The Neural Basis of Cognitive Development: A Constructivist Manifesto

Steven R. Quartz and Terrence J. Sejnowski

Abstract *How do minds emerge from developing brains? According to 'neural constructivism,' the representational features of cortex are built from the dynamic interaction between neural growth mechanisms and environmentally derived neural activity. Contrary to popular selectionist models that emphasize regressive mechanisms, the neurobiological evidence suggests that this growth is a progressive increase in the representational properties of cortex. The interaction between the environment and neural growth results in a flexible type of learning: 'constructive learning' minimizes the need for prespecification in accordance with recent neurobiological evidence that the developing cerebral cortex is largely free of domain-specific structure. Instead, the representational properties of cortex are built by the nature of the problem domain confronting it. This uniquely powerful and general learning strategy undermines the central assumption of classical learnability theory, that the learning properties of a system can be deduced from a fixed computational architecture. Neural constructivism suggests that the evolutionary emergence of neocortex in mammals is a progression toward more flexible representational structures, in contrast to the popular view of cortical evolution as an increase in innate, specialized circuits. Human cortical postnatal development is also more extensive and protracted than generally supposed, suggesting that cortex has evolved so as to maximize the capacity of environmental structure to shape its structure and function through constructive learning.*

1 Introduction

If cognitive and neural processes really do interact, then an added degree of complexity arises in analyzing development, for learning may induce large changes in the very structures involved in learning. This complicates matters, because now learning can alter what can be subsequently learned. To borrow a phrase from physics, systems with time-dependent properties are said to be 'nonstationary.'[1] This term applies to the 'learning mechanism' or 'acquisition device,' roughly defined as the set of processes and structures that transform input data into the steady state representing the mature competence. The nonstationary learner is thus one in which learning causes large scale changes to its basic mechanisms. Pursuing a popular though perhaps misleading metaphor, learning causes major changes to the underlying hardware. Developmental theorists are accordingly confronted with the challenge of characterizing a nonstationary learning mechanism. [537]

1.1 A neurocomputational approach to nonstationarity

(...) We will explore three possible measures for representational change: synaptic numbers, axonal arborization, and dendritic arborization. Applying the above methodology, we will then examine the neurobiology of these changes during development, the involvement of intrinsic and extrinsic factors in regulating them, and their adequacy as indicators of representational complexity.

Two themes emerge from finding a structural measure of representational complexity: 1) development is a progressive increase in the structures underlying representational complexity,

---

[1] S. Pinker, 'Formal models of language learning,' **Cognition** 1 (1979): 217–83.

and 2) this increase depends on interaction with a structured environment to guide development. These form the basis of neural constructivism, the developmental theory we present. This name reflects the Piagetian view that there is an active interaction between the developing system and the environment in which it is embedded. Like Piaget's theory, ours too emphasizes the constructive nature of this interaction, with representational structures progressively added during development. Our primary concern in this target article, however, is to examine the neural processes regulating structural change and their implications for representational change. In particular, dendritic development fulfills important requirements for a nonstationary learning mechanism, suggesting how dendritic development under the influence of environmentally derived activity conforms to cognitive schemes for the construction of mental representations. Although cortical development is often thought limited primarily to the first two years of life, this development is far more extensive, prolonged, and progressive. This indicates that the extent of human cortical postnatal development has been widely underestimated along with its potential role in building mental representations under the guidance of environmentally derived information.

This environmentally-guided neural circuit building is a form of learning, 'constructive learning,' a unique and powerful response to the learning pressures confronting a developing system undermining the central assumptions of classical formal learning theory. The central problem confronting a cognitive system is to find an appropriate class of representations for specific problem domains. Many views suppose that these representations have to be pre-existing, but constructive learning builds these under the influence of the environment, acting alongside the general constraints that are imposed by the neural architecture. As a result, it offers powerful learning abilities while minimizing the need for domain-specific prespecification and so avoiding the heavy burden that nativism places on genetic mechanisms. Ours is not a return to tabula rasa learning, however; learning is a dynamic interaction between a changing, structured environment and neural mechanisms.

Ours is not a return to *tabula rasa* learning, however; learning is a dynamic interaction between a changing, structured environment and neural mechanisms. The neural machinery is extensively shaped by activity stemming from the environment, while its intrinsic properties also constrain this modulation and play an indispensable role in shaping the resulting structures. This interaction, however, is sufficient to determine the mature representational properties of cortex with no need for domain-specific predispositions somehow embedded a priori in the recipient cortex. As a consequence, this makes the relation between environmental changes – whether natural or cultural – and brain structure a direct one. This suggests an evolutionary perspective as a progression to more flexible representations, in contrast to evolutionary psychology.[2] (…) [538-539]

2 Measures of representational complexity The brain is above all a representational device.[3] By 'representation' we mean the neural encoding of environmentally-derived information and transformations resulting from the application of mental operations. The best-known account of mental representation is in terms of language-like systems with a primitive lexicon and syntactic rules corresponding to mental operations.[4] Neural networks offer alternative representational encodings, particularly distributed representations.[5] Although representational complexity can be defined for both types of representations,[6] neural network measures depend on structural properties, making the relationship between complexity and structure a direct one. [539]

Selectionism's learning mechanism operates by reducing an overly complex set of representations to a subset in the mature state. In an important selectionist statement, Changeux and Dehaene bluntly put it this way: 'activity does not create novel connections, but, rather, contributes to the elimination of pre-existing ones.'[7] Indeed, completing the evolutionary analogy, allowing for directed growth in development would be akin to allowing for Lamarkian processes in evolution. For selectionism, then, development marks a reduction in representational complexity. In contrast, neural constructivism sees development as a progressive increase in representational complexity. Since these differ over changes in representational complexity, it is important to ask, what do these accounts use as a measure of representational complexity?

