has eaten too much thanksgiving dinner can relate to the fact that food may eventually turn into an aversive stimulus when too much has been consumed!

Deprivation of a secondary reinforcer, like money, works much in the same way. A person who has |_817,000,000 is not going to be highly motivated to work for money. Someone with only |_83 to his or her name will do almost any kind of work for money if they have no other means for eating or staying warm. You may have noted, however, that the person with only |_83 is probably also deprived of primary reinforcers (like food, shelter, or social contact) as well as money. And the person with |_817,000,000 may work because he/she is deprived of certain social stimuli that money may not buy.

Punishment in Operant Conditioning

Thorndike's earliest studies of cats escaping from puzzle boxes led him to distinguish between two forms of his famous Law of Effect. Thorndike (1898) held that behaviors could be "stamped in" by satisfying consequences or "stamped out" by annoying consequences. This became the basis for his distinguishing between a Strong Law of Effect and his Weak Law of Effect. When behavior is stamped out by annoyers Thorndike felt that a "strong" Law of Effect was at work. He later withdrew this punishment element of his theory, eventually leaving only the "weak" Law of Effect that resulted in a "stamping in" of behavior. His work with human subjects learning verbal behaviors had convinced him that saying "wrong" had less effect than saying nothing, and the most effective response was saying "right" to the learner's responses. Thus Thorndike interpreted these results as arguing against the effectiveness of punishment (Catania, 1998).

Likewise, in his studies of operant conditioning Skinner (1938) described the phenomenon of punishment as well as reinforcement. Positive punishment involves a decrease in the probability of a behavior through presentation of (addition of, and thus the term "positive") an aversive stimulus as a behavioral consequence. Negative punishment describes the removal of a positive stimulus as a behavioral consequence. It is worth re-emphasizing that the stimulus that is presented in positive punishment is usually a painful or otherwise aversive stimulus, while those stimuli that act as negative (removed) punishers are usually sought-after or appetitive stimuli. Sound confusing? Then let's consider these distinctions in more detail.

As noted, punishment is an operant process of decreasing the probability that a particular behavior will occur. According to Skinner (1938) a stimulus cannot be considered a "punisher" if it's presentation (positive punishment) or removal (negative punishment) does not decrease the likelihood of a behavior. For instance, it may seem intuitive that giving extra chores will be a good punishment for a child having drawn on the wall. If, however, the child continues to draw on the wall with the same frequency despite the extra chores, the chores are not punishers and punishment has not occurred.

It is easy to confuse the use of positive and negative to describe types of punishment. As in the case of positive reinforcement, positive punishment refers to the presentation of a stimulus, only now it decreases behavioral probabilities instead of the increase probability that defines reinforcement. But aversive stimuli are used often in operant conditioning procedures. Anything that causes pain, discomfort, high levels of physical and/or mental stress or anything that is undesired is classified as an aversive stimulus. In successful negative reinforcement, their removal results in an increase in the probability of a certain behavior. In positive punishment, however, the presentation of an aversive stimulus results in the decrease of the probability of a certain behavior. But because the stimulus being presented in positive punishment is usually aversive to, in other words, unpleasant for, the organism, some people are inclined to speak of such aversive stimuli as being negative for attracting/repelling the individual. See the potential for confusion? Positive punishment adds negative stimulus consequences.

As with negative reinforcement, negative punishment involves the removal of a stimulus. In this case, the goal is to decrease the probability of a behavior, so the stimulus removed is a desired, pleasant, or "positive" stimulus. Punishing a teenager for missing a curfew by taking away use of the car for a period of time is an example of negative punishment. It is the time during which the stimulus is not available that negative punishment gets its other, more common, name of "time out." Remember, though, "negative" refers to the removal of some stimulus, it is not a value judgment of this type of punishment.

Punishment is generally not a very effective means of behavior control, but there are several punishment factors that will modify how effective it is for decreasing behavioral probabilities. As Skinner noticed, every behavior serves some purpose for the organism (i.e. some children misbehave for attention) and if you decrease the likelihood of a behavior, it will appear again unless you shape a new behavior that will have the same purpose for the organism. Punishment doesn't make behaviors disappear; it just reduces the likelihood that they will appear. Another issue with the use of punishment is what happens to the status of the punisher. A dog (or a child for that matter) may come to find the person who continually punishes it as itself aversive, and it will avoid the individual as it comes to associate him or her with punishment.

Sometimes, punishment is necessary. However, in order to be effective at all, the following factors in punishment must be present. Punishment must occur immediately after the behavior, it must be strong, but not overwhelming and it must consistently follow every instance of the behavior to be reduced. How many times did you hear "Just wait until your father comes home," after you were caught misbehaving as a child? While this is meant to scare you, the punishment is still a long way off if it even comes at all. This type of behavior control doesn't work, much to the frustration of many mothers. In order for punishment to decrease the occurrence of behavior, it must occur immediately after the inappropriate behavior. Too much delay makes any future punishment random and not tied close enough to the behavior that needs to be decreased.

Punishment must not only be immediate, but also must be strong in order to be effective. Telling a child "Stop that!" when he/she is caught hitting another child will not be enough to decrease the behavior. However, the child does not need to have his/her toys taken away for a week for the transgression either. The punishment must fit the crime. A "time out" of about 5 minutes and a lecture of why hitting is wrong is usually aversive enough to a young child to greatly decrease the behavior. As Skinner noted, punishment should not be strong enough to cause harm, but it should be strong enough to be aversive.

Another very important issue in the effectiveness of punishment is consistency. As Skinner noted in his research, punishing behavior only occasionally is not an efficient way to decrease the likelihood of behavior. A child who is punished for hitting needs to be punished every time he/she is caught doing it, otherwise the punishment does not work to reduce the occurrence of this behavior.

Operant Conditioning Procedures

B. F. Skinner's (1938) investigations of operant conditioning introduced a variety of unique experimental procedures as well as demonstrations that various processes observed in Pavlov's classical conditioning also have counterparts in operant conditioning. Skinner's operant conditioning procedures introduce alternative manipulations of operant conditioning variables, such as antecedent stimuli and reinforcement contingency rules. These various operant procedures include extinction, generalization, discrimination, shaping, chaining, and a variety of different schedules of reinforcement.

The processes of extinction, generalization and discrimination that were discussed in the classical conditioning section have counterparts in operant conditioning. In extinction, reinforcement that has been a reliable consequence of a behavior is no longer presented. That is, the behavior no longer generates reinforcing consequences. Skinner noticed that when a behavior is first put on extinction, the organism displays a burst of the behavior and then begins to produce new, but related, behaviors -- a phenomenon called response induction. But eventually the behavior decreases in frequency to the point that it is very rarely emitted. If an instance of the behavior is reinforced again, however, spontaneous recovery will occur.

The operant procedure of discrimination training requires a stimulus be presented before the behavior even occurs, leading to its description as an antecedent to behavior. This antecedent stimulus serves to "set the occasion" that any lever press occurring in the presence of this antecedent will be reinforced. Experimentally such a stimulus may be auditory (i.e. a tone) or visual (usually a light of a certain color) or any other modality.

Skinner illustrated discrimination by reinforcing a rat's lever presses in the presence of an antecedent discriminative stimulus (also called an S|_3 or Sd) and not in it's absence (a condition called S- or Sdelta) . Thus behavior in S- is on an extinction schedule in the absence of the discriminative stimulus. Eventually, rats only pressed the lever in the presence of the stimulus, hence completing the discrimination process. Stimulus discrimination is also used in the process of chaining, where one behavior signals that a different behavior will be subsequently reinforced.

Like stimulus discrimination, stimulus generalization in operant conditioning is only slightly different than its counterpart in classical conditioning. Let's say, for example, that a rat's lever press has been reinforced in the presence of a red light but not it the presence of a green light. The rat will come to press only in the presence of the red light, hence demonstrating stimulus discrimination. If a pink or orange light is shown and the rat presses the lever, stimulus generalization has been demonstrated. As an operant conditioning procedure developed by Skinner, stimulus generalization occurs when an organism performs a behavior under antecedent conditions similar to conditions under which it was reinforced.

Schedules of reinforcement involve procedures whereby not every occurrence of a given form of behavior is followed by a reinforcer. Skinner (cf. Ferster |_2 Skinner, 1957) noted that when every instance of a behavior is reinforced, the animal quickly becomes satiated (has enough of the reinforcer that the stimulus loses reinforcing power) and stops engaging in the behavior. To create more steady and long lasting rates of behavior, Skinner would only reinforce a behavior some of the time. This is called a partial, or intermittent, reinforcement schedule (rather than a continuous reinforcement schedule, or CRF) and there are four major types of partial reinforcement procedural rules: fixed ratio (FR), fixed interval (FI), variable ratio (VR) and variable interval (VI). Each procedure calls for presentation of the reinforcement based on either the number of behaviors produced (ratio) or the timing between behaviors (interval). These schedules of reinforcement each have different effects on behavior and we will see (after discussion of other operant procedures) examples of these schedules in everyday situations.

Skinner eventually became dissatisfied with Thorndike's trial and error learning procedures. Skinner felt that by appropriate manipulation of behavioral consequences an experimenter could lead an individual to a correct or desired behavior much more quickly than it would be discovered by chance occurrences. He was thus interested in finding a much more efficient form of learning than trial and error. Skinner described his alternative process as one of shaping a desired, or target, response through reinforcement of successive approximations to the target behavior (Peterson, n.d.).

In shaping, reinforcement is presented for varying successive approximations in forms of behavior as they approximate the eventual behavior to be learned. Step-by-step, the organism comes to engage in behaviors that more and more closely approximate the target behavior. Eventually only the target behavior is the one reinforced. Shaping usually takes much less time than trial and error learning, where an experimenter must wait for the organism to produce the target behavior and subsequently reward it. Related to shaping is a process called chaining. Chaining is used to condition an individual to produce a specific series, or sequence, of different behaviors before the final behavior is reinforced. The chaining process uses discriminative stimuli presented after each step to "link" the chain of behaviors together.

Extinction in Operant Conditioning

Extinction is as much as an operant conditioning procedure as it is aclassical conditioning one. Extinction is sometimes considered a schedule of reinforcement as it is the process of withholding reinforcement for a previously reinforced behavior. Skinner (1938) noticed that this procedure brings about interesting results in and out of the laboratory. When a rat that has been reinforced for lever pressing is put on extinction, two things will occur: bursts of lever pressing and the appearance of new behaviors. The rat will show and increase in response rate immediately after extinction has begun. The rat will then emit new behaviors that may have been infrequent or not recorded. Each of these are dimensions of what is called response induction. As we have seen, the new behaviors that often follow extinction are key to the shaping procedure.

If a lever press that has been put on extinction is reinforced again, it usually only takes one or two reinforcements before lever pressing returns to it's pre-extinction frequency. This occurs even if extinction lasts days or weeks. This phenomenon (the rapid return of lever pressing) is called spontaneous recovery. As in the case of classical conditioning, the existence of spontaneous recovery suggests that, after extinction, behavior is not extinguished, it is somehow suppressed. The lever pressing did not need to be re-shaped; it emerged quickly after extinction. A human example of extinction can be demonstrated when a soda machine does not give a soft drink even after a person has deposited money into it. Usually, you get response burst, (person pushes the button many times and may deposit more money) and the emergence of new behaviors (kicking, swearing, calling the vendor, etc.)

It is important to note that following the extinction of a reinforced behavior an organism will often display an early increase in the rate of that behavior and then the emergence of new behavior. Skinner called this process of increased response rate and variation "response induction" and it is one effect of extinction. Behavior does not instantly stop as soon as extinction is implemented.

As noted, new behaviors often follow the extinction of a reinforced behavior. Skinner capitalized on this phenomenon when he was developing the operant conditioning procedure of shaping (reinforcing successive approximations and then putting them on extinction in order to draw out new behaviors that would more closely approximate a lever press). This phenomenon may also have some survival value, because if new behaviors were not emitted when reinforcement (especially in the form of food or water) no longer follows a particular behavior, an organism would perish if it simply continued producing the same response over and over again.

After the operant procedure of extinction has been implemented for a previously reinforced behavior and the rate of the behavior jumps initially (bursts) due to response induction, response rates then gradually decline to very low rates. If, however, (even after days of extinction) the behavior is reinforced, the response rate jumps back to near pre-extinction rates. This may happen in only one or two reinforcements. This phenomenon is called spontaneous recovery.

Operant Response Shaping and Chaining

Response shaping is an operant procedure developed by B. F. Skinner to bring about new behaviors in an organism (Peterson, n.d.). This procedure is often used in animal training and usually, but not always, involves positive reinforcement. Shaping procedures also include elements of extinction and is a process whereby the form or function of a behavior is developed into a targeted response. Training a rat to press a lever (target behavior) for food in an operant conditioning chamber is a common example of a shaping procedure. A rat generally does not press a bar very often, if at all, when it is first placed into an operant conditioning chamber (also known as a Skinner box). So how do we get it to do so?

Skinner used the ideas of operant conditioning to find an answer this question. Why not reinforce the rat's behaviors that approximate a bar press? Beginning with what the animal does relatively frequently, say looking at, going over to, and even just sniffing the bar (a behavior that occurs often when a rat is placed into an operant chamber), Skinner reinforced each of these to increase their probability. Then, as each became more likely, Skinner changed the rules of reinforcement to include only those behaviors that more closely resembled or actually were bar presses.

It is important to remember that following the extinction of a reinforced behavior an organism will typically increase the probability of that behavior and also engage in a wider variation of that form of behavior, often resulting in the emergence of new, but related, behaviors. Behavior does not instantly disappear as soon as extinction is implemented but rather reflects this typical "burst" in probability and variability as an early effect of extinction. The appearance of new related forms of behaviors is thus another early effect of extinction.

So after the rat consistently emitted one of the "approximate" behaviors, such as first looking at, or later approaching, and even later for sniffing the bar, it was reinforced (usually with food) for doing so. But soon Skinner would no longer reward the behaviors that least approximated actual bar presses, hence initiating extinction for that behavioral approximation. As soon as that behavior was no longer reinforced, the rat engaged in the behavior even more and emitted variations of the behavior. One variation of sniffing a bar, for example, might be rearing up and placing paws on the bar. When this occurred, Skinner began to reinforce this new behavior. When placing paws on the bar reached a fairly high probability, Skinner would then stop reinforcing paws on the bar and the rat would again begin to emit new variations of such behaviors, one of which typically involves actually scratching at and even pressing down on the bar. Skinner would reinforce this and the shaping procedure would be complete. A bar press behavior had been taught through reinforced successive behavioral approximations to a behavior that might begin with a zero probability of ever occurring.

