

# 1 Lessons From Animal Learning for the Study of Cognitive Development

(1991)

*The epigenesis of  
Mind: Essays on  
Biology & Cognition*

C. R. Gallistel  
*University of California, Los Angeles*

Ann L. Brown  
*University of California, Berkeley*

Susan Carey  
*Massachusetts Institute of Technology*

Rochel Gelman  
*University of California, Los Angeles*

Frank C. Keil  
*Cornell University*

Why open a book on cognitive development with a chapter on animal learning? It is not to argue that an account of cognitive development should use animal models of learning. Rather, it is to take advantage of developments in this area that provide insight into the question of how to characterize cognitive development. We argue that any account of cognitive development is incomplete if it attempts to explain both the what and how of learning solely in terms of general processes—be they associations, prototype abstraction, hypothesis testing, induction, analogical reasoning, assimilation, generalization, or differentiation. We argue that there are reasons for also postulating domain-specific determinants of the nature of cognitive structures. These determinants guide learning, creating structures with their own rules of organization. These ideas about cognitive development have been in the air for sometime; what we offer here is a fleshing out of them. We are emboldened

to do so in part because of what we have learned from considering cases of learning in species other than the human. The message of this chapter is that much learning in animals is best thought of as the product of behavioral mechanisms with elaborated internal structure that have evolved to guide learning of species-relevant features of the environment.

Specification of the conditions under which learning occurs has been a central goal of experimental psychology since its inception as a natural science. A major research agenda from this tradition has been to delineate general laws of learning that hold across species and problems (domains). It was this fundamental commitment to general laws that led major learning theorists (e.g., Thorndike, Hull, Skinner, and Watson) to the behavior of man with laws derived from animal models. Indeed, it has been argued that the main agenda of experimental animal learning was not to explain animal learning but to understand human learning (Schwartz, 1981). From this position, animal-specific solutions are not only irrelevant, they are a positive nuisance in the march towards understanding the general laws of learning that explain human behavior. Animal models, if they are to elucidate human behavior, must assume commonality.

A quite separate intellectual tradition determined the research agenda of the ethologists. Biologists by training or tradition, they sought to uncover the diversity of animal solutions to varying (or identical) environmental pressures. They recognized not only diversity of organisms but also diversity of mechanisms within organisms.

Far from seeking general laws of learning, the main tenet of this creed is innately directed, or preferential, learning (Tinbergen, 1951), which reflects the selective adaptive pressure of a particular environmental niche. This perspective emphasizes the structural integrity of an adaptive complex. Learning, like other adaptive solutions, is embedded within the structure of coherent contexts, and the particular properties of a learning mechanism make sense only with reference to this system. It is the business of ethologists to map the fit of the animal's behavior to its environmental niche and to describe the diversity of such adaptive solutions.

In the latter part of the 1960s, a movement arose within psychology that was heralded as a rapprochement of these two views. Seligman and Hager (1972) called it a "reunion of thought between the psychology of learning and behavioral biology after almost a half-century of separate development" (p. 1). This movement introduced the notion of constraint to deal with the fact that some associations in classical and instrumental conditioning are much easier to form than others. Even here, in the natural territory of general process learning theory, the notion of selectivity was seen as operating pervasively. Animals can form some associations easily, others with some difficulty, and others not at all. Seligman, among others, argued for biological

preparedness for certain forms of learning, stating that "preparedness of organisms reflects the selective pressure that their species has faced" (Seligman & Hager, 1972, p. 464).

As we make clear in this chapter, we prefer the phrase "privileged relationship" to "constraint" or "boundary." A problem with the phrase "constraints on learning" was noted in both of the seminal books that heralded this rapprochement (Hinde & Hinde, 1973; Seligman & Hager, 1972). The phrase suggests that there is a general process mechanism that is somehow too powerful or generative. Uneasiness with this potential interpretation goes back to the beginning of this literature. The terms "boundaries" and "constraints on learning" were called "unfortunate" even in the original Hinde and Hinde (1973) book with that title. The term *predisposition* was added as a subtitle because of a fear that the term *constraint* would lead to the inference that the position could rest "comfortably in the framework of existing learning theories." At the Hinde and Hinde symposium, there was considerable discussion of whether it would still "prove profitable to think in terms of general laws hedged about by constraints, or whether some quite new formulation would seem more profitable" (Hinde & Hinde, 1973, p. 470).

