BEHAVIORAL AND BRAIN SCIENCES (2000) **23,** 781–792
Printed in the United States of America

Continuing Commentary

Commentary on **Steven R. Quartz & Terrence J. Sejnowski (1997). The neural basis of cognitive development: A constructivist manifesto. BBS 20:537–596. {steve; terry}@salk.edu www.cnl.salk.edu/CNL**

Abstract of the original article: 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 the 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 the 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 the cortex has evolved so as to maximize the capacity of environmental structure to shape its structure and function through constructive learning.

From neural constructivism to cognitive constructivism: The steps to be taken

Andreas Demetriou
Department of Educational Sciences, University of Cyprus, 1678 Nicosia, Cyprus. **ademet@ucy.ac.cy**

Abstract: Quartz & Sejnowski's (Q&S's) model for constructive learning agrees with the basic assumptions of mainstream cognitive developmental theories. However, it does not detail the neural equivalents of (1) the process of cognitive change per se, (2) the construction and functioning of thought modules, and (3) the involvement of "mindreading" and "mind-steering" in constructive learning. Specifying these equivalents is necessary if cognitive developmental neuroscience and mainstream cognitive development are to be directly connected.

Quartz & Sejnowski (1997t) (Q&S) advance a number of stimulating assumptions about the neural basis of cognitive development. I will focus only on those assumptions which are important from the point of view of mainstream theory of cognitive development. My aim is to show the problems that Q&S would have to solve before their constructivist manifesto could be adopted by developmentalists primarily interested in the development of thought and reasoning rather than in neural phenomena as such.

Q&S assume that learning is a general-purpose adaptational mechanism that operates on a general-purpose or equipotential substrate (the cortex) to generate structural and representational specificity. Thus, the specificity of environmental structures is reflected in the specificity of cortical areas and neural networks and this is reflected in the specificity of representational structures.

This picture of mind is perfectly compatible with current models of cognitive development. These models recognize that there are general mechanisms of cognitive change that govern the construction of new concepts and mental operations. They also accept that the mind involves specialized thought structures or modules (Case & Okamoto 1996; Demetriou et al. 1993; Karmiloff-Smith 1992).

It is accepted that cognitive change requires—in a hierarchy-like sequence—comparison of representations, abstraction of common elements, reorganization of these elements into new representations, and some kind of symbolic individuation of the new representations so that they can be activated in the future as such. How does Q&S's constructive learning account for these steps? That is, what are the neural equivalents of comparing representations, abstracting their commonalities, assembling new representations, and symbolically individuating them? How are geometric principles of information processing design applied at the successive phases of representational change? These questions need answers if Q&S's general learning theory is to be adopted by cognitive developmentalists.

Explaining modularity is no less important. It was discussed in a number of the first round commentary. However, many crucial questions about modularity still remain unanswered. I agree with Q&S that "it is important to turn attention back to examining environmental structure" (sect. 4.3.1). The mind and the environment are functionally and structurally tuned to each other. As a result, there are thought systems that specialize in the representation and processing of different domains in the environment. Our research (e.g., Demetriou et al. 1993) has identified five thought domains: categorical, quantitative, spatial, causal, and propositional. These domains are computationally specific because they reflect the peculiarities of the elements and relations of the environment domains to which they are affiliated. Admittedly, there is no general agreement as to the identity of domains. Other scholars specify domains on the basis of ontological rather than relational criteria and they speak about thought modules that deal with the physical, the biological, and the psychological world (Karmiloff-Smith 1992). The two kinds of modules may be interrelated. For instance, there are categorical, causal, or mathematical relations in each of these three ontological domains (Demetriou, in press). However, we still do not know how the two types of domains are interwoven in the thinker's mind, because each type of relation functions idiosyncratically in each ontological do-

© 2000 Cambridge University Press 0140-525X/00 $12.50
https://doi.org/10.1017/S0140525X00233451 Published online by Cambridge University Press

main. For instance, physical causality requires the transmission of energy. Psychological causality does not require any energy or the intervention of any medium. Imagine mood variations caused by the memory of an unpleasant encounter.

Q&S's neural constructivism is general enough to be compatible with any of the two types of modularity sketched above. Naturally, we expect more from a cortex-centered theory of learning. Specifically, we expect it to be able to clearly shift the balance in favor of the one or the other of them or to show how the two can be integrated into an overarching model. To do so, the theory would have to provide answers to the following questions:

1. What is it that directs the construction of differentiable neural circuits that correspond to the cognitive modules mentioned above? This question is important because, according to the model, the learning mechanism is general and the brain is initially equipotential. Thus, we need to know how the differentiation is effected at the very beginning. Nativist conceptions do not face this problem because they posit that different circuits are in place from the start.

2. But even if this construction process is satisfactorily understood, we would still need to know the neural parameters of the various operations that are generated by the construction and which define the functioning of the different modules. For example, what does it mean in terms of synaptic communication, conduction velocities and the like to perform a mental rotation as contrasted to the addition of two numbers? Are these mental operations different simply because they are carried out by different neural circuits which function in the same way or are the functional aspects of the various circuits also different? Or, alternatively, one and the same circuit can carry out different mental operations by varying any of the parameters of neuronal functioning?

3. Last but not least, most cognitive developmentalists would accept some kind of nonstationarity and constructivism. However, Q&S's version of them ignores something that is probably very important in human cognitive development. This is mindreading (Baron-Cohen 1995) and, I would add, "mindsteering." People are aware of their own and others' minds and they systematically try to steer minds to preselected directions from a very early age (Karmiloff-Smith 1992). This implies that mindsteering must somehow be part of the neural construction process and that variations in how it is effected causes variations in the generation of environment-specific representations. But where is it located? Is it part of representationally specific circuits or does it possess its own circuits? What neural events does it require? To paraphrase Eccles (1994), how does the self control its brain? Answering these questions would provide Q&S's manifesto with its real constructivist dimension.

Infant perception and cognition and the initial architecture of constructivist models

Peter D. Eimas

Department of Cognitive and Linguistic Sciences, Brown University, Providence, RI 02912 **eimas@magnum.cog.brow.edu**

Abstract: There is a wealth of data on the perceptual and cognitive capacities of infants strongly supporting early nativistic influences on development. Without considering these initial determinants, constructivist models of development are at best incomplete.

Quartz & Sejnowski (1997t) (Q&S) have provided a coherent account of cognitive development based on the interaction between experience with a structured environment and neuronal ensembles. A potential consequence of this interaction is to enhance the complexity of existing neuronal structures, given the constructivist nature of their view of development. These increasingly complex structures are presumed to provide the means for understanding our increasing intellectual capacities with age and experience and to do so with minimal emphasis on nativistic determinants. Constructivist views of development, for example, Piagetian theory, have been the dominant view of cognitive development until the past 30 or 40 years. Of major importance in the shift from constructivism to nativism as the dominant theme of development were the many findings over the past several decades demonstrating a level of perceptual and cognitive sophistication in young infants and neonates that was well beyond the bounds of constructivism of the time (Mehler & Dupoux 1994).

The recent driving force away from nativism has been the development of powerful learning systems that have been shown to accommodate a number of developmental phenomena (e.g., Clark 1993; Elman et al. 1996; Mareschal & Shultz 1996) and assumed to be capable of explaining the very nature of development given a constructivist format. It is argued here, however, that insufficient emphasis has been given to the cognitive abilities of infants by most computational theorists, including Q&S. Indeed, no discussion of a sophisticated innately determined initial state was provided by Q&S and only nativistic principles of development of the most general nature were conceded, thereby casting doubt on the validity of their view of development.

Examples follow of the very young infant's highly sophisticated perceptual and cognitive achievements in their earlier commentary. These findings argue for considerable innate constraints, including those of a maturational nature that are evidenced across the first years of life in terms of processing mechanisms and even the representation of knowledge per se. This evidence is, as noted by Deheane-Lambertz & Deheane, consistent with a selectionist theory of development, although the existence of constructive processes during infancy and other periods of development is not precluded by a highly developed initial state or even by instances of selective processes. The evidence does, however, force a more detailed description of the initial state and the progressive maturation of innate constraints on domain specific processing and knowledge representation; and in so doing it constrains the very architecture of computational models of cognition.

