example, some variants of reading disorders turn out be consequences of genes (or environmental insults) that affect wide-ranging developmental processes, there may still turn out to be others that derive from rarer genes that have purely domain-specific consequences. In a recent study, Ramus and colleagues (Ramus et al., in press) identified a pure form of dyslexia in which only phonological knowledge is impaired (even after compensatory mechanisms are factored out). It would be foolish to rule out a priori the possibility that such a focused deficit might have a specific genetic basis. The same could be said for language disorders: truly domain-specific disorders might be rare in style='color:black'>, but their relative rarity should not lead us to rule out the possibility of modularity.

No simulation or theoretical analysis will save us from the careful empirical work ahead. T&K-S are right to advise caution, but we should be no quicker to dismiss modularity than we should be to accept it.

ACKNOWLEDGMENTS

This work is supported by the NIH, HFSP, the MacArthur Foundation, and the Center for Advanced Study in Behavioral Sciences.

Double dissociations never license simple inferences about underlying brain organization, especially in developmental cases

James L. McClelland[a] and Gary Lupyan[b]

[a]*Department of Psychology and the Center for the Neural Basis of Cognition, Carnegie Mellon University, Pittsburgh, PA 15213;* [b]*Center for the Neural Basis of Cognition, Carnegie Mellon University, Pittsburgh, PA 15213.*

jlm@cnbc.cmu.edu glupyan@cnbc.cmu.edu

http://www.cnbc.cmu.edu/~jlm http://www.cnbc.cmu.edu/~glupyan

Abstract: Different developmental anomalies produce contrasting deficits in a single, integrated system. In a network that inflects regular and exception verbs correctly, a disproportionate deficit with exceptions occurs if connections are deleted, whereas a disproportionate deficit with regulars occurs when an auditory deficit impairs perception of the regular inflection. In general, contrasting deficits do not license the inference of underlying modularity.

Thomas & Karmiloff-Smith (T&K-S) have done the field an important service by calling attention to the fact that lesions to a developing brain can have different consequences from lesions to adult brains. They are entirely correct in pointing out that adaptive learning processes will shape the acquired functions not only of those parts of the brain directly affected by the damage, but also other parts of the brain. In general, the article makes a point we heartily agree with, namely, that the standard logic of neuropsychological interpretation should not be applied uncritically to the interpretation of developmental disorders.

In our view (see also Plaut 1995), the standard logic of neuropsychological interpretation cannot be applied uncritically to the interpretation of any disorders, whether or not they are developmental disorders. By the phrase "the standard logic of neuropsychological interpretation" we mean the reliance on a double dissociation between performance on materials from two different experimenter-defined categories to infer that normal performance relies on separate modules specialized for processing the different categories of materials. This logic has repeatedly been used in both the adult neuropsychology literature and in the literature on developmental disorders discussed by T&K-S (for discussions, see Plaut 1995; Shallice 1988). For example, a double dissociation between living things and artifacts in picture naming and property verification has been used to argue for separate modules for different semantic categories (Warrington & McCarthy 1987), and a double dissociation in reading abstract versus concrete words has been used to argue for separate modules for abstract versus concrete nouns (Warrington 1981). Closer to the example used by T&K-S, a double dissociation in production of the past tenses of exception words versus regular words has been used to argue for separate brain mechanisms for words and rules (Pinker 1991; Pinker & Ullman 2002).

In all these cases, connectionist/parallel-distributed processing models (Rumelhart et al. 1986) have provided alternatives to the standard interpretations of these double dissociations (Farah & McClelland 1991; Joanisse & Seidenberg 1999; Plaut 1995). In general, these models take the following form: A single integrated and interactive system is used for processing items of both categories. Because of item characteristics that covary with category membership, performance on items from one category depends more on one part of the system, whereas performance on items from the other category depends more on another part of the system. For example, Plaut (1995) suggested that concrete and abstract words may differ in the number of semantic features. He trained a network with feed-forward and recurrent connections to map both concrete and abstract words from orthography to semantics. The concrete words contained more semantic features, so they used the recurrent connections more effectively and were less dependent on the feed-forward connections. Lesioning the feed-forward connections produced a relative deficit for abstract words, whereas lesioning the recurrent connections produced a relative deficit for concrete words. Standard neuropsychological reasoning would interpret this double dissociation as evidence of separate modules for abstract and concrete words, but there is no such modularization.