This dissatisfaction with the connotations of the word "constraints" appears repeatedly in the discussion of general process theories, biological constraints positions, and ethological approaches in the issue of *The Behavioral and Brain Sciences* devoted to this question (1981, Vol. 4). Garcia (1981) argued that the notion of biological boundaries suggests that the constraints are "subservient to the general process tradition" (pp. 143-144), and Malone (1981) wrote "in stressing exceptions to the laws of GPLT [General Process Learning Theory], this approach . . . almost gives the laws themselves an added legitimacy" (pp. 151-152). According to Petrinovich (1981), the notion of ". . . constraints involves a tacit acceptance of a uniformity view that holds that there is a general associative mechanism that typifies learning processes, and that this general law is abridged, or amended, by special factors" (pp. 153-154). Furthermore, Shettleworth (1981) asserted that "Biological boundaries to learning have been threatening to become mere adjustments to parameter values of GPLT, while theorists ignore their implicit message that learning is best understood as part of an animal's adaptation to its natural environment" (pp. 159-160).

In this book, we advance the thesis that domain-specific learning mechanisms make learning possible. The focus is on how these mechanisms make learning possible, not on how they constrain a general mechanism from achieving faulty generalizations; hence, our view that it is better to talk about *privileged as opposed to constrained acquisitions*.

We begin with a discussion of biologically specified privileged pairings in what is commonly called associative learning, using as our prime example

conditioned food aversion in rats under laboratory conditions. We then consider examples in which a specific representation-forming learning mechanism makes it possible to acquire crucial information at specific stages in animal development. The argument that complex representation-forming systems are a foundation of development is elaborated by considering the acquisition of bird song, prey recognition in vervet monkeys, and spatial learning in rats.

## ASSOCIATIVE LEARNING

### General Process Learning Theory

The theoretical framework for the traditional study of animal learning shares much in common with the theory of association as developed by the British empiricists. The fundamental assumption is that knowledge about the world and how to respond to it is based on the capacity to form associations in a lawful way. The empiricists formulated two laws of association: (a) the more exposures to a particular association opportunity, the stronger the association (the law of frequency), and (b) the more proximate in time and space the occurrence of the component members of the association, the more likely the association will be formed (the law of contiguity).

There is nothing in the laws of association about the nature of stimuli or responses that can be associated with each other, nor anything that varies with the ontogenetic stage of the individual. The original theory treated all effective sensory inputs and all observable responses as equipotential vis-à-vis the associative process. Learning about language and number should be traceable to the same fundamental laws of association that explain the rat's learning to avoid poisonous food or the pigeon's ability to learn the temporal parameters of a schedule of reinforcement.

Modern animal learning theory and some views of cognitive development retain the general process assumption, though not necessarily the equipotentiality assumption. The position is that all learning is based on the capacity to form associations; there are general laws of learning that apply equally to all domains of stimuli, responses, and reinforcers; the more frequent the pairings between the elements to be associated, the stronger the associative strength; and the more proximate the members of an association pair, the more likely the learning.

The general process theory of learning has been developed around the study of classical and instrumental conditioning and extended in a variety of ways to deal with phenomena of generalization, discrimination, and inhibition and to take into account the central role that reinforcement can often play in whether an association is formed. It has also been significantly revised

in response to experiments on classical conditioning showing that it is not the pairings per se between a conditioned stimulus (CS) and unconditioned stimulus (UCS) but predictiveness that determines whether or not an association forms. If the CS and UCS are independently distributed in time, so that the occurrence of the CS does not predict the occurrence of the UCS, then no conditioning occurs despite numerous fortuitous pairings of the CS and US (Rescorla, 1968). Also, if a new CS predicts what is already predicted by an old CS, no association forms between the new CS and the US, no matter how frequently or closely they are paired (Kamin, 1969). Finally, if another CS accounts for more of the variance in US occurrence than the CS in question, then associations between the CS in question and the US do not form despite frequent pairing (Wagner, Logan, Haberlandt, & Price, 1968).