I first describe evidence from the perception of speech by infants during the first year of life. As is well documented (Jusczyk 1997), the infant is able to perceive very fine differences in the speech signal and this is particularly the case when these differences signal linguistic distinctions, that is, phonetic distinctions that are present in numerous adult languages. The infants are actually assigning the sounds of speech to categories that have or will come to have linguistic significance. Moreover, the processing mechanisms, which have been shown to be quite complex, are not different in kind (i.e., sophistication) from those of adult listeners. And finally, the categorical assignments of speech by infants from a number of language communities are alike, despite considerable differences in the phonetic categories of the parental language. Moreover, this is true until they are about six months of age for vowels (Polka & Werker 1994) and 10-12 months of age for consonants (Werker & Tees 1984). Q&S do note that exposure to the parental language plays a role in the formation of a phonetic repertoire, citing the work of Kuhl et al. (1992) on the development of prototypic values or "perceptual magnets" for certain acoustic parameters of a vowel category, but only for vowels that belong to the parental language. It is worth noting in this regard, however, that Miller and Eimas (1996) have found similar prototypes for consonants in 3- and 4-month-old infants, that is, well before the age at which the native language begins to alter the consonantal repertoire of infants. Such evidence is difficult to reconcile with models based solely on learning from an instructive environment.

There is also at least inferential evidence for maturational influences in the perception of speech, which guide the infant to attend to what is presumably the same acoustic information (e.g., the contour of the fundamental frequency, the presence and duration of pauses, and possibly syllable duration and timing) in a number of instances, but for different purposes at different ages

https://doi.org/10.1017/S0140525X00233451 Published online by Cambridge University Press

(see Jusczyk 1997). Thus, for example, infants from a few days of age on are able to distinguish among the prosodic characteristics of different languages and even quite similar languages after a few months of experience with the parental language (Bosch & Sebastian-Galles 1997). It is not until the age of 7 months, however, that the information permitting the segmentation of words and clauses is used and not until 9 months of age do they differentiate the phonotactic characteristics of different languages as well as segment phrases. It is necessary to account for this differential use of the acoustic characteristics of speech. That the use of these sources of information is governed by experience (learning) is undoubtedly true, and computational learning models of its use are being developed to accommodate the findings (Aslin et al. 1996; Morgan et al. 1996), but as yet the developmental progression has not been shown to be solely a function of this experience. In sum, what information and how and when it is used would seem to be deeply constrained by innate mechanisms that operate in conjunction with environmental inputs to yield a human language.

There is likewise an accumulation of work in other domains that yields similar conclusions, namely, that cognition does not originate solely from general innate constraints that include powerful mechanisms of learning. Rather, highly specific processing and representational constraints are involved that evidence themselves at various points in the earliest years of life. For example, Spelke et al. (1992) and Baillargeon (1995) study infants between 4 and 12 months of age to find the principles underlying the formation of object concepts and how these concepts enter into mental reasoning is likewise better explained by innately determined processing mechanisms evidenced early in life, if not at birth, and that include information regarding the nature of the physical environment (laws of a naive physics). Similarly, Quinn and Eimas (1996, for a review) have found that infants three and four months of age are able to form categorical representations for various animals and artifacts at (approximately) the basic and global levels. Such representations are undoubtedly not innately given, initially or later, but what must be involved are innate mechanisms for category formation and a sensitivity to the appropriate environmental information that will support such representations.

Given these data, constraints on the sensitivities of the input mechanisms of constructivist models must be well rationalized, along with the time at which they gain salience for different aspects of perception and cognition and how these constraints vary across domains of knowledge.

Finally, constructivist models that function by adding hidden units until some criterion of error reduction is achieved can only function if the correct outputs are specified. This specification in very young infants would appear to be possible only on the basis of initial constraints and not by means of an external tutor. If this were not the case, the problem would become one of teaching outcomes and their consequences to a system that often does not as yet have knowledge of the events that lead to these outcomes. This is a seemingly impossible task, as is building a cognitive system in humans or other sentient beings without including a highly developed initial state and a recognition of the manner in which maturation, along with experience, guides the construction of more complex processing procedures and knowledge.

Neural constructivism or self-organization?

Peter C. M. Molenaar and Han L. J. van der Maas

Department of Psychology, University of Amsterdam, 1018 WB Amsterdam, The Netherlands. **op__molenaar@macmail.psy.uva.nl**

Abstract: Three arguments are given to show that neural constructivism lacks an essential ingredient to explain cognitive development. Based on results in the theory of adaptive signal analysis, adaptive biological pattern information and self-organization in nonlinear systems of information processing, it is concluded that neural constructivism should be further extended to accommodate the occurrence of phase transitions generating qualitative development in the sense of Piaget.

The basic assumption of neural constructivism as presented by Quartz & Sejnowski (1997t) (Q&S) is that development is associated with a progressive increase in the neural structures underlying cognitive representations, depending on interactions with a structured nonstationary environment. According to Q&S, neural constructivism provides a neurocomputational approach which is compatible with Piaget's theory of cognitive development.

We agree with Q&S that Piaget's theory of constructive epigenesis is a viable theory of development independent of nativistic and environmental alternatives. In an early paper (Molenaar 1986) we showed that the nonlinear neural dynamics underlying cognitive information processing can explain the occurrence of stagewise development in Piaget's theory. As a simple example, it was shown how increased connectivity in a neural coding field due to continued learning can give rise to a Hopf bifurcation, that is, a phase transition from localized coding to distributed oscillatory coding. This conjecture was corroborated in a recent simulation study using an ART (adaptive resonance theory) architecture (Raijmakers et al. 1996).

Despite our agreement with the basic tenets of Q&S's neural constructivism, we think it is insufficient to explain cognitive development, in particular, the occurrence of stage transitions as empirically observed in several studies (e.g., van der Maas & Hopkins 1998). More specifically, we take issue with three aspects of neural constructivism that all lead up to the same conclusion that one essential factor is missing from Q&S's theoretical model. The three aspects concerned are: the role of learning in structured nonstationary environments, the interaction between genetical and environmental factors affecting development, and the emergence of more powerful cognitive structures.

Nonstationarity. Q&S refer to nonstationary environments and to nonstationary neural structures underlying representations of such environments. Such kinds of nonstationarity are well known in the engineering sciences and have spurred the development of adaptive signal analysis and optimal control techniques (e.g., Sage & Melso 1971). These techniques have proven indispensable in space vehicles functioning in unknown environments. Applications in the social and biological sciences are less common, but nevertheless fruitful (e.g., Molenaar 1949). What is important in the present context is that there is a consistent mathematical-statistical theory for ascertaining the stability of adaptive techniques operating in unknown nonstationary environments. The basic step in this theory consists of the representation of an adaptive estimation technique (and, if present, optimal controller) as a nonlinear dynamical system for which the existence of stable equilibria (yielding the desired stability) is to be ascertained (cf. Ljung & Soderstrom 1987). The upshot of this is that adaptation to nonstationarity does not require qualitative changes in a learning algorithm. On the contrary, much effort in the engineering sciences is devoted to avoiding the occurrence of instabilities (phase transitions) in adaptive estimation and control. Hence the explanation of stage transitions observed in cognitive development requires more than an appeal to nonstationarity; it requires something like the occurrence of bifurcations in the learning algorithm itself.

Genes and environment. Quantitative genetical analyses of developmental processes not only show the (main and interactive) effects of genetic and environmental factors, but also the presence of an independent third source of interindividual differences. The third source gives rise to variations in developmental outcomes even if all genetic and environmental factors are kept constant (either experimentally or in simulation studies). Empirical evidence already occurs in the very first path diagram published by Sewall Wright (1920) under the denotation "developmental noise." We have compiled a number of published quantitative genetical studies showing the substantial impact of third source variation, argu-

https://doi.org/10.1017/S0140525X00233451 Published online by Cambridge University Press

ing that it is caused by the nonlinear dynamical processes underlying much of biological pattern formation (Molenaar et al. 1993). The latter processes are usually represented as nonlinear reaction-diffusion models and occupy a central role in mathematical biology (Meinhardt 1982). The key feature of reaction-diffusion models is their ability to self-organize through bifurcations induced by self-generated variation (originally proposed by Turing), thus providing causal models of adaptive morphogenesis, for example, the growth of neural fields. Consequently, the explanation of biological growth processes, in particular the development of neural fields, appears to require nonlinear dynamical models undergoing phase transitions.

More powerful structures. Paraphrasing Fodor (1980; 1981), a learning device equipped with propositional logic cannot learn a more powerful logic (e.g., including quantification) because it cannot even recognize quantified logic expressions. The issue here is to explain transitions to a more powerful system of information processing, which in Fodor's view can only be accomplished nativistically. Yet, taking into account the nonlinear neural dynamics underlying information processing shows that such transitions too can be explained in terms of phase transitions (Molenaar 1986; Raijmakers 1997). In fact, a complete mathematical-statistical theory of information processing and self-organization in nonlinear dynamical systems is available (Haken 1988), from which general information measures of representational and learning complexity can be derived. Again, it can be shown that the occurrence of bifurcations is necessary to explain transitions to more powerful structures.

