14

The Principle of Species Independent Learning Phenomena

Onur Güntürkün and Daniel Durstewitz

Learning theory provides an exceptionally successful framework to predict the behavior of humans and other animals in learning situations. Unfortunately research has concentrated mostly on the small number of conditions in which learning theory fails without devoting similar efforts to explain why predictions are successful in the vast majority of the remaining cases. Here, we will outline that the success of learning research is possibly related to the fact that learning implies modifications of synapses contingent on the temporal order of spikes. Spiking-time-dependent synaptic modifications reflect the temporal asymmetry of the physical world, a fundamental constraint common to all living beings which might have shaped the molecular architecture from very early on. Resource limitations might have led to additional constraints, producing deviations from general learning theory. As one example, we will discuss the apparent inability of pigeons to associate tones with visual cues. This case illustrates nicely that the specific constraints in pigeons are very likely related to an absence of a common synaptic territory of the auditory and the visual system in higher association areas. Hence, additional constraints might involve local architectural specializations of neural systems without corroborating the general framework of learning theory.

Animals come in a bewildering variety of sizes and shapes. Some have pentachromatic color vision, others are blind. Some hop on two feet, others have several dozen appendages, while again others attach to the substrate for their entire life. But regardless of their different appearances and life styles, all animals are able to learn. Obviously, the ability to extract some lessons from the past has such a tremendous evolutionary advantage that we would expect learning to be a central feature no behaving creature should miss. What we should not easily expect, however, is the fact that the principles of learning seem to be remarkably similar for all living beings. Thus, the variety of forms contrasts with the uniformity of learning principles. During the 40s to 60s, the heydays of Behaviorism, this uniformity was envisaged as a result of a general process learning theory. This theory was seen as covering all aspects of experience dependent behavioral changes, regardless of species, sensory system, or genetic background. The degree of self-confidence of the Behaviorist approach is best appreciated when reading the famous quote of Teitelbaum (Teitelbaum, 1966, pp. 566-567):

> "We arbitrarily choose almost any act from the animal's repertoire and reinforce it with food, or whatever else the animal will work to obtain. Although typically we teach a rat to press a bar or a pigeon to peck a key to obtain a pellet of food, we readily train either to dance around the cage if we choose. We usually use a light to signal the delivery of a pellet but we can use a tone or a buzzer or any other stimulus the animal can detect.... In effect, in any operant situation, the stimulus, the response, and the reinforcement are completely arbitrary and interchangeable. No one of them bears any biologically built-in fixed connection to the other."

It does not take a great prophet to foresee that this extreme position could not go on without serious challenges. And indeed, this belief was called into question during the 70s (Seligman, 1970; Shettleworth, 1972). Criticism of the traditional approach was stimulated by observations of learning phenomena that were contrary to widely espoused general principles of association learning and that appeared to illustrate biological constraints (Domjan, 1983). Discussion of the issue suggested that a revolution in the study of learning was in the making that could concentrate on specialized mechanisms which had evolved to facilitate learning in biologically important situations. In particular, these new studies challenged the assumption that learning occurs the same way regardless of the cues, responses, and reinforcers involved in a learning situation. This equipotentiality principle was viewed as a fundamental assumption of general learning process theory, and the inadequacy of the principle was considered good reason to dramatically alter traditional learning research approaches (Domjan, 1983). However, despite revolutionary proclamations, the ensuing years merely saw the equipotentiality principle to falter without major changes to general learning theory (Shettleworth, 1998).

Rereading the papers with clashes between learning psychologists and ethologists, a glaring omission is obvious: There is a complete neglect of the neural substrate that produces learning (but see Macphail (1993), for a scholarly exception). Ultimately, answers to the question why general learning theory often succeeds but sometimes fails have to be found there. This approach will be the theme of our endeavor. We will first focus on the ubiquity of cases where classic learning theory accurately predicts detailed results of learning experiments. We will argue that the tremendous importance of a high temporal contiguity, which is so characteristic for virtually all learning studies, might go hand in hand with common basic membrane mechanisms of the nervous tissue. Discussing the constraints literature with their emphasis on the inadequacy of the equipotentiality principle, we will discuss the possibility that a lack of common synaptic space between two sensory systems might cause difficulties to associate stimuli perceived by these sensory pathways.

The Power of General Process Learning Theory

Thorndike (1911) was the first to formulate the general process learning theory. Studying cats, dogs, chickens, monkeys, and a fish species he observed only quantitative learning differences. According to his view, the behavior of the animals could be understood as the result of a common set of mechanisms governed by the Laws of Exercise and Effect. In the following account we will follow his distinction between the variable rate of learning on the one hand, and the more invariant laws of learning on the other.

Habituation learning, defined as a decreased response to repeated stimulation, is a good point in case. Although it is a simple non-associative form of learning, it is characterized by nine invariants that typically occur in all preparations and which were first listed by Geer, O'Donohue, & Schorman, (1986) and Thompson & Spencer (1966). Groves & Thompson (1970) formalized this overview further by assuming that all invariants could result from the interaction of two hypothetical processes. One is a decremental (habituation) and one an incremental (sensitization) process that develop independently in the brain. The details of the interaction can, e. g., account for an initial hump in the habituation curve, which, according to Groves & Thompson (1970) can be observed in animals as diverse as newborn humans, spinalized cats, rats, mudpuppies, and quails. Scrutinizing the literature further, this list can easily be expanded to include protozoa (Wood, 1973), anemones (Logan, 1975), leeches (Burrell & Sahley, 1998), and slugs (Burrell & Sahley, 2001). A detailed review of the synaptic mechanisms involved in learning in Aplysia as well as Hermissenda shows that indeed the predictions made by the dual process theory hold up to astonishing detail for the habituation process, while some discrepancies are visible for sensitization (Macphail, 1993). Thus, habituation learning seems to involve very similar mechanisms in all animals studied so far.

The general picture for associative learning is comparable. For example the growth of associative strength during classical conditioning is mostly described using the notation of Rescorla & Wagner (1972). This equation turns out not only to predict accurately the progress of learning in all vertebrates but also in insects like bees (Menzel & Giurfa, 2001). Similar cases of identical mechanisms in vertebrates and bees are seen for second-order conditioning, within-compound association, blocking, and conditional discrimination (for review see Bitterman, (2000). Why are so many aspects of learning theory valid for so many different conditions and species? The answer might be that most learning phenomena depend on basic physical constraints, information-theoretic constraints like maximization of information content, as well as on a small set of physiological key mechanisms reflecting these basic constraints. One of these key aspects could be temporal contiguity (Abrams & Kandel, 1988).

