To appear in S. Wermter, J. Austin &D. Willshaw (eds.) Emergent neural computational architectures based on neuroscience. Springer-Verlag.

# Plasticity and nativism: Towards a resolution of an apparent paradox

Gary F. Marcus[1]

[1] Department of Psychology, 6 Washington Place, New York University
gary.marcus@nyu.edu

**Abstract**: Recent research in brain development and cognitive development leads to an apparent paradox. One set of recent experiments suggests that infants are well-endowed with sophisticated mechanisms for analyzing the world; another set of recent experiments suggests that brain development is extremely flexible. In this paper, I review various ways of resolving the implicit tension between the two, and close with a proposal for a novel computational approach to reconciling nativism with developmental flexibility.

## 1 Introduction: An Apparent Paradox

One strand of contemporary scientific research suggests that human infants are born with sophisticated mechanisms for learning about and analyzing the world. Within the first year of life, human infants can, among other things, anticipate sequences of events [1], keep track of objects that they cannot see [2, 3], discern abstract patterns in artificial languages [4, 5], and discriminate between unfamiliar languages that have different rhythmic properties [6]. In keeping with views advanced by Chomsky [7] and Fodor [8], "nativist" researchers such as Spelke [9] , Pinker [10], Leslie [11] and Crain [12] have taken these studies (and many others like them) to be evidence that the mind is importantly structured in advance of experience.

Another strand of contemporary scientific research suggests that brain development is remarkably flexible (or "plastic") – sizes of some brain areas depends on input [e.g., 13], early in development, some brain cells can be transplanted from one area of the brain to another [14], and certain parts of the brain can even be "rewired" [15, 16].

All this evidence that brain development is flexible has led some to think that nativism is in trouble. How could a newborn be born with language acquisition device if young children with left hemisphere brain injuries can recover language function to a significant extent? If the size of brain regions depends on experience, how could there be a built in module for tracking objects through time? If brain cells are not "born knowing their destinations", how could representations be innate? According to Elman, Bates, Johnson, Karmiloff-Smith, Parisi, & Plunkett [17] "the last two decades of research on vertebrate brain development force us to conclude that innate specification of synaptic connectivity at the cortical level is highly unlikely" (p. 361).

Drawing on similar results, Quartz and Sejnowski [18] concluded that experiments in brain flexibility show that "although the cortex is not a tabula rasa … it is largely equipotential at early stages" (p. 552) and that "nativist theories [therefore] appear implausible" (p. 555). Elman et al argue that "Representation-specific predispositions … may only be specified at the subcortical level *as little more than attention grabbers*" [emphasis added] that ensure the organism will receive "massive experience of certain inputs prior to subsequent learning..." (p. 108).

Do studies of brain development really militate against nativism? Researchers like Elman et al certainly seem to think so, when they make it clear that the target of their attacks is the nativist positions of researchers like Spelke, Pinker, Leslie, and Crain. In the place of these strong nativist positions, Elman et al settle for a sort of stripped-down nativism in which "architectural" aspects of brain organization are innate, but "representations " are not.

While I see the appeal in their position, I think it is ultimately untenable. Evidence that brain development is flexible really does challenge some of the simplest ways in which there could be innate mental structure, but, I will argue, it leaves more sophisticated versions of nativism untouched. Moreover, I will argue that the stripped-down nativism of Elman et al probably relies too much on experience. I will end the paper by sketching a way in which strong nativism might be reconciled with developmental flexibility, proposing a novel computational approach that integrates neural networks simulations with findings in developmental biology.

## 2 Innateness

Before we can get to arguments about why developmental flexibility might be challenging to innateness, it is worth briefly reviewing some of the reasons for believing that significant aspects of mental structure might be innate. (Given space limitations, I do not aim to be comprehensive here; excellent, recent reviews include Spelke and Newport [19] and Pinker [10].)

.

### 2.1 Case study: Objects

One reason for believing that significant aspects of mental structure might be innate is that recent studies of human infants suggest that they are capable of sophisticated analysis of the world. One case study that has received a great deal of attention is infant's understanding of the idea that objects persist in time. Piaget famously noted that 8-month-olds would cease to show interest in a toy if that toy was covered by a blanket; Piaget argued that the child had to *construct* (de novo) the notion of a persisting object. But dozens of recent studies suggest that infants behave as if they know[1] objects persist in time long before they begin to reach for occluded objects. For

[1] I use words like "know" and "understand" loosely here –- I do not mean to say that infants consciously represent knowledge about objects, but rather that their computational systems respond in ways that are consistent with some sort of representation of object permanence.

example, Spelke and Kestenbaum [2] conducted an experiment in which a four-month-old infant was seated at a stage that initially contained two screens. The infant subject would then see an object, in this case a rod, pass behind a screen, a bit of time would pass, and then the infant would then see an identical-appearing rod emerge from behind the other screen. The rod would then go back behind the second screen, some more time would pass, and then the rod would emerge from behind the first screen. This back-and-forth procedure would continue several times until the infant was bored, and then the infant would see the screen lifted, revealing either a single rod or two rods, one behind each screen. Spelke et al. found that infants look longer when they are shown just one rod. Because infants generally look at longer at novel or unfamiliar outcomes, the results suggest that infants were "expecting" to see two distinct rods. Given that the infant only saw only one rod at any given moment, the result suggests that the infants kept track the rods, even when those rods were occluded. While the exact interpretation of these experiments is still open, it seems likely that at least some of the machinery that infants use in this task is innate.