The idea that associations build as a function of frequency is neutral with respect to the size of increments that occur as a function of trials. The Rescorla-Wagner law (Rescorla & Wagner, 1972), sometimes called the "delta" rule in the literature on learning in parallel distributed systems, gives the function for the change in the strength of the  $n$ th association ( $\Delta V_n$ ) as a function of various parameters and of the sum of the pretrial strengths of the other associations:

$$\Delta V_i = K(\lambda - \Sigma V_i)$$

In this equation,  $K$ , the associability parameter, reflects the readiness with which the CS and UCS can be associated. The introduction of this associability parameter was motivated by the evidence against the equipotentiality assumption. Lambda ( $\lambda$ ) is the asymptotic associative strength parameter. The closer  $\Sigma V_i$  is to  $\lambda$ , the less the increment in associative strength ( $\Delta V$ ) on a conditioning trial.  $\Sigma V_i$  is the sum of the strengths of the associative bonds linking the CSs present on a trial to the US. Roughly speaking, the equation captures the idea that the more strongly something (the US) is expected on a given trial, the less its occurrence alters one's expectations.

The parameters  $K$  and  $\lambda$  are CS-UCS pair-specific and are determined empirically. There is no theory of why  $K$  is large for some CS-US pairs and not for others. The Rescorla-Wagner law also says nothing about the time frame over which the associative process operates. General process theorists make pair-specific assumptions about the degree of temporal contiguity that there must be between the CS and UCS for learning to occur. A similar conclusion holds for stimulus sampling models of human concept learning; they either build in factors like selection or use parameter variations to capture stimulus-specific effects (see Atkinson, Bower, & Crothers, 1966, for a review of these models).

The idea that there are structural determinants of learning that go beyond those set by the nature of an organism's sensory and motor endowments has

influenced the study of classical and instrumental conditioning in animals. Where it was once almost universally assumed that research would yield up general laws of learning—ones that would apply to all species, at all ages, in all environments, and to all stimuli and response pairings—there is now much debate. Even those who defend the idea that there are general laws of learning have formulated their theories in such a way as to take account of a pervasive fact: Animals behave as if they treat many stimulus–reward, stimulus–response and/or stimulus–stimulus pairings as privileged. What these privileged relationships are can vary from species to species. These privileged relationships cannot always be explained by an animal's conditioning history; instead, they often make sense only when considered in terms of the evolutionary-functional role they play in the animal's life. The privileged relationships betray the presence of a system that leads the animal to notice and respond selectively to one kind of stimulus as opposed to another under certain conditions. In this sense, they serve as the fundamental evidence for an argument that there are high-level (or nonsensorimotor) determinants of what is learned and how.

### Some Examples of Privilege

*Avoidance Learning.* The equipotentiality assumption that prevailed in learning theory from Pavlov to Skinner asserted that any perceptible stimulus could be associated with any other perceptible stimulus or with any response in the animal's repertoire with equal ease. There are now many and diverse demonstrations that equipotentiality does not obtain. Whereas pigeons readily learn to peck a key in order to obtain food, they have difficulty learning to peck a key to avoid shock (Hineline & Rachlin, 1969). In contrast, they readily learn to flap their wings to avoid shock.

The effect of the reinforcer or US on which responses may be learned can be explained if one assumes that pigeon learning is expressed through response systems tailored to serve specific biological functions in specific contexts, so that the animal does not choose its responses independently of the context. When a bird perceives the threat of shock to its feet, it flies or attempts to fly; it does not peck. The same complex structures that preordain the appropriateness or inappropriateness of certain responses in certain contexts preordain what is likely to predict what. An experiment by Lolordo, Jacobs, and Foree (1982) supports this conjecture, as does Bolles' (1970) account of avoidance learning. Lolordo et al. presented pigeons with a redundant CS made up of a light and tone. Reasoning that pigeons treat sound stimuli as danger signals and visual stimuli as food-related stimuli, they proposed that control of a pecking-for-food operant would be selectively established to the visual component of the CS and that control of avoidance

behavior would be selectively conditioned to the sound stimulus. This was exactly what happened; different components of the CS controlled the target behaviors during transfer tests. Such cross-over effects provide compelling evidence against the view that all stimuli, responses, and reinforcers are equipotential in the creation and maintenance of learned behavior. To account for the privileged treatment accorded certain pairings, it often helps to adopt a biological perspective. Consider Bolles' theory of avoidance learning—one that places such learning in the context of the animal's problems in nature.