In conclusion, we agree with Q&S that the nonlinear dynamical neural substrate underlying coding should be integrated in causal models of cognitive development. We do not agree, however, that the postulation of neural constructivism driven by ongoing interactions with a nonstationary structured environment is sufficient. Nonstationary input itself can be accommodated by qualitative invariant learning, as the engineering sciences teach us. Interaction with environmental inputs is not sufficient to explain the details of the growth of neural fields. Finally, the emergence of more powerful structures, as assessed by general information measures, shows that bifurcations are a necessary ingredient in any explanation of the qualitative change. Hence, we suggest that neural constructivism be further extended to accommodate the occurrence of phase transitions generating qualitative development in the sense of Piaget.

Waiting for Manifesto 2

David Premack and Ann James Premack
Laboratoire du Psycho-Biologie de Developpement, CNRS, Paris, France

Abstract: We suggest that innatism and constructivism may differ only in their time scale.

Bravo to Manifesto 1, "neural development" (Quartz & Sejnowski 1997t) (Q&S). We look forward to Manifesto 2, "cognitive development."

In Manifesto 1, Q&S discuss three different neural measures, all indicating continued growth of the brain over a period of up to 20 years; they see this as inimical to a selectionist view (also to an innatist view) of cognitive development, and as supporting their constructivist view. Unfortunately, however, though detailed in its treatment of neural development, their target article does not provide an equally detailed account of cognitive development. It relies on general remarks.

For example, they highlight the nonstationarity of the learning device they propose; learning transforms the learning device itself, so that what has been learned can then affect future learning. Because of problems inherent in a large network—that instantiates (innatist) guesses about the class of needed target concepts—they start with a small network that adds structure, until a desired error rate is achieved, mapping the structure of the problem space by successive approximations. The effect is to tailor the hypothesis space to the demands of the problem at hand. This general approach is possible, they claim, because the world wears its heart on its sleeve—its basic structure is open to discovery.

Finally, constructivist learning consists of two steps: (1) "a prolonged period of representation construction in which neural structures respond to the . . . structure of the environment," followed by (2) rapid learning made possible by the first phase.

In this commentary we apply their general remarks to some known cases of learning, asking: Do the specifics of these cases conform to the terms of their constructivist proposal? In addition, we describe two general problems that any constructivist program would appear to confront. These problems are caused by: (1) profound differences among domains, and (2) the difference between possible and actual: constructivism *may* work, evolution, that is, innatism patently *does* work.

Before turning to specific cases of learning, consider a key point on which much of our argument turns: A fundamental similarity between constructivism and innatism on which Q&S do not comment. Suppose we ask an innatist, who offers a model explaining one competence or another, "Don't just tell us your model, tell us how the competence evolved." Surely, it did not evolve in one step. The competence that the model purports to explain is based upon detecting (and interpreting) the invariances that define the problem space; doubtless it is the product of a number of small steps. What are they?

Of course, innatists expect consumers to be grateful for the model alone, and do not normally hold themselves responsible for biological origins; but suppose this time we obtain an answer. Notice the relation between the answer, the nativist's evolutionary account of his model, and the constructivist's account of how a child figures out the structure of a problem space. The two could be identical. A series of small steps. One accomplished in phylogenetic time, the other in ontogenetic time. In other words, constructivism and innatism could be identical except for time scale.

What is involved in representing the structure of a problem space? Clearly this, rather than "rapid learning," is the interesting part of constructivism. Indeed, the ability to represent this space is, the authors claim, the special quality of human intelligence. The human brain has not evolved a batch of domain-specific adaptations as innatism claims; rather, they say, it has evolved neural machinery with the capacity to figure out the structure of problem domains.

Do the various human domains that have evolved have enough in common to be dealt with by a unitary mechanism? That is, why do we not have separate neural devices for figuring out how to represent the structure of each domain, for example, a device for language, another for folk psychology, number, folk physics, and so on? If we did, constructivism would greatly resemble adaptationism or evolutionary psychology. A set of pre-adapted tissues each specialized for deciphering the structure of the problems they evolved to solve.

Moreover, when we consider the problems humans are designed to solve we are struck not by their similarities but by their differences. We shall mention only a few in order to make the point; but the topic is worth broader development.

In the case of language, structure involves phonemes on the one hand, and form classes, noun versus verb, on the other. In intuitive psychology, early invariances concern self-propelledness and goal-directedness, with intentionality (or agency) and value+/− as early interpretations on these invariances. Structure concerns physical relations such as containment, support, collision, and the like in intuitive physics. Number presents yet other specialties.

Intuitive physics and psychology at least have motion in common: both are to an uncanny degree analyses not of object quality but of motion. Language and number, however, obviously are not

https://doi.org/10.1017/S0140525X00233451 Published online by Cambridge University Press

analyses of motion; therefore they differ not only from one another but from the others. Can one and the same neural device divine the structure of all these problem spaces, their profound differences in content notwithstanding?

In the key phrase "structure of the problem domain" to what does structure refer? Let us treat category as a poor man's synonym for structure. In this reading, the child figures out specifically what categories are needed to represent a domain.

For example, in language the critical categories are, as we mentioned, the phonemes on the one hand, and the form classes on the other. Infants obviously do not figure out the first of these categories; they distinguish phonemes at birth. Constructivism may do better with the second set of categories; distributional learning may contribute not only to word segmentation but even, surprisingly, to the noun/verb distinction in the developing child (Mintz et al., submitted).

Turn from infants to vervet monkeys. The vervet's problem domain is the predator, the categories of which are: raptor, leopard, and snake, to which it produces three different calls. Does the immature vervet figure out the structure of this domain, that is, does it learn the categories? No, it has the categories; what it learns is how to fine tune the membership of the categories.

For example, a young vervet can mistakenly produce the raptor call to hawks (which resemble the true predator), produce the snake call to inappropriate snakes, and the leopard call to inappropriate ground animals. It corrects these errors, learning to confine the call to the correct member of each category, and to respond more quickly. However, even when the vervet produces its first calls, it does not make between-category errors, for example, issue the snake call to a bird, and so on. Hence, vervets do not "figure out the structure of the problem space." They come with the structure.

When young chimpanzees are taught humanlike words they show what appears to be a perfect example of constructivist cognitive development. We introduced three such animals to plastic words (that adhered to a magnetized writing slate), as names for different fruits. Rather than learn associations between the words and fruit, they divided the problem space into two categories, words and fruit, and learned the class properties of each. That is to say, they learned the class properties of "words," the pieces of plastic; they already knew the class properties of fruit, an innately perceived category.

They next learned this mistaken rule: any member of category "word" can be used to obtain any member of category fruit. For example, although responding at chance level in choosing the correct word for a fruit, they: (1) never placed any object on the writing board except a plastic word, (2) nor accepted any object in exchange for a plastic word other than fruit. Though they preferred chocolate to fruit, after placing a plastic word on the board, they rejected chocolate and took fruit (generally, the wrong fruit). Finally, they inhibited or unlearned the incorrect rule and formed correct specific associations between words and fruit. Henceforth, they learned new words in one trial.

Why, we ask, does chimpanzee word learning appear to conform so closely to constructivist learning? Is it because the case is so completely artificial, there being no innate component except for the category fruit, which is innately perceived?

Consider a more basic version of the same question: Why does distributional learning (possibly) contribute to the noun/verb distinction? Is it because the distinction is not a part of the structure of the world, a biological invariant, but a cultural distinction produced by human behavior?

Suppose constructivism and innatism are the same program operating on different time scales. Innatism would be well-suited for biological invariants, constructivism for cultural invariants. While the vaunted flexibility of constructivism would not benefit biological invariants—they are here to stay—it could benefit cultural distinctions.

This problem remains for constructivism. Was the cultural diversity of the Pleistocene sufficient for constructivism to have evolved?

And why does the fact that the brain continues to develop for up to 20 years automatically support constructivist cognitive development? In one sense it seems an embarrassment: Children acquire adult-level competence in both language and intuitive psychology by about the fourth year.

The above considerations suggest that a unified or monolithic account of cognitive development may not be what we need. There may be cases in which cognitive development definitely entails learning the categories—figuring out the structure. But in other cases, the individual brings the categories, and learning consists merely of fine-tuning the membership.

If there is a place in this world for both innatism and constructivism, we deny that it is entirely our fault. We realize what a sharp disappointment it will be to the adversarial parties if each has to accommodate the other; psychologists have a gourmand's appetite for premature parsimony. There will be no manifest catharsis for any one except for us moderates, and we're too few to count.