### 2.2 Learning

Learning and innateness are often taken to be in opposition, but they need not be: learning mechanisms may themselves be innate. For example, using a variation on the habituation methods of Spelke and others, Saffran, Aslin and Newport [20] recently showed that eight-month-old can detect subtle statistical information from sequences of speech sounds produced in artificial languages. For example, in one experiment Saffran et familiarized infants with a two minute long, unbroken string of "familiarization" syllables such as *tibudopabikudaropigolatupabikutibudogolatu-daropidaropitibudopabikugolatu*. In this familiarization, some sounds are always followed by other sounds (e.g., every occurrence of *pa* was followed by *biku*), whereas other sounds are only sometimes followed by a particular sound (e.g., exactly one third the occurrences of *pi* were followed by *gola* ; other occurrences of *pi* were followed by *daro* or *tibu*). Saffran et al. found that infants attended longer during presentations of sequences like *pigola* than during presentations of words like *pabiku*, showing that infants extracted information about how often particular items follow one another. While it is possible that this statistical learning mechanism is learned, I know of no proposal for how it could be learned; instead, my hunch is that the learning mechanism itself is innate, built in prior to experience.

Similarly, my colleagues and I have shown that seven-month-old infants are able to learn "abstract rules"[4]. For instance, we exposed one set of infants to two minutes of "ABA" with sentences like ga ti ga and li na li. After this two-minute familiarization, we exposed infants to test sentences that were made up entirely of novel words that were either consistent with or inconsistent with the familiarization grammar. The prediction was that if infants can distinguish the two grammars and generalize them to new words, they should attend longer during inconsistent items. For example, if infants that were trained on the ABA grammar, we expected them to attend longer during, an ABB test item like *wo fe fe* than during an ABA test item like *wo fe wo*. As predicted, infants looked longer at the inconsistent items, suggesting that infants able to extract the ABA pattern and use it in evaluating new items.

(Similar results with twelve-month-old children were reported in [5]). Although I cannot prove that the mechanism for rule-learing is innate, I strongly suspect that it is. A true tabula rasa position would be incoherent -- learning must start somewhere.

### 2.3 Learning must start somewhere

Generalizing that point – that learning has to start somewhere – a third reason for believing that something is innate is that there may be no other satisfying account for how a given piece of knowledge could arise. So-called "learnability" arguments are perhaps most often made in the context of language acquisition. For example, Gordon [21] asked children to produce compounds such as *mice-eater and rat-eater*. He found that while children often produce compounds that contain irregular plurals (e.g., *mice-eater*) they essentially never produce compounds containing regular plurals (e.g., *rats-eater*). The way that children behave is consistent with a linguistic distinction that holds in English and perhaps cross-linguistically. But plurals inside compounds are so rare that young children are unlikely to have heard any; their inference thus in some sense probably goes beyond the input. From the fact that all children go beyond the data in a consistent way, Gordon argued that there must be some sort of built-in machinery constraining their learning. More general versions of "learnability" arguments have been made in the domain of language acquisition by Wexler and Culicover [22], Pinker[23, 24], and Crain [12], among others.

Similar arguments have been made in other domains; for example, Spelke [9] suggested that the ability to represent objects may be innate:

> If children are endowed with abilities to perceive objects, persons, sets, and places, then they may use their perceptual experience to learn about the properties and behavior of such entities. By observing objects that lose their support and fall, children may learn that unsupported objects fall... it is far from clear how children could learn anything about the entities in a domain, however, if they could not single out those entities in their surroundings.
>
> …[in contrast] if children could not represent the object-that-loses-its-support as the *same object* as the object-that-falls (and as a different object from the support itself), they might only learn that events in which something loses support are followed by events in which something falls (the object) and something remains at rest (the support).

## 3 Developmental flexibility and DNA as blueprint

If the mind is indeed importantly structured prior to experience, how did it get that way? I ask this not as a question about evolution, but as a question about developmental biology. To the extent the mind is a product of the brain, how could the brain be organized prior to experience?

It would certainly be convenient for nativists if fertilized eggs contained a blueprint for building the brain. Just as an architectural blueprint might specify exactly where every room and corridor in some new office building might be placed, one might imagine the fertilized egg bearing a neural blueprint that would specify where every neuron and connection in the to-be-born child's brain would be placed. This "DNA-as-blueprint" idea would fit nicely with nativism, but, alas, it clearly cannot be right.

For one thing, there just is not enough information in the human genome to specify exact where each neuron and synapse will go [25]. There are about $10^5$ genes which contain about $10^9$ nucleotides, as compared with about $10^{10}$ neurons and about $10^{15}$ or so synapses.