The conventional wisdom of learning studies has been that conditioning requires the associated stimuli to occur immediately after each other and in the correct temporal order. This is shown in a study by Dickinson, Wat, & Griffiths (1992) who examined the effects of delayed reinforcement on learning to press a key. Each time a rat pressed a lever, a food pellet was delivered after a certain delay. Delaying reinforcement for only a few seconds dropped off responding dramatically. Further increasing the delays reduced performance even more until no learning was observable at delays of 64 seconds. Since delayed feedbacks drastically reduce learning efficiency, animal trainers often use immediate secondary reinforcers like verbal prompts (”very good”, ”excellent”) to bridge the delay until primary reinforcers are provided. Given this importance of high temporal contiguity, we will scrutinize in more detail the cellular properties behind this phenomenon to discuss the possibility that it may represent an invariant of nervous systems.

Cellular Basis of Association Learning

In most of the above mentioned situations the learning process is temporally asymmetric. This makes a lot of sense from a functional point of view: If some stimulus S1 temporally precedes another stimulus S2, S1 might be taken as a predictor of S2 but not vice versa. Interestingly, a temporal asymmetry in the cellular mechanisms underlying learning was postulated already by the Canadian psychologist Hebb (1949) in his famous statement (p. 62):

> “When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased. The most obvious and I believe much the most probable suggestion concerning the way in which one cell could become more capable of firing another is that synaptic knobs develop and increase the area of contact between the afferent axon and efferent soma.”

Thus, it is assumed that a change in synaptic strengths between neurons underlies information storage and altered behavior, a view widely adopted by neural network theories which revealed the staggering computational power of such a simple principle (Hertz, Krogh, & Palmer, 1991). However, in many learning theories within cognitive neuroscience and by neural network formal learning rules, Hebb's statement was taken to indicate that coincident pre- and postsynaptic activity would lead to a strengthening of the synaptic weights between neurons (leading to the slogan "Neurons that fire together, wire together"). That is, the temporal asymmetry pointed out by Hebb was largely neglected (Sejnowski, 1999). We will return to the importance of temporal asymmetry on the behavioral level in another example at the end of this chapter.

Since the initial findings by Bliss & Lomo (1973) that post-synaptic potentials can be enduringly increased by high-frequency, tetanic afferent stimulation, physiologists have assumed that long-lasting synaptic changes like long-term-potentiation (LTP) and long-term-depression (LTD) might in fact constitute the cellular correlates of learning as postulated by Hebb (see Dudai, 1989, Chap. 6, for an overview). There are also a number of studies that tried to establish a link between the cellular phenomenon of LTP/LTD and behavioral learning (most recently using gene knockout techniques and transgenic mice) (see Dudai, 1989, Chap. 6, and Squire & Kandel, 1999, Chap. 6, for an overview). Most convincingly this endeavor succeeded in simple organisms like the sea-slug Aplysia where the physiological and molecular cascades underlying habituation and sensitization have been worked out in impressing detail (Kandel, 2001; Squire & Kandel, 1999).

Here we will focus on a number of recent findings that studied the expression of LTP and LTD based on the precise relationship between pre- and postsynaptic spike times. These findings might shed some light on the cellular bases of temporal contiguity. According to Hebb's original statement the synaptic weight from a cell A to a cell B should be strengthened if A tends to fire B - that is the synaptic weight change is supposed to capture a causal relationship between the firing of two cells (not just their mere broad temporal correlation), and hence the precise temporal relationship between spiking times should be relevant. This has indeed been confirmed in various species and various systems like the frog Xenopus retinotectal system (Zhang, Tao, Holt, Harris, & Poo, 1998), the electrosensory lobe of electric fish (Bell, Han, Sugawara, & Grant, 1997), the mammalian hippocampus (Bi & Poo, 1998; Debanne, Gähwiler, & Thompson, 1994, 1998; Levy, Brassel, & Moore, 1983; Magee & Johnston, 1997) and neocortex (Feldman, 2000; Markram, Lübke, Frotscher, & Sakmann, 1997). The strengthening (LTP) or weakening (LTD) of synaptic weights of excitatory synapses onto Xenopus tectum neurons, and onto mammalian hippocampal and neocortical pyramidal cells, in fact obeys Hebb's law as illustrated in Fig. 1a,b: If the presynaptic spike precedes the postsynaptic spike, i.e. if the presynaptic cell could have been involved in firing the postsynaptic neuron, synaptic strength is increased. If, however, the postsynaptic cell fires first, before a presynaptic spike reaches the cell, LTD occurs. Thus, remarkably, a switch of the temporal order of pre- and postsynaptic spiking by just a few milliseconds could lead from strong LTP to strong LTD.

