

multilevel isomorphism is a radical rejection of Marr (1982). Cognition cannot be studied independently of the brain and body. Another important implication is the central role of developmental trajectories in the interpretation of adult cognition. There is no teleology involved in development; mature, normative cognition is an outcome of development, not a pre-specified target (Thomas & Karmiloff-Smith 2003).

We have also argued that brain regions do not utilize complete representations. In support of this view, we discussed, first, how a particular input signal is typically processed separately along distinct dimensions (e.g., shape, color, motion in object perception), and second, that functional brain systems are interrelated with other functional systems, and are located and within a body and an environment. All these contexts provide varying levels and sources of information such that only fragments of information need to be represented. An implication is that cognition will be compositional and systematic (e.g., Fodor 1975; Fodor & Pylyshyn 1988) only if the context permits it.

Progressive specialization, and notions such as emergent modularity (e.g., Karmiloff-Smith 1998a) imply that, as a result of development, the cognitive architecture will exhibit a progressive lack of flexibility to the novel. This can be observed when emergent specialized systems have well-delineated functions (Johnson & Munakata 2005). These specialized systems should not be construed as mere imprinting from environmental pressures and regularities. The child, from birth (e.g., Robertson et al. 2004), is an active contributor to his or her development. Our emphasis on context-dependence may be taken as implying that it is impossible to make general claims about cognition. This is not the case. Instead, we argue that the key to understanding contextualised function is to identify those contexts that are central to the function of interest, while ignoring those contexts that may have a peripheral rather than central role in determining the function of interest.

A recurrent theme in our book is the need for causal theories regarding what makes complex behaviors emerge. These theories need to explain behaviors on multiple time scales. They must explain how and why behaviors unfold as we observe them in real time, as well as how they unfold in developmental time. To do this, we need more than just a very detailed description of the behaviors that can be observed at any point. Certainly, such descriptions are essential for the advancement of causal theories of development, but they are unsatisfactory to the extent that they are unable to explain or predict new behaviors. The clearest example of this is with regards to explaining the behavior of children with developmental disorders. Unless one has a mechanistic theory of what is causing behaviors to unfold, and a causal theory of what is atypical in such children's processing, it is impossible to explain or predict why one set of atypically developing children may show a delay at one behavior and excel at another, while a second set of atypically developing children will excel at the former behavior but have a developmental delay in the latter behavior.

In the companion volume, *Neuroconstructivism, volume 2: Perspectives and Prospects* (Mareschal et al. 2007b), we invited nine research labs with objectives broadly consistent with the neuroconstructivist approach

to present their computational modeling work. The questions the models are built to investigate differ both in the level of description and in time scale over which the relevant behaviors operate. Some models focus on relatively rapid adaptation occurring (perhaps) at the cellular level of description, whereas other models focus on relatively slow adaptation occurring at the cognitive level.

We asked all contributors to the companion volume to emphasize the following aspects of their contributions when describing their work:

What functional brain constraints operate on the process of representation development?

What embodiment or situatedness constraints operate on the process of representation development?

Ultimately, computational models are tools to help us reflect on questions of process and mechanisms. Therefore, we also asked the contributing authors to answer the following questions:

How does the model embody these constraints?

What concrete predictions does the model make?

In *Neuroconstructivism, volume 1: How the brain constructs cognition*, we set out to investigate how the representations that underlie cognition emerge in the brain during development. We argued that the emergence of such representations is the outcome of a constructivist process involving constraints that operate at all levels from the cellular environment to the social environment. To truly understand how these representations emerge, it is necessary to locate our theories at the point where the constraints of brain, body, and environment come together.

ACKNOWLEDGMENTS

The writing of this paper was supported by EC grants 516542 (NEST) and 029088 (NEST), EPSRC Grant EP/D062225/1, ESRC grant RES-061-23-0129, MRC grants G9715587 and G0300188, and British Academy grant SG-40400.

NOTE

1. We acknowledge that the term *neuroconstructivism* has been used by others previously, sometimes with a general meaning very similar to our usage (Karmiloff-Smith 1998a; Quartz & Sejnowski 1997) and sometimes with a differing meaning (Sheridan 1997).

