
Exploring Adaptive Agency I: Theory and Methods for Simulating the Evolution of Learning

Geoffrey F. Miller
Department of Psychology
Jordan Hall, Building 420
Stanford University
Stanford, CA 94305
geoffrey@psych.stanford.edu

Peter M. Todd
Department of Psychology
Jordan Hall, Building 420
Stanford University
Stanford, CA 94305
todd@psych.stanford.edu

Abstract

Psychology construed as the scientific study of adaptive agency can include not only modelling of specific psychological adaptations in particular species, but general exploration of the adaptive processes (including evolution, learning, and computation) that build, modify, and instantiate those adaptations. Connectionist theory has concentrated on understanding the adaptive processes of learning and computation, and has assumed general-purpose learning principles as the prime constructors of psychological adaptations. But connectionism has thereby ignored the central lesson of a century of learning theory in psychology: learning mechanisms must be understood in terms of their specific adaptive functions, just like other psychological adaptations. This paper introduces the notion of psychology as the study of adaptive agency, outlines a hierarchy of adaptive processes underlying adaptive agency, and reviews the history of learning theory and the emergence of ecological and evolutionary approaches to learning. We then develop a taxonomy of adaptive functions that learning mechanisms might serve, and outline a general simulation framework for exploring those adaptive functions. Finally, we present empirical results concerning the simulated evolution of associative learning.

1 INTRODUCTION

Natural selection has constructed animals' minds and behavioral repertoires for adaptive fit to the environmental problems they must face. As the study of such minds and behavioral repertoires, a properly evolutionarily-informed psychology should focus on the notion of *adaptive agency* -- the generation of action in the world in response to challenges to fitness. This framework encompasses many approaches, including (A) analysis and modelling of complex species-typical psychological adaptations (as in human and animal experimental psychology and cognitive ethology), (B) comparison of psychological adaptations across species and consideration of their phylogenetic origins (as in comparative psychology), and (C) general exploration of the adaptive processes themselves that yield adaptive agency (e.g. by simulation methods; including those in the field of artificial life -- see Langton, 1989). Most connectionists doing psychological modelling have contributed primarily to the first of these three enterprises. This paper concentrates on the third.

One of the central applications of connectionist theory has been to develop parallel distributed processing models of psychological mechanisms in humans and, less frequently, other animals (e.g. Rumelhart & McClelland, 1986, Vol. II; Sutton & Barto, 1987; Gluck, Bower, & Hee, 1989). Although most connectionists would deny the charge that they are the 'neo-Behaviorists' of psychological modelling, many seem to adhere to one of the most central, and most problematic, assumptions of Behaviorist psychology: that learning mechanisms can be studied without regard to their specific adaptive functions for the particular species investigated. The very notion of using an 'animal model' (e.g. a rat or pigeon) to investigate human learning assumes the existence of cross-species universal learning principles which, on evolutionary grounds, we may have little reason to expect. We would argue, to the contrary, that in pursuing the first approach to the study of adaptive agency outlined above,

(1) connectionists interested in modelling portions of human or animal minds should attempt to model real domain-specific psychological adaptations (the natural building blocks of minds -- see Cosmides & Tooby, 1987) rather than other units of analysis; (2) connectionists interested in modelling domain-specific psychological adaptations should explore whether, when, and how specific learning mechanisms might be used by those adaptations, rather than simply assuming the adaptive utility of some hypothesized general-purpose learning mechanism. On the other hand, if one is interested in pursuing the third approach to the study of adaptive agency, we would suggest that (3) connectionists *can* usefully explore the general features of adaptive agency by studying how adaptive processes themselves work and interact, without reference to specific psychological adaptations in a particular species. Point 3 seems to conflict with points 1 and 2 because it suggests a less domain-specific, species-specific methodological focus. But just as one can do theoretical astrophysics without constructing specific astronomical models of particular observable celestial bodies, one can explore the dynamics of adaptive processes in general without modelling their specific structural outcomes in the minds and behavioral repertoires of particular species.

Our empirical research concentrates on point 3, because we are currently more interested in how adaptive processes interact than in the particular adaptations those processes happen to have produced on this planet at this time. This paper lays out the theoretical framework and methodological principles guiding our work in exploring adaptive agency. Elsewhere (Todd & Miller, 1990), we describe in more detail the specific methods and results obtained in simulating the evolution of associative learning as a mechanism for 'imprinting' on certain features of one's evolutionary niche; and in a forthcoming paper (Todd & Miller, in press), we extend this simulation method to understanding habituation and sensitization mechanisms as evolutionary adaptations for behaviorally tracking short-term environmental dynamics.

2 NATURAL SELECTION AND THE EVOLUTION OF SUBSIDIARY ADAPTIVE PROCESSES

Evolution as an adaptive process has itself undergone changes: "survival of the stable" probably preceded "survival of the fittest" (Dawkins, 1976). Evolution in the earth's early environment is likely to have selected for physical systems (e.g. autocatalytic sets) with relative stability in the simmering primordial soup. After stability came replication and metabolism: the ability to turn external material into copies and extensions of oneself. Simple physical systems thus evolved into replicating systems. Larger, more complex phenotypes probably evolved to protect the replicators against biochemical breakdown and facilitate their intake of materials for re-

plication and metabolism. These larger phenotypes set the stage for the evolution of behavior-generating systems that could produce innately programmed sequences of activity and movement. (In the land of the sessile, the blind, roving, pre-programmed predator may be king.) Sensory systems could then evolve to guide these behavior-generators more adaptively, based on sensitivity to particular environmental cues. (In the land of the blind, the optic-spotted paramecium may be king.) Thus, blind activity may have preceded reactivity -- behavioral adaptation to the current changing environment on a moment-by-moment basis.

Only after these first two adaptive processes (natural selection with a genotype/phenotype distinction, and behavioral reactivity) had emerged could a third evolve -- 'learning,' defined as the ability to make long(ish)-term adaptive changes in evolved behavior-generators in response to particular environmental conditions and dynamics.¹ (See Shepard, 1987a, 1989, for discussion of these adaptive processes in relation to psychology and connectionist modelling). In this hierarchy, learning emerges not as the primary adaptive force that some theorists (e.g. Behaviorists) have assumed it to be, but rather as a tertiary one, following long-term genotypic evolution and short-term environmental reactivity. Once we reconceptualize 'learning' as merely one process among several that help generate and support adaptive agency, the questions we might ask about this process begin to change as well.

Traditionally, researchers using simulation to explore adaptive agency have started by assuming learning as a primary adaptive process, and then asked how evolution might shape and be shaped by learning. For example, Hinton and Nowlan (1987) and Belew (1990) explicate the Baldwin effect in which learning "guides" evolution; Stork and Keesing (1990) and Belew, McInerney, and Schraudolph (1990) investigate how evolution and learning can combine to affect the initial structure of neural networks. But considering the hierarchy of adaptive processes spelled out above, the question *we* want to ask is, rather, given the already-powerful adaptive processes of genotypic evolution and environmental reactivity, under what conditions would the tertiary adaptive process of learning ever prove useful in terms of increasing individual fitness? Assuming *no* learning, when should learning evolve? We consider learning a mystery to be explained rather than a commonsense explanation for other phenomena.

Fortunately, connectionists are not the first group to grapple with these thorny issues. The history of learning theory in comparative psychology has considered such problems for over a century, and is slowly arriving at a

¹ By this definition, learning includes such processes as experience-guided development not commonly included in this category -- see, e.g., Singer, 1988. We discuss this issue further in section 5.

theoretical consensus that we can take, almost off the shelf, to guide our modelling and simulation efforts. Ignorance of this intellectual heritage may not preclude success, but knowledge of it should help. Connectionists can, at least, hope to avoid recapitulating the pitfalls of certain historical ways of thinking about learning.

3 A BRIEF HISTORY OF LEARNING THEORY IN (COMPARATIVE) PSYCHOLOGY

Although the earliest comparative psychologists recognized learning mechanisms as evolutionary adaptations, the rise of general process learning theory in Behaviorism effectively obliterated consideration of the adaptive functions of learning. With the accumulation of empirical results indicating biological 'constraints' on learning and theoretical arguments for the necessity of innate structure for learning, general process learning theory was gradually abandoned. Yet only recently has an alternative model, ecological learning theory, emerged, to restore emphasis on the adaptive functions of learning.