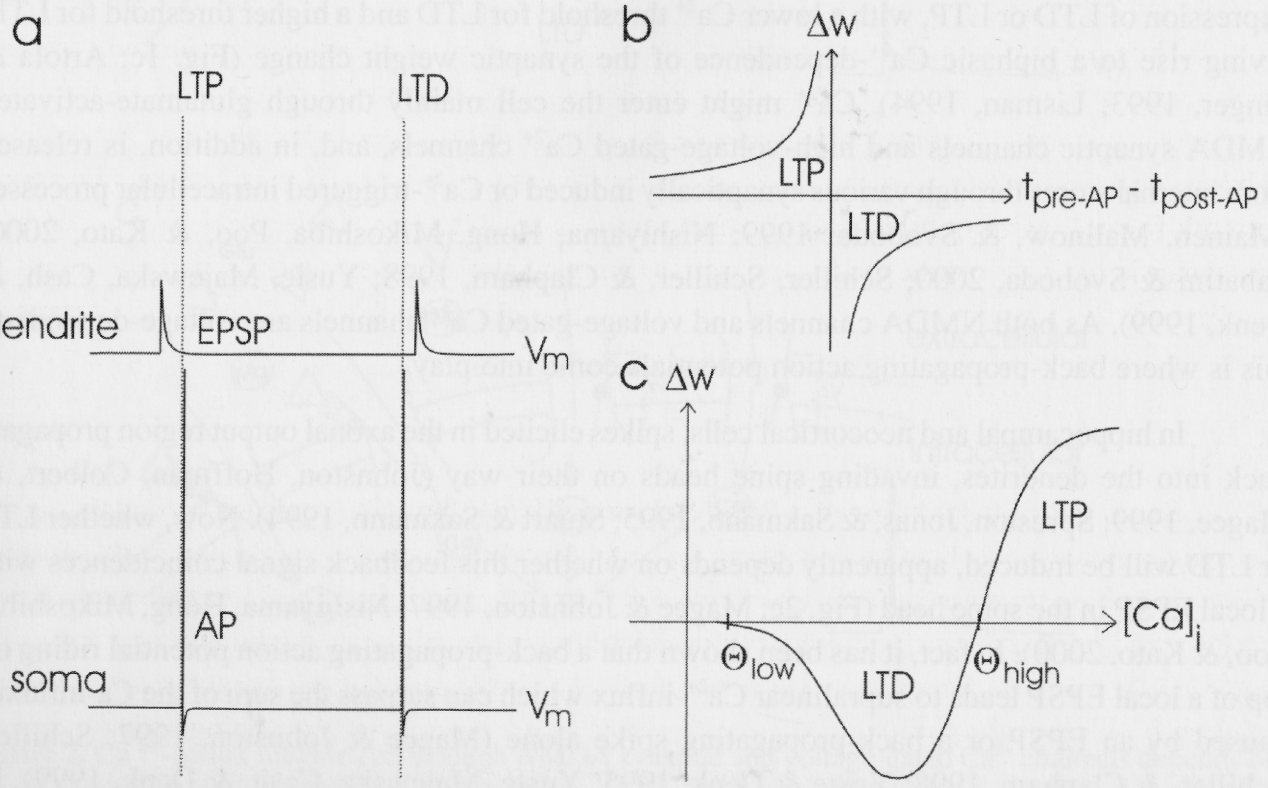

Figure 1. Dependence of the sign of the change in synaptic efficacy on spike-timing and intracellular Ca^{2+}-concentration. **(a)** If an EPSP precedes the postsynaptic spike the synapse is strengthened, i.e. LTP is induced, while the synapse is weakened (LTD) if the order is reversed. **(b)** Dependence of amplitude and sign of synaptic weight change (Dw) on the temporal difference of pre- and postsynaptic spiking (illustration based on results as in Bi & Poo (1998). **(c)** Dependence of amplitude and sign of synaptic weight change (Dw) on intracellular Ca^{2+} concentration with a lower threshold (Q_{low}) for LTD and a higher threshold (Q_{high}) for LTP (illustration based on results as in Artola & Singer, 1993).

Typically the time window for the induction of LTP and LTD is about +/-40 msec wide, but in general the extent of the temporal windows for LTP and LTD as well as their integrals (determining the relative balance between LTP and LTD) can differ (see (Abbott & Nelson, 2000). There are interesting exceptions from the temporally asymmetric rule as depicted in Fig. 1a: For example, in the electrosensory lobe of electric fish the temporal order of LTP and LTD is reversed (Bell, Han, Sugawara, & Grant, 1997). In neocortical pyramidal cells, if GABAergic inputs were elicited while the postsynaptic cell was spiking, LTD of the inhibitory connection was induced, whereas LTP was induced when the GABAergic input occurred more than 100 msec after the last spike in the pyramidal cell (Holmgren & Zilberter, 2001). In the former case, however, the targeted neuron is inhibitory (unlike cortical pyramidal cells), while in the latter case the synaptic input to the postsynaptic neuron is inhibitory rather than excitatory and thus should prevent the cell from firing. Hence, rather than being true exceptions, these kind of mirror-imaged learning rules might reflect different causal relationships in the activity of the two connected neurons.

How is the temporal relation between pre- and postsynaptic spikes detected by the postsynaptic cell and translated into differential regulation of synaptic weights? There may be two important ingredients, back-propagating spikes and local Ca^{2+} influx into the spine-head through NMDA synaptic and voltage-gated Ca^{2+} channels (in addition, Ca^{2+}-release from internal stores is likely to play a role). Ca^{2+} influx triggers the intracellular events that eventually lead the expression of LTD or LTP, with a lower Ca^{2+} threshold for LTD and a higher threshold for LTP, giving rise to a biphasic Ca^{2+}-dependence of the synaptic weight change (Fig. 1c; Artola & Singer, 1993; Lisman, 1994). Ca^{2+} might enter the cell mainly through glutamate-activated NMDA synaptic channels and high-voltage-gated Ca^{2+} channels, and, in addition, is released from internal stores through various synaptically induced or Ca^{2+}-triggered intracellular processes (Mainen, Malinow, & Svoboda, 1999; Nishiyama, Hong, Mikoshiba, Poo, & Kato, 2000; Sabatini & Svoboda, 2000; Schiller, Schiller, & Clapham, 1998; Yuste, Majewska, Cash, & Denk, 1999). As both NMDA channels and voltage-gated Ca^{2+} channels are voltage-dependent, this is where back-propagating action potentials come into play.

In hippocampal and neocortical cells, spikes elicited in the axonal output region propagate back into the dendrites, invading spine heads on their way (Johnston, Hoffman, Colbert, & Magee, 1999; Spruston, Jonas, & Sakmann, 1995; Stuart & Sakmann, 1994). Now, whether LTP or LTD will be induced, apparently depends on whether this feedback signal coincidences with a local EPSP in the spine head (Fig. 2c; Magee & Johnston, 1997; Nishiyama, Hong, Mikoshiba, Poo, & Kato, 2000). In fact, it has been shown that a back-propagating action potential riding on top of a local EPSP leads to supralinear Ca^{2+}-influx which can surpass the sum of the Ca-influxes caused by an EPSP or a back-propagating spike alone (Magee & Johnston, 1997; Schiller, Schiller, & Clapham, 1998; Yuste & Denk, 1995; Yuste, Majewska, Cash, & Denk, 1999). In particular, an EPSP followed by a postsynaptic action potential evokes a larger Ca^{2+} influx than with reversed temporal order (Koester & Sakmann, 1998). This is probably due to the fact that NMDA-receptors activated by the presynaptic release of glutamate exhibit a nonlinear voltage-dependence: They are blocked by Mg^{2+} at low voltages, while at higher postsynaptic membrane potentials the Mg^{2+} block increasingly removes, such that coincident presynaptic transmitter release plus postsynaptic depolarization lead to the most effective opening of this channel, letting