Moreover, as noted in the introduction, brain development is flexible, and this flexibility seems inconsistent with blueprint idea. For example, if the exact structural organization of some brain region were predestined, its size should not depend on the amount of input received. Yet the size of some brain regions does indeed depend on the amount of input [13].

Similarly, if the DNA provided a blueprint, one would not necessarily expect the brain to be able to adapt itself in response to radical "rewiring", yet experiments by Sur and his colleagues [16] show that when visual thalamic inputs are rewired from their usual destination in visual cortex to a novel destination in auditory cortex, the auditory cortex begins to take on some of the properties of visual cortex.

Plainly the DNA does not specify a point-by-point wiring diagram for the human brain. Other evidence further underscores the view of brain development as flexible. O'Leary and Stanfield [14] showed that when visual cortex neurons are transplanted into somatosensory areas, they develop (at least in some respects) as one would expect for somatosensory neurons rather than for visual neurons, projecting not to the visual cortex, but to the spinal cord. Likewise, somatosensory cells transplanted to visual cortex develop projections that are typical of visual neurons. Furthermore, although recovery from brain injuries that occur in adulthood may be quite minimal (although non-zero), recovery from brain injuries in childhood can be much more substantial, with undamaged areas of the brain taking over some of the functions of damaged areas of the brain [e.g., 26].

Where does this leave us?

## 4 The Neo-Constructivist Synthesis

Scholars such as Quartz and Sejnowski and Elman et al see the evidence from developmental flexibility as devastating to nativism. Neither set of researchers wishes to dispense with nativism altogether, but both groups direct their criticism towards researchers such as Chomsky, Fodor, Spelke, Pinker, and Crain, and both put the burden of brain organization primarily on learning, stressing "massive experience" over any kind of significant intrinsic organization. For example, since, as they put it, "neurons can't be born knowing their destinations", Elman et al conclude that strong nativism must be wrong. In its place, they argue that "architectural" aspects of the brain – how many layers there are, how many units are in those layers, and so forth –

are organized in advance, but they suggest that the detailed microcircuitry is not; Quartz and Sejnowski make similar points. Collectively, I will call their position the "neo-constructivist synthesis".

In order to make it plausible that weak initial biases could combine with experience in satisfactory ways, both groups of researchers point to a series of connectionist models or neural networks. *Neural networks* are idealized computer simulations that are intended to tell us something important about how the mind/brain works. Typically, they consist of sets of neuron-like *nodes* interacting in parallel.

The models that they describe typically look something like Figure 1. (I assume most readers have at least a passing familiarity with these networks. In brief, a set of *input nodes* represents the input to the network, the set of *output nodes* represent the output from that network. Intervening between the input and nodes is a set of *hidden units* that re-represent the input. The arrows indicate the extent to which different nodes are connected together. Such models are typically trained on the basis of input-output pairs; during this training, connections between nodes are adjusted in way that attempts to minimize error. For a more in-depth introduction, see, for example, [27].)

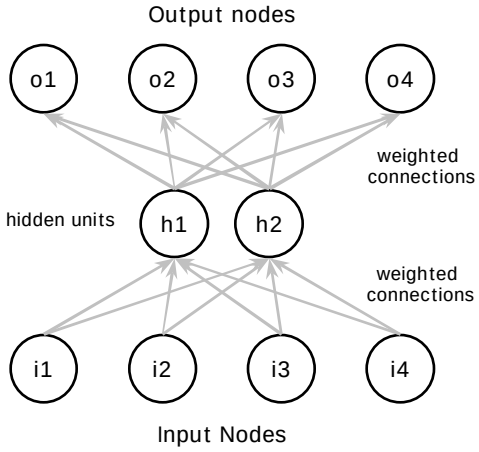


**Fig. 1.** A simple three-layer neural network

There are, in fact, many ways of arranging nodes and connections, and different arrangements have different implications for basic questions in cognition such as whether the mind is like a symbol-manipulating computer; I have written extensively about such issues [28-30] but will mainly skip them here. Suffice it to say here that the neural networks endorsed by Elman et al are, by design, among those with the least innate structure – Elman et al see themselves as providing a computational basis for Piagetian constructivism:

> ... constructivism [considered] development in terms of self-organizing emergent structures arising from the complex interactions between both organism and environment. We believe that the biological-connectionist perspective opens the door to a new framework for thinking about development which embodies some aspects of Piaget's,

> Werner's, and Vygotsky's constructivist intuitions, but which goes beyond them and provides a formalized framework within which to generate empirically testable questions (p . 114)

## 5. Discussion of the neo-constructivist synthesis

Although some may find the neo-constructivist thesis to be appealing, we are by no means forced to adopt it. One reason is that there are serious limits on the particular *models* that researchers like Elman et al have advocated [28-32]. Such models are limited in their abilities to generalize, they have difficulties in representing fundamental notions such as the distinction between individuals and kinds, and they are, I think, too unstructured. But the theoretical position is independent of the models; rather than rehearsing my criticisms of the models here, I want to focus instead on some of the theoretical assumptions that are implicit in the neo-constructivist synthesis.