The general process theory account of avoidance learning has the animal gradually learning to perform a targeted behavior in response to the presentation of a CS. If he does, he avoids shock. For the CS to become effective, it obviously has to be paired with the UCS. Over trials, the conditioned response in question occurs when the CS occurs and thus the animal avoids shock. Bolles pointed out that this account of avoidance learning makes little sense if we consider an animal in the state of nature rather than a laboratory. It will not do to have the animal go through a series of encounters with a predator in order to learn to escape or avoid it. Instead, it seems plausible to take the position that animals are endowed with species-specific defense reactions, such as fleeing, freezing, and fighting—reactions that occur immediately to predators, threats, and the sudden appearance of innocuous stimuli. As Bolles (1970) put it, "no real-life predator is going to present cues just before it attacks. No owl hoots or whistles 5 seconds before pouncing on a mouse. And no owl terminates his hoots or whistles just as the mouse gets away so as to reinforce the avoidance response. Nor will the owl give the mouse enough trials for the necessary learning to occur" (pp. 32–33).

The implications of Bolles' position for the laboratory study of avoidance learning are clear: One should expect differences in the ease with which different classes of stimuli come to control different classes of responses. The case of the differential rate at which a pigeon learns to fly as opposed to peck a key to avoid shock fits this view. So do the experiments done by Bolles that show, for example, that rats learn rapidly to run to avoid shock but learn very slowly, if at all, to rear to avoid shock. Running is a component of a flight response. Rearing is not; it is an exploratory behavior.

*Appetitive Instrumental Conditioning.* It is not just avoidance learning that presents cases of privileged relationships between certain stimuli and certain responses. Shettleworth (1975) reported a similar result in her study of food-rewarded instrumental conditioning of six different golden hamster behaviors. The hamsters learned quickly to dig, scabble, or rear for a food reward. They learned slowly or not at all when required to wash their face, scratch, or scent mark. At first, the rationale for these differences is not

apparent; however, if one considers the way hamsters respond in a natural environment, the differences make sense. Shettleworth observed the way hamsters behave under several conditions including the anticipated presentation of food when they are hungry. The behaviors of digging, scrabbling, and rearing become prevalent under these circumstances, whereas those of face washing, scratching, and scent marking tend to disappear. The former are all exploratory behaviors, the latter are not. The motivational state of hunger recruits those behaviors that might lead the hamster to find food and suppresses other components of the hamsters repertoire (Gallistel, 1980). Given that the exploratory behaviors can aid the animal's search for food, it is no longer surprising that hamsters learn rapidly to perform these responses and have more difficulty learning to perform those that are not hunger related. The preferred relationships betray the presence of a motivational mechanism that organizes the animal's behavior and its learning.

It might be argued that the differential ability to reinforce the two classes of behavior reflects a conditioning history of the hamster (Mackintosh, 1974), rather than a biologically organized motivational-response system that causes the animal to be selective. The problem with this position is that laboratory hamsters are unlikely to have had the opportunity to perform the exploratory behaviors—especially digging—in the presence of food. In fact, when reared in laboratory cages, hamsters spend a great deal of time grooming and little time digging in the wire floor. Despite the presence of food, they apparently do not relate grooming to it. Otherwise, Shettleworth should have had an easy time reinforcing grooming with the presentation of food.

*Learned Food Aversions.* The work by Garcia and his colleagues continues the theme that neural mechanisms that have evolved to serve a particular function determine what can be associated with what. Garcia and Koelling (1966) had rats drink water of a distinct flavor from a licking tube. Via a lickometer connected to the spout, their drinking also activated a flashing light and a noise. When the animals readily drank this funny tasting water from the spout that triggered noise and light, they were punished for doing so either by electric shock delivered through the spout at various latencies after the onset of drinking or by the administration of a poison that made them sick after a latency of about 20 minutes. Both punishments reduced their consumption. In associative terms, one would say that some or all of the stimuli (funny taste, flashing light, noise) had become associated with the punishing US. However, the poisoning was effective even at long delays between drinking and the onset of punishment, whereas the shock was effective only when the punishment was in close temporal contiguity with the drinking. The rats were then tested with water that had the same taste but came from a spout that did not produce the noise and light and with a spout



that produced the noise and light but gave water without the distinctive taste. The animals that had been punished by poisoning avoided the flavored water but drank readily from the spout that caused light flashes and noise, while the animals that had been punished by shock through the spout readily drank the flavored water but avoided the bright, noisy spout.