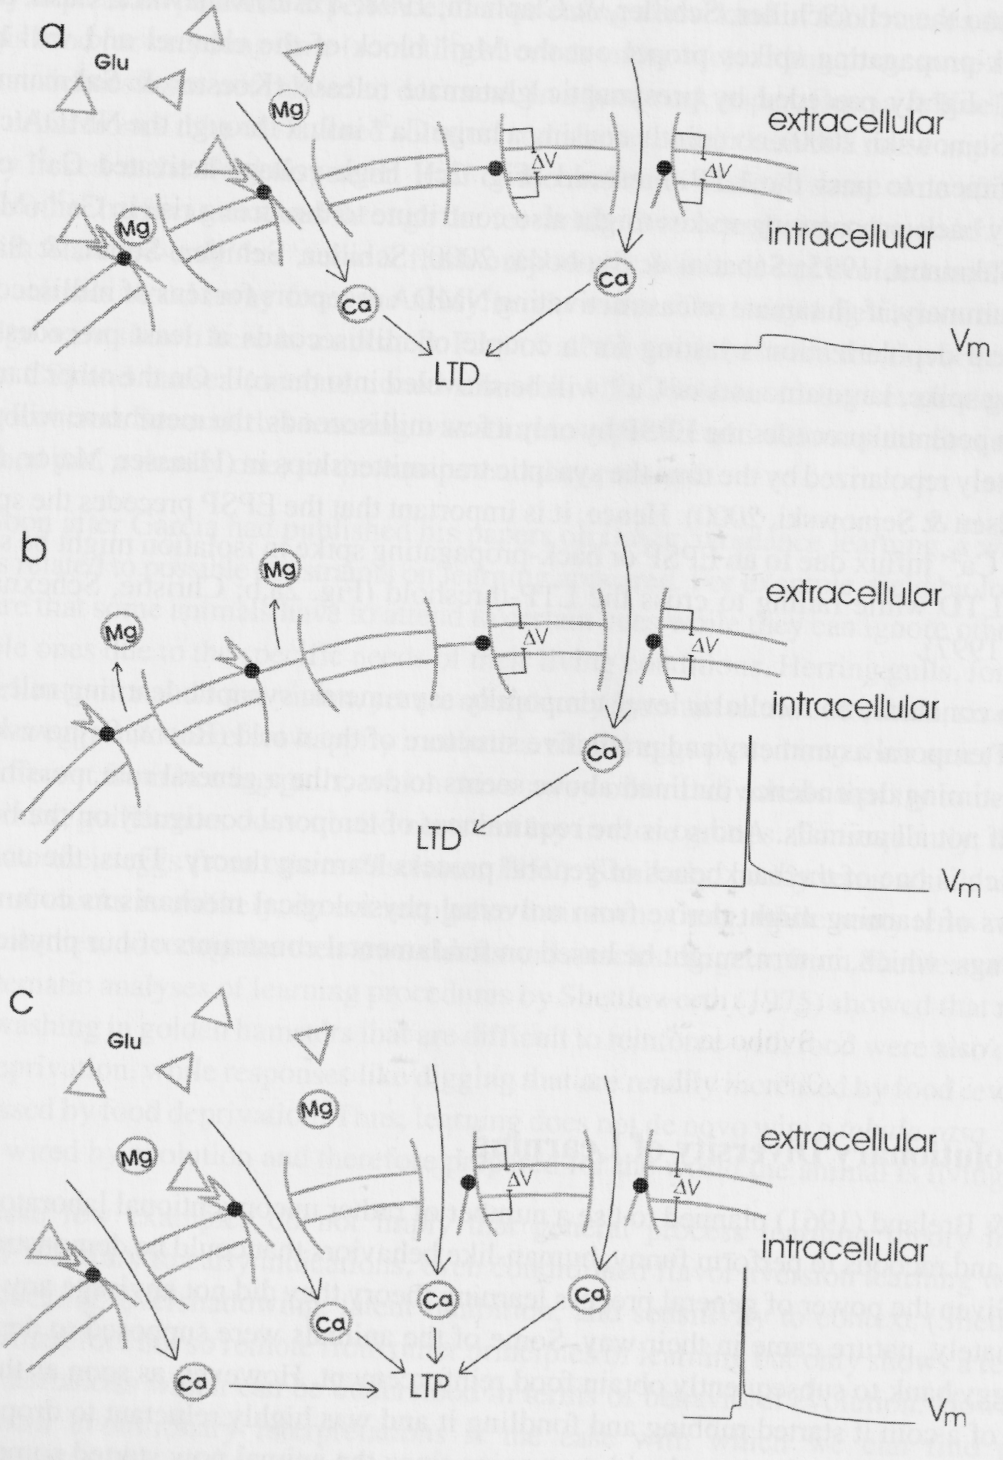

Figure 2. Ca $^{2+}$ -influx into the cell through NMDA synaptic and voltage-gated Ca^{2+} channels depends on nearly coincident timing of presynaptic glutamate release (triangles) and the back-propagating spike (see voltage-traces in lower right corners). **(a)** Glutamate release alone induces an EPSP not sufficient to open all NMDA channels but letting enough Ca^{2+} into the cell to promote LTD. **(b)** A back-propagating spike alone opens only a few Ca^{2+} channels, sufficient to induce LTD but not LTP. **(c)** If a back-propagating spike occurs briefly after glutamate release, the Mg^{2+} block of NMDA channels will be efficiently removed while the receptors are still activated, plus a number of Ca^{2+} channels will open such that large amounts of Ca^{2+} flow into the cell causing LTP.

Ca^{2+} rush into the cell (Schiller, Schiller, & Clapham, 1998; Yuste, Majewska, Cash, & Denk, 1999). Back-propagating spikes propel out the Mg^{2+} block of the channel and will be most effective if slightly preceded by presynaptic glutamate release (Koester & Sakmann, 1998; Paulsen & Sejnowski, 2000), conjointly causing a large Ca^{2+} influx through the NMDA channels that is sufficient to pass the LTP-threshold (Fig. 2c). High-voltage-activated Ca^{2+} channels activated by back-propagating spikes might also contribute to the strong rise in Ca^{2+} (Markram, Helm, & Sakmann, 1995; Sabatini & Svoboda, 2000; Schiller, Schiller, Stuart, & Sakmann, 1997). In summary, if glutamate release activating NMDA receptors for tens of milliseconds and causing local depolarization's lasting for a couple of milliseconds at least precedes a back-propagating spike, large amounts of Ca^{2+} will be shoveled into the cell. On the other hand, if the brief action potential precedes the EPSP by only a few milliseconds, the membrane will probably be completely repolarized by the time the synaptic transmitter ships in (Häusser, Major, & Stuart, 2001; Paulsen & Sejnowski, 2000). Hence, it is important that the EPSP precedes the spike. The amount of Ca^{2+} influx due to an EPSP or back-propagating spike in isolation might be sufficient to trigger LTD while failing to cross the LTP-threshold (Fig. 2a,b; Christie, Schexnayder, & Johnston, 1997).