3.1 EVOLUTIONARY COMPARATIVE PSYCHOLOGY BEFORE BEHAVIORISM

Few modern learning theorists are aware of how evolutionary-minded the earliest comparative psychology was (see Boakes, 1984). Darwin was well-versed in associationist learning theory; Pavlov concerned himself with the adaptive functions of learning when to salivate (Garci y Robertson & Garcia, 1988). Kline (1898), Small (1900), and even the early Thorndike (1911) and Watson (1919) recognized the importance of animals' intrinsic organization for learning. However, these early attempts to understand learning as an evolutionary adaptation were derailed by the commitment of Romanes and other comparative psychologists to Lamarckian inheritance of learned abilities and to a phylogenetic continuity of mind from simpler to more complex organisms that was taken proximally from Spencer's (1855) *Principles of Psychology* but derived ultimately from Aristotle's 'Great Chain of Being'. Thorndike (1898) correctly rejected both of these doctrines, and "viewed the work of his predecessors, studying learning within an evolutionary framework, as unmitigated failure" (Galef, 1988, p. 55). Unfortunately, Thorndike appears to have been unaware of attempts by Baldwin, Osborn, and Morgan to propose alternative conceptions of evolution, inheritance, and learning that might have served as a better foundation for an evolutionary comparative psychology of learning. In rejecting evolutionary thinking, Behaviorists believed they no longer had to consider the specific niches and adaptive problems facing each species, and their doctrines of equipotentiality (equality of all stimuli and responses with respect to their associability) and of the cross-species universality of learning principles gained temporary hegemony.

3.2 BEHAVIORIST LEARNING THEORY

The emerging Behaviorist paradigm viewed learning as a general mechanism for crafting behavioral order out of initial neurological chaos. Morgan (1896) wrote of learning as sculpting directed activity out of initially random movements; Hull (1943) viewed infants as bundles of reflexes generating random movements that slowly get tuned by conditioning to yield adaptive behavior; Skinner (1953) viewed learning as the differential reinforcement of initially undifferentiated behavior. (Parallels to connectionist learning theory should be obvious.) Empiricism, which was initially a theory of knowledge, became first a theory of perception, and then, only with the rise of Behaviorism, a theory of behavior (Bolles, 1988). No one actually believed animals' minds started as *tabula rasae* in Aristotle's strict empiricist sense, but most Behaviorists held that a few basic reflexes and motivations, in conjunction with the ability to form conditioned stimulus-response associations, sufficed to explain all behavior. Species differences in niche and lifestyle were ignored (e.g. Thorndike, 1911). The simple, almost 'atomic' stimulus-response bond was taken as the basic unit of behavioral analysis in an attempt to emulate the theoretical style of physics and make psychology into a "purely objective experimental branch of natural science" (Watson, 1914/1967, p. 1) Behaviorism's central features (as compiled by Davey, 1989) were the principle of instrumental reinforcement (with 'reinforcement' defined circularly -- see Meehl, 1950), the principle of stimulus generalization (theoretically impossible without innate structure -- see Shepard, 1987a,b), temporal contiguity as the prime determinate of association strength (demonstrably false -- see Rescorla & Wagner, 1972), and the equipotentiality of stimuli and responses (also demonstrably false -- see the next section).

3.3 EMPIRICAL RECOGNITION OF 'BIOLOGICAL CONSTRAINTS' ON LEARNING

Some disenchantment with Behaviorist learning theory came from psycholinguistic studies of human language acquisition (e.g. Chomsky, 1957), and from ethological studies of learning in natural settings, which focused on life-cycle patterns of learning, the plasticity of natural behavior, and the dynamics of natural patterns in real environments (e.g. Lorenz, 1937; Tinbergen, 1951). However, the strongest challenges to general process learning theory came from within the Behaviorist tradition of animal learning studies itself. In the 1960's and 1970's, a growing number of studies demonstrated biological 'predispositions' and 'constraints' in learning, which challenged the Behaviorist doctrines of equipotentiality and the universality of learning principles across species. Seligman & Hager (1972) and Hinde (1973) review this literature, which includes demonstrations of conditioned food aversion that violates the principle of contiguity (e.g. Garcia & Koelling, 1966); animal 'misbehavior' and 'instinctive drift', where instinctive behaviors eventually

dominate operantly-conditioned behaviors (e.g. Breland & Breland, 1961); autoshaping, where animals perfect skills without reinforcement (e.g. Williams & Williams, 1969); spontaneous maze learning in rats (Brant & Kavanau, 1965); reward-specific association biases (Jenkins, 1984); species-specific defense reactions (Bolles, 1970; Bolles & Fanselow, 1980); animal's superiority at learning natural over artificial concepts (Hernstein, 1979); preparedness in human phobia learning (Ohman, Dimberg, Ost, 1985); and imprinting to parents (Lorenz, 1952), enemies (Curio & Vieth, 1978), and potential mates (Marler, 1984). In addition, Behaviorists were beginning to recognize that, although they explicitly defined learning in terms of experimental paradigms, they too often were defining it implicitly in terms of highly artificial lab equipment tailored to the constraints and predispositions of each 'animal model' species (Timberlake, 1989). There were 'tricks' to conditioning and shaping (i.e. in selectively avoiding species-typical behaviors) and to setting up lab situations to overcome or circumvent the mechanisms actually used in complex natural learning (Shettleworth, 1984).

Although the empirical proof of biological 'constraints' on learning was overwhelming by the mid-1970's, animal learning theorists simply didn't know what to make of these 'anomalies'. The burden of proof was always against species-specific, domain-specific adaptive learning mechanisms, and for completely general learning principles (Revusky, 1977). But Shettleworth (1984) observed that, from an evolutionary point of view, this bias was absurd, since selective association makes learning more adaptive than it otherwise would be. The burden of proof should, evolutionarily, be against domain-generality and against cross-species universality.

Consternation and confusion accumulated in the animal learning literature, threatening the hegemony of Behaviorist doctrine. Cognitive psychology (which had rejected Behaviorism's philosophy of science and methods of research even as it retained Behaviorism's doctrine of equipotentiality and anti-evolutionary bias) responded by walling itself off from animal learning theory. Where cognitive psychology could not ignore the infiltration of biological 'constraints' into animal learning research, at least it could (and did) contain the damage by maintaining human memory and learning as a separate, pristine field: "the prevailing view of human learning is [still] that it is almost wholly general-purpose in character and can be understood without reference to biological or ecological considerations" (Estes, 1984, p. 626).

3.4 THEORETICAL RECOGNITION OF THE NEED FOR 'BIOLOGICAL CONSTRAINTS'

In fairness, there was some recognition in human cognitive psychology of the need for intrinsic structure to guide learning. Following on Kant's attempt to solve Hume's problem of inferring causality from temporal

succession, Michotte (1954) explored the mind's inherent tendencies to organize perceptual experiences in terms of causal relations. Chomsky (1957) argued that intrinsic structure (e.g. a 'Language Acquisition Device') was necessary for learning grammar given the supposed paucity of linguistic input to children. Shepard (1987a,b) maintained that any cognitive system that lacks innate structure in the 'psychological spaces' in which it organizes perceptual experiences would have no basis for generalizing adaptively to new situations. In general, "non-trivial self-programming can take place only if sufficient knowledge about the world in which the system is to learn is already built into the system" (Shepard, 1989, p. 106). Further, learning principles do not emerge directly from the dynamics of replicating systems in the way the principles of natural selection do. Learning is not a self-organizing adaptive process as evolution is: "the principles that govern learning cannot themselves be learned" (Shepard, 1989, p. 106). Of course, arguments for innate constraints were never very surprising to evolutionary biologists. Emlen (1973), for instance, argued that postulating learning without genetic guidance means postulating the evolution of a mechanism that would allow an animal to arbitrarily change its phenotype without regard to its fitness consequences -- an evolutionary implausibility.

3.5 FROM BIOLOGICAL CONSTRAINTS TO ECOLOGICAL LEARNING THEORY

The accumulation of theoretical arguments and experimental evidence for biological 'constraints' in animal learning did not result in the rapid overthrow of general process learning theories, largely because there was not, until recently, an alternative theoretical framework for understanding learning. Describing selective associations and species-specific or domain-specific learning mechanisms in terms of 'predispositions' and 'constraints' in itself reveals the continuation of a general process view of learning. Constraints could be seen as anomalous biological intrusions on an otherwise normative *tabula rasa* (Shettleworth, 1984): "the implications of the term ['constraints'] ... is that animals would be smart if their genes did not constrain their general ability to learn and thereby make them selectively stupid" (Gould & Marler, 1984, p. 254). Timberlake (1989) suggests that predispositions and constraints are evolutionary outcomes to be explained, not explanations in themselves for failures of general process learning theory, and Revusky (1977) holds that viewing learning in terms of constraints is misleading in that it has fostered blindly empirical investigations of limits on general process learning theory in ignorance of relevant ethological information about animals' niches and behavior.

Research on biological constraints became a simple catalog of anomalies and puzzles without integrative or predictive power, and did not lead to development of a contemporary integrated theory of learning (Davey, 1989). As Cosmides and Tooby (1987) point out, almost

any complex adaptation can be viewed in terms of 'constraints' (e.g. a bird's wings 'constrain' its ability to swim), but most adaptations are better understood as 'enablers' (e.g. a bird's wings *do* enable it to fly). Moreover, the attempt to view constraints as biological boundary conditions on general laws (e.g. Logue, 1979) has not proved fruitful, nor can selective association be accommodated in general process learning theory by adding more parameters, e.g. a scalar 'preparedness' value for every possible stimulus-response association (see Seligman, 1970).