There are three main candidates for this measure: synaptic numbers, axonal arborization, and dendritic arborization. Although these are all related, it is worthwhile to examine them separately, as a change in one does not necessarily imply a change in the others. In the remainder of this section, we consider the support for changes in these measures along with their adequacy as measures of representational power. [540]

2.1.1 Synaptic numbers over development (…) With a larger sample, they found that synaptic density reached a peak around two months of age and did not begin to decline until puberty. Even then, the decline was a gradual one, lasting until 20 years of age (their last sample point). It is hard to reconcile this finding with the claim that selective synapse elimination underlies cognitive development since no such process appears to operate in the prefrontal cortex during the most formative years of cognitive development. Indeed, an additional complication comes from studies showing that brain volume increases during this period, particularly in prefrontal cortical

---

[2] For example: J. H. Barkow, L. Cosmides and J. Tooby, ed.,**The adapted mind: Evolutionary psychology and the generation of culture** (New York: Oxford University Press, 1992); and S. Pinker, **The language instinct: how the mind creates language** (New York: W. Morrow & Company, 1994).

[3] For a detailed discussion, see: Z. Pylyshyn, (1984) **Computation and cognition: Toward a foundation for cognitive science** (Cambridge, MA: Bradford Books, 1984); and P. Churchland and T.J. Sejnowski, **The computational brain** (Cambridge, MA: MIT Press, 1992).

[4] N. Chomsky, 'Rules and representations,' **Behavioral and Brain Sciences** 3 (1980): 1 – 61.

[5] P. Churchland and T.J. Sejnowski, **The computational brain** (Cambridge MA: The MIT Press, 1994)

[6] S.R. Quartz, 'Nativism, neural networks, and the plausibility of constructivism,' **Cognition** 48 (1993): 123–44.

[7] J.P. Changeux and S. Dehaene, 'Neuronal models of cognitive functions,' **Cognition** 33 (1989): 82.

areas.[8] Even if synaptic density remains constant, these volume increases imply synapse addition. [540]

2.1.2 Environmental conditions and synaptic numbers This paradigm has also been used to examine whether the number of synapses changes in mature forms of plasticity[9] have shown that the formation of new synapses accompanies motor learning tasks in behaving animals. To distinguish between the possibility that motor activity and not motor learning caused the increases in synaptic number, a control group underwent exercise on tasks that were easily mastered and required little learning. In these cases, there were no significant increases in cortical synapses. There was, however, angiogenesis (increased capillary support), as one would expect from increased demands. In contrast, cerebellar Purkinje cells in rats that learned complex locomotor tasks showed a 25% increase in synaptic numbers. [541]

2.1.3 Synapse number as a measure of circuit complexity In many real neurons the spatial arrangement of pre- and postsynaptic elements is thought to be crucial to their response properties. One reason for this is the presence of active conductance properties in the cell's membrane; these amplify or otherwise change the incoming signal in nonlinear ways. Nonlinear dendritic conductance properties, now well established,[10] shift the emphasis from absolute synaptic numbers to the spatial arrangement of synapses and the branching patterns of postsynaptic structures. This suggests that axonal and dendritic arborization may be more central determinants of circuit complexity than absolute synaptic numbers. [541-542]

2.2 Axonal arborization Patterns of axonal arborization have been widely used as a measure of representational complexity. Indeed, studies of axonal arborization have largely contributed to selectionism. [542]

2.2.1 Axonal arborization over development The electrophysiological, developmental, and anatomical study of early vision is central to modern neuroscience. In particular, the retinogeniculocortical pathway, and especially the development of ocular dominance (OD) columns of layer IV… (while) ocular dominance columns are alternating, anatomically defined regions of input from primarily one eye… two hypotheses regarding their development have been suggested. One, conforming to selectionism, emphasizes two phases in OD development: a period of exuberant axonal growth followed by selective axonal pruning. The other, more constructivist, hypothesis emphasizes the general expansion of axon collaterals alongside selective pruning.

Hence, although selectionism has dominated modeling work, the experimental work points to more balance between both selective elimination and selective growth, with a progressive increase in the complexity of the arbors underlying OD (right eye) formation. As Antonini and Stryker conclude, '[N] development thus appears to involve both selective elimination of widely extended branches and considerable growth and elaboration.'[11] [543]

2.3 Dendritic development As we noted in section 1, nonstationarity, while perhaps increasing the learning capacities of a developing system, introduces a number of complications. The main reason for this is that changes to the underlying architecture can be unwieldy. Even small architectural changes can have severe consequences for the system's overall performance. One way this may happen is if changes to an architecture undo previous learning by reconfiguring structural elements that represented already acquired knowledge (a phenomenon sometimes referred to as 'catastrophic interference'). Such a process could also have negative consequences for learning if it introduced large-scale changes rather than incremental ones. For example, large-scale changes could make learning too sensitive to small details of training (resulting in overfitting) and would undo the dependency on previous states that makes learning incremental, and thereby stable.