To conclude, on a cellular level temporally asymmetric synaptic learning rules seem to reflect the temporal asymmetry and predictive structure of the world (Rao & Sejnowski, 2001). The spike-timing dependence outlined above seems to describe a general rule, possibly shared by most, if not all animals. And so is the requirement of temporal contiguity on the behavioral level, which is one of the backbones of general process learning theory. Thus, the universality of the laws of learning might derive from universal physiological mechanisms common to all living beings, which, in turn, might be based on fundamental constraints of our physical world.

The Evolutionary Diversity of Learning

Breland & Breland (1961) planned to use a number of rather unconventional laboratory species like pigs and racoons to perform funny, human-like behaviors that could be demonstrated in the public. Given the power of general process learning theory they did not envisage any problems. Unfortunately, nature came in their way. Some of the animals were supposed to deposit coins into a piggy bank to subsequently obtain food reinforcement. However, as soon as their racoon got hold of a coin it started rubbing and fondling it and was highly reluctant to drop it into the slot. This behavior was increased with two coins since the animal now started some 'washing procedure', constantly rubbing the coins against each other. Needless to say, this sequence greatly increased the delay to reward. The hungrier the animals were, the more they showed a tendency to display such apparently counterproductive behavior. The Brelands did suggest a framework within which an animal's behavior in learning situations could be related to its overall behavioral organization, but unfortunately learning theorists largely ignored their work, perhaps because their accounts were mostly anecdotal.

The door to analyze constraints on learning was finally pushed open wide by the work of Garcia and colleagues on flavor aversion learning. This phenomenon is probably known to the reader by unpleasant personal experience: If after eating a new kind of food nausea develops, we avoid this food for very long time (sometimes for our entire life), although on some accounts we know that sickness was not related to our meal. In a landmark paper, Garcia & Koelling (1966) had rats drink from a tube of flavored solution and also exposed them to a noise and a flash each time they licked at the fluid dispenser. Half of the rats were made ill some time after drinking, while the other half was shocked through the feet while they licked at the dispenser. When later tested with 'flavored-only'- or 'noise & flash-only'-water, poisoned rats avoided the flavor, while shocked animals stood away from the noisy bright water. Two major deviations from classic knowledge were shown here in a nutshell. First of all the delay between the flavor (conditioned stimulus, CS) and nausea (unconditioned stimulus, US) could be 12 h or even more and still learning occurred. Second, this learning was stimulus-specific and thus was the first major blow (as it turned out, a deadly one) to the equipotentiality principle.

Soon after Garcia had published his papers on flavor avoidance learning, a whole flood of studies related to possible constraints on learning appeared. For example, field biologists were long aware that some animals have to attend to certain cues while they can ignore other, equally discernible ones due to the specific needs of their living conditions. Herring gulls, for example, have excellent color and form vision, yet are completely ignorant to the coloration, size and form of their own eggs, and therefore happily incubate artificial eggs placed into their nest (Tinbergen, 1953). Different from herring gulls, which drive away other individuals that try to nest in their direct vicinity, guillermots do not build nests but lay in close groups. Consequently, these birds discriminate their eggs from others (Tschanz, 1959). Similarly, noddy terns which nest where exchange of chicks is unlikely, do not recognize their own young, while sooty terns which nest in dense colonies, do recognize their own chicks and attack strangers (from Shettleworth (1972). More systematic analyses of learning procedures by Shettleworth (1975) showed that responses like face washing in golden hamsters that are difficult to reinforce with food were also decreased by food deprivation, while responses like digging that are readily increased by food reward were also increased by food deprivation. Thus, learning does not de novo wire a *tabula rasa.* The table is already wired by evolution and therefore prepared for the world the animal is living in.

These few examples do not imply that general process learning theory has to be abandoned. Contrary to early indications, even conditioned flavor aversion learning was found to show blocking, overshadowing, latent inhibition, and sensitivity to context (Shettleworth, 1993). It is therefore not so remote from other principles of learning but only shows a few, albeit important deviations which can be understood in terms of behavioral evolution. However, the problem with evolutionary interpretations is the ease with which we can find post-hoc explanations of all sorts of phenomena that do not readily fit into a general learning framework. This is due to evolutionary research being an empirical but not necessarily an experimental science. However, in some rare cases post hoc interpretations can be experimentally tested and even neurobiologically analyzed. These examples than offer insights into the neural basis of learning constraints.

Constraints on Learning in Pigeons

If pigeons are trained to depress a treadle in the presence of a compound auditory-visual stimulus either to avoid electric shock or to obtain food reward, the animals easily learn the procedure. If, however, the individual elements of the compound stimulus, red light and pure tone, are presented separately to assess their control of treadle pressing, something strange happens: in the test, the tone turns out to control responding in the shock avoidance condition, while the red light controls performance in the appetitive food reinforcement condition. Thus, the tone dominates aversion learning, while the color stimulus dominates food reward learning (Foree & LoLordo, 1973; LoLordo, Jacobs, & Foree, 1982). Based on these observations, Delius & Emmerton (1978) conducted a meanwhile classic experiment on which we will dwell in some detail.

They used pigeons to run two parallel experiments. In the one involving classical conditioning they implanted shock electrodes and placed the bird onto a foam rubber cradle in the center of a conditioning chamber (Fig. 3). In front of it was a light panel that could be illuminated with red or white light. A loudspeaker suspended from the ceiling provided either a 1 kHz or a 2 kHz tone. During conditioning with auditory stimuli, for half of the animals the high, for the other half the low pitch served as CS- or CS+, respectively. CS+ was followed by a mild electric shock (US). The heart rate change subsequent to CS+ application was the conditioned response (CR). During conditioning with visual stimuli, red and white lights were deemed to be CS+ or CS- in a balanced order. Thus, depending on the experimental group they had been allocated, the pigeons had to associate one of the auditory or one of the visual stimuli with shock.

In the experiment using instrumental conditioning the animals were confined to another conditioning chamber with two pecking keys and a food delivery device (Fig. 3). During visual discrimination the two keys could be illuminated with red or white light. For half of the animals pecking on red (irrespective if presented on the upper or the lower key), for the other half pecking on white was reinforced with access to food. During auditory discrimination both keys were illuminated with green light. If the high pitched tone was delivered, the upper key had to be pecked, while the lower key was correct during low pitch presentation. Note that in this last condition, different from the three other conditioning situations, not only a sensory signal had to be associated with a specific behavioral or physiological response, but also have the different tones to be linked to different visually perceived spatial locations. Taken together, the animals had to associate colors and tones to either shock or to food reward contingent on pecking. Depending on their group, a specific color or a specific tone signaled that either shock had to be expected or a certain key had to be pecked. Since the same colors and tones were used in both paradigms, different procedural outcomes could not be based on perceptual differences.