### Developmental Flexibility

First, the neo-constructivist position seems to rest on the idea that developmental flexibility entails learning. But developmental flexibility does not *entail* learning. While any learning must involve some change of the underlying neural substrate, many changes in the neural substrate probably proceed without anything like learning.

It turns out that developmental flexibility is characteristic of mammalian development in general – we see quite similar flexibility in the development of the heart, the kidney, the eye, and so forth – organs in which learning plays little or no role. Virtually any part of a developing organism can recover from damage if that damage takes place early enough – the recent experiments establishing robustness in brain development are not so different from other experiments from the early days of embryology. At least since Han Spemann's pioneering work in the 1920s [33], developmental biologists have routinely used transplantation as a window into embryology; quite often, if those experiments are done early enough, transplanted tissue takes on some or all of the characteristics of its new host region. In the words of noted embryologist Lewis Wolpert [34, p. 42]:

> In general, if cells of vertebrate embryos are moved from one part to another of the early embryo they develop according to their new location and not from where they are taken. Their fate is dependent on their new position in the embryo: they respond to their new address.

For example, if early in development one takes cells from the region of a frog embryo that normally develops into an eye and transplants them into the gut, they develop into gut cells rather than eye cells, much as a transplanted somatosensory cell

may take on characteristics of its new home.[2] Such flexibility may even be adaptively advantageous; as Cruz [36] put it

> In a rapidly growing embryo consisting of cells caught in a dynamic flurry of proliferation, migration, and differentiation, it would be desirable for any given cell to retain some measure of developmental flexibility for as long as possible. Such would enable an embryo momentarily disabled by cell cycle delay, for instance, or temporarily compromised by loss of a few cells, to compensate for minor disruptions and resume rather quickly the normal pace of development. It is easy to see how such built-in [flexibility] could contribute to the wide variety of procedural detail manifest in nearly every phase of mammalian embryogenesis (p. 484).

One would not want to say that the eye cell *learns* how to be a stomach cell, and one should not assume that a transplanted somatosensory cell learns how to be a visual cell. None of this rules out learning (and there must be important learning eventually), but it does remind us that developmental flexibility does not on its own entail learning.

**DNA as blueprint**

A second problem with the neo-constructivist synthesis is that it seems to equate nativism with the idea of the DNA as a blueprint. In fact, the DNA rarely if ever serves as literal blueprint in any part of biology, but there is no reason that nativism must depend on such a fantastical view of DNA. One need only look to the heart or the eye to see that nature can build highly intricate structure without depending on learning. The idea of DNA as blueprint is really a strawman that makes little sense in any part of biology. As Richard Dawkins [37] has put it, the DNA is much more like a recipe than a blueprint – the DNA gives a set of instructions for building something, not a diagram of what the finished product will look like. But a recipe is enough – the toolkit of biology is sufficiently powerful that it can build bodies without requiring a whole lot of learning, and I suspect that very same toolkit is powerful to build brains as well. For this reason, I believe that a complete account of brain development must make substantial reference to the toolkit of developmental biology.

**Neural activity**

The third serious problem with the neo-constructivist synthesis is that it rests too heavily on learning and neural activity, attributing virtually all detailed brain organization to neural activity. But a number of recent studies in developmental neuroscience suggest that while neural activity is important to brain development, it

[2] Transplants of brain cells, too, seem to be age-dependent, with the chance of a transplanted cell taken on target characteristics greatest earlier in development [35].

may not be essential for early stages of brain development. The idea that an organism's detailed microcircuitry can only be specified on the basis of massive experience simply is not tenable. For example, Crowley & Katz [38] recently demonstrated that the organization of ferret geniculocortical axons into ocular dominance columns could occur even in the complete absence of retinal input. In another set of experiments, Verhage, et al [39] created "knock-out" mice that lacked the gene *Munc-18*, causing a "complete loss of neurotransmitter secretion from synaptic vesicles throughout development" ; their striking finding was that brain assembly was apparently normal, "including formation of layered structures, fiber pathways, and morphologically defined synapses." Considerations like these led Katz, Weliky, and Crowley [40] to conclude that

> "The current emphasis on correlation-based models, which may be appropriate for later plastic changes, could be obscuring the role of intrinsic signals that guide the initial establishment of functional architecture."

## 7. A new approach

In a nutshell, what I think is being left out in the neo-constructivist synthesis is the toolkit of developmental biology. Developing embryos are blessed with an extraordinary array of techniques for organizing themselves, ways of coordinating the actions of simple genes into incredibly complex organisms. It is my hunch that a proper account of brain development should make extensive use of the tools that biology uses when it builds organisms.

All of which is rather vague. To make it more explicit I would like to borrow an idea from the neo-constructivists. Like them, I want to use neural networks as a way of understanding possible mechanisms of development. But unlike them, I want to build neural networks that *grow*, networks that show a good degree of self-organization even in the absence of experience.