The shaping process, because of its use of alternating use of reinforcement and extinction, is often called differential reinforcement of successive approximations in behavior. Successive approximations refer to the different behaviors that lead, step-by-step, to the target behavior (looking at the bar, approaching the bar, the bar sniff, paws on bar, and finally the bar press in this case). Differential reinforcement refers to the fact that we reinforce these approximations until the behaviors are produced reliably and then reinforcement is withheld so that new and different (hence the word differential) behaviors appear that better approximate the target response to be shaped.

The process of shaping also incorporates the creation and use of secondary reinforcers. If you were to shape a dog to "shake hands", you may not want to have to give it food (a primary reinforcer) every time it emits the correct behavior. By the time shaping is half-completed, the dog may be satiated, and food may not work as a reinforcer anymore. Different schedules of reinforcement may not be appropriate in this case, either. What many people do is say "Good, dog!" right before giving it a treat. Eventually, because of the pairing of the praise and food, the praise takes on reinforcing properties (it increases the probability of behavior). Through this classical conditioning procedure of pairing praise with food, you can reinforce the dog less with food and more with praise (now a conditioned reinforcer) and hence complete the shaping process.

In the case of operant chambers rather than dog training, the delivery of food is typically accomplished by a revolving magazine mechanism, much like those that deliver bubble gum one ball at a time from glass vending machines. The sound of this magazine shifting to deliver, in this case, a food pellet serves as a secondary reinforcer much like the praise example above. This allows for behaviors that take place at quite a distance from the actual food dispenser to be reinforced via secondary reinforcers. The establishment of such secondary, or conditioned, reinforcement functions is often referred to as magazine training and the process involves a conditional stimulus (CS is magazine sound) pairing with an unconditional stimulus (UCS is food) relation which is the same as Pavlov's metronome and food in classical conditioning.

Shaping is not limited to use on animals for simple training. Skinner demonstrated the technique had wide applications with his teaching machine, a device that shaped the skills of human students in correctly answering questions in many subjects. Skinner broke down the complex tasks of learning a new subject into small successive units that gradually built into much more complex systems of knowledge. This technique was called programmed instruction and was the basis for how the teaching machine worked. Skinner's teaching machines served as the prototypes for many modern computer-assisted instructional and training programs.

In order to shape very complex behaviors, as is often seen in animal performance shows, an operant conditioning procedure known as chaining must be implemented. In chaining, one behavior is "bridged" or linked to another by use of a discriminative stimulus that is always associated with the next behavior being reinforced. This process can be used to allow many behaviors to follow one another before reinforcement is actually delivered. In certain animals, the "chains" can be very long while in others they are short and reinforcement must be delivered more often. Eventually the discriminative stimuli that bridge each behavior to the next may be gradually "faded" to generalize the discrimination to the behavioral act itself, thus generating a sequence where one behavior sets the occasion for the next behavior, with the eventual end of the chain of different behaviors being the one reinforced.

If a dog trainer wants a dog to learn to "shake hands," then jump through a hoop and then stand on two feet, begging, that trainer will first shape the begging behavior in the presence of some hand signal, such as "hand raised in air." Once this is reliable, the trainer will present hoops (a second discriminative stimulus) and the dog will only be reinforced when it jumps through hoops and sees a hand raised to signal the begging. Finally, the dog will be shaped to shake, which will bring about the hoops, which signal that jumping and then begging will be reinforced. If the dog doesn't shake, the hoops will not appear, and no reinforcement will be given. Eventually, through this chaining procedure, the dog will shake, jump through the hoop and beg in smooth succession (reinforcement being given after the beg only).

Users of the CyberRat laboratory simulation may wish to read a step-by-step description of how best to shape an animal with no prior experimental history. There is such a collection of topics available in the Appendix. These include:

Shaping a New Behavior.
Getting Ready for Shaping.
Understanding the Experimental Chamber.
Getting Your Subject Ready for Shaping: Habituation
Getting Your Subject Ready for Shaping: Magazine Training
Getting Your Subject Ready for Shaping: Observe Behavior Carefully
Begin Shaping (If Operant Level is Low)
Shaping: Not Too Slow, Not Too Fast
Other Factors Involved in Creating New Behavior: Prompting
Other Factors Involved in Creating New Behavior: Discrimination
Factors Involved in Creating New Behavior: Intermittent Reinforcement

Schedules of Reinforcement

One group of procedures Skinner developed in his work on operant conditioning is that involving reinforcement schedules (Ferster |_2 Skinner, 1957). Schedules of reinforcement are simple rules for when reinforcement should be given following a specific behavior. The two main schedule rules are continuous and partial reinforcement. Another word for partial reinforcement is intermittent (less than continuous) reinforcement. The most common intermittent reinforcement rules include four specific types of schedules: fixed ratio , variable ratio, fixed interval and variable interval. Skinner observed that these different schedules have different effects on rates of responding, each of which will be illustrated by the graphics that accompany our more detailed descriptions of each schedule in this or subsequent topical discussions.

In a continuous reinforcement schedule every occurrence of a behavior is reinforced. If a rat is on a continuous schedule of reinforcement (often abbreviated as CRF) for lever pressing, every lever press is reinforced. A child, for example, who gets some dessert every time he or she finishes dinner is on a continuous schedule of reinforcement. As Skinner noted, this schedule produces a relatively moderate and steady rate of responding until the organism becomes satiated (an animal gets so much food as reinforcement that it is no longer hungry or the child has received desserts so often, he/she is tired of them.) This can occur relatively quickly, depending on the size of the

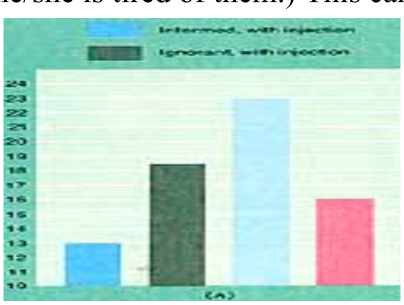

reinforcer, and thus is not an efficient means for maintaining a steady rate of responding over sustained periods of time. The accompanying figure is a simulated graphic illustrating both the relatively steady rate of responding and the slowing or elimination effects of satiation to the reinforcer within a single session.

Both to avoid having to use so much food and to counteract the satiation effects of continuous reinforcement, Skinner used intermittent schedules of reinforcement (Ferster |_2 Skinner, 1957). In intermittent schedules of reinforcement, only certain occurrences of a class of behaviors are reinforced. Sometimes the rule defining which behavioral occurrence should be reinforced is based on time elapsed plus the required response. Thus in what are called the interval schedules a predetermined amount of time must go by before reinforcement is delivered for the first response occurring after the interval of non-reinforcement for responding. Such interval schedules exist as either fixed or variable interval schedules. That is, the amount of time that reinforcement is not delivered for any behaviors is either the same interval following actual reinforcement, or time intervals are randomized durations around some average interval length.

Alternatively, delivery of reinforcement may be based on the number of times a specific class of behavior occurs (called ratio schedules because a particular type of response must occur a certain number of times before reinforcement is given). Such rules include fixed ratio schedules, where the required number of responses stay the same from one reinforced behavior to another, and variable ratio schedules where the number required between reinforcement delivery is some random number around a specific "average" of responses, such a an average of 10-to-1 or 20-to-1 (that is, on average one of 10 or 20 responses will be reinforced, but will randomly vary from 1 to any number, so long as in the long-term, the average of 10 or 20 is maintained).

In laboratory studies using either rats or pigeons, Skinner (Skinner |_2 Ferster, 1957) found that the rates of behavior are different for the various partial schedules of reinforcement and that the schedule chosen is often a function of what type of responding a researcher, or employer for that matter, might desire. Both the interval schedules of reinforcement and the ratio schedules of reinforcement and how they effect the rate of responding effects each type of schedule are covered in more detail in those respective topical discussions.

While conducting research on schedules of reinforcement as variations in operant conditioning procedures, Skinner noticed an interesting phenomenon surrounding the use of partial reinforcement. When a pigeon on continuous reinforcement is subsequently put on extinction (no reinforcement is delivered), the animal emits a burst of responses at first, but then gradually stops responding. In contrast, a pigeon that has been gradually moved to a partial schedule of reinforcement (especially if it is "lean" meaning reinforced rarely in the face of producing lots of responses) will continue responding for a very long time when moved to extinction; often taking multiple sessions before slowing down after extinction is started. This resistance to extinction follows any type of partial reinforcement schedule as long as the schedule is brought on gradually and is a relatively lean schedule, This resistance to extinction phenomenon is thus one of the primary partial reinforcement effects.

Ratio Schedules of Reinforcement

Skinner's research on operant conditioning procedures eventually led him to investigate intermittent, as opposed to continuous, reinforcement schedules (Ferster |_2 Skinner, 1957). Intermittent schedules of reinforcement are simple rules for delivering single reinforcements for multiple occurrences of a specific type of behavior, such as lever pressing. Skinner's original investigations used continuous reinforcement, where each and every lever press was reinforced. But in subsequent research he began to investigate what would happen if not every lever press was reinforced, a practice known as applying rules of intermittent (less than continuous) reinforcement.

One of the simplest, and thus most common intermittent reinforcement rules involves using some ratio of some number of required lever presses for each delivery of one reinforcement. Such ratios may use a "fixed" number, such as FR-10 where every 10th response would be reinforced, or a variable number, such as VR-10, where any constantly varying and random number of responses is used as the criterion for delivering reinforcement for the criterion lever press, so long as a large sample of these ratios average the reference ratio number (in our case, 10).

A special case of the ratio schedule, known as CRF or FR-1, is actually a continuous reinforcement schedule where every occurrence of a behavior is reinforced. If a rat is on a continuous schedule of reinforcement for lever pressing, every lever press is reinforced. A child, for example, who gets some dessert every time he or she finishes dinner is on a continuous schedule of reinforcement. Skinner's earliest work (Skinner, 1938) investigated this schedule almost exclusively, and he observed this schedule to produce a relatively moderate and steady rate of responding until the organism becomes satiated (an animal gets so much food as reinforcement that it is no longer hungry or the child has received desserts so often he/she is tired of them.) This can occur relatively quickly, depending on the size of the reinforcer, and thus is not an efficient means for maintaining a steady rate of responding over sustained periods of time. The accompanying figure is a simulated graphic illustrating both the relatively steady rate of responding and the slowing or elimination effects of satiation to the reinforcer within a single session.

To avoid using so many food pellets, which in his early research he had to hand-manufacture, Skinner eventually investigated intermittent schedules of reinforcement (Ferster |_2 Skinner, 1957). In intermittent schedules of reinforcement, only certain occurrences of a class of behaviors are reinforced. Sometimes the rule defining which behavioral occurrence should be reinforced is based on some interval of time elapsed plus the required response, thus generating what is known as the interval schedules of reinforcement.

Alternatively, delivery of reinforcement may be based on the number of times a specific class of behavior occurs (a particular type of response, such as a lever press or a key peck, must occur a certain number of times before reinforcement is given). Such rules include fixed ratio schedules, where the required number of responses stay the same from one reinforced behavior to another, and variable ratio schedules where the number required between reinforcement delivery is some random number around a specific average number of responses, such a an average of 10-to-1 or 20-to-1 (that is, on average one of 10 or 20 responses will be reinforced, but will randomly vary from 1 to any number, so long as in the long-term, the average of 10 or 20 is maintained).

In laboratory studies using either rats or pigeons, Skinner (Skinner |_2 Ferster, 1957) found that the rates of behavior are different for the various partial schedules of reinforcement and that the schedule chosen is often a function of what type of responding a researcher, or employer for that matter, might desire. A rat which gets reinforced every twentieth lever press is operating under a fixed ratio schedule of reinforcement (actually, an FR-20). As Skinner's research illustrated, this version of partial reinforcement produces very steady rates of responding, but only after a brief break after the reinforcement is delivered-a pattern often referred to as break run (see the accompanying illustration). Factory workers who get a certain amount of money for, say, every 5th completed unit of product are working under a fixed ration schedule of reinforcement. In the workplace fixed ratio schedules are known as piecework schedules wherein pay is based on a fixed number of components produced (an example of measuring behavior by its effects on environment, which is the defining feature of operant behavior -- it "operates upon environments" to produce some effect -- in this case, a part of some sort). In his laboratory, Skinner discovered that this schedule produced a fairly predictable brief break followed by a steady rate of subsequent responding. Employers often use piecework schedules because they usually result in relatively high productivity.

The variable ratio schedule is also a partial schedule of reinforcement. A variable ratio schedule of reinforcement, like a fixed ratio, involves the delivery of reinforcement based on the number of behavior occurrences. In a variable schedule, however, it is an average number, not a fixed number, of responses that are reinforced. A rat on this schedule may get reinforced, on average, for every ten responses. But because it is an average, reinforcement may come after two responses or after twenty. Reinforcement is not delivered every ten responses, although there may be a time when a tenth response is reinforced. Skinner noted that this is a very powerful schedule and it produces very high and quite constant rates of responding.

Gambling on slot machines is a clear example of a variable ratio schedule of reinforcement. For instance, on average, every 50th hit on the machine will being about a jackpot (reinforcement). This means that the jackpot could occur in two hits or two hundred hits. Gamblers find this possibility (two hits until jackpot) irresistible and many develop gambling problems because this type of reinforcement schedule produces high and consistent rates of responding. Many individuals continue gambling despite increasing debt because the indeterminate predictability of reinforcement and also because of the unique resistance to extinction caused by intermittent reinforcement.

As with interval schedules of reinforcement, a pigeon or a person that has been gradually moved to a ratio schedule of intermittent reinforcement (especially if it is a lean schedule, meaning reinforced on average or on every, say, 50th response rather than, say, every 5th response) will continue responding for a very long time when moved to extinction; often taking multiple sessions before slowing down after extinction is started. This resistance to extinction follows any type of partial reinforcement schedule as long as the schedule is brought on gradually and is relatively lean. This phenomenon is thus called the resistance to extinction effect of intermittent schedules of reinforcement.