Authors' Response

Constraining constructivism: Cortical and sub-cortical constraints on learning in development

Steven Quartz[a] and Terrence Sejnowski[b]

[a]*Division of Humanities and Social Sciences, California Institute of Technology, Pasadena CA 91125 and* [b]*Howard Hughes Medical Institute, The Salk Institute, Department of Biology, University of California, San Diego, La Jolla CA.* **steve@hss.caltech.edu** **terry@salk.edu**

Abstract: It is becoming increasingly clear that acquiring cognitive skills is feasible only with significant developmental constraints. However, recent research provides the strongest evidence to date for constructivist development. Here, we examine how these two apparently conflicting perspectives may be reconciled. Specifically, we suggest that subcortical and cortical structures possess divergent developmental strategies, with many subcortical structures satisfying Fodor's criteria for modularity. These structures constitute an early behavioral system that guides the construction of later emerging cortical structures, for which there is little evidence for modularity. Thus, we focus on how the dynamic time course of development itself implicitly constrains the emergence of cortical representations, reducing the requirement for built-in encodings of knowledge in cortical circuits, as on the traditional nativist conception.

We would first like to thank the commentators for their insightful remarks regarding our target article and for providing us with an opportunity to clarify a number of issues raised by our target article (Quartz & Sejnowski 1997t). The commentary process has been extremely helpful in raising substantive issues in our neural constructivist position and in our presentation of it. In this reply, we address both matters as they pertain to the issues raised in the commentaries.

R1. Balancing neural and cognitive perspectives on development

As both **Demetriou** and **Premack & Premack** point out, our target article contained a more extensive treatment of

https://doi.org/10.1017/S0140525X00233451 Published online by Cambridge University Press

neural development than cognitive development, leaving out many fundamental concepts in cognitive development. Our emphasis on neural development followed from the methodology that we outlined in section 1.1 of our target article. That methodology prescribed first identifying the basic neural mechanisms of development and then relating these to processes of cognitive development. As Demetriou discusses, in using this methodology we stopped short of detailing two fundamental processes: (1) the process of cognitive change, and (2) the construction and functioning of thought modules. Demetriou rightly observes that in neglecting these two processes we stopped short of presenting a full account of cognitive development in neural terms. Similarly, Premack & Premack await a second, cognitive, manifesto, in which we more fully link neural and cognitive processes under a unified framework. Such a manifesto must await a number of basic advances before it can be attempted. That is, despite much progress in developmental cognitive neuroscience, a number of fundamental issues must first be addressed before successfully relating high-level developmental phenomena to their neural substrates. For example, work in connectionist modeling has attempted to relate some of the cognitive processes Demetriou discusses to neural processes, and in so doing has yielded intriguing clues regarding the dynamics of development. However, it remains unclear how many properties of such models relate to biological substrates and their developmental processes. **Eimas** also notes that supervised constructive neural network algorithms are developmentally unrealistic without a highly developed initial state, effectively requiring replacing an external teacher with internal constraints. These and other considerations suggest that relating high-level cognitive processes of development to its neural basis via connectionist modeling, while promising, encounters a number of basic obstacles.

How can the gap between the dynamics of connectionist modeling and the neural basis of cognitive development begin to be closed? One approach, which we adopted in our target article, is to constrain cognitive theory by what is known regarding the mechanisms of neural development. We thus began by asking what processes guide the construction of the neural structures underlying cognition. The reason why we consider this an important matter for cognitive theorists lay in the possibility that patterns of neural activity, reflecting informational structure of the distal environment, have a central role in constructing neural circuits. Traditionally, neural development was regarded as an internal affair, a process of intrinsic maturation of little relevance to cognitive theory, which dealt with learning, not maturation. However, the emerging view from developmental neurobiology regarding the central constructive role of activity suggested a basic shift in classic debates, such as the nativist-selectionist one. Indeed, it suggested that the traditional distinction between learning and biological maturation might not stand up to empirical scrutiny. Since the separation of developmental psychology and developmental neurobiology rests in part on this distinction, its breakdown would provide a principled motivation for attempting to integrate developmental psychology and developmental neurobiology.

An initial research strategy investigating the neural basis of cognitive development in terms of neural mechanisms may be a more productive methodology than attempting to identify the neural substrates of high-level cognitive phenomena. Progress in developmental cognitive neuroscience depends on characterizing the mechanisms of development in neural terms and building up from an account of these mechanisms to capture cognitive processes. We thus chose to remain as close as possible to the neural mechanisms of development and relate these to processes of cognitive development to the extent possible. Building on a well-specified account of what is sometimes referred to as the learning mechanisms underlying cognitive development, it should then be possible to formulate a constrained account of the higher cognitive processes discussed by **Demetriou** and **Premack & Premack,** although this remains largely a promissory note at this stage of developmental cognitive neuroscience.

Since the appearance of our target article, a number of studies have appeared that provide the strongest evidence to date for the role of activity-dependent neural constructivist processes in development. In particular, this work clarifies the issue of whether neural activity has a permissive or instructive developmental role, an issue that generated substantial debate in the original commentaries. Among the most pertinent new studies is that of Maletic-Savatic et al. (1999), who investigated the constructive role of activity in immature organotype culture preparations. These preparations culture a thin slice of the hippocampus for several days, preserving many of the neurons and their connections. Highly-localized stimulation elicited the sprouting of filopodia, thread-like protrusions which likely later convert to dendritic spines. This growth was not evident in the presence of NMDA-receptor blockers, suggesting that the properties of this receptor, a neural substrate of Hebbian learning, is required for eliciting this growth. These results are a direct empirical confirmation of the neural constructivist mode of development we posited in section 2.3.5 of our target article.

In another important study, Engert and Bonhoeffer (1999) report that dendritic spines, indicative of new connections, emerge in a highly restricted part of the dendritic tree of hippocampal cells after long-term potentiation in the same region. This is extremely pertinent, since the highly localized structural changes, a few synapses on a small stretch of an active dendrite, illustrates the highly instructive nature of neural activity. Evidence for neurogenesis in the adult brain also continues to expand, with the discovery that exercise enhances the survival of newly generated granule cells in the dentate gyrus of the hippocampus and increases the strength of long-term potentiation in the dentate by 100% (van Praag et al. 1999). Far from playing a permissive role, activity appears to be instructive as Hebb (1948) envisioned, a role that is inconsistent with previous selectionist models (Changeux & Danchin 1976; Changeux & Dehaene 1989). Indeed, along with the results we cited in our target article, these recent results provide sufficient evidence to rule out strict selectionist models of development.

A consensus is also emerging regarding the centrality of progressive constructivist processes in development. Although the visual pathway was once regarded as a model system for selectionist development, in a recent review of its development, Katz and Shatz (1996, p. 1134) conclude, "the strategy for forming adult circuits involves a local control of sprouting and synaptogenesis rather than selection from a large pre-existing repertoire." While this more constructivist account holds for projections to the visual cortex,

https://doi.org/10.1017/S0140525X00233451 Published online by Cambridge University Press

approximately 90% of synapses in the cortex originate from other cortical neurons. Among these, pyramidal cells in layers 2/3 have long, intrinsic horizontal axon collaterals within both layer 2/3 and layer 5, which form periodic "clusters" of finer axon branches. The development of this system also appears to involve a progressive increase in structural complexity. In a recent review, Katz (1997, p. 534) concludes, "this 'constructivist' view of the role of activity (Purves 1988) seems to more plausibly account for the development of the system of clustered connections than mechanisms based on large-scale regressive events."

The combination of extremely fine-grained studies at the level of individual dendritic segments and systems-level neuroanatomical studies represents an important advance in our understanding of the mechanisms and processes of neural development. With an emerging understanding of the cellular mechanisms of constructive development and the centrality of this mode of growth in the cortex, it should now be possible to begin to construct highly constrained computational models of development and explore their implications for the higher level cognitive processes of interest to developmental psychologists.