These concerns lead to the following two related conditions that a nonstationary mechanism must satisfy:

1 The Locality Condition. The addition of structure must be at the appropriately local scale and must not result in wholesale changes in representation with each new elemental change;

2 The Stability Condition. Under normal circumstances, local changes must not undo previous learning. [545]

2.3.1 Dendrites as learning structures Dendritic growth satisfies the locality condition in that, unlike the retraction of an axon which might alter the connectivity of thousands of cells, dendritic segments are local processing elements, whose elimination will not cause large-scale changes in the overall pattern of connectivity. Hence, the constructivist learning scheme operates at the level of individual dendritic segments, rather than the whole cell, thereby building in these conditions. Hence, the constructivist learning scheme operates at the level of individual dendritic segments, rather than the whole cell, thereby building in these conditions.

It is also important that dendrites grow much more slowly than do axons. Axon's grow at

---

8 See: A.S. Dekaban and D. Sadowsky, 'Changes in brain weights during the span of human life: Relation of brain weights to body heights and body weights,' **Annals of Neurology** 4 (1978): 345–56; T.L. Jernigan, S.L. Archibald, M.T. Berhow, E.R. Sowell, D.S. Foster, and J.R. Hesselink, 'Cerebral structure on MRI, Part I: Localization of age-related changes,' **Biological Psychiatry** 29 (1991): 55–67 (and references therin).

9 J.E. Black, K.R. Isaacs, B.J. Anderson, A.A. Alcantara and W.T. Greenough, 'Learning causes synaptogenesis, whereas motor activity causes angiogenesis, in cerebellar cortex of adult rats,' **Proceedings of the National Academy of Sciences USA** 87 (1990): 5568–72.

10 G.J. Stuart and B. Sakmann, 'Active propagation of somatic action potentials into neocortical pyramidal cell dendrites,' **Nature** 367 (1994): 69–72; and Z.F. Mainen, J. Joerges, J.R. Huguenard and T.J. Sejnowski, 'A model of spike initiation in neocortical pyramidal neurons,' **Neuron** 15 (1995): 1427–39.

11 A. Antonini and M.P. Stryker, 'Development of individual geniculocortical arbors in cat striate cortex and effects of binocular impulse blockade,' **Journal of Neuroscience** 13 (1993): 3572.

approximately 500 microns/day compared to 15-35 microns/day for dendrites.[12] This suggests that the two are sensitive to statistical structure at different timescales and that dendrites are extracting this structure on a much longer temporal scale.

A more important and general reason for examining the growth of dendrites in relation to the construction of representations is that dendrites are the primary receptive surface of a neuron. Moreover, since dendrites do not just conduct passively, but amplify the incoming signal in nonlinear ways, their processing properties make them central to how information is processed by neural systems. (...) [545]

2.3.2 **Dendritic arborization** (...) As dendrites grow, the integration of synaptic activity is altered in ways that depend on the geometry of the branches and the placement of synapses (Segev et al., 1995 contains an excellent collection of classic and contemporary views on dendritic function).[13] (...) It should be borne in mind that although there is a reduction in synapses per micron of dendritic length, the increasing dendritic surface area of these cells implies an overall increase in the number of synapses. [545]

2.3.3 **Environmental conditions and dendritic development** Whereas studies such as Valverde's illustrate the dependence of dendritic form on afferent pathways, the study of Mooney et al. (1992) illustrates the striking malleability of developing dendrites. Mooney et al. (1992) examined the effects of neonatal enucleation on the dendritic morphology of superior collicular (SC) neurons.[14] Like Valverde, they found that the dendrites of SC neurons were redirected toward sources of residual input, the deep layer of the SC, whose input is from somatosensory axons. But when they examined these cells' physiological response properties they found that a majority of them were no longer visually responsive, as in the normal case, but now had somatosensory response properties.

This leads to an important result: in the enucleated animal, the dendrites redirected their growth to find active afferents; where these were of a different modality, the cells changed their response properties to reflect this residual source. So, these response properties corresponding to the cell's function are not predetermined, but depend on interacting with the information modality latent in the pattern of incoming activity. [546]

2.3.4 **Directed dendritic development and patterns of activity** The finding that dendrites actively seek out incoming activity and shape their responses to mirror that activity calls for a closer look. (...)

What is the signal that regulates this development? As Katz et al. (1989) note, one likely source of this signal derives from correlated activity within a column, since it originates from one eye, but is discontinuous at the borders between stripes from different eyes.[15] This change in correlated activity might therefore underlie the bias away from the neighboring region if the postsynaptic cell maximized the amount of correlated input it received. What would the role of such a developmental signal be? The most direct role would be in the development of the response properties of the cell. Cells [of a specific] layer are almost exclusively monocular, that is, they respond to stimulation from only one eye. So, by maximizing correlated input and avoiding uncorrelated input, a cell's dendrites would come to arborize within a single column, and would thus help to maintain monocularity. In addition, by taking advantage of a signal that was intrinsic to the afferents, this organization would come about without the need for pre-specifying it. Similar themes of dendritic development in the somatosensory cortex have also been observed.[16]

(...) As in the Katz et al. (1989)[17] study, it is the degree of correlation in the afferent activity rather than simply the presence of activity that underlies dendritic organization.[18]

An interpretation of these results is that dendritic segments function as detectors of correlated activity and grow preferentially in such regions. Support comes from Tieman and Hirsch's (1982) finding that exposure to lines of only one orientation during development has specific effects on dendritic development.[19] The dendritic field orientations of cells from cats raised with exposure to lines of a single orientation were significantly elongated in conformity with this shift in the visual environment.