The Garcia and Koelling experiment reveals the presence of two distinct learning processes operating in accord with different principles about what goes with what and over what temporal intervals. One process treats the noise and flashes that result from touching tongue to spout as much more likely predictors of the shock than is the distinctive flavor of the water, but it recognizes this predictive relation only when the temporal link is tight. The other process treats the distinctive flavor of the water as a much more likely predictor of the subsequently experienced illness than are the noise and light that accompany drinking, and it picks out this predictive relation between a particular taste and illness even when the two experiences are separated by intervals measured in hours rather than seconds (Revusky & Garcia, 1970).

The differences between these associative mechanisms are treated as parameter differences in the values of the constants in the Rescorla-Wagner model. We have trouble with this, not because it is wrong to say that there are parameter differences; there obviously are, but the question remains "Why?". What is it about the relation between taste and illness that makes them readily associable over long delays? What is it about taste and illness that allows for what is essentially a violation of the law of temporal contiguity?

The principles by which these learning mechanisms operate reflect fundamental facts about the world in which these animals have evolved. Both taste and toxicity derive from the chemical composition of substances and tend, therefore, to covary, the more so in that many of the poisons an omnivore is apt to ingest have evolved in forage plants and prey animals as defenses, making it of adaptive value to the poison source to have a marked and distinctive taste. On the other hand, nothing in the structure of the everyday world makes the toxicity of a food covary with lights and sounds that accompany ingestion but do not emanate from the substance itself. Also, toxins generally take some while to make their effects felt. A learning mechanism sensitive only to tight temporal contiguity would never associate the distinctive tastes of most toxins with their delayed but potentially fatal effects. On the other hand, a mechanism for detecting predictive relations among external variables like lights, sounds, and shocks cannot readily be tuned to pick out predictive relations operating with temporal lags of hours, because in any one hour there are too many potentially predictive experiences of changes in these kinds of variables. It must require close temporal linkage to solve the problem of what predicts what. Thus, the characteristics that are optimal for one mechanism are precisely the wrong characteristics for the

other mechanism. The domain-specific tuning of the distinct learning mechanisms revealed by the Garcia experiments makes adaptive learning possible.

Rozin and Kalat (1971) and Rozin and Schull (1988) point out that the special tuning of the learning mechanism that mediates bait shyness is but one aspect of a behavioral complex that shows many special adaptations that promote healthy food selection in the rat. For example, rats have an innate bias for things that are sweet, which tend to be high energy compounds, and an innate aversion for things that are bitter, which tend to be alkaloid poisons. These biases alone will not guarantee success. Not all bitter things are toxic, and not all sweet things are nutritious. What helps the rat sort out his food environment are three feeding habits. The first is to eat familiar foods and avoid novel foods. The second is, when eating a novel food for the first time, to eat only small amounts. If the novel food contains a toxic substance, the rat is not likely to eat enough of the food to die but it eats enough to experience the ill effects. The third is to wait a long time between meals involving novel foods, allowing time for illness to develop without producing confusion over the source of illness.

Recently, Galef and his associates (1987; Galef, McQuoid, & Whiskin, 1990) have shown another aspect of this adaptive specialization. Part of what contributes to a rat's sense of familiarity with a food is smelling it on the breath of fellow rats. In Galef's experiments, an "observer" rat is exposed to a "demonstrator" rat that has recently eaten one of two novel diets. Seven or 8 days later, the observer ingests both novel diets and subsequently becomes ill. When next tested with the two diets, the observer avoids the novel diet that it did not smell on the breath of another rat a week earlier. Smelling a substance on the breath of a conspecific tends to prevent that substance's being perceived as the source of illness, even when the experience of ingestion and illness occurs long after the experience with the demonstrator rat. It does not matter whether the "demonstrator" was or was not ill during the demonstration. This aspect of the poison-avoidance mechanism operates on the implicit principle that other rats know what they are eating. Noting that a demonstrator ate something gives that something the seal of safety no matter what the condition of the demonstrator rat.

The storing of memories of the food odors it has detected on the breath of conspecifics for use—often much later—in deducing which foods have caused it to become ill is analogous to the young song bird's storing memories of the songs of conspecifics for later use in developing its own song (see Marler, this volume). It is another example of Marler's "instinct to learn," that is, of the innate foundations that determine what is stored and how that information is employed to direct subsequent learning and behavior.