Open Peer Commentary

Are interactive specialization and massive redeployment compatible?

doi:10.1017/S0140525X08004081

Michael L. Anderson

Department of Psychology, Franklin & Marshall College, Lancaster, PA 17604.
michael.anderson@fandm.edu www.agcognition.org

Abstract: I offer a simple method for further investigating the Interactive Specialization framework, and some data that may or may not be compatible with the approach, depending on the precise meaning of "specialization." Findings from my lab indicate that, while networks of

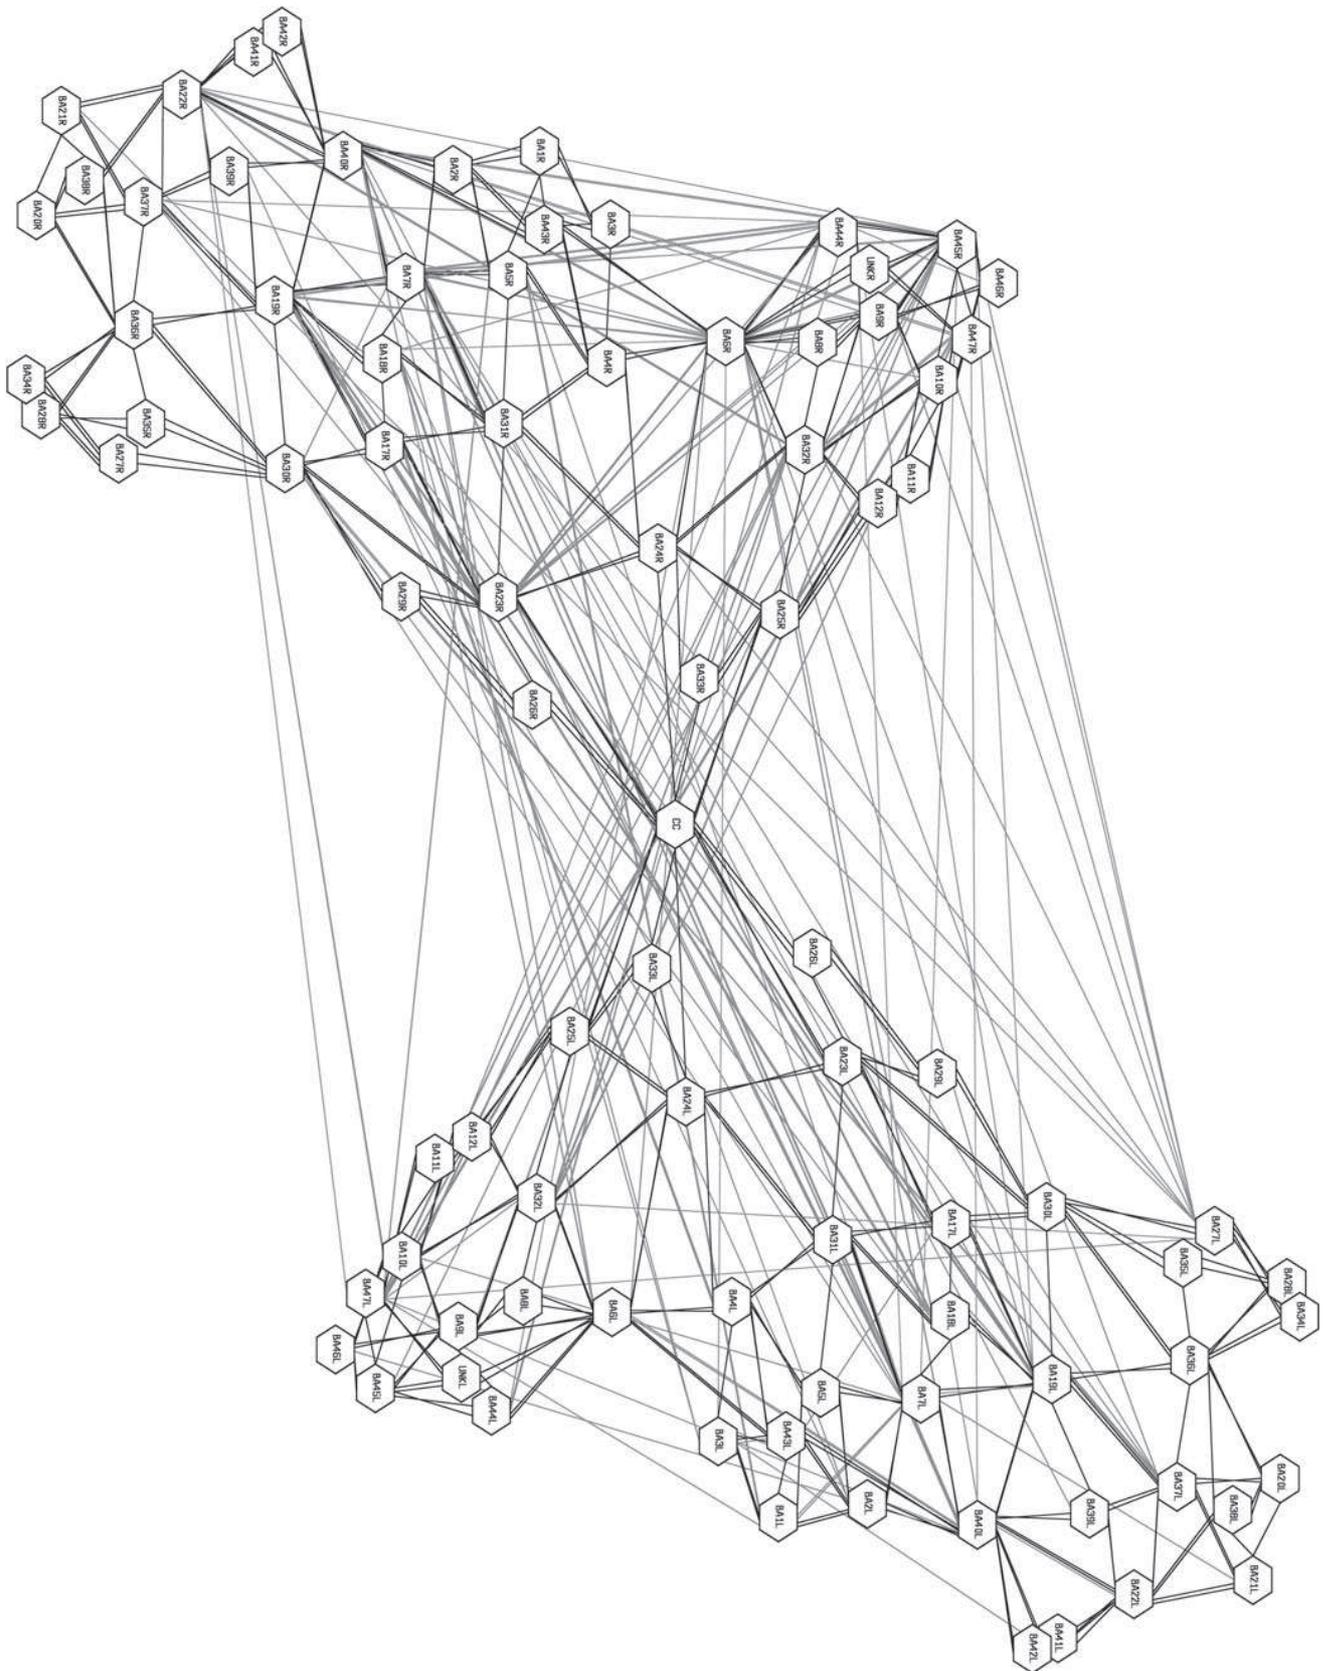


Figure 1 (Anderson). Cortex represented as adjacency + co-activation graphs. Here the Brodmann areas are nodes, with black lines between adjacent areas and gray lines between areas showing significant coactivation. The graph on the left shows coactivations from 56 action tasks, and the graph on the right shows coactivations from 77 attention tasks. Graphs rendered with aiSee v. 2.2.

