

Evolution might select constructivism

James Hurford, Sam Joseph, Simon Kirby,
and Alastair Reid

Department of Linguistics, University of Edinburgh, Edinburgh EH8 9LL,
United Kingdom. jim@ling.ed.ac.uk www.ling.ed.ac.uk/~jim

of information processing limitations that change over the course of learning. Contrary to expectation, they suggest that limitations can actually work to the learner's advantage, essentially by helping the learner to decompose complex tasks into simpler ones. In their view, processing limitations can act as filters that "select" input appropriate to the learner's level of expertise. Moreover, if a domain is hierarchical, the learner can extract useful components of the more complex complete task from this selected input. Both Elman (1993) and Goldowsky and Newport (1993) have demonstrated that neural network models can be made more effective in artificial language tasks by *limiting* their capacities in realistic ways.

An illustration might help clarify the basis of this surprising finding. A prelinguistic child with adult memory and attention capacities (a stationary learner) would be faced with long strings of phonemes that could be parsed and related to each other and any potential referents in a staggering number of ways. This profusion of possible relations is discussed by Q&S as an intractably large hypothesis space (sect. 4.2), that remains large due to the "poverty of the stimulus" (sect. 4.2.1, para. 3). In contrast, an actual child with limited memory and attention can only consider the relations between those phonemes that fall within its limited span of memory and attention. These limitations focus the child on relations at the phrase level, a more approachable problem, with a much smaller hypothesis space. In addition, understanding the structure of simple phrases helps the learner understand more complex sentences.

This work provides a psychological mechanism by which hierarchical representations might be built. Processing limitations initially select for the child the most basic, tractable components of the task domain. Once these components are mastered, the child can apply them to understanding the increasingly large portions of the domain to which it is exposed as the limitations abate. Because of this process, the child might be expected to build complex representations using simpler ones that were mastered earlier. This prediction adds a psychological dimension to the neural details provided by Q&S. It suggests that the growing neural structures described by Q&S might first represent simple components and then progressively combine and interlink these building blocks to build more complex representations.

It is important to note that the processing limitations are manifestations of neural constraints just as representational limitations are. Thus, neural development establishes the timecourse of the structure in the input that reaches the brain and the structure that can be represented by the brain. In this way, the representational capacity of the brain could be tightly coupled to the complexity of the input that impinges upon it, with both increasing gradually over neural development. Q&S highlight the "dynamic interaction between a structured environment and the neural mechanisms that are responsive to that structure" (sect. 4.3, last para.). The addition of gradually abating processing limitations to their framework allows the neural mechanisms not only to respond to the structure in the environment but actually to shape and amplify the environmental structure to which it will react. Though the flexibility of filtering is limited, it probably contributes to the development of complex and powerful neural representations.

Quartz & Sejnowski's target article is part of an exciting trend in studying development embracing the complexity of a dynamically changing system that shapes and is shaped by its environment. It seems that the time is ripe for neuroscience and psychology to enrich each other's understanding of development.

Abstract: There is evidence for increase, followed by decline, in synaptic numbers during development. Dendrites do not function in isolation. A constructive neuronal process may underpin a selectionist cognitive process. The environment shapes both ontogeny and phylogeny. Phylogenetic natural selection and neural selection are compatible. Natural selection can yield both constructivist and selectionist solution to adaptive problems.

Ontogeny. In order to promote their hypothesis of neural constructivism and demote the process of neural selectionism Quartz & Sejnowski (Q&S) make several unjustified attacks on the role of the synapse in development. To begin with, they point out that the studies of pruning in human prefrontal cortex (sect. 2.1.1 para. 5) are inadequate and do not have data covering ages 1–5 and 8–15. It is unfair of Q&S, however, to demand human data in this instance when almost all of their other references to neuro-anatomical studies rely on nonhuman data. A more important point, which they avoid making, is that studies (including Rakic et al. 1994, whom they cite) of other brain regions show a definite increase, and then decline in synaptic numbers, with different regions having different time courses. Since cognitive processes are not restricted to the prefrontal cortex and develop at different rates, the aforementioned study contradicts the claim that synapse elimination does not underlie cognitive development.