In summary, a complex of special adaptations, including adaptations of the learning mechanism, make it possible for a rat to learn to avoid foods that are bad for it. Specialized learning mechanisms with implicit commitments to the

nature of the world they must adapt to also make it possible for young birds to adjust the parameters of the circuitry that extracts from sounds the angular position of their sources, and they make it possible for migratory birds to learn as nestlings the facts about the night sky they will need to know to maintain their orientation during the migratory flight they will first make months later as young adults.

## NON-ASSOCIATIVE LEARNING

The preceding section gave examples of relationships whose learning is privileged. These examples were developed within the conceptual framework provided by the associative analysis of learning. Gallistel (1990) has questioned whether this conceptual framework is the correct one for understanding the nature of the learning that occurs in these and other examples of classical and instrumental conditioning. He gives a simple but powerful analytic model of classical conditioning in which the notion of an association plays no role. He shows that this model, which has no free parameters and is computationally simpler than the Rescorla-Wagner model, gives a more adequate account of the salient findings from the last 20 years of classical conditioning experiments. In the Gallistel model, classical conditioning experiments do not study a general learning process; rather, they study a learning mechanism that is specifically dedicated to computing a representation of the temporal dependence of variation in one variable on variation in one or more predictor variables—multivariate time series analysis.

Whether or not the associative framework is the appropriate framework in which to view the examples in the previous section, there are many examples of animal learning that do not fit this conceptual framework. These examples make it very clear that for most kinds of animal learning to occur there must be a specific learning mechanism that makes that particular kind of learning possible. All of these learning mechanisms have been shaped over evolutionary time so that the structure of the learning mechanism—what information it processes and how—reflects the structure of the problem that has shaped its evolution.

### How the Barn Owl Learns Where Sounds Come From

Localizing the source of sounds accurately is important to the adult barn owl, who relies on the sounds rodents make to orient its predatory attacks. It can pluck a scurrying mouse off the forest floor in total darkness, guided only by the rustle the mouse makes as it moves through the litter. Sound localization is based on frequency-specific binaural differences in the intensities and arrival times (phases) of sound waves. The owl is born with circuitry that

processes these differences in order to extract the angular direction of the sound source, but the innately given parameters of this specialized circuitry are necessarily only approximately correct, because the interaural differences in intensity and phase for a sound of a given frequency from a given direction depend on the idiosyncrasies of the animal's head shape. They vary from individual to individual, and, of course, they change as the head of an individual grows larger during maturation.

The barn owl has a specialized learning mechanism that recalibrates its sound localization circuitry during development (Knudsen, 1983; Knudsen & Knudsen, 1990). Like the other specialized mechanisms we have been considering, this mechanism is dedicated to constructing a particular kind of representation—in this case, a representation of the angular direction of a stimulus source. And again, some principles about the relevant aspects of the world are implicit in the innately given structure and functioning of the learning mechanism. In this case, one may recognize two such principles:

1. The angular position of the source for auditory and visual inputs originating from one object is a unique location in one and the same space. A mouse does not reflect light from one location in space while simultaneously causing rustles from a different location. Thus, the location assigned to the origin of the auditory input from a source should be the same as the location assigned to the origin of the visual input from that source.

2. Vision is inerrant; the location assigned to a source by the visual system can be used to calibrate the circuitry that computes a location from auditory input.

The eyelids and the auditory canals of baby barn owls open when they are about two weeks old. If the correspondence between the location assigned to a source by the auditory system and the location assigned to the same source by the visual system is altered at this age, either by putting a plug in one ear canal or by fixing displacing prisms in front of the eyes, the birds adjust the direction in which they orient their head in response to a sound by as much as 20°. Thus, the orientation of the head in response to a sound from a given direction matches the orientation of the head in response to a flash from that direction. If the alteration has been produced by an ear plug, this adjustment means that the resulting orientations to sound have been corrected so that the head is now oriented toward the source, but if the alteration has been produced by prisms, this adjustment means that the resulting orientations to sound are now in error by the same amount as the visual orientations. The visual orientations are off by nearly the amount by which the prisms deviate the visual input. The barn owl visuo-motor systems cannot learn to correct for the effects of deviating prisms (unlike the human visuo-motor systems),