Similar problems arise in a developmental context. Different impairments imposed on the system at the beginning of training can differentially impact *learning* to perform correctly with items of different types. Again, standard neuropsychological reasoning would incorrectly imply that the underlying organization is modular.

To illustrate this point, we revisited Simulation One from T&K-S. We trained a single, three-layer, feed-forward network on both regular and exceptional past-tense forms using the same architecture and training patterns (kindly provided by Michael Thomas). Performance after training the intact network was 100% correct for items of both types, in line with the proposal (Rumelhart & McClelland 1986) that a single integrated system might underlie the processing of both regular and exceptional forms.[1] We replicated their "intact" condition, which produced 100% correct performance on both regulars and exceptions, and a fairly severe (80%) "starting point" lesion, which produced fairly good performance on regulars but poor performance on exceptions. We included another condition expected differentially to impair performance on regulars (following Hoeffner & McClelland 1993). This condition builds on the suggestion (Leonard 1998; Leonard et al. 1992) that the English regular past-tense inflection may be weakly represented in the speech signal and therefore difficult to perceive for children with certain forms of language impairment (see also Tallal 1995). Specifically, the nonsyllabic forms of the English past tense (/t/ and /d/, as in "liked" and "loved") involve very slight additions that agree in voicing with the preceding sound and can be very difficult to detect (Bird et al., in press), and the syllabic form (/^d/ as in "hated"‘) is unstressed. To simulate a deficit in perceiving these inflections, the units representing the past-tense inflection were sometimes set to 0 in the target past-tense pattern that the network is given as a model for what it should learn. This is based on the idea that children learn from what they hear and that the perceptual impairment makes the inflection sometimes imperceptible. This condition was otherwise identical to the intact condition. What we see in Table 1 is that in this new condition there is a disproportionate deficit in processing regular past tenses.[2]

Looking at Table 1, we clearly see a pattern of double dissociation. If we found two groups of children who exhibited the two patterns seen in the table and then employed standard neuropsy-

Table 1 (McClelland & Lupyan). *Performance inflecting regular and exception words after training under two different forms of developmental abnormality*

	Item Type	
Deficit Type	Regular	Exception
Intact Network	100%	100%
Sparse connectivity (80% of connections removed)	75%	22%
Perceptual Deficit (60% failure to perceive the regular inflection)	38%	97%

chological reasoning, we would conclude that the normal brain contains separate systems for processing regular items and exceptions. This inference would be incorrect, however, because in this case we know that intact performance is generated by a single system that processes both regular and exceptional forms correctly. We simply have two different deficits that differentially impair learning to process the different types of items. A lesion that produces sparse connectivity reduces the ability of the network to become sensitive to particular combinations of input phonemes that must be considered simultaneously to inflect an exception correctly. An impairment that impacts perception reduces the network's exposure to the information that indicates the correct pronunciation of the regular past tense.

In summary, T&K-S have sounded an important note of caution, indicating that standard neuropsychological reasoning cannot be applied uncritically to the interpretation of patterns of deficits seen in developmental disorders. We hope we have underscored their point by noting that this caution is important in interpreting adult as well as developmental cases. Our simulation suggests that developmental "double dissociations" can be especially misleading because developmental disorders can produce contrasting differential deficits in a single integrated mechanism, but the application of standard neuropsychological reasoning would interpret this pattern as evidence for a two-part system.

NOTES

1. The simulation reported by T&K-S showed relatively poor generalization to novel items. There are several possible reasons for this: (1) the training corpus (which is based on one used previously by Plunkett & Marchman 1993) employed a relatively large number of exceptions compared to regular forms; and (2) the patterns used for the variants of the regular past tense inflection do not reflect its phonological characteristics or its systematic relation to the phonological features of the stem. We would expect that a corpus that more realistically reflected the frequency structure of the language and the phonology of the regular inflection would produce a higher level of generalization.