Interval Schedules of Reinforcement

One group of procedures Skinner developed in his work on operant conditioning is that involving reinforcement schedules (Ferster |_2 Skinner, 1957), or simple rules for which occurrences of a behavior will be reinforced. If reinforcement is not continuous --

delivered for every occurrence of a type of behavior -- then a partial, or intermittent, reinforcement schedule is in effect. Common rules for scheduling intermittent reinforcement include fixed ratio, variable ratio, fixed interval and variable interval. In this section we will consider the unique response rate patterns generated by either fixed or variable interval schedules of reinforcement.

In interval schedules of reinforcement, the amount of time that reinforcement is not delivered for any behaviors is either the same duration following a reinforcement (fixed interval schedules), or time intervals are randomized durations around some average interval length (variable interval schedules, which are sometimes also called random interval schedules). If, as a researcher, you were to reinforce a rat for the first lever press that occurred after one minute following a previous reinforcement, you would be using a fixed interval schedule of reinforcement called an FI-60 (for the 60 seconds that elapses after the last reinforcement before the next response will generate another reinforcer). As a partial schedule, only the lever press that follows a one-minute interval is reinforced, and then a new one-minute interval is reset. Any lever press occurring during this one minute interval fails to bring reinforcement.

In laboratory studies using either rats or pigeons, Skinner (Skinner |_2 Ferster, 1957) found that the rates of behavior generated by interval schedules of reinforcement are uniquely but predictably variable across time, thereby creating a predictable "pattern" in the rate of responding for fixed vs variable interval schedules. A fixed interval schedule, as can be seen in the accompanying graphic illustration, produces one of the more unique response patterns. This pattern is typically described as a scalloping change in the rate of responding. What happens is that immediately after being reinforced, the rat stops responding. But as time passes, responding begins, at first slowly, and then the rate increases until it is very high near the end of the minute. This means that when the interval expires, the animal is very likely to respond, and the first such response is reinforced. As you can see, because of the scalloping effect, this schedule is not efficient at producing steady rates of responding.

Many teachers give weekly exams to test what students have learned in a given course. The opportunity to be reinforced for studying, by receiving an "A", is on a fixed (because it is every week) interval schedule of reinforcement. As Skinner noted in the laboratory (and as many parents and teachers know) the rate of responding under this schedule is scalloped. This means that right after a quiz, responding (studying) drops to nothing. Responding then slowly increases three nights away from the next quiz, until students are "cramming" the night before. This schedule is ineffective at maintaining a steady rate of responding (studying) both in and out of the laboratory.

The variable interval schedule uses rules similar to the fixed interval schedules, except the duration of the interval constantly changes in a random fashion around some average. Thus a VI-60 schedule would result in an average non-reinforcement availability period of 60 seconds, but each specific period could be any duration so long as a large sample of these intervals results in an average of 60 seconds. So under a variable, or random, interval schedule, a rat would be reinforced, on average, for the first response occurring after one minute from the last response. This means that the rat may receive reinforcement after, say, twenty seconds between behaviors or after four minutes between behaviors.

Has a teacher/professor ever given pop quizzes in a course you have taken? If he/she did, they would be using a variable interval schedule of reinforcement. What was the effect on the way members of the class prepared for the exam? As demonstrated by the accompanying graphic illustration, Skinner observed that the rate of responding under a variable interval schedule, although not as high in ratio schedules, is very steady and consistent.

So your instructor who gives pop quizzes is using a variable interval schedule of reinforcement in an effort to maintain high and consistent rates of studying. Perhaps the quizzes are given weekly, on average, but a subsequent quiz could come, one or two days or even two weeks after the last quiz. Since a student has no idea when the next quiz will come, studying is much more consistent than with weekly (fixed interval) exams.

While conducting research on schedules of reinforcement in operant conditioning, Skinner noticed an interesting phenomenon accompanying the use of partial reinforcement. When a pigeon on continuous reinforcement is subsequently put on extinction (reinforcement is no longer delivered for responding a specific way), the animal emits a burst of responses at first (a phenomenon known as response induction), but then gradually stops responding. In contrast, a pigeon that has been gradually moved to a partial schedule of reinforcement (especially if it is "lean," meaning, for example, the animal is reinforced for the first response following on average an interval of 180 seconds) will continue responding for a very long time when moved to extinction. In fact, it often takes multiple sessions before responding begins slowing down after extinction has started. This is known as resistance to extinction and it follows any type of partial reinforcement schedule as long as the schedule is brought on gradually and is relatively lean. Resistance to extinction explains why very lean schedules of reinforcement result in behaviors that are very persistent as well as occurring at a high rate. So both high rates of responding and resistance to extinction are often known as partial reinforcement effects.

Applications of Operant Conditioning

B. F. Skinner's operant conditioning principles have been applied in many areas. In education operant procedures have been used to develop programmed instruction, a teaching technique based on elements of shaping and chaining (Skinner, 1968). Programmed instruction was used extensively in Skinner's early development of teaching machines, which today have been replaced by computers. In fact, the text you are reading right now incorporates mainly operant principles in its approach to computer-assisted instructional design. Its reliance on adaptive adjustments to changing learner skills is a direct application of shaping procedures applied to higher-level reading comprehension and learning skills (Ray, 2004).

Operant principles are also used in various therapies. Miller introduced operant procedures as the defining technique in biofeedback (Miller, 1969); a therapy designed to reduce stress and physical reactions to stress. Alternatively, behavior modification for misbehaving children is often implemented in the home as well as in the school. Through secondary reinforcement offered through token economies (Ayllon |_2 Azrin, 1968) a client's good behavior may be positively reinforced and inappropriate behaviors ignored.

Animal trainers use operant conditioning procedures when training performance animals. Circuses, marine parks and zoos use shaping, chaining, and other procedures in order to teach animals to perform both for display and entertainment as well as for routine animal care, such as presenting a leg for drawing a blood sample. People also use operant procedures, sometimes unwittingly, when training their own pets (Pryor, 1985). Giving a dog a treat when it "shakes hands" is using the principle of positive reinforcement.

Programmed Instruction

Whether teachers recognize it or not, operant conditioning principles are often incorporated into the classroom environment (Skinner, 1968). In traditional teaching, instructors lecture and students take notes. Or perhaps students work either independently or in small groups on what has been previously presented by the teacher. But because there is only one teacher in a classroom, students do not get immediate feedback as to the accuracy of their work. Skinner's approach to this problem was to develop what is called programmed instruction.

Skinner also designed an apparatus called a teaching machine (Skinner, 1989) as a key element in the delivery of programmed instruction. The machine was a box with a window and a scrolling knob. The student was presented with some introductory material and then questions about that material for the student to answer. The student then scrolled the knob to reveal the answer to confirm whether they had learned the material presented. The student was to continue this process until he/she reached mastery of a series of such programmed sequences, or "frames" of material.

As the student progressed, the machine could be fitted to present more difficult material. Alternatively, it could also be taken down a level of difficulty if extra review was necessary. Through reinforcement (feedback indicating that the student answered correctly) and successive approximations (increasing difficulty of materials) students are shaped and taught mastery in any subject the instructor uses the machine to teach.

Does this sound familiar? It should, the operations in the teaching machine are the forerunners of computer aided instruction, and are the very basis for how the tutoring component of the software you are currently using works (Ray, 2004). Teaching machines never really caught on in mainstream classrooms, mostly because people feared that they were impersonal and lacked the warmth of teachers. Nevertheless, many students progressed far more rapidly and with far fewer errors in their learning using such teaching designs. Thus programmed instruction has evolved into the increasingly popular computer assisted instruction of today and research has found such instruction to be highly effective.

As noted, Skinner's teaching machine used the procedure of shaping, or successive approximations, to assist students in their learning process. Successive approximations to a final learning goal is often the foundation in programmed instruction. But not all forms of computer assisted instruction are based on this operant principle. Most of computerized teaching programs may look similar to programmed instruction, in that they ask students questions after readings and then give immediate feedback as to the accuracy of answers. But few such programs require that the student progress toward less and less dependency upon the programmed strategy. This is a shortcoming of all but the most sophisticated of programs that incorporate what is called "adaptive instruction" designs (Ray, 2004). That is, the learning goals are changed to adapt to the individual learner's changing skills and knowledge.

All computer-assisted instruction allows the teacher to spend time with students who are having difficulty while allowing more advanced students to continue and excel with immediate feedback. But few are designed to give and then fade supportive prompts and to present successively more difficult questions as the adaptive instructional software you are currently using does when used in "Tutor" mode of presentation.

Thus the computer assisted instruction material you are currently studying uses principles and procedures based originally on Skinner's programmed instruction (Ray, 2004). The tutor mode of the MediaMatrix software program turns your computer into a more modern and sophisticated version of Skinner's teaching machine. When you are in the tutor mode, the system helps prompt you as to the most important concepts and properties of those concepts, then asks you questions at the end of each segment of material presented. The system begins with the highest density of prompting, the smallest frame of content, and the easiest form of question, multiple choice. As you progress, the prompts are gradually faded, the unit or frame of content presented gets larger, and the questions become more difficult if you answer a series of questions accurately.

The MediaMatrix adaptive instructional system gradually moves from multiple choice questions to the less prompted fill in the blank, association recognition and, finally, minimally prompted verbal associates questions as you become more proficient in learning the material with less help and greater accuracy. If you begin to have difficulty with specific content or at a current level of difficulty, the program will successively drop levels until you are succeeding again. Because all lower level (multiple choice and fill-blank) questions are also represented in the association form of questioning, the system shapes the user into being able to answer accurately the more challenging association questions that depend upon total recall, as opposed to mere recognition, of the material.

This programmed instructional format, called adaptive instruction, relies upon artificial intelligence to compare the students' growing verbal or semantic networks of terms to an expert's network to adjust all of its varieties of presentations, including which questions you are asked. Such adaptive programmed instruction is designed to eventually wean the student from the need for programmed formats, thus teaching the student how to learn more traditionally presented materials through a shaping of that reading comprehension skill.

Such computer-based adaptive instruction can assist instructors and students alike (Ray, 2004). Students who have tutored on materials assigned prior to a class that intends to cover much the same topic of material find they are well prepared in the fundamentals, which allows the instructor to take more time teaching other dimensions, such as ethics, applications, or research foundations, and less time on simple definitional, conceptual, and review of fundamentals. The goal is to create prepared learners and to allow everyone to be at much the same level of understanding when the class begins. These were the goals of Skinner when he first designed programmed instruction and teaching machines. But it took the development of modern personal computers and sophisticated software development to achieve the real aspirations of Skinner's inventions based on the application of operant principles in and out of the classroom.

Therapeutic Applications of Operant Conditioning

Skinner's operant conditioning principles also are the foundation of various therapeutic applications (Skinner, 1972). Behavior modification, or the process of changing responses through stimulus control and token reinforcement economies, is one such application. Token economies (Ayllon |_2 Azrin, 1968) work to reinforce positive behaviors while simultaneously placing inappropriate behaviors on extinction. In psychiatric institutions token economies help maintain appropriate behaviors by reinforcing those behaviors directly. Likewise, disruptive classroom behaviors can be reduced through extinction and the reinforcement of more appropriate behaviors as well (Swiezy, Matson, |_2 Box, 1992). Even maladaptive physiological responses, such as anxiety or migraine headaches, can be addressed by operant techniques that use additional feedback, known as biofeedback (Miller, 1969), to help an individual know the state of normally unconscious bodily processes. It is thus worth considering each of these types of therapeutic applications of operant conditioning in a bit more detail.

Behavior modification is an operant approach to overt behavioral therapy and education (Bellack, Hersen, |_2 Kazdin, 1982).. The responses targeted for change are usually maladaptive for the individual and/or are inappropriate is given situations. For example behavior modification processes are often used in schools to help children whose behaviors have become disruptive and harmful to themselves or others. Behavior modification is also used in psychiatric institutions or institutions for the severely mentally challenged. The goal of behavior modification is to teach appropriate behaviors that serve the same function as maladaptive or absent behaviors (i.e. functions such as getting attention, help, praise, food, relief from boredom, etc.) This is usually accomplished with token economies being the specific type of consequential stimulus (reinforcement)

Token economies rely upon conditioned reinforcers, such as poker chips, points, or stars on a chart, used in a rule-guided process. A token economy utilizes both positive reinforcement and secondary reinforcers. Behavior modification is accomplished through the use of tokens to reward certain behaviors that occur during certain situations (stimulus discrimination) and thus represent a popular method of both antecedent and consequential stimulus control (Ayllon |_2 Azrin, 1968). Once an individual has accumulated a certain amount of these tokens (conditioned reinforcers), the tokens can be traded in for more direct and tangible reinforcers in the form of toys, favorite snacks, time with a computer game or anything that can be presented that will increase the probability of a behavior (positive reinforcement).

For example, consider a first grade classroom where playing with blocks is not a maladaptive or harmful behavior unless it occurs during the class time when a teacher is trying to instruct her class in preparation for some activity. Some children may start or continue playing with a toy during a time or circumstance such as this when it is inappropriate to do so. Instead of punishing the child to eliminate such undesirable behaviors, token economies can be used to create more desirable alternatives. For example, the child can receive a gold star every time they pay attention when it is appropriate in class. Then, at recess they can play freely and safely for even more gold stars. After so many stars, the child may then receive a favorite snack or time with a favorite game or computers. If the right reinforcement is used for the right behavior in the right setting, the child will begin to play with blocks only during play time and to pay attention during the class time that requires attending. This application of Skinner's operant procedures has been found to be very effective across many different situations, both educational and institutional.

Stimulus control in behavior modification may refer not only to the process of reinforcement, but also the process of controlling for antecedent stimulus discriminations in such a way that maladaptive behaviors become more appropriate and acceptable responses in specific situations. We cannot completely rid a person (or any organism) of a particular behavior. It may reappear anytime circumstances permit. But by controlling reinforcement and the discriminative settings where behaviors are appropriate, it is possible to create environments where maladaptive behaviors have no function and where new and more acceptable behaviors do. Eventually, acceptable responses replace those that are inappropriate. This is the essence of antecedent and consequential stimulus control, and it represents an application of Skinner's operant procedures.