An insight from this study is that a dendritic tree samples its input space actively in response to the environmental structure. A similar result has been obtained for layer IV stellate cells by Coleman et al. (1981), who suggest: 'If an alteration of the spatio-temporal pattern of the afferent activity is sufficient to lead to dendritic alterations during development, this implies that dendritic

---

[12] See: H.B.M. Uylings, C.G. Van Eden, J.G. Parnavelas and A. Kalsbeek, 'The prenatal and postnatal development of the rat cerebral cortex,' in **The cerebral cortex of the rat**, eds. B. Kolb & R. C. Tees (Cambridge, MA: MIT Press, 1990); and M.J. Katz, E.B. George and L.J. Gilbert, 'Axonal elongation as a stochastic walk,' **Cell Motility** 4 (1984): 351–70.

[13] I. Segev, J. Rinzel and G.M. Shepherd, **The theoretical foundations of dendritic function: Selected papers by Wilfrid Rall with commentaries** (Cambridge, MA: MIT Press, 1995).

[14] R.D. Mooney, M.M. Nikoletseas, T.D. King, S.V. Savage, M.T. Weaver and R.W. Rhoades, 'Structural and functional consequences of neonatal deafferentation in the superficial layers of the hamster's superior colliculus,' **Journal of Comparative Neurology** 315 (1992): 398–412.

[15] L.C. Katz, C.D. Gilbert and T.N. Wiesel, 'Local circuits and ocular dominance columns in monkey striate cortex,' **Journal of Neuroscience** 9 (1989): 1389–99.

[16] W.T. Greenough and F.L. Chang, 'Plasticity of synapse structure and pattern in the cerebral cortex,' in **Cerebral cortex**, vol. 7: **Development and maturation of cerebral cortex**, ed. A. Peters & E. G. Jones (New York: Plenum Press, 1988).

[17] Katz, et al., 'Local circuits and ocular dominance columns in monkey striate cortex.'

[18] Differences in the degree of rostral-caudal dendritic bias between normal and stripe-induced cells support the view that this development involves progressive growth rather than elimination of exuberant structure, see L. C. Katz and M. Constantine-Paton, 'Relationships between segregated afferents and postsynaptic neurones in the optic tectum of three-eyed frogs,' **Journal of Neuroscience** 8 (1988): 3160–80 (3178). The conclusion according to L. C. Katz et al. is that, 'the pattern of afferent segregation has played a significant role in shaping the structure of the postsynaptic dendritic field of cortical neurons,' 'Local circuits and ocular dominance,' 1393. (Editorial comment: this refers to note 4 in the original text).

[19] S.B. Tieman, S. Hirsch, 'Exposure to lines of only one orientation modifies dendritic morphology of cells in the visual cortex of the cat,' **Journal of Comparative Neurology** 211 (1982): 353–62.

trees may develop in a form that will optimize spatio-temporal summation for the postsynaptic neuron.'[20] [546-548]

3 Directed dendritic development and representational change

What sort of representations does the brain use? One of the most important principles of cortical representation involves 'geometric principles of information processing design.'[21] Mead states:

> Computation is always done in the context of neighboring information. For a neighborhood to be meaningful, nearby areas in the neural structure must represent information that is more closely related than is that represented by areas further away. Visual areas in the cortex that begin the processing sequence are mapped retinotopically. Higher-level areas represent more abstract information, but areas that are close together still represent similar information. It is this map property that organizes the cortex such that most wires can be short and highly shared; it is perhaps the single most important architectural principle in the brain.

From this principle, the physical structure of a neural area corresponds to a representational space. In this representational space, nearby things are more closely related semantically than things that are far apart. This map property is extremely powerful as a representational system. When brain areas can communicate, increasingly rich representations can be successively built up. Each area is a layer in an increasingly abstract feature space. Just as information in a map is held by such spatial properties as physical distance, the physical structure of cortex encodes information. With geometric principles of information processing the information is held in the three-dimensional pattern of neural connectivity. As constructive factors play a central role in building this physical structure, they also shape the representational properties of cortex. Building neural circuits with directed growth thereby builds the brain's representational properties.

These spatial properties of representation are largely lost in the traditional connectionist network because of the way the connectionist neuron integrates information, typically summing its input and sending a (perhaps graded) output if some threshold is exceeded. This makes the entire cell the basic computational unit. In contrast, biological neurons are thought to segregate into sub-regions that function as autonomous processors. Local dendritic segments might be the brain's basic computational units.[22] Dendrites are not simple signal integrators with passive conductance properties, as in classical cable models.[23] Imaging studies have found that some dendritic systems (for example, CA1 hippocampal neurons) have a heterogeneous distribution of voltage-gated Ca2+ channels, suggesting nonlinear membrane properties.[24] Intradendritic recordings in these cells likewise reveal strong nonlinearities in their electrical properties. [25] (...)

Nonlinear properties give a cell more computational power than traditionally thought. [26] A cell with active dendritic segments can perform the nonlinear discrimination that requires a hidden-layer network of connectionist neurons. The spatial properties of a cell may also determine many of its functional properties. Connecting this back with our earlier discussion of geometric principles of information processing, when such a cell is embedded in a representational space, its spatial structure takes on additional significance. A phenomenon referred to as the 'clustering' of related inputs onto dendritic segments results in a pattern of termination mirroring the informational structure of input: electrotonically close synapses encode common features of the input space and effectively fire the cell.[27] The presence of cluster-encoded features significantly alters both the representational properties of cortex and its computational power. [549]

3.2 Hierarchical representation construction

Much of non visual cortical development, in contrast, displays an extensive and more protracted development. Cells in frontal cortex are far

---

[20] P.D. Coleman, D.G. Flood, M.C. Whitehead and R.C. Emerson, 'Spatial sampling by dendritic trees in visual cortex,' **Brain Research** 214 (1981): 19.

