

# Vision, Learning, and Development

Christopher M. Brown

The University of Rochester  
Computer Science Department  
Rochester, New York 14627

Technical Report 492

May 1994

## **Abstract**

It seems to be a common feeling that animals learn to see, and this feeling, together with the re-emergence of computer learning paradigms that mimic many forms of human learning, has raised hopes that learning is the key to the computer vision problem. Indeed, it seems clear that Nature does not “program” all our visual capabilities into the genome, and we certainly know that programming a computer with a closed-form solution to the vision problem is a daunting task.

This aim of this informal and elementary report (basically a term paper) is to cast doubt on the idea that biological systems learn to see. The complex process of development, beginning at fertilization and ending with a mature individual, could be considered to have genetic (“nature”) and learning (“nurture”) processes as logical endpoints or opposite poles. This report mostly considers what goes on between those endpoints, and is meant to raise the possibility that some of the least understood processes in biology are responsible for visual capabilities.

# 1 Argument and Apology

In the mid to late 1980s, “computer learning” emerged from a decade of neglect and became one of the hot research areas of the 1990s. For the purposes of this paper and to save argument, I am willing to stipulate that work in this area has largely succeeded, by which I mean and that several good computational models of learning exist. I am concerned here only with one putative application of learning, *viz.* the application of learning in “learning to see”. The idea of learning to see is popular with us computer vision researchers for two reasons. First, we have not been able to solve the vision problem in closed form, and the idea of explicitly programming in (into some genome, as it were) all the necessary visual capabilities seems increasingly daunting. Second and therefore, there seems to be a strong urge among us to believe that successful (i.e. biological) vision systems learn to do vision and therefore that computer learning can play a central, helpful role in the development of computer vision systems.

This report is to suggest that the sort of learning that computer scientists have been modeling, and indeed that which psychologists usually study, may be inadequate to learn vision. The specialized neural structures that support vision certainly do *develop* in conjunction with experience of one sort or another. However, development is a powerful and little-understood set of interacting processes that is only influenced by the genome and that only includes learning as a subprocess. Learning theory is about individual-specific experience, but many results of development are only indirectly determined by individual-specific experiences. If biological vision is a developmental phenomenon rather than a learning phenomenon, then most current computer learning models are irrelevant to acquiring biological vision capabilities. Further, the processes possibly most relevant are among the least well understood in all of biology.

Learning is indeed useful in computer vision. If one can already see, then learning and vision work together nicely. For instance, with vision we can learn model parameters that can be used for matching and object recognition – there are endless examples of this, from ancient [1] to modern [13]. I reemphasize that this report concerns *learning vision*, not *learning using vision*.

Do not expect me to clear up any issues, to explain anything, to put forward a new theory of vision, to present any new results, nor least of all to write as an authority on learning, development, or neuroscience. This report is mostly at the level of a term paper in a course in development. My brief re-packaging of these complex subjects is bound to be frustrating or maddening to almost everyone, especially developmental psychologists, ethologists, computer visionaries, and computational learning scholars of all flavors. I can only hope that it might provide some new facts or unfamiliar ideas to a forgiving subset of the readership in the computer vision community, and that it might challenge those readers to think, criticize, read further, and quickly go beyond the confines of this report.

## 2 Antithesis and Thesis

The idea that humans learn to see seems pervasive. In numerous talks in different contexts I have heard words to the effect that babies are “learning to see” during their first months or years of life. There have been at least three recent workshops on vision and learning ([73, 10], plus one I organized for ARPA-IU PIs). It is not hard to find quotations like the following.

*It is known that learning takes place over a long period and plays a central role in the development of such a capability [vision] in humans. . . .*

*The evidence for learning in vision includes even low-level vision. For instance, a common visual experience, overhead light source, is learned and used to perceive shape from shading [reference to Ramachandran] (1)<sup>1</sup> [38].*

We shall return to this work in Section 6, but I should note here that the learning model proposed in [38] is one that indeed extends the architecture of the basic neural net model by processes that add structure to the net, so it incorporates an important aspect of development into its interpretation of “learning”.

It seems to me that more vision researchers than not believe, or assume, or are forced into the position, that we learn all sorts of basic visual capabilities. There are two problems: First, it is not at all clear that the existing models of learning in psychology and the neurosciences have anything to do with learning vision, and second there is plenty of evidence that learning (as normally understood) is not necessary for vision.

Of course we “learn” to see in the sense that something changes in the brain to make it happen. But what change is it? Is it learning in the sense of conditioning, which is the learning that most psychologists study and that computer scientists have spent much effort modeling? This paper claims not, and claims that developmental psychology provides interesting (or disturbing) insights into the mechanisms underlying vision and other basic capabilities such as speech and motor control.

Both psychology and computer science are quite comfortable with learning (as practiced by mature organisms) and actually know very little of the mechanisms of development, and the natural inclination has been to try to apply the tools available to the problem — to apply “learning” to “vision”. Fig. 1 illustrates a highly simplified view of how one might come to the conclusion that vision is learned. First, one might believe that anything not genetically determined is learned. Then one would take what is known about learning from psychology, which for various understandable reasons has concentrated on basically the type of learning involved in conditioning in mature animals. *Voilà!* The result (I claim) is that one confronts the problem of acquisition of vision with a limited paradigm, inadequate mechanisms, and unrealistically high expectations for learning.

Visual capabilities develop in a cortex that is highly plastic and whose differentiation derives from experience as well as from the genome. Learning and other developmental processes can occur simultaneously. Fig. 2 is meant to illustrate my position. There is much that happens during development with results that are not by any means genetically determined and also are not determined by individual-specific experiences. Second, there are many sorts of learning and memory that do not get much attention in psychology (or anywhere else), and there is a possibility that such mechanisms to play an important part in acquiring vision.

## 3 What I Mean by Vision, Learning, and Development

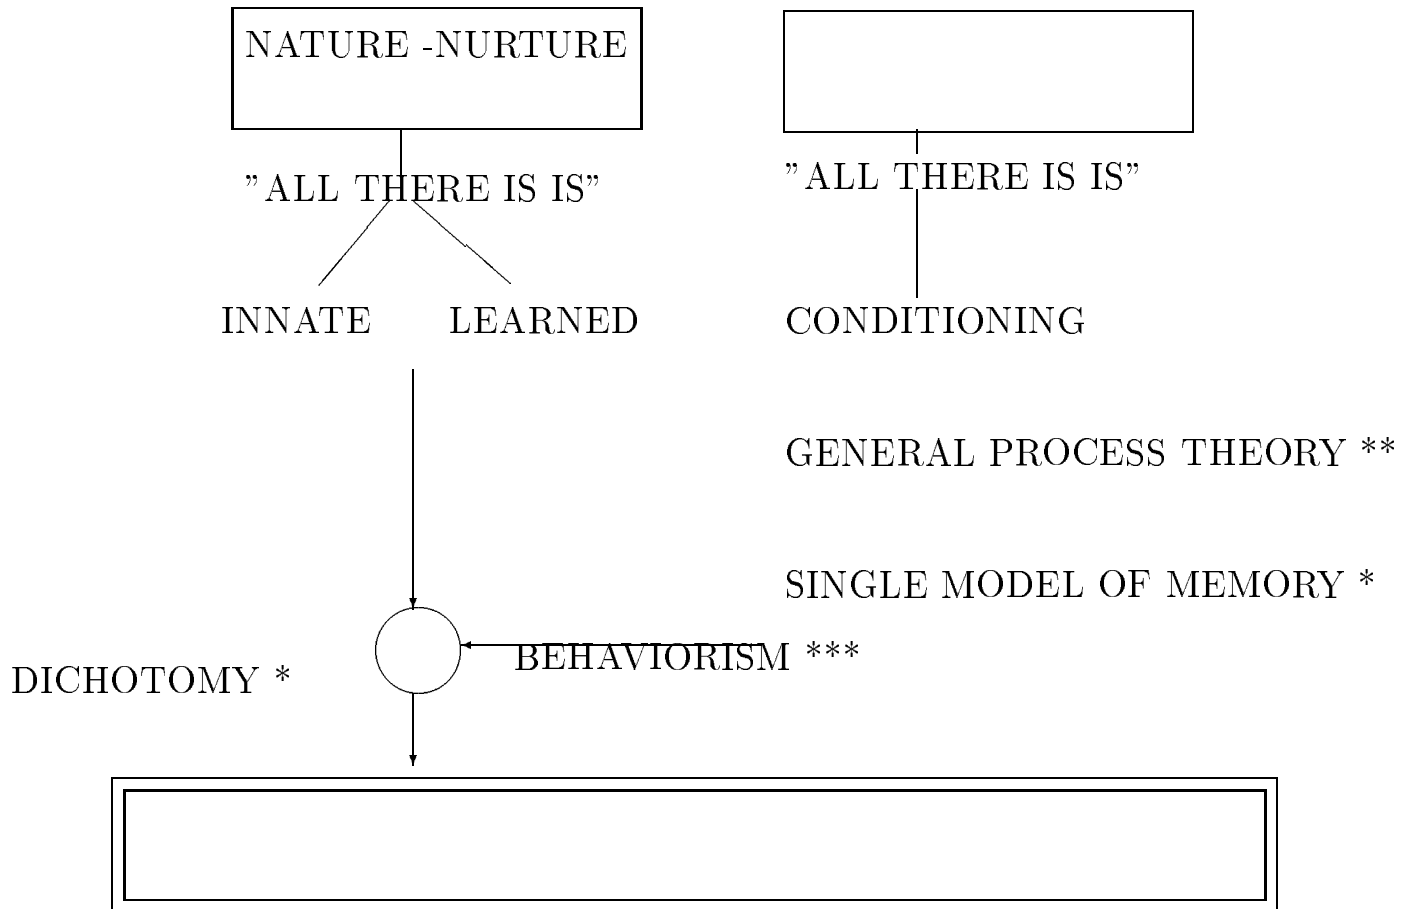
### 3.1 Vision

By vision I mean capabilities like those in the following incomplete list.

- “Segmentation” (flexible, salience-driven, task-dependent grouping; figure-ground; “object-hood” recognition).
- Invariances and constancies (color, perspective, occlusion, lighting).
- Stereo fusion.

---

<sup>1</sup>Numbered notes are found at the end of the paper.



EQUIPOTENTIALITY \*\*\*  
 \*\* STATE OF ART, RIGHT OR WRONG

\*\*\* NOBODY BELIEVES THIS ASSUMPTION \*\*

"LOTS OF VISION IS LEARNED LIKE ASSOCIATIONS"

\* PSYCHOLOGY KNOWS BETTER – DO WE?