Biological constraints research eventually crippled general process learning theory, but left a huge theoretical gap. Only recently has ecological learning theory (Davey, 1989) risen to take its place. Ecological learning theory's strategy is to start, like all adaptationist accounts in evolutionary biology and behavioral ecology, with consideration of what adaptive functions might be served by the biological structure or process in question -- in this case, the processes of learning. In this view, learning theory must be linked to a consideration of the adaptive pressures bearing on the evolution of learning mechanisms, and of the phylogenetic resources (e.g. 'pre-adaptations') available for constructing learning mechanisms (Timberlake, 1989). Just as evolution does not 'build down' bodies or organs from more general-purpose designs, evolution would not be expected to somehow 'constrain' general-purpose learning principles. We might rather expect evolution to generate specific learning mechanisms attuned to particular ecological problems. Among the central tenets of ecological learning theory are: recognizing learning mechanisms as evolutionary adaptations (Dawkins, 1983; Shettleworth, 1983); recognizing the possible biological utility of learning (Kamil & Roitblat, 1985; Lea, 1984); considering the ecological problems facing organisms, (Plotkin & Odling-Smee, 1979); attending to relevant ethological information (Johnston, 1981a); taking an evolutionary view of reinforcement (Vaccarino & Glickman, 1989); appreciating that animals are often rather specifically adapted to their niche (Slobodkin & Rapoport, 1974); and understanding the evolution of learning in the context of already-functioning behavioral systems (Mayr, 1974).

3.6 TOWARDS UNDERSTANDING THE ADAPTIVE FUNCTIONS OF LEARNING

An evolutionary and ecological perspective on learning gives rise to very different questions, not only about the proximate mechanisms of learning (what the Behaviorists investigated almost exclusively), but about the ultimate adaptive functions of different learning mechanisms. For example, Davey (1989, p. xiv) asks "What is the biological function of learning? How does it contribute to inclusive fitness? What selection pressures bear on the evolution of learning processes? Could generalized learning processes ever be selected for? Have basic learning processes evolved separately in different species?"

Questions of adaptive functions ('why?') have logical priority over analyses of proximate mechanisms ('what?' and 'how?') for any evolved biological system (Davey, 1989); but adaptive mechanisms and adaptive functions do illuminate each other. Analyses of proximate learning mechanisms alone cannot constitute a complete psychological learning theory: "The common belief that 'learning' is an alternative to an evolutionary theory of adaptive function is a category error. Learning is a cognitive process. An adaptive function is not a cognitive process. It is a problem that is solved by a cognitive process" (Cosmides & Tooby, 1987, p. 292). The guiding question in ecological learning theory thus becomes "what kind of learning mechanisms would natural selection have produced?" (*ibid.*).

Ecological learning theory suggests that we answer this question first by considering what specific kinds of ecological problems might be solved by the evolution of learning mechanisms. Among the central problems facing terrestrial organisms are finding food, finding mates, allocating reproductive effort, caring for offspring and kin, avoiding predators and parasites, and navigating through the environment. We might expect learning mechanisms to be organized around these adaptive problems in the context of *behavioral systems* (Timberlake, 1989), *Darwinian algorithms* (Cosmides & Tooby, 1987), or *psychological adaptations* (our preferred term), each containing cognitive, motivational, emotional, volitional, learning, and memory components. For example, learning has been investigated specifically as an aid to foraging behavior (e.g. Lea, 1984; Staddon, 1980). These considerations support points (1) and (2) advocated in section 1.

4 HOW ECOLOGICAL LEARNING THEORY CAN INFORM CONNECTIONIST LEARNING THEORY

Behaviorists attempted to see how much of real human and animal behavior they could explain just by reference to general principles of learning interacting with environmental contingencies and conditioning paradigms. Connectionists all too often attempt to see how much of human mental life can be explained just by reference to general principles of learning interacting with the 'statistical regularities of the environment'. Arguments from parsimony can be dangerous. The history of learning theory in comparative psychology indicates that in both Behaviorism and Connectionism the burden of proof against evolved psychological adaptations has been misplaced. Evolutionary considerations suggest that we should reverse this traditional burden of proof and assume that most psychological adaptations will include either no learning mechanisms, or very finely-tuned learning mechanisms with quite specific functions, e.g. to promote experience-sensitive development of behavior-generators, to track changes in body shape and size, to al-

low spatio-temporal integration of certain kinds of information (i.e. 'memory'), or possibly to track certain environmental dynamics. Learning is a subsidiary rather than an autonomous adaptive process, because learning mechanisms evolve to serve particular adaptive functions defined in ecological and evolutionary terms. Ecological learning theory suggests, then, that connectionist *learning* theory per se cannot serve as the core theoretical framework for connectionist modelling of psychological adaptations. Only evolutionary psychology (Cosmides & Tooby, 1987), appropriately extended and modified, can fill that role.

4.1 ECOLOGICAL LEARNING THEORY, CONNECTIONISM, AND EVOLUTIONARY PSYCHOLOGY: A HAPPY RECONCILIATION?

Evolution, learning, and computation can all be construed as *adaptive processes* (Holland, 1975). Connectionism has concentrated almost exclusively on the adaptive processes of learning and computation, but perhaps evolution could be added in as just another process at a longer time scale. This would result in a tidy kind of 'evolutionary connectionism' where every connectionist model of a human psychological adaptation would consider three adaptive processes at different time scales (as outlined by Shepard, 1989). First, at the shortest time scale, computation would allow the connectionist system to adjust its internal representations and overt responses to the requirements of the current environment, e.g. by perceptual completion, interpretation, categorization, prediction, and inference. An activation dynamics equation governs the network's relaxation in state space to fulfill the hard and soft constraints set by the current environmental input. Second, at an intermediate time scale, learning processes would adjust connection weights and biases, perhaps by gradient descent in weight space according to some connection dynamics equation, e.g. error back-propagation. Third, at the longest time scale, a simulated process of natural selection (e.g. a genetic algorithm) could evolve network designs by performing a stochastic search through an architecture space (e.g. Miller, Todd, & Hegde, 1989; Belew, 1990; Belew, McInerney, & Schraudolph, 1990). In the limit, we could simply 'evolve' connectionist models of psychological mechanisms in an abstract 'econiche' composed of experimental results and theoretical heuristics which the models should fit. Alternatively, consideration of the adaptive tasks that certain psychological adaptations might have been designed to solve evolutionarily could help guide modeling of those adaptations by human psychologists. According to this view, if we make our computational tasks just a little more ecologically valid, and our learning processes a little more biologically plausible, we'll have a flexible, powerful paradigm that updates connectionist modelling to fit better with ecological learning theory (Davey, 1989) and with the emerging field of *evolution-*

ary psychology (Cosmides & Tooby, 1987).

4.2 TURNING THE TABLES ON LEARNING

By prematurely adopting the sort of 'evolutionary connectionism' outlined above, however, we may be missing a valuable opportunity for re-conceptualizing learning itself. Evolutionary theory has traditionally been dominated by learning theory in psychology; what would happen if we momentarily inverted this dominance relation and asked: given the already powerful adaptive process of evolution by natural selection, what could learning really add?

Johnston (1981b) analyzed the relative costs and benefits of learning from an evolutionary perspective, and concluded that learning is not always an adaptive thing to have. Fitness costs of learning may include longer infancy and adolescence, with delayed reproductive maturity (as Staddon, 1983, p. 1, observes, it is "sometimes better to be dumb and fast than intelligent and slow"), increased juvenile vulnerability during learning, increased parental investment during learning, the neural 'bookkeeping' cost associated with memory storage and the possibly greater connection complexity and density required for learning, and, perhaps most importantly, the developmental fallibility of learning: "the importance of *not* learning maladaptively is underestimated" (Shettleworth, 1984, p. 448) In particular, not learning the wrong things at all may be more important than learning the right things quickly (Revusky, 1984). Proposed benefits of learning include being able to adapt to changes and fluctuations in the environment, particularly when the environment may change unpredictably during the animal's lifetime (Slobodkin & Rapaport, 1974; Plotkin & Odling-Smee, 1979), and being able to exploit new niches when required (Davey, 1989).

All adaptive costs and benefits must be understood relative to evolutionarily available alternatives. There may be no a priori need for learning to evolve if other psychological mechanisms suffice to generate adaptive behavior. Why insert an intermediate adaptive process (learning) between the evolution of psychological adaptations and the online functioning of those adaptations in generating behavior contingent on current environmental input? Hardwiring may suffice. Descartes (1662/1972), for example, ignored learning and viewed all animal behavior as reflexive responses to current environmental events. Davey (1989, p. xiii) observed "it is surprising how relatively few species have abandoned fixed behavioral patterns in favor of learning abilities", and Mayr (1974, p. 652) noted "considering this great [supposed] advantage of learning [i.e. adaptability to changing environments], it is rather curious in how relatively few phyletic lines genetically fixed behavior patterns have been replaced by the capacity for the storage of individually acquired information".