[21] R. Durbin and G.J. Mitchison, 'A dimension reduction framework for understanding cortical maps,' **Nature** 343 (1990): 644–47; Churchland and Sejnowski, **The computational brain**; C. Mead, **Analog VLSI and neural systems**. (Reading, MA: Addison-Wesley, 1989); and G.J. Mitchison and R. Durbin, 'Optimal numberings of an N X N array,' **S.I.A.M. Journal on Algebraic and Discrete Methods** 7 (1986): 571–82. (Editorial comment: extended citation refers to Mead, 277).

[22] See also: C. Koch, T. Poggio and V. Torre, 'Retinal ganglion cells: A functional interpretation of dendritic morphology,' **Philosophical Transactions of the Royal Society of London, B: Biological Sciences** 298 (1982): 227–63, and 'Nonlinear interactions in a dendritic tree: Localization, timing, and role in information processing,' **Proceedings of the National Academy of Sciences USA** 80 (1983): 2799–802; G. M. Shepherd and R.K. Brayton, 'Logic operations are properties of computer-simulated interactions between excitable dendritic spines,' **Neuroscience** 21 (1987): 151–65; B.W. Mel, 'NMDA-based pattern discrimination in a modeled cortical neuron,' **Neural Computation** 4 (1992):502–17, and 'Information processing in an excitable dendritic tree,' **CNS Memo 17**, Computational and Neural Systems Program, California Institute of Technology(1992), and 'Information processing in dendritic trees,' **Neural Computation** 6 (1994): 1031–85; S.W. Jaslove, 'The integrative properties of spiny distal dendrites,' **Neuroscience** 47 (1992): 495–519; I. Segev, et al., **The theoretical foundations of dendritic function.**

[23] W. Rall, 'Theoretical significance of dendritic trees for neuronal inputoutput relations,' in **Neural theory of modelling**, ed. R.F. Reiss (Stanford: Stanford University Press, 1964).

[24] O.T. Jones, D.L Kunze and K.J. Angelides, 'Localization and mobility of omega-conotoxin-sensitive Ca21 channels in hippocampal CA1 neurons,' **Science** 244 (1989): 1189–93; and W.G. Regehr, J.A. Connor and D.W. Tank, 'Optical imaging of calcium accumulation in hippocampal pyramidal cells during synaptic activation,' **Nature** 341 (1989): 533–36.

[25] R. K. Wong, D. A. Prince and A.I. Basbaum, 'Intradendritic recordings from hippocampal neurons. **Proceedings of the National Academy of Sciences USA** 76 (1979): 986–90; and L.S. Bernardo, L.M. Masukawa and D.A. Prince, 'Electrophysiology of isolated hippocampal pyramidal dendrites,' **Journal of Neuroscience** 2 (1982): 1614–22.

[26] J.A. Feldman and D.H. Ballard, 'Connectionist models and their properties. **Cognitive Science** 6 (1982): 205–54; R. Durbin and D.E. Rumelhart, 'Product units: A computationally powerful and biologically powerful extension to backpropagation networks,' **Neural Computation** 1 (1989): 133–42; B.W. Mel and C. Koch, 'Sigma-Pi learning: On radial basis functions and cortical associative learning,' in **Advances in neural information processing systems**, ed. D. S. Touretzky (San Francisco: Morgan Kaufmann, 1990); and C. Koch and T. Poggio, 'Multiplying with synapses and neurons,' in **Single neuron computation**, ed. T. McKenna (Boston: J. Davis, S. Zornetzer, Academic Press, 1992).

[27] Mel, 'NMDA-based pattern discrimination in a modeled cortical neuron,' and 'Information processing in an excitable dendritic tree.'

slower to develop and undergo the majority of their growth after two years of age.[28] In addition, the extent of their postnatal development is dramatic – they grow to over thirty times their dendritic length at birth. Scheibel (1993) likewise reports a long period of dendritic development in Broca's area in which mature forms emerge only after six to eight years.[29] Why, then, is human non-visual cortical development so slow to develop and so extensive?

Our view is that the human brain's development is a prolonged period in which environmental structure shapes the brain activity that in turn builds the circuits underlying thought. In place of pre-wired modules, patterned activity builds up increasingly complex circuits, with areas staging their development. Cortical areas further away from the sensory periphery wait in anticipation of increasingly complex patterns of activity resulting from development in lower areas. As this development proceeds, areas of the brain become increasingly specialized for particular functions, reflecting a cascade of environmental shaping. Some brain circuits close to the sensory periphery, such as in our early visual system, are in place by six months of age; but those in language areas, further away from the sensory periphery, do not begin to complete their development until the eighth year of life. [550]