In contrast to the neural networks of the neo-constructivists, the neural networks that I aim to build will integrate ideas about nodes and connections with some of the basic principles of developmental biology, including the following:

- Basic processes such as cell division, cell migration, and cell death.
- Gene expression. Genes can either be "expressed" or "repressed". What governs whether a particular gene is on or off is (among other things) the presence or absence of specific *regulatory* proteins that serve as enhancers or repressors for that gene [41]. When a gene is on, it sets into motion a transcription process that ultimately yields a particular protein. In the simulations, genes are rules with preconditions (which correspond to promoter sequences) and actions (which correspond, for example, to the construction of various proteins).

- Cell-to-cell communication. Many of the regulatory proteins that serve as triggers can pass from one cell to another; in this way, and also by means of electrical signaling, cells can communicate with each other. Mechanisms for both chemical and electrical signaling are included in the simulations.
- Cascades. Because an expressed gene can yield proteins, and proteins can trigger the expression of genes, one gene can trigger the action of another, or even several others, each of which in turn might trigger several others, and so forth -- what we might call a *cascade* [41-43]. These cascades, sometimes described as *regulatory networks* or *gene hierarchies,* are critical, because they provide a way for a complex coordinated actions to emerge. A particularly vivid example of this comes from the work of Walter Gehring; he and his collaborators have shown that a simple fruit fly gene known as *pax 6* triggers the action of (at least) three other genes, each of which in turns launches the action of still more genes, about 2500 in all [43-45]. What is special about *pax 6*, which Gehring calls a master control gene, is that it sits atop a hierarchy of genes that lead to the construction of an eye. Fruit flies that lack this gene generally do not have eyes; even more striking is the fact that if pax-6 is expressed (turned on) artificially in other parts of the body, eyes may grow in those regions; for example, Gehring and his collaborators were able to induce fruit flies to grow eyes on their antennae. The lesson here is not that there is a gene for building eyes – *pax 6* cannot do this by itself – but rather that the action of a single gene can through the process of cascading snowball into tremendously complex machinery. (Machinery for building cascades emerges in the simulation – and in nature – automatically, in virtue of the mechanisms that control gene expression.)

Whereas most work in developmental biology works bottom-up, by testing what happens if particular genes are “knocked-out” or artificially expressed, I aim to work in a more top-down fashion, asking how brains with particular properties could be assembled by genetic-like processes such as those mentioned above.

It is still very early days for this project, much too early to report any concrete results. Because the project is almost entirely new, I have spent the initial stages developing a prototype simulator that could be used to support this kind of work. A screen shot of the simulator is shown in Figure 2.

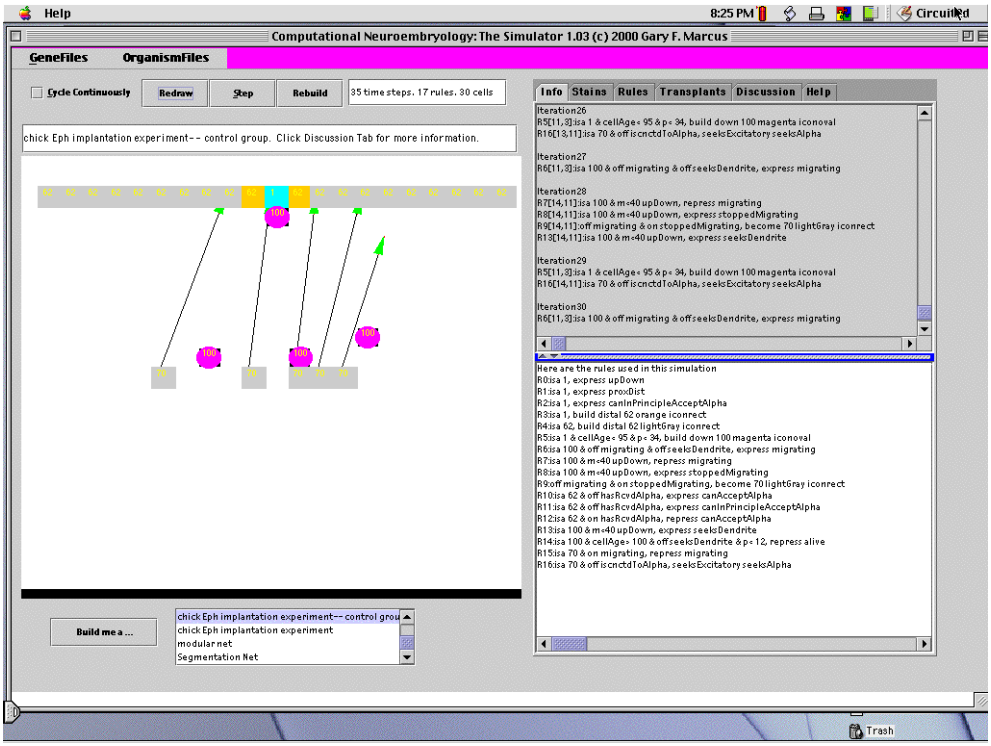


**Fig. 2.** Simulator

The main window shows the current state of a given embryo (what cells there are, what genes are expressed in those cells, etc.); another window shows the "genome" for that organism, a third shows what genes were expressed in a given time step.