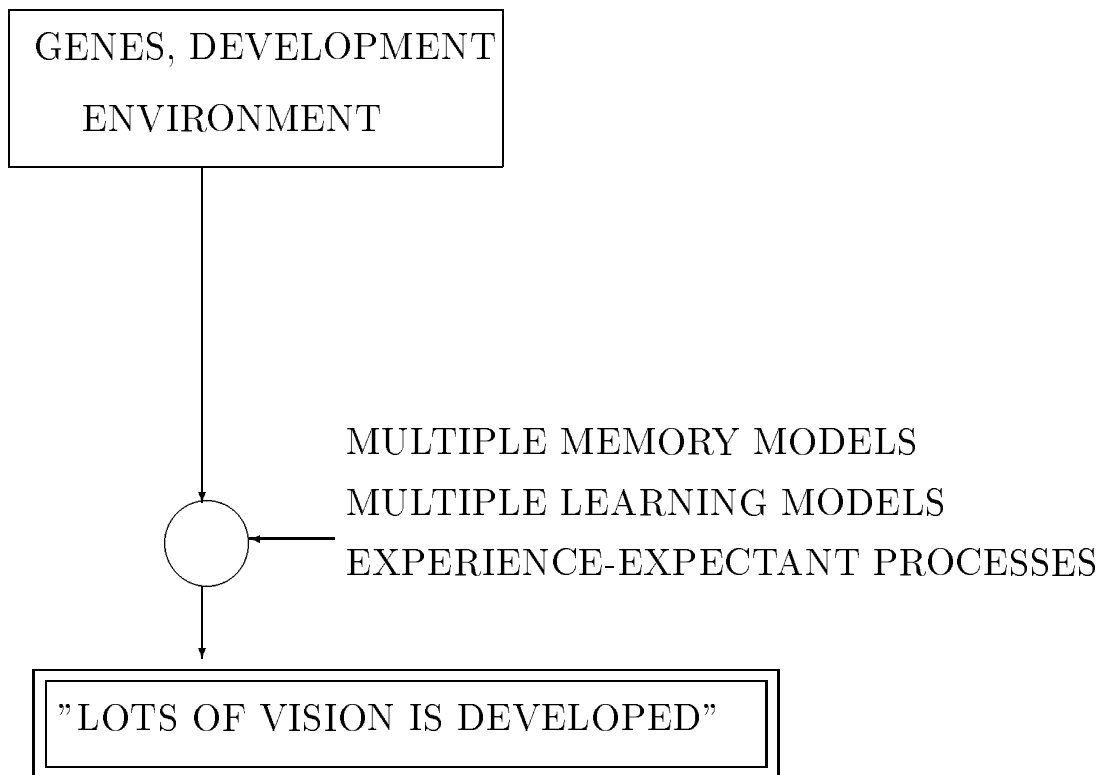


Figure 2: Phenomena and a claim that are not currently mainstream concerns in computer vision research.

- Integration of visual cues (for object recognition, say).
- Visuo-motor “reflexes” and controls.
- Object recognition.
- Reconstruction (depth, egomotion, scene parameters).
- Examination of the percept (overcoming constancies).
- Motion field processing.
- Hand-eye coordination, navigation.
- Building relational descriptions.

## 3.2 Learning

It seems ill-advised to attempt a definition of learning in animals, since it is a complicated proposition. A good start at defining learning in animals is [52]. Broad definitions such as “intelligent systems that learn – that is, that improve their performance as the result of experience” [75] include many developmental and purely maturational phenomena that are not usually called learning.

However, one rather striking difference between what is usually called learning and the sort of developmental phenomena I have in mind is the difference between *experience-expectant* processes and *experience-dependent* processes [28]. The former depends on experience that tend to be species-specific, ubiquitous, and could have played an evolutionary role. Cognitive mechanisms that result from interaction of the animal (sometimes even before birth or hatching) to non-specific aspects of its species-typical environment are sometimes called *primal* (not *innate*, since they need certain sorts of experience). Experience-dependent plasticity stores information that is individual-specific, and it is this sort that is involved in what is usually called learning (2).

For this report I am stipulating that with a few exceptions the AI community has successfully modeled human and animal learning (i.e. experience-dependent learning). There are journals and an annual conference for machine learning; in fact there are several non-overlapping AI communities doing several sorts of learning. My feeling is that we computer visionaries, like psychologists, generally think of experience-dependent plasticity when we think of learning, and indeed we mainly think of conditioning in mature organisms, but the word is used rather widely in AI literature. The following is an (overlapping and likely incomplete) list of categories for computer learning as related to computer vision.

Randal Nelson points out a basic dichotomy in learning algorithms which it is well to keep in mind because I am only interested in half of it. The issue is whether the learning is based on new information acquired from outside the system. In the learning that interests me, it is. The other type of learning involves optimization of functions, reorganization of a data set, and so forth with no need for extra “experience”.

- Numerical Optimization. The analogy here is that the algorithm learns the best solution. It is not clear how to relate this to learning vision since it assumes a parametric problem statement.
- Memorization. Computers are good at rote memorization. Let us ignore the fact that some people may be reluctant to equate rote memorization with learning. One theory of vision involves (or simply is) matching into a database of object “views” (e.g. [37], among many others). If rote memorization can solve the vision problem then learning is indeed used, and an experience-dependent and simple form of learning it is, too. However, many of the

technical details of the solution need to be filled in, such as technical solutions to problems like segmentation from background, obscuration, and other invariances. Are these solutions learned to be learned using another technique? The whole history of computer vision is the search for invariances that can overcome the brittleness of a template-matching approach.

- Interpolation and (rarely) Extrapolation. Related to memorization is the idea of building an interpolated lookup table from input data. This idea has been dignified with various impressive names (such as radial basis functions), and modernized by calling it a neural net technique, and oversold by calling interpolation “generalization”. Extrapolation is a rarer and much more interesting (and challenging) kind of generalization [76].
- Adaptive methods. Simple gradient-space methods have been used for a long time in areas such as adaptive control, where the “MIT rule” (basically gradient search) for gain adjustment anticipated today’s rules for updating connectionist unit activation levels. In psychology, adaptation is a short-term, reversible, ongoing activity (think of pupillary diameter, for instance), whereas learning denotes a longer time-scale, often irreversible process. Low-level adaptive thresholding and other parameter adjustment is a commonplace in working vision programs, but in the right hands can be used to recover visual information [46, 91].
- Reinforcement Learning. This symbolic technique for solving credit assignment problems claims (possibly optimistically) a meaningful relationship with dynamic programming [88]. It is much in the news since it attacks a fundamental problem. It could be useful for learning vision control structures. Its inherent computational disadvantages (only small action and state spaces are practical) can be ameliorated by mixing its learning techniques with other model-building techniques. An elegant example is a hybrid reinforcement learning, neural net and explanation-based scheme for robot control [57].
- Parameter and parameter range learning in algorithms and prior structures. This is a useful technique, in which parameter values or ranges, probabilities (prior or conditional), statistics, etc. are computed from experience and are saved in a data structure (e.g. a decision tree, neural net, Markov random field, object description template, relational structure with properties and relations) that is used for recognition.
- Classification (Supervised) Regularity Discovery (Unsupervised). Clustering techniques have been a mainstay in computer vision ever since statistical pattern recognition was applied to visual data in the 1950’s. They recently reemerged in the form of neural net implementations. In the latter guise they take the form of parameter learning in prior structures, accomplished by gradient descent techniques (implemented with the backpropagation algorithm, simulated or mean-field annealing, or other optimization techniques.)
- Connectionist learning, pattern association, function approximation, self-organization, competitive learning, hetero-synaptic learning. One can look at this huge body of work as instances of parameter learning for prior structures, classification, or function approximation. Many people nowadays only know the classical topics through their “neural” net instantiations, and there are some new wrinkles. Kohonen nets, for example, establish nonlinear mappings between features; the computational model was perhaps inspired by biological findings about cortical self-organization (3).
- Genetic Algorithms. A problem solution can be represented by a string, and then the string can be modified to search the solution space. Modifications are inspired by simplified versions of DNA mutation, splitting, and combination, with the hope of propagating good sub-solutions and evolving good solutions (4).

- Symbolic (Concept, Decision tree, Explanation-based, Case-based) learning. This is a large field dedicated to rearranging and digesting symbolic input into symbolic data structures. It deals with higher-level representations but there are many techniques in common with other learning methods, including clustering and rote memorization. These techniques, or rather techniques at this level, have not been much applied to perceptual domains.
- Learning Theory. Rigorous mathematical models of what can be learned have begun to emerge, often within the context of what formal languages can be learned. A very simple sort of question would be: Can a particular language be learned with only positive instances? A recent reference for this field of work is [87]. This is a fairly esoteric field whose relevance to vision is little explored.

In summary, many aspects of high- and low-level experience-dependent learning and adaptation have been modeled in AI. If learning was directly applicable to learning vision then I am stipulating that some existing computational method should just be rolled out. On the other hand, if learning is directly applicable to learning vision and a new learning paradigm is needed, then that would be interesting in itself.

Perhaps it is apposite to note that gaps in computational learning applications appear in (at least) motor skill learning and language learning. Perhaps for the same reason that (I'm claiming here) learning is not applicable to acquiring sight, there is not much serious work on computational models of natural language acquisition. Learning motor skills is another mysterious process. So far motor control learning been attacked with limited success as a parameter-learning task for a set of pre-wired basis functions governing joint angle velocities [77] and more successfully by memorization and table interpolation and extrapolation techniques that look more like conditioning than anything else [59, 76, 78].

### 3.3 Development

Finally, by development I simply mean all changes to an individual that happen from fertilization to maturity. One simple view of the world is that development consists of “programmed in” (genetic) capabilities and “learned” capabilities. If this is true, then it would seem there is no new problem in learning to see, only perhaps technical ones – we computer scientists certainly know how to program things in, and I have stipulated that we know how to do computational learning. However, things do not seem that simple, and this paper is mostly to argue for the possible importance of experience-expectant processes and other non-learning phenomena in the acquisition of vision.

### 3.4 Learning vs. Development

My purpose is not to quibble over vocabulary, it is to establish that the current metaphors, techniques, and mathematics of learning may not be adequate to “learn” vision.