Another problem arises with Q&S's argument that it is the dendritic arbor which provides the necessary measure of representational complexity. Dendrites do not function in isolation; the representational power of a neuron must be a composite of dendrite and synapse. Thus factors affecting synapse formation and location are just as important as those affecting dendritic growth. Following from this is a more general point that the addition of structure does not necessarily change the one-to-one connectivity that may have arisen from pruning. For example, in the neuromuscular junction, axons show an increase in branching and synapse number but this certainly does not represent an increase in representational complexity (J. Lichtman, personal communication). A further illustration is seen in the cerebellum where a *single* climbing fibre connects with a *single* Purkinje cell despite the complicated arborisations of both (Purves & Lichtman 1985). Our main point here is that increase in structural complexity (e.g., dendritic arborisation) does not necessarily give rise to an increase in representational complexity.

Finally there is the issue of levels of processing. Q&S wish to discuss the ways in which "cognitive and neural processes interact during development." It is reasonable to say that the acquisition of cognitive processes and their performance are inseparable components of cognitive development and that "learning guides brain development." This does not mean that learning and brain development proceed by the same process. It is perfectly feasible that a constructive process at the neuronal level (e.g., extension of a dendritic arbor) may underpin a selectionist process at the cognitive level (e.g., hypothesis elimination; Levine 1966). Thus, although neural constructivism and constructive learning are both valid concepts, neither one entails the other. The interaction between neural and cognitive processes in development is unlikely to be so simple as to be described by a single concept, be it constructivism or selectionism.

Phylogeny. The target article is negative about the prospects of understanding the human brain through understanding its evolution. Q&S's only positive statement is that there is a direction to evolution, that the "hallmark of cortical evolution is . . . an increasing representational flexibility." Q&S do not actually provide any evidence for this view; they simply argue that the only alternative is

to see the brain as a “hodgepodge of specialized circuits,” which is not in accord with the neurobiological evidence.

Although Q&S's view places prime importance on the structure of the environment in ontogeny, it demotes its importance in phylogeny. We believe this is an impoverished position. Instead, evolutionary psychologists (and evolutionary linguists) should be encouraged to take up Q&S's timely challenge to be responsive to neurobiological constraints, and explore the way in which the structure of the environment could influence the evolution of constructivist learning itself. Q&S do suggest that there are still some initial biases “that are consistent with a largely equipotential cortex,” but it is crucial to relate these biases with the environment in which humans evolved.

It is likely that the constructivist approach to learning in particular problem domains is a particular evolutionary response to those problems. We need to understand which features of the environment and the problems the environment poses lead to the evolution of a constructivist strategy. For example, the “long period of dendritic development in Broca's area” is an evolutionary response to the pressure to learn language. As Q&S themselves point out, this constructivist response is bought at the cost of “increased vulnerability during a protracted developmental period.”

This leads to a view of domain specific innateness in which development plays a central role. What is coded for genetically is a maturational program which leads to certain areas of the developing brain being particularly sensitive (in a constructivist sense) to input. We can now talk about a new kind of neural “specialisation”: a specialisation or dovetailing of developmental strategies to the environment in which organisms evolve. Extended cortical dendritic growth will dovetail with problem domains that require flexible representation construction. This will not be true for all problem domains, in which case it should be unsurprising that this extended development is not uniform over all cortical areas. Kirby and Hurford (1997) show that this type of evolutionary adaptation of developmental programs can be modelled computationally. For language learning they predict an early phase of “constructivist” development leading to a critical period for acquisition.

Topics for evolutionary research in the light of Q&S's paper could include an attempt to understand the relationship between constructivist learning and the Baldwin Effect (Baldwin 1896), in which flexible learning should be selected for when the environment is unpredictable but should be selected against when it is not, *whenever increased plasticity has a cost*. The effect predicts that learning can guide the evolution of innate solutions (see, e.g., Hinton & Nowlan 1987) by increasing the phenotypic plasticity of organisms and broadening the search space of a particular genome. Without further understanding, however, of the parameters along which constructivist development can vary (Q&S list: “generic initial cortical circuitry, conduction velocities, subcortical organization, learning rates, and hierarchical development”), we cannot model exactly how learning and evolution will interact.