3.3 What is the role of regressive events in development? The evidence we have examined demonstrates that the popular view of development as largely a regressive event must be reconsidered. We suggest that regressive events are simply the consequence of reduced neural specificity, as indicated by the counter-evidence to Sperry's chemoaffinity hypothesis. Any theory, whether selectionist or constructivist, that rejects a strong view of neural specificity will thus need to posit regressive events. If cells do not bear nearly unique molecular addresses, then stochastic sampling mechanisms must be posited. These will by their very nature introduce some structure into a system that will later be eliminated. Neural constructivism allows these sampling mechanisms to be directed, but they are still stochastic. Structural elimination, or error-correction, are likewise required, but this does not mean that error-correcting processes are the only developmental mechanisms, or that developmental selection occurs only among intrinsically generated structures.[30] Rather, selection is only one kind of process in a dynamic interaction between environmentally derived activity and the neural growth mechanisms that activity regulates. This changes the emphasis from synaptic elimination to synaptic turnover. New techniques that allow the same structures to be examined successively over time[31] reveal that constructive and selective events co-occur as required for directed sampling mechanisms. [550-551]

4 A learning-theoretic approach to development

4.1.1 Adding a neurobiological constraint Once we are talking about natural systems, it is worthwhile to raise a neurobiological constraint. So far, this discussion has proceeded as though the only significant factors were learning-theoretic pressures, but it is particularly important to consider whether the views coming out of learning theory are consistent with neurobiological constraints on development. For natural systems, the constraint that a learning theory should be consistent with underlying neural mechanisms has been severely underestimated. Indeed, in our opinion this biological constraint has equal footing with the learning-theoretic one and both must be viewed as complementary constraints that developmental systems must satisfy.

As suggested by Quartz and Sejnowski (1994),[32] the view that strong, domain-specific knowledge is built into cortical structures runs into severe difficulties from developmental neurobiological evidence. Although we will not review that material in detail here, recent experiments on heterotopic transplants,[33] cross modal rewiring[34] and clonal analysis of cell migration[35] all establish that

---

[28] J.P. Schade and W.B. van Groenigen, 'Structural organization of the human cerebral cortex. I. Maturation of the middle frontal gyrus,' **Acta Anatomica** 47 (1961): 72–111.

[29] A.B. Scheibel, 'Dendritic structure and language development,' in **Developmental neuro-cogniton: Speech and face processing in the first year of life**, ed. B. de Boysson-Bardies (Dordrecht: Kluwer Academic Publishers, 1993).

[30] A widely used metaphor to describe this process is that of error-correction. It should not be assumed, however, that the exuberant connections are strictly in error, since they may serve a useful purpose in instances in which a changes in connectivity is required (as in the case of blindness). (Editorial comment: this refers to note 7 in the original text).

[31] N.A. O'Rourke, H.T. Cline and S.E. Fraser, 'Rapid remodeling of retinal arbors in the tectum with and without blockade of synaptic transmission,' **Neuron** 12 (1994): 921–34; N.A. O'Rourke and S.E. Fraser, 'Dynamic aspects of retinotectal map formation revealed by a vital-dye fiber-tracing technique,' **Developmental Biology** 114 (1986): 265–76, and 'Dynamic changes in optic fiber terminal arbors lead to retinotopic map formation: An in vivo confocal microscopic study,' **Neuron** 5 (1990): 159–71; D. Purves and R.D. Hadley, 'Changes in the dendritic branching of adult mammalian neurones revealed by repeated imaging in situ,' **Nature** 315 (1985): 404–06; D. Purves, R.D. Hadley and J.T. Voyvodic, 'Dynamic changes in the dendritic geometry of individual neurons visualized over periods of up to three months in the superior cervical ganglion of living mice,' **Journal of Neuroscience** 6 (1986): 1051–60; D. Purves, J. Voyvodic, L. Magrassi and H. Yawo, 'Nerve terminal remodeling visualized in living mice by repeated examination of the same neuron,' **Science** 238 (1987): 1122–26; and M.E. Dailey and S.J. Smith, 'The dynamics of dendritic structure in developing hippocampal slices,' **Journal of Neuroscience** 16 (1996): 2983–94.

[32] S.R. Quartz and T.J. Sejnowski, 'Beyond modularity: Neural evidence for constructivist principles in development,' **Behavioral and Brain Sciences** 17 (1994): 725–26.

[33] B.B. Stanfield and D.D.M. O'Leary, 'Fetal occipital cortical neurons transplanted to the rostral cortex can extend and maintain a pyramidal tract axon,' **Nature** 313 (1985): 135–37; B.L. Schlaggar and D.D.M. O'Leary, 'Potential of visual cortex to develop an array of functional units unique to somatosensory cortex,' **Science** 252 (1991): 1556–60; and reviewed in D.D.M. O'Leary, B.L. Schlaggar and B. B. Stanfield, 'The specification of sensory cortex: lessons from cortical transplantation,' **Expermental Neurology** 115 (1992):121–26.

[34] D.O Frost, 'Anomalous visual connections to somatosensory and auditory systems following brain lesions in early life,' **Brain Research** 255 (1982): 627–35; M. Sur, P.E. Garraughty and A.W. Roe, 'Expermentially induced visual projections into auditory thalamus and cortex. **Science** 242 (1988): 1437–41; A.W. Roe, S.L Pallas, S. J. Hahm and M. Sur, 'A map of visual space induced in primary auditory cortex,' **Science** 250 (1990): 818–20; A.W. Roe, S.L. Pallas, Y.H. Kwon and M. Sur, 'Visual projections routed to the auditory pathway in ferrets: Receptive fields of visual neurons in primary auditory cortex,' **Journal of Neuroscience** 12 (1992): 3651–64; S.L. Pallas, A.W. Roe and M. Sur, 'Visual projections induced into the auditory pathway of ferrets. Novel inputs to primary auditory cortex (AI) 59 from the LP/pulvinar complex and the topography of the MGN-AI projection,' **Journal of Comparative Neurology** 298 (1990): 50–68; C.J. Schatz, 'How are specific connections formed between thalamus and cortex?,' **Current Opinion in Neurobiology** 2 (1992): 78–82; and reviewed in M. Sur, S.L. Pallas and A.W. Roe, A. W. 'Cross-modal plasticity in cortical development: Differentiation and specification of sensory neocortex,' **Trends in Neuroscience** 13 (1990): 227–33.