The main difference between learning and development is probably not time-scale. If one accepts that learning is Hebbian, then one relevant time constant is perhaps the synaptic firing rate, which is on the order of milliseconds. Development seems to take longer, on the order of months perhaps. But recall that learning (certainly as simulated by back-propagation in multi-layer perceptrons, and often in real life as well) is not instantaneous. Even fact learning can require repetitions, and skills must not only be learned over years but skills must be practiced continuously. Developmental progress, on the other hand, can be amazingly rapid (once the stage is set) and need not be practiced once attained.

Certainly one thing that separates “learning” from “development” is the idea of parameter-setting in a fixed architecture versus (if you like) parameter-setting in an architectural space. We

perhaps need new computational models to address the latter question, although some good work has been done in using existing learning-like models to guide wiring operations (Section 6). Taking this idea seriously would mean coming up with a brain-wiring model that can explain the various data presented in Section 5. Strengthening the argument are the data that may indicate everyday neural plasticity in mature individuals (Section 5.4). If plasticity accompanies everyday activities, this is all the more reason for thinking about mathematical models of plasticity and growth to supplement the existing models of learning.

## 4 Beyond Conditioning

First I'd like to make the case that there are several types learning, and that not all learning can be explained by simple paradigms based on conditioning.

Behaviorism seems to be fairly thoroughly consigned to the past, but it had some influence over learning theory that it might be as well to make explicit, since the conditioning paradigm is rather close in some ways to some current learning paradigms in that one uniform mechanism is taken to be responsible for a vast range of phenomena (e.g. [74]).

In classical conditioning: E1 is a stimulus from experimenter (bell, light), E2 is a reinforcer (food, shock). Conditioning makes the subject pair these two stimuli (respond similarly to either). In instrumental conditioning, E1 is a behavioral response from animal, E2 a reinforcer; the subject learns to produce the behavior in order to obtain (or avoid) the reinforcer. The view that all learning is conditioning leads to several consequences, including the possibility of general laws of learning (e.g. general process theory) and predictions about what can be learned (e.g. equipotentiality). Roper [72] examines the case for both general process theory and equipotentiality theory, and the next two sections are taken from his chapter.

### 4.1 General Process Theory

General process theory holds that all instances of associative learning involve the same basic underlying mechanism or process, which leads directly to the hope of some general "laws of learning". Although such laws are not ever presented together as such, the following sorts of laws have been proposed and defended.

- Associative learning only occurs when the time interval between E1 and E2 is short.
- In classical conditioning, speed or strength of learning increases with intensity of E1 (within limits).
- In associative learning, speed or strength of learning increases with size of reinforcing event E2 (within limits).
- If the reinforcer is withheld, learned response declines in frequency or intensity.
- When E1 is a compound stimulus whose elements differ in intensity, conditioning occurs only to the more intense one (overshadowing).
- Prior conditioning of one element of a compound stimulus prevents subsequent conditioning of other elements (blocking).

One of the cornerstone laws of learning, the first in the list above, was removed by the work of Garcia [24] on food aversion learning, in which a rat learns to avoid food that makes it sick, even

with a delay of up to twelve hours between the rat’s ingesting food and the subsequent illness effect. These results have been duplicated under various stringent control regimes and appear solid [19]. The rest of the above “laws” have not been severely damaged by psychological experiments. The highly stylized methods of these experiments possibly have not tested the theory across a broad spectrum of conditions. Here the difficulty of reductionist approaches is in evidence: rats in cages may not demonstrate general principles of learning. So general process theory is not heavily damaged by the evidence so far, and different theoretical approaches to laws of learning and conditioning are still being followed (5).

## 4.2 Equipotentiality

Equipotentiality, or the claim that all pairs of events E1 and E2 can be associated with equal ease, in any species, would be a compelling consequence of a simple conditioning learning model. However, equipotentiality does not exist; in fact, learning abilities are nonuniform and constrained within and across species.

An elaboration of Garcia’s flavor-illness delay experiment leads to the flavor-illness specificity effect, which is a good example of constraints on learning. Rats can learn to associate taste with illness but not taste with electric shock. Contrariwise, they can associate flashing lights and noise with shock, but not flashing lights and noise with illness [25]. These “tasty-water” and illness, “bright-noisy water” and shock results have been extensively and carefully reduplicated [19].

It can be shown that differential learning of particular cue-reinforcer pairs happens in instrumental as well as classical conditioning. In other words some jobs are harder to condition either classically or instrumentally. For example, it is hard to train pigeons to peck to avoid a shock or to train rats to approach a stimulus to avoid shock; it seems that in “lower” species (e.g. rats and pigeons) it is easier to condition responses that are like normal behavior than those that aren’t (6).

Some responses cannot be conditioned; in particular, face-washing in hamsters and scratching in rats [81, 60]. The difficulty of conditioning on involuntary scratching is an interesting story, which however raises the basic question of deficits of learning versus deficits of performance.

Further evidence is provided by ethological studies. Of course ethologists have always known that there are huge inter-species differences in learning behavior. Chicks will not imprint on just any stimulus; herring gulls can recognize their chicks but not their eggs. Surprisingly accurate and long-lasting learning happens in several species for species-specific goals. One may cite salmon returning to their birth streams many years later to spawn, birds navigating, and bees learning routes to food. In every case one can make a species-specific argument for why it is useful for the animals to learn these particular things.

Birdsong learning has a large and fascinating literature. The learning capabilities vary vastly between species. Usually birds only learn songs constrained by their species. Within those limits there are different learning mechanisms for the songs that a bird will sing (learned early, not forgotten, not modified or adapted) and for those that it will not sing but must recognize, like those of its changing but competing neighbors (learned throughout life, adaptable) (7).

Equipotentiality is a corollary to a theory of one simple, uniform learning mechanism. If equipotentiality existed, it would lend weight to such a theory, but it doesn’t.

## 4.3 Types of Memory

If learning changes memory with experience, then the existence of different types of memory should make different types of learning seem more likely. Janowsky [39] does a thorough job of reviewing studies of memory differences, and presents the following taxonomy.

- Declarative
  - Semantic (facts)
  - Episodic (events)
- Nondeclarative
  - Skills (Motor, Language, Perceptual, Cognitive)
  - Priming (Perceptual, Semantic)
  - Dispositions (Classical and Operant conditioning)
  - Nonassociative (Habituation, Sensitization)

It is interesting to me that language and visual and perceptual memories are grouped under “skills”. One could easily imagine a richer taxonomy.

Song learning, food caching, and imprinting can be reliably studied in the laboratory, and much of the results in memory systems come from them. Sherry and Schacter [80] define multiple memory systems as those that have different operational principles. They then take an evolutionary approach to demonstrate the likelihood of multiple memory systems. Their mainstay is *functional incompatibility*, which implies that it often makes sense to have different implementations of a system so that it can operate for different purposes or operating ranges. The concept of *exaptation* is related to and important in this argument. It is strongly associated with the writings of S. J. Gould, for example, and it states that existing biological structures can be taken over for different purposes. A prime example is wings and feathers; it is hard to argue that wings and feathers are an adaptation “for flying” – partly-evolved wings or feathers do not produce advantages conferred by flying.

The evolutionary argument goes as follows. Memory capabilities are heritable: (*dunce* mutation in *Drosophila* cannot pair shock and odor, but normal flies can.) Further, memory helps in survival: the results of animal contests, success of foraging, and cooperation and altruism between animals are important to remember. Last, functional incompatibility gives a reason that one mechanism, such as “memory”, should be adapted or exapted to entirely different functions. Evolution often produces multiple systems with related functions (e.g. batwing circulation, which uses blood for heat regulation as well as to carry metabolic products). In birds, song learning vs. food caching shows two opposed uses for the memory system. In caching (in marsh tits and chickadees, say) learning goes on throughout life with no sensitive periods, there is large “generality” (a widely varying set of site types can be remembered), as well as perfect memory of filled caches and forgetting of empty ones [82]. The learning of songs to be sung, on the other hand, has sensitive periods, is usually restrictive in what can be learned, and the memory is not adaptable or forgotten. In fact generally certain sorts of learning (e.g. imprinting, salmon stream memory) take place over restricted timescales and are not forgotten. Other sorts (e.g. songs not sung, foraging successes) take place throughout life and exhibit both forgetting and sometimes impressive generality. Thus the conclusion is that different, functionally incompatible memory systems are useful for survival, and that they may have evolved by exaptation.

In humans, it seems to me that there are indeed several sorts of memory, but I do not find Sherry and Schachter’s arguments compelling as to why declarative and nondeclarative systems are functionally incompatible. They say that skill learning generalizes over situational, episodic details but that fact learning keeps them: but neither skill nor fact learning should remember episodes. (Probably the merging of fact with episodic memory is a mistake). Skills and facts are both “forgotten” in their own ways unless continuously rehearsed. For me the different learning techniques for facts and motor skills are more compelling evidence of qualitative differences (8).

There is anatomical evidence for different structures supporting gradual, incremental, “nondeclarative” learning (as in skill learning) and rapid, declarative learning (as in fact learning). Amnesiacs

often fail at declarative learning, but can acquire and perform nondeclarative skills [56]. Hippocampus, amygdala, and other limbic structures in the medial temporal region are involved in declarative, but not non-declarative, memory [86].

#### 4.4 Conclusions

The consensus in the psychological community seems to be the following.

- Behaviorism is wrong, and the conditioning paradigm is limited.
- Equipotentiality theory is wrong.
- General laws, or a real theory, for associational learning are still being sought.
- Associative memory is one limited type of many.
- Hebbian, associative learning is one limited type of many.
- The paradigms for studying memory (rat labs, ethologists, developmental vs. psychological...) are all limited in their own ways and the communities do not communicate much one with the other.

In this section I have tried to review some evidence that there may be more than one type of learning, and memory, and to explore limitations of the conditioning paradigm of learning. The next section broadens the scope to include phenomena that are not even called learning, but which may have more to do with acquiring vision than learning does.

### 5 Beyond “Nature or Nurture”

Any binary distinction between nature and nurture (genetic determinism versus learning) is convincingly refuted by ample evidence of an entire range of processes that operate in more or less mysterious ways to structure both brain and behavior [48, 6, 42, 66].

Early on it was realized, though not universally, that heredity and environment were not “either-or” causes for any identifiable behavioral trait.

*The only reasonable answer is that the two proportions [of influence] are one hundred percent environment, one hundred per cent heredity. They are not additive; any bit of behavior whatever is fully dependent on each. [34]*

Later, the role of other processes in what was previously lumped under the “genetic” rubric was acknowledged by some workers.