The results showed that the colors came to control the key pecking in the appetitive instrumental paradigm but not the heart rate in the aversive classical procedure. At the same time tones controlled cardioacceleration in the aversive classical but not key pecking in the appetitive instrumental procedure (Fig. 3). The authors interpreted this set of data as a learning constraint shaped by the specific evolutionary needs of a granivorous animal which lives under the selection pressure created by various animals of prey. In the normal environment of pigeons, sounds are unlikely predictors of food, since grains are silent companions. However, in order to discriminate grains against a scattered background, pigeons need excellent vision including color processing. Therefore, the inability to discriminate tones and the ease to discriminate colors for food reinforcement is understandable. On the other hand, aversive heart-rate conditioning is likely to be related to natural avoidance behavior, and indeed the shocked pigeons not only displayed cardioacceleration but also struggling. In the wild, avoidance behavior is mainly a response to predators. Sound is an excellent predictor of stealthy predators, while under many circumstances like during the night or in a bushy environment, color is not. Therefore, the double dissociation between aversive and appetitive conditioning with tones and colors might be understood as a compromise of an animal with just two milliliters of brain. Pigeons might be too short of neuronal space to be able to relate and store all the possible associations between the multitude of events they can perceive. They might be forced to select for associations with a higher probability of occurrence.

Up to now all this sounds like one of these post-hoc 'just so' story that have plagued evolutionary interpretations. But it is not. Specific testable extrapolations derive from this approach. First of all, bird species that feed on noisy prey should differ from pigeons. Indeed, barn owls easily associate auditory and visual cues for food reward (Wagner, 1993). Additionally, there is an extrapolation that is related to pigeons. As shown by Delius (1985), grains may be silent by themselves, but generate a good deal of noise when they are pecked at. Any peck that reaches its target yields massive acoustic feedback, partly mediated by bone-conducted

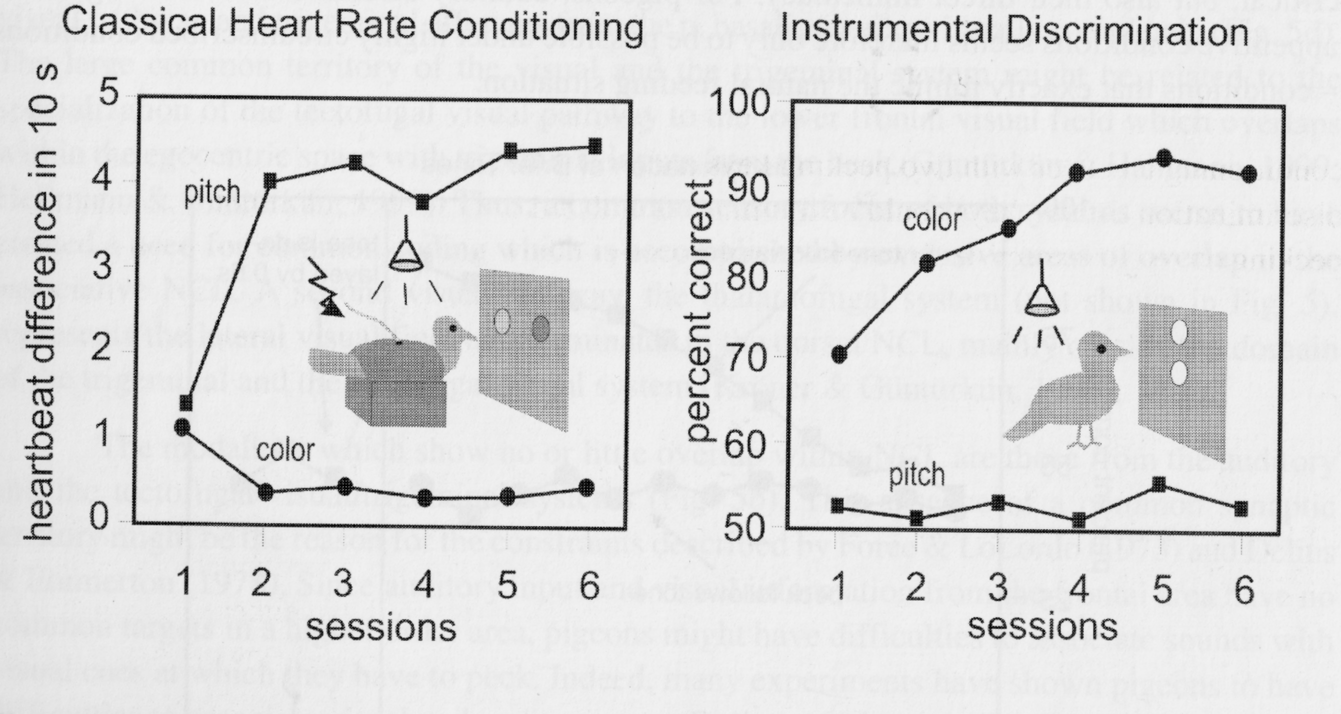

Figure 3. Schematic depiction of the experimental setup and main results of Delius & Emmerton (1978). In the upper panel the classical heart rate conditioning is shown. In the auditory test, the animal has to associate one of two tones delivered by a loudspeaker suspended above the birds with shock (lighting arrow). In the visual test, one of two color cues in front of the animal are to be associated with shock. Heartbeat differences after S+ delivery are the depending variable. As shown, pigeons are able to learn a pitch but not a color discrimination in this context. In the lower panel the instrumental discrimination is shown. In the auditory test, animals have to learn to peck a certain key after hearing a certain tone. In the visual test, they have to choose the key with the correct color to obtain food reward. As shown, under these conditions pigeons learn the color and not the pitch discrimination.

sound. The avian skull is especially suited for such bone-transmission as the jaw bones link up with the middle-ear bones that support the eardrum. It is likely that the pecking related auditory feedback is exploited as a cue for the properties of the grain the animal is pecking at. However, this auditory feedback has a highly specific temporal order: the auditory feedback is generated subsequent to the visually guided peck ('first see the target then hear the sound') and never the other way round. But in the experiment of Delius & Emmerton (1978) the tone was given first and the pigeons then had to peck one of the keys ('first hear the sound then see the target'). Thus, it is conceivable that pigeons can learn auditory discriminations during appetitive conditions only if the tones are delivered after their pecks. Delius (1985) tested this hypothesis. In this experiment, the pigeons had to discriminate between two tones of different lengths (10 and 90 ms). In a two key Skinner-box, within a trial, each pecking response on a given key immediately yielded an S+ tone, while pecks on the other key always triggered S- tones. 16 consecutive pecks on the S+ key yielded reinforcement. Responses on the S- key were never rewarded. After a reinforcement the allocation of S+ and S- to the two keys was determined according to a quasi-random sequence. Pigeons learned this procedure very well (Fig. 4). This means that simply reversing the sequence of events to the natural order (first visually guided peck, then emitted tone) enables learning, while during the procedure used by Delius & Emmerton (1978) (first delivered tone, then visually guided peck) the animals were unable to learn the association. During a control procedure, the discriminative tones triggered by the key pecks were delayed by 600 ms. Immediately learning broke down. Thus, not only the temporal order of the events was critical, but also their direct immediacy. For pigeons, auditory discrimination learning under appetitive conditions seems therefore only to be possible under highly circumscribed conditions – conditions that exactly mimic the natural feeding situation.