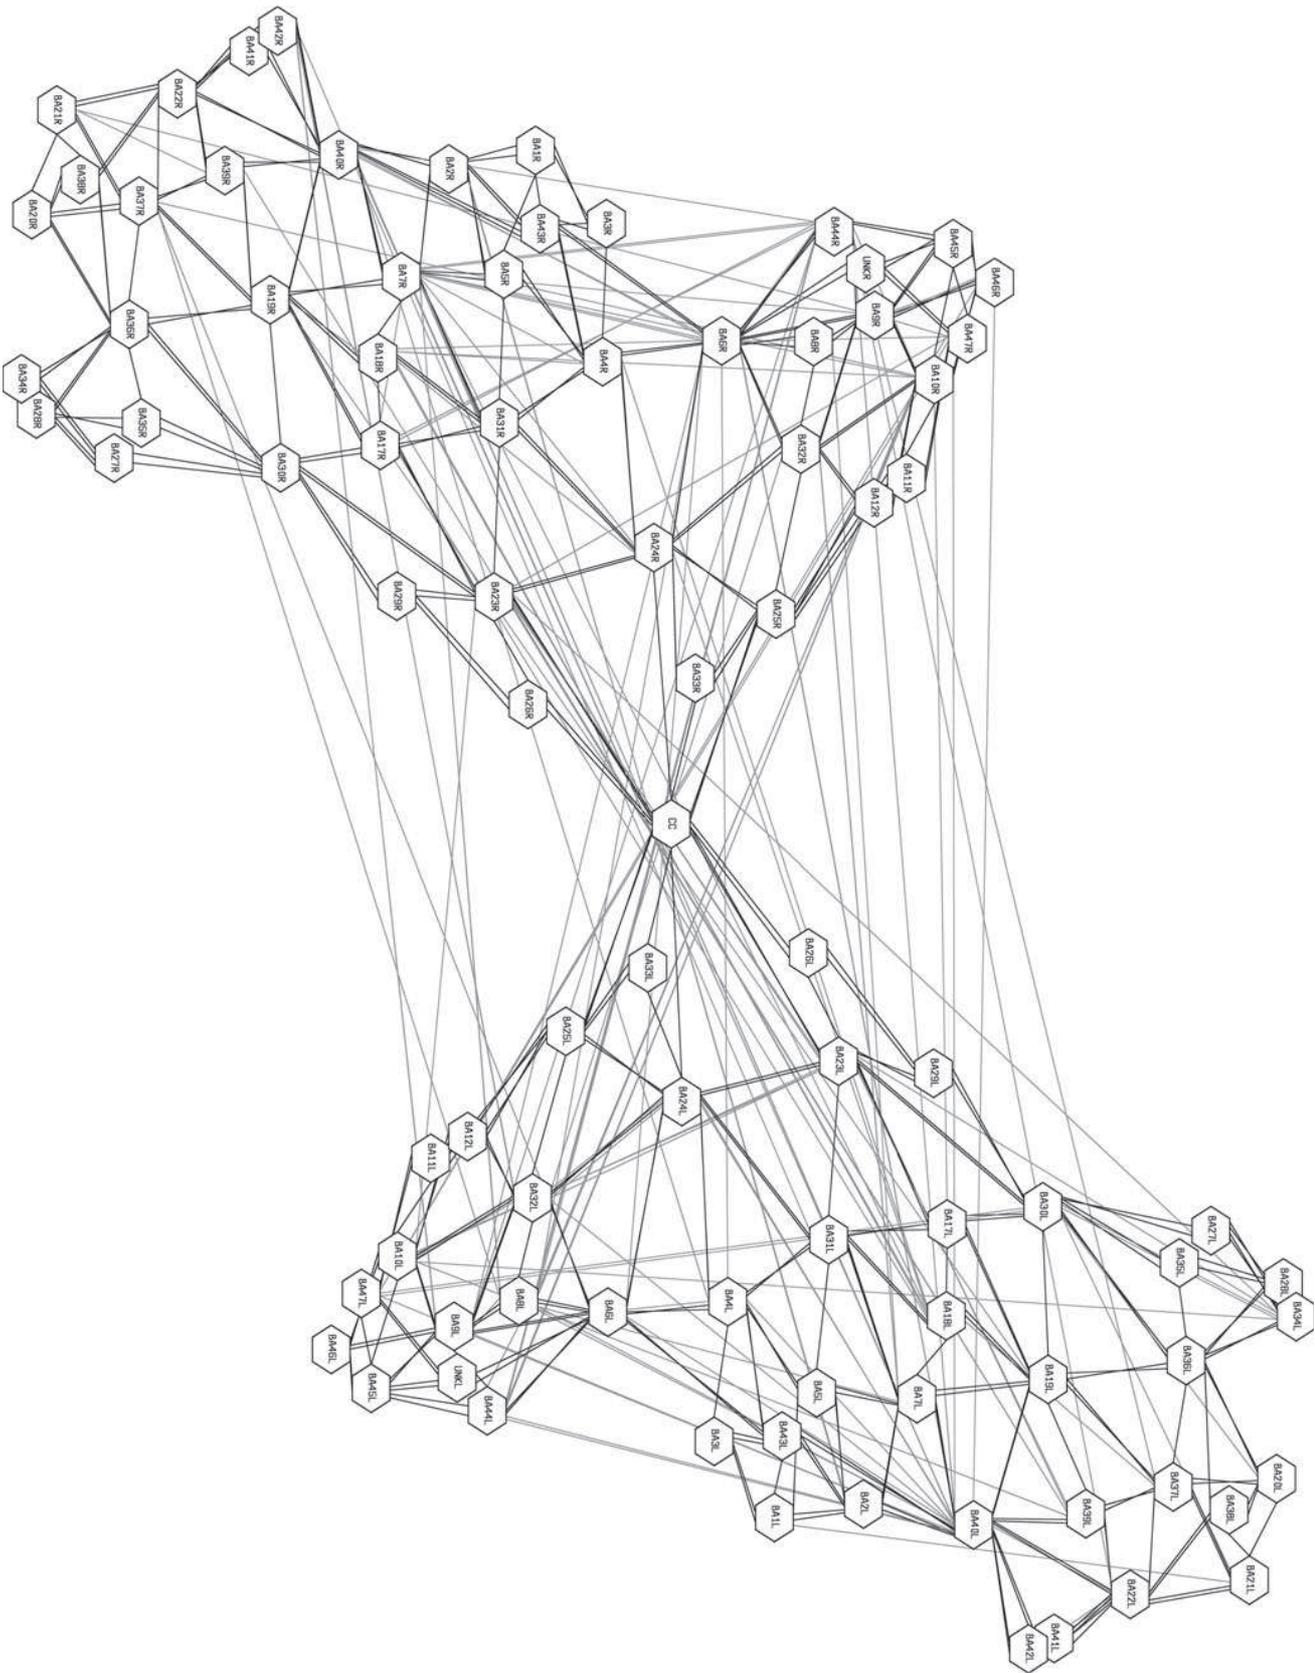


Figure 1 (Anderson). Continued.

brain areas cooperate in specialized ways to support cognitive functions, individual brain areas participate in many such networks, in different cognitive domains.

Much as I would love to have the time and space to comment more fully on this fascinating and comprehensive book – for example, the issues of embodied cognition in general, and embodied representations in particular, are topics near and dear to my heart – I will conform to the *BBS* mold and comment on only the one aspect of the book that is most relevant to my most recent work: the framework of *Interactive Specialization* (IS). In what follows I will offer: (1) a method for investigating this framework, and (2) some data that may or may not be consistent with it (depending on what is meant by “specialization”).

In contrast to localization-based approaches, which at the extreme tend to expect a 1:1 correspondence between cognitive functions and brain areas, the IS framework in *Neuroconstructivism, Vol. 1: How the Brain Constructs Cognition* (Mareschal et al. 2007a), treats cognitive functions as the “emergent product of interactions between different brain regions” (p. 60). For IS, the targets of functional brain mapping should be collections of cooperating neural circuits, not individual brain areas; and the focus of developmental neuroscience should be on changes in interregional connectivity, rather than on the maturation of individual regions.

I wholeheartedly applaud this shift in focus. Networks of cooperating brain regions are *prima facie* far more plausible targets of (high-level) functional attributions than are individual areas, so much so that I find it something of a puzzle that there hasn't been far more effort along these lines (which is not to imply that there has been none). I think one reason is that the techniques typically used in cooperation-focused investigations are complex, they are hard to master, and they produce results that can be difficult to interpret. Thus I would like to take a moment to outline a *very simple* analytical technique we have been using in my lab (Anderson et al., in press).