In the particular simulation that is illustrated here, the top row depicts one layer of cells (to be thought of as "nodes" or neurons"), the bottom depicts another layer of cells; the dark gray cells represent a set of cells that are migrating from top to bottom; the arrows represent "axons" that are growing along gradients attempting to connect the top and bottom layers.

Various buttons allow the user to modify the genome, step forward in time, rebuild the organism, and so forth. Other controls allow the user to display the concentrations of diffusing "morphogens", stain cells according to their patterns of gene expression, and so forth. It is also possible for users to selectively lesion particular cells, transplant cells from one location to another, test the effects of knocking-out particular genes, etc..

Figures 3 and 4 give some more examples.[3] Figure 3 is a sort of time-lapse illustration of the growth of a simple organism, which we might think of a ladybug. In the ladybug simulation, there is only structure, not function, and nothing neural.

[3] Further examples may be found on my web site, http://www.psych.nyu.edu/gary.

Still, several interesting points are captured. First, the development of the ladybug proceeds in parallel – the *developmental program* that builds the ladybug is like a standard computer program in that it is made up of rules with preconditions and actions (each gene is essentially an IF-THEN rule) but whereas standard computer programs proceed serially, one step at a time, the developmental program (like those in biology) proceeds in parallel – each cell at the same time.

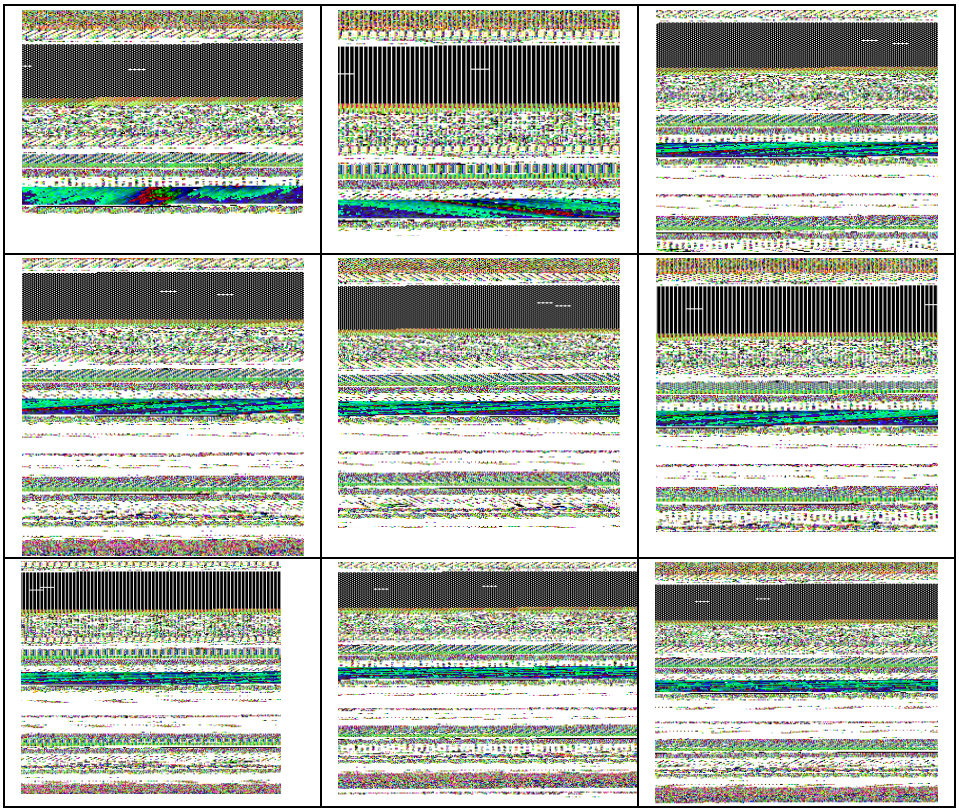

**Fig. 3.** Stages in the development of a simulated embryo

Second, the ladybug program illustrates, in a tiny way, the notion of compression; the complete ladybug has 79 cells, but just 32 rules; the developmental program is thus much more efficient than a literal blueprint would be.

Third, the ladybug program is developmentally robust; portions of the ladybug can be lesioned or amputated, and, salamander-like, they will grow back. The ladybug is thus a primitive illustration of a system that is innately organized (learning plays no role) yet developmentally robust.

Examples in Figure 4 show some simple of the neural structures that can be built in the prototype simulator.