Quite a different form of application of operant conditioning was developed largely by Neal Miller (1969) and is called biofeedback. Biofeedback is an operant approach to therapy that uses visual and/or auditory signals to reflect some internal state of the patient -- states that he or she would otherwise not be aware of. These signals, or feedback, serve as positive reinforcement when they indicate that the individual has successfully changed his or her internal responses in some target direction or amount. For example, someone who gets anxious in crowds may wear a heart rate monitor in a crowded situation. They then may read the monitor and use relaxation techniques to keep their heart rate under a certain level. The same can occur with high blood pressure. An individual can wear a blood pressure monitor at work and learn to keep it under a certain level by relaxing when a stressful situation presents itself. Since successful readings serve as positive reinforcers, people learn to relax in anxiety or stress provoking situations. The research on the effectiveness of Miller's biofeedback therapy is mixed, but it has been shown to be useful under various conditions, including control of migraine headaches (c.f., Sturgis, Tollison, |_2 Adams, 1978).

Operant Procedures in Animal Training

Response shaping is an operant procedure developed by B. F. Skinner to bring about new behaviors in an organism (Peterson, n.d.). This procedure is often used in animal training and usually, but not always, involves positive reinforcement (Skinner, 1951). Shaping procedures also include elements of extinction and is a process whereby the form or function of a behavior is gradually developed into the desired (target) response. Training a rat to press a bar (the target behavior) for food in an operant chamber is a common example of a shaping procedure. A rat generally does not press a bar very often, if at all, when it is first placed into an operant conditioning chamber (also known as a Skinner box). So how do we get it to do so?

Skinner used the processes of operant conditioning to find an answer this question. Why not begin by reinforcing the rat's behaviors that approximate a bar press, even if they are remote from actual bar presses, and then gradually shift the criteria for reinforcement to only those behaviors that more closely resemble bar pressing? Beginning with what the animal does relatively frequently, say looking at, going over to, and even just sniffing the bar (a behavior that occurs often when a rat is placed into an operant chamber), Skinner reinforced each of these to increase their probability. Then, as each of these behaviors became more likely, Skinner changed the rules of reinforcement to include only those behaviors that more closely resembled or actually were bar presses.

It is important to remember that following the extinction of a reinforced behavior an organism will typically increase the probability of that behavior and also engage in a wider variation of that form of behavior, often resulting in the emergence of new, but related, behaviors. Behavior does not instantly disappear as soon as extinction is implemented but rather reflects response induction, which is an increase in probability and variability as an early effect of extinction. The appearance of new, but somewhat similar or related forms of behaviors is thus another early effect of extinction.

So after the rat consistently emitted one of the "approximate" behaviors, such as first looking at, or later approaching, and even later for sniffing the bar, it was reinforced (usually with food) for doing so. But soon Skinner would no longer reward the behaviors that least approximated actual bar presses, hence initiating extinction for that behavioral approximation. As soon as that behavior was no longer reinforced, the rat engaged in response induction by emitting the behavior even more frequently and engaging in variations on that behavior.

One variation of sniffing a bar, for example, might be rearing up and placing paws on the bar. When this occurred, Skinner reinforced this new behavior. When placing paws on the bar reached a fairly high probability, Skinner would then stop reinforcing paws on the bar and the rat would again begin to emit new variations of such behaviors, one of which typically involves actually scratching at and even pressing down on the bar. Skinner would reinforce this and the shaping procedure would be complete. A bar press behavior had been taught through reinforced successive behavioral approximations to a behavior that might begin with a zero probability of ever occurring.

The response shaping process, because of its use of alternating use of reinforcement and extinction, is often called differential reinforcement of successive approximations in behavior. Successive approximations refer to the different behaviors that lead, step-by-step, to the target behavior (steps such as looking at the bar, then approaching the bar, then bar sniffs, paws on bar, and finally the bar press in our example). Differential reinforcement refers to the fact that we, at first, will reinforce any or all of these variations until one of the behaviors is produced reliably and then reinforcement is withheld so that new and different (hence the word differential) behaviors appear that more closely approximate the target response being shaped.

The process of shaping also incorporates the creation and use of secondary reinforcers. If you were to shape a dog to "shake hands", you may not want to have to give it food (a primary reinforcer) every time it emits the correct behavior. By the time shaping is half-completed, the dog may be satiated, and food may not work as a reinforcer anymore. Different schedules of reinforcement may not be appropriate in this case, either. What many people do is to use a child's toy "cricket" to produce a click, or to say "Good, dog!" right before giving it a treat (Pryor, 1985). Eventually, because of the pairing of the click or praise and food, the sound takes on reinforcing properties (it increases the probability of behavior). This is a classical conditioning, or stimulus contingency, procedure involving the pairing of the previously neutral sound (NS/CS) of the clicker or praise with food (UCS). This allows you to be able to reinforce the dog less with food and more with clicks or praise (now conditioned reinforcers) and hence to complete the shaping process without the animal becoming satiated on food.

In the case of operant chambers rather than dog training, the delivery of food is typically accomplished by a revolving magazine mechanism, much like those that deliver bubble gum one ball at a time from glass ball vending machines. The sound of this magazine shifting to deliver, in the rat's case, a food pellet serves as a secondary reinforcer much like the clicker or praise example above. This allows for behaviors that take place at quite a distance from the actual food dispenser to be reinforced via the secondary or conditioned reinforcer of the sound. The establishment of such secondary conditioned reinforcement functions is often referred to as magazine training.

There is a highly sophisticated computer simulation program available (at www.cyberrat.net) called CyberRat, that allows students who do not have access to live animal laboratories to experience both magazine training and shaping dynamics via simulation. This simulator uses an extremely large array (over 1800) of very brief digital video clips of live animals in a traditional operant chamber to produce the highly realistic illusion of a seamless real-time video feed showing a live animal being placed into an operant conditioning chamber and behaving exactly as real animals behave in this environment.

Through artificially intelligent algorithms the sequences of these clips may be altered through a student's reinforcement button by delivering simulated "water reinforcements" to any selected individual animal for its successive approximations to bar pressing. If you reward the animal appropriately, the video clips alter their sequences to simulate actual changes in behavior that simulate the entire operant response shaping process with highly realistic results. If you have access to CyberRat, visit the Appendix Elaboration on Shaping linked here for more details on how you can learn to shape a laboratory rat just as Skinner did! To access any or all of the Appendix Topics on how to shape a naive rat in CyberRat, click on the topic of interest:

Shaping a New Behavior.
Getting Ready for Shaping.
Understanding the Experimental Chamber.
Getting Your Subject Ready for Shaping: Habituation
Getting Your Subject Ready for Shaping: Magazine Training
Getting Your Subject Ready for Shaping: Observe Behavior Carefully
Begin Shaping (If Operant Level is Low)
Shaping: Not Too Slow, Not Too Fast
Other Factors Involved in Creating New Behavior: Prompting
Other Factors Involved in Creating New Behavior: Discrimination
Factors Involved in Creating New Behavior: Intermittent Reinforcement

Of course, shaping techniques are not limited to use on animals for simple training. Skinner (1989) demonstrated that the technique has wide applications with his teaching machine, a device that shaped the skills of human students in correctly answering questions in many subjects. Skinner broke down the complex tasks of learning a new subject into small successive units that gradually built into much more complex systems of knowledge. This technique was called programmed instruction and was the basis for how the teaching machine worked. Skinner's teaching machines served as the prototypes for many modern computer-assisted instructional and training programs.

The Cognitive Perspective

Many psychologists do not totally agree with an operant or behavioral interpretation of learning. They argue that classical and operant conditioning processes overly simplify how organisms, and especially humans, interact with their environments. These psychologists believe that you cannot dismiss cognitive (mental) processes when studying learning, as they believe the operant approach seems to do. This alternative focus on mental activity defines what is typically referred to as the cognitive perspective on learning.

Cognitive interpretations of learning have their roots in Greek philosophy, but resurfaced in the 17th century in the work of the associationists philosophers. British associationistic philosophers such as John Locke and David Hume, also collectively known as the British empiricists philosophers, believed that our experiences throughout life are critical in forming mental associations that define who we are and what we believe. Modern cognitive theories emerged when psychologists, rejecting exclusively consequence-based explanations of behavior, began to elaborate their own interpretations of learning.

Cognitive processes and activities such as information processing, mental representations, predictions, and expectations are central to the cognitive interpretation of learning. Cognitive psychologists don't completely discount the findings of the operant and behaviorally oriented scientists, they merely believe that there are also cognitive events involved in how organisms learn. For cognitive scientists, these events include internal processes which translate into a modern interpretation of a rather ancient concept: the "mind." Thus cognitive events are mental events.

For example, one of the early German gestalt psychologists, Wolfgang Kohler (1925) took a cognitive perspective when he explained the problem-solving behavior he observed in chimpanzees. He believed that these animals, as well as humans, could learn to solve problems though rather sudden insights about the character of a problem and its alternative potential solutions. He gave chimps such objects as boxes and sticks and then hung bananas out of reach to see what the chimps would do to get the bananas. Eventually they stacked the boxes and extended their reach using the stick to knock down the bananas. Kohler thought this sudden assimilation of objects was due to the animal's having a mental insight into possible arrangements that would solve their problem.

Likewise, Edward Tolman (1948) concluded through his work with rats learning to navigate mazes that animals learned about the structure of their environments without the required presence of reinforcement. He let rats explore mazes without the presence of goal boxes where reinforcing consequences were available. Later, these animals with non-reinforced experiences of mazes were compared with rats with no maze experience for their speed of learning when reinforcing consequences were now available. Rats with prior exploratory experience learned more quickly. Tolman thus proposed that animals as well as humans acquire a "cognitive map" which represents their surroundings mentally by direct associative experience. Based on this work Tolman discounted the need for direct contact with behavioral consequences as a necessity of learning. And because you can only see such learning when rewards are made available, he called the learning acquired from mere exploration "latent" learning. While cognitive psychology had its origins in this early research conducted during the 1930's and 1940's, it has become even more popular today.

For example Albert Bandura (1977) has rejected a strict behavioristic view by demonstrating that individuals can learn without coming into direct physical contact with behavioral consequences in their environments. His research identified the phenomenon called observational learning; a type of learning that occurs by imitating others who serve as models that we see being reinforced, and thus this form is learning is also called imitative learning or modeling. Bandura's best known example of imitative learning involved children watching other children on film acting aggressively during play (Bandura, Ross, |_2 Ross, 1963). Later, when the children who viewed the film were allowed to play in a room that included a "bobo doll" (a child-sized inflated standup doll that children can knock over and it bounces back to an upright position) these children were more aggressive in their play than children who had not watched aggressive models on film. The interpretation was that they had learned to be more aggressive because they had observed other children behaving in an aggressive way.

More recently, Robert Rescorla (1988) has suggested a cognitive interpretation of Pavlovian classical conditioning by suggesting an "expectancy" interpretation of the phenomenon. His theory rests upon the idea that experiencing two stimuli occurring closely together in time leads one to "expect" the occurrence of a subsequent stimulus whenever the first one occurs, and that this expectation accounts for conditioned reflexes.

Another research program that exemplifies a cognitive interpretation are the developments in what has been called "learned helplessness" and its role in the development of depression. Seligman and his colleagues (Maier, Seligman, |_2 Solomon, 1969) were the first to demonstrate that dogs denied opportunities to avoid shocks while in an experimental chamber called a shuttle-box later failed to learn to take advantage of avenues of escape or avoidance when such opportunities were made available. If dogs had such escape and avoidance possibilities from the beginning, they easily learned to avoid their shock presentations by jumping over a low barrier between themselves and the safe (no shock) portion of the shuttle box. But those who had, from the beginning, no avenue of escape never learned to jump to the safe side even when escape or avoidance was subsequently possible. Seligman went on to use this "learned helplessness" as a model for interpreting human depression, where little effort to exert counter-control over the negative events is a person's life is apparent.

Those who have participated in what has been called the "cognitive revolution" against behavioral interpretations, and thus have taken a cognitive perspective on understanding behavior, focus not on observable responses but on the inferred mental processes involved in learning. Stimuli in the environment serve as signals and the prediction of what follows, and this is an essential mental activity. Some cognitive psychologists believe that learning occurs through information processing activity that is exclusively mental, while others focus on the roles of mental representations in the learning process. Theorists who conduct cognitive research look to discover and identify the mental processes that occur when an organism is behaving and learning.

Cognitive theories of learning and behavior have practical applications, especially in therapy. Systematic desensitization was discussed in the section on classical conditioning, but cognitive psychologists have added their own procedures to this approach by using visual imagery rather than physical stimuli. Albert Ellis' rational emotive therapy (Ellis, 1973; 1993) is also a clinical application of cognitive principles blended with behavioral principles. This cognitive-behavioral therapy rests on the idea that inappropriate and self-defeating beliefs are the root of psychological disorders. Aaron Beck, another cognitive psychotherapist, has a similar view on therapy based on the belief that anxiety promoting patterns of thinking are what cause anxiety and depressive disorders (Beck, 1976; 1993).

Early Cognitive Ideas: Associations and Insight

The associationistic philosophers of 17th and 18th century Britain were some of the first philosophers to take a cognitive perspective on learning and behavior. John Locke, George Berkely and David Hume were such associationists and they rejected the then-prevalent notion of innate ideas. Instead, they theorized that humans form individual personalities through mental associations made through experiences with the environment. Locke (1690/1959) proposed that we are all born as a tabula rasa or "blank slate". Through experience, we form mental relations between contiguous (close in time and/or space) events and their effects on us. For example, a person who has a suspicious nature has become that way, perhaps, because when they began to trust someone, they were soon after deceived. These philosophers were known collectively as the British Empiricists as well as the Associationists and they believed that we make mental representations and process information in our minds as we grow and gain experience. The ideas and principles generated by the associationists form a foundation of modern cognitive theory.

One psychologist who later helped to define the early cognitive perspective was Wolfgang Kohler. Kohler is perhaps best known for his studies of problem-solving behavior in chimpanzees (Kohler, 1925). Originally trained as a Gestalt psychologist in Germany, Kohler proposed a very cognitive explanation of a chimp's behavior. In his research, Kohler would set up puzzles for the chimps to solve. One of the most famous was to tie a bunch of bananas to a string and then to tie the string to the ceiling of a large enclosure, thus placing the bananas out of reach of the chimps. There were, however, a large box and a long stick inside the enclosure. Kolher observed the chimps in this situation and found that the chimps made few false trials or errors. They would simply jump at the fruit a couple of times, pace around the enclosure, move the box under the bananas, get a hold of the stick and then knock them down. Kohler labeled this solution to a problem that had suddenly emerged "insight".