R2. Modules, domain specificity, and the initial state

Both **Premack & Premack** and **Eimas** presented a number of persuasive and balanced arguments regarding the need for constraining cognitive development in terms of a richly structured initial state. We are in basic agreement with these arguments and the moderate position both commentaries chart between the poles of radical empiricism and nativism. Indeed, in many ways, our intention in the target article was to sketch a response to a developmental paradox whose resolution bears on the need for balancing initial constraints and developmental plasticity. On the one hand, learning-theoretic and other formal results indicate that learning is impossible without highly constraining the learner. Indeed, although neural networks were originally lauded for their general-purpose learning capacities, subsequent analysis revealed that they too could only feasibly learn when highly constrained. In an important review, Geman et al. (1992) state, "learning complex tasks is essentially impossible without the a priori introduction of carefully designed biases into the machine's architecture." We believe these results are incontrovertible. What remains to be discovered, however, is what form these constraints take in terms of the brain's structural properties. The most straightforward and popular proposal regarding constrained learning is, as both Premack & Premack and Eimas point out, via modules, which encode domain-specific representations. Since knowledge in the brain is believed to be encoded by the pattern and weights of synaptic connections, the most straightforward and popular translation of such representational encodings in terms of the brain involves encoding information in the microstructure of neural circuits (see Pinker [1994] for statements of this strategy). This is, of course, only one of any number of concrete proposals regarding how domain-specific information could be innately encoded in the brain (for discussion, see Elman et al. 1996), but its popularity merits attention. Indeed, if the claim is about what is innately given as domain-specific representations, then it is difficult to know how such representational knowledge could be encoded if not via patterns of connectivity.

One of the virtues of the representation-microstructure proposal is that it results in testable predictions: (1) that regions of cortex, to the extent that they instantiate different domain-specific modules, should differ in terms of their microstructure, and (2) that a representationally constrained module should be limited in terms of the sorts of domains it can represent. The latter is simply another way of saying that since modules are required in order to constrain possible representations, such a constrained device is not general or multi-purpose. If a module is general-purpose, then it undermines the reasons for positing it in the first place.

The emerging evidence from developmental neurobiology points toward broad regional differences between different areas of cortex that control the differentiation of cells in concert with patterns of input activity. If there is some form of modularity in the cortex, it is unlikely to be the type found in subcortical structures, as discussed below. More important, the nature of the differences in intrinsic properties of neurons is probably graded such that different cortical areas are "tuned" to processing information at an appropriate rate and over the appropriate time window.

R3. Can developing cortex represent multiple domains?

Premack & Premack review a number of cognitive domains to highlight their striking differences. Whereas the basic structure of the language domain involves concepts such as phonemes, form classes, and noun versus verb, intuitive psychology involves concepts such as self-propelledness and agency. **Demetriou** also suggests that there are five irreducible thought domains: categorical, quantitative, spatial, causal, and prepositional. Premack & Premack ask whether the same neural device can uncover the structure of these disparate problem spaces. The suggestion is that these problem spaces require their own specialized representations, at least some of which must be innately encoded. This specialization amounts to (2) above.

The question raised by **Premack & Premack,** "whether the various human domains that have evolved have enough in common to be dealt with by a unitary mechanism" can only be settled by considering each domain on a case by case basis. However, recent neuroimaging studies are beginning to produce results that bear on this issue. For example, Sadato et al. (1996; 1998) have investigated the neural basis of Braille discrimination tasks in congenitally blind patients. They determined that the visual cortex of these blind subjects mediates the sensory processing demanded by the Braille discrimination tasks. The question of whether visual cortex is functionally involved in the task itself was determined by transcranial magnetic stimulation of occipital cortex in blind subjects (Cohen et al. 1997). A magnetic pulse delivered to visual cortex during the discrimination task produced processing errors, indicating the functional involvement of visual cortex in the task. At the neural level, it appears that the tactile processing pathways usually linked to the secondary somatosensory area are rerouted in blind subjects to the ventral occipital cortical regions originally reserved for visual shape discrimination.

Although these results do not in themselves address the issue of representational biases in various cortical regions,

https://doi.org/10.1017/S0140525X00233451 Published online by Cambridge University Press

the capacity of developing cortical regions to become functional within a novel sensory domain indicates that cortical regions may develop representations for multiple problem spaces, as presumably tactile discrimination and visual shape discrimination require their own representations. Indeed, it is interesting that this representational flexibility is seen in primary sensory areas. One might suspect that if a high degree of domain-specific representational knowledge were built in anywhere in the brain, it would be in primary sensory areas. For one, these structures have a long evolutionary history, with homologous structures across a wide variety of species. Also, the structure of the visual and tactile world is relatively stable, suggesting that a priori representations would have little probability of mismatch with the world. This bears on **Premack & Premack**'s interesting discussion of the relation of constructivist development to evolutionary strategies. As they note, at a different time scale the successive stages of constructivist development are analogous to evolutionary change. Therefore, what one could build constructively during development one could also build innately during evolution. The cross modal plasticity exemplified by the Braille discrimination tasks illustrates the major advantage of building representations during development rather than transferring them to innate encodings. The extreme robustness of constructivist developmental themes and its capacity to accommodate contingencies echo recent work in evolutionary developmental neurobiology that point to the robust flexibility of developmental programs as exploratory processes (Gerhart & Kirschner 1997). Therefore, there is relatively little support for the proposal that representational knowledge is encoded a priori in the microstructure of cortical circuits.

R4. An overlooked modularity: Subcortical organization

The lack of support for the a priori encoding of knowledge in the microstructure of cortical circuits does not mean a return to tabula rasa learning. In retrospect, we should have more explicitly stated what we believed to be the implications of this conclusion. A relatively unconstrained *cortex* in terms of its microstructure in no way implies that the *brain* is an unconstrained learning device. This fallacious implication relies on conflating the brain and the cortex. Our focus was almost exclusively on what is known about the way that the cortex is organized and whether that organization supports a modular knowledge thesis. However, cortical organization is different from other parts of the brain, leaving it a separate issue whether there is support for modularity in noncortical structures. Indeed, there is increasing evidence that part of the hypothalamus is highly modular in terms of function. For example, Balaban (1997) was able to transfer the chicken crow vocalization to the quail by transplanting a part of the hypothalamus. As Balaban (1997, p. 2006) notes, "a simple model in which crowing differences are due to evolutionary changes in a single higher brain area is not tenable." Instead, major subcomponents of the behavioral difference were independently transferred with interspecies transplantation of separate subcortical regions. In contrast, no one has yet succeeded in transferring any species behavior between animals by transplanting part of the forebrain. Therefore, modularity is alive and well, but living in the hypothalamus and other subcortical structures. We believe this is significant for cognitive theorists and developmental scientists. Although the cortex has received the most attention by cognitive neuroscientists, the subcortical generation of complex behavior is an often overlooked constraint on cortical representation. For example, Dixson (1998) reviews at length the contribution of subcortical structures in the generation of complex primate sexual behavior. From a developmental perspective, subcortical structures, which are developmentally early, may have a central role in providing early behavioral and cognitive competences (see Johnson 1999 for a specific proposal regarding subcortical contribution to face recognition). As Johnson suggests, such subcortical constraints likely play a central but overlooked role in directing, or bootstrapping, the emergence of cortical representations.

The highly localized functions attributable to specific regions of the hypothalamus and many other subcortical structures is in marked contrast to the way that functions in the cortex are widely distributed. For example, vision is distributed over half the cortex and polysensory and association areas are the hallmarks of cortical organization in mammals and especially primates. The contrast between cortical and subcortical organization suggests that biology can produce modularity when nature wants to and that the anatomical correlate is highly conserved dedicated circuits. Interestingly, these species specific differences are expressed by the same set of genes, which are common from frogs to humans. What appears to vary is the subset of cells that these genes are expressed in and how they are modulated by other signals. For example, the species specific affiliative behavior of an asocial species of vole can be transformed to that of the highly social prairie vole by changing the pattern of expression of vasopressin receptors in a transgenic vole to that found in prairie voles (Young et al. 1999). If the cortex were organized in this manner, one would suspect that experimental results would have established such themes, as it has for subcortical structures. However, what has been seen so far is broad areal specification (Rubenstein et al. 1999), suggesting that the themes of cortical organization are very different indeed.

The striking contrast between cortical and subcortical organization shifts the focus to a variety of other sources of constraints, in so doing shifting away from the traditional emphasis on static representational encoding to exploring how the constructive developmental process itself can be a highly constrained process in terms of its systemic properties and in how its dynamics constrain the emergence of representations. In our target article (sect. 4.3) we highlighted the importance of characterizing the biases that are consistent with an initially pluripotent cortex and listed the following features: generic initial cortical circuitry, conduction velocities, subcortical organization, learning rates, and hierarchical development. As research on the genetic control of cortical regionalization makes increasingly clear, genetic mechanisms play a key role in determining the gross organization of the cortex (Rubenstein et al. 1999); however, these findings appear fully consistent with a role for afferent signaling in determining mature cortical function. To this we would add that the contribution of intrinsic properties of cells in different parts of the cortex and differences in patterns of innervation needs to be evaluated. For example, circuits in the prefrontal cortex maintain activity

https://doi.org/10.1017/S0140525X00233451 Published online by Cambridge University Press

during delay periods even when the input is removed, forming a neural correlate of working memory. In contrast, in the sensory areas in the early stages of processing information are reset by new information as soon as it arrives. This is reflected in different intrinsic properties of the neurons and circuits that predispose them to support persistent attractor states in the prefrontal cortex, and rapid switching in primary sensory areas (Tsodyks & Sejnowski 1995). Another difference between different areas of cortex is the pattern of neuromodulation. The frontal cortex, but not sensory cortex, is innervated by dopamine, and this modulates the robustness of the delay period activity, as has been recently shown in a network model of frontal cortex (Durstewitz et al. 1999). It is worth emphasizing that these regional response properties are not indicative of encoded knowledge representations, but instead reflects more general processing constraints whose contribution to developmental constraints remains to be assessed.