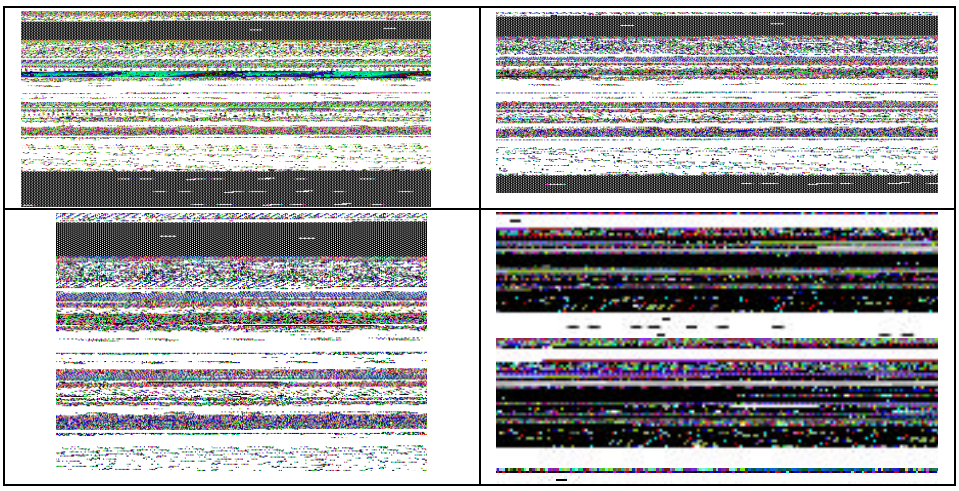

**Fig. 4.** Four networks "grown" in the simulator. See text for further details.

The top left panel illustrates a developmental stage in a two layer network that is topographically organized, such that connections maintain the relative left-right ordering. The top right panel illustrates a more complex network with several layers, and different types of connections. The bottom left panel depicts the core of a network that would solve a very simple visual segmentation task; the bottom right panel depicts a three-dimensional multilayered version of that visual segmentation network. Each of these models is to some degree developmentally robust; none depends on learning for its basic organization.

There is, to be sure, not enough known yet about developmental neuroscience to fully constrain this modeling enterprise; but important new discoveries are being made at impressive pace; biologists are learning, for example, a great deal about how axons are guided to their destinations. The mechanisms underlying what Roger Sperry [46] dubbed *chemoaffinity* are now being understood at a genetic level [e.g., 47], Such insights can be directly integrated into the modeling framework outlined here. My hope is that such modeling can ultimately help us to tie together discoveries in developmental biology, developmental neuroscience, and cognitive development.

## References


1. Haith, M.M., C. Hazan, and G.s. Goodman, *Expectation and anticipation of dynamic visual events by 3.5-month-old babies.* Child Development, 1988. **59**: p. 467-479.
2. Spelke, E.S. and R. Kestenbaum, *Les origins du concept d'object.* Psychologie Francaise, 1986. **31**: p. 67-72.
3. Wynn, K., *Addition and subtraction by human infants.* Nature, 1992. **358**: p. 749-750.
4. Marcus, G.F., S. Vijayan, S. Bandi Rao, and P.M. Vishton, *Rule learning in 7-month-old infants.* Science, 1999. **283**: p. 77-80.
5. Gomez, R.L. and L.-A. Gerken, *Artificial grammar learning by 1 year-olds leads to specific and abstract knowledge.* Cognition, 1999. **70**(1): p. 109-135.
6. Nazzi, T., J. Bertoncini, and J. Mehler, *Language discrimination by newborns: Towards an understanding of the role of rhythm.* Journal of Experimental Psychology: Human Perception and Performance, 1998. **24**: p. 1-11.
7. Chomsky, N.A., *Rules and representations.* 1980, New York: Columbia University Press.
8. Fodor, J.A., *The language of thought*. 1975: New York: T. Y. Crowell.
9. Spelke, E.S., *Initial knowledge: Six suggestions.* Cognition, 1994. **50**: p. 431-445.
10. Pinker, S., *How the mind works*. 1997, NY: Norton.
11. Leslie, A.M., *Pretense, autism, and the "Theory of Mind" module.* Current Directions in Psychological Science, 1992. **1**: p. 18-21.
12. Crain, S., *Language Acquisition in the Absence of Experience.* Behavioral and Brain Sciences, 1991. **14**: p. 597-650.
13. Kennedy, H. and C. Dehay, *Cortical specification of mice and men.* Cerebral Cortex, 1993. **3**(3): p. 171-86.
14. O'Leary, D.D. and B.B. Stanfield, *Selective elimination of axons extended by developing cortical neurons is dependent on regional locale: Experiments using fetal cortical transplants.* Journal of Neuroscience, 1989. **9**: p. 2230-2246.
15. Sharma, J., A. Angelucci, and M. Sur, *Induction of visual orientation modules in auditory cortex [see comments]*. Nature, 2000. **404**(6780): p. 841-7.
16. Sur, M., S.L. Pallas, and A.W. Roe, *Cross-model plasticity in cortical development: differentiation and specification of sensory neocortex.* Trends in Neuroscience, 1990. **13**: p. 227-233.
17. Elman, J.L., E. Bates, M.H. Johnson, A. Karmiloff-Smith, D. Parisi, and K. Plunkett, *Rethinking innateness: A connectionist perspective on development*. 1996, Cambridge, MA: MIT Press.
18. Quartz, S.R. and T.J. Sejnowski, *The neural basis of cognitive development: A constructivist manifesto.* Behavioral and Brain Sciences, 1997. **20**: p. 537-56; discussion 556-96.
19. Spelke, E.S. and E.L. Newport, *Nativism, empiricism, and the development of knowledge*, in *Handbook of Child Psychology (5th ed.), Vol. 1: Theories of development*, R.M. Lerner, Editor. 1998, Wiley: NY. p. 275-340.
20. Saffran, J., R. Aslin, and E. Newport, *Statistical learning by 8-month old infants.* Science, 1996. **274**: p. 1926-1928.
21. Gordon, P., *Level-ordering in lexical development.* Cognition, 1985. **21**: p. 73-93.
22. Wexler, K. and P. Culicover, *Formal principles of language acquisition*. 1980, Cambridge, MA: MIT Press.
23. Pinker, S., *Formal models of language learning.* Cognition, 1979. **7(3)**: p. 217-283.
24. Pinker, S., *Language learnability and language development.* 1984, Cambridge, MA: Harvard University Press.