Kohler began studying the chimps because he was dissatisfied with Thorndike's theory of trial and error learning (Thorndike, 1898). Kohler's observation of his chimps in problem solving situations strengthened his conviction that his departure from the law of effect was a more accurate interpretation of how learning worked. Kohler believed that the solution came to the chimps as a mental representation of what would be successful. This is the process of insight and Kohler (1959) believed that animals as well as humans could overcome obstacles in this fashion rather than by Thorndike's trial and error. Mental representation is a key concept in Kohler's theory, although today many psychologists disagree with the notion that insight is as simple as a solution that just "comes to mind."

A good illustration of this more modern interpretation is Epstein's rather well-known demonstration that pigeons could be shaped using operant conditioning procedures to push small boxes from one place to another, and also to pick up and wield small sticks. When Epstein (1981) then gave these pigeons such objects in a replication of Kohler's earlier experiments where the box had to be moved under a reinforcer that could only be "knocked down" with a stick, pigeons with these prior experimental histories quickly did just as Kohler's chimps had done. Epstein points out that this "generative" process, as he called it, of generalizing prior training to slightly modified situations was simply just that: response generalization and multiple-response combinations to generate what appear to be novel behaviors but really aren't novel at all.

Place Learning and Latent Learning

A psychologist operating from the cognitive perspective during the middle 20th century was Edward Tolman. Dissatisfied with operant explanations of learning, Tolman (1930a) focused his studies on how rats learn to navigate through mazes. Behavioral psychologists who emphasized operant conditioning and its stress of reinforcement believed that rats learn to get through mazes because of reinforcing consequences resulting from winding their way through such mazes, learning turn by turn as discriminative cues. Tolman rejected the necessity of reinforcement (Tolman |_2 Honzik, 1930b) and developed many original research designs to test his theories emphasizing what he called place learning and latent learning. Tolman believed that learning occurred (in animals as well as humans) through mental activity such as insight and the formation of mental representations of the environment he referred to as cognitive maps. Tolman thus believed that animals learn more about "place" rather than how to engage in "habits" (Tolman, 1948).

To test his cognitive theory of place learning, Tolman created a maze with three different routes. One route was a straight path to the goal box where a reward of food was present. The second route was to the left of the first and was slightly longer because it had a small "c" shape in the beginning portion. The third route was to the left of the first and was the longest of the three as it had a half-square shape to it.

With experience in the maze, the rats came to prefer the first route and would regularly take it when placed in the maze. Tolman then blocked route one, and left only two and three as options. When the rats came to the blockade, they immediately turned around and took route two (the second shortest) with absolutely no training to do so. When Tolman blocked routes one and two, the rats would come to the blockade of route two and immediately take route three to get to the reward. From these results, Tolman concluded that, with experience, rats form mental representations or models of the maze in the form of cognitive maps. Tolman believed that the rats' movements in the mazes were not directed by discriminative stimuli, but guided instead by cognitive maps the rats had formed (Tolman, 1948).

Cognitive maps, or mental representations of the spatial layout of the environment, form the critical element in Tolman's theory of place learning. You probably have a cognitive map of how to get to your favorite restaurant. You do not get there driving "landmark-to-landmark" once you have learned the route. The only time you think about the route "landmark-to-landmark" is when you are telling someone who has never been to this restaurant how to get there. With experience, and through the reinforcement of enjoying a favorite meal, you now have a cognitive map of where other buildings and landmarks are in relation to it. You can automatically follow the route you have in this map to get to your restaurant instead of always reading a map or having to travel step by step.

While Tolman's theory of place learning is essentially cognitive, it has an ecological application to it. Many scientists working from Tolman's research have applied place learning in understanding how animals navigate their environment and remember important features, such as where to find food, where predators often hide, etc. They also apply this theory to understand how birds that nest in a community know how to find their "home nest" amongst many other nests that are in close proximity and look very similar.

Taking the cognitive concept of place learning a step further, Tolman wanted to show that animals could learn to navigate their environment without receiving any reinforcement. Tolman used a complex maze and three groups of rats in what he referred to as latent learning research. The first group had one trial in the maze per day for 11 days. These rats received no reward for navigating the maze. The second group also had one trial in the maze per day for 11 days, but these rats were rewarded for navigating the maze. The rats in the final group went 11 days (at one trial per day) with no reward, but were then rewarded on the trial of the 12th day.

The results supported Tolman's theory. Rats who never received rewards for completing the maze improved only slightly (improvement was measured in number of errors) over the course of 11 days. The rats that were consistently rewarded improved quickly until they reached a maximum efficiency toward the end of the study. The results of the third group of rats (those that had been rewarded only on the 12th day) were the most striking. After the initial rewarded trial, these rats were just as efficient at completing the maze as those rats that had been rewarded the entire time (Tolman |_2 Honzik, 1930b)!

Tolman explained these results as latent learning. The rats had learned to navigate their environment all along, but this learning did not emerge until it was reinforced. It took only one reinforcement for the rats to reach maximum efficiency. The learning was latent or hidden from view until reinforcement brought it out. It did not take reinforcement to learn the behavior; the behavior was simply observed and strengthened with reinforcement. This phenomenon sometimes explains how a small child will divulge knowledge on how to do something only when the time is appropriate, leaving the parent or teacher to ask, "Where did you learn that?" The child may have learned the information from TV or some other source and will only display the knowledge when it is appropriate or when they will be reinforced by praise and/or attention.

Observational Learning

Albert Bandura (1977) is a cognitive theorist who contributed throughout the latter half of the 20th century and continues as a strong force in cognitive psychology. He especially does not agree with Skinner's ideas about shaping and reinforcement as the primary way that new behaviors are acquired. Bandura often asked how it could be possible for people to imitate others and thus learn from mere observation (Bandura, Ross, |_2 Ross, 1961; 1963) if shaping and direct contact with reinforcement or punishment is necessary. Using this problem as a springboard, Bandura conducted many studies and identified the phenomenon of observational learning, or learning new responses by observing and modeling the behavior of others (Bandura, 1965). Bandura's studies demonstrated how humans, and eventually animals as well, can learn by watching others behave, and how mere observation of, not physical contact with, behavioral consequences in the form of reinforcement and/or punishment is sufficient for learning to occur.

An illustration of Bandura's concept of observational learning is one of his studies he conducted using kindergarten aged children. All the children in one study watched a film portraying an adult engaged in aggressive behavior (Bandura, Ross, |_2 Ross, 1963). The adult served as a simple model for such behavior, in that this adult was in a room full of toys and was seen verbally insulting, hitting, kicking, throwing and hammering a large plastic inflatable bobo doll. For one group of children, the model was reinforced for the assaults with candy and soda. For a second group, the model was punished verbally and then received a "spanking." The third group of children viewed the aggressive model where no consequences were given. The results demonstrated that humans could learn by simply witnessing the consequences of others.

When left in the room alone with a similar bobo doll, those children in the first group, (where the model was rewarded for aggression) displayed many aggressive behaviors by imitating what the model had done as well as showing novel aggressive actions. Those in the group witnessing the model being punished for aggressive behavior were much more gentle with the doll and displayed few, if any, aggressive acts toward it. Children in the group who witnessed the model receiving no consequences for aggression where more ambiguous in their behavior, showing some aggressive and some gentle behaviors.

Bandura emphasized the role of observation, attention, imitation and expectation in this process. First, an individual must observe and pay attention to another person (serving as a model illustrating the behavior). You cannot learn from someone else if you are daydreaming or paying attention elsewhere. You simply will not be able to see or remember their behavior. The next requirement is a physical ability to imitate the observed behavior – an individual must be able to imitate the behavior of the model. We'd love to learn to fly by watching an eagle, but this simply can't work. Someone suffering from paralysis cannot learn to walk by watching someone else. Finally, there needs to be some type of expectation of consequences. If someone witnesses their friend being reinforced for volunteer work, that person is likely to imitate this and try serving their community as well. Behavior can still be imitated and initiated if there are no consequences, but observational learning is much more efficient if the model's behavior results in some type of consequence; whether punishment or reinforcement.

According to Bandura, behaviors can be learned simply by observing a model being reinforced for a behavior, and then by imitating that model's behavior. So the question becomes, do children who watch violence on TV learn to commit crimes and become violent criminals? The answer is yes and no. It is true, as clearly seen in Bandura's studies, that children can learn to be aggressive and to perform acts of violence by watching and imitating others. They may even try these behaviors as a means to acquire what they want. If, however, children are consistently punished for violent behavior while simultaneously reinforced and praised for appropriate behavior, they may not become violent individuals. In this case, the child learns that violence is not the means to get what one desires. If the child does obtain reinforcement (perhaps by gaining attention) for using violence, then yes, they may very well develop aggressive and violent behaviors and may become much more likely to be involved in crime. While anyone can learn behaviors through imitation and observation (Bandura, 1977), operant conditioning can still have effect whether those behaviors become frequent or not. The converse is also true as an individual can learn a behavior through operant conditioning, but their behavior frequency can be affected by the observation of others being reinforced or punished by a particular behavior.

Learned Helplessness

Learned helplessness is a phenomenon Martin Seligman and his colleagues (Maier, Seligman, |_2 Solomon, 1969) identified in his studies of negative reinforcement and punishment. Seligman negatively reinforced one group of dogs for jumping over a barrier in an apparatus called a shuttle box. This apparatus is little more than a cage divided into two separate sections by a barrier wall that, usually, may be jumped over to escape from one side of the box or the other.

Seligman's procedures involved trials that began by turning on a warning stimulus. This warning was quickly followed by a brief presentation of electrical shock delivered to a dog through the floor of one compartment of the shuttle box. Only one side of the two-compartment shuttle box was ever electrified at a given time, and this was always the side the animal was standing in when a trial began. If the dog jumped over the low barrier wall to get to the other chamber of the shuttle box, it escaped the shock. Soon this negative reinforcement resulted in the animal jumping the barrier as soon as the warning stimulus comes on, thereby avoiding the shock altogether. Thus this group of dogs showed no ill effects of the procedures and quickly learned to jump as an avoidance response as soon as the shock began.

Another group of dogs experienced the same trials of warning stimulus followed by shock, but the barrier in their case was too high to jump over. Thus, no matter what they did they could not escape nor avoid the shock. Seligman subsequently lowered the barrier in these dogs' shuttle box to the same height as used for the group of dogs who easily learned escape-avoidance, and set the experiment so they could escape the shock. But Seligman (Seligman |_2 Maier, 1967) found that these dogs did nothing to attempt to escape their shocks! They would just cower in their cages, whimpering and taking the shocks. He referred to this failure to learn the escape-avoidance that normal dogs easily learned as a form of an acquired cognitive state of "helplessness" and thus these dogs had learned to "give up trying."

These "helpless" dogs became very inactive, lethargic and would sometimes stop eating when they were not in the experimental conditions. Seligman labeled this state as learned helplessness (also sometimes referred to as conditioned helplessness) and concluded that it occurs in humans as a form of depression. Many people become overly dependent or depressed because all of their attempts to escape or avoid negative situations have failed. Eventually, people give up and an attitude of learned helplessness develops.

Applications of Cognitive Learning Theories

The cognitive perspective on learning has many applications that go beyond TV violence, especially in therapeutic situations. Therapists specializing in systematic desensitization sometimes apply cognitive principles, such as visual imagery and mental representation, to the counter-conditioning process. Ellis' rational-emotive therapy (Ellis, 1973; 1993) works to change a patient's beliefs about a certain situation in order to alter perceptions of negative consequences and hence, negative feelings. Beck's (1976; 1993) cognitive therapy, which is similar in many ways to rational-emotive therapy, leads patients to understand their patterns of inaccurate and anxiety provoking patterns of thinking. Beck believes that by changing these patterns, patients can overcome feelings of depression and anxiety.

Outside of the therapeutic environment, cognitive research in learning is often applied to the computer sciences as guides to developments in various forms of artificial intelligence. For example, artificial intelligence research on informational input includes pattern recognition problems, such as interpreting hand writing and spoken language inputs so that they may be translated into computer codes to make them useful as control commands or text production (e.g., dictation input programs). On the processing side, artificial intelligence includes neural network simulations, programs that can learn based on user feedback, problem solving and simulation algorithms, and even programs that can play world-class chess against human competitors. On the output side, artificial intelligence research includes the development of fabrication "printers" that can manufacture objects directly from digital blueprint images, robotic and neurally controlled prosthetics for amputees, and even attempts at spoken conversational language where the computer participates as a "social" entity.

The concept of systematic desensitization is discussed in detail in the applications of classical conditioning section. While desensitization is very much a behavioral therapy, therapists who take a more cognitive perspective on learning often add visual imagery or mental representation in the desensitization process. For example, instead of presenting either a real snake or even a picture of a snake to a patient who fears snakes, a therapist might begin with instructions to imagine seeing a snake at a distance, then to also imagine gradually walking nearer to it while staying very relaxed. This application of cognitive psychology's emphasis on mental activity is a way for patients to practice desensitization without coming into contact with the actual feared stimulus or even a physical representation of it.

Visual imagery also allows patients to learn how to react to a feared stimulus by imagining what they might do if they were to abruptly come into contact with it. For example, a person with arachnophobia can go home after training in a therapy session and practice by imagining what it would be like to encounter a large spider. They can imagine spiders and experience the emotions at a more acceptable or manageable level as well as mentally formulating a plan of action without having to actually be in contact with a real stimulus.

This visual imagery also works well for those with a phobia of flying. It is impractical for the therapist to continually go on therapeutic flights with the patient. So during therapy sessions, before an actual flight is set up, the patient "practices" by visually imagining being on an airplane; letting themselves feel the emotions they will experience while they plan appropriate ways to react to those emotions.

Cognitive Applications in Therapies

Cognitive approaches to therapy emerged from relevant research on learning and problem solving . The most common psychological disorders that are treated from a cognitive perspective include depression and anxiety. Cognitive explanations for these conditions emphasize a person's negative beliefs and irrational interpretations of situations. One of the most prominent of such an approach is Ellis' Rational-Emotive Therapy (Ellis, 1973; 1993). Beck's approach to cognitive therapy also emphasizes negative patterns of thought (Beck, 1976; 1993). Research has found that both therapies are effective in treating many disorders including depression and anxiety.

Ellis' rational-emotive therapy is designed to change maladaptive behavior by changing irrational interpretations that individuals make in certain situations. The skeleton of rational-emotive therapy is Ellis' ABC (Activating event, Belief, Consequence) model of psychological disorders. As a cognitive perspective on therapy, this model uses the concepts of mental belief, interpretations and emotions. The therapist must break into the ABC model and change these beliefs before the patient's emotions and behavior can be altered.