R5. The role of dynamical systems theory in development

To understand how the constraints we have enumerated above operate, it will be necessary to characterize the dynamics of development and its time-dependent properties. That is, it will be necessary to understand how these constraints shape the developmental trajectory. For example, what are the implications of a limited initial architecture for the acquisition properties of a learning system? Whereas traditional accounts suggested that these limitations weakened the learning system, neural network modeling casts these limitations in a new, advantageous, light. An initially restricted network must pass through a phase of limited representational power during early exposure to some problem and then build successively more powerful representational structures. Thus, these early limitations may actually help the system first learn the lower-order structure of some problem domain and subsequently use what it has learned to bootstrap itself into more complex knowledge of that domain (reviewed in Plunkett et al. 1997).

Such considerations suggest that the developmental trajectory that is determined by the dynamics of the constraints enumerated above play an important but neglected role in constraining development. This brings us to the comments of **Molenaar & van der Maas** on the use of dynamical systems theory (DST) in developmental theory. As van Gelder (1998) and others have reviewed, there has been increasing interest in DST in cognitive science. DST has been advanced as a framework for understanding the dynamics of cognition, particularly as an alternative to syntactically governed representations of traditional cognitive science. However, DST is probably most illuminating for cognitive science in its study of self-organizing systems (see Mareschal & Thomas, in press). In particular, as Molenaar & van der Maas point out, DST provides a number of conceptual tools for understanding the hallmark of self-organizing systems, namely the nonstationarity that is at the core of developmental systems.

Nonstationarity simply refers to the fact that the properties of developing systems are highly time-dependent. Thus, understanding the process of developmental change, rather than simply its initial state or final outcome, is paramount in developmental science. It is here that the study of self-organizing systems provides a number of important insights. For example, self-organizing systems have helped to explicate the developmental role of spontaneous neural activity, which is known to play a role in constructing neural circuits (reviewed in Wong 1999). As Linsker demonstrated (reviewed in Linsker 1990), randomly generated activity, what appears as essentially noise, can create feature filters given the functional properties that neural circuits possess in combination with their geometrical properties (e.g., interaction functions so that nearby activity is excitatory and inhibitory with increasing distance). Ordered structure is thus an emergent property of the dynamics and geometrical organization of such systems.

The role of randomly generated activity in the construction of neural circuits provides a conserved developmental mechanism when such activity shifts to experience-dependent activity with the onset of postnatal experience. Whether the early emergence of structure under the constraints enumerated above is sufficient to construct the rich initial state both **Premack & Premack** and **Eimas** point out as requisite for cognitive developmental remains an open research question. However, we believe this is a much stronger source of constrained development than previously indicated and one worth further investigation.

Continuing research on self-organization has also contributed to a more sophisticated understanding of learning. **Eimas** is right to point out that a constructivist neural network that employs a back-propagation type algorithm is confronted with the problem of who plays the role of a tutor. Over the last few years, significant progress has been made in exploring unsupervised learning algorithms for neural network models (Hinton & Sejnowski 1999). Unlike the earlier supervised learning algorithms, which required a detailed teacher to provide feedback on performance, the goal of unsupervised learning is to extract an efficient internal representation of the statistical structure implicit in the stream of inputs. Babies are barraged by sensory inputs already in the womb and their environment after birth is filled with latent information about the environment. During development unsupervised learning could shape circuits in the early stages of sensory processing to more efficiently represent the environment; in the adult brain, similar forms of implicit learning could provide subconscious cues to help guide behavior. Recently, Markram et al. (1997) have shown that neurons in layer 5 of cerebral cortex have recurrent connections that support a temporally asymmetric form of Hebbian learning: If a presynaptic input occurs 40 milliseconds or less before a spike occurs in the postsynaptic cell, the strength of the synapse increases, but if the input occurs in a 40 milliseconds window after the spike, the strength decreases. This form of plasticity is Hebbian because it depends conjointly between activity in the pre- and postsynaptic cells; but because of the temporal asymmetry, it implements a form of temporal difference learning (Montague & Sejnowski 1994) that allows the cortex to learn temporal sequences and predict future sensory inputs. In particular it has been shown that directional selectivity of cortical cells can arise from experience with moving visual stimuli during development (Rao & Sejnowski, in press). Temporally asymmetric Hebbian plasticity has also been found in the hippocampus, which may support the acquisition of conscious, declarative

https://doi.org/10.1017/S0140525X00233451 Published online by Cambridge University Press

knowledge about the world, including the order of sequential events (Bi & Poo 1998).

R6. Piagetian stages and dynamical systems theory

Molenaar & van der Maas point out that an essential feature of constructivism is the phase transitions that Piaget believed underlie the stage-like properties of cognitive development. Molenaar & van der Maas also rightly note that we did not provide a treatment of this essential feature of constructive development, making our neural constructivism incomplete in this regard. In contrast, they suggested that the stage transitions of Piaget's constructivism, should be characterized in DST terms, such as Hopf bifurcations. It is worth mentioning that although Piaget's stage theory has been extensively criticized there is much evidence supporting the dimension-like expansion of cognitive capacities. For example, Luciana and Nelson (1998) recently examined the developmental emergence of functions involved in prefrontally-guided working memory systems in four- to eight-year-old children. The development of prefrontally-guided working memory systems, which is thought to involve particularly the dorsolateral region of PFC, in four- to eight-year-old children appears to proceed dimensionally, beginning with the refinement of basic perceptual and sensorimotor functions and culminating in the emergence of distributed networks that integrate complex processing demands.

Cognitive discontinuity could be mediated by a number of neural mechanisms. For example, Shultz et al. (1994) and Mareschal and Shultz (1996) used the constructivist neural network cascade-correlation to learn the Piagetian balance scale task. In this task, a child is asked to predict the outcome of placing various numbers of equal weights at various distances to the left or right of a fulcrum. Between the ages of four and twelve, children appear to pass through four well-defined, discrete stages determined by sequence or rules as they master this task (Siegler 1981). The Cascade-correlation networks provided better fits to human data than did previous rule-based and connectionist models (Shultz et al. 1994). In particular, the model displayed quantitative changes as incremental improvement, which was likely mediated by weight changes. However, the model also displayed instances of qualitative restructuring analogous to the discrete stages children pass through as they master this task. In the model, the emergence of discontinuous, stage-like behavior was mediated by the addition of new units, which expanded the representational capacities of the network. This model therefore suggests how quantitative changes to the learning mechanism as an activity-dependent increase in representational power may underlie cognitive stages manifested as the behavioral level. Of course, it is important to bear in mind that this is only one possible mechanism and that other strategies, such as the recruitment of new regions, may also mediate such stage-like shifts.

As **Molenaar & van der Maas** indicate (see also Molenaar 1986), the dynamics of self-organizing systems suggest a response to Fodor's (1980) anti-Piagetian argument regarding the impossibility of acquiring more powerful conceptual structures. They rightly point out that the phase transitions of some dynamical systems may provide the conceptual tools and informational measures of representational complexity to refute Fodor's challenge. Quartz (1993) proposed a specific form of this argument, in suggesting that activity-dependent synaptogenesis corresponds to increases in an information-theoretic measure known as the V-C dimension. Using this measure, it can be shown quite straightforwardly that a system can acquire more powerful conceptual structures as activity-dependent synaptogenesis constructs more powerful representations.