*Now, it may be comforting, in the sense that it gives us the feeling that we have increased our understanding of the problem, to say that a behavior pattern (or a structure) is innate if it is... “encoded in the DNA”. ... I believe that the comfort and satisfaction gained from disposing of the problems of ontogenetic development by the use of such concepts are misleading, and are based upon the evasion or dismissal of the most difficult and interesting problems of development. ... The problems of ontogeny and differentiation of structures in complex organisms, however, have hardly been touched as yet by the recent massive advances in molecular biology. [48]*

Indeed several factors undeniably influence the development of an individual, including the following.

- Genetic – properties of fertilized ovum.
- Chemical, prenatal.
- Chemical, postnatal.
- Sensory, shared; pre- and post-natal experience normally inevitable for species members (or living things).
- Sensory, individual; member-variant experiences.
- Traumatic; “abnormal”, cell-destroying events.

The following observation unpacks some of these factors a bit. Its implications are thought-provoking if (as I’m claiming) vision capabilities are largely formed by developmental processes and one wishes to make a computational model.

*...Sharp lines cannot be drawn, in early stages of development, between: the effects on neural development of nonbiological conditions (temperature, light, chemical conditions in the environment); nonspecific effects of gross stimulus input; the developmental effects of passive practice forced during ontogeny; the developmental effects of practice resulting from spontaneous activity of the nervous system; links and integrations between behavioral elements, resulting from early, nonfunctional partial performances; interoceptive conditioning resulting from inevitable tissue changes and metabolic activities; simple conditioning to stimulation resulting from spontaneous movements; and simple instance of conventional conditioning and learning. [48]*

The rest of this section is devoted to development and neural plasticity, mostly in an anecdotal way and at varying levels of detail. My goal is to establish non-genetic, non-learning factors as important, and to make a case that vision does not need to be learned.

## 5.1 Process of Development

Developmental processes can have dramatic effects. Fig. 3 shows ants who are identical genetically but whose form and whose role in life have been manipulated in the nest by different treatments when larvae.

Eye-color can be a developmental characteristic (Mutant chocolate-eyed *Crustacean Gammarus* eye color is normal if reared in lower temperatures.) The ability to fly is influenced by developmental conditions; wings of *Drosophila* reared in cold conditions are unable to support flight. One would naively consider binocularity or monocularity as something that was purely determined by genetics. Wrong: varying the amount of magnesium chloride leads the fish *Fundulus heteroclitus* to develop with one functioning eye as opposed to the normal two [71]. Is the number and location of eyes an adaptive characteristic, and is information about such number and location in the genome or in its relationship with its chemical environment?

The mechanisms of development are hard to understand but at least one debate seems to have been sorted out [65]. Preformationist theories (of the seventeenth century, say) held that a completely formed being exists in either the egg or the sperm, and the other partner’s contribution diffuses

Figure 3: Morphological effects of larval development in ants. Social roles and behavior of course vary widely as well [6]

into this being. *Neopreformationism* explains cellular differentiation by the mechanism that the full complement of hereditary material present after conception is reduced during development, so that each cell type has a specific subset of genes, which determines its type. The alternative to neopreformationism is *neopigenesis*, which holds that all cell at all stages have a full complement of genetic material which is *expressed* at different times, or *sensitive periods*. This latter view has been upheld by many experiments in which complete embryos develop from nuclei transplanted from individuals after differentiation is complete [29] (Fig. 4).

Further, if the two cells resulting from the first division of a fertilized egg are separated into separate cells (blastomeres), either suffices to create a complete individual [17] (Fig. 5).

Developmental phenomena exhibit *sensitive periods*, which are times during development that the right sort of experience produces developmental results, and that deprivation of experience can cause deficits (e.g. Fig. 6). Several good examples can be found in [72, 6, 41]; most humans find language learning (especially accents) harder after puberty, for instance. Social behavior, love, and sexual behavior are strongly influenced in monkeys by early deprivation [32, 31].

The effects of experience in development can have any of several effects.

- Maintain existing capability
- Facilitate its development
- Induce one that only develops with experience
- Attune, or materially improve its performance

Fig. 7 illustrates these ideas graphically.

Developmental phenomena are highly interactive, and take place in a predictable course in which one effect initiates or allows the next; developmental effects “set the stage” for later ones, sometimes in trivial ways, sometimes in non-trivial ways.

Figure 4: Nuclear transplantation into amphibian eggs [29]. The nucleus of cells from embryos and tadpoles at different stages of development when placed into egg cells support the development of normal adult frogs.

Figure 5: Cells from the two-cell embryo are adaptable. The first two blastomeres are separated and re-grown in (a) echinoderm, (b) amphibian, (c) mammal [17].

Figure 6: Sensitive period in song-learning. (a) shows a sample of songs learned by a single marsh wren from a tape-recording tutor. (b) shows data for nine birds exposed to a total of 44 sing types, each for only a few days. Frequency of song-learning is taken to indicate sensitivity [45].

Figure 7: Different effects of experience in development [4].

Some behavior seems as if it is learned but is not – it results from maturation or development. A male puppy injected with male hormones will promptly (without being shown) begin to cock its hind leg instead of squatting to urinate. Nestlings behave in a way that looks like they are learning to fly by practicing flying movements, but when this behavior is prevented they fly on schedule anyway.

Sometimes developmental factors seem quite unrelated to the developmental outcome. Activity in chicks results in testosterone production that is essential for imprinting – imprinting experiments therefore are conducted with chicks on running wheels. A mallard duckling needs the prenatal, embryonic “contact-contentment” calls of its siblings (i.e. the embryo in the egg needs to hear from the surrounding eggs) in order to develop postnatal preferences for the maternal calls of its own species later (e.g. [27]).

Chick activity and imprinting must of course take place during proper sensitive periods [41], and several fascinating imprinting studies have revealed intriguing interactions between experience-dependent and experience-expectant processes. For instance, consider two groups of chicks, one exposed to a rotating red box (which will cause imprinting) and the other to a rotating hen. Two hours after training each group preferred the trained stimulus (experience-dependent learning), but 24 hours after training both showed increased preference for the hen (experience-expectant ‘learning’). Again, chicks exercised on running wheels had no preference for either test object two hours after simply being exposed to light but 24 hours after exposure to light had a strong preference for the stuffed hen (innate or primal process). Somewhat surprisingly, species-specific considerations may not play a role in chicks’ preference; hen’s heads and duck’s heads are equally preferred, and there is no preference between an entire (stuffed) hen and an entire (stuffed) polecat.

At birth, development has proceeded to different stages in similar species. For instance, chicks and pheasants are *precocial*, in that they are moving, seeing, and pecking immediately after emerging from the shell. Robins are *altricial*, in that they develop more outside the shell. Mammals vary from precocial species like lambs to altricial ones like rats. Even among ungulates there are altricial “hidiers”, like deer, whose young stay in a nest for some weeks, and precocial “runners”, like bison, whose young are up and around within minutes. Humans rate as fairly precocial visually. Prairie dogs are interesting, in that they must cope with a visual world suddenly after natural “dark-rearing”. They are raised in their burrows until they can move around well and then emerge into the light. A few days later they are chased away by their mothers, to find an abandoned burrow or die. But as we see below, and as we should expect if developmental processes are more important than learning, artificially dark-reared kittens and monkeys have surprising visual powers.

## 5.2 Brain Development

Development of the brain is marked by an initial overproduction of neurons, some thirty percent of which (in rodents) quickly die off through time and experience [15, 64]. Not just neurons, but synapses are initially overproduced both in the muscles and the central nervous system. These too are tuned, or specialized, by elimination [68].

Cortical neurons are born in the interior of the developing brain and migrate outwards, using amoeboid processes, to increasingly outer cortical layers. The familiar cortical areas, which seem fairly uniform between individuals, emerge from a developing cortex that differs greatly from the final product, being populated by transient structures and by being much less differentiated (Fig. 8). It seems the entire neuroepithelium can generate most if not all of the classes of cortical neurons, both transient and permanent. Boundaries between cortical areas are not fixed, as can be shown by studies with visual areas and sensory areas generally. Not only that, but transplantee pieces of “protocortex” take up the area-specific properties (projections, axon growth) common to their new location. Thus long after neurons are formed their roles are being determined by a large set of developmental factors.

Figure 8: Origin, differentiation, migration of neurons and glia [14].

## SUMMARY

Age	Functional Anatomy	Behavior
Newborn	SC Pathway, +layer 5 and 6 pyramidal output to LGN and SC	Saccadic pursuit tracking Perferential orienting to temporal visual field “Externality Effect”
One Month	As above + inhibitory pathway to SC via BG	As above + “obligatory attention”.
Two Months	As above + MT (magnocellular) pathway to SC	Onset of smooth pursuit tracking and increased sensitivity to nasal visual field
Three Months and over	As above + FEF (parvocellular) pathway to SC and BS	Increase in anticipatory tracking and sequential scanning

Externality effect: don't attend to small patterns within big ones. SC, superior colliculus, SN substantia nigra, BG basal ganglia, LGN lateral geniculate nucleus, BS brainstem, FEF frontal eye fields, MT middle temporal area, VI primary visual cortex.

Figure 9: Summary of suggested correlations between development and behavior [40].

Visual acuity in human newborns is low, at about 20/600, and contrast sensitivity also is poor; both improve significantly over the first six months. By three months, infants can distinguish patterns on the basis of their shape and form [21]. In normal early development, a (speculative) emerging “wiring diagram” can be (speculatively) related to anatomical structure and even behavior [40]. Figs. 9 – 12 show Johnson's theory of how these pieces might fit together.

### 5.3 Deprivation Experiments

Almost everyone has heard of the famous and elegant sort of visual deprivation experiments, usually with kittens, in which visual experience is constrained in some way, with resultant behavioral, neuroanatomical and neurophysiological effects [35, 8]. This work is much too voluminous to cite here, but a start would be [89, 28], and references given therein. This work is often taken by laymen to imply a learning model for vision acquisition, but that is not the correct message.

Kittens raised in visual environments lacking certain edge orientations, for instance, show a reduced number of neurons sensitive to that orientation [35, 8] (Fig. 13). The mechanisms at work are competition and cooperation, not atrophy. It is not lack of visual experience that leads to differential orientation sensitivity, it is that experience leads to related functionality taking over the available brain resources. Monocular deprivation experiments have the most dramatic effects: after only six hours of monocular deprivation, kittens show effects of loss of cortical binocularity. If they are then given normal binocular experience during the sensitive period for binocular development, they develop normally.