Ellis feels that if you alter an individual's irrational belief about some event you also change the consequences of such beliefs, which take the form of negative emotions. With these negative emotions gone the maladaptive behaviors should also disappear. Ellis emphasizes that a therapist's job in rational-emotive therapy is to illuminate the maladaptive mental processes that occur in a patient and then to teach alternate ways of looking at a situation. Under this therapeutic approach, the therapist demonstrates to patients how negative their beliefs are and helps them to change their beliefs so they no longer feel such negative emotions, and thus they no longer exhibit maladaptive behaviors.

An activating event, (representing the A in Ellis' ABC model of psychological disorder) is any event or situation in the life of a patient that causes that patient to develop a negative belief or irrational interpretation. This belief, in turn, leads to negative emotions and maladaptive behavior. Activating events can be virtually anything in the life of the patient from missing a bus to a death in the family. These events can sometimes be under the control of the patient, but sometimes they are beyond anyone's control. The therapist's job is to change the beliefs that correspond to these events; it is not the therapist's place to interfere with the events themselves.

The B in Ellis' ABC Model in rational-emotive therapy stands for the beliefs that patients form as a result of their experiences with activating events. These beliefs are the key to understanding the nature of the patient's problem and are the starting point of therapy. Such beliefs often take the form of irrational interpretations of certain situations in the patient's life. It is only when the therapist leads the patient to understand and change these beliefs that the patient can begin to reduce the negative emotions and alter the maladaptive behaviors he or she experiences.

The C in Ellis' ABC model represents the consequences a patient experiences because of the negative beliefs he or she holds regarding a specific activating event. In

Ellis' cognitive model, consequences are often in the form of negative emotions experienced by someone suffering the psychological disorder brought about by these negative beliefs. A therapist using rational-emotive therapy reduces or eliminates these consequences indirectly by changing the patient's beliefs about the events in his or her life.

Aaron Beck's therapy (Beck, 1993) is also an example of a cognitive approach to therapy that is quite similar to Ellis' rational emotive therapy. Beck's cognitive therapy rests on the premise that the difficulties and disorders that people experience are due to anxiety-promoting patterns of thinking. Beck believes that individuals who suffer from psychological disorders, such as anxiety and depression (depression being the original disorder the therapy was designed for), are constantly thinking about themselves and events in their lives in very negative ways. These negative thoughts, according to Beck, are the source of the problem rather than any physical event or flaw.

Therapists using Beck's cognitive therapy are taught to use Socratic questioning; a method of questioning that leads patients to identify these negative thought patterns. Once the patient identifies their anxiety promoting thought patterns, they are taught how to alter this and begin to see themselves and events for what they truly are. Minimizing the positive events in one's life is a major form of this negative thinking. For example, a person suffering from depression might be given a birthday party. When asked about it he or she may say, "They just did it because they feel bad for me. I could tell no one really wanted to be there." This individual is minimizing a very positive event in their life. A therapist using Beck's cognitive therapy would then, through Socratic questioning, illuminate this pattern of thinking for the patient and help to change it.

Another form of an anxiety-promoting pattern of thinking is maximizing the negative events that occur in one's life. In Beck's cognitive therapy this is highly maladaptive and can lead to psychological disorders. An example of maximizing the negative may be a patient in cognitive therapy who states, "I had an argument with my best friend. I am no good to anyone and I will never have any more friends."

Mis-attributing fault to oneself is another common anxiety-promoting pattern of thinking that, according to Beck, can lead to psychological disorders. Someone who tries to accept blame for a friend losing his or her job, when the individual had absolutely nothing to do with the decision is mis-attributing fault to the self. As with other faulty thoughts, Socratic questioning illuminates this pattern of thinking for the patient and helps to change mis-attributing fault.

Ecological Perspectives on Learning

While studying various procedures in learning, such as classical and operant conditioning, some scientists have questioned the role of artificial laboratory models and have thus assumed a more ecological perspective in understanding the learning process. Some behaviors that are more complex than Pavlov's simple reflexes appear to require no learning at all, such as beavers building dams or birds building nests. Other behaviors appear to be extremely easy or difficult for a given species to learn (Seligman, 1970). Still other behaviors can be demonstrated to develop, and sometimes very quickly or easily (Seligman, 1971), with some developing at any time while other developing only during "critical periods" of an organism's development.

Such variations call into question a key assumption of early learning theorists: that all forms of behavior are governed equally by the broadly applicable principles of learning (Seligman, 1970) -- whether those principles are based on classical, operant, or even cognitive procedures and interpretations. Those taking a more ecological perspective on learning focus on the effects of environmental context as well as the characteristics of a given species being conditioned. As such, ecological researchers assert that the generalized principles of learning must be contextually interpreted and are thus more limited in how, and to what behaviors, such principles apply.

For example, early operant scientists (Breland |_2 Breland, 1961) noticed that when shaping or training organisms to do complex tasks, many would revert to natural (that is, apparently unlearned) behaviors seen in all members of the species. Thus when teaching pigeons to pull a string for food, many would sporadically peck the string instead of pull it. From a review of such literatures Seligman (1970) concluded that this is due to the fact that a pigeon is much more biologically "prepared to learn" (thus defining what Seligman calls "preparedness") to peck at something than to grasp and pull it with its beak when food is the consequence. Reasoning much as the ethologists might, Seligman also asserts that in addition to preparedness there are certain biological constraints in the pigeon's natural environment and physiology that make a pigeon's use of its beak more successful for finding food by pecking than by pulling. Seligman's concept of behavioral preparedness thus includes the notion that an organism can be prepared, unprepared and even contra-prepared for learning a specific form of behavior.

Another ecologically effected learning phenomenon is bait shyness, or conditioned taste aversion. For example, Garcia and his colleagues (Garcia, Kimeldorf, Hunt, |_2 Davies, 1956; Garcia, McGowan, |_2 Green, 1972) found that a rat stops eating a given type of food if it later experiences nausea. In other experiments thirsty rats were given saccharin-sweetened water to drink. All animals were presented combinations of external stimuli that accompanied their drinking. These stimuli included a click and a flash of light as well as the taste of saccharin-flavored water each time the rat licked at the water dispenser. One group of Garcia's rats received a painful shock after the presentation of the click, the light and the water independently. The other group of rats received X-irradiation that would elicit nausea after experiencing each of the stimuli independently.

When Garcia tested for associations, he made a surprising discovery. The rats that were given the shock after each of the stimuli displayed aversion only to the click and the light, not to the flavored water. Those rats that received x-rays after each stimulus only displayed aversion to the flavored water and not to the click or the light. Garcia concluded that organisms are biologically predisposed to develop certain associations between stimuli and that these take precedence over other relations. The rats in his study were prepared for the association between click/light and shock as well as taste and nausea. These associations are successful for survival in the natural environment and are easier to make because the organism is biologically constrained or prepared to do so.

Ethologists are behavioral biologists who study complex behavior patterns specific to a given species. In some cases, such patterns are consistent from one member of the species to all other members of the same species. But in other cases, there may be slight or even dramatic variations in either the behavioral pattern or the stimuli which elicit the pattern. Ethologists thus also take an ecological perspective on learning because they seek answers to questions regarding genetic contributions to complex behaviors that are clearly impacted by learning as well. For example, Thorpe studied how European chaffinche birds acquire unique dialects in their otherwise species specific songs. He raised some hatchlings where they could hear the natural songs of their species, and another group in complete auditory isolation. He found that the song patterns in each group had some components in common, but that the group exposed to natural examples of their species had other components to their songs that their isolated cohorts lacked (Thorpe, 1956).

Other ethologists, such as Konrad Lorenz (1935/1970), explored the role of critical stages of maturation and development in the phenomena of imprinting. Imprinting involves learning to stay close to or to follow movements of a mother figure rather than other adult members of a group. Lorenz even discovered that recently hatched geese goslings would attach and follow an adult chicken rather than their real mother if they experienced that chicken as a dominant substitute for their mother. He also found that hatchlings would follow him everywhere as a surrogate version of their mother if he had been the dominant figure in their environment at a critical stage of their young lives. Both Thorpe (1956) and Lorenz (1955) illustrate the complex interaction between phylogenic predispositions to acquire behaviors and ontogenic associations between behaviors and discriminative stimuli that make genetically shared behaviors for all members of a species nevertheless unique in their form of expression for individuals of that species.

APPENDIX TOPIC: Shaping a New Behavior

Now that you have studied some of the principles of operant reinforcement, extinction, and stimulus control, you are ready to put these principles to work. One good way to see the power of reinforcement in our lives is to practice the use of a reinforcer to train a new behavior. Coaches do this all the time. For example, they give feedback for skillful performance hoping this feedback will encourage the athlete to repeat that skill the next time the opportunity arises. In the CyberRat shaping exercise you will act as a coach - for a rat. Not a rat you can actually hold in your hand, but perhaps the next best thing, a simulation that will give you an experience very similar to what you would have if you were to coach a real rat to press a real bar (the skilled performance) in a real operant chamber.

Have you have ever wanted to "talk to the animals?" In this exercise we hope you will see that reinforcement provides one way to communicate with another individual - including an individual such as a laboratory rat. When you give your simulated rat a drop of water as it presses the bar in an operant chamber, you will be playing one role in a two-individual conversation that takes place between the rat and you. If your drop of water qualifies as a reinforcer for the rat's behavior, soon the rat will be pressing, you will be giving "feedback," and your conversation will be going back and forth like that of a good coach and a star athlete. As with good coaching, success will depend on your timing and the correct selection of behavior to emphasize in your conversation. With practice, you and your rat will both succeed.

This appendix was written to provide specific discussions of the processes involved as you begin to train new behavior. We will use a specific example -- shaping CyberRat, the simulation based on a real laboratory rat, to press a lever for water reinforcers. The process and principles of shaping, however, are more general and the discussion should help you prepare for other examples of training as well. We suggest, however, that you should study your CyberRat User's Manual both before you read this appendix and also that you keep a copy of the CyberRat User's Manual nearby for cross reference. So, to learn how to shape, let's consider some specific tips, including:

Getting Ready for shaping.
Undertanding the Experimental Chamber.
Getting Your Subject Ready for Shaping.
Begin Shaping
Other Factors Involved in Shaping.

APPENDIX TOPIC: Getting ready for shaping.

Before you start shaping CyberRat or any other individual, you will need to make some selections concerning the relevant experimental and subject variables involved in the process. In this section we will consider issues on selecting your rat, which include considering both family and individual subject histories as well as manipulated conditions such as deprivation.

First, choose your subject and know something about his or her family history. We have given you a bit of a challenge by inviting you to shape the behavior of a laboratory rat. What do you know about such an individual? We should emphasize that the simulated CyberRat you will get to know is actually made up of videos taken of a real laboratory rat. Therefore CyberRat has mannerisms and behavioral patterns that are real - not just drawn from the imagination or expectations of a cartoonist. You will work well with CyberRat if you are aware of what makes a real laboratory rat "tick."

Laboratory rats are not as fearful of humans as are their wild counterparts. In fact, laboratory rats were genetically selected many generations ago because they were tameable and, in fact, friendly when well treated. Today's laboratory rats continue that tradition. Second, laboratory rats are generally hearty and healthy. They keep their bodies quite clean by "grooming" their hair with their tongue and front paws. We think you will see your CyberRats doing this.

They are sociable, and once acquainted quite playful with each other and with their caretakers. They explore their environment and readily learn "what leads to what" (a skill we fully plan for you to observe!). They are capable little creatures -- strong enough to push with a force greater than their weight (if they have the right place to stand), jumping (with a vertical jump that is many times their height!) and climbing well and possessing excellent balance for running along narrow ledges if given a chance.

They hear very well (detecting much higher pitches than humans, for example), are good at identifying odors of interest to them (probably better than are we), and see well enough to navigate about a room (remember, in their normal living conditions they are active at night and hence less "visual" than daytime animals such as humans). They tend to poke their nose into small openings and sniff the air coming through and they seem to use their touch-sensitive "whiskers" (called vibrissae) to inspect objects closely.

They balance well on their "haunches" (they often rear up on their hind legs to inspect objects above their head) and they use their front paws to hold and manipulate things (like the bar we will offer them). Sometimes they use their teeth to "explore" or even gnaw on things, holding and "shaking" objects as if to discover their properties, or holding objects to "drag them away." One of the authors has, in fact, had his hand pulled into the cage by a rat with whom he was quite friendly -- all done quite carefully, as though the rat sought the person's company.

You should know also, of course, that a rat's teeth are very effective incisors that can pierce and cause damage when the situation calls for it (or when inept mistakes are made in handling the animal without due support, care, or respect).

Second, know as much as possible about what your selected rat's specific individual history is. Is your rat "experimentally naïve"? Researchers often divide the life of their rat participants into two stages -- before they became subjects in scientific research, and after their career has begun It is doubtful that laboratory rats make this same division. For them, they are learning all the time, always "building on" previous learning. Still, once your rat enters the operant chamber for the first time, it will start accumulating experiences that will have much more specific influence on its future behavior in that chamber, so in this case naïve simply means no prior experience in a chamber like the one we will use for experimentation, and no history of being taught any behaviors like those we will attempt to teach. Once it has learned to press the bar for water, for example, it will ALWAYS be easier to retrain this behavior, even if years have gone by and many other things have been learned and unlearned.

For this reason, we encourage you to start by selecting your CyberRat animal quite specifically for what you want your rat to accomplish in the operant chamber. As you will read below, any early experience in the chamber provides some "habituation" to such chambers. Do you wish to observe that process? As rats learn about the water delivery and its associated sounds (what we will call "magazine training" below) they will be much more trainable when you seek to use water reinforcers to influence their behavior. Do you wish to carry out magazine training? Your instructor may actually suggest that you select a rat that has already been magazine trained so you will be able to more quickly shape bar pressing. But most will request that you start with an experimentally naïve animal so that you can experience the ENTIRE process of training an animal, not simply one stage or dimension of that process. Even if you start with a rat that has been magazine trained, however, please remember that successful training always "builds on" prior experiences that are important in preparing the individual to learn "the next step." And, if you have time, we also recommend that you start with an experimentally naïve rat so you can participate in developing this prior experience.