R7. The terms of the debate

For most of its history, the study of cognitive development has been burdened by a debate that has caused more polarization than productive research. We are of course referring to the nativist-empiricist debate that continues to this day. Faced with the apparently conflicting results from learning-theory and developmental neurobiology, it would be productive for developmental theorists to go beyond the terms of the traditional nativist-empiricist debate to investigate how these apparently conflicting perspectives can be reconciled. To that end, empiricist-minded theorists would do well to acknowledge the need for domain-specific constraints. As we stated in our target article, tabula rasa learning is infeasible. However, the inference from the need for developmental constraints to symbolic encoding of domain-specific information in cortical structures, culminating in the modular mind (e.g., Hirschfeld & Gelman 1994), appears equally flawed. Constraints on learning may take surprising form and may be developmentally emergent properties of the brain as a representational and dynamic system. In this spirit, Eimas's comment that "what must be involved are innate mechanisms for category formation and a sensitivity to the appropriate environmental information that will support such representations" is on the mark. As **Eimas** also notes, a constrained initial state should be regarded as the starting point of development rather than its endpoint. The developmental perspective that flowed from Chomsky's linguistic speculations hindered an appreciation for both the richness and protractedness of development and its far-reaching interactions between the world and developmental mechanisms.

The rich interaction between the world and neural mechanisms also extends to gene expression. The polar debate between a nativist brain and tabula rasa learning hinges in part on a now outdated understanding of gene regulation. It now appears that gene regulation in the brain is an extraordinary complex and subtle process in which neural activity, in part regulated by experience, affects which genes are being expressed in a given cell at a particular time. Indeed, activity from the world regulates genes whose products are transmitter receptors; turns on genes that underlie the branching and growth of brain cells; and, regulates early intermediate genes following long-term potentiation (reviewed in Quartz 1999). As an appreciation of the centrality of activity-dependent gene expression becomes clearer, the brain and world will likely be seen as a interacting complex system, one in which the dichotomies of traditional nativist-empiricist debates no longer adequately frame the problem.

Premack & Premack suggest that protracted development is in a sense an embarrassment. Adult level competence in such domains as language is essentially reached by

https://doi.org/10.1017/S0140525X00233451 Published online by Cambridge University Press

four years of age, so why should the brain continue to grow well into adolescence? We believe part of the answer stems from a point raised by **Demetriou** regarding theory of mind. As he notes, the capacity to understand others in terms of their internal states is essential to human cognition (Premack & Woodruff 1978). Demetriou asks how such a capacity emerges and what are its neural substrates. Stone et al. (1998) discuss the distributed neural circuits that may mediate theory of mind, particularly the orbitofrontal cortex. As they note, theory of mind emerges gradually, with more subtle forms of theory of mind reasoning not emerging until adolescence. Russell et al. (1998) report developmental delays in theory of mind in deaf children, presumably because their early opportunities for learning about mental states are relatively restricted and that the normal development of theory of mind is dependent upon such opportunities. We find it striking that the prefrontal cortex appears to be the most protracted developmentally (reviewed in Quartz 1999) and that its mature capacities depend critically on interacting in a social environment. The protractedness of the constructivist strategy appears particularly sensible in the social domain, as it allows for increasingly complex strategies to develop in response to the contingencies of one's social environment.

Perhaps the most important emerging lesson in developmental science, then, concerns the rich and protracted interaction between a structured environment and activity-dependent, constructivist neural mechanisms. A richly structured initial state in no way diminishes the richness of this interaction or its central importance in understanding the origin of knowledge. Indeed, since much of the initial structure appears to take the form of constraints on the dynamics of this interaction, it highlights the importance of understanding the origin of knowledge and its neural representation as a dynamic constructive interaction with the world. This is in marked contrast to the traditional emphasis on static innate knowledge representations and its marginalization of the process of developmental change. As the commentaries so strongly illustrate, there is an increasing dialogue between developmental psychologists and developmental neurobiologists, suggesting that the developmental sciences are transforming into an integrated approach that will reveal a far richer account of how cognitive life is constructed from the developing brain.

References

Aslin, R.N., Woodward, J. Z. LaMendola, N. P. & Bever, T.G. (1996) Models of word segmentation in fluent maternal speech to infants. In: *Signal to syntax: Bootstrapping from speech to grammar in early acquisition,* ed. J. L. Morgan & K. Demuth. Erlbaum. [PDE]

Baillargeon, R. (1995) A model of physical reasoning in infancy. In: *Advances in infancy research.* Vol. 9, ed. C. Rovee-Collier & L. P. Lipsitt. Ablex. [PDE]

Balaban, E. (1997) Changes in multiple brain regions underlie species differences in a complex, congenital behavior. *Proceedings of the National Academy of Sciences USA* 94:2001–2006. [rSQ]

Baron-Cohen, S. (1995) *Mindblindness: An essay on autism and theory of mind.* MIT Press. [AD]

Bi, G. Q. & Poo, M. M. (1998) Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience* 18:10464–72. [rSQ]

Bosch, L. & Sebastian-Galles, N. (1997) Native-language recognition abilities in 4-month-old infants from monolingual and bilingual environments. *Cognition* 65:33–69. [PDE]

Case, R. & Okamoto, Y. (1996) The role of central conceptual structures in the development of children's thought. *Monographs of the Society for Research in Child Development:* 6:1–2. [AD]

Changeux, J. P. & Danchin, A. (1976) Selective stabilisation of developing synapses as a mechanism for the specification of neuronal networks. *Nature* 264:705–12. [rSQ]

Changeux, J. P. & Dehaene, S. (1989) Neuronal models of cognitive function. *Cognition* 33:63–109. [rSQ]

Clark, A. (1993) *Associative engines.* MIT Press. [PDE}

Cohen, L. G., Celnik, P., Pascual-Leone, A., Corwell, B., Falz, L., Dambrosia, J., Honda, M., Sadato, N., Gerloff, C., Catalá, M. D. & Hallett, M. (1997) Functional relevance of cross-modal plasticity in blind humans. *Nature* 389:108–83. [rSQ]

Demetriou, A. (in press) Nooplasis: 10 + 1 postulates about the formation of mind. *Learning and Instruction: The Journal of the European Association for Research on Learning and Instruction.* [AD]

Demetriou, A., Efklides, A. & Platsidou, M. (1993) The architecture and dynamics of developing mind: Experiential structuralism as a frame for unifying cognitive developmental theories. *Monographs of the Society for Research in Child Development:* 58:5–6. [AD]

Dixson, A. F. (1998) *Primate sexuality: Comparative studies of the prosimians, monkeys, apes, and human beings.* Oxford University Press. [rSQ]

Durstewitz, D., Seamans, J. & Sejnowski, T. J. (1999) Dopaminergic modulation of activity states in the prefrontal cortex. 25:1216. [rSQ]

Eccles, J. C. (1994) *How the self controls its brain.* Springer-Verlag. [AD]

Elman, J. L. Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D. & Plunkett, K. (1996) *Rethinking innateness.* MIT Press. [PDE, rSQ]

Engert, F. & Bonhoeffer, T. (1999) Dendritic spine changes associated with hippocampal long-term synaptic plasticity. *Nature* 399:66–70. [rSQ]

Fodor, J. A. (1980) Fixation of belief and concept acquisition. In: *Language and learning: The debate between Jean Piaget and Noam Chomsky,* ed. M. Piatelli-Palmerini. Routledge & Kegan Paul. [PCMM, rSQ]

(1981) *Representations: Philosophical essays on the foundations of cognitive science.* Harvester. [PCMM]

Geman, S., Bienenstock, E. & Doursat, R. (1992) Neural networks and the bias/variance dilemma. *Neural Computation* 4:1–58. [rSQ]

Gerhart, J. & Kirschner, M. (1997) *Cells, embryos, and evolution: Toward a cellular and developmental understanding of phenotypic variation and evolutionary adaptability.* Blackwell Science. [rSQ]

Haken, H. (1988) *Information and self-organization: A macroscopic approach to complex systems.* Springer. [PCMM]

Hebb, D. O. (1949) *The organization of behavior: A neuropsychological theory.* John Wiley and Sons. [rSQ]

Hinton, G. E. & Sejnowski, T. J., ed. (1999) *Unsupervised learning: Foundations of neural computation.* MIT Press. [rSQ]

Hirschfeld, L. A. & Gelman, S. A. (1994) *Mapping the mind: Domain specificity in cognition and culture.* Cambridge University Press. [rSQ]

Johnson, M. H. (1999) Ontogenetic constraints on neural and behavioral plasticity: Evidence from imprinting and face processing. *Canadian Journal of Experimental Psychology* 53:77–90. [rSQ]

Jusczyk, P. W. (1997) *The discovery of spoken language.* MIT Press. [PDE]

Karmiloff-Smith, A. (1992) *Beyond modularity: A developmental perspective on cognitive science.* MIT Press. [AD]

Katz, L. C. (1997) Linking layers and connecting columns. In: *Molecular and cellular approaches to neural development,* ed. Cowan, W. M., Jessell, T. M. & Zipursky, S. L. Oxford University Press. [rSQ]