Finally, Q&S state that “[neural] selectionism . . . is strictly incompatible with the claim that evolutionary pressures have picked out specialized circuits.” Although those advocating neural selectionism may emphasise its epigenetic component, it seems clear that it is not incompatible with natural selection. Moreover, natural selection could pick out the initial set of neural circuits (or even the mechanism for their generation) over phylogenetic time, which would then be subject to selectionist processes in the ontogenesis of the individual.

In summary, we are left with a view that the environment can impinge on the cortex on both the ontogenetic timescale (through several possible developmental processes at both a neural and cognitive level) and a phylogenetic timescale (through the selection of developmental processes which are appropriate to the environment). The important contribution of the target article is in putting the structure of the environment and the dynamics of development back on centre stage.

ACKNOWLEDGMENT

The authors would like to acknowledge the members of the Adaptive Systems Workshop at the University of Edinburgh, whose discussion of the target article clarified many of the issues raised. Corresponding author: JH.

“Differentiationism” can reconcile selectionism and constructivism

G. M. Innocenti

Institut de Biologie Cellulaire et de Morphologie, 1005 Lausanne, Switzerland. giorgio.innocenti@ibcm.unil.ch

Abstract: Increased complexity of representations in development probably results from the differentiation of distributed neural circuits. Axonal differentiation plays a crucial role in this process. Axonal differentiation appears to be achieved in stages, each involving combinations of constructive and regressive events controlled by cell intrinsic and cell extrinsic information.

Quartz & Sejnowski (Q&S) emphasize the notion that, beyond the empire of the genes, development and learning interact with each other. This is a lovely idea, developed with “mucho gusto” and coming at the right time. But that is not all. They address other fundamental issues in brain sciences, including the relation between neural structure and representations (sect. 2), the nature of the processes underlying the development of neural structure and the role of intrinsic versus extrinsic information in these processes (sect. 3). In this theory-laden journey through the neurosciences, they lost me in a few places. I shall speak to some of them.

Some measure of neural structure, possibly dendritic complexity, ought to correlate with representational complexity.

Why dendrites? And which dendrites? I was totally lost here. Most of the arguments advanced to support the role of dendrites could be extended to axons, when the complexity of their terminal arbors is considered, along with their possible computational properties (Innocenti 1995). Furthermore, there may be few, if any, crucial differences in the way dendritic and axonal arbors develop. Both differentiate through (often) exuberant growth and regression, show directed growth, respond to environmental conditions and can undergo local changes over extended life spans, including adulthood. Neither of them seems to call the shots in development; rather, they adapt to each other. Indeed, representations might not relate to any neuronal compartment in isolation, but rather to distributed neural circuits where axons, dendrites, and synapses combine interactively with each other, and with structural parameters, including the number of neurons and biochemical determinants of synaptic transmission. If increased complexity of representations means more numerous representations, each involving more elements, then it might be related to the number of *different* neural circuits available, to the number of *different* neurons belonging to each of these circuits, and perhaps to some measure of their connectiveness. As Edelman (1987) stressed the notion of “different” is important; I will come back to it. For now, there can be no doubt that:

Regressive events do play a crucial role in the development of neural circuits.

On this point, the target article is biased in its presentation and quotation of the literature. More complete accounts of regressive events in the formation of neural circuits can be found (e.g., Clarke 1994; Innocenti 1991; O'Leary 1992). In short, development involves massive reduction in the number of neurons, dendrites, dendritic spines, axonal branches, and synapses and, in addition, elimination of receptors, of neurotransmitters, and so on. There are some methodological difficulties here. One has to do with the quantification of neural structures. As rightly noticed by Q&S, the assessment of synaptic densities in a changing volume does not provide results unequivocally interpretable in terms of connectivity, and sometimes not even in terms of total numbers. In general, regression of synapses (or of other