Finally, a word about prior water deprivation. One condition you will select as a parameter of one very important experimental (independent) variable after you choose your rat is the level of water deprivation it will bring to the experimental session. Whether a drop of water is or is not a reinforcer depends on the current level of water deprivation, or to be more precise, the time since your animal last had water available for drinking. Within limits, of course, the longer the deprivation the more effective the water is as a reinforcer. Rats generally drink each day, often alternating between small drinks and small bites of food during their "meals". If a rat has access to water for 15 minutes in a day, however, it will generally maintain a good water balance unless the temperature is hot and very dry. If you choose a rat that is 23 hours water deprived, then you will be able to use drops of water as an effective reinforcer without doing any harm to your animal. This also allows one hour each day to conduct your experiment and to follow with a period of free-access to water after the session. That is why CyberRat defaults to a setting of 23 hours water deprivation.

At this point we may now turn our attention from understanding the rat we have selected to a better understanding of the experimental chamber that also has been selected as the environment in which you will train your animals.

APPENDIX TOPIC: Understand the experimental chamber.

The CyberRat simulation takes place inside a real "Skinner box" displayed via video (B.F.Skinner actually preferred a more generic name for the simple space that he designed. He might want us to call it an "operant conditioning chamber."). Usually such a chamber is a small cubic space, perhaps 1 foot on a side. This chamber also contains several specific and unique elements that you need to understand, including the water delivery mechanism, or magazine; the manipulandum; the stimulus lights; as well as some construction features that are relevant.

The water dipper, or magazine. As pictured below, the wall of our chamber has a square opening on its lower left side that leads to a small enclosure. In some chambers this enclosure has a small hole on its floor and a small metal cup can be raised through this hole to provide a drop of water to the rat. In other variations on water delivery an electrical solenoid, or valve, is used to allow one drop of water to be transported from a bottle through a tube, with the drop being deposited on the floor of the cup that fills the back of the opening. This is the form of delivery used in the films upon which CyberRat is based, and it makes a practical difference. In the dipper delivery of water, only one drop of water is ever present, even with repeated operations of the dipper. That is, the dipper cup only holds one drop and submersing it while that drop is still there only replaces that drop with a new one. With the solenoid/valve delivery, each delivery accumulates, thus giving the animal a potential "reserve" of several drops if several deliveries have been made since the last visit to the water delivery area. Again, CyberRat uses this solenoid delivery, and thus drops of water accumulate. In either case, the size of the drop can be experimentally altered, thereby changing the "amount of reinfocement" given each time. As you read the CyberRat User's Manual you will find that you can change the size of the water drop you deliver within CyberRat's experimental parameters screen. This is one of several details of the chamber that is available for you to vary.

The manipulandum (or operandum). A rectangular bar, or lever, is mounted in the middle of this same wall, at a height that allows the rat to place its paws on the top and press down. This "lever" or "bar" moves about 2 cm when the rat presses down on it. Only a small force is required to press the bar, and pressing it causes a switch to close so the equipment can record that a press has occurred. Because we will be training the rat to "manipulate" the bar and because we will most often take pressing this bar as the operant behavior we will shape and study, this bar is sometimes referred to as "the manipulandum" or as "the operandum."

Stimulus lights. Besides the water dipper and the bar, this same wall contains two stimulus lights. If you are studying how a behavior comes to be emitted only in some situations and not in other siguations (e.g., light on rather than light off, left-light vs right-light), you may turn these lights on and off as signals to the rat in the chamber.

The rest of the chamber. A plain metal wall is opposite the one with the bar and the other equipment. The floor is composed of metal rods placed close enough together so that the rat is comfortable moving around but far enough apart to allow urine and feces to fall through into the waiting pan of wood shavings below (for easy removal and cleaning). The ceiling and the remaining two side walls of the chamber are made of clear Plexiglas plastic. The chamber itself is placed within a smooth plastic "shell" that provides a quiet, evenly lighted spot for the chamber. A fan in this shell keeps the air fresh inside.

Now, we have described this setting as though it is not very interesting, but you will notice that your rats available in CyberRat will spend many minutes moving around inside the chamber, seeming to sniff at one thing (especially at corners), nibbling on small details such as the bars on the floor or screws holding things together, rearing up to the ceiling, etc. The rat will also spend some time grooming his/her fur and face, much like a cat does. You will have many kinds of behavior to observe in even this simple operant conditioning chamber. As you train your rat to press the lever, however, you will notice that these other possible activities become far less frequent. So let us now move to considering how to best get your subject ready for shaping bar pressing.

APPENDIX TOPIC: Get your subject ready for shaping: Habituation

There are several elements of preparation that are important to accomplish properly if you wish to shape your rat as quickly and efficiently as possible. These include habituating your subject to the novelty of being in a totally foreign environment such as the operant chamber, magazine training the animal to respond to the sound of water delivery, cautions against over-watering your rat in a single session, learning to observe all variations of behavior very carefully, and learning to measure the operant (before conditioning) rate of the behavior you wish to modify or train. So let's look at each element in more detail.

Habituation of exploration, of startle. When you first put your rat into the operant chamber, you will probably see it move about, sniffing and touching all parts of the chamber. You might describe the rat as "a bit on edge, or very alert." After a while, however, it will move around less and you might describe it as being "more comfortable." In this quieter state, your water reinforcers will have a better chance to influence the rat's behavior. There is still a time, however, when a novel noise will possibly evoke a jump or startle reaction. When you first operate the water dipper, for example, you may notice this reaction -- more in real animals than the simulated rats in CyberRat's colony. Don't worry. After this sound is paired a few times with the arrival of water, there will be no startle, just eager movement to the water dipper. You should probably give the rat a few minutes to become less reactive to the environment (we could describe this phase as "habituation" to the chamber - a kind of learning we might describe as "settling down"). In fact, CyberRat allows you to run an entire "before-conditioning" habituation session where no water is available for bar pressing or manual delivery by you. We recommend you conduct such a session for anywhere between 20-60 minutes as an entire, separately identifiable "habituation" session so that measures of all the behaviors prior to shaping will always be available for later comparisons as you conduct experiments to change these behaviors.

Thus, if you wish to measure operant level of all forms of behavior without the complications of presenting drops of water or of having bar presses produce drops of water, you need to set the experimental condition (Schedule) to "Habituation" before you begin your experimental session. Your CyberRat User's Manual will illustrate how to accomplish this.

After running a prior habituation session, it is still a good practice to allow at least a few minutes of additional habituation in the next session prior to operating the water dipper. To setup CyberRat to allow you to deliver water reinforcers for training purposes, select Manual Reinforcement as your schedule. If you are going to attempt to shape bar pressing you will also want the bar to deliver water reinforcers if it is pressed by your rat. To setup CyberRat parameters to do this, as you select the Manual Reinforcement schedule, select the sub-menu item of "Bar ON". If you intend to shape some behavior other than bar pressing, you would want Manual Reinforcement with Bar OFF.

Within 3-4 minutes of your new session your animal should be ready to learn to find the water if and when it is presented. Begin by delivering one drop of water manually when the rat is a very short way away from the dipper or has poked her nose into the reservoir (this will make the sound of delivery far less startling to the animal). But, after the first 8-10 deliveries, you might start pressing your reinforcement button only as the rat is moving her head either out of the reservoir or toward the water delivery area. The next section that explains the concept of magazine training will tell you why this is a good idea.

APPENDIX TOPIC: Get your subject ready for shaping: Magazine training

Establishing a location where each reinforcer will be provided along with a sound as a signal that the reinforcer has arrived (conditioned reinforcer) is an important step in allowing you to shape behavior elsewhere in the operant chamber. An old-time meaning of the word "magazine" is "a reservoir or storage place, especially for provisions" (or for gunpowder, but that's another story). "Magazine training," therefore, became the phrase used to describe teaching the individual you are about to train as to where it can find the reinforcers you will use to accomplish your shaping. Often the arrival of a reinforcer at that spot ("the magazine") produces a specific sound that signals its arrival (in our case a click as the solenoid delivers a drop of water).

For a water-deprived individual, the association between these clicks and the arrival of water that they signal is the basis for a variation of Pavlovian conditioning that establishes the click as a reinforcer -- a conditioned reinforcer. This magazine training is an important prerequisite to successful shaping. We will emphasize below that reinforcers should be given without delay when the rat emits a response that you are hoping to reinforce. You need to have the "click" of the dipper be firmly associated with the rat finding water in the dipper, so that this sound will be as effective a reinforcer as the water itself.

Beware of Satiation--not too many "free" drops of water, please. If you provide water too rapidly, the rat will stop drinking temporarily. We say that the rat is satiated. Notice, however, that all we really know is that the rat will not approach and drink from the reservoir. Water delivery no longer increases the frequency of approach to the reservoir after its "click" (i.e., clicks or water no longer reinforce approach). On the other hand, if you separate the deliveries of water by several seconds, the rat will continue to approach and drink after each click of the dipper. Of course, after perhaps a hundred deliveries the rat will be satiated even if the rate of delivery is slow. Good shaping requires that you achieve a balance between giving enough water deliveries to keep the rat engaged in the behavior you are shaping, but not so many that deliveries lose their ability to reinforce behavior. With some practice you should become expert at achieving this balance.

Once you have accomplished magazine training, of course, you have added another "reason" for the rat's behaviors in the chamber - the rat might be "water seeking." In fact, the goal of shaping is to develop a specific way for the rat to obtain water -- by pressing the bar. When a behavior consistently produces a consequence, that consequence may positively reinforce the behavior -- that is, increase its future frequency. Another way to describe such an increase is that you have provided another "reason" for emitting that behavior.

Now you may begin the REAL process of shaping some new behavior, such as pressing the bar. But the process of shaping starts by observing the existing and ongoing behavior of the rat and then selecting certain of these actions by consistently following them with a click and a drop of water. The actions you select should be those that move the behavior closer to "pressing the bar for water" - the target behavior we are hoping to develop. This makes observation skills critical in successful training and shaping.

APPENDIX TOPIC: Get your subject ready for shaping: Observe Behavior Carefully

We have discussed magazine training and the dangers of satiation. Now we are ready to get down to the work of shaping the rat's behavior. It is time to watch carefully what the rat is "doing." We put that word in quotes, since it is possible to think that all behaviors the rat emits are controlled in some way by the consequences of its behavior -- by what the behavior accomplishes or produces. These consequences of the behavior (the "reasons" for the behavior) can be viewed as natural reinforcers for that behavior. As you first observe your rat, you will not really know what these reasons are. The rat moves around the chamber. It stops and sniffs in a corner. It moves toward the protruding bar touching it with its nose. We ask what is the rat really "doing." Though we can loosely say that the rat appears to be "exploring" its environment, it is always risky to offer reasons for a behavior until we have studied that behavior thoroughly. As such, exploring may be more a description of the FORM of behavior than its purpose or outcome. It is important to describe the behavior merely as movements and to keep an open mind regarding the specific "reasons." Learning to describe behavior "neutrally" (without drawing conclusions regarding its reasons) is an important skill for psychologists to develop. In clinical practice, for example, a psychologist should keep an open mind about the reasons for a problem behavior until the repeating pattern of that behavior reveals why it keeps occurring (what it "accomplishes." What the person is "doing."). There is a complete and highly sophisticated coding trainer built into CyberRat that is accessed via the Multi-Behavior Analysis section of the upper-right menu. All behaviors in the "coding system" offered there are clearly defined by the form of behaviors, not their functions or reasons for occurrence. If you have significant difficulties in successfully shaping your rat, you may want to practice coding behavior to acquire better observation skills as they relate to how rats behave (see the CyberRat Manual regarding "coding"). From such descriptions and their quantification, you will be able to determine the "operant" (pretraining) level of each form of behavior that is emitted in the operant chamber environment (see Unconditional Probabilities graph in CyberRat's Multi-Behavior Analysis section).

Measure operant level. Does the rat press the lever before you begin to use your water reinforcers? If the rat is already pressing the lever "for other reasons" it will be quite easy to reinforce these presses with water. Every coach secretly hopes that his or her player already shows the skills needed to be great. Then all the coach need do is teach when each skilled behavior is appropriate. When a skill is high before specific training, we say that it has a "high operant level." So--does your rat already press the bar? Frequently? If so, you can wait until a bar press occurs "for other reasons" and then deliver your water reinforcer. A specific process of shaping would not be required since the target behavior could be reinforced directly.

However, we expect that your CyberRat will not have a high operant level of bar pressing unless you selected a subject that has prior training on this behavior. Most laboratory rats will occasionally rest their paws on the bar as they move about the chamber, or they will occasionally push at the bar with their nose. Usually these presses do not occur often enough to encourage you to merely sit and wait for them. It is a good idea, however, to be aware of how often such bar presses do occur and to take advantage of them when they do (that is, quickly deliver a reinforcer!!). CyberRat simulations keep bar press operant levels extremely low so that animals don't just learn by "trial-and-error" on their own, but rather require you to train each rat.

APPENDIX TOPIC: Begin shaping (if operant level is low)

OK, you have hopefully conducted a session using the Habituation schedule in CyberRat so you have a good operant level measure for each class of behaviors emitted within the operant chamber. You also have started to conduct one or more sessions where you have used magazine training techniques to establish a reliable "go-to-water" reaction within a few seconds after delivering water (and its sounds associated with delivery). Now you are ready to reinforce the closest available behavior that looks like it might be a component of bar pressing. Perhaps you have learned that your rat does not press the bar very frequently "for other reasons." It has a low operant level for bar pressing. Instead, it is doing other things - sniffing one place or another, moving from the back of the chamber to the front, etc. By using principles of shaping, however, you can help move the behavior toward the target of bar pressing for water. Your first step should be to reinforce the ongoing behavior of the rat that is the closest "approximation" to bar pressing.

Look at what your rat is doing. Since you have already accomplished magazine training, you may accidentally have already increased the frequency of some behavior (such as approach to or perseverating at the water delivery area). Watch your rat for a minute or two and decide which of its current behaviors is the one closest to bar pressing. Is it removing its head from the water reservoir and turning toward the bar? Then wait for that behavior to occur and immediately deliver a reinforcer. The timing of your delivery will be very important. A reinforcer affects most strongly the behavior that occurs IMMEDIATELY before its delivery. If you delay your delivery, you will actually reinforce the behavior that followed the one you were hoping to select! You need to be ready to deliver the reinforcer quickly. Thus you will need to learn to predict what behavior the rat is actually emitting and will emit next. Is it turning toward the bar? A reinforcer will increase the rat's tendency to do that again. Or is it actually about to return to the back wall of the chamber? A reinforcer will increase that tendency and the process of shaping bar pressing may be set back. Learning to accurately reinforce the behavior that is closest to the desired performance makes a good coach successful. Shaping your rat to press the bar is an example of good coaching. It may take some practice, but we encourage you to keep improving your shaping skills until your rat subjects in CyberRat show you are a successful coach. To become this successful coach, there are four principles you should keep in mind that govern successful shaping.