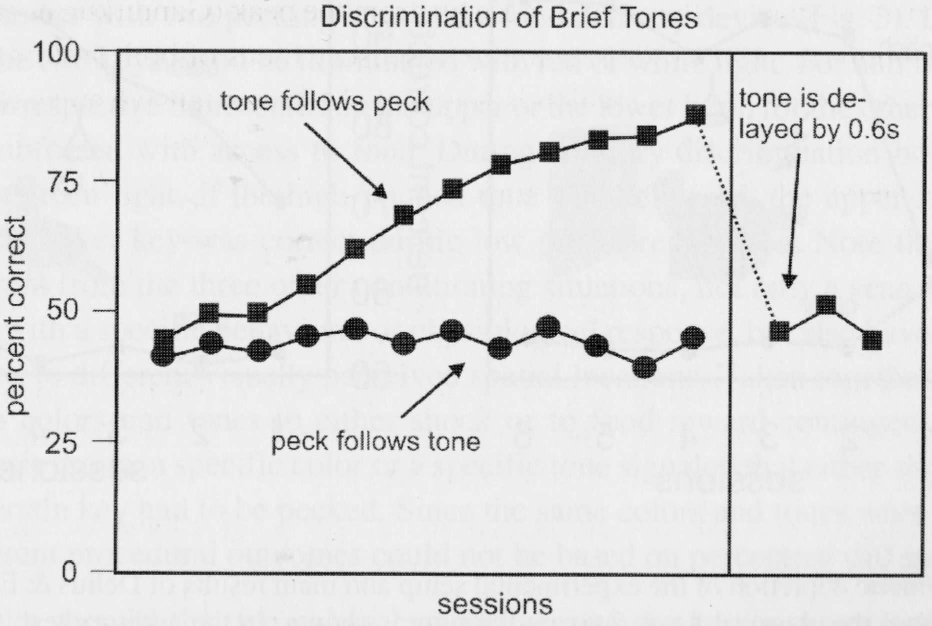

Figure 4. Auditory discrimination learning by pigeons under two different conditions. Under one condition ('peck follows tone') the animals had to choose a pecking key after hearing a certain tone. As shown, pigeons were unable to learn this association. Under the other conditon ('tone follows peck') the animals had first to peck a key to evoke a given pitch. Depending on the pitch they had to go on pecking or had to switch to another key. Pigeons learn this quite well. If the tone is delayed by 0.6 s, they fail. Data based on (Gebauer, 1980) and referred by Delius (1985).

The Neural Basis of Learning Constraints

What could be the neural basis of these constraints? The critical aspect here might be that the animals in the condition they could not accomplish had to associate two different tones with pecking responses to two different visually perceived spatial positions, i.e. had to integrate auditory and visual information for guidance of the spatially directed pecking response. Obviously, this association requires the visual and auditory modality to come together somewhere in the brain to be synaptically associated. Up to now only a single forebrain area in birds could be shown to provide extensive room for multimodal associations. This is the neostriatum caudolaterale (NCL).

The NCL is a semilunar area in the caudal forebrain of birds that has reciprocal connections to all sensory systems (Kröner & Güntürkün, 1999). The inputs are relayed via secondary sensory areas of the different modality-specific systems. Cells in the NCL respond to various modalities and seem to code mainly for the relevance of stimuli and not for their sensory features (Diekamp, Kalt, & Güntürkün, 2002; Kalt, Diekamp, & Güntürkün, 1999). NCL-lesions cause deficits in cognitive domains related to the holding of information for subsequent action sequencing and the establishment of higher-order associations (Diekamp, Kalt, Ruhm, Koch, & Güntürkün, 2000; Güntürkün, 1997).

The pattern of sensory afferents within NCL might provide clues to the functional architecture of the pigeon's nervous system. In ventral NCL visual inputs from the tectofugal visual pathway and trigeminal afferents from the n. basalis overlap virtually completely (Fig. 5d). The large common territory of the visual and the trigeminal system might be related to the specialization of the tectofugal visual pathway to the lower frontal visual field which overlaps within the egocentric space with trigeminal inputs from the beak (Güntürkün & Hahmann, 1999; Hellmann & Güntürkün, 1999). Thus, a common sensory focus of both systems seems to have created a need for common coding which is accomplished by extensive areas of overlap in the associative NCL. A second visual pathway, the thalamofugal system (not shown in Fig. 5), represents the lateral visual field and terminates in the dorsal NCL, mainly outside the domain of the trigeminal and the tectofugal visual system (Kröner & Güntürkün, 1999).

The modalities which show no or little overlap within NCL are those from the auditory and the tectofugal-visual/trigeminal systems (Fig. 5b). This absence of a common synaptic territory might be the reason for the constraints described by Foree & LoLordo (1973) and Delius & Emmerton (1978). Since auditory input and visual information from the frontal area have no common targets in a higher order area, pigeons might have difficulties to associate sounds with visual cues at which they have to peck. Indeed, many experiments have shown pigeons to have difficulties to associate visual and auditory cues (Delius, 1985).

However, the trigeminal pathway from the n. basalis also transmits auditory information, since the n. basalis not only receives trigeminal projections but also some auditory afferents from the lemniscal nuclei (Schall, Güntürkün, & Delius, 1986; Wild & Farabaugh, 1996). Hence, pigeons have two auditory systems with terminals within NCL, the main auditory pathway terminating in dorsal NCL and the small auditory system terminating together with trigeminal information in ventral NCL. The auditory neurons within n. basalis are very likely activated by the auditory feedback during pecking. This means that most acoustic input into NCL is transmitted by the main auditory system delivering input into dorsal NCL (outside the tectofugal visual domain). Only during pecking a second auditory system is switched on providing acoustic information into ventral NCL (right into the tectofugal visual representation). Therefore, associations between tectofugal visual information and auditory stimuli are only easily accomplished during pecking bouts. Thus this pattern of wiring reflects specifics of the causal structure of the pigeon's sensory world: "first you see a target to peck at – then you feel and hear the impact of your own peck". From these conditions, Schall & Delius (1991) formulated the prediction that the critical association between frontally presented visual targets and different tones which are elicited by the pecks of the animal are formed via the basalis-system. To test their hypothesis, they conditioned pigeons to discriminate between two acoustic signals generated by the animals own pecks. Then they lesioned the basalis. Indeed, the discrimination of the two sounds delivered after pecking was lost, although the animals main auditory pathway was still intact.