Figure 10: Suggested oculomotor pathways in the human newborn [40].

Figure 11: Suggested oculomotor pathways in the human, age one month [40].

Figure 12: Suggested oculomotor pathways in the human, age two months [40].

Figure 13: The Blakemore and Cooper [8] apparatus and representative findings. The lower diagrams show the orientations of stripes to which units of the kittens reared in horizontal (left) and vertical (right) stripes are sensitive.

More evidence that the orientation sensitivities arise through competition: dark-reared kittens (not dark-reared too long) are able to see almost immediately upon being brought into the light. How do dark-reared cats develop binocularity without outside experience? Through spontaneous neural activity: The extended sensitive period presumably helps. In general, the visual cortex of dark reared kittens develops normally, but for some fraction of unresponsive cells [36, 9, 79, 83, 12].

The visual accomplishments of dark-reared cats are surprising if we think cats learn to see. On the other hand, regaining binocularity or even exhibiting near-normal vision are examples of the many developmental deficits that can be overcome with no bad effects (short-term nutritional deficits are another example). As long as the critical period is still in force, “normal” experience at least partially overcomes earlier deficits like those caused by visual deprivation. After the critical period, damage is much more permanent. One ameliorating effect, consistent with the competitive nature of this sort of development, is that dark rearing can extend sensitive periods [16]. Interestingly, there does not seem to be evidence for cat-like orientational sensitivity in monkeys [90]. Dark rearing only degrades monkey spatial vision an octave or so. Spatial vision development correlates with gestational time, not postnatal time (i.e. not visual experience) (Boothe, in [3].)

I'll say more about cooperation in Section 5.4, but one argument for the sort of plasticity that leads to orientational specificity is that it allows calibration that leads to sharper stereo vision. Plasticity is needed to adjust to individual variations and to variations due to growth and maturation. Dramatic evidence for such plasticity exists in frogs, in which surgically rotating an eye 180 degrees in its socket is overcome by plasticity that leads to points in the tectum being stimulated by points in equivalent positions in the two eyes (which get stimulation from the same world points) not opposite points in the two eyes (where the connection would have been in an un-modified frog) [44].

The deprivation experiments reveal a competitive mechanism for developing certain visual capabilities. To my mind the more interesting spin-off result is that dark-reared cats and monkeys exhibit effective vision very quickly upon being allowed into the visual world. These cases, along with precocial animals like chickens, bison, and lambs, are for me a strong argument against a necessary, central role for learning (via teaching, conscious experience, structured experience) in vision.

## 5.4 Plasticity and Cooperation

I have mentioned plasticity in the protocortex, and alluded to continuing plasticity in adults.

Zeki's recent book [93] provides good pointers into some of the literature on plasticity in mature individuals (monkeys), which has begun to quantify the physical limits to certain sorts of expansions due to competition for sensorially de-afferented cortex. In one case the nerves from the forelimbs were severed, and 12 years later it was found that the lower face representation area had expanded between 10 and 14 mm into the "vacated" area [67].

Some of the most dramatic demonstrations of brain plasticity in mature individuals is the remapping of cortical representations based on sensorimotor experience [53]. The method is either to train monkeys in sensori-motor tasks or to make neuro-vascular skin transfers or to use finger casts. The pre- and post condition cortical skin-surface or movement representations are shown to change as a result (Figs 14 – 16).

The principle that seems to be operating here is that the co-occurrence of stimuli is responsible for determining the cortical map. This principle of forming connectivity is a different form of Hebbian change, since it involves several synapses at once: it is a *hetero-synaptic* learning rule. A prediction is that cortical representations of skin surface are continuously remodeled by experience, and the process goes on throughout life. Thus individuals going about their daily lives will have rather different cortical maps depending on their recent idiosyncratic activities, as in Fig. 17.

Merzenich's work argues that learning should be studied using dynamic behavior, and one should be on the look out for dynamic effects in the brain even in adults. These findings are rather a switch from the idea that a fixed brain architecture is laid down over a limited time in a fixed way by experience-expectant processes exposed to any of a broad range of stimuli. The hetero-synaptic learning rule and competitive models inspired the idea of "Neural Darwinism" ([22] and Section 6).

## 5.5 Conclusion: Vision as Development

Development encompasses many powerful and surprising mechanisms. The aspect of development that I concentrate on is that which takes Greenough's expected experiences (possibly very minimal, such as random spontaneous neural activity) and produces individuals surprisingly competent and seemingly sharing species-wide visual capabilities.

I am claiming that basic vision capabilities seem more like the result of experience-expectant development than experience-dependent learning.

**No reinforcement or supervision.** For example, by the time one starts teaching children to name the beasts of the field (say by pointing at a dog and saying "doggie"), one is assuming

Figure 14: After training, finger surface areas are mapped differently into the cortex [53].

Figure 15: After transplanting enervated skin, finger surface areas are remapped into the cortex as part of the new location [53].

Figure 16: Wearing finger casts causes the finger surface areas to be remapped in the cortex [53].

Figure 17: Individual variation in owl monkey cortical maps, possibly caused by differing recent individual experiences [53].

that lots of difficult segmentation and constancy problems are solved. One doesn't get a chance to stick baby's primal sketches on the refrigerator and critique them. When one teaches "object recognition" in this manner one is teaching a use of vision; it is too late to teach the vision process. I don't think Piaget argued that anyone "teaches" his various conservation laws, and it isn't clear that positive or negative reinforcements from nature would be enough to induce an infant to "get it right". Further, infants (in this culture) lead a pretty sheltered life, are protected from most forms of negative reinforcement, and do not have to do succeed at visual tasks to get positive reinforcement.

**Shared Visual Experiences.** Sensory experiences seem to immerse us, reliably and predictably, in the universe. Further, we seem to share our perceptual universe, in important ways, with our conspecifics (modulo various low-level physical or neurological limits imposed by peripheral and central differences like cataracts, myopia, strabismus, color-blindness, etc.). We can successfully assume that everyone sees the world in the same way, something like the way we can assume that everyone in our language community understands our language, (and nothing like the way we might assume that everyone shares our politics). Why should sense experiences be in common? Perhaps a better evolutionary solution would have been to allow individual perceptual differences the way that we have individual physical and mental differences. One can imagine a society that uses such differences to its advantage (blind mind-readers are a science fiction favorite). However, perceptual experiences certainly seem to be shared between individuals to a degree that would be surprising if we each made up our own rules for it. If we all do vision the same way despite a huge differences in early visual experience, traditional conditioning learning may not be the right metaphor (9)(10).

**Timing: The Emergence of Vision:** The development of vision happens when other early developmental processes do, unlike the formation of political opinions, learning to play the piano, or learning the capitals of the 50 states. Vision is not easy to modify; many experiments having to do with "learning to compensate" for image displacements are inconclusive and more careful ones reveal that some (perhaps linear) distortions can be compensated but that others cannot, and in general vision is not very adaptable [84]

**Vision Without Experience Is Common:** Several animals do some sort of useful vision almost purely without experience. "Visual cliff" experiments show depth-perception in the form of cliff avoidance as soon as an animal can see (1 day in goats, 4 weeks in rats), and independent of experience (dark-reared rats first exposed to light at 3 months) [26]. There are chickens and related species like pheasants, the naturally reared-in-the-dark prairie dogs, the ungulate "runners" such as bison and lambs, and there are artificially dark-reared cats. Probably the cats are the most surprising example since we seem to get vision almost without natural experience-expectant processes being driven even by randomly patterned light.

**Vision Without Much Experience Is Common:** Even in humans, relatively sophisticated visual abilities are attributable to infants:

*Yet infants of only 3-5 months do seem to know that a hidden object is still there when looking rather than reaching is the dependent measure [cites three references]. [18]*

This last quotation is from a paper in which Diamond makes experiments to demonstrate that hippocampal development in infants is more complete than previously thought; she does this, as referred to in the quote, by using a measure of infant "knowledge" based on looking, not reaching. To me this quote is interesting in two ways. First, the short time needed to evince this constancy knowledge makes it necessary to postulate a rather efficient learning mechanism if learning is heavily involved, especially since in early infancy the visual system is not very good. Second, a respected infant development researcher is talking about object-hood recognition in three month old infants. As I have said before, recognizing that an object **is** an object is a huge part of the vision problem, and that central capability is here blithely passed over without a comment; it is as if we take it for granted that infants have solved the object-hood problem at a very young age. Elizabeth Spelke makes very much the same point about the sophistication of vision in young infants:

Before human infants can reach for objects to manipulate them, they can already perceive objects as bounded, complete, and as persisting over occlusion, under certain conditions [84]

Does one learn to solve the object-hood problem, and if so how? Spelke's chapter considers this problem, and concludes that indeed infants are *not* segmenting a scene the way adults do, and that they might in fact "learn" (although the learning mechanism is mysterious) the usual Gestaltist grouping laws over time by observing properties of objects and non-objects.

**Cognitive Impenetrability:** Visual perceptions are cognitively impenetrable and not easy to modify, unlike things we learn associatively. True, to some extent some constancies can be attenuated. Overcoming constancies can happen by "attending to the percept", which is a fairly sophisticated and unusual operation (which I'd wager very few mammalian species can do). The goal is to perceive sense impressions in as raw a form as possible. Then one might notice that haystacks at certain times of day look purple, as Monet did. Or one might notice all the colored reflections on a surface that would be labeled simply "shiny black" in most everyday perceptual situations (and remembered as such). Some capabilities are harder to override: it is hard just to see color patches and not "things" or "objects", it is hard not to do stereo fusion with both eyes open and focussed, it is impossible to overcome simple optical illusions (Fig. 18).

## 6 Towards Developmental Models

*The most spectacular plasticity in the nervous system occurs during its development. The learning which so fascinates and mystifies us when we contemplate acquisition of general and episodic knowledge, motor and cognitive skills, language and reading, pales by comparison with the electrifying conundrum that is development* – T. Sejnowski [14].

This short section contains what I have found in the way of computational models for development as they relate or could relate to vision. My apologies to those not represented here through my ignorance. I hope alert readers can point me to other relevant work.