Katz, L. C. & Shatz, C. J. (1996) Synaptic activity and the construction of cortical circuits. *Science* 274:1133–38. [rSQ]

Kuhl, P. W., Williams, K. A., Lacerda, F., Stevens, K. N. & Lindblom, B. (1992) Linguistic experience alters phonetic perception in infants by six months of age. *Science* 225:606–608. [PDE]

Linsker, R. (1990) Perceptual neural organization: Some approaches based on network models and information theory. *Annual Review of Neuroscience* 13:257–81. [rSQ]

Ljung, L. & Soderstrom, T. (1987) *Theory and practice of recursive estimation.* MIT Press. [PCMM]

Luciana, M. & Nelson, C. A. (1998) The functional emergence of prefrontally-guided working memory systems in four- to eight-year old children. *Neuropsychologia* 36:273–93. [rSQ]

Maletic-Savatic, M., Malinow, R. & Svoboda, K. (1999) Rapid dendritic morphogenesis in CA1 hippocampal dendrites induced by synaptic activity. *Science* 283:1923–27. [rSQ]

Mareschal, D. & Thomas, M. S. (in press) Self-organization in normal and abnormal cognitive development, In: *Brain and behaviour in human development. A source book,* ed. A. F. Kalverboer, & K. Gramsbergen. Kluwer Academic Publishers. [rSQ]

Mareschal, E. & Shultz, T. R. (1996) Generative connectionist networks and constructivist cognitive development. *Cognitive Development* 11:571–603. [rSQ]

https://doi.org/10.1017/S0140525X00233451 Published online by Cambridge University Press

Mareschal, M. H. & Shultz, T. R. (1996) Generative connectionist networks and constructivist cognitive development. *Cognitive Development* 11:571–603. [PDE]

Markram, H., Lubke, J., Frotscher, M. & Sakmann, B. (1997) Regulation of synaptic efficacy by coincidence of postsynaptic action potentials and excitatory postsynaptic potentials. *Science* 275:213–15. [rSQ]

Mehler, J. & Dupoux, E. (1994) *What infants know.* MIT Press. [PDE]

Meinhardt, H. (1982) *Models of biological pattern formations.* Academic Press. [PCMM]

Miller, J. L. & Eimas, P. D. (1996) Internal structures of voicing categories in early infancy. *Perception and psychophysics* 58:1157–67. [PDE]

Molenaar, P. C. M. (1986) On the possibility of acquiring more powerful structures: A neglected alternative. *Human Development* 29:245–51. [PCMM, rSQ]

(1994) Dynamic latent variable models in developmental psychology. In: *Analysis of latent variables in developmental research,* ed. A. von Eye & C. C. Clogg. Sage. [PCMM]

Molenaar, P. C. M., Boomsma, D. I. & Dolan, C. V. (1993). A third source of developmental differences. *Behavior Genetics* 23:519–24. [PCMM]

Montague, P. R. & Sejnowski, T. J. (1994) The predictive brain: Temporal coincidence and temporal order in synaptic learning mechanisms. *Learning Memory* 1:1–33. [rSQ]

Morgan, J. L., Shi, R. & Allopenna, P. (1996) Perceptual bases of rudimentary grammatical categories. In: *Signal to syntax: Bootstrapping from speech to grammar in early acquisition,* ed. J. L. Morgan & K. Memuth. Erlbaum. [PED]

Pinker, S. (1994) *The language instinct.* W. Morrow. [rSQ]

Plunkett, K., Karmiloff-Smith, A., Bates, E. & Elman, J. L. (1997) Connectionism and development psychology. *Journal of Child Psychology and Psychiatry and Allied Disciplines* 38:53–80. [rSQ]

Polka, L. & Werker, J. F. (1994) Developmental changes in perception of nonnative vowel contrasts. *Journal of Experimental Psychology: Human Perception and Performance* 20:421–35. [PDE]

Premack, D. & Woodruff, G. (1978) Does the chimpanzee have a theory of mind? Behavioral and Brain Sciences 1. [rSQ]

Purves, D. (1988) *Body and brain: A trophic theory of neural connections.* Harvard University Press. [rSQ]

Quartz, S. R. (1993) Neural networks, nativism, and the plausibility of constructivism. *Cognition* 48:223–42. [rSQ]

(1999) The Constructivist brain. *Trends in Cognitive Sciences* 3:48–57. [rSQ]

Quartz, S. R. & Sejnowski, T. J. (1997) The neural basis of cognitive development: A constructivist manifesto. *Behavioral and Brain Sciences* 20:537–96. [AD]

Quinn, P. C. & Eimas, P. D. (1996) Perceptual organization and categorization in young infants. In: *Advances in infancy research,* Vol. 10, ed. C. Rovee-Collier & L. P. Lipsitt. Ablex. [PDE]

Raijmakers, M. E. J. (1997) Is the learning paradox resolved? *Behavioral and Brain Sciences* 20:573–74. [PCMM]

Raijmakers, M. E. J., van der Maas, H. L. J. & Molenaar, P. C. M. (1996) Numerical bifurcation analysis of distance-dependent on-center off-surround shunting neural networks with external input. *Biological Cybernetics* 75:495–507. [PCMM]

Rao, R. P. N. & Sejnowski, T. J. (in press) Predictive learning of temporal sequences in recurrent neocortical circuits. *Advances in Neural Information Processing Systems* 12. [rSQ]

Rubenstein, J. L. R., Andrson, S., Shi, L. M., Miyashita-Lin, E., Bulfone, A. & Hevner, R. (1999) Genetic control of cortical regionalization and connectivity. *Cerebral Cortex* 9:524–32. [rSQ]

Russell, P. A., Hosie, J. A., Gray, C. D., Scottm C., Hunter, N., Banks, J. S. & Macaulay, M. C. (1998) The development of theory of mind in deaf children. *Journal of Child Psychology and Psychiatry and Allied Disciplines* 39:903–10. [rSQ]

Sadato, N., Pascual-Leone, A., Grafman, J., Deiber, M. P., Ibañez, V. & Hallett, M. (1998) Neural networks for Braille reading by the blind. *Brain* 121:1213–29. [rSQ]

Sadato, N., Pascual-Leone, A., Grafman, J., Ibañez, V., Deiber, M. P., Dold, G. & Hallett, M. (1996) Activation of the primary visual cortex by Braille reading in blind subjects. *Nature* 380:526–28. [rSQ]

Sage, A. P. & Melso, J. L. (1971) *Estimation theory with applications to communications and control.* McGraw-Hill. [PCMM]

Shultz, T. R., Mareschal, D. & Schmidt, W. C. (1994) Modeling cognitive development on balance scale phenomena. *Machine Learning* 16:57–86. [rSQ]

Siegler, R. S. (1981) Developmental sequences within and between concepts. *Monographs of the Society for Research on Child Development* 46. [rSQ]

Spelke, E. S., Breinlinger, K., Macomber, J. & Jacobson, K. (1992) Origins of knowledge. *Psychological Review* 99:606–23. [PDE]

Stone, V. E., Baron-Cohen, S. & Knight, R. T. (1998) Frontal lobe contributions to theory of mind. *Journal of Cognitive Neuroscience* 10:640–56. [rSQ]

Tsodyks, M. V. & Sejnowski, T. J. (1995) Rapid state switching in balanced cortical network models. *Network* 6:111–24. [rSQ]

Van der Maas, H. L. J. & Hopkins, B. (1998) Developmental transitions: So what's new? *British Journal of Developmental Psychology* 16:1–13. [PCMM]

Van Gelder, T. (1998) The dynamical hypothesis in cognitive sciences. *Behavioral and Brain Sciences* 21:625–28. [rSQ]

Van Praag, H., Christie, B. R., Sejnowski, J. J. & Gage, F. H. (1999) Running enhances neurogenesis, learning and long-term potentiation in mice. *Proceedings of the National Asociation of Sciences USA.* [rSQ]

Werker, J. & Tees, R. C. (1984) Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development* 7:49–63. [PDE]

Wong, R. O. (1999) Retinal waves and visual system development. *Annual Review of Neuroscience* 22:29–47. [rSQ]

Wright, S. (1920) The relative importance of hereditary and environment in determining the piebald pattern of guinea pigs. *Proceedings of the National Academy of Science USA* 6:320–332. [PCMM]

Young, L. J., Nilsen, R., Waymire, K. G., MacGregor, G. R., Insel, T. R. (1999) Increased affiliative response to vasopressin in mice expressing the V1a receptor from a monogamous vole. *Nature* 400:766–68. [rSQ]

https://doi.org/10.1017/S0140525X00233451 Published online by Cambridge University Press