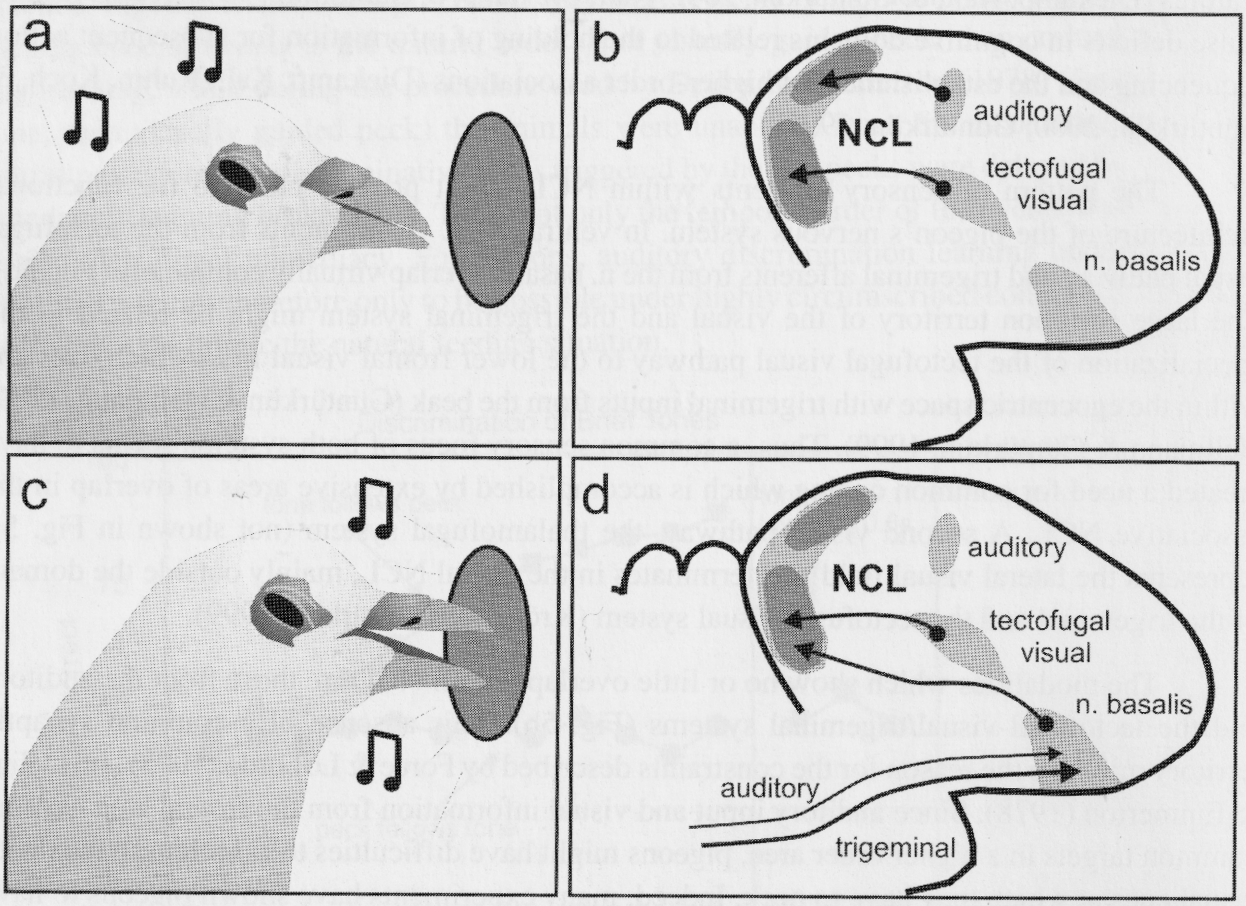

Figure 5. Schematic diagram showing the possible neural basis of learning constraints in the example outlined in the chapter. If pigeons **(a)** have to associate a tone with a visual cue (here a pecking key), **(b)** the projections of the auditory and the tectofugal visual system (were the frontal visual field is represented) have no common area of terminals in the higher-order associative forebrain structure NCL (neostriatum caudolaterale). Consequently, no learning occurs. If however, **(c)** pigeons elicit a certain tone with their own peck, **(d)** auditory processes to the nucleus basalis are evoked which overlap in NCL with the projections of the tectofugal visual system. This association is learned.

Taken together, these findings suggest that the constraints on learning in this case might be related to the absence of a common synaptic territory between the main auditory and the visual system, preventing pigeons from coupling visual cues to temporally preceding acoustic signals not processed by the basalis pathway. As long as the tones are delivered without the animal pecking a key, pigeons have a hard time to learn that they have to choose a certain key after hearing a specific tone (Figs. 5a and b). If however, the tone is elicited by the peck of the animal, auditory input is relayed via the basalis-system into ventral NCL. There it finds extensive common synaptic space with visual information coding for one or the other key (Figs. 5c and d).

Of course, our interpretation of the neural basis of learning constraints in this model rests on the implicit assumption that diverse terminal fields within NCL curtail synaptic associations. In principle, it is conceivable that indirect synaptic links exist between the visual and auditory modalities. It is also possible that sensory representations converge in structures further downstream. Our knowledge on the avian brain is presently too limited to exclude these possibilities. However, learning constraints are very likely related to local neural architectural adaptations of a given species as outlined in this example. They confine the set of all possible temporal associations to those most likely in the sensory environment of a given species, implementing a priori probabilities, without necessarily violating – in this sense – general learning theory. Such restrictions could save anatomical and energetic resources, and possibly also enable faster responding or buildup of most relevant associations.

The tremendous power of general learning theory, on the other hand, probably is a direct result of fundamental physical constraints of the world in which all species live. Neural systems had to incorporate these invariances to achieve predictability. Any animal that consistently failed to reliably predict the near future lost the game of evolution. Those who could cope well were probably those that inherited a neural system with learning principles that matched the physical reality.