Belew considers that a computational model of ontogenic development can be used as a bridge between genetic algorithms and neural networks [7]. The idea is to generate neural nets using genetic algorithms, and the nets go on to learn by normal Hebbian principles. The genome in Belew's work is expressed by way of a context-sensitive grammar. The context sensitivity of gene expression can capture the basic properties of neopigenesis and of course timing can be built into the grammar via intermediate forms. This line of work is quite interesting; clearly the problem is somehow to get at the interaction of genome and environment. One thing real organisms have going for them is that all the necessary chemical and physics mechanisms come for free in the real universe, whereas they have to be simulated or made up from whole cloth in the world of the artificial.

Johnson's book on development has a defense of connectionism as a formalism for studying developmental-type calculations [5], and a PDP volume makes similar arguments [74]. If one were to study the generation of structure by some developmental-like mechanism, then LISP code or simple computing networks are obvious primitives.

A recent example of a system that "develops" from nothing, adding layers of network as it gains experience, is the Cresceptron [38]. It configures itself as a result of its experiences, adding neural planes as necessary (after learning to recognize 25 objects it has 4133 such planes). This is ambitious work that tries to develop the capabilities of segmentation, object recognition, and concept formation, all simultaneously. To my mind it is just the sort of research that needs to be carried out, but I feel that the tools available today are frustratingly weak for such work.

Figure 18: The Wundt-Herring illusion. Why not teach vision that all the horizontal lines are parallel?

Churchland and Sejnowski's excellent book [14] mentions probably the most successful case history of developmental modeling – at least the one that exploits Hebbian (connectionist) principles – Miller and Stryker's model of ocular dominance column development [54, 55]. The computer modeling is complemented by laboratory tests on kittens. The model encompasses the Lateral Geniculate Nucleus, which appears as two two-dimensional arrays of computing units (each 25 units square). As does Edelman's Darwin system (described next), the primary mechanism at work is that the correlation of inputs (here from the left and right eyes) is responsible for strengthening connections. However the model is fairly complex, incorporating versions of four biological parameters that can be experimentally measured in animals. The parameters describe initial connectivity, the correlation of activity in the right and left sides, the lateral influences whereby activity competes, and the total synaptic strength supported by each unit. The trans-cortical influence is interesting; the strengthening or growth of synapses at one point is influenced by activity somewhere else. The model incorporated some 61,000 synapses for the total of 1875 units in the model (the two arrays of 625 each and some other units).

In one experiment with the simulation, the parameters were set to those reasonable for kitten cortex, with random initial synaptic weights. Training this network with something like 200 trials resulted in adjacent stripe-like groups of simulated cortex acquiring ocular dominance – they are dominated by the same eye. The model could be tuned and compared to many known phenomena (e.g. deprivation experiments), and also used to explore what happens in cases that are impossible to duplicate in nature (e.g. different cell connectivity).

The Miller and Stryker work is among the best examples of biologically-motivated simulation of development. It raises two questions, however. First, Could we not have pre-programmed the result? Granted that plasticity is a nice engineering trick for accomplishing the wiring, is it necessary in a theoretical sense? Second, what can we say about development (or learning) beyond this low-level wiring based on concurrent inputs? E.g. once we have binocularity, where does segmentation come from?

The basic idea of Hebbian weight-modification is very attractive in its simplicity, generality, and seeming relation to biology. As a result there has been much recent work in the neural net community to show how certain neural structures like orientation columns, or certain phenomena like receptive fields (demonstrated by hidden units) or even algorithms (such as shape from shading [47]) can emerge from unsupervised exposure to visual stimuli [51, 92, 50, 50, 49, 62, 63, 58]. The annual proceedings of the Neural Information Processing Systems conference (e.g. [30]) is a reliable place to find new work in this area. Parameterized models of particular neural architectures that can be compared with biology are a good thing. One thing that is usually missing is a recipe for the autonomous *formation* of the architecture: once the connections are in place it seems that, to some degree, their weights can be adjusted so as to produce various phenomena that can be related to what we think we know about the brain. Whether the training times and regimes are biologically realistic is an issue, and so is the place of the results in a systematic description of how we learn to do vision.

Edelman's work on the Darwin architectures (e.g. [22]) does not attempt to reproduce parameterized aspects of biological systems, but tries to use self-organizing neural techniques to produce a (simulation of a) seeing, behaving system. The learning rules are at least partly inspired by Merzenich's work on stimulus co-occurrence as a signal for synapse development [53]. The Darwin architecture has a substantial and complex hierarchy of modules, as well as adaptive connections that are governed by the hetero-synaptic learning rule that uses temporal co-occurrence to signal associations (Fig. 19). The Darwin program learns to control a simulated 2-D jointed arm to grasp food and push away noxious pests. Although Edelman does not put forth Darwin as a model of development, it certainly is a model of computational plasticity and self-organization, and it is perhaps elegant that Edelman appeals to the same competitive principle to explain both brain plasticity and biological evolution.

Figure 19: Darwin III architecture [22].

Self-organization has of course been a vital topic since Turing and von Neumann, and is at the heart of development. People have noticed the parallels between ocular column development and the development of physical structure in propagating chemical reactions and random spatial processes. An interesting treatment of the self-organizing properties of dynamic systems with special emphasis on developmental aspects is given in [85].

Generally there is more to be done in the modeling and simulation of development to see what can be accomplished with what mechanisms. I should particularly like to see attention paid to something besides the lower-level visual capabilities (say spatial vision). For instance, linking the visual areas would be an interesting problem (11).

## 7 Conclusions

*When something hasn't been formulated rigorously, justification becomes a formidable task.* — S.J. Gould

In this rambling report I have tried to raise the issue of experience-expectant developmental processes as an alternative to experience-dependent learning as a major determiner of our basic visual powers. I have shown how, for quite understandable reasons, most learning research and knowledge is confined to what is learned in basically very simple conditioning situations by adult animals. I have alluded to the current view that there are many powerful processes that determine an individual's capabilities besides "nature" (the genome) and "nurture" (learning). I'm claiming that vision is not "innate", but that it is very likely "primal" in that it emerges approximately the same in all conspecifics who have some minimum and as yet undetermined amount of some undetermined set of experiences, visual and nonvisual. Learning paradigms based on association and synapse-strengthening do not seem adequate to explain the development of vision structures in the cortex, which however do arise partly as the result of experience, and which involve neural growth and death, synapse growth and death, and various sorts of plasticity that are not well understood or modeled.

Currently, we seem to have the following situation.

- Psychological Learning Research
  - Still based around conditioning paradigm.
  - Deals with adult or fully-functional animals.
  - Small number of species (overbred rats, cats, monkeys).
  - Limited range of tasks (ones doable in a cage).
  - Ethological (Natural) Learning Research –
- Psychological Learning Research – lots known about birdsong and imprinting.
- Development Research
  - Prenatal: large amount known.
  - Behavioral: small number of species, small range of behaviors.
- Computer Modeling of Learning, Development
  - Several learning models exist.
  - Models of development are rare to nonexistent.

- Computer architecture may impose constraints (e.g. memory addressing) on what we can model easily.

What recommendations for further research might emerge from these various observations (12)?

- Consider development as a possible inspiration for self-organization. Immediately we are faced with Marr’s Level I question: what needs to be computed? How do we start a computational model?
- Just what stimuli are needed, if any, to develop?
- What individual differences are there in vision? Certainly there are skilled vision tasks (e.g. photo interpretation) in which adult training makes a difference. Are there experience-dependent effects in our more basic visual capabilities?
- Just what primal visual capabilities **do** young animals (e.g. lambs) have? What capabilities do dark-reared kittens lose? Psychophysical facts about, say, “spatial vision” are fairly easy to come by, and vaguely interesting (I guess) but behavioral and cognitive visual capabilities are more what I have in mind.
- What role is played by the order of development, of “stage-setting”? For instance, Johnson’s pairing of wiring and capabilities in Figs. 9 to 12 could lend itself to a computational theory.
- What about neural cell diversity? Don’t we need a larger range of neural models? I can’t help but wonder why there are so many types of neuron in the brain and so few types in simulations.
- Non-associative learning: What are, and what enables, motor skill learning and language learning? What are good computational models for them and do these alternative learning models throw indirect light on development?
- Alternative memory models will be needed to go with alternative learning models.
- Consider self-organization of multi-level hierarchies; how to model the stage-setting phenomenon?
- How does the idea of continuous experience-dependent brain rewiring fit into a computational theory of vision?

## 8 Acknowledgments

Thanks to Profs. E. Newport, K. Nordeen, and E. Adams for inspiration and encouragement. Needless to say, they are not responsible for any of this. Dana Ballard and Randal Nelson critiqued an early draft. Virginia deSa helped with many references to the neural net literature. Oscar Firschein provided encouragement, criticism, and some good references.

## 9 Notes

1. I cannot verify that the Ramachandran reference states or implies the latter conclusion. As far as I can tell the citation refers to an article that does not contain the word “learn” nor any claim involving learning. Rather there are references to evolution and an observation about a connectionist simulation.

... we come to perceive solidity and depth. ...  
 ... the visual system seems indeed to assume that only one light source illuminates the entire image. This may be because our brains evolved in a solar system that has only one sun. ...  
 ... One reason may be that our primitive primate ancestors, who resembled tarsiers ...  
 ... [describes a connectionist implementation of shape from shading] Intriguing as this computer simulation is, its biological relevance is still not clear ... [70].

In a related article we find another reference to evolution but none to learning.