The prevalence of hardwiring among terrestrial organisms has several explanations. Staddon (1983, p. 1) suggests "direct stimulus-response mechanisms, plus some sensitivity to rates of change, are sufficient for a wide range of surprisingly intelligent behavior" (see also Braitenberg, 1984). This is particularly true for small, fast-breeding organisms, whose short generation time facilitates rapid evolutionary change in response to environmental change and limits how much time they have to exploit learned information during their lifespan. Staddon (1983, p. 395) also remarks "the longer an animal's life span, and the more varied its niche, the more worthwhile it is to spend time learning". He goes on to note that, given most animals have rather short life spans, "It is not surprising, therefore, that learning pays a rather small part in the lives of most animals" (*ibid.*). Moreover, "animals in invariant environments can rely on equally invariant patterns of behavior" (Mackintosh, 1987, p. 336).

Some science seeks to reduce the strange to the familiar; our goal in this section has been the reverse. We sought in turning the tables on learning to make an apparently commonsense adaptive process seem strange and problematic from an evolutionary point of view. Perhaps by turning up the heat of evolutionary theory on the cauldron of learning theory, we can perform the kind of theoretical annealing that has been so successful in other areas of adaptationist biology. One of the hottest questions to ask becomes, not why have any biological constraints on learning evolved, but why isn't all initially 'learned' behavior canalized into genetically hardwired psychological adaptations? What can a few years of learning really buy a cognitive system already fine-tuned by millenia of natural selection? Whereas connectionists have taken learning as the ultimate adaptive process, real evolving organisms always have an alternative: hardwire the knowledge.

5 TOWARDS A TAXONOMY OF ADAPTIVE FUNCTIONS FOR LEARNING

The most salient aspect of learning to Behaviorists was its dynamic ability to bring organismic behavior into a better fit with the current conditions of the environment. Learning was, quite intuitively, assumed to have a kind of environment-tracking function. It was a way for organisms to adapt to environmental changes faster than evolution could. This view carries over into many recent justifications of learning as an adaptive adjunct to natural selection (e.g. Belew, 1990). But we still wondered why an organism would evolve to allow environmental conditions to change how its behavior-generating mechanisms work (by 'learning'), rather than allowing natural selection to optimize those mechanisms (by 'hardwiring') just as it has optimized so many other physical adaptations (see Mayr, 1974; Staddon, 1983; Menzel, 1984)? How can we clarify and extend these intuitions about the evo-

lution of learning?

We see three main adaptive functions for learning. First, and perhaps most importantly, 'learning' may serve to increase an organism's 'developmental leverage', allowing it to build a larger, more complex, more finely organized phenotype than it otherwise could, given a certain size genotype. Sensitivity to certain predictable environmental regularities during neural development, and the resulting sensory activation patterns, could guide the self-organization of an animal's behavior-generating mechanisms (e.g. see Singer, 1988). Learning may allow the genotype to 'store information in the environment' and let environmental regularities do much of the hard work of wiring up adaptive behavior-generators.² The environmental regularities used in this way may take a rather abstract form. For example, parental 'imprinting' in birds (Lorenz, 1937) can be viewed as a way of building a behavior-generator sensitive to the appearance of one's parent, based on the following environmental regularity: the first large moving thing one sees after hatching is very likely one's parent. Of course, the particular behavior-generator constructed by different birds for recognizing their parents will be different (the birds will 'learn' different parental images), but the species as a whole relies on the same environmental regularity when doing the construction.

Second, 'learning' construed as 'memory encoding' can assist in the spatio-temporal integration of environmental information. Behavior-generators guided only by environmental cues in the here and now may be inferior to behavior-generators sensitive to the relevant environmental cues from the distant and past.³ Animals may evolve to construct the functional equivalent of variable-delay neural delay lines (i.e. 'episodic memories') from certain sensory systems to certain behavior-generators, to expand the temporal scope of their sensitivity to environmental cues. That is, they may evolve to be able to 'bring to mind' information they recruited in the past that is not currently available in the environment. For mobile animals, broadened temporal sensitivity can translate into broadened spatial sensitivity as the animal moves about and recruits environmental information. (Constructing a

² There have been many demonstrations that experimentally depriving developing nervous systems of certain environmental regularities (e.g. certain kinds of visual stimulation) results in maladaptively organized topographic maps (e.g. in striate visual cortex -- see Hubel & Wiesel, 1965; Wiesel & Hubel, 1965). Such experiments should not be construed as demonstrations of the *impotence* of genetic programming and the *importance* of 'experience' and 'learning', but of how efficient evolved developmental mechanisms are at recruiting environmental regularities to assist in self-organization. The fact that neurologists can selectively eliminate those regularities in laboratories should not make us doubt their reliability and ubiquity in nature.

³ However, Gibson (1966, 1979) and Johansson, von Hofsten, and Jansson (1980) warn against underestimating the informational richness of the present, proximal environment, or underestimating the range of adaptive behavior that it can guide.

'mental map' of one's environment from experience gathered during sequential exploration is a paradigmatic example of using 'learning' to integrate information across space and time.) This function of learning neither constructs behavior-generators (function 1 above) nor modifies their online functioning (function 3 below), but simply expands the range of environmental information to which they are sensitive.

Third, as suggested earlier, 'learning' may allow organisms to adjust the online functioning of their behavior-generators faster than natural selection would allow, by conferring on those behavior-generators some sensitivity to dynamic changes in environmental conditions during an organism's life. In this case, learning adapts phenotypes to ongoing changes and *particularities* of the environment rather than depending on environmental *regularities* during phenotype-construction. For example, animals equipped with a special mechanism for learning motor skills to remove new varieties of parasites from their bodies may fare better than competitors lacking such a mechanism if a new species of parasite migrates to the area. Different parasites may have different modes of attachment to the animals' bodies, so may require different removal methods. A learning mechanism that allowed an animal to infer, practice, and perfect an appropriate removal method given, e.g., a visual assessment of the parasite's attachment method, might prove adaptive. In this case, the animals would not be using an environmental regularity to construct a behavior-generator during development (the new parasites weren't around then), or integrating information about the parasites across space and time (parasites currently attached to the animals' bodies are very much in the here and now), they are simply developing a new behavior-generator (e.g. a new method of parasite removal) adapted to a new environmental problem.

This proposed environment-tracking function includes many sub-functions not usually considered in learning theory. One reason for modifying the operation of behavior-generators during the lifetime of an animal is that the animal's body may be growing and changing -- i.e. the spatial relations among its sensory transducers and motor effectors, and between those and its surrounding environment, may be changing. Certain learning mechanisms may evolve to track such relations between the animal's gradually changing body and relatively stable aspects of the environment, rather than tracking 'objective' environmental dynamics per se. Another way of expressing this would be to say learning tracks not just the external environment, but also the corporeal environment of bone, sinew, and flesh in which the animal's behavior-generators (i.e. brains) are embedded.

A second reason for modifying the operation of behavior-generators in response to certain environmental conditions is that animals may need to track changes not only in their individual bodies (their basic phenotypes)

but in their 'extended phenotypes' (Dawkins, 1982), including the location, health, and reproductive status of their kin and offspring. Animals may be able to rely on relatively fixed developmental sequences and internal clocks to modify the operation of their behavior-generators as they grow and age. But to track the whole of their extended phenotype, they need to actually observe when other copies of their genes are being instantiated (i.e. when kin and offspring are born), when those copies are gaining access to metabolic and genetic resources (i.e. growing, eating, and mating), and when they are being threatened (i.e. injured or dying). Animals may evolve learning mechanisms that permit acquisition and maintenance of an ongoing cognitive model of one's kin and social-exchange network, including how to recognize and assist them, and how to request assistance of them. Many learning mechanisms may be rather specifically tuned to promote this type of kin recognition and kin selection (see Hamilton, 1964).

Such complexities aside, environment-tracking is probably the adaptive function of learning most familiar to adult humans (e.g. learning a new restaurant location, or a new person's name), so it has been more commonly studied by psychologists. Yet it is likely to be a less common use of learning than experience-guided phenotype-construction throughout the animal kingdom. Several issues of scale in time and space arise in considering the adaptive functions of learning; these can illuminate why experience-guided development may be more widespread than environment-tracking learning. We might expect small and large organisms to use experience-guided development to almost the same extent, i.e. as much as possible. Big animals, by definition, must build large and often incredibly complex phenotypes given moderately sized genotypes. Although small animals do not have as much of a phenotype to build, they must build their phenotypes from very small gametes that contain very little genetic material.

However, larger phenotypes generally take longer to build, implying longer generation time and a slower rate of genetic evolution (the greater extinction rate of large species is generally thought attributable to their difficulty in genetically adapting to changing environments). Thus, larger phenotypes will generally have a harder time tracking environmental change, and their genotypes will generally lag farther behind being adapted to the current environment, so in general there may be unusually strong adaptive pressures for larger animals to evolve environment-tracking learning mechanisms. The salience of these adaptive pressures to humans, one of the larger and longer-lived species on the planet, may lead us to overestimate the importance and popularity of environment-tracking learning. But, as Davey (1989, p. 20) suggests, "learning ... is likely to evolve only when more fundamental processes of information gain (such as phylogenesis [i.e. evolution]) have reached an upper limit to the amount or rate of change that they can cope with".