2. We do not wish to defend the particulars of this simulation as an adequate model of the phonological impairment of any real children; specifically, we believe that in reality, such impairments affect perception of some aspects of exception items as well as regular items, so that the differential would not be as extreme.

Weak evidence for a strong case against modularity in developmental disorders

Ralph-Axel Müller

Department of Psychology, San Diego State University, San Diego, CA 92120. **amueller@sciences.sdsu.edu**
http://www.sci.sdsu.edu/~amueller/index.html

Abstract: Thomas & Karmiloff-Smith (T&K-S) provide evidence from computational modeling against modular assumptions of "Residual Normality" (RN) in developmental disorders. Even though I agree with their criticism, I find their choice of empirical evidence disappointing. Cognitive neuroscience cannot as yet provide a complete understanding of most developmental disorders, but what is known is more than enough to debunk the idea of RN.

The saddest part about this target article is its apparent necessity. It is surprising enough that some cognitive neuropsychologists discount plasticity in adults with acquired lesions. Brain reorganization in recovering aphasics has been shown in many studies, which contradicts Residual Normality (RN). Arguably, the precise links between language recovery and reorganization are under debate for adult acquired lesions (Rosen et al. 2000; Thulborn et al. 1999). In *developmental* populations, however, these links cannot be denied. Loss of the left hemisphere in childhood is often accompanied by good language outcome (Vargha-Khadem et al. 1997), and evidence suggests more pronounced interhemispheric reorganization in children than in adults (Müller et al. 1999). Even for comparatively well-defined structural lesions, the RN assumption is therefore inaccurate: The remaining neurofunctional system *changes* after damage. Can RN be appropriate for developmental disorders that are not even fully understood pathogenetically?

As Thomas & Karmiloff-Smith (T&K-S) argue convincingly, RN in developmental disorders requires strong assumptions that are probably untrue. They present evidence from neural network modeling, contending that cognitive neuroscience has no definitive answers. I disagree: First, computational models are informed by highly oversimplified properties of neural architecture and function, and results may not fully apply to biological brains. For example, the biological meaning of added noise or reduced discriminability in T&K-S's lesion models remains fuzzy. More importantly, neuroscience provides a wealth of evidence showing that RN cannot be expected in developmental disorders.

For a trivial start, genes do not code for cognitive modules (Gottlieb & Halpern 2002). As long as there is loose talk about "genetic double dissociations" (Pinker 1999), an innate "language acquisition device," and "the genetically determined component of the brain . . . that is dedicated to . . . language" (Chomsky 2002, pp. 83–85), ad nauseam repetition is unavoidable: The genome may be a "code," but one for proteins, not cognitive subsystems. The RN assumption could be reasonable if modular cognitive dysfunctions were linearly caused by aberrant genes. There are indeed developmental disorders with known genetic causes (e.g., phenylketonuria or fragile-X syndrome), but these affect multiple brain and body systems in pleiotropic ways. The same applies to language impairment in family KE, caused by a defective FOXP2 gene (Lai et al. 2001). Behavioral patterns and brain morphometry suggest a broad deficit spectrum with pronounced motor involvement, not a specific language impairment (Watkins et al. 2002a; 2002b).

I will discuss the limits of the RN assumption, taking autism as an exemplary developmental disorder. Modularity has been claimed regarding theory-of-mind (ToM) deficits in autism (Leslie 1992; Scholl & Leslie 2001), which are supported by behavioral and neuroimaging evidence (Baron-Cohen et al. 1999; Castelli et al. 2002; Rutherford et al. 2002). Neuroimaging and lesion studies have also identified brain sites in the temporal and frontal lobes that are normally involved in ToM (Calder et al. 2002; Stone et al. 1998). ToM in autism can therefore serve as a test case for the RN assumption in similar ways to past-tense formation in Williams syndrome and developmental language impairment, as discussed in T&K-S.