The first principle emphasizes the importance of timing. We have already warned that you should avoid delaying the delivery of a reinforcer, because this will accidentally reinforce the behavior that occurs just before the reinforcer is delivered.

The second principle is that by becoming familiar with the behavior of your rat in CyberRat, you should become skilled at predicting an ongoing SEQUENCE of behaviors. One way to summarize this principle is that a good shaper knows the individual being shaped extremely well and is ready to reinforce the behavior that is, in fact, closest to the performance desired.

The third principle addresses the sequence of "approximations" that you choose in the shaping process. Shaping is often referred to as the "Method of Successive Approximations." You will succeed as a trainer when you apply the two principles above across a sequence of steps that moves the behavior ever closer to the desired target performance. Planning such a sequence is the key to successful shaping when the ultimate performance is complex (e.g., teaching a child to tie his or her shoes). Teaching CyberRat to press the bar for water actually involves teaching the rat to carry out a series of steps involving approach to the bar (often from the location of the water reservoir), rising up and placing its paws on the bar, and then pressing down on the bar. As you carry out shaping it is useful to think that you are training this sequence.

A fourth principle: Move your behavioral criterion for reinforcer delivery at the "right pace" -- not too slow, and not too fast. This principle addresses when you should shift from one step to the next in your sequence of approximations. Suppose you have shaped the behavior of consistently turning toward the bar. The rat turns, and you deliver your "click" with the water reinforcer. The rat then turns toward and collects its drop of water. But you want to move to the "next step" in your series of successive approximations: "turning toward AND APPROACHING the bar." When should you change your "rule" and wait until the rat not only turns toward the bar but actually moves toward it or sniffs it? The next section provides some added information to help you make that decision.

APPENDIX TOPIC: "Not too slow"/ "Not too fast"

"Not too slow" -- As you continue reinforcing a specific behavior during the shaping process, that behavior has a tendency to become more "fixed" - i.e., less variable from one occurrence to the next. The rat develops a consistent and often simplified way of meeting this criterion. As the behavior becomes less variable, the variations that are still "closer to the target behavior" become less frequent. As the behavior becomes stereotyped, it becomes more difficult to shape the next step. Conclusion: don't wait too long. This part of principle four might be called "continued reinforcement restricts behavior."

Another part of the fourth principle and another reason for not waiting too long is that sometimes a little "extinction" can help generate new behavior that meets the requirements you have set for that "next step" in your sequence. Just as continued reinforcement reduces variability, ceasing to reinforce a particular behavior increases variability. Once you stop reinforcing a particular behavior (e.g., turning toward the bar), you can bet that the variety of behaviors the rat emits will increase. After you stop reinforcing its turn toward the bar, for example, instead of alternating between "turning toward the bar" and "turning back toward the water dipper" as it has been doing, you may find the rat turning in circles, turning and rearing, turning and moving forward -- a variety of patterns of behavior. Some of these behavioral sequences, in fact, may be different from any you have seen before. Remember that the procedure of stopping the reinforcement of a behavior is called "extinction." This part of principle four might be called "extinction at first promotes new behavior."

"Not too fast" -- But, just as you should not move your criterion too slowly (in order to avoid stereotypy and to gain the benefits of a little "extinction"), you should also continue to reinforce behavior at each step of your sequence long enough that the effects of your reinforcers produce an orderly change in the frequency of behavior you are now reinforcing. That is, don't change criteria too quickly. So the last part of principle four is that you should see the change in behavior become somewhat predictable before you shift to the next step.

We have now completed the description of principles that we think will help you shape your rats in CyberRat to press the bar for water. The next section, however, lists three additional factors that might be added to these principles of response shaping when you take your newly acquired talents and apply them to other types of behavioral change -- behavior that can be prompted, that can be brought under discriminative control, and that can be developed to persist even when reinforcement is intermittent. Prompting is not a part of CyberRat, but discriminative control and intermittent reinforcement are.

APPENDIX TOPIC: Other Factors Involved in Creating New Behavior: Prompting

Sometimes the target behavior you wish to reinforce during shaping can actually be evoked or guided. We have described how you can carefully reinforce successive approximations to a target behavior in order to create a new behavior that had not previously been observed (e.g., a rat pressing a bar). Let's think now about how we might speed up such a process. Could we speed it up, for example, if we found a way to directly evoke a bar press? If we could get a rat to press the bar for "another reason," and then we gave a drop of water for each of these evoked bar presses, perhaps the drops of water would act as reinforcers to strengthen the bar pressing. Then we would not need to carefully select each successive approximation. That might work, though we would then need to remove this "other reason" if we hoped to have the bar pressing continue for the water reinforcers alone. Would this approach work? Well, the answer is "Maybe."

Consider an example in order to see the issues. Suppose, for example, you smeared a little peanut butter on the bar and then put the rat in the chamber. Suppose further that your rat quickly approached and nibbled at this peanut butter and, in the process of nibbling, pressed the bar. Suppose further that you followed this bar press with a drop of water, which the rat drank. Will that drop of water reinforce bar pressing (i.e., increase the frequency of bar pressing)?

Perhaps yes, perhaps no. For this "short cut" to work to strengthen bar pressing, several things would need to be true. First, the nibbling would need to produce bar presses several times so that the drop of water could have its reinforcing effect. Second, the variety of bar presses would need to fit well with the kind of pressing you have as a target behavior. Since what the rat is doing is nibbling, the bar presses that occur will probably be due to movements made by its head rather than its paws. Will reinforcing head-presses also reinforce other kinds of bar presses? Third, the water reinforcer would have to "compete" effectively with the peanut butter reinforcer for control of the ongoing behavior around the bar (nibbling, pressing, etc.). That is, if the peanut butter is a very strong reinforcer and the water a very weak reinforcer, we would probably not see much control by the water reinforcer. Fourth, the behavior of bar pressing would need to continue even after the peanut butter had been consumed. When the peanut butter is present, it undoubtedly has a strong smell and changes how the bar looks.

As you remember from your prior reading and as we will emphasize below, a reinforcer changes behavior "in a specific context." When the peanut butter is gone (consumed), that may change the situation enough that the prior reinforcers don't apply to behavior emitted in the presence of a bar-without-peanut butter. So, can you reinforce bar presses that are evoked by nibbling at peanut butter? For these four reasons, our answer is "maybe." Try it on a real animal (CyberRat doesn't "do" peanut butter or anything similar), but don't be surprised if you need to fall back to shaping successive approximations. One of the more effective prompting techniques often seen in home dog training is the use of a reinforcer (such as a bit of food) as a prompt, say to evoke a "sit up" by holding the reinforcer above the dog's head, thus evoking a sniffing and reaching or lifting of the body to approximate a "sit up." Eventually that reinforcer needs to be faded as the prompt or the animal won't situp without first seeing it.

Of course, when we want a human to perform a new behavior, we often model the performance and then ask them to imitate what we did. Even small children have a strong tendency to imitate what they see. We can say, for example, "Here's how a good golf swing looks (demonstrate a good swing)... Now you do it." And, the listener probably will produce some kind of golf swing. This is a little like putting peanut butter on the bar in order to get the rat to "give it a try." Your modeling has produced a swing.

For this to be a good first step in creating a new golfer, however, the same four kinds of worries list above apply here as well. For your coaching to effectively improve their golf swing, your modeling and requests must continue to promote attempted swings, these swings must be like the kinds of swings the person would make "on their own" when you are no longer providing the model, the outcomes of their swings will need to provide strong enough reinforcers for the control to switch from your encouragement to "playing golf," and the behavioral change produced by these swings will need to persist even after you are no longer present. The "lesson of the peanut butter" might be, then, that just prompting or evoking a behavior will not be enough. For a prompted behavior to continue without the prompt, you will need to work to gradually shift control from the prompt to the stimuli and outcomes you hope will eventually be the context and the reasons for the behavior.

APPENDIX TOPIC: Other Factors Involved in Creating New Behavior: Discrimination

Remember that reinforcement occurs in a context and its effect is limited to this context (discriminative control, attention). We emphasized above that reinforcement changes behavior "in a situation" or "in a context." If you change the situation (context) you may no longer see the effects of prior reinforcers that were given in the prior situation (context). Said another way, the effects of reinforcement generalize only to some situations. So, if you have shaped a skill in one context (say the safe environment of a psychologist's office), you may well have to gradually change the situation, reinforcing the behavior in each of these different situations, until finally the behavior occurs in the situation you are hoping to influence (e.g., giving the speech in front of a live audience).

Now, this control by context (which we will now call "discrimination" as you learned to do in the text) is often an important part of what you are teaching. You may be teaching the individual, for example, to choose the correct one of two alternatives such as in answering a True/False or a multiple choice question. Or you may be training them to respond quickly when a danger light is illuminated. In many cases, the new performance we want to train involves both learning what to do (behavior) and when to do it (discrimination training). Is there a way that the principles of "successive approximations" may be used for such discrimination training? Our answer is a firm "Yes." We have given a couple examples below.

The "Easy to Hard effect" provides one example of this teaching of generalization. Suppose you wanted to teach an individual to perform a difficult listening task, for example, to understand native speakers conversing in a language with which the individual was not familiar. Rather than starting by exposing the individual to this difficult task, you might start by having the individual speakers speak slowly and use only standard "textbook" expressions. Once the individual is able to understand this simplified conversation, you could gradually shift toward understanding normal conversation of native speakers. The basic rule is, training a difficult discrimination is easier if the individual can already carry out a simpler, related discrimination. Now, we include this example of the benefit of slowing down a language when a person is learning to listen, even though we expect the example will generate some controversy. Language instructors often emphasize the dangers of "slowing down" speech (slowing speech down changes it) and they encourage their students to listen to normal-paced speech. Well, they are correct. We would agree that students must push themselves to achieve understanding of the normal-paced speech, but we suggest that early in language learning it is important to have exposure to easier material, with a gradual shift toward the normal pace and complex phrases of native speech.

Transfer from control by one (easier) dimension to a different (less controlling) dimension by fading out the first dimension. The example above involved transfer from an easy discrimination to a more difficult discrimination, but both discriminations were drawn from the same "dimension," that is, the same kind of judgment (e. g. recognition of slower to recognition of faster speech). Sometimes it is useful to establish an easy discrimination and then gradually shift to a harder discrimination even though the easy and the hard discrimination each require control by different "dimensions" of the situation. An example might be teaching a person who is developmentally delayed to distinguish between different coins, say nickels, dimes, and quarters. Suppose you wish to teach this individual to pick the higher-value coin from each possible pair that is presented. We would suggest that you start by highlighting the correct choice, perhaps by placing that coin on a white piece of paper while the incorrect coin is placed on a dark paper. It will be easy to teach the person to choose the coin on the light paper over the one on the dark paper. As the training continues, however, you could gradually change the brightness of the papers to make the brightness discrimination more difficult. At some point, control may transfer from the brightness of the paper on which the coin is placed to properties of the coins themselves.

A variation on this transfer procedure is possible if you place the coins on surfaces that can be back-illuminated with various intensities of light. Start, as above, with the light turned on behind the surface with the correct coin and dark behind the surface with the incorrect coin. After the individual learns to choose the brighter side, instead of varying the brightness of the surfaces to produce a shift to properties of the coins, instead, delay turning on the light for a brief period after the coins are shown. Gradually increase this delay. At first the person will wait until the light is turned on and will then make their choice. But, this period before the light is illuminated is, as you can imagine, a perfect time for the individual to study the coins, seeking to anticipate which one will be the lighted coin (correct). As they come under control of the properties of the coins in this prediction, they may well make an early selection during the pre-light period. By reinforcing correct "early" choices you will strengthen control by properties of the coins, thereby shifting from lighting to coin characteristics.

APPENDIX TOPIC
Other Factors Involved in Creating New Behavior: Intermittent Reinforcement

Intermittent reinforcement increases persistence (perseveration) by reinforcing different variations of behavior. There is still one more dimension to explore for the idea of "successive approximations" in training new behavior. Often we want to train persistent behavior. "When the going gets tough, the tough get going," says the football coach. That is, "if at first you don't succeed, then try, try again." How do we encourage such persistence? In all our examples above, we encouraged the use of consistent, immediate reinforcer delivery in order to efficiently train a new behavior. But, consistently reinforced behavior (often called "continuously reinforced behavior," or CRF) is not very persistent. When the Soda Machine stops giving you a soda when you put in your money, you don't usually try, try again. That would be "throwing good money after bad." But, life, the bard might say, is not a soda machine. If she doesn't say "Yes" to your first request to go out for coffee, perhaps she WILL say yes to your second, third... Persistence might even pay off.

Now, how might we train such persistence? Once a behavior is established (reinforced) you can shift from CRF to using intermittent reinforcement. But it is important to do so in relatively small successive stages. Persistence is demonstrated when a behavior continues to be emitted while it is only occasionally reinforced. You can generate persistent behavior by gradually shifting from CRF to intermittent reinforcement.

We will give an example that hopefully will improve your skills as a parent when and if those skills are needed. We suppose you too have been annoyed by standing in a grocery store check-out line behind a child who keeps pleading with their parent to get them some candy. This asking for candy can be very persistent, can't it? The parent at first says "No" and continues to say "No" but the child escalates by making the requests louder and more insistent until the parent finally gives in.

Think about it. Intermittent reinforcement. And, notice the escalation? Doesn't that fit with our description above of the effect of extinction on behavior? When you stop reinforcing a behavior that has previously been reinforced (here, requesting candy) you will often find that the forcefulness and variety of the behavior increases (technically called response induction). On a football field that may be the desired result. In a grocery store check out line, we suffer from it. But, in either case, the persistence and forcefulness follow from a history of intermittent reinforcement. Now, experiment with this phenomenon using CyberRat, but be aware that the rules of "not-to-fast" apply to leaning out how dense the delivery of reinforcers are just as they apply to shaping new behaviors through successive approximations. That is, there is a successive approximation to more and more intermittent schedules to be considered, lest you extinguish the behavior before it becomes more persistent. (This extinction during schedule transitions is sometimes referred to as "ratio strain").