Thus, we needn't suppose that small animals lack environment-tracking learning abilities because they are 'less advanced' than larger animals. We must see 'learning' as (1) as way of developmentally fleshing out the genotype during development, (2) informationally fleshing out the animal's current *Umwelt* (view of the world) over space and time, and (3) temporally filling in the adaptive function of environment-tracking in between generation times, when natural selection operates. Small, fast-lived organisms simply don't need much of the third function. Since small-bodied species as collective entities out-number, out-weight, out-reproduce, and typically out-last large-bodied species, they objectively instantiate most of the adaptive agency on this planet. The adaptive functions of learning that are central to them must be considered central to psychology as the (non-anthropocentric) study of adaptive agency.⁴

6 A SIMULATION FRAMEWORK FOR EXPLORING ADAPTIVE AGENCY

A general understanding of adaptive agency cannot rest on experimentation, observation, and theory alone. Some adaptive processes happen on time-scales that preclude experimental manipulation or direct observation. Also, the canalization of terrestrial evolution along certain lines (e.g. using DNA for genotype material) makes us wonder about alternative possibilities not directly observable. And the stochastic, complex nature of adaptive processes make strict experimental control of real physical systems very difficult. For these reasons, researchers have recently turned to computer simulation as the most tractable way of exploring certain adaptive processes. Simulation allows strict control over specified parameters, exploration of alternate phylogenies and developments, and rapid observation of processes that take eons in the real world. Simulation bears on psychology not only as a method for modelling specific psychological adaptations and specific historical phylogenies, but as a way of generally exploring the adaptive processes that produced those phylo-

genies and adaptations. However, the facility with which such simulations can be developed and explored, and the inherent appeal of watching adaptive systems develop, can allow the proliferation of studies not well-grounded in a theoretical framework. Without theoretical grounding, results are difficult to interpret and to assimilate into a coherent picture. We hope the notion of psychology as the study of adaptive agency can help to unify and direct all such studies (observations, experiments, and simulations).

To capture all of the adaptive processes discussed earlier in a simulation, we must have methods of simulating genotypic evolution, the generation of behavior, and the ability to learn new behaviors. (We ignore 'cultural' transmission for now, but see Belew, 1990). Specifically, we use a genetic algorithm to evolve successive generations of a population of neural network architectures, which in turn control the behavior of simple creatures which can learn as they live in a simulated environment. This is an extension of earlier work exploring the use of genetic algorithms to design network architectures capable of learning specific input/output mapping tasks (Miller, Todd, and Hegde, 1989); here, the algorithm's measure of fitness depends not on learning an arbitrary task, but on behaving adaptively in the simulated environment.

In simulating the evolution of further adaptive processes, for instance specific behavior-generators and specific learning abilities, we must first specify some environment and what defines fitness in that environment. We then observe, through the course of the evolutionary simulation, which adaptive processes are most important for maximizing individual fitness by "solving" the relevant environmental problems. Ackley and Littman's (1990) sophisticated simulations, for example, show the utility of evolving motivational systems to guide learning. Nevertheless, their simulations are pre-set to operate either with learning or without, rather than set up with an environment whose adaptive problems allow the evolution of learning itself to be studied. We strive to create simulations in which the subsidiary adaptive processes that evolution can spawn are as open-ended as possible.

6.1 TYPES OF LEARNING

Since we are primarily interested in using simulation to explore issues in the evolution of learning, we must address the varieties of learning we could investigate. Tolman (1932) asked "Is there more than one type of learning?" The answer, of course, depends on what one means by 'type'. Different kinds of learning might be distinguished by at least three different criteria: (1) type of experimental paradigm used to investigate it; (2) neural mechanism implementing it; (3) adaptive functions served by it. Behaviorists concentrated on the first criterion, defining classical versus operant conditioning, for example, according to the sorts of environmental contingencies the experimenter sets up for the laboratory an-

⁴ The suggested logical and historical primacy of experience-dependent development over environment-tracking learning is a conjecture not easily supported or refuted. It is appealing to us because learning mechanisms that evolved originally for experience-guided development of behavior-generators could conceivably serve as pre-adaptations for the evolution of environment-tracking learning. That is, given a learning mechanism that recruits patterns of environmental stimulation to help in construction of a psychological adaptation, one could easily imagine how prolongation of that mechanism's sensitive period could confer longer, perhaps lifelong, adaptive flexibility to that mechanism. Selective pressures to prolong such sensitive periods, we conjecture, may underlie the evolution of most environment-tracking learning mechanisms. Johnston (1981) discusses how potentially adaptive concomitants of evolved learning abilities might serve as pre-adaptations for later learning abilities by conferring on animals certain "ecologically surplus abilities", i.e. learning abilities not directly selected for. Rozin (1976) also suggests that learning mechanisms may have evolved first as isolated specializations and later became available to a wider range of behavior-generators.

imal. Connectionists view the first criterion in terms of structuring the training set and test set for the network, and the second in terms of the specific mathematical learning algorithm implemented in the network. Rarely has a taxonomy of learning been conceptualized in terms of adaptive function. The theoretical and empirical study of learning could concentrate on any of these three ways of distinguishing learning processes, or it could pursue a fourth strategy of attempting to elucidate general principles of operation governing all learning processes, understood in some sufficiently abstract way -- as in Shepard's (1987a) work towards a universal law of generalization.

Fortunately, there may be some correspondence between learning mechanisms construed in terms of experimental paradigm used to investigate them and learning mechanisms construed in terms of their adaptive functions. Learning mechanisms evolved to solve particular ecological problems; to the extent that different experimental paradigms present ecologically valid, adaptively isomorphic problems, they may map onto real adaptive functions of learning. But the fit between adaptive functions and neural mechanisms may be much looser: "analysis of learning in terms of functional problems does not map directly onto the learning theorists' analysis in terms of [neural] mechanisms" (Shettleworth, 1984, p. 431). Different adaptive functions might be implemented by similar neural mechanisms for changing synaptic weights, or the same adaptive function might be implemented in very different neural mechanisms in different species.

For these reasons, we chose to categorize learning mechanisms by adaptive function rather than by experimental paradigm or neural mechanism. But the problem remains: which adaptive function should be explored first? To investigate the evolution of learning for exploiting environmental regularities specifically as a means of maximizing phenotype size and complexity given limited genotype size, a simulation must include adaptive pressures or constraints on genotype length or specificity. Without such pressures or constraints, the genotype may simply expand to accurately specify (hardwire) the appropriate phenotype, rather than evolving developmental tricks that depend on internalizing environmental regularities during development. Clearly, such pressures depend on what developmental mechanisms exist or can evolve for generating phenotypes from genotypes. Since the vagaries of real neural development and of adaptive pressures on genotype size are still poorly understood, we have not found a satisfactory general method for simulating neural development or for imposing pressures on genotype length. So we have avoided simulating the evolution of learning as a way of achieving developmental leverage.

Exploring the third 'adapting to particularities' function of learning requires setting up an adaptive problem with environmental changes too rapid for genotypic evolution

to track. But if the environment changes *too* rapidly, then extensive simulation of learning during an organism's lifespan would be required. Likewise, exploring the second information-integration function of learning would require simulating extensive interactions between creatures and worlds, including recurrent network dynamics or memory encoding and retrieval systems. We did not want to become mired in simulating constantly-changing environments or sophisticated dynamic learning mechanisms, so we opted for a kind of imprinting scenario, where experience-guided development uses an abstract environmental regularity to help build behavior-generators adapted to unpredictable particularities of the niche, as follows.

The simplest way to defeat natural selection is to make the genotype unable to know ahead of time which one of two alternate niches it will find itself in during phenotypic development. If one niche requires one kind of behavior-generation mechanism and another niche requires a different kind, natural selection alone will be unable to select the proper mechanism to guide the phenotype's behavior. Natural selection must instead select for the evolution of a more general mechanism that can flip into one of two states depending on some assessment of which niche it finds itself in. Thus, we chose a kind of "imprinting" or parameter-setting based on the early environment as the simplest possible case in which learning, construed as adaptation to specific environmental regularities, could evolve.

6.2 INCLUSIVE FITNESS AS THE ONLY NATURAL 'SUPERVISOR' FOR LEARNING

Many previous attempts to use genetic algorithms to evolve neural network architectures have evaluated architecture fitness by training the networks with a supervised learning procedure, (i.e. one with an externally provided "target vector" the network is to produce given each input vector), typically back-propagation (Belew, McInerney, & Schraudolph, 1990; Miller, Todd, & Hegde, 1989; for a review, see Weiss, 1990). While supervised learning paradigms may be appropriate in evolving connectionist systems for particular commercial applications, they are problematic and perhaps misleading in scientific studies of adaptive agency. In particular, to be biologically plausible, the source of the "targets" or other supervising feedback must be justified. Organisms as whole functioning agents in real environments rarely receive patterns of information analogous to training signals in back-propagation. Although the distinction between supervised and unsupervised learning procedures can be blurred, we have chosen to focus on the more defensible latter end of the spectrum, including self-organizing, associative, and simple feedback-based mechanisms. But even if we sidestep the issue of target-based training, the concept of feedback still raises problems.