the regional characteristics of mature cortex depend fundamentally on interaction with afferent input. While the cortex is not a tabula rasa, as there may be a common basic circuitry and repetitive arrays of cell types,[36] it is largely equipotential at early stages (in agreement with studies on cortical plasticity and early brain damage.[37]

Consistent with this view, O'Leary (1990)[38] refers to the immature cortex as protocortex, which shares a common laminated structure, cell types, and basic circuitry but which diminishes the need for prespecification. It is the differing pattern of afferent activity, reflective of different sensory modalities, that confers area-specific properties onto the cortex – not predispositions that are somehow embedded in the recipient cortical structure. In addition, the fact that many of these processes operate before birth, as in the case of spontaneous visual activity,[39] suggests that cortical specification could begin by the very mechanisms that will be used postnatally through interaction with an environment. Hence, the fact that various regions of cortex receive different patterns of afferent termination and activity seems to be the prime determinant of specialized cortical functions. A system in which the cortex is 'enslaved by the periphery' has a number of clear advantages in terms of responding flexibly to varying environmental pressures, plasticity, and changing body size.[40] In the previous section, we tried to suggest how this interaction between developing cortex and environmentally derived activity builds up the neural circuits underlying cognition.

Adding the neurobiological constraint to the learning-theoretic one results in yet another impasse. From the perspective of learning theory, it appeared that the only response to the learnability problem was to build in much of the problem domain a priori in the form of highly specialized structures. Yet, from the perspective of biological constraints it appeared that cortical structures do not build in this knowledge, but rather allow both pre- and post-natal activity to determine features of the cortex. In the following section, we suggest that the neural constructivism offers a powerful means of escaping this impasse. [552-553]

4.2 Constructive learning The constructivist learner shows its advantages here. It does not involve a search through a pre-defined hypothesis space, and so it is not one of selective induction (also known as model-based estimation, or parametric regression). Instead, the constructivist learner builds its hypothesis space as it learns. This has shifted the problem from one of parameter estimation to a nonparametric regime. We must accordingly ask, what is the effect of allowing a system to add structure – to build representations – as it learns?

Here again nonstationarity poses a challenge since we are asking about the effects of building representations according to the features of the learning problem. Neural network research has been particularly helpful in characterizing this sort of nonstationarity because the close relation between a network's architecture and its representational properties provides a constrained framework for looking at representational change during learning.

An increasingly sophisticated literature on the formal properties of neural networks has emerged. For example, a number of general results on the approximation properties of neural networks have been established.[41] From a nonparametric framework,[42] has demonstrated that a network that adds units at an appropriate rate relative to its experience is what statisticians call a consistent nonparametric estimator. This asymptotic property means that it can learn essentially any arbitrary mapping. The intuition behind this result, which will play a central role in characterizing constructive learning, follows a general nonparametric strategy: slowly increase representational capacity by reducing bias at a rate that also reduces variance. Since network bias depends on the number of units, as a network grows, its approximation capacities increase. The secret is regulating the rate of growth so that variance's contribution to error does not increase. (...)

The general strategy of constructivist learning is this. Rather than start with a large network as a guess about the class of target concepts, avoid the difficulties associated with overparameterized networks by starting with a small network. The learning algorithm then adds appropriate structure according to some performance criterion and where it is required until a desired error rate is achieved. Since the construction of the learner's hypothesis space is sensitive to the problem domain facing the learner, this is a way of tailor making the hypothesis space to suit the demands of the problem at hand. This allows the particular structure of the problem domain to determine the connectivity and complexity of the network. Since the network has the capacity to respond to the structure of the environment in this way, the original high bias is reduced through increases in network complexity, which allows the network to represent more complex functions. Hence, the need to find a good representation beforehand is replaced by the flexibility of a system that can respond to the structure of some task by building its representation class as it samples that structure to learn any polynomial learnable class of concepts. Research on constructive algorithms has

---

[35] C. Walsh and C.L. Cepko, 'Clonally related cortical cells show several migration patterns,' **Science** 241 (1988): 1342–45, 'Widespread dispersion of neuronal clones across functional regions of the cerebral cortex. **Science** 255 (1992): 434–40, and 'Clonal dispersion in proliferative layers of developing cerebral cortex,' **Nature** 362 (1993): 632–35.

[36] See, O'Leary, et al., 'The Specification of sensory cortex: lessons from cortical transplantation.'

[37] For example, H. Neville, 'Neurobiology of cognitive and language processing: effects of early experience,' in **Brain maturation and cognitive development**, eds., K. R. Gibson & A. C. Peterson (New York: Aladine de Gruyter Press, 1991).

[38] D. D. M. O'Leary, 'Do cortical areas emerge from a protocortex?,' **Trends in Neurosciences** 12 (1990): 400–406.

[39] L. Maffei and L. Galli-Resta, 'Correlation in the discharges of neighboring rat retrial ganglion cells during prenatal life,' **Proceedings of the National Academy of Sciences USA** 87 (1990): 2861–64; M. Meister, R. Wong, D. Baylor and C.J. Shatz, 'Synchronous bursts of action potentials in ganglion cells of the developing mammalian retina,' **Science** 252 (1991): 939–43.