The technique involves the simple expedient of choosing a spatial segmentation of the cortex (current analyses use Brodmann areas, but any consistent segmentation scheme will work) and analyzing large numbers of fMRI studies in some specific cognitive or behavioral domain to identify statistically significant instances of coactivation. We recently compiled a database of fMRI studies containing 530 experiments in 18 cognitive domains, with the results of each experiment coded primarily in terms of which Brodmann areas contained post-subtraction activation(s). The baseline chance of activation for each area is determined by dividing the number of experiments in which it was activated by the total number of experiments in the database. Then, for each pair of Brodmann areas, we use a chi-square measure to see if their observed degree of coactivation (in a given domain) was significantly different from what would be predicted by chance. With these coactivated pairs, we also perform a binomial analysis to obtain directional information. (It is sometimes the case that, while area A and area B are coactive more (or less) often than would be predicted by chance, the effect is asymmetric, such that area B is more active when area A is active, but not the reverse.)

The results of such analyses are conveniently represented as a graph, where the nodes are Brodmann areas, and edges between the nodes indicate significant coactivation. Figure 1 shows the graphs from one such analysis, for a set of action and attention tasks. We hypothesize that the network of coactivated areas revealed by such analysis represent those areas of the cortex that cooperate to perform the cognitive tasks in the given domain.

With the data in this format, it becomes possible to formulate some very simple questions, and use some well-understood methods to answer them. For instance, in graph theory a

clique is a set of fully interconnected nodes that are sparsely connected to the rest of the graph (Alba 1973). Neural cliques, which in this context represent small networks of coactive areas operating independently of other parts of the graph, seem likely to correspond to the neural components that support a set of closely related cognitive functions or subfunctions. One can also look at other features of the graphs, such as local topography, to help make plausible inferences about underlying function. For example, a hub-and-spoke pattern of coactivation may indicate broadcast or information consolidation functions; in contrast long strings of connected nodes might indicate serial processing. Since one can also look at the emergence of such structures over time, this seems an ideal tool for investigating the IS framework. The basic data for such investigations will become increasingly available as results from more longitudinal fMRI studies are released. More generally, I think that graph theory is an underutilized tool in cognitive neuroscience, and coactivation graphs in particular offer the promise of making cooperation-sensitive investigations into neural function more broadly intelligible.

This brings me to some promised data. We generated coactivation graphs in eight cognitive domains (action; attention; emotion; language; memory; mental imagery; reasoning; visual perception) from 472 experiments in our database and evaluated how much edge overlap and how much node overlap there was among the various domains. In this context, node overlaps indicate Brodmann areas that support tasks in both domains, whereas edge overlaps indicate similar patterns of cooperation. Given the focus on the developmental emergence of networks of brain areas to generate cognitive functions, the IS framework would seem to predict little edge overlap, as differences in function would presumably result from differences in interregional cooperation. This is just what we found; using Dice's coefficient as our measure ($D = 2(o_{1,2}) / (n_1 + n_2)$, where o is the number of overlapping elements, and n is the total number of elements in each set) we found very little edge overlap between the domains (Mean(D) = 0.15, SD 0.04). However, we found a *great deal* of node overlap (Mean(D) = 0.81, SD 0.04). It is not clear if the IS framework is compatible with this latter finding, because Mareschal et al. do not distinguish between increasing *computational* specialization in development (whereby a given region comes to have an increasingly well-defined subfunctional role in the networks in which it participates) and *domain* specialization (whereby increasing specialization of functional networks also implies increasing dedication of participating regions to a narrow and domain-restricted range of cognitive functions). The result reported is just one among a number of findings that suggests that brain regions are *not* domain specialized entities, but in fact typically support functions across many different cognitive domains (Anderson 2007a; 2007b; 2007c). I would welcome the authors' comments and clarifications on this issue.

A good approach to neural and behavioural development but would be even better if set in a broader context

doi:10.1017/S0140525X08004093

Patrick Bateson

Sub-Department of Animal Behaviour, University of Cambridge, Cambridge CB23 8AA, United Kingdom.

ppgb@cam.ac.uk www.zoo.cam.ac.uk/zoostaff/bateson.htm

Abstract: An attractive feature of *Neuroconstructivism, Vol. 1: How the Brain Constructs Cognition* is its emphasis on the active role of the