Years of learning by "being taught" instill in us intuitions about the utility of corrective feedback to guide learning. But such intuitions make it easy to overlook the fact that it is at least as difficult for organisms to evolve the ability to perceive feedback signals from the environment to guide their learning, as it is to evolve the perception of any other complex external cues. Consider for example the complexities involved in registering the information that one has just been rebuffed in a social exchange. Feedback signals cannot be assumed to be just somehow "provided" to an organism for it to use in adjusting its behavior. Instead, feedback systems must be understood as special *sensory* systems evolved to provide information to special learning mechanisms that in turn adaptively change the functioning of certain behavior generators (e.g. in the simulations of Ackley and Littman, 1990). Feedback systems, whether motivational, emotional, volitional, or proprioceptive, evolve just like other aspects of adaptive agency -- by cumulative selection of incrementally better-adapted designs.

Ideally, an organism might prefer to guide its learning with direct information about how its inclusive fitness changes as a result of its behavior. But there is no such thing as an inclusive fitness transducer that can be used to supervise learning. Organisms must instead evolve to sense inclusive fitness indirectly, through whatever proximal sensory cues have been reliably associated with increased fitness in their environment. Thus, natural selection itself is ultimately the *only* source of 'supervision' for learning systems.

The indirectness of natural selection's supervision of learning leads to the complexity inherent in real evolved learning mechanisms. Humbled by this complexity, we decided not to clutter our initial simulations with the requirement of evolving a motivational system to provide supervising feedback during learning, *in addition* to evolving the learning system itself. So instead we chose to start by exploring the simplest set of unsupervised associative learning mechanisms we could conceptualize, as will be described below.

6.3 THE GENETIC ALGORITHM FOR EVOLVING NEURAL NETWORKS

To simulate the evolution of learning in our explorations of adaptive agency, we use a relatively standard form of Holland's (1975) genetic algorithm, combined with a simple "developmental" method which translates genotypes into neural network architecture phenotypes. In this method, a strong genetic specification scheme (as defined by Miller, Todd, and Hegde, 1989) interprets each genotype as a *connectivity constraint matrix* that directly specifies the nature of each unit and connection in the network architecture.

Once a network, instantiating the behavioral mechanisms of an individual creature, has been so constructed, it is evaluated in the simulated world over several time-steps representing the creature's lifespan. During each time-step, a creature's network receives sensory input based on the current external environmental cues available, processes that input according to its architecture and current weights, generates behavior based on the activation of its output units, and changes its connection weights based on an unsupervised learning rule (e.g. Hebbian association). The effects of the creature's behavior on the world and on its own fitness are then registered, and the next time-step begins.

7 A SIMPLE SCENARIO FOR THE EVOLUTION OF UNSUPERVISED LEARNING

For our first exploration of adaptive agency, we attempted to devise the simplest, cleanest scenario in which learning could prove adaptive, focusing on a kind of imprinting function. After analyzing the building blocks needed for associative learning, we analyzed what sorts of environments might exert adaptive pressures to evolve that type of learning. Finally, we constructed an appropriate simple world to see if learning would spread through a population of simulated creatures behaving in that world. We explain the scenario used by outlining a biological metaphor that specifies the structure of the econiche and the nature of the adaptive problem. (The scenario and results are described in more detail in Todd and Miller (1990); space constraints preclude making this section much more than an overview.)

Our scenario can be imagined as an underwater realm, in which parents emit eggs randomly into two different types of feeding patches: those where food is green and poison red, or vice-versa. Each creature in this world lives a fixed lifespan, eating or ignoring food and poison at each life-step, and amassing energy which determines its eventual number of offspring in the next generation. Eating food raises energy; eating poison drains energy. Food smells sweet and poison smells sour across all creatures, but with some perceptual error rate -- the smell-sense accuracy -- determined by the turbulence of the water in this world. Food and poison each have characteristic fixed colors *within* one creature's life, but the meaning of each color varies *between* creatures, food being red for some and green for others, depending on their patch as mentioned above. The color-sense is 100% accurate. Thus natural selection can 'predict' the association between smell and object, but not between color and object -- this will be the task for learning.

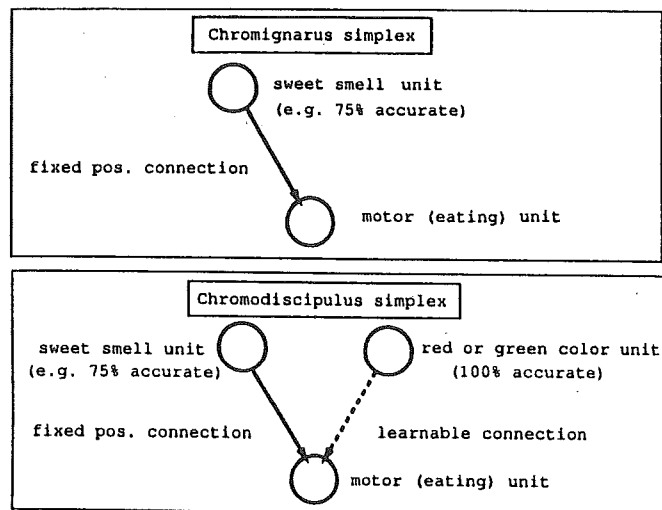


Figure 1. The network designs of *Chromignarus* (without color learning), top, and *Chromodiscipulus* (with color learning), bottom.

7.1 EVOLVING ADAPTIVE NETWORKS

As the genetic operators process genotypes again and again through many generations, recombination and mutation will sometimes produce particularly fit network designs. One reasonable design we can expect consists of a sweet smell sensing unit connected to a motor unit (which controls the creature's eating) by a fixed positive weight. A creature equipped with this nervous system will depend purely on smell to decide when to eat, with its behavioral accuracy dependent on the inherent noisiness of the smell sense in its world (i.e. by amount of turbulence). Although this design will sometimes make mistakes (that is, ignoring food or eating poison that smells wrong because of the turbulence), still on average it will eat more food than poison. Thus these creatures' energy, and number of offspring, will be higher than if they were just eating randomly. We call this design the *color-blind eater*, or, more fancifully, *Chromignarus simplex*, and it is shown in the top portion of Figure 1.

The best evolved creature design though is an elaboration on the color-blind eater chassis. Creatures of this type gain an adaptive edge by including a red or green color sensing unit, along with a learnable connection to the motor (eating) unit. This creature design we name *Chromodiscipulus simplex*, color-learning eaters, as shown at the bottom of Figure 1.

With this design, an appropriately excitatory link between the color unit and the eating unit will be built up over successive time-steps in the creature's life by a Hebbian correlational weight-change mechanism. Eventually, this weight will be large enough that the color unit alone can cause the eating unit to come on, *regardless* of what

the sweet smell unit says to do -- the creature has now learned that a particular color means food. The creature can now rely on this completely accurate visual cue, rather than the inaccurate smell cue, and always choose to eat properly, thereby increasing its fitness further.

7.2 INITIAL SIMULATION RESULTS

We continued our investigation of the conditions under which learning could evolve with the more interesting question of how *quickly* learning would evolve, given various smell accuracies in different worlds. By tracking population average fitness values, it is possible to tell when the use of learning has spread through the population. Initially, the average fitness quickly rises to a plateau at which the fixed sweet-smell to eating unit connection is present in most of the creatures (*Chromignarus* temporarily dominates). After remaining fairly level at this fitness value for possibly many more generations, the average population fitness again jumps, indicating that the learnable connection from a color unit has penetrated the population (*Chromodiscipulus* ultimately rules). Average fitness then levels out again, this time around its final highest value. Thus recording when fitness jumps occur can tell us when the different creature designs predominate in the population.

For each of 17 smell accuracies between 50% and 100% we ran 20 populations of 100 individuals for 1500 generations each. Figure 2 shows how many generations it took each population to make each of the two jumps to new fitness-plateaus. These two jumps correspond to the widespread appearance of *Chromignarus* (without color learning) -- indicated by asterisks -- and of *Chromodiscipulus* (with color learning) -- indicated by bullets. The bottom curve shows the average number of generations taken to evolve *Chromignarus* across the 20 runs at each accuracy level, and the top curve indicates the analogous average generations to evolve *Chromodiscipulus*, and thus learning itself.

The fixed smell connection (the *Chromignarus* design) evolves rapidly, in less than 100 generations for most accuracy levels. The greater the accuracy of smell, the more quickly the fixed smell connection spreads, because the adaptive advantage to be gained from evolving it (i.e. the adaptive pressure) increases. More interesting is the effect of smell accuracy on time taken to evolve color learning (*Chromodiscipulus*). Here we found an unexpected U-shaped relationship: color learning evolved most quickly for smell accuracies around 75%, and took longer and longer for accuracies diverging on either side of that middle range (as shown by the upper, solid, curve in Figure 2).