[40] See: Walsh and Cepko, 'Widespread dispersion of neuronal clones across functional regions of the cerebral cortex,' and 'Clonal dispersion in proliferative layers of developing cerebral cortex.'

[41] For example: G. Cybenko, 'Approximation by superpositions of a sigmoid function,' **Mathematics of Control, Signals, and Systems** 2 (1989): 303–14; K. Hornik, M. Stinchcombe and H. White, 'Multilayer feedforward networks are universal approximations,' **Neural Networks** 2 (1989): 359–66; and F. Girosi and T. Poggio, 'Networks and the best approximation property,' **Biological Cybernetics** 63 (1990): 169–76.

[42] H. White, 'Connectionist nonparametric regression: Multilayer feedforward networks can learn arbitrary mappings,' **Neural Networks** 3 (1990): 535–49.

become increasingly sophisticated, and the results with constructive learners are impressive.[43]

The research we have just examined indicates a fundamental distinction between the constructivist strategy and models of selective induction. For the latter to have any chance of learning, the network must build in much of the problem domain a priori. Besides the neurobiological implausibility of this strategy, there are more general reasons why using highly biased networks is not a sound strategy in the biological case. Primary among these is that the highly biased network will only work for the specified application, but if the nature of the application is not properly predicted, the network will be a poor performer. Hence, tailor-making network architectures to suit the particular demands of some problem domain may be a useful heuristic strategy for artificial networks whose problem space is defined, or at least delimited, in advance by the designer. Biological learners, however, face an additional problem: not only is the problem space not defined beforehand, it is changing – the environment is highly nonstationary. Systems that are highly specialized for the anticipation of a particular problem domain will fail in the event of significant changes to that domain. The upshot is that specialization may bring efficiency, but it comes at the expense of flexibility. Although most natural systems are only confronted with ecological change, human cognition requires highly flexible and adaptive representations to accommodate both cultural and technological innovations. We doubt that the pace of this change can be met by a representational scheme requiring a major intrinsic specification. [553-554]

5 Conclusions

Although psychologists and neurobiologists both study development, communication and collaboration between fields have been limited. Reasons for this vary. Until recently, there was a lack of pertinent neurobiological data. In addition, reductive works such as Lenneberg (1967)[44] viewed advances in the biological basis of development as lessening the cognitive contribution. So, where connections were made, they reinforced the opposition of neural and cognitive descriptions of development, an opposition that was perhaps most strongly made in the functionalist contention that neural descriptions were irrelevant for cognitive explanations (the so-called arguments from 'multiple instantiability').

The extent and duration of large-scale brain changes during development have also been underappreciated. Whereas many researchers believe that the major events in brain development end by 2 years of age, the evidence we have reviewed illustrates these continue well past the first decade of life. Rather than being strictly reductive, neural constructivism points to the interaction between cognitive and neural processes in development, suggesting that cognitive and neural levels of description will need to be integrated into a single explanatory framework to explain this prolonged growth. Neural constructivism thus provides a meeting ground for cognitive scientists and neuroscientists. Although we are only beginning to understand how the world and brain interact to build the mind, the story that is unfolding already makes nativist theories appear implausible. What lies ahead promises to be an exciting – and far richer – account in which the mind emerges from a prolonged interaction with a structured world. [555]

* The editors would like to thank Dr. Sejnowski for generously allowing us to publish a highly abridged version of this seminal paper of 1977. The full text can be found under the same title in the journal *Behavioral and Brain Sciences*, 20, pages: 537-596. There are also numerous online versions easily accessible.
It is the hope of the editors that this excerpted version will provide the major conceptual foundation for a greater understanding of the other articles both scientific and philosophical found in the *Cognitive Architecture* volume.
Please note:

- Excerpts are indentified in hard brackets by page reference to the original version.
- Figures and Tables have been excluded.

---

[43] For example: M.R. Azimi-Sadjadi, S. Sheedvash and F.O. Trujillo, 'Recursive dynamic node creation in multilayer neural networks,' **IEEE Transactions on Neural Networks** 4 (1993): 242–56; S.E. Fahlman and C. Lebiere, 'The cascade-correlation architecture,' in **Advances in neural information processing systems**, ed. D. S. Touretzky. (San Mateo: Morgan Kaufmann, 1990); M. Frean, 'The upstart algorithm: A method for constructing and training feedforward neural networks,' **Neural Computation** 2 (1990): 198–209; Y. Hirose, K. Yamashita and S. Hijiya, 'Back-propagation algorithm which varies the number of hidden units,' **Neural Networks** 4 (1991): 61–66; V. Kadirkamanathan and M. Niranjan, 'A function estimation approach to sequential learning with neural networks,' **Neural Computation** 5 (1993):954–75; J.C. Platt, 'A resource-allocating network for function interpolation,' **Neural Computation** 3 (1991): 213–25; Y. Shin and J. Ghosh, 'Ridge polynomial networks,' **IEEE Transactions on Neural Networks** 6 (1995): 610–22; T.R. Shultz, D. Mareschal and W.C. Schmidt, 'Modeling cognitive development on balance scale phenomena,' **Machine Learning** 16 (1994): 57–86; and M. Wynne-Jones, 'Node splitting: A constructive algorithm for feedforward neural networks,' **Neural Computing and Applications** 1 (1993): 17–22.

[44] E. H. Lenneberg, **Biological foundations of language** (New York: Wiley, 1967).