*We believe perception of apparent motion is controlled in the early stage of visual processing by what is in effect a bag of tricks, one the human visual system has acquired through natural selection during millions of years of evolution. [69].*

2. One way to think about experience-expectant processes is that they are just regular learning with some sort of “bias”. So for instance if I am a chick, shown a rotating red cube, who retains an imprinting preference for a chicken over a red cube, maybe I learned to imprint and the bias took care of what I ultimately imprinted on (cf. Section 5.1). The issues then become how to implement the biases technically, and then what are the rules for “biased learning” and how does it fit with the (experience-dependent) “learning” we know about.
3. I’ll restrain myself from citing several neural net references, but I will put in a plug for the *AI Encyclopedia*, which is a surprisingly good reference for AI topics (e.g. [2].)
4. Gradient descent, simulated annealing, and genetic techniques are all ways to search an “energy landscape” of solutions, with genetic techniques (GAs) having the advantage that they are the least well understood. With their zeroth-order model of DNA combination, GAs have nothing much to do with biology and nothing at all to do with evolution (Darwin knew, but perhaps someone forgot, that evolution does not maximize anything). GA success is particularly dependent on the energy landscape of the problem type and even problem instance, and has the same problems as other numerical optimization techniques in being applied to vision.
5. There have been several modern approaches to reformulating learning laws, which I have not followed up. Some names are Rescorla and Wagner, who seem to have several enhancements to classical theory. One that attracts me is that conditioning gets at underlying co-occurrence or correlation, with the aim of getting at causality. Thus one might predict that conditioning could take place with negative time displacements, not just positive ones: that temporal linking could be quite flexible. Another approach is that of Gallistel [23], whose representational theory grants many explicit representations for directionality, distance, time, etc. to even the simplest brains, and thus predicts a highly non-uniform neural architecture model.
6. In fact, pigeons show an interesting “autoshaping” of behavior that, interpreted in an extreme way, could indicate that they are not experiencing instrumental conditioning at all, only classical conditioning. It turns out that pigeons peck a button differently depending on whether the reinforcement is food or water. They have characteristic ways of pecking at food and drinking water (eyes open or closed, tongue position, etc.) and they assume these types of pecks for the associated reward. Further if the reward is sexual they make sexual displays to the button. The micro-behavior thus can be interpreted as “mistaking” the button for food, water, or a potential sex partner. In general this autoshaping of behavior points out differential ease of training and inter-species differences for acquiring arbitrary instrumental responses [11, 33].

7. I cannot resist this story which argues that in one species the sung song is learned in a highly unstructured, “equipotential” way. The marsh warbler *Acrocephalus palustris* learns all its songs in its first autumn and winter and thereafter does not change its songs. However, there is a small problem: during this time of their lives the birds are migrating from Europe to East Africa, and have no adult warblers to learn from. So they simply learn sounds as they go. As a result the average marsh warbler’s song incorporates an average of 76 other birdsongs, and the total number of species copied in marsh warbler songs is over 212 [20]. In fact, to the extent that subsongs are not identified, they can be taken as evidence for the existence of as-yet-uncataloged African species.
8. There is still ample room to wriggle, and to some extent these systems can cover for each other in humans. E.g. it is certainly not true that facts are always grasped at once. Grandmaster chess players are often very finicky about what chess pieces look like when they are competing. I learned to spin a top by lots of practice; my brother (the physics major) looked at the arrangement of string and top, thought, and succeeded first time. Walking seems to be learned by humans and not by lambs but it isn’t clear that rug rats need much practice to become tots. Bike-riding is interesting; it has aspects of a practiced motor skill but also aspects of an “aha” experience or sudden fact-learning. Unlike many skills it is not “forgotten”, it can likely be learned throughout life, and doesn’t get much better with practice (except for competition metrics). Genius-level talents involving chess, music, and mathematics might involve new sorts of learning and memory, but a simpler explanation is that these talents are exaptations of language systems. As we know, there is plenty of evidence for inter-individual variability in these abilities. Last, it is perhaps ironic that the best mechanisms we have now for motor learning look like the most straightforward conditioning paradigms [59, 76, 78].
9. Individual differences must exist throughout the visual system, and I am under the impression that their study is of interest to psychophysics. Individual differences may provide some fascinating data with which to assess the limits of innate, primal, developed, and learned vision.
10. One can’t help but wonder if there are not serious commonalities between species, too. If there were not substantial overlap then there would be more exploitable differences like the day-glo and black “camouflage” that manages to increase one’s visibility to one’s fellow hunters and decrease it to deer. (Actually it appears that the white deer may have quite serviceable color vision and the efficacy of camouflage is just to break up the silhouette or outline. It appears that deer act as if two sensory modalities have agreed that danger is present before the deer will take action [43].) Besides, even this example relies on the sort of physical difference modulated out of consideration above. As it is, dogs do fine in cars, they catch frisbees, they deal with plate glass, they recognize manufactured objects, they have identifiable visual aversions, and so forth. One does not get the impression that mammals parse the world up in ways that are all that different from us. Of course I know about “what is it like to be a bat” (This 1974 article has been reprinted in an interesting and amusing anthology [61]) I am not trying to wax philosophical, but I am interested in the practical situation.
11. I am not convinced that we really understand what the different visual areas do. We are still at the stage that a substantial part of the community seems to think that simple cells are (i.e. have the function of) edge detectors. Attributing an easily-articulated functionality to a complex thing like a brain area on the basis of a few neurons studied in isolation may be wishfully simplistic.
12. If anyone can indulge risky ideas and holistic theories, artificial intelligence researchers can. Let us try to avoid being overly-influenced either by popular paradigms or by blind spots in the natural sciences. Remember we’re in the sciences of the artificial.

## References

- [1] A. P. Ambler, H. G. Barrow, C. M. Brown, R. M. Burstall, and R. J. Popplestone. A versatile computer-controlled assembly system. *Artificial Intelligence*, 6(2):129–156, 1975.
- [2] M. A. Arbib and J. Buhmann. Neural networks. In S. C. Shapiro, editor, *Encyclopedia of AI, 2nd. Ed.*, pages 1016–1060. Wiley, 1992.
- [3] R. N. Aslin, editor. *Development of perception: psychobiological perspectives*. Academic Press, 1981.
- [4] R. N. Aslin. Experiential influences and sensitive periods in perceptual development: a unified model. In R. N. Aslin, J. R. alberts, and M. L. Peterson, editors, *Development of Perception, psychobiological perspectives: Volume 2, the visual system*, chapter 2, pages 45–93. Academic Press, 1981.
- [5] E. A. Bates and J. L. Elman. Connectionism and the study of change. In M. H. Johnson, editor, *Brain Development and Cognition: A Reader*, chapter 26, pages 623–642. Blackwells, 1993.
- [6] P. Bateson. Genes, environment and the development of behaviour. In T.R. Halliday and P. J. B. Slater, editors, *Genes, Development and Learning*, pages 52–81. Freeman, 1983.
- [7] R. K. Belew. Interposing an ontogenic model between genetic algorithms and neural networks. In S. J. Hanson, J. D. Cowan, and C. L. Giles, editors, *Advances in Neural information processing systems 5*, pages 99–106. Morgan Kaufmann, 1993.
- [8] C. Blakemore and G.F. Cooper. Development of the brain depends on the visual environment. *Nature*, 228:477–478, 1970.
- [9] C. Blakemore and R. C. von Sluyter. *Journal Physiology, London*, 248:663, 1975.
- [10] K. W. Bowyer and L. Hall. Working notes, AAAI fall symposium 1993: Machine learning in computer vision: What, why, and how. Technical Report FSS-93-04, AAAI, 1993.
- [11] K. Breland and M. Breland. The misbehavior of organisms. *American Psychologist*, 16:681–684, 1961.
- [12] P. Buisseret and M. Imbert. Visual cortical cells: their developmental properties in normal and dark-reared kittens. *Journal of Physiology*, 255:511–525, 1976.
- [13] X. Chen and F. Schmitt. Vision-based construction of CAD models from range images. In *International Journal of Computer Vision*, pages 129–136, 1993.
- [14] P. S. Churchland and T. J. Sejnowski. *The Computational Brain*. Bradford Book, MIT Press, 1992.
- [15] W. M. Cowan. The development of the brain. *Scientific American*, 241:113–133, 1974.
- [16] M. Cynader and D.E. Mitchell. Prolonged sensitivity to monocular deprivation in dark-reared cats. *Journal of Neurophysiology*, 43:1026–1040, 1980.
- [17] E. Deucher. *Cellular interactions in animal development*. Chapman and Hall, 1975.
- [18] A. Diamond. Rate of maturation of the hippocampus and the developmental progression of children’s performance on the delayed non-matching to sample and visual paired comparison tasks. *Annals, New York Academy of Sciences*, pages 394–426, 1990.

- [19] M. Domjan. Ingestional aversion learning: unique and general processes. *Advances in the Study of Behavior*, 11:276–336, 1980.
- [20] Dowsett-Lemaire. The imitative range of the song of the marsh warbler *acrocephalus palustris*, with special reference to imitations of african birds. *Ibis*, 121:453–468, 1979.
- [21] J. H. Flavell. *Cognitive Development, 2nd Ed.* Prentice-Hall, 1985.
- [22] G. N. Reeke G. M. Edelman. Neural darwinism. In M. A. Arbib and J. A. Robinson, editors, *Natural and Artificial Parallel Computation*, pages 211–250. MIT Press, 1990.
- [23] C. R. Gallistel. *The organization of learning.* MIT Press, 1990.
- [24] J. Garcia, F. R. Ervin, and R. A. Koelling. Learning with prolonged delay of reinforcement. *Psychonomic Science*, 5:121–122, 1966.
- [25] J. Garcia and R. A. Koelling. Relation of cue to consequence in avoidance learning. *Psychonomic Science*, 4:123–124, 1966.
- [26] E. J. Gibson and R. D. Walk. The 'visual cliff'. *Scientific American*, 202:64–71, 1960.
- [27] G. Gottlieb. Development of species identification in ducklings IV: Change in species-specific perception caused by auditory deprivation. *Journal of Comparative and Physiological Psychology*, 92:375–387, 1978.
- [28] W. T. Greenough, J. E. Black, and C. S. Wallace. Experience and brain development. In M. H. Johnson, editor, *Brain Development and Cognition: A Reader*, chapter 14, pages 290–322. Blackwells, 1993.
- [29] J. B. Gurdon. *The control of gene expression in animal development.* Harvard University Press, 1974.
- [30] S. J. Hanson, J. D. Cowan, and C. L. Giles, editors. *Advances in Neural Information Processing Systems 5.* Morgan Kaufmann, 1993.
- [31] H.F. Harlow. Theheterosexual affectional system in monkeys. *Am. Psychol.*, 17:1–9, 1962.
- [32] H.F. Harlow and M. K. Harlow. Social deprivation in monkeys. *Scientific American*, 207:68–74, 1962.
- [33] E. Hearst. Stimulus relationships and feature selectionin learning and behavior. In S. H. Hulse, H. Fowler, and W. K. Honig, editors, *Cognitive Processes in Animal Behavior*, pages 51–88. Erlbaum, 1978.
- [34] D. O. Hebb. Heredity and environment in mammalian behaviour. *British Journal of animal behaviour*, 1:43–47, 1953.
- [35] H.V. B. Hirsch and D. N. Spinelli. Visual experience modifies distribution of horizontally and vertically oriented receptive fields in cats. *Science*, 168:869–871, 1970.
- [36] D. H. Hubel and T. N. Wiesel. Receptive fields of cells in striate cortex of very young, visually inexperienced kittens. *Journal of Neurophysiology*, 26:994–1002, 1963.
- [37] K. Ikeuchi. Generating an interpretation tree from a CAD model for 3-D-object recognition in bin-picking tasks. *International Journal of Computer Vision*, 1(2):145–165, 1987.
- [38] T. S. Huang J. J. Weng, N. Ahuja. Learning recognition and segmentation of 3-D objects from 2-D images. In *International Journal of Computer Vision*, pages 121–128, 1993.