Figure 2. Generations to evolve *Chromignarus* (without color learning) and *Chromodiscipulus* (with learning) plotted against smell-sense accuracy. Asterisks indicate time to evolve *Chromignarus* for each of 20 runs at each of 17 smell-sense accuracy levels; the dotted line indicates average time across the 20 runs at each smell accuracy. Bullets indicate time to evolve *Chromodiscipulus* for the same 20 runs at each accuracy level; the solid line indicates average time across the 20 runs at each accuracy. Note that average time to evolve *Chromignarus* decreases monotonically as smell-sense accuracy increases, but average time to evolve *Chromodiscipulus* follows a U-shaped Learning Evolution Curve.

7.3 THEORETICAL INTERPRETATION OF THE U-SHAPED EVOLUTION FUNCTION

We view the U shape as emerging indirectly from a trade-off between the phylogenetic adaptive pressure to evolve learning (during species-wide evolution), and the ontogenetic ease of learning (during each individual creature lifetime). These forces interact at the various smell sense accuracy levels as follows.

At low smell accuracies, where *Chromignarus* does quite poorly, there is great phylogenetic adaptive pressure to evolve color learning, because it would add significantly to this creature's fitness by overcoming its color-blind, smell-guided error-prone behavior. However, this large potential benefit is offset by the ontogenetic difficulty of actually accomplishing learning at noisy low smell accuracies. In fact, learning can be so slow in this case that a learning creature's lifespan may elapse before it gains any benefit from this ability. Thus the ontogenetic difficulty of learning offsets its high phylogenetic adap-

tiveness at low smell accuracies, and learning will take a long time to evolve.

At high smell accuracies, in contrast, color learning would be easy to perform ontogenetically, because perceived smell, color, and substance type will be highly correlated, and the associations between them could build up quickly. However, there is little phylogenetic adaptive pressure to evolve color learning in this case, because the smell sense alone suffices to guide highly adaptive eating behavior. Since natural selection cannot distinguish *Chromignarus* from *Chromodiscipulus* if they are both doing almost perfectly, this 'ceiling effect' will keep *Chromodiscipulus* from proliferating. So again color learning will take a long time to evolve.

But for middle smell accuracies, color learning is relatively adaptive and relatively easy. Color learning gives a significant fitness increase over using smell alone, and learning can occur fairly quickly, since the eating unit comes on rather more often to food than to poison. Mid-level smell accuracy represents a happy medium between phylogenetic adaptive pressure and ontogenetic ease of learning, leading to the rapid evolution of color learning and its spread through the population.

8 PLANNED EXTENSIONS AND FUTURE RESEARCH

Our theoretical motivation will continue to be the exploration of adaptive agency and interactions among adaptive processes; our methodological strategy will continue to focus on the search for simple, elegant scenarios that reveal potentially general patterns and dynamics underlying adaptive agency. Given this orientation, we have gradually abandoned our earlier ambitions to create a general-purpose system for investigating the evolution of very complex nervous systems in very complex environments. The rush to build as much biological realism as possible into our simulations as quickly as possible, can, we fear, obscure those features of simulation that make it so useful in other sciences: parametric control, replicability, conceptual clarity, ease of analysis, and speed. Thus, we hope to develop more simple scenarios that not only capture the central features of certain adaptive problems, but that can reveal unanticipated patterns and complexities.

More specifically, we intend to develop a series of slightly more complex learning scenarios to investigate how natural selection, associative learning, and environmental dynamics interact. One could imagine that, given a series of results from such scenarios, a more general theory concerning the interaction of adaptive processes might emerge -- not a formalistic model in terms of dynamical systems or information theory, but a concrete understanding of the interactions among adaptive pressures, cue structures in different environments, genetic representations and operators, developmental mechan-

isms, learning, behavior-generation, and information-processing. Later, we intend to address the adaptive problems of foraging, communication, and protean behavior. At each step, we hope to keep our motivations for simulation closely tied to resolving theoretical issues in the study of adaptive agency, while remaining sensitive to the sorts of unanticipated phenomena, patterns, and dynamics that simulation research so often reveals. In allowing research to be guided so strongly by a clearly articulated conceptual framework, we may give up some of the immediate richness and appeal of simulation-for-its-own sake, but we hope to achieve a theoretical depth and breadth, and a connection to major issues and perennial questions, that will, we believe, be more satisfying in the end.

Acknowledgements

This research was partially supported by National Science Foundation Graduate Fellowships to both authors. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation.

We wish to thank Dave Rumelhart, Roger Shepard, John Tooby, Leda Cosmides, David Ackley, and Rik Belew, who have influenced and guided this research along the way.

References

- Ackley, D.H., & Littman, M.S. (1990). Learning from natural selection in an artificial environment. In *Proceedings of the International Joint Conference on Neural Networks* (pp. 189-193). Hillsdale, NJ: Erlbaum.
- Belew, R.K. (1990). Evolution, learning and culture: Computational metaphors for adaptive search. *Complex Systems*, 4, 11-49.
- Belew, R.K., McInerney, J., & Schraudolph, N.N. (1990). *Evolving networks: Using the genetic algorithm with connectionist learning* (Tech. Rep. No. CS90-174). San Diego: University of California, Computer Science and Engineering Department.
- Boakes, R.A. (1984). *From Darwin to Behaviorism*. Cambridge, UK: Cambridge University Press.
- Bolles, R.C. (1970). Species-specific defense reactions and avoidance learning. *Psychological Review*, 77, 32-48.
- Bolles, R.C. (1988). Nativism, naturalism, and niches. In R.C. Bolles & M.D. Beecher (Eds.), *Evolution and learning* (pp. 1-15). Hillsdale, NJ: Erlbaum.
- Bolles, R.C., & Fanselow, M.S. (1980). A perceptual-defensive-recuperative model of fear and pain. *Behavioral and Brain Sciences*, 3, 291-301.
- Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press/Bradford Books.
- Brant, D.H., & Kavanau, J.L. (1965). 'Unrewarded' exploration and learning in complex mazes by wild and domestic mice. *Nature*, 204, 267-269.
- Breland, K., & Breland, M. (1961). The misbehavior of organisms. *American Psychologist*, 16, 681-684.
- Chomsky, N. (1957). *Syntactic structures*. The Hague: Mouton.
- Cosmides, L., & Tooby, J. (1987). From evolution to behavior: Evolutionary psychology as the missing link. In J. Dupre (Ed.), *The latest on the best: Essays on evolution and optimality*. Cambridge, MA: MIT Press/Bradford Books.
- Curio, E., & Vieth, W. (1978). Cultural transmission of enemy recognition. *Science*, 202, 899-901.
- Davey, G. (1989). *Ecological learning theory*. New York: Routledge.
- Dawkins, R. (1976). *The selfish gene*. Oxford, UK: Oxford University Press.
- Dawkins, R. (1982). *The extended phenotype*. Oxford, UK: Oxford University Press.
- Dawkins, R. (1983). Universal Darwinism. In D.S. Bendall (Ed.), *Evolution from molecules to men*. Cambridge, UK: Cambridge University Press.
- Descartes, R. (1662/1972). *Treatise on man*. Cambridge, MA: Harvard University Press. Translated by T.S. Hall.
- Emlen, J.M. (1973). *Ecology: An evolutionary approach*. Reading, MA: Addison-Wesley.
- Estes, W.K. (1984). Human learning and memory. In P. Marler & H.S. Terrace (Eds.), *The biology of learning* (pp. 617-628). New York: Springer-Verlag.
- Galef, B.G. (1988). Evolution and learning before Thorndike: A forgotten epoch in the history of behavioral research. In R.C. Bolles & M.D. Beecher (Eds.), *Evolution and learning* (pp. 39-58). Hillsdale, NJ: Erlbaum.
- Garcia, J., & Koelling, R.A. (1966). Relation of cue to consequence in avoidance learning. *Psychonomic Science*, 4, 123-124.
- Gibson, J.J. (1966). *The senses considered as perceptual systems*. Boston: Houghton-Mifflin.
- Gibson, J.J. (1979). *The ecological approach to visual perception*. Boston: Houghton-Mifflin.
- Gluck, M.A., Bower, G.H., & Hee, M.R. (1989). A configural-cue network model of animal and human associative learning. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society* (pp. 323-332). Hillsdale, NJ: Erlbaum.
- Gould, J.L., & Marler, P. (1984). Ethology and the natural history of learning. In P. Marler & H.S. Terrace (Eds.), *The biology of learning* (pp. 47-74). New York: Springer-Verlag.
- Hamilton, W.D. (1964). The evolution of social behavior I and II. *Journal of Theoretical Biology*, 7, 1-52.
- Hernstein, R.J. (1979). Acquisition, generalization, and