25. Edelman, G.M., *Topobiology : an introduction to molecular embryology*. 1988, New York: Basic Books. xv, 240.
26. Vargha-Khadem, F., D.G. Gadian, K.E. Watkins, A. Connelly, W. Van Paesschen, and M. Mishkin, *Differential effects of early hippocampal pathology on episodic and semantic memory [see comments] [published erratum appears in Science 1997 Aug 22; 277(5329):1117].* Science, 1997. **277**(5324): p. 376-80.
27. Bechtel, W. and A. Abrahamsen, *Connectionism and mind: An introduction to parallel processing in networks*. 1991: Cambridge, MA: Basil Blackwell.
28. Marcus, G.F., *Can connectionism save constructivism?* Cognition, 1998. **66**: p. 153-182.
29. Marcus, G.F., *Rethinking eliminative connectionism.* Cognitive Psychology, 1998. **37**(3): p. 243-282.
30. Marcus, G.F., *The algebraic mind: Integrating connectionism and cognitive science*. 2001, Cambridge, MA: MIT Press.
31. Marcus, G.F., U. Brinkmann, H. Clahsen, R. Wiese, and S. Pinker, *German inflection: The exception that proves the rule.* Cognitive Psychology, 1995. **29**: p. 186-256.
32. Marcus, G.F., S. Pinker, M. Ullman, J.M. Hollander, T.J. Rosen, and F. Xu, *Overregularization in language acquisition.* Monographs of the Society for Research in Child Development., 1992. **57**(4, Serial No. 228).
33. Spemann, H., *Embryonic development and induction*. Yale University. Mrs. Hepsa Ely Silliman memorial lectures [1933]. 1938, New Haven,: Yale University Press. xii, 401.
34. Wolpert, L., *The triumph of the embryo*. Repr. (with corrections) ed. 1992, Oxford England ; New York: Oxford University Press. vii, 211.
35. Levitt, P., *Molecular determinants of regionalization of the forebrain and cerebral cortex*, in *The new cognitive neurosciences*, M.S. Gazzaniga, Editor. 2000, MIT Press: Cambridge, Mass. p. 23-32.
36. Cruz, Y.P., *Mammals*, in *Embryology : constructing the organism*, S.F. Gilbert and A.M. Raunio, Editors. 1997, Sinauer Associates: Sunderland, MA. p. 459-489.
37. Dawkins, R., *The Blind Watchmaker*. 1987, NY: Norton.
38. Crowley, J.C. and L.C. Katz, *Development of ocular dominance columns in the absence of retinal input.* Nature Neuroscience, 1999. **2**(12): p. 1125-1130.
39. Verhage, M., *et al.*, *Synaptic assembly of the brain in the absence of neurotransmitter secretion [In Process Citation].* Science, 2000. **287**(5454): p. 864-9.
40. Katz, L.C., M. Weliky, and J.C. Crowley, *Activity and the development of the visual cortext: New perspectives*, in *The new cognitive neurosciences*, M.S. Gazzaniga, Editor. 2000, MIT Press: Cambridge, Mass. p. 199-212.
41. Jacob, F. and J. Monod, *On the regulation of gene activity.* Cold Spring Harbor Symposium on Quantitative Biology, 1961. **26**: p. 193-211.
42. Gilbert, S.F., *Developmental biology*. 5th ed. 1997, Sunderland, Mass.: Sinauer Associates. 1 v. (various pagings).
43. Gehring, W.J., *Master control genes in development and evolution : the homeobox story*. The Terry lectures. 1998, New Haven: Yale University Press. xv, 236.
44. Halder, G., P. Callaerts, S. Flister, U. Walldorf, U. Kloter, and W.J. Gehring, *Eyeless initiates the expression of both sine oculis and eyes absent during Drosophila compound eye development.* Development, 1998. **125**(12): p. 2181-91.
45. Halder, G., P. Callaerts, and W.J. Gehring, *Induction of ectopic eyes by target expression of the eyeless gene in Drosophila.* Science, 1995. **267**: p. 1788-1792.
46. Sperry, R.W., *Chemoaffinity in the orderly growth of nerve fiber patterns and connections.* Proceedings of the National Academy of Sciences, 1963. **50**: p. 703-710.
47. Brown, A., *et al.*, *Topographic mapping from the retina to the midbrain is controlled by relative but not absolute levels of EphA receptor signaling.* Cell, 2000. **102**(1): p. 77-88.