- [39] J. S. Janowsky. The development and neural basis of memory systems. In M. H. Johnson, editor, *Brain Development and Cognition: A Reader*, chapter 28, pages 665–678. Blackwells, 1993.
- [40] M. H. Johnson. Cortical maturation and the development of visual attention in early infancy. In M. H. Johnson, editor, *Brain Development and Cognition: A Reader*, chapter 10, pages 167–194. Blackwells, 1993.
- [41] M. H. Johnson and J. Morton. *Biology and Cognitive Development: The case of face recognition*. Blackwells, 1991.
- [42] T. D. Johnston. The persistence of dichotomies in the study of behavioral development. *in press, Developmental Review*, 1994.
- [43] K. Karn. Personal communication, 1994.
- [44] M. J. Keating. Visual function and binocular visual connexions. *British Medical Bulletin*, 30:145–151, 1974.
- [45] D. E. Kroodsma. Aspects of learning in the ontogeny of bird song: where, from whom, when, how many, which and how accurately? In G. Burghardt and M. Bekoff, editors, *Development of behavior: comparative and evolutionary aspects*, pages 215–230. Garland, 1978.
- [46] E. P. Krotkov. *Active Computer Vision by Cooperative Focus and Stereo*. Springer-Verlag, 1989.
- [47] S. Lehky and T.J. Sejnowski. Network model of shape-from-shading: Neural function arises from both receptive and projective fields. *Nature*, 333:452–454, 1988.
- [48] D. Lehrman. Semantic and conceptual issues in the nature-nurture problem. In L. Aronson, E. Tobach, D. Lehrman, and J. Rosenblatt, editors, *Development and evolution of behavior*. Freeman, 1970.
- [49] Ralph Linsker. From basic network principles to neural architecture: Emergence of orientation columns. *Proc. Natl. Acad. Sci. USA*, 83:8779–8783, November 1986.
- [50] Ralph Linsker. From basic network principles to neural architecture: Emergence of spatial-opponent cells. *Proc. Natl. Acad. Sci. USA*, 83:7508–7512, October 1986.
- [51] Ralph Linsker. Self-organization in a perceptual network. *Computer*, pages 105–117, March 1988.
- [52] N. J. Mackintosh. General principles of learning. In T.R. Halliday and P. J. B. Slater, editors, *Genes, Development and Learning*, pages 149–177. Freeman, 1983.
- [53] M. M. Merzenich, G. H. Recanzone, W. M. Jenkins, and R. J. Nudo. How the brain functionally rewires itself. In M. A. Arbib and J. A. Robinson, editors, *Natural and Artificial Parallel Computation*, pages 177–210. MIT Press, 1990.
- [54] K. D. Miller, J. B. Keller, and M. P. Stryker. Ocular dominance column development: analysis and simulation. *Science*, 245:605–615, 1989.
- [55] K. D. Miller and M. P. Stryker. Ocular dominance column formation, mechanisms and models. In S. J. Hanson and C. R. Olson, editors, *Connectionist Modeling and Brain Function: The Developing Interface*, pages 255–350. MIT Press, 1990.

- [56] B. Milner, S. Corkin, and H. L. Teuber. Further analysis of the hippocampal amnesic syndrome: 14 year follow-up study of H.M. *Neuropsychologia*, 6:215–234, 1968.
- [57] T. M. Mitchell and S. B. Thrun. Explanation-based neural network learning for robot control. In S. J. Hanson, J. D. Cowan, and C. L. Giles, editors, *Advances in Neural information processing systems 5*, pages 287–294. Morgan Kaufmann, 1993.
- [58] P. R. Montague, P. Dayan, S.J. Nowlan, A. Pouget, and T.J. Sejnowski. Using aperiodic reinforcement for directed self-organization during development. In S. J. Hanson, J. D. Cowan, and C. L. Giles, editors, *Advances in Neural Information Processing Systems 5*, pages 969–976. Morgan Kaufmann, 1993.
- [59] A. Moore. *Efficient Memory-Based Learning for Robot Control*. PhD thesis, University of Cambridge, November 1990.
- [60] M. J. Morgan and D. J. Nicholas. Discrimination between reinforced action patterns in the rat. *Learning and Motivation*, 10:1–22, 1979.
- [61] T. Nagel. What is it like to be a bat? In D. R. Hofstadter and D. C. Dennett, editors, *The Mind's I*, chapter 24, pages 391–403. Bantam, 1982.
- [62] K. Obermayer, H. Ritter, and K. Schulten. A principle for the formation of the spatial structure of cortical feature maps. *Proc. Natl. Acad. Sci.*, 87:8345–8349, November 1990.
- [63] K. Obermayer, K. Schulten, and G. G. Blasdel. A comparison between a neural network model for the formation of brain maps and experimental data. In J. E. Moody, S. J. Hanson, and R. P. Lippmann, editors, *Advances in Neural Information Processing Systems 4*, pages 83–90. Morgan Kaufmann, 1992.
- [64] D. D. M. O’Leary. Do cortical areas emerge from a protocortex. In M. H. Johnson, editor, *Brain Development and Cognition: A Reader*, chapter 15, pages 322–337. Blackwells, 1993.
- [65] R. W. Oppenheim and L. Haverkamp. Early development of behavior and the nervous system: An embryological perspective. In E. Adler, D. Pfaff, and R. W. Goy, editors, *Handbook of behavioral neurobiology*. Plenum, 1985.
- [66] S. Oyama. *The ontogeny of information*. Cambridge University Press, 1985.
- [67] T. P. Pons, P. E. Garraghty, A.K. Ommaya, et al. Massive cortical reorganization after sensory deafferentation in adult macaques. *Science*, 252:1857–1860, 1991.
- [68] D. Purves and J. W. Lichtman. Elimination of synapses in the developing nervous system. *Science*, 210:153–157, 1980.
- [69] V. S. Ramachandran. The perception of apparent motion. *Scientific American*, 254(6):102–109, June 1986.
- [70] V. S. Ramachandran. Perceiving shape from shading. *Scientific American*, 256(8):76–83, August 1988.
- [71] K. T. Rogers. Optokinetic testing of cyclopean and synophthalmic fish hatchlings. *Biological Bulletin*, 112:241–248, 1957.
- [72] T. J. Roper. Learning as a biological phenomenon. In T.R. Halliday and P. J. B. Slater, editors, *Genes, Development and Learning*, pages 178–212. Freeman, 1983.

- [73] Y. Aloimonos R.S. Michalski, A. Rosenfeld and P. Pachowcz. Machine learning and vision: Research issues and directions. Technical Report AIC/GMU Report in Preparation, George Mason University and University of Maryland, October 1992. NSF/ARPA Workshop at Harper's Ferry.
- [74] D. E. Rumelhart and J. L. McClelland. *Parallel Distributed Processing*. MIT Press, 1986.
- [75] J. C. Schlimmer and P. Langley. Learning, machine. In S. C. Shapiro, editor, *Encyclopedia of AI, 2nd. Ed.*, pages 785–805. Wiley, 1992.
- [76] J. Schneider. Learning open loop control of complex motor tasks. In *AAAI Fall Symposium Series Working Notes: Machine Learning in Computer Vision: What, Why, and How?*, pages 130–134, October 1993.
- [77] J. Schneider and C. Brown. Robot skill learning and the effects of basis function choice. Technical Report 437, Department of Computer Science, University of Rochester, September 1992.
- [78] J. Schneider and C. M. Brown. Learning complex motor skills. In *Proceedings: IEEE International Conference on Robotics and Automation*, Submitted 1994.
- [79] H. Sherk and M. P. Stryker. *Journal of Neurophysiology*, 39:63, 1976.
- [80] D. Sherry and D. Schacter. The evolution of multiple memory systems. *Psychology Review*, 94:439–454, 1987.
- [81] S. J. Shettleworth. Reinforcement and the organization of behaviour in golden hamsters: differential overshadowing of a CS by different responses. *Quarterly Journal of Experimental Psychology*, 33B:241–256, 1981.
- [82] S. J. Shettleworth and J. R. Krebs. How marsh tits find their hoards: the roles of site preference and spatial memory. *J. of Exp. Psychology: Animal Behavior Processes*, 8:354–375, 1982.
- [83] W. Singer and F. Tretter. *Journal of Neurophysiology*, 39:613, 1976.
- [84] E. S. Spelke. Origins of visual knowledge. In D. N. Osherson, S. M. Kosslyn, and J. M. Hollerbach, editors, *Visual cognition and Action (An Invitation to Cognitive Science Vol. 2*, chapter 4, pages 99–128. MIT Press, 1990.
- [85] E. Thelen. Self-organization in developmental processes: can systems approaches work? In M. H. Johnson, editor, *Brain Development and Cognition: A Reader*, chapter 24, pages 555–591. Blackwells, 1993.
- [86] G.J. Thomas. Memory: time binding in organisms. In L. R. Squire and N. Butters, editors, *Neuropsychology of memory*, pages 374–384. Guilford, 1984.
- [87] M. K. Warmuth and L. G. Valient, editors. *4th Workshop on Computational Learning Theory*. Morgan Kaufmann, 1991.
- [88] C. J. C. H. Watkins. *Learning from Delayed Rewards*. PhD thesis, Cambridge University, 1989.
- [89] T. N. Wiesel. Postnatal development of the visual cortex and the influence of the environment. *Nature*, 299:583–591, 1982.
- [90] T. N. Wiesel and D.H. Hubel. Ordered arrangements of orientation columns in monkeys lacking visual experience. *Journal of comparative neurology*, 158:307–318, 1982.
- [91] L. B. Wolff, S. A. Shafer, and G. E. Healey, editors. *Radiometry*. Jones and Bartlett, 1992.

- [92] A.L. Yuille, D.M. Kammen, and D.S. Cohen. Quadrature and the development of orientation selective cortical cells by hebb rules. *Biol. Cybern.*, 61:183–194, 1989.
- [93] S. Zeki. *A Vision of the Brain*. Blackwells Scientific, 1993.