- discrimination reversal of a natural concept. *Journal of Experimental Psychology: Animal Behavior and Processes*, 5, 116-129.
- Hinde, R.A. (1973). Constraints on learning: An introduction to the problems. In R.A. Hinde & J. Stevenson-Hinde (Eds.), *Constraints on learning: Limitations and predispositions*. London: Academic Press.
- Hinton, G.E., & Nowlan, S.J. (1987). How learning guides evolution. *Complex Systems*, 1, 495-502.
- Holland, J. H. (1975). *Adaptation in natural and artificial systems*. Ann Arbor, MI: University of Michigan Press.
- Hull, C.L. (1943). *Principles of behavior: An introduction to behavior theory*. New York: Appleton.
- Jenkins, H.M. (1984). The study of animal learning in the tradition of Pavlov and Thorndike. In P. Marler & H.S. Terrace (Eds.), *The biology of learning* (pp. 89-114). New York: Springer-Verlag.
- Johansson, G., von Hofsten, G., & Jansson, G. (1980). Event perception. *Annual Review of Psychology*, 31, 27-63.
- Johnston, T.D. (1981a). Contrasting approaches to a theory of learning. *Behavioral and Brain Sciences*, 4, 125-173.
- Johnston, T.D. (1981b). Selective costs and benefits of learning. In J.S. Rosenblatt, R.A. Hinde, C. Beer, & M.C. Busnel (Eds.), *Advances in the study of behavior*. New York: Academic Press.
- Kamil, A.C., & Roitblat, H.L. (1985). The ecology of foraging behavior: implications for animal learning and memory. *Annual Review of Psychology*, 36, 141-169.
- Kline, L.W. (1898). Suggestions toward a laboratory course in comparative psychology. *American Journal of Psychology*, 10, 399-430.
- Langton, C.G. (Ed.). (1989). *Artificial life*. New York: Addison-Wesley.
- Lea, S.E.G. (1984). *Instinct, environment, and behavior*. London: Methuen.
- Logue, A.W. (1979). Taste aversion and the generality of the laws of learning. *Psychological Bulletin*, 86, 276-296.
- Lorenz, K.Z. (1937). The companion in the bird's world. *Auk*, 54, 245-273.
- Lorenz, K.Z. (1952). The past twelve years in the comparative study of behavior. In C.H. Schiller (Ed.), *Instinctive behavior*. New York: International Universities Press.
- Mackintosh, N.J. (1987). Learning. In D. McFarland (Ed.), *The Oxford guide to animal behavior* (pp. 336-346). New York: Oxford University Press.
- Marler, P. (1984). Song learning: Innate species differences in the learning process. In P. Marler & H.S. Terrace (Eds.), *The biology of learning* (pp. 289-310). New York: Springer-Verlag.
- Mayr, E. (1974). Behavior programs and evolutionary strategies. *American Scientist*, 62, 650-659.
- Meehl, P.E. (1950). On the circularity of the law of effect. *Psychological Bulletin*, 47, 52-75.
- Menzel, R. (1984). Biology of invertebrate learning (Group Report). In P. Marler & H.S. Terrace (Eds.), *The biology of learning*. Berlin: Springer-Verlag.
- Michotte, A. (1954). *La perception de la causalite*, (2nd ed.). Louvain: Publications Universite de Louvain.
- Miller, G.F., Todd, P.M., & Hegde, S.U. (1989). Designing neural networks using genetic algorithms. In J.D. Schaffer (Ed.), *Proceedings of the Third International Conference on Genetic Algorithms* (pp. 379-384). San Mateo, CA: Morgan Kaufmann.
- Morgan, C.L. (1896). *Habit and instinct*. London: Edward Arnold.
- Ohman, A., Dimberg, U., & Ost, L.-G. (1985). Animal and social phobias: Biological constraints on learned fear responses. In S. Reiss & R.R. Bootzin (Eds.), *Theoretical issues in behavior therapy*. New York: Academic Press.
- Plotkin, H.C., & Odling-Smee, F.J. (1979). Learning, change, and evolution: An enquiry into the teleonomy of learning. *Advanced Studies of Behavior*, 10, 1-41.
- Rescorla, R.A., & Wagner, A.R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A.H. Black & W.F. Prokasy (Eds.), *Classical conditioning II: Current research and theory*. New York: Appleton-Century-Crofts.
- Revusky, R. (1977). Interference with progress by the scientific establishment: Examples from flavor aversion learning. In N. Milgram, L. Krames, & T. Alloway (Eds.), *Food aversion learning*. New York: Plenum.
- Garcia y Robertson, R., & Garcia, J. (1988). Darwin was a learning theorist. In R.C. Bolles & M.D. Beecher (Eds.), *Evolution and learning*. Hillsdale, NJ: Erlbaum.
- Rozin, P. (1976). The evolution of intelligence and access to the cognitive unconscious. *Progress in Psychobiology and Physiological Psychology*, 6, 245-280.
- Rumelhart, D.E., & McClelland, J.L. (Eds.). (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press/Bradford Books.
- Seligman, M.E.P. (1970). On the generality of the laws of learning. *Psychological Review*, 77, 406-418.
- Seligman, M.E.P., & Hager, J.L. (1972). *Biological boundaries of learning*. Englewood Cliffs, NJ: Prentice-Hall.
- Shepard, R.N. (1987a). Evolution of a mesh between principles of the mind and regularities of the world. In J. Dupre (Ed.), *The latest on the best: Essays on evolution and optimality*. Cambridge,

- MA: MIT Press/Bradford Books.
- Shepard, R.N. (1987b). Toward a universal law of generalization for psychological science. *Science*, 237, 1317-1323.
- Shepard, R.N. (1989). Internal representation of universal regularities: A challenge for connectionism. In L. Nadel, L.A. Cooper, P. Culicover, & R. M. Harnish (Eds.), *Neural connections and mental computation*. Cambridge, MA: MIT Press/Bradford Books.
- Shettleworth, S.J. (1983). Function and mechanism in learning. In M. Zeiler & P. Harzen (Eds.), *Advances in analysis of behavior: Biological factors in learning (Vol. 3)*. New York: Wiley.
- Shettleworth, S.J. (1984). Natural history and evolution of learning in nonhuman mammals. In P. Marler & H.S. Terrace (Eds.), *The biology of learning* (pp. 419-434). New York: Springer-Verlag.
- Singer, W. (1988). Ontogenetic self-organization and learning. In J.L. McGaugh (Ed.), *Brain organization and memory: Cells, systems, and circuits*. New York: Oxford University Press.
- Skinner, B.F. (1953). *Science and human behavior*. New York: Macmillan.
- Slobodkin, L.B., & Rapoport, A. (1974). An optimum strategy in evolution. *Quarterly Review of Biology*, 49, 181-200.
- Small, W.S. (1900). An experimental study of the mental processes of the rat. *American Journal of Psychology*, 11, 131-165.
- Spencer, H. (1855). *Principles of Psychology* (Vols. 1-2). New York: Appleton.
- Staddon, J.E.R. (1980). Optimality analyses of operant behavior and their relation to optimal foraging. In J.E.R. Staddon (Ed.), *Limits to action: The allocation of individual behavior* (pp. 101-141). New York: Academic Press.
- Staddon, J.E.R. (1983). *Adaptive behavior and learning*. Cambridge, UK: Cambridge University Press.
- Stork, D.G., & Keesing, R.C. (1990). *Evolution and learning in neural networks: The number and distribution of learning trials affect the rate of evolution*. Paper presented at Neural Information Processing Systems: Natural and Synthetic, Denver, CO.
- Sutton, R.S., & Barto, A.G. (1987). A temporal-difference model of classical conditioning. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society* (pp. 355-378). Hillsdale, NJ: Erlbaum.
- Thorndike, E.L. (1898). Animal intelligence: An experimental study of the associative processes in animals. *Psychological Review Monographs*, 2, 1-109.
- Thorndike, E.L. (1911). *Animal intelligence*. New York: Macmillan.
- Timberlake, W., & Lucas, G.A. (1989). Behavior systems and learning: >From misbehavior to general principles. In S.B. Klein & R.R. Mowrer (Eds.), *Contemporary learning theories* (pp. 237-275). Hillsdale, NJ: Erlbaum.
- Tinbergen, N. (1951). *The study of instinct*. Oxford: Clarendon Press.
- Todd, P.M., & Miller, G.F. (1990). Exploring adaptive agency II: Simulating the evolution of associative learning. In *Proceedings of the International Conference on Simulation of Adaptive Behavior: From Animals to Animats*. Cambridge, MA: MIT Press.
- Todd, P.M., & Miller, G.F. (in press). Exploring adaptive agency III: Simulating the evolution of habituation and sensitization. In *Proceedings of the First International Conference on Parallel Problem Solving from Nature*. Berlin: Springer-Verlag.
- Vaccarino, F.J., Schiff, B.B., & Glickman, S.E. (1989). Biological view of reinforcement. In S.B. Klein & R.R. Mowrer (Eds.), *Contemporary learning theories: Pavlovian conditioning and the status of traditional learning theory* (pp. 111-142). Hillsdale, NJ: Erlbaum.
- Watson, J.B. (1914). *Behavior: An introduction to comparative psychology*. London: Holt, Rinehart, & Winston.
- Watson, J.B. (1919). *Psychology from the standpoint of a behaviorist*. Philadelphia: Lippincott.
- Weiss, G. (1990). *Combining neural and evolutionary learning: Aspects and approaches* (Tech. Rep. No. FKI-132-90). Munich: Technical University of Munich.
- Williams, D.R., & Williams, H. (1969). Auto-maintenance in the pigeon: Sustained pecking despite contingent non-reinforcement. *Journal of the Experimental Analysis of Behavior*, 12, 511-520.